THE TECH CAPITAL

THE SOVEREIGNTY MIRAGE

THE ILLUSION OF INDEPENDENCE IN A CONNECTED WORLD

EC's Ursula von der Leyen's vision of a sovereign Europe collides with a global reality that won't bend to borders.

Netherlands
AMSTERDAM
AMS
Est. 2016

California
Est. 2022
ARCATA
ACV

Georgia
ATLANTA
ATL
Est. 2019

Spain
Est. 2021
BARCELONA
BCN

Massachusetts
BOSTON
BOS
Est. 2015

Belgium
BRUSSELS
BRU
Est. 2022

Argentina
BUENOS AIRES
BUE
Est. 2019

China
BEIJING/TIANJIN
TJ1
Est. 2022

Illinois
Est. 2017
CHICAGO
CHI

India
CHENNAI
CHN
Est. 2022

Colorado
DENVER
DEN
Est. 2015

edgeconnex®

Hyperlocal to Hyperscale
Customer Defined Data Centers

Ireland
DUBLIN
DUB
Est. 2017

Germany
FRANKFURT
FRA
Est. 2025

Virginia
HERNDON
HQ
Est. 2009

Texas
HOUSTON
HOU
Est. 2013

India
Est. 2021
INDIA
IND

Turkey
ISTANBUL
IST
Coming Soon

Florida
JACKSONVILLE
JAX
Est. 2015

Indonesia
JAKARTA
JKT

Malaysia
Johor Bahru
JHB
Coming Soon

Malaysia
KUALA LUMPUR
KUL
Est. 2025

Nevada
LAS VEGAS
LAS
Est. 2015

Wisconson
MADISON
MAD
Est. 2015

Tennessee
Est. 2015
MEM
MEMPHIS

Arizona
MESA
AZ

Mexico
MEXICO CITY
MEX
Coming Soon

Florida
MIAMI
MIA
Est. 2015

Germany
MUC
MUNICH
Est. 2017

India
NOIDA/DELHI
NDA

Japan
OSAKA
OSA
Coming soon

Arizona
PHOENIX
PHX
Est. 2014

India
PUNE
PNQ
Est. 2025

Enabling High Performance Compute & AI Solutions

edgeconnex.com

Virginia
Est. 2015
RICHMOND
RIC

California
SACRAMENTO
SAC
Est. 2014

UTAH
SALT LAKE CITY
SLC
Est. 2014

California
SAN DIEGO
SDG

California
SANTA CLARA
SVC
Est. 2015

Washington
SEATTLE
SEA
Est. 2014

China
SHANGHAI
SHA
Est. 2022

Singapore
SGP
Regional HQ
Est. 2021

Florida
TALLAHASSEE
TAL
Est. 2014

Israel
TEL AVIV
TLV
Est. 2021

Poland
WARSAW
WAW
Est. 2018

GLOBAL AWARDS
MAY 11, 2026
LONDON

THE TECH **CAPITAL**

IFF APAC
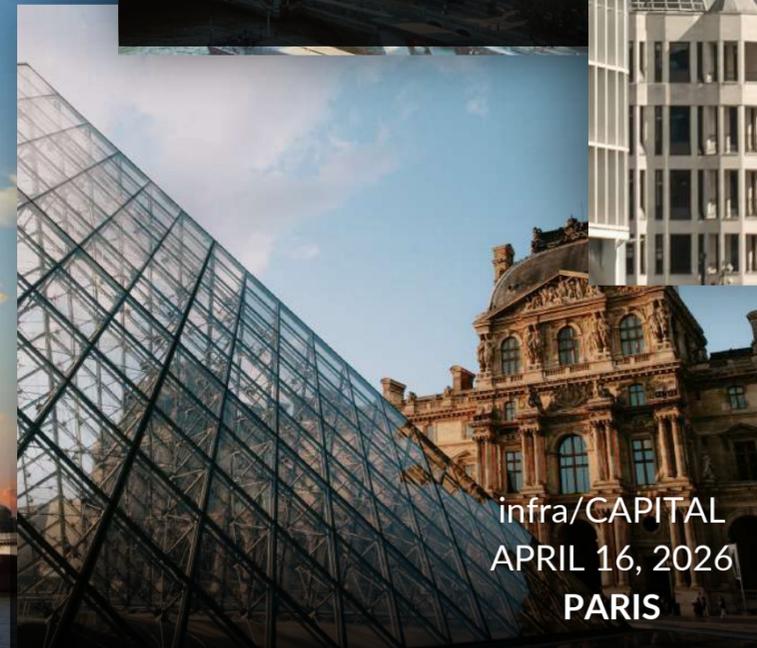SEPTEMBER 2026
SINGAPORE

IFF LATAM
JUNE, 2026
SAO PAULO

IFF GLOBAL
MAY 11-12, 2026
LONDON

IFF MEA
NOVEMBER 24, 2025
ABU DHABI

infra/CAPITAL
APRIL 16, 2026
PARIS

INFRA|AI GLOBAL SUMMIT
MARCH 30-APRIL 1, 2026
ATHENIAN RIVIERA

REGISTER TODAY
www.events.thetechcapital.com

IN THIS ISSUE

WINTER 2025 | ISSUE 04

THE TECH CAPITAL
www.thetechcapital.com

IN THIS ISSUE

WINTER 2025 | ISSUE 04

THE TECH CAPITAL

The freshest content in one scan

IN THIS ISSUE

WINTER 2025 | ISSUE 04

A FRAGMENTED WORLD

As 2025 comes to an end, it feels as though the world's digital backbone has matured - and fractured - all at once. Connectivity, data, and computation have become instruments of both economic growth and national strategy. Across boardrooms and ministries, questions of access, control, and investment now share the same table.

The sector has grown fast, but unevenly. Markets wrestle with capital costs, regulation, and geopolitical friction, yet ambition remains undimmed. What began as an industry of engineers and financiers has become one where lawyers, policymakers, and sovereign investors shape its next phase. That's why this issue carries particular weight for us.

At the centre of this edition is The Legal 50, the first-ever ranking of legal professionals focused on digital infrastructure. These are the individuals structuring deals and frameworks that sustain the world's data economy.

From London to Lagos and Dubai to Seattle, dealmakers are adapting to a more complex financial landscape. In the Middle East - one of our core focuses this season - Islamic finance is emerging as a structural pillar of digital asset investment. Rooted in Shariah principles that prohibit interest and speculation while favouring real economic activity, Islamic instruments are proving compatible with the sector's tangible and predictable nature. Sukuk (Islamic bonds) and commodity murabaha (cost-plus sale) are increasingly paired with conventional debt to finance AI-ready data centres, fibre, and platforms. As one banker in Riyadh told us, this is "no longer an alternative channel - it's a parallel system of value creation."

Policy remains the other frontier. We look at the debate over "tech sovereignty" and its rise in national investment agendas. European Commission President Ursula von der Leyen - depicted on our cover - once said: "The trick is to build public digital infrastructure that is interoperable, open to all and trusted." Her words capture the tension now shaping everything from data localisation to AI governance. Sovereign wealth funds channelled US$9.4 billion across 53 digital infrastructure deals, including US$5.4 billion in data centres - proof that states now view connectivity and computation as strategic assets.

Innovation remains the sector's pulse. Photonic computing - once confined to research labs - is edging towards commercial reality. Our team explores how the use of light instead of electrons for data processing could lower costs, cut heat, and reshape the energy profile of AI computing.

The year ahead brings uncertainty. The OECD forecasts global GDP growth slowing to 2.9% in 2026, as tariffs, fiscal constraint, and policy risk weigh on markets. Yet, constraint has a habit of breeding innovation - and infrastructure has always thrived in periods of reinvention.

As we unveil this edition at The Tech Capital's Middle East & Africa Finance Forum in Abu Dhabi, I am reminded that finance remains the connective tissue of transformation. For four years, our International Finance Forum has convened leaders shaping the world's digital economy. Bringing that dialogue to the Middle East and Africa signals the region's growing role in defining the next chapter. Thank you for joining us on this journey. I wish you a productive close to 2025 and a successful year ahead.

João Marques Lima

THE TECH CAPITAL

GLOBAL

AWARDS

ACHIEVEMENTS OF EXCELLENCE

The world's first truly dedicated digital infrastructure finance awards

Visit
www.events.thetechcapital.com
for more information.

SUBMISSION CLOSE JANUARY 2026 | GALA NIGHT - MAY 11, LONDON

MEET THE TEAM

The Tech Capital

João Marques Lima
Managing Editor

Jack Haddon
Editor

Abigail Opiah
Production Editor

Khamila Mulia
APAC Editor

João Roque Marques
Newsroom Associate

Lilou Cortesi
Business Development Manager

Anaelle Bach
Marketing Executive

Ana Catarina Vaz
Digital Media Assistant

Faith Apejoye
Production Assistant

João Lopes
Graphic Designer

Joanne Dixon
Executive Assistant

Mary Calçada
Finance Director

Next issue March 2026

FINDING BALANCE IN A TIME OF CONSTRAINT

HOW POLICY STRAIN AND TRADE FRICTIONS WILL SHAPE THE 2026 ECONOMIC OUTLOOK

by João Marques Lima

As growth cools and policy uncertainty deepens, investors in digital infrastructure face a shifting landscape where tariffs, fiscal restraint and green regulation decide who keeps pace - and who falls behind. João Marques Lima writes.

The global economy enters 2026 on a weaker footing. The OECD projects world GDP growth to ease from 3.3% in 2024 to 3.2% in 2025 and 2.9% in 2026, as higher tariffs, reduced front-loading of production and policy uncertainty weigh on trade and investment.

While early 2025 showed resilience-driven by pre-tariff stockpiling and AI-related investment, the second half marked a turn. The full weight of tariff increases, particularly from the United States, is now filtering through to supply chains and consumer prices. The OECD warns the effects are "becoming increasingly visible in spending choices, labour markets and consumer prices".

Slowing but uneven growth

The United States is expected to see GDP growth slow from 2.8% in 2024 to 1.8% in 2025 and 1.5% in 2026, as consumption softens and the effective tariff rate, which reached 19.5% by late 2025, constrains trade. Fiscal support and high-technology investment will provide only partial offset.

Europe's outlook remains subdued with the euro area forecasted to expand by 1.2% in 2025 and 1.0% in 2026, as easier credit conditions are balanced by higher trade frictions. Germany's modest fiscal expansion contrasts with consolidation in France and Italy. In the United Kingdom, growth is expected to ease from 1.4% to 1.0% amid tighter fiscal policy and weaker external demand.

Asia continues to provide relative momentum, though the pattern is uneven. China's growth is projected to slow from 4.9% to 4.4% as fiscal support fades and tariff effects deepen. India remains a regional bright spot, growing 6.2% in 2026 on the back of policy easing and reforms. Indonesia holds steady at 4.9%, while Japan's expansion cools from 1.1% to 0.5%.

Elsewhere, Brazil slows to 1.7%, Mexico edges up slightly to 1.3%, and Saudi Arabia sustains output above 3.5% on firm energy investment. What unites these cases is a loss of global demand momentum, with investors confronting weaker trade flows and tighter financial conditions.

Monetary and fiscal tightrope

Headline inflation across the G20 is forecast to fall from 3.4% in 2025 to 2.9% in 2026, though divergence remains. The United States may face renewed price pressures from tariffs, while the euro area approaches 1.9%. Japan returns to target near 2%.
Disinflation has stalled in many economies as food and services costs

persist. "Rising food prices are behind a resurgence of goods inflation," the OECD notes. At the same time, wage growth has moderated without fully restoring real incomes, particularly in the United Kingdom, Canada and Japan.

Financial markets have stayed buoyant, led by technology stocks and a record US$3.9 trillion in crypto-asset capitalisation. Yet the OECD cautions that "stretched asset values and growing concern about future fiscal risks" could expose vulnerabilities. Sovereign spreads in Europe have widened, US term premia remain high, and gold prices have risen by around 40% since January 2025-signals of latent unease.

Central banks began easing in late 2025. The US Federal Reserve and Bank of Canada cut rates, while the ECB held steady and the Bank of Japan maintained gradual tightening. Emerging-market central banks, including Brazil and Türkiye, are expected to follow with modest reductions.

Fiscal space, however, is narrowing. Persistent deficits, rising debt service costs and structural pressures from defence, ageing and climate adaptation are eroding flexibility. The OECD urges credible medium-term frameworks to safeguard debt sustainability. Coordinated policy credibility-where monetary easing is matched with disciplined fiscal management-is now critical. Without it, further rate cuts could unsettle rather than stabilise bond markets.

Emerging market debt maturities by the end of 2026 add further strain. Refinancing needs and exchange-rate volatility threaten to raise project financing costs, particularly where revenues are denominated in local currency. Governments with limited foreign-exchange buffers could face higher risk premiums, curbing access to long-term capital.

Investment caution amid structural drags

Private investment remains below pre-pandemic levels across most economies. Uncertainty around trade, taxation and regulation continues to dampen capital formation. Demographic ageing in advanced economies is constraining labour supply, while productivity gains have yet to materialise.

The OECD expects global potential growth to slow from roughly 3% in 2025 to about 2% per cent in the early 2030s. Without reform, that deceleration could harden into a structural constraint. Weaker-than-expected growth or persistent inflation could also trigger risk repricing in financial markets, where valuations, particularly in technology equities, are already stretched.

The OECD identifies faster artificial intelligence adoption as one of the few potential upsides to this baseline, capable of lifting productivity and moderating inflationary pressures. That same dynamic reinforces long-term demand for digital infrastructure, from cloud networks to high-capacity data processing, supporting selective capital deployment even in a slower economy.

Digital infrastructure finance in 2026

The digital economy remains one of the few areas of sustained capital inflow, but 2026 marks a shift from expansion to consolidation. With global growth cooling and financing costs adjusting after years of tightening, investors are reassessing risk exposure to large-scale data and connectivity projects.

Power constraints, fragmented regulation, and sustainability compliance are now central to investment planning worldwide. In North America and Europe, hyperscale cloud campuses are integrating renewable energy and exploring modular designs to reduce environmental impact. Publicly listed operators face stricter ESG reporting and energy-efficiency benchmarks, rewriting financing and risk assessment.

In Asia Pacific, projected data centre capacity will exceed 37,000 MW by 2030, though capital discipline is tightening. Southeast Asia could reach 5.2-6.5 GW, yet financing increasingly depends on green benchmarks and policy support. For example, Singapore links approvals to efficiency and carbon-performance thresholds, while Japan and Korea pilot hydrogen-powered and renewable-driven facilities.

Middle East and Africa are emerging as new hubs for cloud infrastructure, leveraging abundant solar potential and improving connectivity to support regional AI workloads and 5G rollout.

Capital formation is also evolving with strategic partnerships and joint ventures now the preferred route to share cost and regulatory exposure. KKR's US$800 million acquisition of a 20% stake in Singtel's Digital Infrastructure Co. and Equinix's partnership with Astra International in Indonesia reflect this model. Such collaborations blend global capital with domestic presence, aligning commercial and policy priorities.

Yet valuation discipline is tightening. Publicly listed data-centre operators saw steep gains through 2025, fuelled by AI-driven demand projections. A reassessment could follow if slower growth or higher funding costs persist. Emerging-market refinancing pressures could further test leveraged projects reliant on dollar funding.

Nevertheless, structural demand for digital capacity endures. Cloud migration, AI workloads and 5G rollout continue to drive network expansions. The focus in 2026, however, is shifting from growth at scale to sustainable yield. Private equity investors are emphasising operational efficiency, power optimisation and modular design to secure long-term returns.

A Singapore-based infrastructure fund manager noted: "The question for 2026 isn't whether data demand will grow - it's how to finance that growth responsibly in a world of slower GDP and stricter energy policy."

The Search for stability

For policymakers, 2026 is a test of balance. The OECD urges governments to cooperate on trade transparency and resist further fragmentation, arguing that open and predictable trade remains vital for confidence.

For investors, the year will hinge on selectivity and risk discipline. Slower growth, tariff tensions and fiscal constraints will keep volatility elevated, but opportunities in green and digital infrastructure anchored in policy stability and productivity enhancement-offer resilience.

Beyond 2026, the OECD anticipates a world of slower but more differentiated growth, where competitiveness will depend on productivity reform, energy transition and credible policymaking. For digital infrastructure, success will belong to those combining financial execution with stability-balancing ambition with restraint in an era defined by limits.

Caspar David Friedrich's The Wanderer above the Sea of Fog (1818) - a timeless reflection on navigating uncertainty and opportunity in 2026.

C-LEVEL VIEW ———————————

THE 2026 DIGITAL SURGE ACROSS MEA
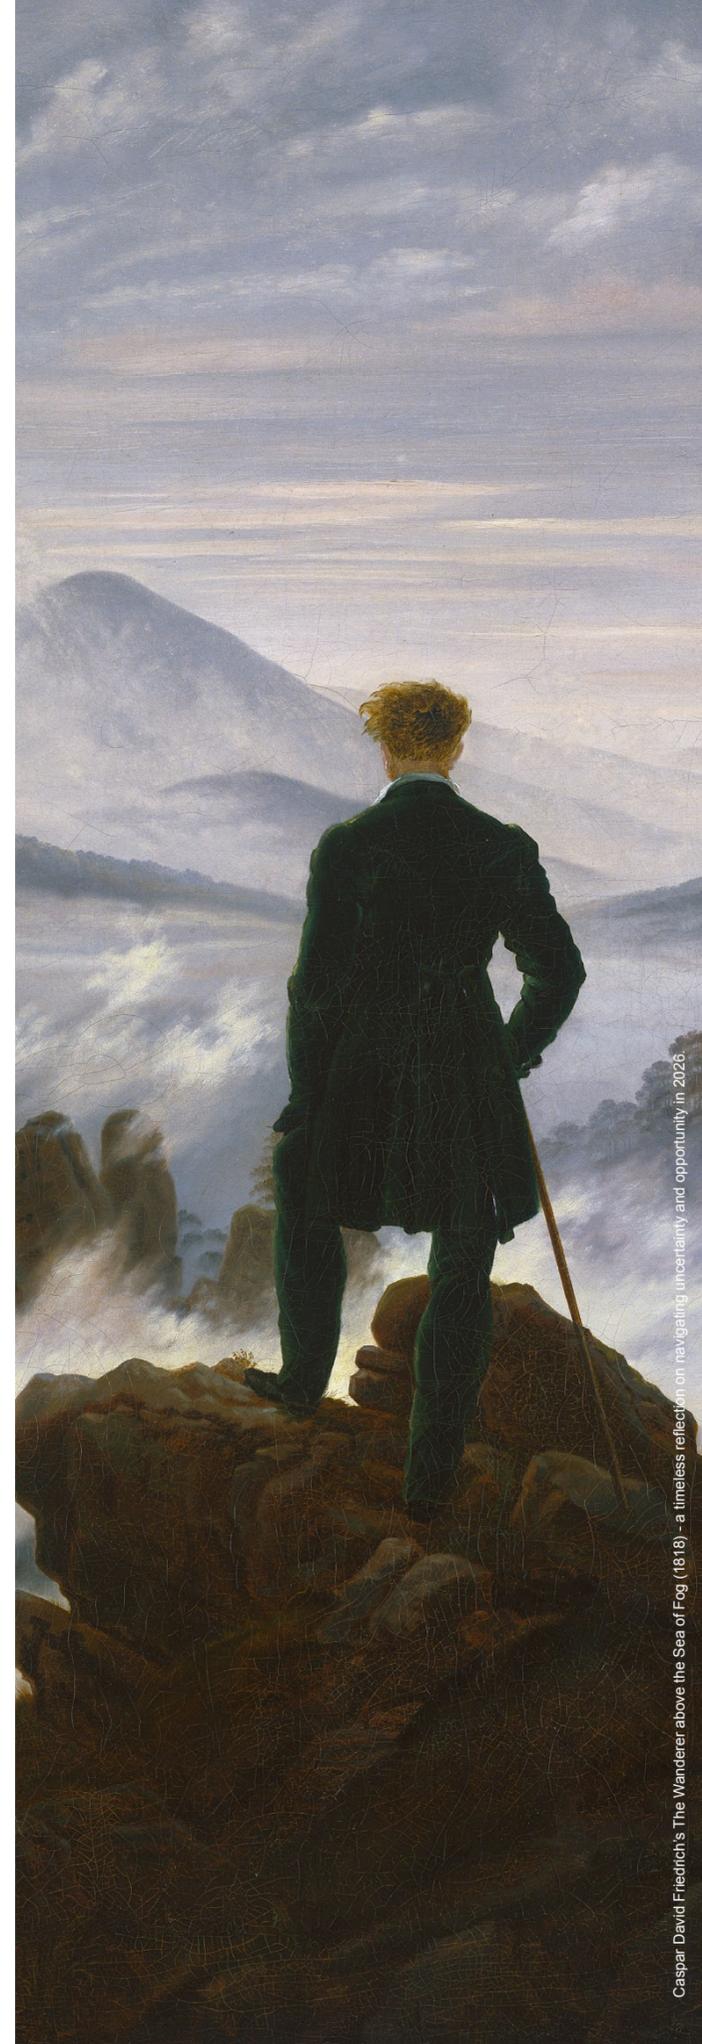
AI is booming, cloud is expanding, and data hubs are springing up across deserts and coastlines. The Middle East and Africa are gearing up for a digital revolution - and 2026 is when it all clicks into gear. Khamila Mulia reports.

The Middle East and Africa are no longer catching up - they're setting the pace. Africa's internet economy has already blown past US$180 billion, while the Gulf is fast-tracking its ambition to become the world's next AI powerhouse. The result? An insatiable demand for scalable, sovereign, and sustainable digital capacity.

But the road ahead isn't frictionless. Energy constraints, new sovereignty mandates, and the relentless need for real-time, low-latency connections are forcing providers to rethink what digital resilience really means.

The Tech Capital brought together some of the region's most influential voices to map the year ahead. What they reveal is a market brimming with possibility - one where bold partnerships, sharp investments, and visionary infrastructure will decide who leads the next wave of MEA's digital revolution.

Paula Cogan
EMEA Managing Director, Digital Realty

Looking ahead to 2026, we see demand for digital infrastructure across the Middle East and Africa continuing to grow at pace, driven by the twin forces of demographic growth and digital transformation. Africa in particular stands at a defining moment: a young, connected population and rapid innovation across sectors are fuelling an internet economy projected to already be worth close to $180 billion. None of this progress is possible without the right foundation; reliable, interconnected, and sustainable data centre capacity.
As this demand grows, we'll see the competitive landscape become more dynamic. Hyperscalers are deepening their regional footprints, telcos are expanding into cloud and edge services, and local operators are scaling fast to meet new requirements. The opportunity is vast enough to support collaboration across these ecosystems, and the winners will be those who can combine global reach with local insight and long-term partnership.
The main challenges next year will be power and permitting, with talent also an emerging pressure point.

Steven Marshall
Chairman - Infrastructure, Cordiant Capital

We expect hyperscale data centres to continue attracting the most funding. Interconnection hubs and fibre networks are becoming more strategic, while edge infrastructure will see targeted investments in markets driven by mobile, fintech, and smart-city applications.
Investors are concentrating on assets that combine scale, resilience, and multi-service capabilities. AI-driven demand is amplifying financing constraints: with national grids under strain, power availability and grid access are critical. Proximity to reliable, affordable energy has become a key determinant of bankability. Capital costs and permitting remain hurdles, especially in emerging markets.
To navigate these constraints, deal structures are becoming more flexible, combining long-term capital, strategic partnerships, and risk-mitigation instruments. Power-purchase agreements and inflation-linked tariffs manage energy and operational exposure. Tight alignment among developers, operators, and investors is essential to withstand volatility.
Competition is maturing. Governments, hyperscalers, and investors are racing to secure energy-reliable sites, fuelling joint ventures and sale-and-leasebacks that let hyperscalers scale quickly without overextending balance sheets. Investors are competing for pre-leased assets that deliver immediate yields. Execution is now a differentiator.

Dr. Thomas King
CTO, DE-CIX

We are seeing continued AI sector growth in the Middle East, with recent moves by OpenAI and Anthropic adding momentum to high-performance computing data centre projects. In Africa, development is advancing more gradually, led by local providers, while hubs such as Johannesburg, Cape Town, Nairobi, and Lagos are attracting international investment.
Across both regions, enterprises and governments need high-performance, low-latency connectivity for AI inference and data-sovereignty needs, and internet exchanges (IXs) help by aggregating local networks and edge providers to bring traffic closer to end users.
A key challenge is the distance to major cloud and content platforms, which adds latency when traffic transits via Europe. The spread of carrier-neutral IXs, increased peering, and more cloud on-ramps is shortening paths. Cloud-exchange functions on IXs, including multi-cloud routing, give operators greater control over data flows and can improve resilience via dedicated connections.

Snehar Shah
CEO, IX Africa

We are seeing "reverse migration" from clouds hosted outside the continent, where high costs in hard currency, plus the IP transit costs, are offset by colocation deployments with local data centres. We are also starting to see the first AI deployments, with niche players providing "GPU as a service" to financial institutions, research organisations, and so forth.
South Africa has been the most advanced market on the continent, with some 400 MW of deployed capacity. Relatively, Kenya and Nigeria, which are seen as the next frontier markets, stand at less than 5% of that. In Kenya, both Oracle and Microsoft have publicly announced their cloud deployments; we should soon see active deployments in 2026.
There are other non-US sovereign cloud players also planning their entry into our market. As infrastructure deployed by banks over the last decade comes to the end of its life, we are starting to see migration to the cloud, and, coupled with data sovereignty requirements, this is accelerating the need for local cloud zones in the region.

Dr. Ayotunde Coker
CEO, Open Access Data Centres

Digital infrastructure challenges vary across the continent. Power is the main global data centre issue. In Africa, the grid rarely meets all requirements, so captive power solutions are essential.
Accessing the right capital structure is a continuous challenge for CEOs, given the high capital needs for infrastructure. Complicated supply chains and long equipment lead times are also issues. Attracting, developing, and retaining skills is a global challenge we face in Africa, though local talent is strong and responsive to training.
I expect rising demand in Tier 1 markets - South Africa, Nigeria, Egypt, Kenya, and Morocco - with Tier 2 growth in Ghana and Côte d'Ivoire as regulation and infrastructure mature. Kinshasa, DRC, is also gaining momentum.
Key drivers include hyperscale cloud rollout, mobile broadband adoption, digital transformation, and data sovereignty. AI and machine learning are accelerating edge and interconnection investments, with low-latency inference workloads driving demand for GPUs as a service.
South Africa already has ample hyperscale capacity, while Nigeria's new projects exceed 80 MW. Subsea connectivity is in place, and scaling will continue through 2026-27 as localisation and competition intensify.

Daniel Efrati
CEO, NED Data Centers

The cloud remains the dominant force, expanding at around 30% annually across Tier 2 and Tier 3 markets, as enterprise clients continue to migrate from on-premises environments to collocation and cloud-based solutions. Inference-led AI deployments are also gaining traction, supported by government efforts to establish sovereign AI capabilities and national computing strategies. This is creating demand for smaller, distributed "mini AI factories" and sovereign computing projects, such as Israel's national supercomputer tender, which was recently awarded to Nebius.
While large-scale AI training remains concentrated in a few established hubs, particularly in the Gulf where collaborations between the UAE, Saudi Arabia, and US partners are driving investment in regional AI training infrastructure, the broader EMEA landscape is being reshaped by inference-driven demand.
As international hyperscalers expand their presence and regional telcos move further into the data centre and cloud ecosystem, competition will intensify.

Matt Evans
CEO, Lennox Data Centre Solutions

The immediate pressure point is the supply chain. If demand spikes and manufacturing cannot keep pace, we are exposed. At current scales, only a few companies can deliver what is needed, which strains the ecosystem.
The winners will be those who engineer for resilience and adaptability from day one. The only way to future proof a data centre is to build it into the DNA, from design to delivery, so it can evolve when the industry shifts again.
The skills shortage is an even bigger concern. There are not enough experienced people for the demand ahead. If acceleration arrives as predicted, projects will slip, builds will be pushed back, and some may pause. We need to keep existing knowledge in the industry and pass it on.
Large scale retrofits are the next opportunity. Many 10- to 15-year-old sites were built for a different era. Retrofitting to support modern workloads will deliver efficiency gains, with cloud demand likely to be the main catalyst.

THE SOVEREIGNTY MIRAGE

THE ILLUSION OF INDEPENDENCE IN A CONNECTED WORLD

by Abigail Opiah

AS EUROPE RACES TO RECLAIM CONTROL OF ITS DIGITAL FUTURE, THE WORD "SOVEREIGN" HAS BECOME BOTH A SHIELD AND A SLOGAN. YET BENEATH THE RHETORIC, THERE IS AN AWKWARD TRUTH: NO ONE CAN QUITE AGREE ON WHAT "SOVEREIGNTY" MEANS - OR WHO GETS TO DEFINE IT. ABIGAIL OPIAH REPORTS.

"The trick is to build public digital infrastructure that is interoperable, open to all and trusted."

These, now famous, words from European Commission President Ursula von der Leyen have become something of a mantra for Europe's digital policymakers, a vision of sovereignty without isolation, and openness without vulnerability. But the deeper the EU moves into debates over data localisation, AI governance and cloud control, the harder it becomes to define what that balance looks like.

That same tension is now visible in global investment patterns. According to the IFSWF Annual Review, in the full year of 2024, sovereign wealth funds channelled $9.4 billion across 53 deals in digital infrastructure, with $5.4 billion flowing into data centres alone, a signal that nations are treating connectivity, storage and computation not just as commercial assets, but as strategic instruments of resilience in an era of climate and geopolitical strain.

When Europe's digital ministers met in Amsterdam in early 2025, their closing statement contained a phrase that is fast becoming unavoidable in technology and infrastructure circles: "tech sovereignty."

The D9+ group, which is a coalition of pro-tech European nations that once prided itself on resisting protectionism, called for policies to "strengthen technological capabilities" and "increase tech sovereignty." That shift may sound procedural, but it marks a political and philosophical turning point. For years, the D9+ had defended open digital markets and warned against regulatory overreach. Now, faced with new US tariffs on European industries, the same countries are rethinking the wisdom of relying on America's tech giants.

At the 2025 State of the Union Address, President von der Leyen, confirmed massive investment in European AI gigafactories.

"They support our innovative start-ups to develop, train, and deploy their next-generation AI models. When we called on the private sector to join forces with us, the response was overwhelming," she said.

"I will meet CEOs from some of the largest European tech champions. They will hand over their European AI & Tech Declaration. This is their commitment to invest in Europe's tech sovereignty."

She also sought to dispel the notion that Europe has fallen behind in the global AI race. "Too often, I hear that Europe is late to the race, while the US and China have already gotten ahead. I disagree. Because the AI race is far from over," von der Leyen stated.

"First, we want to accelerate innovation. Europe has some of the world's fastest public supercomputers. We are now putting them at the service of our best start-ups and scientists, so they can forge the AI we need," she added.

Her remarks highlight a theme that runs through Europe's approach to digital sovereignty, which is the belief that open collaboration and public infrastructure can coexist with strategic autonomy. Sovereignty, in this view, is not about closing borders, it is about building capacity.

Yet beneath the rhetoric, there is an awkward truth: no one can quite agree on what "sovereignty" means.

In the data centre and cloud world, "sovereign" has become the new "sustainable." The word is plastered across press releases, keynote slides and product roadmaps, which is a shorthand for autonomy, trust and national control. But ask five executives what it means, and you will get five different answers.

"The conversation around sovereign infrastructure often needs more nuance. There is an important distinction between data locality, data residency and data sovereignty. Data locality, or ensuring data stays within a specific region, is something cloud providers can already guarantee," says Kevin Dunn, Vice President and General Manager, EMEA at Wasabi.

"True sovereignty, however, implies every component in the technology stack would need to be designed and manufactured in a specific country to be deemed acceptable, and this is a reality far from being realised."

For some, it is about data residency, with the aim of ensuring that information stays physically within national borders. For others, it is about jurisdiction in order to make sure that no foreign government can compel access to that data, even if it is stored locally. Some even see sovereignty as the ability to design and operate digital infrastructure, from semiconductors to AI models, without depending on another country's supply chain.

Peter Gadd, Vice President, International at VAST Data believes that true sovereignty is not about where data resides, it is about control.

"It is the ability to decide who accesses your data, where it is processed, and under what legal framework. Many vendors talk about 'keeping data in-country,' but if you cannot audit, encrypt, and govern it under your own rules, you don't have sovereignty, you have a postcode."

That ambiguity makes the term politically useful and commercially convenient. "Sovereign" sounds defensive but visionary, pragmatic yet patriotic. It allows European governments to push back against American dominance without explicitly calling it protectionism and lets cloud providers repackage compliance controls as strategic independence.

But the vagueness also risks hollowing the word out completely. If every service can be branded as "sovereign," then the concept itself loses power, and the security and policy implications get lost in the noise.

"The word 'sovereign' is being thrown around like confetti right now. But if everything is sovereign, nothing is," says Gadd.

"The danger is we dilute the term before we have properly defined it. Real sovereignty demands transparency, accountability, and technical enforcement, not just a marketing label."

The Cloud Act and the jurisdictional fault line

The conversation around sovereignty sharpened in 2018, when the US Congress passed the Clarifying Lawful Overseas Use of Data Act, better known as the Cloud Act. The legislation obliges American companies to provide data to US authorities, even if that data is stored on servers in other countries.

"For instance, should the US government request data from a US company which operates in the EU, things can get complicated, putting regulations like GDPR and the US Cloud Act into conflict with one another. Without bilateral agreements, companies risk breaching local laws and face conflicting obligations. Ultimately without any overarching international frameworks it begs to question whether any entity can truly 'own' data," says Robbert Santifort, Principal Associate, Eversheds Sutherland.

"In an AI-driven economy, the concept of 'sovereignty' may evolve from exclusive control to a new kind of shared stewardship where value is co-created, but responsibility is collective."

For European governments and regulators, the implications were alarming. The law effectively extends US jurisdiction into foreign soil, directly clashing with the EU's own privacy regime under the General Data Protection Regulation (GDPR). It means that a European public agency could, in theory, find its citizens' data handed over to a foreign government, without consent, and without the host country having any say.

"The Cloud Act exposed a hard truth: physical data residency does not equal legal sovereignty. You can store your data in Frankfurt or Paris, but if it is sitting on infrastructure owned by a US company, you are still under a level of US jurisdiction," says Gadd.

"That is an uncomfortable paradox most organisations have only recently started to understand. But there are ways to mitigate this and regain control."

That tension has never fully gone away. Instead, it has deepened. With an increasingly unpredictable White House and the possibility of another Trump presidency, European policymakers have grown nervous about the potential misuse of that power. The question is not just can the US access European data, but would it, and under what circumstances?

For governments building AI systems that depend on sensitive or classified data, that risk is not theoretical. Defence departments, energy ministries and health agencies are all starting to explore large-scale language models and predictive algorithms. Training those systems often involves ingesting vast datasets of citizens' information, exactly the kind of material no state wants slipping beyond its jurisdiction.

In Europe, President von der Leyen confirmed the set-up of 12 AI factories at the AI action summit in February. "Now we are also bringing AI factories to the next level, with AI gigafactories. Very large data and computing infrastructure, to train very large models," she said.

And that is where the notion of sovereign infrastructure re-enters the picture.

Fiona Ghosh, partner at law firm Ashurst, said this reflects a wider regulatory balancing act taking shape in the UK and beyond.

"This is a good example of the UK government balancing the tightrope of refusing (for now) to regulate AI en masse with the creation and need for guardrails around safe and responsible innovation. It is a well-trodden route for UK plc to deploy sandboxes when encouraging new use cases and pushing at the boundaries of regulatory perimeters; indeed, we lead the world in this regard.

"What is most interesting, is that the government's chosen sectors are those that bring the most immediate benefits to the widest range of people namely, healthcare, transport, and housing. The inclusive nature of the government's approach is to be lauded and echoes the hallmarks of its key policy objectives.

"Therefore, the blueprint aligns with the government's wider 'Plan for Change' where the goal of regulatory reform is for 'growth' as opposed to just risk-mitigation. Expect to see more of the same soon."

Her remarks capture the tension at the heart of modern digital governance, between freedom to innovate and the need to contain systemic risk. The UK's "sandbox" model has become a proving ground for testing that balance, offering a controlled environment for experimentation without surrendering oversight. Similar logic now applies to digital sovereignty: governments are seeking to build trusted systems that remain open enough to evolve, but secure enough to protect their citizens' data and interests.

Infrastructure under your own flag

If sovereignty once referred mainly to the data layer, it now extends all the way down the technology stack, to the physical servers, the fibre routes, even the power supply. A government that truly wants to protect its digital assets would need to control not just the code but the compute: the machines, the cooling systems, the locations and the networks through which information flows.

That is easier said than done, and here is why. The hyperscale cloud model was never built for jurisdictional purity. By design, these platforms replicate data across regions to guarantee uptime and resilience. The default is global redundancy, information duplicated and distributed across data centres in multiple countries.

"Data ownership has shifted beyond the simple question of possession to a legal and technical puzzle that's moving faster than regulation can keep up with. Jurisdictions are defining 'ownership' differently under their own laws. It's therefore up to tech companies to put processes in place to protect customer data," says Santifort.

Not every customer has the knowledge of where their data sits at any given moment, and working around those defaults is complex, time-consuming and expensive.

Mark Chester, Co-head of Data Centres, Eversheds Sutherland seems to think that ownership of physical assets such as land, data centres and hardware is relatively clear-cut. However, the software layer is much more complicated.

"Questions of control, access, and genuine ownership (as distinct to a licence to use it) are complex, particularly when you add the operational services into this mix," he adds.

To address those concerns, the major hyperscalers in the form of AWS, Microsoft, Google and others alike, have launched "sovereign cloud"

offerings, promising that data will stay within a defined geography and that access will be restricted to in-region personnel. These models often rely on strict technical and legal controls to demonstrate compliance with local laws.

The underlying technology stacks, control planes and update mechanisms still depend on global systems. In practice, that means a degree of dependency, and trust, remains unavoidable.

"Building and operating sovereign AI infrastructure requires vast GPU capacity, reliable power and advanced cooling - resources that are costly and scarce in many regions. Governments and enterprises are navigating competing legal frameworks, such as the Cloud Act and GDPR, which create uncertainty around where data can reside and who can access it. Many AI workflows still depend on cross-border data exchanges for training, labelling and optimisation, so imposing strict territorial limits can disrupt operations or require fundamental redesigns," says Jon Healy, Managing Director EMEA at Salute.

"Operationally, the shortage of skilled talent adds another layer of difficulty. Hyperscalers are adding local controls, yet legal exposure and vendor dependency persist. Emerging neocloud providers offer flexibility and clarity, but not always at scale."

The rise of the neoclouds

That dependency has opened a gap in the market for what some are calling "neoclouds." These are newer, smaller providers building infrastructure from the ground up with sovereignty in mind.

According to JLL, the neocloud ecosystem encompasses around 190 individual operators, with leading players such as CoreWeave, Nebius, and Crusoe.

These companies deliver specialised GPU resources through Graphics Processing Units-as-a-Service (GPUaaS), supporting demanding workloads in AI, blockchain, gaming, and scientific computing.

The market has shown rapid growth: CoreWeave's share price has nearly tripled since its March IPO, Nebius has raised $700 million in 2024 followed by another $1 billion in 2025, and S&P Global estimates that more than $10 billion was invested in this sector last year.

According to a recent analysis in the Washington Examiner, these companies which specialise in AI-optimised compute, are positioning themselves as an alternative to the legacy hyperscale model.

"The challenge is that sovereign has become a marketing label. We now see hyperscalers offering 'sovereign clouds' that promise data locality, but the underlying control planes, update mechanisms, and support contracts often still depend on foreign entities. That's not sovereignty; that's segmentation," says Dom Wilde,
SVP & Genegral Manager of Core Business, Mirantis.

"Newer, open-source-based 'neocloud' providers are starting from a different premise. They are building systems that can be verified, audited, and operated by local teams, with transparent code paths and

the option to run entirely disconnected if needed. It's sovereignty by design, not by region."
Instead of layering new controls on top of a globally distributed platform, neoclouds are designing for localisation from the start. They focus on contained geographies, transparent supply chains and energy-efficient operations, often powered by renewable or stranded energy sources.

"New AI cloud providers and national platforms have a clean slate; they can design for sovereignty from the outset. They can build governance and auditability into the fabric of the system, rather than bolting it on as an afterthought. That's where we're seeing real innovation today," adds Gadd.

In theory, this makes neoclouds more agile, less politically exposed and better suited to customers who prioritise jurisdictional control over global elasticity.

It is a compelling pitch, especially for European governments seeking both digital autonomy and credible AI capability. Yet neoclouds also face a reality problem: they lack the scale, ecosystem and maturity of the incumbents. Building a truly sovereign cloud is not just a technical exercise; it is an economic and diplomatic one. It requires collaboration across energy, security, procurement and policy, areas where the EU still struggles to act as a unified entity.

A word that says too much… and too little

"Sovereignty" has become a linguistic mirror, reflecting each constituency's fears and ambitions. For For von der Leyen, it is a way to argue for investment and independence without descending into isolationism. For cloud vendors, it is a market differentiator. For regulators, it is a rallying cry for accountability.

But as the term spreads, it risks losing its edge. Just as "green" once became a label for everything from carbon offsets to office recycling bins, "sovereign" now risks being stretched to the point of meaninglessness. Not every workload needs national custody, and not every data centre with a flag outside is truly sovereign inside.

"Ultimately, the question every organisation should ask is: If the legal, political, or commercial relationship with my cloud provider changed tomorrow, would I still be able to operate? If the answer is no, then you don't truly have sovereignty, you have dependency," adds Wilde.

In a nutshell, sovereignty in digital infrastructure is not about slogans or server locations. It is about trust, who holds it, who enforces it, and who benefits when it is broken.

"I believe that true sovereignty will depend on deeper collaboration between policymakers, technologists and infrastructure providers to balance control, capability and competitiveness," says Healy.

The harder question is not what sovereignty means, but who gets to define it. And right now, in Europe's scramble to protect its digital borders while staying open for business, even that answer remains up for negotiation.

FUNDING AI'S CORE

HOW InfraAI 2026 BRINGS THE WORLD TO GREECE

by João Marques Lima

The sun rises over the Athenian Riviera, casting its light across azure waters and centuries-old olive groves. Here, where democracy, philosophy, and the foundations of human inquiry were first conceived, the world will gather once more - not to debate politics or ethics, but to shape the infrastructure of the future.

From March 30 to April 1, 2026, The Tech Capital's InfraAI Global Summit will convene technologists, operators, policymakers, and investors from every corner of the globe to confront one of the most consequential challenges of the 21st century: building the backbone that will power artificial intelligence worldwide.

Last year, InfraAI 2025 brought together 150 stakeholders from over 20 countries to explore emerging trends in AI infrastructure, often referred to as "neoclouds" - next-generation ecosystems of high-capacity data centres, fibre networks, satellite connectivity, and cell towers. In 2026, the summit expands from a single-day forum to a three-day gathering, creating space for deeper dialogue, cross-border collaboration, and high-level networking in a moment when the global AI ecosystem is accelerating at unprecedented speed.

Artificial intelligence is no longer just a modernisation gamble - it is a geopolitical instrument and an economic lever. Across the United States, hyperscale GPU farms are growing at breakneck pace, while China deploys massive state-backed AI clusters and satellite-based connectivity to power everything from autonomous vehicles to national research initiatives. Europe, meanwhile, is strategically positioning itself, pledging hundreds of billions in public and private investment to close the gap.

AI infrastructure isn't just about computing power; it's about national competitiveness, economic resilience, and global influence. Countries that can efficiently deploy scalable infrastructure will attract talent, capital, and innovation. Those that lag risk losing a generation of AI-driven growth.

At the heart of this race are neoclouds and AI factories - integrated ecosystems combining computational power, connectivity, and real estate capable of supporting the demands of modern AI. Market estimates point out that these assets alone could attract upwards of $7 trillion globally over the next decade or less.

InfraAI 2026 will highlight how these infrastructures are being designed, financed, and deployed. Sessions will explore investment strategies, sustainable energy solutions for high-capacity data centres, and emerging technologies such as AI-optimized hardware and satellite networks. By connecting finance, engineering, and policy in a single forum, the summit provides a rare opportunity to understand how global AI infrastructure can scale responsibly and efficiently.

Europe at the Crossroads

While InfraAI is a global event, Greece's selection as host is symbolic. Europe's AI ambitions are gaining momentum: the European Union has

pledged over €200 billion to advance AI capabilities, complemented by private investment in large-scale "gigafactories" for AI computation. In addition, the EU is funding gigawatt-class data centres designed to house hundreds of thousands of GPUs, aiming to close the gap with the U.S. and China.

Hosting the summit on the Athenian Riviera, at the Asteria Domus Glyfada, merges strategic discussions with a setting steeped in history and humanistic values - a reminder that technological progress must intersect with culture, ethics, and the global common good.

InfraAI distinguishes itself with its global reach. In 2026, attendees will come from North America, Europe, Asia, the Middle East, and Latin America, creating opportunities to exchange knowledge, identify investment opportunities, and coordinate on standards and policies that can accelerate AI adoption responsibly.

Strategic partnerships with organisations such as the European Data Centre Association and Structure Research ensure that discussions translate into actionable collaboration.

The summit's structure - combining keynote panels, roundtables, and networking sessions - encourages dialogue across disciplines. Participants leave not only with insights into infrastructure deployment and financing but also with a clearer understanding of how policy and technology intersect in a rapidly evolving AI landscape.

Athens: Where Strategy Meets Opportunity

The choice of Athens is more than picturesque. This is the cradle of democracy, philosophy, and inquiry - a place where ideas were tested, debated, and refined. Just as ancient thinkers shaped the foundations of modern society, today's summit participants are shaping the infrastructure that will underpin AI for generations to come.

The serene Riviera, with its luxurious venues and historical gravitas, provides the ideal environment for high-stakes discussions and the kind of informal networking that often drives deals and collaborations.

As AI continues to reshape industries, economies, and societies, the role of infrastructure becomes increasingly visible and consequential. The InfraAI Global Summit 2026 offers a rare forum where investors, infra builders, and policymakers can explore scalable solutions, anticipate emerging challenges, and forge global partnerships. By blending international perspectives with in-depth discussions on neoclouds and AI infrastructure, the summit illustrates how strategic gatherings can influence the trajectory of technology and the global economy.

For those seeking to navigate the complex interplay of capital, technology, and regulation, Athens in 2026 offers more than a meeting - it offers a lens into the future of AI infrastructure.

For more information,
visit www.events.thetechcapital.com/page/infraai26

The Middle East & North Africa data centre market is entering a rapid growth phase, with industry estimates placing total IT load capacity at roughly 1.48 GW in 2025, expanding to around 2.98 GW by 2030 as demand surges from cloud, AI, and edge workloads. Major capacity builds, hyperscaler projects, and neocloud commitments are accelerating this expansion. Hyperscalers are anchoring the boom: AWS has pledged multi-billion-dollar investments in Saudi Arabia, while deals such as Microsoft's ~$1.5 billion partnership with UAE AI group G42 underscore current strategies.

LOCATION FOCUS

MENA

Teleography
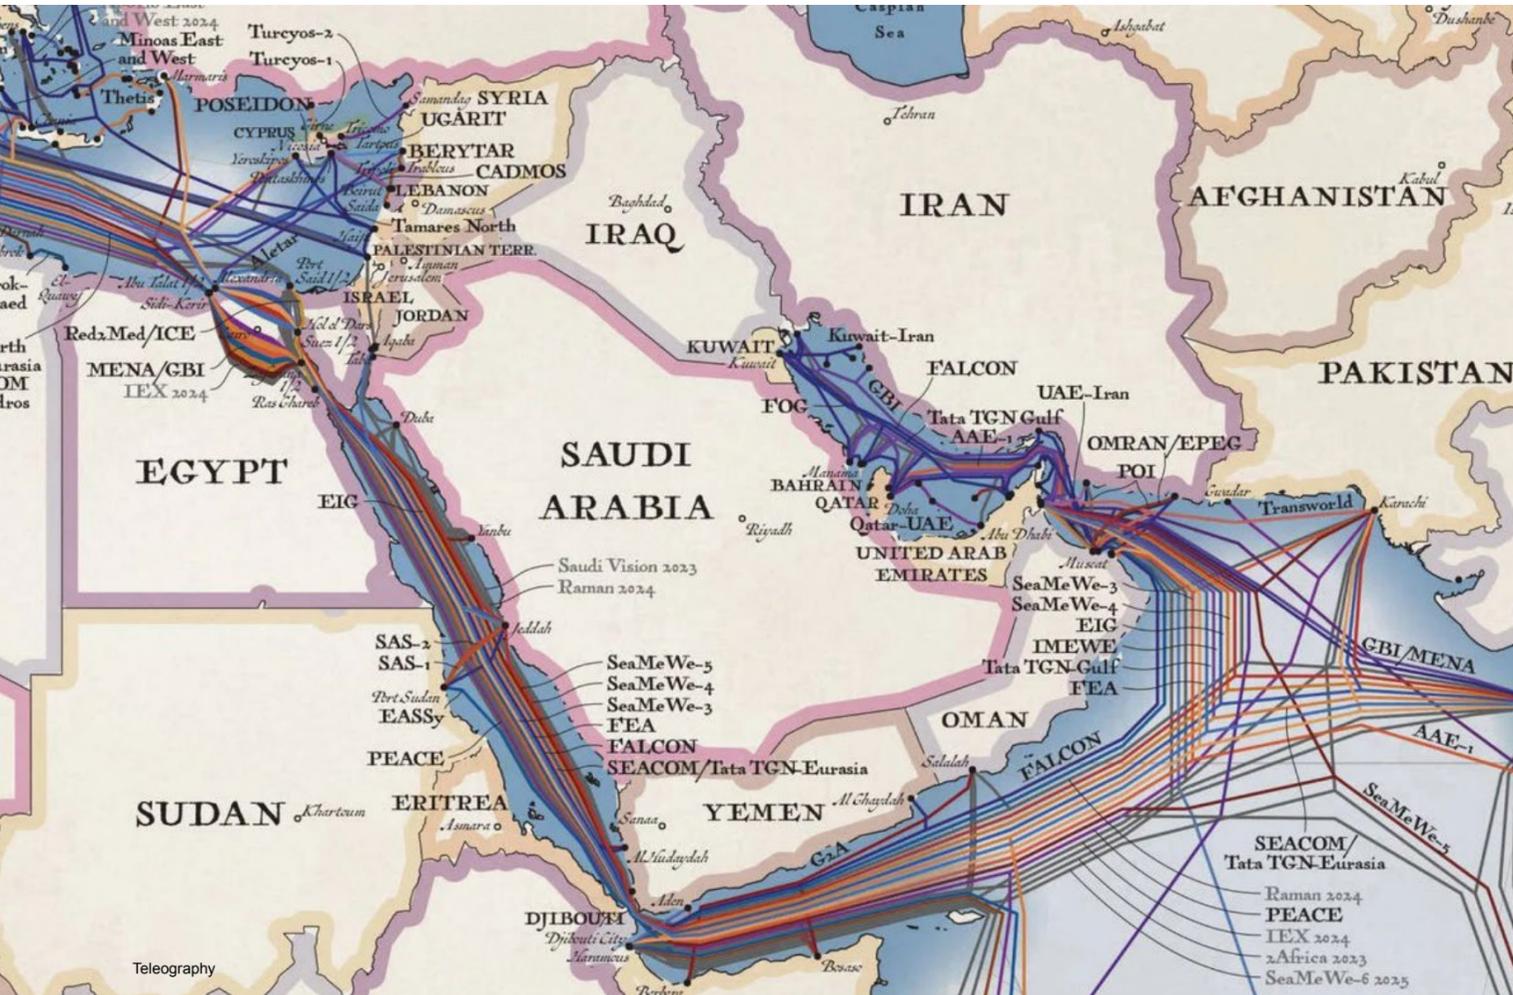
MORE CABLE, MORE CHALLENGES

INSIDE MEA'S SUBSEA PUSH

by Khamila Mulia

New capacity from projects like 2Africa, SEACOM 2.0, and Medusa could cut latency and lower prices across the Middle East, African and more. But recurring Red Sea disruptions mean stakeholders still have work to do to build faster, safer routes.

Over the past few years, the Middle East and Africa (MEA) have emerged as new frontiers for subsea connectivity. A wave of systems, including the newly launched SEACOM 2.0, the Medusa Submarine Cable System which is entering service on the western Mediterranean segment, and Meta's Pan-African Submarine Cable, as well as cross-border infrastructure such as the proposed India–Saudi Arabia and India–UAE power interconnectors, are redrawing digital routes between Africa, the Gulf, India, and Europe.

This push is driven by hyperscaler demand, energy diversification, and regional ambitions to become AI and data centre hubs. However, the region's geography and geopolitical risks, particularly in the Red Sea, are forcing builders to rethink traditional routes and build in more resilience.

The current wave of MEA cable builds is not just about moving more traffic. Governments in the Gulf and parts of Africa are pouring money into AI, cloud regions, and power generation, creating pull for high-capacity international routes.

At the same time, operators and hyperscalers want to de-risk the Europe-Asia path after repeated Red Sea disruptions, so they are looking for more Mediterranean, Atlantic, and Indian Ocean options. "One inefficiency everyone is tackling is route concentration through the Red Sea, and the lack of genuinely diverse landings. That is why we now see redesigns that push traffic through the Mediterranean, the Atlantic, or the Indian Ocean," says Dani Simis, CTO of Israel-based data centre operator NED.

Market observers say Gulf states, especially Saudi Arabia and the UAE, have become preferred locations for energy-hungry AI workloads because they can pair abundant, cost-effective energy, from both fossil fuels and renewables such as solar, with aggressive digital policies. That, in turn, pulls traffic east to Asia and west to Europe and North America.

Claude Achcar, founder and managing partner at Actel Consulting, notes that the same patterns, at a smaller scale, are now visible in African data centre markets, as cloud adoption and sovereignty rules push for more local hosting. "This is leading to the proliferation of subsea cables landing all around the continent and across connecting East and West," he tells The Tech Capital.

Friction points

Despite the scale of investment, the region still faces significant obstacles to delivering subsea cable projects. These include political and security risks, complex regulatory regimes, limited supporting infrastructure, environmental constraints, and basic logistics.

"Permitting and security challenges in high-risk corridors, particularly around the Red Sea, and the need for landing-site approvals now pose greater bottlenecks than wet-plant manufacturing," Simis says.

In practice, the slowest parts of deployment are gaining repair-ship access, securing local permits, and working within narrow marine safety windows. Even on routes that avoid the Red Sea, cross-border coordination for new landings and terrestrial backhaul remains highly complex.

Recent incidents like undersea cable cuts in the Red Sea, that briefly disrupted access in parts of Asia and the Middle East, have delayed several projects. "Geopolitical constraints are dramatic in obtaining the best cable route decision and cable laying operations," says Simis.

Governments and operators have responded by improving protection and surveillance, rerouting traffic, and pursuing diplomatic channels to keep links open. Even so, vulnerabilities in the region persist, which means stronger, coordinated action from all stakeholders is still needed.

If the projects such as ongoing 2Africa landings, SEACOM 2.0, and Medusa's MEA touching segments stay on schedule, the next 12 to 18 months could mark a turning point for MEA's digital economy. The additional capacity should ease congestion, lower latency between the Gulf, India, and Europe, and push wholesale prices down as competition intensifies. But these benefits will only hold if the region can protect and repair the network fast enough. With every new landing station and corridor, the exposure surface widens.

"Some risks are typical of undersea cable projects. Unintentional damage from commercial shipping like anchor dragging and natural hazards have traditionally been the major causes of damage to these cables," says David Hunt, partner in the London office of Boies Schiller Flexner. "The new risk is deliberate acts by both state and non-state actors. Expansion of the cable network also puts more pressure on the limited ships and equipment to repair these cables when something goes wrong."

In this context, he says it is worth noting the increase in state investments in having more resources for cable laying and maintenance, citing Japan's decision to subsidise NEC's purchase of undersea cable vessels.

A basic level of protection is essential, and surveillance and monitoring are becoming more important as threats from hostile actors grow, especially in politically sensitive areas. When incidents occur, diverse routes and networks help contain the impact, with satellites available as temporary supplements. Continued investment in repair vessels and equipment will also be needed to keep these systems running.

"Participating with key regional and international stakeholders will also be crucial to ensure these projects can be completed and repaired without unwanted interference," Hunt continues. "Pressure could also be applied on the maritime industry and the insurers that support them to adhere to stricter maritime standards to reduce incidents of anchor dragging and unintentional damage of cables."

THE HUMAIN CONDITION

SAUDI ARABIA'S GRAND AI EXPERIMENT

by Jack Haddon
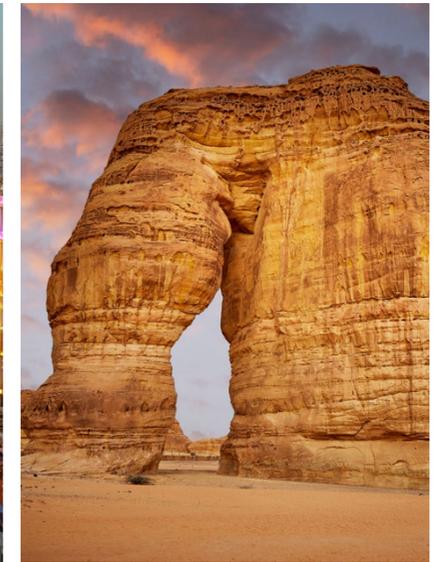
Under the guidance of Crown Prince Mohammed bin Salman, HUMAIN isn't just another tech venture - it's a national project to build the entire AI stack, from silicon to software. With billions in investment and global partnerships, the Kingdom is racing to become an AI superpower. Jack Haddon reports.

In May 2025, Saudi Arabia's Public Investment Fund (PIF) launched a new vertically integrated business designed to build the "entire AI stack", from hyperscale data centers and cloud platforms to native Arabic AI models and applications. In practice, HUMAIN is seeking to establish Saudi Arabia as both a creator of advanced Arabic AI, and potentially a global and regional exporter of AI compute power.

Infrastructure Scale and Partnerships

HUMAIN has already revealed several plans for the build-out of AI computing infrastructure. Alongside its launch in May 2025 it announced two significant partnerships with AMD and NVIDIA to deploy up to 500 MW of capacity with each of the chip designers

GPUs. The initial NVIDIA phase alone includes an 18,000 NVIDIA GB300 GPU cluster. Meanwhile its deal with AMD set an initial target for both parties invest up to US$10 billion to deploy 500MW of AI compute capacity over the next five years.

Separately, Amazon Web Services (AWS) has committed over US$5 billion to develop a Saudi "AI Zone" that bundles AWS cloud services with on-site AI hardware, and HUMAIN has also signed MoUs with Cisco and Qualcomm and is deploying Groq AI accelerators for inference, reflecting an open, multi-vendor ecosystem.

CEO Tareq Amin has outlined targets of roughly 1.9 GW of data-centre capacity by 2030, growing to 6GW by 2034. Construction on two campuses in Riyadh, each with eleven 200MW data centres has already begun, and the AMD/HUMAIN joint venture is aiming to have 50MW live by the end of 2025, with at least 50MW added each quarter for the foreseeable future.

Crucially, domestic power costs will be extremely low. Electricity prices in Saudi Arabia are already the lowest in the world, leveraging the kingdom's already cheap fossil fuels and growing renewable energy

mix. In a bid to lure data centre developers to the Kingdom, Riyadh is making power even more cost effective by subsiding rates for data centres, a policy HUMAIN can expect to benefit from.

Capital Structure and Leadership

HUMAIN is wholly owned by PIF, but Saudi leaders are courting external capital to share risk. Aramco, another PIF and Saudi Government controlled organisation, and one of the largest oil and gas companies in the world, plans to take a minority stake, following a non-binding term sheet signed in late October. That deal will see the PIF consolidate its AI plans across HUMAIN and Aramco, with the former acquiring talent,

Reports say HUMAIN is in preliminary talks with US private-equity giants Blackstone and BlackRock to co-invest in its data centers, and Amin is also seeking equity partners among major global data centre operators, per an interview with The Financial Times.

It is also expected to begin a process of consolidation in the Saudi Arabian and wider Gulf region markets. The Tech Capital has learnt of

plans for HUMAIN to roll up existing operators to acquire the skills, experience and landbank of already established operators.

According to sources who preferred to remain anonymous, the first of these is Quantum Switch, which has two 9MW data centres in Dammam.

At the same time, Amin has said that HUMAIN will launch a US$10 billion "HUMAIN Ventures" fund this year to invest in AI startups worldwide.

HUMAIN's governance underscores the initiative's national importance.

Its board is chaired by the Crown Prince and includes senior ministers for communications and the economy, the PIF governor and technology executives.

Amin is a well-recognised tech CEO with local experience, having previously led Aramco digital in the Gulf and Japanese media and telco disruptor Rakuten.

RED CARD FOR LAG

HOW MOROCCO IS KICKING LATENCY OUT OF THE GAME

by João Marques Lima

WHILE 2030 FIFA WORLD CUP FANS CHEER ACROSS CASABLANCA AND TANGIER, MOROCCO MAY BE PULLING OFF A BIGGER WIN BEHIND THE SCENES - USING HYPERSCALE DATA CENTRES AND GREEN ENERGY TO REWRITE AFRICA'S DIGITAL PLAYBOOK. JOÃO MARQUES LIMA FINDS OUT.

When the ball rolls in Casablanca, Rabat, and Tangier for the 2030 FIFA World Cup, Morocco will stand before the world as one of football's grandest stages. The cameras will focus on stadiums, crowds, and the shimmering new trains and highways connecting them.

Yet far from the spotlight, behind the walled perimeters of high-security campuses outside Casablanca and Settat, a quieter form of construction is reshaping Morocco's economic destiny. These are the country's new data centres - the digital foundations of a future economy that may, in time, prove even more transformative than the World Cup itself.

Across Africa, a continent long associated with infrastructure deficits rather than surpluses, data centre investment is accelerating at a pace few anticipated even five years ago. Research from Mordor Intelligence values Africa's hyperscale data centre segment at nearly US$4.9 billion by 2025, rising to US$15 billion by 2031, driven by surging demand for cloud storage, digital services, and artificial intelligence.

According to IJGlobal, more than US$3 billion in new projects were announced or under construction across sub-Saharan Africa in just the past 24 months. In this race to power the continent's digital leap, Morocco has quietly emerged as one of Africa's most strategically placed contenders - geographically, economically, and politically.

A Northern Gateway to the Cloud

For decades, South Africa dominated Africa's data infrastructure map, controlling more than 60% of installed capacity. Johannesburg was the gravitational centre for cloud investment. But the rise of hyperscalers such as Amazon Web Services, Microsoft Azure, Google Cloud, and Oracle is shifting the landscape. Geography, proximity to major exchange points, and low-latency networks now matter as much as legacy scale.

Few African markets meet this test as well as Morocco. Situated at the confluence of Europe, the Middle East, and West Africa, the country benefits from undersea fibre connections that link it directly to France, Spain, and Portugal. These routes place Moroccan data hubs mere milliseconds from European exchange points, giving them a latency advantage few competitors can match.

This proximity is particularly critical as the continent's internet traffic grows exponentially, and African businesses increasingly rely on cloud services to support e-commerce, fintech, and AI applications.

The strategic positioning has not gone unnoticed by cloud operators. AWS recently partnered with Orange Middle East & Africa to expand services in Morocco and Senegal.

"Customers of all sizes and all industries in Morocco and Senegal will be able to access local AWS compute and storage for data residency, low latency, and security needs for applications," says AWS Vice President Jan Hofmeyr.

The announcement marks a broader push by hyperscalers to bring services closer to users, particularly as AI, fintech, and digital services drive demand for near-instant access.

The contrast between Morocco's World Cup stadiums and the server halls powering the country's digital future could not be more striking. Large, high-security campuses outside Casablanca and Settat house rows upon rows of servers, air-conditioning units, and redundant power systems, humming continuously under the desert sun.
Power reliability remains a challenge across Africa, but Morocco's commitment to renewable energy and grid stability gives its operators a rare advantage. According to the International Renewable Energy Agency (IRENA), Morocco had 4.105 GW of renewable capacity in 2023, ranking among the top five in Africa.

The government aims to reach 8.5 GW through new photovoltaic farms and energy storage projects, with the ability to transfer 4.2 GW of solar energy to Europe via subsea cables.
These renewable initiatives are critical for hyperscale computing, where electricity can account for up to 60% of operating costs. Low-carbon power also aligns with ESG mandates increasingly demanded by international investors, particularly for AI-intensive workloads.

Subsea interconnections from Morocco to southern Europe could add up to 23 GW of renewable power, ensuring that the growth of digital infrastructure does not come at the expense of sustainability.

In addition to energy considerations, the physical design of Morocco's data centres reflects global best practices. Facilities combine four essential infrastructure pillars: IT systems (servers, storage, and networking), electrical systems (power supply and backup generators), mechanical systems (cooling and ventilation), and robust building infrastructure to protect against environmental risks. This multi-layered design ensures operational resilience, enabling Moroccan centres to maintain service levels comparable to facilities in Europe or North America.

AI, Investment, and Market Context

Artificial intelligence is reshaping the global data infrastructure market. By 2030, spending on AI-related compute and storage is projected to far exceed US$500 billion, with hyperscalers controlling more than 60% of global capacity. Africa's share remains small but is growing rapidly, driven by demand for cloud-based AI services, fintech applications, and other digital services that require low-latency access.

In Morocco, the AI revolution is materialising through ambitious projects such as the Naver-Nvidia Nexus Core Systems-Lloyds Capital consortium, which announced plans to build a 500 MW AI data centre. The facility is designed as a sovereign AI hub for Europe, the Middle East, and Africa (EMEA), powered entirely by renewable energy.

Its location - just 15 kilometres from Europe and linked via multiple subsea cables - emphasizes Morocco's geographic advantage. Engineered for GPU-dense, high-performance AI workloads, the project positions Morocco as a potential hub for the next generation of AI services on the continent, signalling Africa's shift toward infrastructure designed specifically for AI rather than repurposed legacy data centres.

The Moroccan data centre market is among the fastest growing in Africa, expected to expand at 6% annually through 2026, with total investments reaching US$328 million, according to the Kingdom's Competition Council. Supported by the government's digital transformation strategy, sectors ranging from banking to telecommunications and public administration are driving demand for modern, secure data infrastructure.

Key Moroccan players - including Maroc Telecom, Inwi, Medasys, and N+One - are helping establish hyperscale and sovereign cloud facilities, often through joint ventures with international companies such as Zircom, DXC Technology, and Orange. Beyond infrastructure, digital transformation is projected to positively impact six million jobs, highlighting the sector's broader socio-economic significance.

To fully appreciate Morocco's rise, it must be seen alongside other African markets. West Africa, for example, is expanding rapidly: Nigeria, Ghana, Côte d'Ivoire, and Senegal together have around 120 MW of planned or under-construction Tier III data centre capacity through 2030.

Operators like Africa Data Centres, Raxio, PAIX, and MainOne/Equinix are active across the region, focusing largely on Lagos, Accra, and Abidjan. South Africa remains the continent's historical leader, with Johannesburg and Cape Town anchoring most installed capacity, while Egypt's Smart Village and Telecom Egypt facilities are expanding the North African footprint.

Morocco distinguishes itself with a combination of geographic proximity to Europe, renewable energy, regulatory clarity, and open - access fibre networks, which together reduce latency, energy costs, and compliance risk. Its legal framework ensures data sovereignty, mandating that sensitive information be stored within national borders.

The National Commission for the Control of Personal Data Protection (CNDP) enforces these standards and evaluates foreign countries for adequate protection, supporting the growth of sovereign cloud services and reassuring both local and international investors.

Energy policy is closely linked. Morocco currently derives nearly 40% of electricity from renewables, with the government targeting 52% by 2030. Large-scale solar and wind projects not only supply the domestic grid but also enable cross-border power transfer to Europe via subsea cables. This combination of energy reliability, renewable sourcing, and regulatory clarity strengthens Morocco's appeal for capital-intensive data centre investments.

Africa's digital infrastructure opportunity is enormous and between now and 2030, cumulative investment requirements for "cloud and

Population: Approximately 37 million (2025 estimate).

Capital: Rabat (administrative capital), while Casablanca is the largest city and economic hub.

Landmass: About 446,550 km², making it the 57th largest country in the world.

GDP: Roughly US$150–160 billion (nominal GDP, 2024 estimate).

GDP per capita: Around US$4,000–4,500 (nominal).

Internet penetration: Approximately 70% of the population uses the internet (2024 estimate).

Mobile users: Around 50 million mobile subscriptions, exceeding 100% penetration due to multiple SIM ownership.

Official languages: Arabic and Berber (Amazigh); French widely used in business and government.

Head of state: King Mohammed VI (reigning since 1999).

Major industries: Tourism, agriculture, phosphate mining, automotive, aerospace, renewable energy, and IT/digital services.

AI readiness" are estimated at over US$50 billion. In West Africa alone, data centre capex is projected to exceed US$1 billion, with an additional US$2.5 billion needed for complementary IT infrastructure.

Morocco's stability, renewable energy profile, and proximity to European capital markets make financing easier. Development institutions such as the International Finance Corporation (IFC) and Proparco have already backed Moroccan digital projects, and ESG-linked infrastructure funds are increasingly participating.

These flows, combined with private-sector deals like AWS and the Naver-Nvidia consortium, illustrate that the market is maturing and capable of supporting hyperscale deployments.

Furthermore, connectivity underpins Morocco's edge. The country is linked to 12 undersea cables, including 2Africa and the Atlas Offshore system, providing redundancy and direct access to European IXPs.

This compares favourably to West African markets, which, despite improvements, still rely on limited terrestrial fibre networks.

Senegal and Côte d'Ivoire have seen significant fibre rollout over the past decade, yet access remains concentrated among a few operators, driving up costs. Morocco's liberalised, open-access model allows multiple operators and hyperscalers to interconnect freely, fostering competition and resilience.

Beyond Borders

While Morocco is gaining attention, the African data centre market is competitive and diverse. South Africa remains strong, Egypt is building AI and hyperscale capacity in Smart Village, and Kenya's Nairobi corridor is emerging as an East African tech hub. In West Africa, MainOne, Africa Data Centres, and Raxio are growing rapidly.

Morocco's advantage lies in its combination of geography, renewable

energy, and regulatory clarity, making it appealing for projects that require cross-border reach and sustainable energy sourcing. Its approach - modular, scalable, AI-ready campuses - could serve as a model for other African nations seeking to attract high-quality infrastructure investment.

Africa's digital future is continental. Morocco's approach - leveraging geographic proximity to Europe, renewables, and open-access fibre - highlights the importance of cross-border connectivity. Subsea cables and renewable energy interconnections are enabling power and data to flow north and south, creating a more integrated EMEA digital ecosystem.

Challenges remain: power reliability in West Africa is fragile, energy costs can exceed US$0.20 per kWh, and permitting delays slow new projects. Yet governments across the continent are rolling out national digital strategies, AI policies, and sovereign cloud frameworks, often accompanied by incentives for renewable-powered infrastructure.

The 2030 FIFA World Cup offers Morocco a public stage, but the digital infrastructure growing alongside it may prove even more transformative. The same networks that will stream 8K live matches will carry cloud services, AI computations, and e-commerce transactions, linking the country's digital future to its physical infrastructure. Just as sport requires long-term planning, coordination, and investment, so too does digital infrastructure.

Morocco's combination of renewable energy, subsea connectivity, regulatory clarity, and AI-ready data centres demonstrates that Africa's next growth story can be powered by computation rather than extraction.

When the final whistle blows in 2030, Morocco's true victory may not be counted in goals scored, but in gigawatts built and terabytes processed-a testament to the continent's evolving digital ambition.

NEWS

FROM OXAGON TO RIYADH

INSIDE SAUDI'S DATA CENTRE BUILD FRENZY

by Khamila Mulia

Rajit Nanda, DataVolt CEO

In a desert kingdom racing to lead the AI revolution, DataVolt's CEO Rajit Nanda reveals the blockbuster plan: harnessing the sea, sun, and sand to power a digital future - one cooled by seawater and fueled by green hydrogen. By Khamila Mulia.

S audi Arabia is expanding its role in the data and digital economy. The country couples abundant, low-cost energy with improving network connectivity, which can "reduce the cost of operating large-scale compute.

At the centre of this push is DataVolt, a Riyadh based sustainable data centre operator and a wholly owned subsidiary of Vision Invest, the Saudi development and investment holding company. Earlier this year, DataVolt partnered with NEOM to develop a net zero AI factory campus in Oxagon.

"As of today, we are on track to meet the 2028 milestone for Phase 1 of a 1.5 gigawatt campus. The AI factory will be built in three phases; 300 MW, 600 MW, and 600 MW," DataVolt CEO Rajit Nanda tells *The Tech Capital*.

He adds that Phases 2 and 3 are targeted for completion in 2030 and 2032 respectively. The first phase involves an initial investment of about US$5 billion for the energy and data centre infrastructure, with plans for an additional US$10 billion of compute infrastructure.

Saudi's geoeconomic position is relevant to compute economics, Nanda believes. The kingdom combines relatively low-cost power with improving international connectivity on Red Sea routes that link Europe, Africa and Asia.

That mix supports large-scale data centre builds and can lower delivered compute costs. Policymakers have expanded renewables through competitive tenders, while large green hydrogen and grid projects are advancing, which together can improve long term power availability for high density workloads.

According to Data Center Maps, the country has at least 51 facilities from 29 local, regional, and global operators including Saudi Telecom Company (STC), Khazna Data Centres, Equinix, Amazon's AWS, Google Cloud, and so forth.

Project details

DataVolt's Oxagon facilities are on the Red Sea coast (and form part of NEOM's planned industrial city), with access to three subsea cables that provide high-capacity fibre connectivity. This position enables reach into Africa, Asia, and Europe. According to Nanda, workloads can be transferred from Saudi Arabia with end-to-end latencies of about 15 milliseconds to Africa and around 35 milliseconds to Europe.

The project uses a mix of cooling methods, combining seawater and liquid cooling for high density deployments.

On energy, DataVolt plans to operate on renewables on a 24 by 7 basis, with green hydrogen as secondary source, delivered through a, co-located setup rather than the Saudi National grid. The company is securing supply through PPAs, including advanced discussions with the NEOM utility, and expects additional resilience from a nearby green hydrogen facility described as one of the world's largest, about one kilometre from the data centre.

The combination of renewables and green hydrogen targets high energy availability, with hydrogen fuelling gas turbines when renewable output is insufficient.

Oxagon is a planned floating industrial city within Saudi Arabia's wider NEOM project. Announced in 2017 as part of Vision 2030, NEOM aims to diversify the economy through new urban and industrial developments.
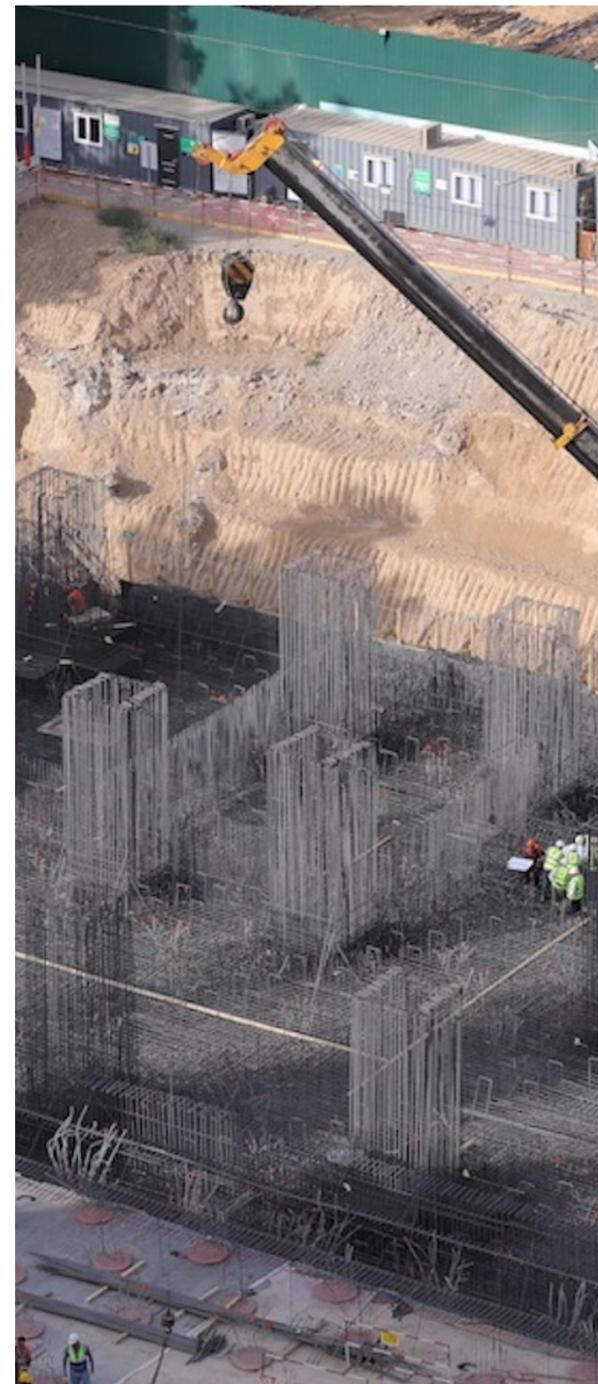
It includes several planned cities, among them The Line, a linear, car free, carbon neutral urban area envisaged at about 170 kilometres long, 500 metres high and 200 metres wide. Coverage of NEOM's progress has been mixed, with scepticism and criticism alongside periodic progress announcements. However, Nanda says Oxagon is moving forward.

"I cannot comment on NEOM overall, as it comprises multiple projects and I do not have full visibility. I do, however, have visibility on Oxagon, which is progressing well. Our extended group is building a green hydrogen facility valued at about US$9 billion; it is roughly 85% complete and expected to begin production next year, with around 20,000 workers on site," he says.

Nanda adds that three other factories are under advanced development and shortly targeting construction. On the renewables side, a wind turbine facility is due to start in about four months, and both a solar PV module plant and a solar ingot and wafer plant are moving into construction. He estimates Oxagon is attracting approximately US$30 billion to US$40 billion of advanced manufacturing investment, including DataVolt's next generation data centre.

Saudi and beyond

DataVolt is also building an AI ready data centre in Riyadh: the first phase will be 16MW with power available for a further two expansion phases in the near future. The facility will use grid power and a mix of cooling methods, primarily liquid cooling with some air cooling. It is designed for high density racks to support AI workloads.

Full scale construction of the first 16 MW phase is under way, with completion targeted for December 2026, according to Nanda. "We are discussing with multiple customers for the site, and we hope to announce customers for the first phase in the coming months," he says.

The company is also pursuing international expansion. In September, it signed an agreement with the Ministry of Digital Technologies of Uzbekistan and IT Park to establish a Tier 3+ data centre with an IT load of 12MW that will operate on renewable energy sources and be integrated with a Battery Energy Storage System.

Nanda says the move is intended to support the next phase of development and could attract hyperscalers, noting Uzbekistan's growing role in Central Asia's digital market.

"Another market that we continue to look at is Africa. We want to create a connectivity corridor between the Kingdom of Saudi Arabia and Africa in order to provide compute capacity," Nanda says. He adds that much of the computing would be carried out in Saudi Arabia, with results delivered to African markets over existing subsea cables.

Looking Ahead

Entering 2026, the Middle East's data centre market is becoming more exciting, with domestic demand to grow fastest in Saudi Arabia and the United Arab Emirates.

In Saudi, Nanda estimates about 300 MW of capacity is built or under construction today, rising to roughly 600 to 750 MW by 2030, largely driven by cloud requirements. In parallel, he sees a separate case for multiple gigawatts of AI workloads located in the kingdom to serve global demand.

DataVolt plans to focus on executing the strategic projects identified over the past twelve to eighteen months, with an emphasis on capacity built for AI workloads. "We want to make the kingdom an important hub for training and inferencing workloads. The priority is to deliver large scale, high efficiency data centres at speed to meet rising compute needs," Nanda says.

THE BOTTOM LINE

MIDDLE EASTERN NATIONS STAKE THEIR CLAIM IN THE AI ERA

by João Marques Lima

The Middle East is fast becoming one of the world's hottest battlegrounds for digital infrastructure dominance. Once seen primarily as an oil and logistics hub, the region is now rewriting its playbook - pouring billions into hyperscale data centres, sovereign AI initiatives, and global technology alliances. The result: a powerful new nexus in the global digital economy where artificial intelligence (AI) workloads are redefining how infrastructure is designed, financed, and deployed.

UAE and Saudi Arabia Lead the Charge

The United Arab Emirates stands at the epicentre of this transformation. Dubai's established players - like Khazna and Equinix - are expanding rapidly to meet surging enterprise and hyperscale demand, while Abu Dhabi's G42 is anchoring a new wave of AI-focused infrastructure. Together, they're turning the UAE into a dual-core digital economy, underpinned by massive capacity and state-backed innovation.

Saudi Arabia is following a similar trajectory, with Riyadh emerging as a regional data hub and Dammam leveraging its energy-sector ties to attract data-intensive industries. The Kingdom's strategy is clear: spread computing power across its economic corridors while fostering partnerships like AWS and HUMAIN's US$5 billion AI collaboration - one of the largest in the region.

The Gulf's smaller nations are not sitting still. Qatar, through MEEZA and Ooredoo, has cultivated a strong, investment-backed ecosystem. Bahrain, meanwhile, is carving a niche as a secure haven for financial and sovereign cloud workloads. Oman's Sohar and Barka are building state-led edge and connectivity projects, broadening access to distributed computing.

These emerging nodes reflect a strategic shift: smaller markets are specializing, while larger ones scale out, creating a complementary digital fabric across the region.

AI Takes Centre Stage

The next big leap? Sovereign AI.

Governments across the region are betting that local control of AI infrastructure will define their digital futures. The UAE's Stargate UAE project - a one-gigawatt AI data centre cluster - is the clearest sign yet. The first 200-megawatt phase comes online in 2026, part of a vision that could scale to five gigawatts.

Oman's Digital Triangle initiative and Saudi Arabia's AI-driven partnerships follow the same pattern: domestically anchored, globally connected, and strategically aligned with national development goals.

The Stargate concept - originally a U.S. initiative linking AI and supercomputing infrastructure - has found fertile ground in the Gulf. In 2025, the UAE and the U.S. launched a joint AI campus spanning 10 square miles, backed by G42, OpenAI, Oracle, NVIDIA, Cisco, and SoftBank. It's the first global expansion of the model and could reshape international AI capacity distribution.

The UAE's strategic position - within a 2,000-mile radius of half the world's population - turns this into more than a domestic project; it's a global network node connecting Europe, Africa, and Asia.

What's striking is how quickly regional operators are evolving. No longer content with simply providing power and racks, they're building AI factories - integrated ecosystems for training, experimentation, and enterprise AI adoption. The goal: become not just infrastructure providers, but technology creators.

As GPU architectures shift and multi-cloud demands intensify, Middle Eastern operators are adopting multi-tenant platforms and self-service AI models - mirroring cloud giants while keeping operations sovereign and secure.

The Middle East's digital rise isn't a story of catch-up - it's one of leapfrogging. With sovereign AI strategies, hyperscale infrastructure, and deep-pocketed investment funds, the region is positioning itself as an indispensable pillar of the global AI economy.

The oil era built its wealth; the data era may define its legacy.

Alisha Mehta Adams
Partner, Jackson Walker

Alisha Mehta Adams is a Partner at Jackson Walker, where she focuses on environmental, regulatory and digital infrastructure advisory. She counsels data centre developers, operators and investors on permitting, sustainability compliance and operational risk across large-scale facilities. Her work integrates site selection, power and water strategy, and lifecycle regulatory risk for the digital infrastructure sector in the United States, enabling efficient advancement of data-centre projects in a complex regulatory environment.

Amanda Ludlow
Partner, Ashurst

Amanda Ludlow is a Partner and Global Co-Head of Ashurst's Digital Economy practice, based in London. She advises on major international digital infrastructure projects including subsea fibre, fibre networks, towers, data centre investments and telecoms roll-outs. Her standout engagements include the creation of Australia's national broadband provider and Meta's 2Africa subsea cable project. She has also advised nbn on its acquisition of Telstra's HFC network She guides clients through M&A, joint ventures and regulatory strategy across the TMT sector.

Ari Fritzerald
Partner, Hogan Lovells

Ari Q. Fitzgerald is a Partner in Hogan Lovells' Washington, D.C. office and leads the Communications, Internet & Media practice. He provides strategic regulatory, policy and transactional advice in telecoms, spectrum and digital infrastructure sectors, including acting for major funds and operators on data centre acquisitions. His prior government experience strengthens his counsel on complex global connectivity, digital infrastructure and communications-policy matters.

Benjamin Heriaud
Partner, Vinson & Elkins

Benjamin N. Heriaud is a Partner in the Capital Markets & M&A group at Vinson & Elkins in New York. He advises corporates, private equity sponsors and investment banks on capital raising, acquisition-financing and high-yield debt offerings, with a focus on technology and industrial-sector clients. His recent work includes advising Coatue as the lead investor in a US$1.1 billion Series C funding round in CoreWeave, Inc., a specialised GPU cloud provider, and advising on Switch, Inc.'s US$600 million debut offering of senior notes and its US$500 million acquisition financing. His experience positions him well for TMT-adjacent infrastructure financings.

Brittney Sandler
Counsel, Pillsbury Winthrop Shaw Pittman

Brittney D. Sandler is Counsel at Pillsbury in Washington, D.C. and co-leads the firm's data centres team. She advises on commercial and technology transactions including data centre development, cloud/SaaS outsourcing, tenancy and colocation agreements for financial services and tech industry clients. Her work spans contract structuring and operational risk mitigation for digital infrastructure projects. Her practice is well recognised in the data centre infrastructure space.

Christopher Fisher
Former Managing Partner and Chair, Digital Infrastructure Practice

Christopher B. Fisher is former Managing Partner and currently Chair of the Digital Infrastructure Practice at Cuddy & Feder. He practices out of New York and focuses on land use, zoning, permitting, commercial real estate and telecoms infrastructure, including wireless carriers, tower and small-cell deployments. His recent work includes precedent-setting zoning matters and digital infrastructure siting across New York and Connecticut, aiding carriers and infrastructure operators in complex deployments.

Barry Gross
Partner, Simmons & Simmons

Barry Gross is a Partner at Simmons & Simmons, specialising in data centres and real estate-intensive businesses. He advises developers, investors and operators on high-value data centre transactions, including development projects, leasing, M&A and financing across multiple jurisdictions. Notable work includes cross-border acquisitions and service agreements for hyperscale and enterprise clients. He also advises on restructurings and real estate-intensive enforcement matters, co-leading Simmons' multidisciplinary data centre team.

Ben Henry
Partner, Troutman Pepper Locke

Ben Henry is a Partner at Troutman Pepper Locke with offices in Chicago and London, specialising in the international data centre and digital infrastructure sector. He has advised on over 1,500 transactions in 80 jurisdictions, including a major sale of eight European data centres by Colt Technology Services to NorthDC. He has also facilitated on behalf of a Singaporean sovereign wealth fund the acquisition and development of a Dublin site leased to a hyperscale cloud company, marking a significant investment in cloud infrastructure. His practice supports developers, end-users and investors on cross-border operational, real-estate and regulatory matters in the data centre ecosystem.

Connen O'Neill
Partner, Gibson Dunn

Connell O'Neill is a Registered Foreign Lawyer (NSW) of Gibson, Dunn & Crutcher (Hong Kong) and Partner of Gibson, Dunn & Crutcher LLP, leading the firm's Asia TMT & Technology practices. He advises on digital infrastructure M&A, towers, data centres, subsea cables, fibre and regulatory matters across Asia and the Middle East. His recent engagements include representing e& on its US$2.2 billion divestment of a 40% stake in Khazna Data Center in 2025, guided PLDT on its 2024 network modernisation with Amdocs' cloud-based OSS platform, and acted for Metro Pacific Tollways Corporation on its multi-lane free flow tolling program contracts with Egis in 2024.

Conrad Andersen
Partner, Latham & Watkins

Conrad Andersen is a Partner in the London infrastructure practice of Latham & Watkins and Global Vice-Chair of the firm's Energy & Infrastructure Industry Group. He advises infrastructure funds, investors and financial institutions on acquisition and financing transactions, including digital-infrastructure. His deal work includes advising on the acquisition and financing of six European data centres by STACK Infrastructure, covering Germany and the Netherlands, as well as several fibre transactions for Gigaclear, EuNetworks, and CityFibre.

Conrad Purcell
Partner, Haynes Boone

Conrad Purcell is a Partner in Haynes Boone's London office, specialising in energy and infrastructure projects across EMEA, including digital infrastructure. He led the firm's advice to a pan-African data centre developer on a US$100 million debt-financing package from the IFC to support data centre expansion in sub-Saharan Africa. He regularly advises sponsors, lenders and developers on finance, commercial agreements and project-deployment strategy in telecoms and digital-infrastructure.

Daniel Campbell
Partner, Hunton

Daniel M. Campbell is a Partner at Hunton in Richmond, whose practice focuses on commercial real estate transactions for technology companies and investors, including development, acquisition and financing of data centre facilities. He supports clients in deploying new-build digital infrastructure campuses, leveraging land-assets and executing real estate strategy in the TMT sector, guiding site-selection, funding and transactional structuring for hyperscale and colocation developments.

David Coulling
Partner, Herbert Smith Freehills Kramer

Since joining Herbert Smith Freehills Kramer in September 2003, David Coulling has worked on a range of transactions for major technology companies. He advised VMO2 on its mobile network deployment and sale of its stake in UK towerco Cornerstone, and AustralianSuper on its €1.5 billion investment in Vantage Data Centers EMEA. In the satellite space he worked with Bharti Enterprises on its bailout of OneWeb and has since advised OneWeb on numerous issues including its secondary listing in London, and its acquisition by Eutelsat.

David Penna
Partner, Davis Polk

David Penna is the co-head of Infrastructure Finance, advising on complex project development, financing and M&A transactions. He focuses on the energy, power and infrastructure sectors. His practice encompasses the development and financing of data centres, towers, fibre, wind, solar, battery storage, biomass and power transmission projects, as well as natural gas, oil, LNG and petrochemicals projects.

Dino Barajas
Partner, Baker Botts

Dino E. Barajas is a partner and the firmwide Chair of the Project Finance Practice Group. He specialises in domestic and international project development and finance, with a focus on Latin American infrastructure, debt financings, and mergers and acquisitions. Dino advises lenders, investors, and developers across the energy, power, and infrastructure sectors. His clients include commercial lenders, investment funds, and project sponsors, and he has led transactions across Latin America, the United States, the Caribbean, and the Middle East.

Filip Van Elsen
Partner, A&O Shearman

Filip Van Elsen is Technology Sector Lead. He advises leading corporations, private equity funds, technology companies and emerging growth companies on strategic technology transactions. He has particular experience in assisting established and emerging companies on digital transformation projects in the pharmaceutical, banking, energy, mining and digital infrastructure sectors.

David Hart
Partner, Baker McKenzie

David is a partner in Baker McKenzie's corporate group, based in London. He joined the Firm as a trainee in 2007 and became a partner in 2019. He is a member of the Firm's TMT Industry Group, where he advises corporate and financial sponsor clients on cross-border mergers and acquisitions, joint ventures and investments. This includes advising Equinix on the formation of its US$1 billion European hyperscale data centre JV with GIC.

David Martin
Corporate Partner, Linklaters

David Martin is a corporate partner in the London office's financial sponsors group and co-head of the firm's Telecoms, Media and Business Services sector. David regularly works on transactions in the infrastructure sector, with a particular focus on digital infrastructure. He advises on mergers and acquisitions, consortium deals, co-investments, financial sponsor exits (including equity capital markets transactions), joint ventures and carve-outs, often with a multi-jurisdictional, cross-border and competitive element. He was worked on multiple deals with DigitalBridge, including its acquisition of Yondr Group, and co-investments in Vantage Data Centers.

Gabriel Silva
Partner, Simpson Thacher & Bartlett LLP

Gabriel Silva is a Partner in the Energy and Infrastructure Practice at Simpson Thacher, based in New York, and was named "Lawyer of the Year" at the 2025 TMT Finance Awards USA. He advises infrastructure clients and private equity investors on mergers, acquisitions, dispositions, and digital infrastructure transactions, including data centres, telecom towers, and fibre networks. Gabriel has led over 60 transactions across the U.S., Latin America, and Europe in the past three years and is recognised as a leading lawyer by Law360, Chambers Global, and Legal 500.

Ibaad Hakim
Partner, Paul Hastings

Ibaad Hakim is an Abu Dhabi-based Partner in Paul Hastings' Energy and Infrastructure practice. He advises on high-value oil and gas, power, and infrastructure projects, many project-financed, covering procurement, contract drafting, negotiation, claims, and dispute resolution. His expertise spans EPC, turnkey, BOT/BOOT, DBOM, PFI and standard form contracts including FIDIC, NEC, JCT, ICE and IChemE. He is qualified in England & Wales and Pakistan.

Jacqueline Chan
Partner, Milbank LLP

Jacqueline Chan is a partner in Milbank LLP's Singapore office, formerly based in Hong Kong. She advises on international corporate finance, M&A, acquisition finance, and debt restructurings across Asia, representing sponsors, borrowers, lenders, private equity funds, sovereign wealth funds, corporates, and leading banks. Recognised by Chambers Asia, IFLR1000, Legal 500, and Best Lawyers, she has received numerous awards for banking, M&A, restructuring, and project finance. Jacqueline is consistently ranked among Asia's leading lawyers in corporate, finance, and energy sectors.

James Crooks
Partner, Sidley Austin

James Crooks specialises in leveraged finance and general debt transactions, advising private equity sponsors, corporate borrowers, credit funds, and financial institutions across mid- and large-cap financings, asset-backed deals, stressed and distressed situations, and regulated sectors. He has expertise in financial services, insurance, healthcare, pharma, technology, and sports, delivering innovative, market-leading solutions. Recognised by Chambers UK, Legal 500, IFLR1000, and The Lawyer's Hot 100, James combines market insight with commercial pragmatism on high-profile transactions.

James Dodsworth
Global Head of Real Estate, White & Case

James Dodsworth is Global Head of Real Estate at White & Case. He advises clients on investment, structuring, leasing and financing of commercial real-estate assets, with a particular specialism in the data-centre and digital-infrastructure sector. For example, he led White & Case's role advising Yondr Group on its sale to DigitalBridge. He is widely acknowledged as an authority on global data-centres and regularly speaks at industry events.

Jeeseon Ahn
Partner, Bracewell

Jeeseon Ahn advises financial institutions and sponsors on domestic and international project and structured financings, investments, and transactions across the energy and infrastructure sectors, including renewables, conventional power, oil and gas, mining, data centres, and digital infrastructure. She has significant experience representing export credit agencies and multilateral development banks on cross-border infrastructure projects, particularly in Latin America and Asia.

Jessica Melville
Partner, Norton Rose Fulbright

Jessica Melville is a project and structured finance lawyer based in London. She focuses on advising parties on the financing of energy and infrastructure projects in emerging markets.
Jessica advises commercial lenders, development finance institutions, sponsors and governments on the various forms of lending, including advising on complex multisourced financings. Jessica commenced her career in the Norton Rose Fulbright Perth office before transferring to London in 2016.

Jitendra Soni
Partner, Argus Partners

Jitendra Soni is a Partner in Argus Partners' corporate and technology practice, based in Bengaluru, leading teams across Bengaluru and Delhi. He advises clients on technology, fintech, e-commerce, data protection, cloud, blockchain, NFTs, AR/VR, and emerging digital sectors. Jitendra represents startups, investors and corporates in acquisitions, joint ventures, private equity and venture capital, including Amazon, Arkam Ventures, iDO Devices, and Infinx, as well as cross-border data governance and regulatory compliance matters.

James W Grice
Chair, Data Centers and Digital Infrastructure Practice, Akerman

James W. Grice is Chair of Akerman's Data Centres and Digital Infrastructure Practice, advising investors, developers and operators on data centre acquisitions, dispositions, development and leasing. He has led transactions including Mapletree's $1.3 billion portfolio purchase from Sila Realty Trust and strategic investments for Macquarie Capital USA and Landmark Dividend. Grice also guides clients on power procurement, tax incentives, renewable energy, and regulatory matters, supporting hyperscale and enterprise data centre projects across North America and Europe.

Jane Reuger
Partner, Perkins Coie

Jane Rueger helps clients negotiate power purchase agreements (virtual and physical), gas purchase and sale agreements, mergers and acquisitions transactions, project financings, and interconnection agreements. Jane's work includes representing developers and operators of data centers in the negotiation of multifaceted transactions combining power infrastructure and real estate. She also counsels market participants on creative solutions to participation in complex regional transmission organization (RTO) and independent system operator (ISO) markets across the United States.

Jocelyn Paulley
Partner, Gowling WLG

Jocelyn Paulley is a Partner at Gowling WLG, specialising in technology and data law. She advises on IT agreements, data protection, data centres and telecommunications. As co-lead of the firm's Data Protection and Cyber Security team, she supports clients across diverse sectors on privacy compliance and digital transformation. Known for her pragmatic and commercial approach, Jocelyn helps clients navigate complex technology, cloud and e-commerce arrangements with confidence.

John Creedon
Partner, Ropes & Gray International

John M. Creedon is global co-head of Ropes & Gray's real estate investments and transactions group and co-lead of its infrastructure industry group. He advises private equity funds, hedge funds, REITs, endowments, and institutional investors on joint ventures, co-investments, debt financings, distressed acquisitions, leasing, and development projects. Recognised by Chambers USA, Legal 500, and Real Estate Forum's 50 Under 40, John has extensive domestic and international experience and supports pro bono real estate initiatives with Pine Street Inn.

Justin Haan
Partner, Morrison Foerster

Justin Haan is a Partner in Morrison Foerster's Technology Transactions Group in San Francisco. He advises technology industry leaders on intellectual property and technology matters, including licensing, technology transfers, joint development, and M&A transactions. Justin's work spans AI, machine learning, cloud computing, semiconductors, autonomous vehicles, and data centres. He structures and negotiates commercial agreements, software and content licences, service arrangements, and collaboration, manufacturing, and distribution contracts.

Kemal Hawa
Partner, Kirkland & Ellis

Kemal Hawa is a corporate Partner in Kirkland & Ellis' Washington, D.C., office, advising leading private equity firms, infrastructure funds and real estate investors on digital infrastructure transactions globally. He specialises in M&A, financing, joint ventures, and development projects for data centre operators, telecoms carriers, cloud providers, fibre and submarine cable companies. He represented U.S. Internet Corp., a Minnesota-based fiber and wireless provider, in its sale to Metronet.

Marco Nicolai
Partner, Bird & Bird

Marco Nicolai is a Partner in Bird & Bird's Real Estate Group and leads the UK Construction Group in London. He advises on large-scale property and infrastructure projects in the UK and internationally, with particular expertise in data centres, energy infrastructure and complex developments. Drawing on extensive UK and Middle East experience, Marco acts for investors, developers, lenders and public bodies on all aspects of construction procurement and delivery.

Matt Kirsch
Partner, Arnold & Porter

Matt Kirsch concentrates his practice on the acquisition, financing, ground leasing, development, and sale of real property. As part of Matt's data centre work, he has provided counsel to various clients related to the development of hyperscale data centre campuses across the US and internationally, aggregating over 5,000MW of power capacity, in various stages of planning, construction and operation. As part of these projects, Matt leads the negotiation of purchase and sale agreements, construction financing documents, joint venture arrangements, and a variety of development specific agreements.

Michael Pearson
Partner, Clifford Chance

Michael Pearson is a Partner in Clifford Chance's Energy and Infrastructure Group and Co-Head of the Infrastructure sector. He advises on major infrastructure projects in the UK and internationally, specialising in concession arrangements and project finance structures. Michael's work spans public-private partnerships, complex development projects and cross-border financing, supporting investors, sponsors, lenders and government bodies on structuring, documentation and execution of high-value infrastructure transactions.

Mike Conradi
Partner, DLA Piper

Mike Conradi is co-chair of the international telecoms practice, advising governments, regulators, and businesses on regulatory and commercial matters. Recognised as a leading telecoms lawyer worldwide, he is praised by Chambers for his strong market knowledge, negotiation skills, and ability to navigate complex regulatory environments. Mike combines deep sector expertise with strategic insight to support clients across global markets. He holds a degree in Politics, Philosophy and Economics from Oxford University.

Marcus Bagnall
Partner, Wiggin

Marcus Bagnall is a telecoms and technology specialist with over 15 years' experience advising TMT sector clients and sponsors on the commercial and regulatory aspects of major digital infrastructure and industry digitalisation projects. He advises across Europe, Australia and the Middle East on complex telecommunications, technology and next-generation network arrangements. Qualified in both England & Wales and New South Wales, Marcus holds an LLM from the University of New South Wales and previously worked with Linklaters in London and Webb Henderson in Sydney.

Maria-Leticia Ossa Daza
Partner, Paul, Weiss, Rifkind, Wharton & Garrison

Maria-Leticia Ossa Daza is head of the Latin America Practice and a partner in the Mergers & Acquisitions Group, advising clients on complex cross-border corporate transactions for over 20 years. She has extensive experience in financial institutions, technology, infrastructure, and consumer goods, representing international and domestic clients. A native of Colombia, Leticia speaks English, French, Italian, Spanish, and Portuguese. She has been repeatedly recognised among Latin America's leading corporate lawyers and actively participates in professional and leadership networks.

Mitch Holzrichter
Partner, Mayer Brown

Mitch Holzrichter is a Partner at Mayer Brown, advising investors, sponsors, contractors and lenders on complex infrastructure and digital infrastructure transactions. He has represented clients on equity investments, financings, acquisitions, and design-build agreements, including Macquarie Capital's investment in Prime Data Centers, and the acquisition of a Virginia-based data centre operator with development rights. His work spans airports, toll roads, vertiports, and revenue-risk concession projects globally.

Nancy Lamb
Head of Data Centers - Construction Solicitor, arch.law

Nancy Lamb is a leading expert in data centre construction contracts and dispute resolution, known for her straight-talking, no-nonsense approach. She started her career as a construction solicitor at international law firms Pinsent Masons LLP and Hill Dickinson LLP, in the UK and the Middle East before immersing herself into the data centre sector working for both contractors and clients. Nancy also works in the pharmaceutical and residential and commercial property sectors. She is also the lead consultant on the Conexus Law collaborative contracting service.

Patrick Horan
Partner, Arthur Cox

Patrick Horan is a Partner in Arthur Cox's Competition and Regulated Markets Group, advising on EU and Irish competition law, foreign investment screening, and sectoral regulation. He counsels clients on merger control, investigations, State aid, compliance programmes, and competition litigation, with extensive experience in telecoms, transport, and energy sectors. Patrick is recognised by Chambers Europe, Legal 500, and Who's Who Legal as a Thought Leader in competition law.

Paul Graham
Partner, Fieldfisher

Paul Graham is a Partner at Fieldfisher, recognised as a Leading Individual for TMT work by Legal 500 and Chambers. He holds a Certificate in Digital Transformation from MIT and Cybersecurity for Business Leaders from Saïd Business School, University of Oxford. With over 25 years' experience, he advises telecoms companies, energy suppliers, governmental bodies and IDNOs on network-security matters. In 2024, he advised a global telecoms company on a major ransomware attack affecting over 20 jurisdictions.

Peter Bergan
Partner, Orrick

Peter C. Bergan Jr. advises investors, developers, and owners on a broad range of real estate transactions, from large-scale commercial developments to complex infrastructure projects. With over a decade of experience, he has led joint ventures, financings, and acquisitions across commercial, industrial, retail, hospitality, and digital infrastructure sectors. Peter has expertise in data centre and digital infrastructure assets, including hyperscale and colocation facilities, and is an active member of the Urban Land Institute and the International Council of Shopping Centres.

Sean McDevitt
Partner, Arthur D. Little

Sean McDevitt is a Partner at Arthur D. Little, based in Boston, where he leads the firm's TIME (Telecommunications, Information Technology, Media and Electronics) and Private Equity practices in the US. With nearly 30 years of consulting and operational experience, he advises on M&A, market entry and spectrum strategy across the communications sector. Sean previously held senior roles at Cartesian and CSMG and holds an MBA from Northwestern University.

Sebastien Bonneau
Partner, McDermott Will & Schulte

Sebastien Bonneau is dual-qualified in the UK and France, with over 20 years' experience advising on complex cross-border M&A, private equity, joint ventures, and real estate transactions across Europe and the US. He is a recognised expert in the data centre sector, having advised on many of its most transformative deals since 2006. Sebastien also advises clients in healthcare, energy, and real estate, working with investors, developers, and corporates on strategic acquisitions, investments, and development projects.

Stefan Kilgus
Partner, Watson Farley & Williams

Dr. Stefan Kilgus is a partner in the Watson Farley & Williams project and structured finance group. He regularly advises banks and other financial institutions, institutional investors, sponsors and developers on project financing of offshore and onshore wind farms, solar parks and other energy and infrastructure projects in Germany, Europe and Latin America.

Mark Chester
Partner, Eversheds Sutherland

Mark Chester is a UK lawyer with over 20 years' experience in the data centre sector. He advises developers, owners, operators, funders, and enterprise users on a wide range of legal issues in this rapidly growing industry. With deep sector knowledge, Mark provides commercially focused guidance, helping clients navigate complex transactions and operational challenges. His experience across corporate, finance, and regulatory matters enables him to deliver practical, relevant advice that supports the development and operation of data centres underpinning the digital economy.

Priscilla (Fasoro) Combari
Partner, Covington & Burling

Priscilla Combari advises clients on complex commercial and technology-driven transactions, including software development and implementation, cloud and SaaS licences, professional services, commercial partnerships, reseller arrangements, and distribution deals. She represents clients across technology, life sciences, utilities, consumer goods, airlines, hospitality, automotive, and banking sectors. With in-house experience at a technology company, Priscilla combines legal expertise with commercial insight, delivering solutions that are legally sound, operationally practical, and strategically aligned with clients' business objectives.

Susan Booth
Partner, Holland & Knight

Susan J. Booth is a Los Angeles-based real estate attorney with nearly three decades of experience advising on acquisitions, sales, and capital markets transactions, particularly in data centres, hotels, office buildings, multifamily, industrial, and mixed-use projects. She represents domestic and international banks, pension funds, private equity, REITs, and family offices, with extensive experience in CMBS, portfolio loans, joint ventures, and loan workouts. Susan combines deep sector knowledge with strategic guidance to help clients navigate complex transactions and troubled assets.

Tim Cline
Partner, BCLP

Tim Cline is a US-based real estate lawyer with a nationwide practice across traditional and alternative asset classes, including multi-family residential, warehousing, logistics, wind and solar farms, data centres, medical offices, and high-value residential. He represents institutional investors, REITs, sovereign wealth funds, utilities, and enterprises in acquisitions, sales, leasing, and financing. Tim leads BCLP's Data Centres & Digital Infrastructure Team and is active in mentoring, diversity initiatives, and internal engagement, supporting the development of junior lawyers and promoting inclusive practices.

POWER, PRECISION, AND AI

THE NEW PRIORITIES FOR DATA CENTRE GROWTH

by João Marques Lima

Bain & Company's new forecast to 2030 suggests the global data centre industry is shifting from rapid expansion to a power-constrained, efficiency-driven phase shaped by AI workloads and regional diversification. João Marques Lima reports.

For more than a decade, the world's data centres have been built at an astonishing pace - warehouses of servers quietly underpinning the digital economy. But the next stage of that growth, Bain & Company suggests, will look very different.

In its latest forecast through 2030, the consultancy finds that the early scramble to meet surging demand for generative artificial intelligence (AI) is giving way to a more deliberate, capacity-constrained chapter. The emphasis, it says, is shifting from scale at any cost to a focus on precision, power management, and long-term execution.

"The winners will be defined not by scale alone but by their ability to navigate complexity with precision," Bain says in its report.

The analysis, built on scenario modelling and sector insights from Bain's Energy & Natural Resources and Technology practices, paints a picture of an industry settling into a new rhythm. Hyperscale technology companies - those that design and operate massive computing infrastructure - are continuing to absorb the bulk of new capacity. At the same time, large enterprises are starting to roll out production-grade AI systems, embedding the technology into daily operations rather than treating it as an experiment.

Although the rate of expansion varies by region and project, Bain's forecast indicates that overall momentum will remain strong. What is changing is how growth happens. New construction is increasingly selective, shaped by the availability of power, land, and local regulation rather than by pure demand.

Among the most significant developments Bain identifies is a shift in the balance of AI workloads. Training the largest AI models once dominated infrastructure planning; now, inference-the process of using trained models to generate results-is becoming the centre of gravity.

This change, driven by rising enterprise adoption, is altering data centre economics and architecture. The growing demand for inference means data must be processed quickly and closer to users, increasing interest in smaller, distributed networks rather than vast single-purpose facilities.

Operators are weighing whether to build proprietary sites or use colocation facilities shared with others. Silicon diversity-supporting multiple types of chips-and flexible cooling systems are also becoming strategic considerations.

Capital Efficiency Replaces Expansion Fever

Many industry watchers expected hyperscalers to rein in spending in 2025. Instead, Bain reports that their investments increased, though with a noticeable shift in focus. Rather than racing to build new sites, companies are concentrating on improving capital efficiency and optimising the assets they already control.

"The pace of new construction growth has begun to stabilise," Bain says. "Hyperscalers are focusing more on capital efficiency and getting more selective with new deployments."

This pragmatism reflects a maturing market: rapid expansion has given way to a phase where cost, sustainability, and reliability are as important as raw capacity.

Geographically speaking, North America remains the world's largest hub for data centre capacity, powered by the capital expenditure of US-based technology giants. Yet Bain notes that growth is now spreading to new regions as governments and corporations alike develop "sovereign AI" strategies-efforts to ensure local control over data and computation.

Markets in Europe, the Middle East, and parts of Asia are drawing new investment as operators seek geographic flexibility. Factors such as latency, energy sourcing, and data sovereignty are shaping decisions about where to locate the next generation of facilities.

Even as the market globalises, data centres themselves are becoming larger. Bain expects mega-campuses-sites with power capacity of at least one gigawatt-to become standard for training the most advanced AI models. Only a limited number of these large facilities will be needed worldwide to meet the most demanding workloads.

Yet for inference and other less intensive tasks, smaller and more distributed centres are proving sufficient. The result is a more complex network: a few vast hubs linked to many smaller nodes.

Power: The New Bottleneck

If GPUs were once the scarcest resource in AI infrastructure, electricity is quickly taking that role. Bain describes power availability as "the critical gatekeeper of growth". Even as supply constraints for hardware and materials ease, securing reliable and affordable energy has become the defining challenge for developers.

Behind-the-meter (BTM) power generation-where operators produce electricity on site-is increasingly shaping where and when new projects are built. These projects are most common in the United States and currently rely heavily on independent gas-powered sources.

Utilities, regulators, and developers are under pressure to coordinate as energy demand surges. In some regions, utilities are already working directly with operators to plan for large future loads.

The picture that emerges from Bain's analysis is one of transformation rather than slowdown. AI continues to drive strong global demand, but the priorities of the sector are shifting: from speed to strategy, from scale to efficiency, and from abundance to constraint.

Over the coming decade, the most successful data centre operators may not be those that build the most space, but those that manage to balance performance, sustainability, and power in a world where e very watt counts.

BEVAN
SLATTERY
ENTREPRENEURSHIP FOUNDER

PHOTO: THE TECH CAPITAL'S APAC FINANCE FORUM 2025, SEPTEMBER, SINGAPORE

FINANCE & INVESTMENT

Digital infrastructure investment is fueling the next wave of globalisation. As cloud computing, edge networks, and AI systems become more deeply integrated, capital is flowing toward the platforms that enable intelligent connectivity. Investors in this space aren't just financing real estate - they're building the foundation for smarter industries, faster economies, and a more connected world where data moves with purpose and power.

While hundred-billion-dollar deals capture the headlines, Legal & General's approach to digital infrastructure takes a more measured, but equally as important approach. Matteo Colombo tells Jack Haddon how the UK titan is turning data centres, fibre, and cloud platforms into a strategic arena where capital, sovereignty, and policy collide.

I
n early September Matteo Colombo told the world that Legal & General Asset Management (L&G) had raised US$702 million in the first close of its new digital infrastructure fund.

A notable development for one the pre-eminent asset managers in the UK, but Colombo had some stiff competition for headline grabbing news. The digital infrastructure industry has grown accustomed to talking in the hundreds of billions range in the back half of 2025, and L&G's latest accomplishment flew somewhat beneath the radar.

But perhaps it shouldn't have. The fund's design, investor composition, and geographic reach perfectly illustrate a new dynamic that is emerging in digital infrastructure investment, where sovereignty, geopolitics and government policy matter just as much as the nuts and bolts of physical infrastructure and balance sheets.

Speaking with The Tech Capital from L&G's headquarters in the City of London, Colombo's explanation of his investment philosophy demonstrates how one of the UK's largest institutions is approaching digital infrastructure as both a national interest and a global market.

Sovereignty, capital, and the role of domestic investors

Colombo clarifies that a second close of the fund is expected by the end of the year, expanding the capital that the fund will target seven to ten assets spanning data centres, fibre, towers, and cloud services is designed to spread exposure across operating models and customer bases. Approximately half of the fund's portfolio is expected to be in data centres as a result of the relative size of that market and the firm's prior balance sheet experience in the segment.

Reflecting a wider strategy of the overall company, assets that L&G has invested in from its own balance sheet are being moved into the fund, including stakes in the UK's KAO Data, and US Aligned Data Centers.

A powered land development alongside fellow UK asset manager Goldacre in Newham, London, is also being considered as an asset that could be moved to the fund (via an independent valuation and the fund essentially purchasing it from L&G's balance sheet), although Colombo insists a decision on this is not final.

This approach allows institutional investors to capitalise on the strategies that Legal & General was previously applying to its own capital.

Crucially, the fund is structured with an 80% focus on the UK and Europe and around 20% in North America and other markets. Within that allocation, the UK remains the largest target, reflecting both the scale of its data economy and policy drivers favouring domestic ownership.

Discussing the rationale behind the geographic mix, Colombo notes that the UK remains one of the world's largest data centre markets by installed capacity.

"For diversification, we're not focused on one market. But the UK is still the largest digital economy in Europe, and that's where many of our institutional investors have natural alignment," he says.

Furthermore, as data sovereignty has emerged as a strategic concern for many governments, the physical assets that support digital economies, such as data centres, fibre networks, and cloud platforms are increasingly viewed as critical national infrastructure.

Questions over who should own and operate this infrastructure are

starting to be asked with increasing frequency and volume, with domestic capital increasingly encouraged to take a larger role in financing assets that underpin national connectivity.

Colombo believes L&G's position as a FTSE 30 company places it among a limited group of UK institutions capable of meeting that expectation.

For governments and end users at enterprises, domestic ownership offers a degree of policy and operational continuity that overseas capital can find harder to achieve. Pension schemes and insurers, L&G's traditional investor base, are being encouraged to channel more of their capital into UK infrastructure as both a financial and strategic opportunity.

Public–private alignment and institutional leverage

L&G's presence as a large domestic investor offers advantages beyond capital. As one of the UK's most significant property owners, Colombo says it is able to engage in local partnerships and policy discussions that smaller or foreign funds may find harder to access in the UK.

Digital infrastructure projects, and particularly large-scale data centres, require collaboration with local authorities on planning, energy, and connectivity.

The firm's development in Newham illustrates this approach, with L&G's relationships with local councils developed through its real estate and regeneration activities in the borough, providing "value-add" to the project.
This is something that L&G expect from its co-investors as well.

Data centre projecto in London joint-funded by Legal and General

Matteo Colombo

DOMESTIC CAPITAL, GLOBAL SCALE

DIGITAL SOVEREIGNTY AS AN INVESTABLE ASSET

by Jack Haddon

Data centre project in London joint-funded by Legal & General and Goldacre

Kao Data's Harlow campus backed by Legal & General
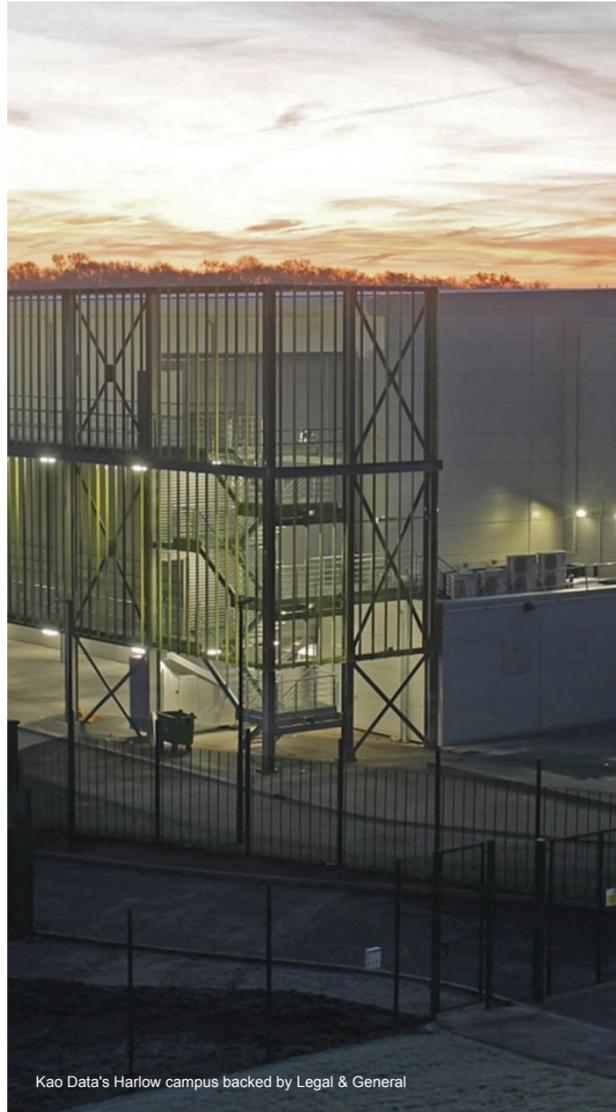
"We often work in consortia," Colombo says. "Alongside bringing capital, we can bring value-add through local partnerships or real estate expertise. That helps unlock opportunities that might otherwise be complex to deliver. But we are also looking other shareholder partners to bring their own value-add."

This approach is important to L&Gs focus on supporting sovereign infrastructure deployments.

Each country in Europe will have different data sets with different levels of sovereign sensitivity, with specific firms and jurisdictions that will be more appropriate for investing in the infrastructure that supports its processing and storage.

"We have investments in Europe that will be attractive to local pension funds from those countries. If an investor's mandate is to invest in the UK or other specific geographies, we can support that. Equally, if they have a mandate for diversification, we can accommodate that too," Colombo explains.

L&G's structure also supports cross-sector collaboration, especially in a sector closely tied to digital infrastructure and with its own set of sovereignty risks emerging this decade.

Colombo explains how the firms digital and clean energy investment teams sit closely together, with the digital infrastructure fund and the L&G NTR Clean Power (Europe) Fund, which was launched in partnership with leading renewable energy specialist NTR earlier this year, offering opportunities for collaboration.

"We're looking at concepts like power purchasing agreements (PPAs), which are an attractive thing to bundle with the digital proposition," he reveals.

This is just one of the ways that L&Gs new digital infrastructure fund demonstrates the shift in how institutional investors are approaching the intersection of technology, sovereignty, and policy.

By anchoring the fund in UK and European assets while maintaining a measured global reach, Matteo Colombo and his team have positioned L&G to play a central role in shaping the ownership and governance of critical digital infrastructure.

The first close may have flown under the radar. But the implications of investing in digital infrastructure in this manner are sure to be more than a minor blip.

DIGITAL GARDEN SUMMIT

WHERE LEADERS SHAPE THE FUTURE OF DIGITAL INFRASTRUCTURE IN JAPAN AND APAC

MARCH 25TH **- 26**TH
TOKYO, JAPAN

Gold Sponsors

CUSHMAN & WAKEFIELD   AIRTRUNK

Silver Sponsors

ada infrastructure   ONION technology   edgeconnex   MORRISON FOERSTER   Goodman   STTelemedia Global Data Centres   STACK INFRASTRUCTURE

Bronze Sponsors

GT GreenbergTraurig   AAR+ ASIA KEN ARCHITECTURAL RESEARCH INC.   CyrusOne KEP   IDCFrontier   JLL   OPTAGE   Princeton Digital Group

Strategic Partners

WOMEN IN DATA CENTERS   THE TECH CAPITAL   APREA   STRUCTURE RESEARCH   DC NEWS ASIA   APDCA

ReData

THE RISE OF BRAZIL'S BILLION-DOLLAR DATA DREAMS

by Jack Haddon

Chris Torto

With ReData poised to slash import taxes on critical server equipment, Brazil is suddenly the frontier for hyperscale investment. Chris Torto, CEO of Ascenty, breaks down with Jack Haddon how this seismic shift is creating billion-dollar projects, attracting US and Chinese tech titans, and rewriting the rules of capital in Latin America.

''We will definitely see more international banks in the financing of these projects, as you see it globally today.''

Despite being introduced as a provisional measure to Brazil's parliament in September, ReData, a legislative move from President Lula da Silva's government to slash import taxes on equipment vital to the operation of data centres, is already having an impact on the country's data centre market.
"Five years ago, a typical data centre we would build would cost about US$100 million," says Chris Torto, CEO of Ascenty, the oldest and one of the largest hyperscale focused data centre platforms in the Brazilian market.§
"Now, we're not looking at anything below US$1 billion."

This increase in scale shows just how serious global technology companies are about investing in their digital infrastructure footprint in Brazil.

And it really is global. Torto says that Ascenty's current development pipeline is a 50/50 split between Chinese and American

hyperscalers, with both sides looking to work heavily with third party operators due to the complexities of the Brazilian market.

The sudden influx of demand unlocked by ReData will fundamentally change Brazil's data centre market in a number of ways; including the way individual facilities and data centre platforms are financed, the topology of the hyperscale-dedicated development company landscape and Brazil's data sovereignty environment.

Financing implications

The increase in the size of the projects that Ascenty now sees in its pipeline requires a new model of financing.

Its two main shareholders are private equity giant Brookfield and global data centre REIT Digital Realty, which Torto (who owns the other 2% of the business) acknowledges puts the firm is in a beneficial position to access deep pools of capital.

In addition, the firm has a roster of 15 domestic and international banks that have provided term loans to fund the construction of its existing portfolio, which spans Brazil, Mexico, Chile and Colombia.

But for projects of a US$1 billion or more, Torto expects to be introducing project financing with existing and new banking partners into the capital

mix, as well as more advanced financing methods like Asset Backed Securitisation that is being observed more commonly in mature markets like the US and Europe.

Moving from term loans to project financing means that the mix of capital providers is also likely to change.

"We do have a healthy dose of the local banks participating with us, but we will definitely see more international banks in the financing of these projects, as you see it globally today," Torto explains.

Torto also expects to see longer terms on project finance style loans, rather than the term loans employed to date.

"If we have a 15-year deal, I think we'll be able to do financing of seven-to-ten-year terms on those kinds of projects, eliminating the refinancing risks."

These longer-term loans are not possible under the current structure, due to variety in contract terms that Ascenty has across its portfolio. Using these leases to underwrite term loans means an average of the existing contracts has to be used, preventing the longer-term financing that would be possible on larger projects if they are financed individually.
"It also depends on who the customer is," Torto adds. "We're talking about the major hyperscalers that have incredible credit ratings. So that's a plus for us as well."

Impact on the Brazil market

The potential to work on these larger projects, in the US$1-5 billion range initially, is a result of the ReData initiative not only providing incentives for data centres' industry but representing a structural push for Brazil's entire digital infrastructure ecosystem.

By reducing barriers, fostering investment, and aligning energy and connectivity demands, the provisional measure creates the foundation for large-scale, long-term growth.
"The approval of ReData provides us with greater regulatory clarity and long-term predictability, which are essential for strategic planning," Torto explains.

For ongoing projects, this means reduced operational risks and a more favorable environment for the data centre sectors expansion.

"Regarding our pipeline, we now have the confidence to accelerate

deployments, expand capabilities, and evaluate new regions of interest together with our clients and prospects."

It also opens to door to international investors entering the Brazil market, launching new platforms to capture demand from US and Chinese tech firms that are themselves attracted to the market.

Ascenty has already been serving this demand in Brazil for 15 years, which Torto believes will help the company maintain its status as a market leader, even with increased competition.

"We've built 30 data centres that are in operation today. So, when you start talking about data centres that are a hundred megawatts and up, who has the most experience in building and operating data centres is us," he says.

In addition to its ability to access equity capital from Brookfield and Digital Realty, and its existing relationship with the banks, Ascenty has another strategic advantage.

"We own 4,000 kilometres of our own fibre within Brazil that connects our data centres with everybody else, as well as the cable landing stations. No other data centre operator can offer both data centre capacity and connectivity. And that's a very strategic advantage for us."

Even before ReData, new entrants were coming to Brazil to capture this demand and provide fresh competition to Ascenty.

Examples include Arch Capital's 247 Data Centres, CloudHQ, and Patrea Investments new platform, Omdia.

While these firms are starting to build facilities already, Torto is confident that the likes of Ascenty, and its peers with hyperscaler experience such as ODATA and Scala Data Centers will continue to thrive.

"Maybe some [of these new firms] will be successful, but I don't know if the hyperscalers, and there's only a half dozen of them, are going to be willing to risk a 100MW+ project on a newcomer coming into a country that's very difficult to do business," he predicts.

"You get a lot of people coming down, they put their sign out on the front of their shingle and say, we're ready to do business. I think they're in for a rude awakening of how difficult it is to do business and to get that first deal. I think it would be very difficult for a hyperscaler to choose a new cover to build 100-megawatt data centre."

Even with international experience, Brazil presents an environment with unique permitting, licensing, construction and labour laws.

Torto himself has spent 40 years in Brazil. Prior to Ascenty, he was the founder and CEO of Vivax, the second-largest cable TV carrier in Brazil, and also led Voyager, an internet company that went public in 1999 and was acquired by a US telecom group in 2000.

"It's not easy. Even today, after 40 years, there's things that I'm still earning after all of this time of how to operate," he adds.

This complexity, and the availability of firms like Ascenty with the necessary experience to navigate it, means that hyperscalers from the East and West are both planning on partnering with local firms for their deployments.

"Some of them have tried it and they've had a lot of mishaps," he says of the US hyperscalers going at it alone.

"If we talk to the Chinese hyperscalers, they're not even thinking about that. They're only going to use third parties."

But despite this, Torto believes that there is space for new firms and investors to make their way into Brazil – it will just require a more patient approach.

"Start a little bit smaller," he advises. "Do some of the smaller cloud deployments in the 2-10MW range to get operating experience and show that you can build."

Beyond practical considerations and local operating experience, the size of these companies is also important, especially to US hyperscalers, Torto says.

"If you talk to the major US ones, you have to have certain market capitalisation requirements to be able to even win their business. One US hyperscaler, you have to be worth at least a billion dollars. So that just eliminates a lot of the newcomers from doing some of the major deals."

Why is Brazil attractive to hyperscalers?

The reason that ReData is seen as such a landmark legislation for Brazil is its removal of a major bottleneck that was blocking investment in what is otherwise an ideal market for US or Chinese hyperscalers to do business.
"Brazil is the perfect spot to be the hub of AI for Latin America," Torto explains.

It has a surplus of power, that power is cheap, at just US$0.05 kw/h and its renewable, with 90% of the grid powered by hydro, wind and solar.

Its population also offers a compelling reason to invest.
In October, TikTok owner ByteDance was reported to be moving ahead with a giant campus in the northeastern state of Ceara.

This location is important, Cerea is one of Brazil's free trade zones, and from a data centre perspective acted as scaled down and localised version of the tax environment that is now extended to the rest of the country under ReData.

In May, ByteDance's development partner on the project, wind farm developer Casa dos Ventos, announced it obtained two key approvals to move forward with a 300MW data center project at the Pecem port complex, exemplifying the scale of the commitments that a more favourable tax environment could unlock.
"There are 170 million TikTok users in the United States. There are 180 million in Latin America," Torto says, further demonstrating why large commitments make sense for Brazil.
"The largest users of Netflix outside of the United States is in Brazil. The largest users of a number of these social media platforms, and AI platforms, will be in Latin America," he continues.

With an established data centre economy and supply chain, Brazil makes sense as the Latam hub for these digital infrastructure investments.

But it could also be a hub for US focused demand as well. Torto notes that many of the hypersaclers own their own subsea infrastructure connecting Latin America with the US. For the training of AI models,

more readily available, cheaper and cleaner power could incentivise US hyperscalers to commit to capacity deployments in Brazil to develop their models.

Building AI dedicated sites, whether for the local market or for the US, could further boost the Brazil data centre market from a cloud perspective.

Torto notes that they are not separate investments in the minds of the US hyperscalers.

"I've been to a number of the hyperscalers offices, and they've talked about this. It's not going to be, oh, well, we got AI over here, we got cloud over here. They're going to be combining a lot of it," he reveals.

With a 50% discount on server equipment compared to the previous tax environment, assuming ReData is permanently established in Brazilian law, data that is currently being stored and processed in the US could slowly move back to Brazil.

Torto doesn't think this will be immediate, but as the hyperscalers reinvest in server equipment, doing so in Brazil becomes an option, which could lead to lower prices, more features and data coming back to Brazil, boosting its digital economy and data sovereignty.

What are the risks?

Of course, all of this depends on the provisional measure being approved. Under Brazil's legislative system, a provisional measure becomes law immediately, but only until its congress approves it and establishes it permanently. This has to happen within 60 days, with potential for a further 60-day extension.

Sentiment from the industry is that ReData will be approved, and Torto is no exception:
"What we can say is that we are very optimistic about Redata and hope that the measure is approved, not only for the benefit of the data centre industry, as mentioned previously, but also for the country as a whole."

But it's not a dead cert. "The absolute worst case would be Congress

doesn't pass the law," he adds.

Beyond that, he admits that investments in Brazil's transmission infrastructure could also lead to the expected boom in data centre development not coming to fruition.

"The second worst case scenario is some of the investments being made in bringing the power to certain locations, upgrades of substations and upgrades of transmission lines take longer than projected," Torto says. These investments are currently projected to take between two to three years, as large quantities of the surplus energy are in the North of the country, while the data centre hotspot is in the southeast around Sao Paulo.

In contrast, the larger sites Ascenty is seeing demand for are projected to take 18 to 24 months, so any delays in the former would be "significant".
However, Torto is confident that this disaster scenario will be avoided. While there are certain government approvals that are needed for the work to be completed, the power industry in Brazil is largely in the hands of the private sector, which the Ascenty boss sees as an advantage.

If Torto is right, it won't be long before we see billion-dollar projects becoming the norm, as Brazil is poised to rival the world's most advanced data ecosystems, attracting equal interest from both US and Chinese hyperscalers.

Established operators like Ascenty, with deep local experience, robust capital backing and integrated connectivity assets, are well-positioned to lead this next chapter.

However, the opportunity is not without risks.

The provisional nature of ReData and the need for parallel investment in Brazil's transmission infrastructure remain critical variables that could shape the pace of growth. Yet, if the legislation is approved and supporting infrastructure keeps pace, Brazil's mix of abundant clean energy, strategic geography and growing digital demand could establish it as Latin America's AI and cloud hub, and one of the most dynamic data centre markets in the world.

MEGA MOVES IN DATA INFRASTRUCTURE

THE TOP 10 DATA CENTRE DEALS BY VALUE

As BlackRock and MGX make history with their $40bn takeover of Aligned Data Centers, The Tech Capital looks back at the top 10 largest deals by value.

05

CoreSite
An American Tower Company

$10.1bn: American Tower Corporation Buys CoreSite (2021)

American Tower Corporation (NYSE: AMT) expanded its digital infrastructure portfolio by **acquiring Colorado-based colocation provider CoreSite**. The acquisition enhanced American Tower's capabilities in edge computing, supporting the growing demand for low-latency data services, a key addition to its vast cell tower portfolio. The acquisition was primarily funded through a combination of cash on hand and new debt issuance, allowing American Tower to maintain financial flexibility. In 2022, AMT sold a US$2.5 billion stake in CoreSite.

06

QTS

$10bn: Blackstone Infrastructure Partners Buys QTS Realty Trust (2021)

In its second blockbuster deal of the decade, **Blackstone Infrastructure Partners acquired mega-scale operator QTS Realty Trust**, consequently delisting the former REIT from the New York Stock Exchange. The acquisition was executed through a leveraged buyout, with Blackstone Infrastructure Partners utilising a mix of equity and debt. More recently, in August 2024, Blackstone's REIT Blackstone Real Estate Income Trust, Inc. (BREIT) said QTS Data Centers has a development pipeline valued at more than US$92 billion.

01

Aligned
Adaptive Data Centers

$40bn: AI Infrastructure Partnership to acquire Aligned Data Centers (2025)

A powerhouse investment consortium including BlackRock, Global Infrastructure Partners (GIP), and Abu Dhabi's MGX – joined by Nvidia, Microsoft, and xAI – has agreed to a **US$40 billion takeover of Aligned Data Centers**, one of the world's largest data centre operators. The deal, the largest in the data centre space to date, marks the first major move from the group's US$100 billion AI Infrastructure Partnership, designed to finance and build the next generation of data centres powering artificial intelligence.

02

AIRTRUNK

$16.1bn: Blackstone and CPPIB Acquire AirTrunk (2024)

In a landmark transaction, Blackstone (NYSE: BX), in partnership with the Canada Pension Plan Investment Board (CPPIB), **acquired Australian hyperscaler AirTrunk**. The acquisition underscores Blackstone's strategic focus on expanding its footprint in the Asia-Pacific region. The acquisition was financed through a combination of equity from Blackstone and CPPIB, alongside debt financing. This structure allowed both entities to leverage their financial strengths while mitigating risk. The acquisition also represents Blackstone's largest investment in Asia Pacific to date.

07

interxion
A DIGITAL REALTY COMPANY

$8.4bn: Digital Realty Buys Interxion (2019)

Colocation giant Digital Realty's (NYSE: DLR) acquisition of Interxion, a European data centre provider, marked a significant expansion of its global footprint and was a transaction that reshaped the geographic board across the continent – even as the REIT had to divest certain assets due to Brussel's fair competitiveness policy. Digital Realty utilised a combination of equity and debt financing, including a public offering of shares to support the acquisition. The transaction has resulted **in growing revenues across the portfolio** and turned Digital Realty into one of the largest players in Europe.

08

DFT
DuPont FABROS TECHNOLOGY

$7.8bn: Digital Realty Buys DuPont Fabros Technology (2017)

Mid-2010s, **Digital Realty acquired DuPont Fabros Technology**, a provider of wholesale data centre solutions. The acquisition was financed through a mix of cash, stock, and debt, allowing Digital Realty to leverage its strong balance sheet.

03

CyrusOne.

$15bn: KKR and Global Infrastructure Partners Buy CyrusOne (2021)

Three years ago, KKR (NYSE: KKR) and Global Infrastructure Partners – itself acquired by BlackRock (NYSE: BLK) in January 2024 for US$12.5 billion – **joined forces to acquire CyrusOne**, a global data centre provider. The acquisition aligned with the investors' strategy to capitalise on the burgeoning demand for data centre services which skyrocketed during the COVID-19 pandemic. The acquisition was structured through a leveraged buyout, with significant equity contributions from both KKR and Global Infrastructure Partners, complemented by debt financing. The deal is anticipated to yield substantial cost synergies and revenue growth through expanded service offerings and enhanced market presence in the years to come.

04

switch

$11bn: DigitalBridge and IFM Acquire Switch Inc (2022)

DigitalBridge (NYSE: DBRG), in collaboration with IFM Investors, acquired Switch Inc., a Nevada-based player known for its innovative technology and sustainable practices deployed across its megascale data centre campuses. The acquisition was an early sign of the increasing emphasis on sustainability and energy efficiency in data centre operations. The deal was financed through a mix of equity and debt, with DigitalBridge and IFM Investors each contributing significant capital.

09

Global Switch

$7.7bn: Jiangsu Shagang acquires Global Switch over several deals (2016-2019)

Rather than a headline-grabbing single deal, Chinese steel company Jiangsu Shagang's acquisition unfolded in phases: an initial 49% stake was acquired as part of a consortium in 2016, followed by incremental purchases and, finally, the acquisition of the remaining 24% stake for $2.4 billion in 2019.

10

COMPASS
datacenters.

$5.5 billion: Brookfield and OTTP acquire Compass Data Centers (2023)

While unconfirmed, insider reports state that the **sale of Compass Data Centers** was conducted at an enterprise value of $5.5 billion. Previous investors RedBird Capital Partners and the Azrieli Group exited their position in the company, with new investor Brookfield taking a stake via its infrastructure group and OTTP transferring ownership from its private equity fund to its real estate fund.

FROM NICHE TO NORM

HOW ISLAMIC FINANCE IS WIRING THE GULF'S DIGITAL FUTURE
by Khamila Mulia

Once a specialist tool, Islamic finance is now powering the region's data centres, fibre networks, and AI platforms - often alongside conventional debt. Hybrid structures are cutting costs, extending maturities, and turning Shariah-compliant capital into a mainstream force for digital infrastructure.

In the Middle East, Islamic finance is an essential terrain for anyone raising capital. It is already shaping the capital stack behind new digital infrastructure in the region, and more issuers are taking notice.

Rooted in Shariah principles that prohibit interest, speculation, and prioritise real economic activity, Islamic finance appeals beyond Muslim investors because it mirrors many ESG practices.

It is also increasingly used to fund sustainable projects across the region, including digital infrastructure, where tangible assets and predictable cash flows fit naturally with Shariah requirements; barring riba (interest), curbing gharar (excessive uncertainty), and discouraging maysir (speculation).

What's changing is how sponsors blend it. Issuers pair sukuk (Islamic bonds) and commodity murabaha (cost-plus sale) with conventional loans to finance AI-ready data centres, fibre, and platforms.

The hybrid approach can lower the cost of capital, lengthen maturities, and smooth covenant issues, while tapping deep Gulf liquidity and staying within Shariah rules on asset backing and risk sharing.

Recent deals show the mix in action. Qatari data centre operator Meeza raised QAR 800 million (US$219 million) via a commodity murabaha from Dukhan Bank; regional telco Ooredoo Group secured a 10-year, QAR 2 billion hybrid facility from QNB, Doha Bank, and Masraf Al Rayan to scale its data centre and AI business; and Oman telco Omantel issued a US$500 million sukuk in 2024 with a seven-year tenor at 5.375% maturing on 16 January 2031.

Together these deals point to Islamic finance moving from a specialist tool toward a standard option for digital infrastructure funding in the region, with hybrids like Ooredoo's gaining traction for their flexibility.

Why does it matter?

When Islamic finance meets modern digital infrastructure, the fit is more practical than theoretical: These tools favour real assets and clear cash flows, which are the same features that underpin data centres, fibre, and platforms. Advances in software, data, and workflow automation are also making documentation, approvals, and audits more efficient, improving transparency without changing the core rules.

"Shariah scholars play a critical role, issuing fatwas that validate structures and ensure alignment with principles like risk sharing and asset backing," says Maya Marissa Malek, CEO of UAE-based global Islamic finance consultancy Amanie Advisors.

Islamic and conventional bond prices tend to move together, and large data-centre financings in the Gulf now commonly carry tenors close to a decade.

"Mixing Islamic and conventional tranches yields a competitive blended cost. In the Gulf, sukuk generally price in line with comparable bonds and can tighten when demand is strong, so adding an Islamic tranche rarely lifts the all-in rate, and can even lower it," Mindaugas Suklevicius, founder and fund manager at HF Quarters, tells The Tech Capital.

Stronger liquidity also reduces refinancing risk. He points to Ooredoo's QAR 2 billion, 10-year hybrid (Islamic and conventional) facility for data-centre and AI expansion as a case in point.

Hybrid financing expands the investor pool and lets issuers optimise structure. "Incorporating Shariah-compliant elements attracts liquidity from Shariah-compliant and ethical investors, increasing capital supply and potentially reducing pricing through competitive bidding and diversified risk," says Malek.

She adds that Islamic investors favour stable, asset-linked investments, which can extend maturities beyond the 5–7 years typical of bank debt, often to 10–15 years or more, depending on appetite and deal structure.

The rise of hybrids

Hybrid Islamic-conventional financing is now common across infrastructure, power, transport, and real estate, where segregated tranches bridge funding gaps while meeting Shariah requirements.

Hafeet Rail, the UAE-Oman link, closed about US$1.5 billion last year with both conventional and Islamic tranches in AED and OMR.

Saudi Arabia's megaproject NEOM has also tapped Islamic financing, such as sukuk and murabaha, alongside separate conventional facilities, to support its pipeline.

In digital infrastructure, demand has moved beyond regional Islamic banks to include sovereign wealth funds, pensions, and family offices, according to Malek.

What's speeding execution is the harmonisation and standardisation of documents and Shariah precedents.

"Experienced Shariah experts recommend the best-fit structure and validate compliance from the outset, reducing revisions and disputes and, with proactive engagement, potentially shortening timelines by weeks," she says.

A hybrid transaction starts by mapping real assets to an Islamic structure and everything else to a conventional facility.

Both sit under one umbrella agreement so covenants, events of default, and payment priorities match. The Islamic side delivers asset-backed funding and wider investor reach; the conventional side adds flexibility. Together they bring price tension, longer tenor, and smoother execution.

"Timelines can be compressed by appointing a lead Shariah adviser and involving the board early, running approvals in parallel with drafting as well as pre-clearing the tangible asset pool and having substitution mechanics ready," says Suklevicius.

Bridging the divide: Aligning lenders

Certainly, blending isn't seamless. Suklevicius explains that Islamic and conventional lenders often set different rules for borrowers, which can complicate deals if not managed upfront. In ijara (leasing) structures, for instance, the financier owns and leases the asset outright, shifting maintenance, insurance, and enforcement duties compared to conventional security over borrower-owned collateral.

Penalty interest is not allowed under Shariah and late-payment charges flow to charity, creating mismatches in default mechanics. Therefore, sponsors lean on a Common Terms and Intercreditor Agreement to level the playing field, with shared security pools and pro-rata payment waterfalls ensuring no side gets shortchanged.

"Cashflows are synchronised, aligning lease rentals to interest periods, and risk management harmonised under the ISDA/IIFM Tahawwut Master Agreement (global master agreement for transactions in Islamic derivatives)," says Suklevicius.

This setup not only prevents clashes but amplifies hybrid appeal, as lenders compete on equal terms for data centre financings that blend hard assets like servers with softer ones like fibre optics. In the Gulf's liquidity-rich market, it is becoming table stakes for issuers chasing that decade-plus tenor without covenant creep.

Tailoring for tech

Digital infrastructure mixes solid hardware with fast-changing tech demands, so Shariah-compliant tools are adapting quickly, although challenges remain, Suklevicius believes. For new data centres, funding often uses istisna-to-ijara: build the asset first, then lease it. Existing sites lean on simple ijara sale-and-leaseback deals for fast cash.

Meanwhile, quick needs like IT upgrades tend to rely on murabaha facilities (a cost-plus purchase model) that renew easily.

To speed things up, follow a clear process: Pick the right structure and match it to assets, then bring in Shariah experts early for approval drafts. Next, handle paperwork using regulatory standards while sorting regulations like Islamic insurance (takaful) and taxes.

"Regulatory, tax, and insurance setup, including takaful, have to be taken into account," Suklevicius notes, adding that delays from repeated scholar reviews can add weeks if steps aren't coordinated.

Run tasks in parallel can be the answer: Vet assets ahead, add flexible clauses for swapping out old tech, and appoint one lead Shariah advisor to oversee it all. Gaps linger, though. Funding pure intangibles like software is tough, often needing clever fixes such as musharaka partnerships that share risks based on real performance, he adds.

Still, major deals that happened recently show these issues are fading, making Shariah rules a real advantage for massive AI projects. With regional digital spending projected to contribute US$320 billion from AI alone by 2030, Islamic financing is essential for staying ahead.

Closing the loop: Myths, tips, and the horizon

As Islamic finance cements its role in powering the Middle East's digital boom, from AI hubs to resilient fibre networks, it's worth clearing the air on lingering doubts and arming issuers with actionable steps. "Islamic finance is an ethical, asset-based system accessible to all, emphasising universal values like fairness and sustainability," says Amanie's Malek.

Far from being "only for Muslims," it's a broad-appeal toolkit that aligns with global ESG mandates, drawing in ethical capital worldwide. On cost, the myth of it being more expensive doesn't hold. Products price competitively, especially for scaled projects, and the diverse investor pool, from Gulf sovereigns to impact funds, often drives rates down further, says Malek.

There's also concern about complications and slow processes. Malek advises three top things to address these issues; appoint a reputable Shariah scholar or advisor early, engage a hybrid team, and focus on the underlying assets and cash flow to unlock the best-fit Shariah solution and seal the deal.

Looking ahead to the next twelve months, Suklevicius sees steady tailwinds for hybrid adoption: Regional liquidity is flush, policies bolster Islamic markets, and ESG synergies are supercharging sustainable sukuk pipelines, which is ideal for energy-thrifty data centres.

Regulatory digitisation adds momentum, from the UAE SCA's freshly minted framework for security and commodity tokens (adopted mid-2025 after consultations) to ADGM's FSRA ongoing market tweaks on staking and fiat-referenced tokens.

Moreover, AAOIFI's Shariah Standard 62 on sukuk, now eyeing finalisation by early 2026 after extended drafts. Its rollout could ripple through structures, but proactive jurisdictions will adapt swiftly.

"To capitalise on this growing market, international lenders should prepare by building in-house Shariah expertise, setting up hybrid deals processes, and pre-packaging dual ESG plus Shariah products to access the broadest investor base possible," Suklevicius concludes.

Maya Marissa Malek, CEO at Amanie Advisors

Mindaugas Suklevicius, founder and fund manager at HF Quarters

---

THE ESSENTIALS, DECODED

• **Sukuk (Islamic bonds):** Ownership interests in specific assets or asset-backed cash flows. Common for portfolio scale or refinancing.

• **Murabaha (cost-plus sale):** The financier buys an asset and resells it at a disclosed mark-up on deferred terms. Useful for equipment purchases and working capital.

• **Ijarah (leasing):** Pay to use an asset, with an option to transfer ownership later. Works well for land rights, Mechanical, Electric, and Plumbing (MEP) plant, and sale-and-leaseback of existing assets.

AI'S NEXT BOTTLENECK

WHY POWER AND LAND ARE THE NEW DATA GOLD

by Joe Carbery, Associate Director, Corporate Finance at Heligan Group

The Nasdaq's 50% surge over the past 18 months has been driven by AI. Yet beneath the rise of GPUs and large language models lies a more pressing story, the infrastructure required to power them.

This is no longer a race for silicon; it's a land grab for power, grid access, and build capability. A single NVIDIA H100 draws up to 1,000 watts. A rack of AI servers can exceed 100 kW. These systems don't just need space, they demand reinforced floors, industrial-grade cooling, and direct grid access.

The bottleneck isn't chips. It's physical readiness. Data centres are delayed not by semiconductors but by substations, transformers, and cooling systems. Today, infrastructure defines the limits of AI deployment.

From Chips to Capacity: The Shift in Value

Global data-centre M&A is set to exceed $80 billion in 2025, extending the record set in 2024. Private equity continues to dominate, accounting for up to 90% of deal value, but hyperscalers are reshaping the field. Microsoft, Amazon, Alphabet, and Meta are on track to invest more than $400 billion this year to secure capacity. Deals over $10 billion, once rare, are now routine.

The $40 billion acquisition of Aligned Data Centers, backed by BlackRock, Microsoft, and Nvidia, marked a turning point. Control is shifting beyond compute to encompass land, energy, and thermal systems. Other landmark transactions reflect the same momentum:

• Blackstone and CPPIB's $16 billion acquisition of AirTrunk secured grid-adjacent campuses across Asia-Pacific.

• DigitalBridge and Silver Lake's $9.2 billion commitment to Vantage Data Centers will deliver 3 GW of future capacity.

• Schneider Electric's acquisition of Motivair and Daikin's purchase of DDC Solutions show industrial moves into liquid cooling, now essential for high-density AI workloads.

This is consolidation with strategic intent. Platforms are being valued not only on installed capacity but on expandable megawatts, retrofit potential, and proximity to reliable power.

Infrastructure Must Evolve – Not Just Scale

AI has redefined what infrastructure must deliver. Where legacy data centres were engineered for predictable loads, AI demands thermally intense, rapidly scaling environments. Cooling is shifting from traditional air to liquid loops, immersion tanks, and submerged modules. Some operators are experimenting with offshore locations and deep-sea water for thermal efficiency.

Power architecture is evolving, with vertical substations, on-site generation, and dynamic load balancing replacing legacy cabling and redundancy systems. The most valuable platforms are built to adapt, able to retrofit and reconfigure for changing chip designs, sustainability mandates, and regulation.

Engineering has become a source of strategic edge. The ability to design, build, and optimise now underpins asset value. Firms such as Accenture, WSP, and Turner & Townsend are expanding in digital infrastructure, recognising that delivery expertise now sits at the core of asset value.

Scarcity and the New Risk Stack

Infrastructure is increasingly constrained by external forces. In markets such as Northern Virginia, Dublin, and Frankfurt, grid strain has triggered moratoria on new developments. Construction timelines now exceed five years in many regions.

Sustainability is no longer optional. With cooling accounting for up to 40% of total energy use, investors are targeting platforms with waste-heat recovery, district heating, and renewable baseloads. The ability to demonstrate low-carbon delivery is emerging as a pricing premium.

The risk stack has shifted. Energy access, permitting, and build viability now outrank tenant demand. Infrastructure is hard to build, power, and deliver. Capital is chasing platforms that can cut through constraint with speed, certainty, and scale.

Infrastructure as Competitive Advantage

Infrastructure is no longer a support function; it is the battleground of the AI economy. The platforms that outperform control three levers:

• Megawatts of scalable power
• Megabytes of high-speed bandwidth
• Mechanical resilience and operational depth

This dynamic favours operators with integrated models combining land acquisition, grid negotiation, and engineering delivery. These are high-spec, high-return platforms engineered for perpetual reinvention.

For investors, focus is shifting from capital deployment to operational control. Success depends on assets built for higher densities, stricter regulation, and sustainability scrutiny. The winners won't compete on capital alone but on control of land, grid access, engineering execution, and performance excellence.

Why Infrastructure Grounds the AI Bubble

Talk of an AI bubble is growing as software and silicon valuations stretch historic norms. Yet unlike earlier tech cycles driven by speculation, today's AI expansion is anchored to hard infrastructure. Every model and GPU run is constrained by physical throughput -power, cooling, and concrete.

This is not a weightless bubble; it's a build cycle measured in megawatts and materials. While software hype may fade, infrastructure endures. Investors backing grid-adjacent land, thermal systems, and engineering-led platforms are underwriting the foundations of the digital economy.

Investor Outlook

Data centre scarcity is rewriting valuation logic. Assets will be priced not by current utilisation but by latent capacity, grid proximity, and retrofit potential. Physical constraints have become the new metric of innovation, setting the speed and scale of AI deployment.

At Heligan, we see this as the defining battleground of the AI decade, where those who control infrastructure will set the pace, the price, and the possibilities of artificial intelligence.

Space-based data centers are emerging, and next-generation satellites are the next frontier in digital infrastructure. Positioned in Low Earth Orbit, they promise near-zero-latency cloud and AI services while leveraging the cold of space for energy-efficient cooling. With multi-billion-dollar investments underway, orbital infrastructure is set to redefine connectivity and computing geography.

SPACE INFRASTRUCTURE

STAR WARS

RISING RIVALS IN THE LEO SATELLITE RACE

by Jack Haddon

KNOWN, STARLINK'S DOMINANCE IN LEO SATELLITES IS. BUT BENEATH THE SURFACE, AWAKENING, A NEW COMPETITIVE FORCE IS. JACK HADDON REPORTS.

"This transformed at least the vision of Starlink from a satellite ISP to a satellite MNO, putting them in direct conflict with the US wireless carriers."

In rural Texas, a cattle rancher returns home after a long day's work and enjoys a cold beer while watching the latest Star Wars spin-off on Disney+.

Cruising above, at 36,000 feet, an engineer is on their way back from a trip to see family to the mine they work at in Brazil, catching the same show on their mobile. When they land, their role will look slightly different – autonomous machinery will be improving efficiency and safety, and, for this season, they'll be able to stay in touch with family and friends back home.

As they pass over the Caribbean, passengers on a cruise ship sailing the blue seas below will still be able to take work calls and emails, even though they should be relaxing on their hard-earned vacation.

But while they're relaxing, on the other side of the world, a drone operator in the Donbas region of Ukraine is doing anything but. They're coordinating a sophisticated attack on a crucial Russian asset, part of a deadly conflict reshaping the global security paradigm.

All of this is only possible because of one technology: low Earth orbit (LEO) satellite internet.

The LEO satellite space has undergone a profound transformation in the past decade.

The market is now forecast to grow at a compound annual growth rate of 21.9% from 2024 to 2029, reaching US$25 billion in annual spending, market research from Gartner shows.

And according to Goldman Sachs Research, the market could grow by seven times over the next decade, reaching US$108 billion by 2035.

As many as 70,000 LEO satellites are expected to be launched over the next five years, installed at altitudes of 160–2,000km, and circling the Earth every 90 minutes.

"Elon Musk's mercurial behaviour regarding Ukraine, threatening to switch off service, makes customers think twice about relying on Starlink."

Mapping the battlefield

In the West, the current market leader, SpaceX's Starlink, will contribute significantly to these new launches, as will its competitors, including Amazon Kuiper and the UK's OneWeb, now a part of French satellite group Eutelsat following an acquisition that closed in September 2023.

Much of the satellite fever that is gripping the connectivity space is a result of the scale that has been achieved by Starlink since it began building its LEO constellation in May 2019.

It currently operates the vast majority of operational LEO satellites in orbit, with over 8,500 in lower LEO orbit (<600 km).
This represents over two-thirds of all operational satellites in existence and over 75% of active lower LEO satellites (11,226), according to data from Jonathan's Space Pages.

Data on the number of satellites in space, as well as the projected launches from these players and a multitude of others, varies, but Jonathan McDowell, the curator of Jonathan's Space Report and an astrophysicist working at the Harvard-Smithsonian Center for Astrophysics, has been widely considered the journal of record for the space economy for over three decades.

The data in this article will be taken from Jonathan's Space Pages, his online database, unless stated otherwise. It's also worth noting that launches are increasingly frequent (weekly at this point), so by the time you read this, these numbers will have changed - but the scale of Starlink's constellations relative to everyone else will stay the same.

Eutelsat's OneWeb constellation is currently nearing completion, with 641 working satellites, and after a series of delays, Kuiper has launched 155, with 41 operational.
Roughly 53,000 of the estimated 70,000 launches over the next half-decade are likely to be from China, writes Allen Chang, Head of the Greater China Technology Research team at Goldman Sachs.

The largest Chinese planned constellations include Guowang, Qianfan, and Honghu-3, which between them have announced plans to launch approximately 23,000. To date though, only Qianfan has active satellites - and just 67 of them at that.
In total, considering mega-constellations (which Jonathan defines as having plans for more than 1,000 satellites) and large constellations (50–1,000 planned), there are over 555,000 planned LEO satellites across 67 planned constellations.

Currently, 48 of these constellations have at least one active satellite, although it should be noted that some companies, such as Starlink, have multiple constellations, and 337,323 of the "planned" satellites are from E-Space, a U.S. company with just nine launches and five operational.

The plan of attack

This data from Jonathan comprises all LEO satellites, not just those that offer internet services. Other use cases include Earth observation, weather forecasting, navigation, scientific research, and military operations. But even within the satellite internet sector, there are different business models emerging.

Starlink has primarily focused on the B2C (business-to-consumer) space, connecting users in remote locations, offering connectivity for people living in areas not covered by fixed or fixed wireless internet from terrestrial players, or else offering a faster and higher-bandwidth alternative. OneWeb, on the other hand, does not play in the consumer space, focusing instead on business and government customers, as Head of Investor Relations Joanna Darlington explains:

"The big difference is that Starlink is 90% B2C. You order a terminal

online and set it up yourself. We don't touch B2C at all; we're purely B2B (business-to-business) and B2G (business-to-government)." However, Starlink has started addressing B2B applications, particularly in mobility.

"They've signed contracts with airlines like Air France, while we have Air Canada and United Airlines. Where Starlink enters B2B territory, they're definitely a competitor for our business model," Darlington continued.

This also makes Starlink a more direct competitor to the legacy satellite industry, which operates in geostationary (GEO) and mid-Earth orbit (MEO).

GEO and MEO satellites require less upfront investment to cover a specific area of the Earth. Because GEO satellites orbit at the same rate that the Earth rotates, a single satellite is always facing the same part of the planet. But the drawback here is that latency and bandwidth are not comparable with LEO solutions.

Kuiper, meanwhile, will likely combine both business models, tapping into Amazon's existing AWS cloud platform to integrate connectivity solutions with cloud computing in the B2B space, while using its customer-facing subscription businesses such as Prime Video to enter B2C.

A subset of the B2C sphere is direct-to-device (D2D) services, which instead of connecting to a terminal, offer satellite internet, calls, and messaging to smartphones or other mobile devices. This is an area that Starlink is also growing in, alongside its investment in B2B and B2G.

Starlink vs. everyone: The new frontlines

The impact of this latest market segment entry from Starlink should not be understated. It redraws the battle lines and changes the question of who Starlink is competing with altogether. Not only does it have Amazon, Eutelsat, and Chinese constellations snapping at its heels, it has opened a new front with the major US wireless carriers.

In September 2025, SpaceX acquired US$17 billion of spectrum licenses from EchoStar, marking a major step for Starlink into the US wireless industry. This marked an evolution rom its previous strategy in the US, which relied on a partnership with T-Mobile. T-Mobile officially launched its satellite texting service with SpaceX in July 2025.

According to Carlos Placido, an independent consultant in the satellite space, this transaction "essentially transformed at least the vision of Starlink from a satellite ISP to a satellite MNO, putting them in direct conflict with the US wireless carriers."
Without high-density cell tower infrastructure, Starlink won't be

competing in major urban areas, but Placido adds that Starlink doesn't need to dominate in any one market to become a relevant global player. They can still be dominant in rural areas.

"MNOs might initially be okay with that," he says, "because those are areas that are not profitable for them anyway to deploy cell towers. But losing dominance in those areas is not something that they're going to like."

In an interview on the All In Podcast, SpaceX CEO Elon Musk made a comment that is likely to further put U.S. MNOs on the defensive, stating that acquiring Verizon was not out of the question.

"I think it was a mistake to openly announce that they're going to be an MNO," Placido says. "That puts MNOs on the defensive, and maybe they'll think twice or three times before signing a distribution deal with Starlink."
Instead, Placido believes Starlink should have taken a more measured approach, perhaps indicating that it needed EchoStar's spectrum to deal with exclusion zones and investing in an ecosystem approach.

Kuiper's alliances

Building an ecosystem-led strategy may be even more important for Starlink given the strategy it seems that Kuiper is deploying.

Kuiper's integration with AWS represents an ambitious bid to "platformize" space connectivity, embedding satellite infrastructure into the broader cloud computing stack.

Placido can imagine a world where Kuiper's ecosystem-friendly approach would leverage AWS for data routing, AI applications, and even content delivery.

Early partnerships, such as with NBN in Australia and DirecTV in Latin America, suggest a collaborative model with other satellite players, rather than Starlink's stand-alone dominance.

For Amazon, LEO connectivity is less about direct profit and more about reinforcing the gravitational pull of its broader digital infrastructure. Kuiper could become the satellite edge of AWS, a logical extension of the cloud's global footprint.

"Looking ahead, Amazon could further leverage AWS to enhance Kuiper's service layer, potentially enabling advanced features such as edge computing, dynamic network management, or bundled cloud services for enterprise and government clients," says Garinder Shankrowalia, Senior Analyst, Broadband Access at Omdia.

If Amazon does take this approach, it might force Musk's hand to follow a similar path.

"If that's the case, what would be the role of Google? Because Google is a major investor in SpaceX," Placido considers.

Kuiper will inevitably start to compete with AWS in the consumer market, potentially creating a price war in what is effectively now a monopoly, or stepping in to pick up demand in markets where Starlink has sold all of its capacity and would have to lower service levels to bring new users online.

This infringement means Starlink could be incentivized to compete in Amazon's converged ecosystem approach, moving from pure connectivity to networking.

If such a development is to take place, Musk will need to build stronger and more collaborative partnerships beyond those he has already established with major telcos worldwide.

From the commercial battlefield to the actual battlefield

But the intrinsic nature of Musk's approach to business and diplomacy, in light of his political foray this year alongside and then against Donald Trump, seems to be turning businesses and governments away from Starlink rather than towards it.

This is something that Eutelsat sees as an advantage.

"The geopolitical situation over the last six months has been super important for us for many reasons. Elon Musk's mercurial behaviour regarding Ukraine, threatening to switch off service, makes customers think twice about relying on Starlink," Darlington tells The Tech Capital.

Conversely, Darlington claims Eutelsat is "not tied to a single nation's interests and has good governance."

"The conflict in Ukraine has shown that space has become a battlefield, thrusting satellites into the spotlight," she adds.
Keeping Ukraine connected is essential to its survival, as mobile and satellite networks coordinate military operations and enable a new breed of modern, drone-led warfare. While the post-apocalyptic horrors of this technology have shocked millions on social media, their effectiveness cannot be denied. Ukrainian engineers are experimenting with Starlink's Direct-to-Cell (D2C) technology for military communications, and Kyivstar completed its first live test of D2C satellite technology in August.

Starlink in particular has been critical in areas where traditional

internet infrastructure is disrupted by the war, providing connectivity for military operations, including drone reconnaissance, artillery targeting, and troop communication.

While the relationship has not always been perfect (Starlink was accused of cutting access to its services during a planned operation to target Russian assets in Crimea), Ukrainian officials acknowledge that Starlink has saved "thousands of lives" and can cut down the time from identifying to striking a target with a drone from 20 minutes to just one.

Ukraine has explored other options, including OneWeb, but technical challenges, cost, and network capacity have stood as barriers.

"Donald Trump saying Europeans need to protect themselves has changed quite a lot of things," Darlington says. "We signed a €1 billion framework agreement with the French army in June, and practically all our capacity over Ukraine is sold."

Darlington adds that Eutelsat is gaining traction with other entities that have complex relationships with the U.S. These include a big contract with Greenland, significant contracts with Taiwan (which doesn't want to rely on Elon Musk "due to his relationship with China"), and growing interest in Canada "because of how they've been treated over tariffs by the U.S."

This pipeline of deals for Eutelsat spotlights the interest from governments in satellite sovereignty. While by their nature, LEO constellations need to be global to offer uninterrupted connectivity, governments want as much control over the services that cover their territory as they can get , or, failing that, a service provider from an ally they can trust.

This push for sovereignty, control, and an alternative to US or Chinese LEO alternatives is part of the rationale for a planned financing deal that will see the French and UK governments anchor a capital investment of €1.5 billion in Eutelsat.

Under the plan, France's state shareholding agency will inject €750 million and bring its stake in the satellite operator to 29.65%, while Britain plans to invest €163.3 million.

In addition to this direct financing, Eutelsat will also participate in the EU's Iris2 constellation, the first satellites of which will be launched in 2029.
Iris2 will combine LEO satellites from OneWeb with MEO and GEO solutions from SES led consortium SpaceRISE to build a multi-orbit sovereign satellite communications network dedicated to the EU.

The EU itself will provide €6 billion of the funding for the project, with SpaceRISE and Eutelsat proving around €2 billion each.

The EU is mainly investing in developing technology for IRIS². There are two main reasons the EU is building IRIS²: first, they want sovereign capacity without relying on America or China; second, they want to promote their own space industry, supporting both legacy players like Thales, Ariane, and Leonardo, as well
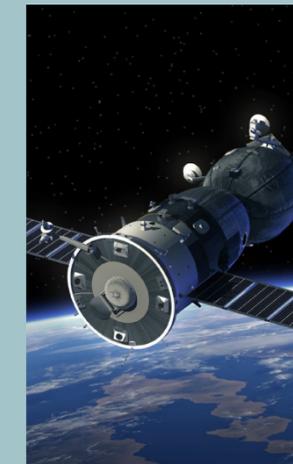as new space companies."

The funding structure for Iris2, driven by this focus on sovereign space tech, will give Eutelsat a significant advantage in reducing how much it has to pay in an especially Capex intensive industry.

"When the constellation is operational, the EU won't need all its capacity. For the LEO portion, the EU will keep 25% for itself, and 75% will be given to us to distribute commercially. This will effectively replace us building our own capacity"

What began as an effort to connect remote users has evolved into a geopolitical race between Starlink, Kuiper, Eutelsat, and emerging Chinese constellations, each vying for dominance in space, business, and influence.

Starlink's aggressive push into the mobile and enterprise markets, Kuiper's ecosystem-driven integration with AWS, and Eutelsat's bid for neutrality and to become a European champion all highlight how LEO constellations have become politically sensitive critical infrastructure, not just commercial ventures.

As nations seek satellite sovereignty and companies blur the line between connectivity and the cloud, the stakes extend far beyond streaming Netflix, towards national security, and the future of enterprise computing architecture.

STARLINK
Leading global LEO constellation with thousands of satellites focused on broadband communication services

OneWeb
LEO-focused constellation, now operated under Eutelsat following a merger, targeting remote and underserved regions

EUTELSAT
Major GEO operator with expanding LEO capabilities via OneWeb acquisition and integration of multi-orbit solutions

amazon | project kuiper
Deploying a global LEO constellation of over 3,200 satellites to deliver high-speed, low-latency broadband to underserved regions and enterprise customers, with commercial service rollout beginning in 2026

SES
Operates both GEO and a leading MEO constellation following acquisition of Intelsat, providing services to government, telcos, and enterprise customers globally.

TELESAT
Canadian operator with longstanding GEO fleet and the developing Lightspeed LEO constellation for enterprise and carrier markets.

Viasat
Primarily GEO (with ViaSat-3), serving aviation, maritime, and residential broadband markets; expanding global reach

inmarsat
Legacy GEO mobile satellite provider for aviation and maritime, expanding to hybrid arrangements via partnerships in LEO/GEO

SPACE 42
UAE-based SpaceTech integrating LEO/GEO satellite communications, Earth observation, and AI, expanding global geospatial intelligence with five SAR satellites and advanced connectivity solutions for government and enterprise clients.

China's Guowang, Qianfan, and Honghu-3S plan to deploy tens of thousands of satellites between them, building national mega-constellations for global broadband and data services comparable in scope to Starlink and Kuiper.

WHO RULES THE ORBIT?

THE POWER PLAYERS SHAPING SATELLITE COMMUNICATIONS

FROM GROUND TO ORBIT

THE 5G/6G SATELLITE SYNERGY SHAPING GLOBAL NETWORKS

by Kevin Cobley, RF Systems Director at Plextek and Richard Jacklin, Commercial Lead at Plextek

While the potential of IoT is theoretically limitless, real-world applications are still constrained by the capabilities of current connectivity infrastructure. Despite ongoing efforts by telecom providers to deliver seamless global communication, extending cellular coverage to commercially unviable or remote regions, where building terrestrial base stations isn't feasible, remains a major hurdle.

Non-Terrestrial Networks (NTNs) are well-positioned to overcome these challenges. However, technical limitations have historically held them back. That's now changing. With the advancement of 5G and 6G technologies, we're entering a new phase where investment in satellite communications could soon enable seamless, standardised cellular connectivity- even in the most isolated locations.

In this context, Kevin Cobley and Richard Jacklin of Plextek explore how the integration of 5G and 6G with NTNs, particularly through Direct-to-Device (D2D) communication, has the potential to revolutionise global connectivity…

From Base Stations to Sky Stations: The Direct-to-Device Disruption

Despite the fact that cellular technologies and satellites have existed since the early days of 2G and GSM, satellites have historically been primarily used for backhauling.

This process involves transporting traffic from the base stations to the mobile switching centre via satellite links, microwave radio links, or other methods of communication. However, current satellite coverage for mobile devices usually requires bulky antennas or offers very low data rates.

With 5G and 6G, satellites are now being considered for direct connectivity between base stations and user devices. This shift to D2D communication means devices – like smartphones or IoT sensors – can connect directly to a satellite without a ground-based intermediary. Connectivity around the world, no black spots in coverage, higher throughput, and even access to the most remote areas are just a few of the advantages. But it also means overcoming significant engineering challenges.

Breaking Barriers: Solving the Technical Challenges of Next-Gen Connectivity

One of the main challenges in space and satellite communication is efficient frequency usage. High spectral efficiency is a hallmark of 5G, enabling more effective use of the frequency spectrum. However, achieving this in space-based systems and user devices is anything but straightforward. Large antennas are often impractical, making it essential to adopt lightweight, compact, and efficient solutions that maintain performance. Another critical factor is the integration of antennas and radios – particularly in massive MIMO (multiple-input, multiple-output) systems, which require numerous antenna elements. This has driven a shift toward fully integrated units, where the antenna becomes a core part of the electronics.

Another major hurdle is the power budget for handset-to-space communications, especially considering the constraints of both devices and orbiting base stations. The distance between a mobile device and a satellite is significantly greater than between a terrestrial base station and a mobile device (300 to 500 kilometres versus 15 kilometres). That alone creates a major power challenge. Satellites often spend extended periods in eclipse, during which they cannot generate power because they are on the dark side of the Earth. This limitation necessitates careful power management strategies.

And finally, higher latency as a result of distance between device and satellite compared to device and terrestrial network, will need to be mitigated so as not to impact the user experience. While it is not possible to eliminate latency, it is possible to minimise its impact. One way is to perform edge processing in space instead of relying on the bent pipe method that relays data back to and from Earth for processing. This approach provides the best user experience but is power-intensive.

A system-level approach is essential for managing and balancing these interrelated challenges - one that combines technical expertise, especially in RF front-end design, low-loss and low-power solutions, with design-for-excellence principles like scalability, manufacturability, and testability..

And Finally..

The integration of 5G and 6G technologies with satellite networks has the potential to reshape global connectivity, but in practice the possibilities of IoT are still limited by current forms of connectivity. D2D communication could ensure reliable, low-power data transmission from remote sensors without requiring a local network. In defense, it could offer resilient and rapidly deployable networks. And for consumers, the dream of flawless mobile coverage – whether on land, at sea, or in the air - would finally become a reality. However, D2D is not without its challenges. Businesses need to invest in specialist engineering expertise to navigate the harsh realities of space - such as radiation, extreme temperatures, and launch-related stresses - while designing products that are scalable and cost-effective.

Kevin Cobley, RF Systems Director at Plextek

Richard Jacklin, Commercial Lead at Plextek

SPACEBOUND AI

WHEN CLOUD AND DATA CENTRES TAKE ORBIT

by João Marques Lima

As AI's energy appetite grows, engineers are taking the cloud into orbit. From orbital data centres to lunar storage vaults, humanity's information is reaching for the stars, powered by constant sunlight, cosmic cooling, and a vision of digital resilience beyond Earth. João Marques Lima reports.

Far above the deep indigo of Earth's horizon, where sunlight slices through the void, engineers are imagining a new frontier for humanity's information: moving the cloud itself into orbit. The idea is as audacious as it is logical. With AI's energy demands growing faster than terrestrial grids can handle, and land for large data campuses becoming scarce, why not take the problem to space?

What was once confined to academic panels and speculative tech conferences is now becoming a commercial reality. Orbital data centres are moving from concept to planning tables as computing companies and governments scramble to find sustainable ways to power AI and store humanity's ever-growing digital memory. In this silent, sunlit expanse, sunlight is constant, cooling comes free, and the political and legal disputes over land become irrelevant.

Jeff Bezos, founder of Amazon and Blue Origin, has come forward as one of the most high-profile advocates of this shift. Speaking at Italian Tech Week in Turin in October, he said: "Moving energy-intensive compute into orbit is both technically possible and economically rational." Elon Musk, CEO of SpaceX, echoed the sentiment, describing orbital compute as a "really good use of space to help save the planet".

Power, Space, and the Energy Imperative

As widely reported, data centres already consume around 2% of global electricity, and that figure is rising steeply. AI models, increasingly sophisticated and data-hungry, push these numbers even higher. Sam Altman, CEO of OpenAI, warns that "we still don't appreciate the energy needs of this technology… we need fusion or radically cheaper solar plus storage".

Starcloud, a California-based aerospace and infrastructure firm formerly known as Lumen Orbit, believes the solution is off-planet. The startup estimates that orbital data centres could operate at roughly US$0.002 per kilowatt-hour-nearly twenty times cheaper than average wholesale electricity in the US or UK.

The physics of orbit make these savings possible. For instance, Starcloud's proposed solar arrays would operate at a 95% capacity factor-compared with around 24% for US terrestrial solar farms and generate over five times the energy of identical arrays on Earth. With no day-night cycle, no weather interruptions, and negligible atmospheric losses, orbital solar arrays could supply energy almost constantly.

Cooling, normally a costly portion of terrestrial data centres, becomes almost effortless in space. Radiators can dump waste heat directly into the -270°C background of the cosmos. Starcloud calculates that a 40-megawatt cluster launched for roughly US$5 million could operate for a decade with negligible energy cost and zero water usage. In contrast, terrestrial data centres require vast amounts of water daily for heat exchange, particularly in hotter climates.

The European Union's Advanced Space Cloud for European Net Zero Emission and Data Sovereignty (ASCEND) program, a 16-month, €2 million ($2.1 million) study investigating the feasibility of launching data centres into orbit, has reached a "very encouraging" conclusion. The study found that "deploying data centres in space offers an eco-friendlier solution for hosting and processing data."

Bezos, Musk, and the industrial logic

For Jeff Bezos, the rationale is both environmental and industrial. He has long argued that heavy industry should eventually migrate off Earth, leaving the planet "for life and light industry." Blue Origin engineers are now designing modular orbital platforms that could host compute nodes powered directly by integrated solar fields.

Analysts close to Amazon Web Services (AWS) note that the project aligns with the company's long-term energy strategy. AWS, one of the world's largest renewable power buyers, faces increasing scrutiny over AI's energy intensity. Placing compute in orbit could decouple future expansion from terrestrial grids while tapping virtually
unlimited solar energy.

Launch economics are improving rapidly as reusable heavy-lift rockets such as SpaceX's Starship and Blue Origin's New Glenn have driven launch costs down to roughly US$30 per kilogram. Starcloud estimates that a 5-gigawatt orbital complex-larger than any terrestrial power plant-could be assembled in fewer than 100 launches.
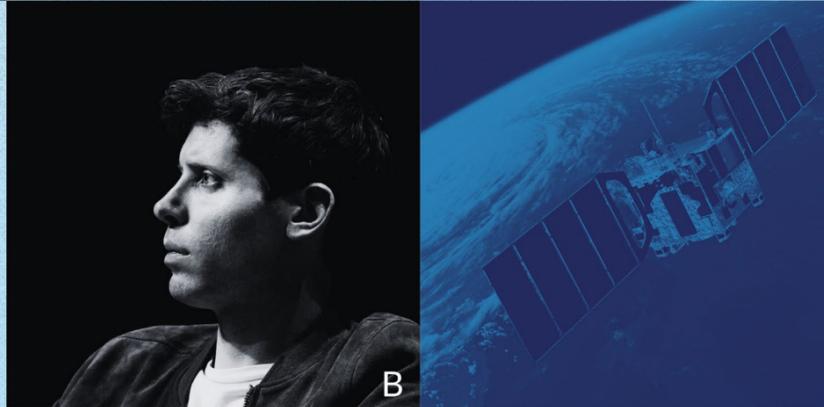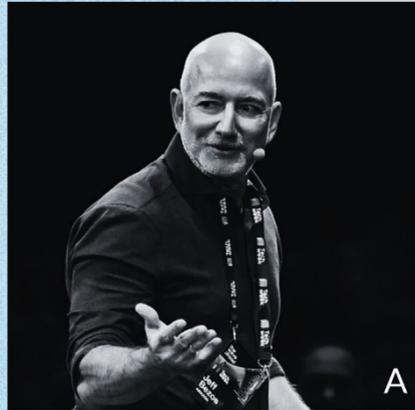
Bezos's vision positions orbital infrastructure as a natural extension of the cloud, while Musk emphasizes the environmental case and practical use of excess orbital capacity.

As private companies like SpaceX and Blue Origin develop reusable rockets and push lunar and Martian ambitions, investors are paying attention. Hines, a global real estate firm, recently acquired a fully leased 250,000-square-foot industrial property in Florida's Space Coast, catering to aerospace tenants.

"A real revolution has happened in the industry, and as things start to get unlocked, companies are looking for how they can monetize space more broadly, and there's a lot of pieces to that," said David Steinbach, Hines' global chief investment officer.

He points to both Earth-based infrastructure and lunar manufacturing. "We are in the early days of something that will be some major investments… think about all the nodes that are going to get developed and created. It's exciting, and I think investors need to be thinking that way," Steinbach said.

"One of those rails is data centres. There is unlimited power in space because of the sun, unlimited cooling with the vacuum of space, and unlimited real estate in terms of where you can put
these things," he added.

Moon as a Data Vault

Further from Earth, Christopher Stott, CEO of Lonestar, sees the Moon as a haven for data rather than people. His company, based in the Isle of Man and the US, is building off-world storage designed to serve as a digital archive immune to terrestrial disruption.

"I look at the Moon and I just see it… it's so logical to use Earth's largest satellite," Stott said. "A bunch of customers came to us and said, 'Please help us. We need our data to be safe, secure, accessible, and sovereign.' We couldn't find a perfect place on Earth… so we started looking at low orbit and GEO, and for various reasons, we decided to go up to the Moon."

For Stott, permanence is the key selling point. "We're a backup, a resilient backup. No matter what happens on Earth-storms, climate change, wars, human errors-the Moon is still there. My data is safe."

Lonestar's systems run on solar power and generate Type 1 carbon credits. "Thanks to Musk and Bezos, rockets are finally affordable," Stott noted. "This infrastructure has to be far more resilient than anything on Earth."

He also emphasizes scale. Humanity now creates 2.5 quintillion bytes of new information every day, doubling every two years. "We're building living digital twins of entire nations. The Moon is perfect for data storage; if it weren't there, we'd have to build it."

Sovereignty, Modularity, and Security

Space offers more than energy efficiency - it offers legal and political

clarity. "There are data sovereignty laws in space-space is the most regulated of all human activities," Stott explained. Under the Outer Space Treaty, nations retain jurisdiction over objects they register in orbit, allowing them to apply national frameworks to onboard data.

This opens possibilities for secure, extraterritorial storage. Assets can be split so that each country has dedicated drives, never shared with others. "It's sovereignty in silicon," Stott said.

Modularity is central to these designs. Self-contained compute or storage units can dock and undock independently. Starcloud calls this "incremental scalability"-expanding from a single module to multi-gigawatt clusters without the regulatory friction that slows terrestrial projects. Governments and corporate clients alike see value in modular orbital data nodes that can grow alongside their needs.

Engineering Challenges and Risks

Orbital data infrastructure is not without hurdles. Radiation, debris, and maintenance logistics demand sophisticated engineering and strict oversight. Over 30,000 pieces of space debris larger than 10 cm orbit Earth, posing collision risks. Radiation shielding must protect sensitive electronics from high-energy particles, while autonomous repair systems handle micro-damage. IBM researchers note that "the environmental case is strong, but the engineering challenge remains herculean."

Nevertheless, commercial logic is tightening. Lonestar has already achieved revenue on its first mission and is planning pilot projects to further prove the viability of orbital storage. With orbital centers, expansion is theoretically unlimited, unconstrained by terrestrial land or grid capacity.

Economics, Scale, and the Long View

The economic case is compelling as launch costs are dropping, solar power in orbit is abundant, and modularity allows clusters to grow incrementally. A single 5-gigawatt orbital cluster could eclipse the world's largest terrestrial data campuses. Investors, insurers, and governments are beginning to view orbital data centres not as speculative ventures but as viable infrastructure.

Bezos envisions space as the "industrial zone of the future", a place where high-energy computation, power generation, and heavy industry could migrate off Earth. For companies like Lonestar and Starcloud, the immediate goal is more grounded: creating redundancy for civilization's digital memory.

"This idea that we can protect data, meet data sovereignty, and do it with existing technology-it's evolved tools preventing another 'Library of Alexandria,'" Stott said.

As night falls across the planet and sunlight glints off a solar array orbiting above the terminator line, the cloud begins to take on a literal meaning. Humanity's information may soon orbit the world it describes, safe, resilient, and powered by the sun that never sets.

ORBITING THE FUTURE

MEGA-CONSTELLATIONS IN 2026

by João Marques Lima

The space surrounding Earth is no longer a quiet expanse reserved for national agencies. As 2026 begins, low Earth orbit has become a crowded highway of private satellites, with over 10,000 active units currently circling the planet and projections suggesting that number could triple in the next five years, based on the latest market data by Novaspace.

Low Earth orbit is particularly appealing for these mega-constellations because its lower latency enables high-speed broadband, real-time gaming, and financial transactions that demand split-second communication.

Among the leaders in this expanding domain is SpaceX's Starlink, which operates more than 7,300 satellites serving roughly 5 million users in 125 countries.

According to the Financial Times, the company anticipates generating $12 billion in revenue in 2025 and plans to expand its constellation to 40,000 satellites.

OneWeb, with 648 satellites, is focusing on connecting remote regions and generated €187 million in revenue in the year ending June 30, 2025, representing a significant portion of Eutelsat Group's sales, according to CircleID.

Amazon's Project Kuiper has already produced over 100 satellites and aims to deploy thousands more to deliver affordable broadband to underserved markets, based on reporting by TechTarget.

The proliferation of these networks is not merely about bridging digital divides. In regions lacking robust fibre or cellular infrastructure, satellite broadband is fostering telemedicine, remote education, and e-commerce.

Beyond connectivity, these satellites are enabling real-time Earth observation, providing high-frequency imaging for agriculture, disaster response, and environmental monitoring. They also support Internet of Things networks that connect shipping containers, farmlands, and industrial sensors in remote locations, as well as emergency communication systems that operate when terrestrial networks fail during natural disasters.

According to Juniper Research, revenue from fixed satellite broadband alone is expected to double from $10 billion in 2025 to $20 billion by 2030, underscoring the economic weight of this industry. Grand View Research further projects the satellite internet market to grow at a 13.9% compound annual rate through 2030, emphasising investor confidence in the sector's potential.

Yet with growth comes risk. As orbital traffic intensifies, so does the danger of collisions and the creation of debris that could imperil both commercial and scientific satellites.

Space agencies and operators are increasingly collaborating on tracking systems, collision avoidance protocols, and satellite deorbiting plans. Sustainability remains a pressing concern, and the concept of "orbital highways" - designated lanes for satellites in LEO - is emerging as a way to reduce the risk of catastrophic collisions while maintaining long-term usability.

Regulatory frameworks, however, remain fragmented across nations, presenting both challenges and opportunities for standardisation, liability, and international cooperation.

Looking ahead, the next decade may transform LEO into an integral layer of global infrastructure. Those keeping an eye in this space predict that satellite networks will increasingly integrate with terrestrial services, enabling seamless communications, AI-driven environmental monitoring, and interconnected transportation systems.

The satellites above could become as indispensable as the servers below - just notice the growing number of satellite dishes being installed at data centre campuses across the globe -, redefining industries from logistics to climate science.

For operators, investors, and regulators alike, the pressing question is no longer whether to deploy these constellations, but how to do so responsibly. According to industry experts, the space surrounding Earth is now a critical domain for commerce, invention, and strategic infrastructure, and the way it is managed today will shape humanity's reach beyond the planet for decades to come.

SATELLITE MEGA CONSTELLATIONS MARKET

FORTUNE BUSINESS INSIGHTS

Satellite Mega Constellations Market to grow at **25.53% CAGR** during 2025-2032

- $4271.4 Million — 2024
- $5559.2 Million — 2025
- $27309.0 Million — 2032

BY APPLICATION
- Navigation & Positioning **15.94%**
- Broadband Connectivity
- Earth Observation
- Others

BY CONSTELLATION PROGRAM
- Starlink
- Oneweb
- Kuiper
- Others

BY ORBIT TYPE
- LEO
- MEO
- GEO
- Others

BY CONSTELLATION SIZE
- Small (100-500)
- Medium (501-1000)
- Large (1001-3000)
- Very large (Above 3000)

INDUSTRY DEVELOPMENT
EnSilica declared that it received funding from the U.K. Space Agency through its connectivity in Low-Earth Orbit ("C-LEO") programme.

NORTH AMERICA
$1573.8 Million 2023
$2084.5 Million 2024

Middle East | Asia Pacific
Europe | Rest of the World

DRIVERS
Rise in Commercial Broadband Demand* along with Government Initiatives

TRENDS
Integration of Advanced Technologies in Satellite Mega Constellation

FOCUS

People.
Places.
Perspectives.

BRIDGE GAMES, CEO GAINS

TRANSFERRING SKILLS FROM THE FELT TO THE BOARDROOM

by Abigail Opiah

At first glance, it is just a card game. But for CEO Rajiv Ramaswami, Bridge is a battlefield of strategy, trust, and quiet calculation - a lens through which he navigates the high-stakes world of global business. Every hand reveals more about strategy, trust, and the art of winning. Abigail Opiah steps onto the playfield to find out.

When Rajiv Ramaswami, president and CEO of Nutanix, talks about strategy, logic and communication, his mind often drifts to an unlikely place… the Bridge table. Unusual for some, but for Ramaswami, the centuries-old card game is more than a pastime; it is a mirror reflecting the nuances of leadership, decision-making and teamwork.

"I discovered Bridge when I went to IIT," he recalls, referring to his time at the Indian Institute of Technology, one of India's most prestigious universities. "The college was full of nerdy types who were very analytical. Bridge therefore appealed to this group. It is a game of logic and strategy."

Decades later, Bridge remains a thread woven through Ramaswami's personal and professional life. "Half of my classmates from electrical engineering and computer science at IIT are now in the Bay Area," he says. "We still play bridge together every week."

Ramaswami's story begins in India in the 1960s, in what he describes as a "somewhat atypical family." His mother, who had advanced degrees, was a working professional at a time when few women were.

"For me, that was a bit unusual back then," he says. "The values early on were around education, how important it is, and because both my parents were busy with their careers, my sister and I learned to be fairly independent."

Frequent moves across cities also taught him adaptability. "Every couple of years we would move. New city, new school, new friends. I think I must have gone to seven or eight schools during my K through 12," he says. "You grow up learning to adapt."

After earning a place at IIT, a feat he describes as "lucky", Ramaswami went on to build a career spanning IBM Research, Cisco, VMware, and now Nutanix. But despite his achievements, he remains reflective about the people and experiences that shaped him.

"I was fortunate early on to have a phenomenal mentor at IBM," he says. "He taught me how to take care of people and lead with empathy."

Lessons from the card table

He explains that Bridge offers profound parallels to leadership. "Bridge is a team sport, and you are only as good as everyone else in your team. You might be phenomenal, but if your partner is not aligned, it does not work," Ramaswami says.

"That is very similar to business. You have to collectively boost the capability of the team if you want to win."

Partnership, in bridge and in business, is built on trust and communication. "There are set rules about what you can communicate and what you cannot," he says. "If your signals get crossed, it can lead to a bad outcome. The same goes for leadership, you need to communicate clearly, align on objectives and listen carefully."

He pauses, then adds, "People often underestimate the need to communicate. Whether it is aligning everyone on a vision or understanding issues, communication is key."

HOW TO PLAY BRIDGE:

Players: 4, forming 2 partnerships (North-South vs East-West).

Deck: Standard 52-card deck.

Card ranking: Ace (high), King, Queen, Jack, 10 down to 2 (low).

DEALING

Each player is dealt 13 cards.

Cards are usually dealt clockwise.

BIDDING (AUCTION)

Players bid in a clockwise rotation to declare the number of tricks their partnership will try to take above six.

Bids consist of a number (1–7) and a trump suit (♠, ♥, ♦, ♣) or No Trump (NT).

Example: "2 Hearts" = attempt to take 8 tricks with hearts as trump.

Bidding continues until three consecutive passes.

The final bid becomes the contract. The partnership that wins the bid becomes the declarer.

PLAYING THE HAND

The player left of the declarer leads the first card.

Players must follow suit if possible; if not, they can play any card (including trump).

TRICK WINNER

Highest card of the suit led, unless trump is played.

If trump is played, highest trump wins.

The winner of the trick leads the next trick.

SCORING

If the declarer's partnership makes the contract, they earn points.

If they fail, the defending side earns points.

Additional points are awarded for overtricks, undertricks, slams (bidding 6 or 7 tricks above six), and honours.

Beyond communication, Ramaswami sees Bridge as a masterclass in logic and strategy. "For every hand, you create an upfront plan, but you must reassess it with every card that is played. You are trying to create a mental map of everything around you; what the other players might have, what information you can infer."

In business, he says, it is much the same. "You never have all the data or all the facts. You make the best educated guesses based on the information available, and you adapt as things change."

Ramaswami laughs when he mentions one of Bridge's more delicate lessons. "People say you should not play bridge with your spouse, it can lead to conflict," he says. "Because when one person makes a mistake, emotions can run high. So you learn the value of constructive criticism. You figure out how to provide feedback without blowing up the relationship."

It is a lesson he carries into his leadership philosophy. "Whether at work or in life, people will make mistakes. The key is learning how to respond; calmly, constructively, and without ego."

For Ramaswami, Bridge is more than a strategic exercise; it is a form of decompression. "I did not start playing bridge because I thought it would make me a better leader," he says. "I play because I enjoy it. It is my outlet."

Beyond the boardroom

Over a career that spanned continents and technological revolutions, Ramaswami's approach to leadership has evolved. "Hire the best people and enable them to be successful," he says. "If they succeed, you succeed."

He also champions listening, a skill he admits he had to learn. "A lot of CEOs talk all the time. I used to. But you have to listen, especially to your harshest critics. If you speak first, you kill the discussion."

Adaptability has been another lifelong lesson for Ramaswami. He likes to say: "be global, think local", when he reflects on understanding where you are and adapting accordingly, sometimes in small ways, sometimes in big ones.

He brings that mindset to Nutanix's culture, too. "We have a return to office policy like most other companies, and now we understand that there are certain countries where commuting is horrendous. In India, there are a lot of traffic jams in Bangalore. We try our best to make sure everybody has cabs for people to commute in comfort."

When he is not leading a global technology company or bidding a hand in Bridge, Ramaswami enjoys music, tennis, and the daily New York Times crossword. He also writes reflective emails to Nutanix employees after every trip, sharing what he has learned from each country or city he visits.

"It is about accountability," he says. "Employees drive the company, and I owe it to them to share what I am doing."

Looking ahead, he hopes to be remembered not just as a CEO, but as someone approachable and fair. "Someone willing to help and wanting to do some good work," he says simply.
w
Bridge, after all, has taught him that success, in cards or in life, depends less on the hand you are dealt, and more on how you play it.

TOP 10 PEOPLE ON THE MOVE

Adam Selipsky
KKR

KKR has appointed Adam Selipsky, who was the former chief executive of AWS, as senior technology and AI strategy adviser. He will provide guidance on the firm's global digital infrastructure investments, including data centres, energy, and artificial intelligence.

Hamish Smith & Oliver Weston
Newmark

Hamish Smith and Oliver Weston have been appointed partners to Newmark's European data centre team. The double nomination will see the two executives focus on capital markets activity for data centres and powered land, including investment, development, refinancing, and site selection for new facilities.

Stephen Liang
Vertiv

CTO Stephen Liang will retire on 1 January 2026 after more than 30 years with the company and its predecessor, Emerson. Following Liang's retirement, Scott Armul will take on an expanded role as chief product and technology officer.

Marcel Louw
Digital Realty

Marcel Louw has been named Digital Realty's new managing director for Africa, with responsibility for operations in Kenya, Mozambique, Nigeria and Ghana. Louw has more than two decades of experience in infrastructure investment and fund management.

Chia Wee Luen
Microsoft

Microsoft has named Chia Wee Luen as its new managing director for Singapore. Wee Luen has more than 20 years of experience in enterprise technology and has previously held senior leadership positions at Red Hat, Qlik and Oracle.

Paul Finch
Kao Data

Paul Finch, the company's long standing Chief Technical Officer (CTO), has stepped down from his role to pursue new opportunities, while continuing to support the Board in a new capacity as a Senior Advisor.

Ken Marron
FLOW Digital Infrastructure

Ken Marron has been named FLOW Digital Infrastructure's Chief Technology Officer. Based in Tokyo, he leads the company's technology function across the Asia–Pacific. Marron has more than 30 years' experience in data centre construction and architecture across Japan, Korea, Taiwan, Singapore, and Europe.

Wayne Searle
AVK

AVK has appointed Wayne Searle as chief people officer. Searle has close to 20 years of senior human resources experience, including roles at Vertiv, where he led HR across Europe, the Middle East and Africa, and was involved in global mergers and acquisitions.

Kevork Kechichian
Intel

Kevork Kechichian has been appointed EVP and GM of Intel's Data Centre Group. He will oversee the company's data centre operations across cloud and enterprise. Kechichian has more than 30 years of industry experience, having previously worked at Arm, NXP Semiconductors and Qualcomm.

Gene Reilly
American Tower

American Tower Corporation has appointed Gene Reilly to its board of directors. Reilly was most recently vice chairman of Prologis, Inc., which he joined in 2003, when it was known as AMB Property Corporation. He has four decades of experience in real estate investment and management.

The Tech Capital brought its APAC Finance Forum 2025 to Capella Singapore on September 24-25, drawing senior figures from digital infrastructure, finance, investment, and policy. Set against the backdrop of Asia-Pacific's fast-evolving digital economy, the event unpacked how capital, AI, and sustainability are converging to shape the region's next wave of growth.

48 HOURS IN ABU DHABI

As The Tech Capital makes its debut in the Middle East and Africa, the spotlight turns to a region shaping the next era of digital growth. The MEA Finance Forum 2025 will bring together international investors, regional financiers, and digital infrastructure leaders to explore how finance can drive transformation across data, connectivity, and energy ecosystems. And while you're in Abu Dhabi for the conference, here's how to make the most of your time in the city.

1. Manarat Al Saadiyat
Manarat Al Saadiyat is a cultural and artistic center that hosts exhibitions, workshops, and events focused on the arts. It is an important venue for promoting creativity and cultural exchange in Abu Dhabi.

2. Corniche Beach
Corniche Beach is a public beach featuring crystal-clear waters and soft white sands. It's an ideal place for relaxation and offers stunning views of Abu Dhabi's skyline and the Arabian Gulf.

3. Li Beirut
Li Beirut is a Lebanese restaurant offering a fusion of traditional Lebanese flavors with a modern presentation. The restaurant combines authentic dishes with a contemporary dining experience, bringing the rich heritage of Lebanese cuisine to Abu Dhabi's culinary scene.

4. Ray's Bar
A sophisticated bar located on the 62nd floor of the Conrad Abu Dhabi Etihad Towers, offering breathtaking city views.

5. Qasr Al Hosn
Qasr Al Hosn is the oldest building in Abu Dhabi, originally serving as the residence of the ruling family. Now transformed into a museum, it tells the story of the city's rich history and cultural heritage. The fort showcases exhibits on Abu Dhabi's transformation from a small fishing village to a modern metropolis.

6. Louvre Abu Dhabi
Louvre Abu Dhabi is a world-class museum of art and civilization, displaying masterpieces of historical, cultural, and sociological significance from around the world. Designed by architect Jean Nouvel, the museum is known for its striking dome structure and diverse collection that spans centuries and civilizations.

7. The National Aquarium Abu Dhabi
The largest aquarium in the Middle East, home to over 300 marine species, offering immersive experiences and educational programs.

8. Ferrari World
A Ferrari-themed amusement park featuring high-speed roller coasters, interactive attractions, and driving experiences, including the world's fastest roller coaster, Formula Rossa.

9. Yas Waterworld
A massive water park offering thrilling slides, wave pools, and family-friendly attractions, inspired by Emirati culture and the legend of the Lost Pearl.

10. Abu Dhabi MAll
One of the city's premier shopping malls, offering over 200 international brand stores, dining options, and entertainment facilities. It is an ideal spot for fashion lovers and families looking for a variety of activities.
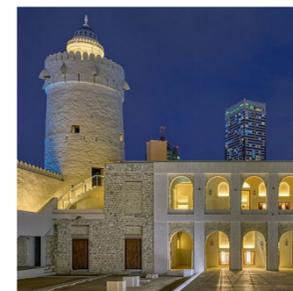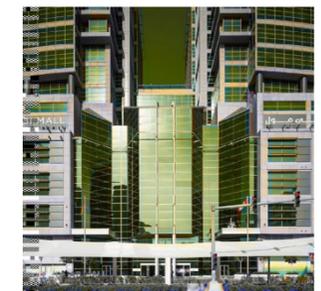

1

2

3

4

5

6

7

8

9

10

# MADE IN MENA

SEVEN GADGETS DRIVING THE NEXT WAVE OF REGIONAL TECH

by Abigail Opiah

International products
Here are 3 gadgets not necessarily made in MENA, but useful and
currently available online.

1. CardoO Smartwatch (Egypt)
A smartwatch by Egyptian IoT startup CardoO. Typical features include
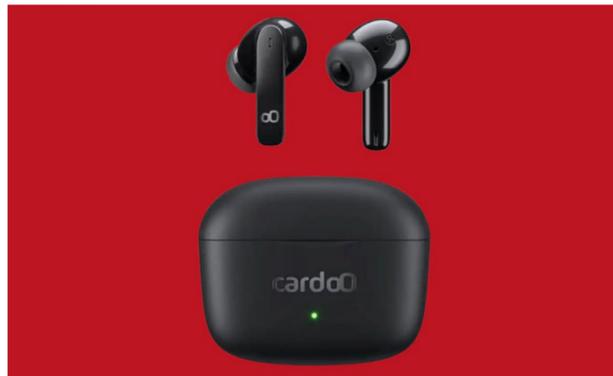health sensors (heart rate, SpO etc.), call and notification support, and
sporty design.

Price: EGP 2,375 for the standard model. That converts to roughly
US$75-80.
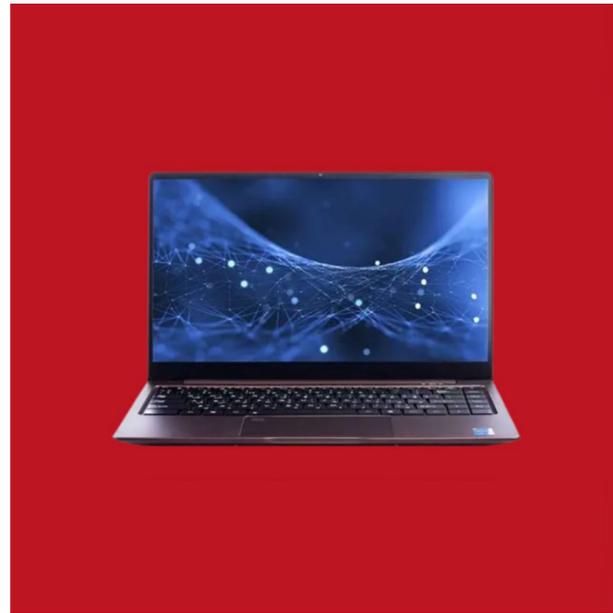
3. TAGITOP-UNI Laptop (Jordan)
A laptop made by TAGTech (Talal Abu-Ghazaleh for Technology). The
"UNI" model has an Intel i5-1135G7 CPU, 8GB RAM, 256GB SSD +
500GB HDD, 14.1″ Full HD screen. Designed for education / students.

Price: Dinar 596 (JOD 596) (before tax etc.). Roughly US$840-900.

5. Plaud Note Pro (AI Note Taker)
A digital note-taker device that records audio, syncs with transcripts,
etc. Good for students, meetings, etc.

Price: US$179.

7. OBJLGEV 3-Axis Gimbal Stabiliser
A stabiliser for smartphones (iPhone/Android), with anti-shake, LED ring etc.
Useful for vloggers.

Price: US$330.

2. CardoO Buds (Egypt)
Wireless earbuds ("smart earbuds") by the same company, with features
like multi-connection, 3D sound, noise cancellation etc.

Price: EGP 2,500. Roughly US$85-100.

4. TAGITOP-EDU Laptop (Jordan)
A more budget-oriented model by TAGTech meant for educational use.
It has an Intel i3 10th gen CPU, 4GB RAM, 128GB SSD, 14″ Full HD
screen.

Price: JOD 267 (Jordanian dinar 267). Roughly US$375-400.

6. Rabbit R1 AI Companion
An AI companion device, voice, "smart speaker" style interactions, etc.

From Cairo to Amman, homegrown tech innovation is
quietly reshaping how the region connects, studies,
and creates. Therefore, in this magazine edition,
I thought it would be interesting to dig out a few
tech gadgets made in the MENA region. Plus some
international tech that we could not
leave off the list.

PHOTON & FORTUNE

HOW THE FUTURE OF AI MAY RUN ON LIGHT

by João Marques Lima

As computing seeks speed without the heat, the age of light has arrived. João Marques Lima examines how photonic computing could cut costs, curb energy use, reduce the cloud's carbon footprint, and reshape the economics of global data infrastructure.

In the race to sustain ever-growing demands on computing infrastructure, photonic computing is emerging as a credible alternative to purely electronic systems. It exploits light (photons) rather than electrons for data transmission and processing, thereby offering superior speed, lower thermal dissipation, and energy gains. The potential shift is not trivial: if realised, photonic computing could reconfigure the economics of data centres, influence capital investment in chip and network fabrication, and become a cornerstone of sustainable computing.

What is photonic computing, and why it matters

In essence, photonic computing uses light to carry, switch, and sometimes compute information - either in hybrid architectures or, in the longer term, in circuits dominated by optics. Standard electronics are limited by resistance, capacitance and heat dissipation; photons, by contrast, travel at the speed of light and avoid many of those constraints.

According to the World Economic Forum, photons can transmit data faster than electrons and operate in low-energy environments, making them well-suited for intensive computational workloads.

In recent years, researchers have demonstrated prototype photonic processors designed for artificial intelligence (AI) and other high-performance tasks. For example, at the Massachusetts Institute of Technology (MIT), scientists have developed a photonic chip architecture capable of ultrafast AI calculations, pointing to potential throughput advantages where optical techniques may outstrip traditional electronics for specific operations.

Most photonic systems today are hybrid. Control logic, memory, and general-purpose instructions remain electronic, while photons handle the data-heavy tasks such as interconnects or matrix operations. The goal is to remove bottlenecks in speed, energy consumption and heat generation - the key constraints in scaling modern computing.

A greener alternative

One of photonic computing's most compelling promises lies in energy efficiency and thermal management. Conventional data centres expend enormous energy on cooling, often matching or exceeding the energy required for compute operations. Optics, in contrast, generate far less waste heat.

According to Cambridge Consultants, replacing electronic links with photonic ones can increase energy efficiency by up to tenfold and expand bandwidth per link by 10-50 times. Similarly, the UK's Photonics 2035 projection suggests that integrated photonics could reduce total data centre energy consumption by more than half by 2035. Reduced heat output also means lower water usage - a growing environmental concern in regions facing climate
and infrastructure stress.

With fewer active electronic components, resistive losses fall dramatically. Photonic pathways can often be passive or require minimal electrical drive, so power requirements
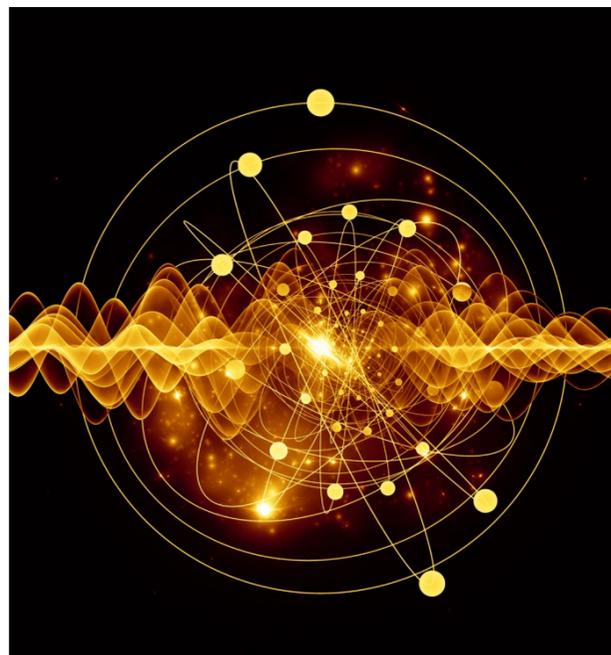
drop significantly. For large-scale data centres and AI clusters, these cumulative reductions in energy demand and cooling overhead could be substantial.

There are, however, balancing factors. Photonic subsystems introduce their own energy loads, particularly from lasers, modulators and optical amplifiers. But at scale, and with improving device efficiency, these costs tend to be lower than those of equivalent electronic interfaces.

The combination of higher speed, lower power, and reduced heat points to a model of greener computing infrastructure - a strategic advantage for operators under increasing regulatory and environmental scrutiny.

Opportunities and challenges

Photonics presents several distinct opportunities in such that in high-performance computing and AI systems, interconnect bandwidth has become a limiting factor. As AI models scale, inter-node communication - such as synchronising parameters or distributing weights - creates bottlenecks. Photonic interconnects address this by providing enormous bandwidth with very low latency.

Optical links are already standard for long-haul and rack-to-rack communication but extending these benefits within the rack - across chips and servers - is the next frontier. Some chipmakers are developing co-packaged silicon photonics, integrating optical transceivers directly into switch chips to reduce energy and distance losses. According to NVIDIA, this approach delivers up to 3.5 times lower power consumption compared with traditional pluggable modules.

However, the technical and industrial challenges remain significant. Photonic systems must overcome issues of manufacturing scale, precision alignment, yield and cost of materials. Integration with existing Complementary Metal-Oxide-Semiconductor (CMOS) processes is complex, and design tools for optical circuits are less mature than for electronics. Components such as lasers and modulators are sensitive to temperature and fabrication variability, requiring new standards for reliability and packaging.

For now, most photonic computing designs target domain-specific functions or acceleration of select kernels rather than full general-purpose workloads. Yet the progress is rapid. Some early pilot systems have reportedly matched or even surpassed existing processors while operating at dramatically lower temperatures. Industry observers expect the technology to begin commercial disruption of high-performance computing within the next five to seven years.

Impact on the data centre market

The data centre market - particularly hyperscale and AI-focused operators - faces severe pressures from power ceilings, cooling costs and capital-intensive network upgrades. Photonic computing offers a potential remedy to several of these challenges.

Power constraints are foremost. Many facilities already draw close to their permitted energy limits. AI workloads, with their heavy computational density, push these boundaries further. By reducing the energy required per bit transferred, photonic architectures enable operators to expand capacity within existing power envelopes.

Cooling and infrastructure costs also stand to benefit. Less heat generation means smaller or more efficient cooling systems, lowering both capital expenditure on chillers and operational costs for air-

handling and pumping. Over a multi-year horizon, these savings contribute meaningfully to total cost of ownership.
Network fabric and latency could see further improvement. Optical links, when tightly integrated, offer minimal signal degradation and faster switching. For large-scale AI clusters, reducing communication delay can be as critical as raw compute speed.

Ecosystem shifts will follow. Established switch and transceiver vendors will need to adapt or partner with emerging photonics specialists. Providers of photonic integrated circuits, hybrid optical-electronic modules, and design services are expected to increase their influence. Hyperscale operators such as Equinix, Digital Realty, and major cloud providers including AWS, Microsoft Azure and Google Cloud may become early adopters, co-developing photonic infrastructure with chip vendors to gain efficiency advantages.

New commercial models could emerge, from modular photonic compute pods to photonic interconnect services offered on a pay-as-you-use basis.

Economics of photonics

From a market perspective, photonics is already a substantial global industry. Grand View Research estimates the worldwide photonics market at US$1.03 trillion in 2025, projected to reach about US$1.3 trillion by 2030, a compound annual growth rate (CAGR) of 4.8%. Within that, specialised segments such as photonic integrated circuits (PICs) are expanding much faster.

The PIC market, valued at roughly US$14.1 billion in 2024, is expected to rise to US$25.8 billion by 2030 - an annual growth rate of around 10.8%. Other consultancies forecast that by 2034 the market could reach close to US$100 billion, reflecting its importance in AI, telecoms, and data centre applications.

Silicon photonics, which allows optical components to be fabricated on conventional silicon wafers, is projected to grow from US$2.59 billion in 2025 to US$8.13 billion by 2030, representing a CAGR of about 26%. Market researchers view this as a core enabler of scalable and cost-efficient photonic devices.

Capital investment in the sector is rising sharply with the European Union recently approving €133 million (US$142 million) in funding for photonic semiconductor pilot facilities in the Netherlands as part of its strategic semiconductor initiative. Meanwhile, in February 2025, European multinational semiconductor contract manufacturing and design company STMicroelectronics announced a partnership with Amazon Web Services to develop photonics chips for data centre transceivers - a segment expected to grow from around US$7 billion in 2024 to US$24 billion by 2030.

Venture capital is equally active. Lightmatter, a US-based photonics start-up, raised US$400 million in late 2024, lifting its valuation to US$4.4 billion. The company intends to integrate photonic chips directly into partner data centres and has signalled plans for an initial public offering.

From a cost perspective, photonic components remain more expensive than mature electronic alternatives, largely due to lower manufacturing volumes and higher precision requirements. However, as fabrication scales and yields improve, the cost per gigabit transmitted or processed is expected to decline sharply - and we have been here before. For large operators, the total-cost-of-ownership (TCO) calculation increasingly favours photonics when energy and cooling savings are included.

wFor a hyperscale facility, the question becomes whether the capital premium for photonic infrastructure is outweighed by deferred upgrades to power and cooling systems, denser compute deployments, and long-term operational savings. As the economics shift, adoption could accelerate rapidly.

What it means for the industry

Ultimately, photonic computing has the potential to reshape supply chains, industrial partnerships and regional strategies. Traditional transceiver and optical-component vendors will need to pivot toward integrated photonics or risk being displaced by newer entrants. Semiconductor foundries may expand into photonic process technologies, while materials such as gallium arsenide, indium phosphide and lithium niobate could become strategically vital.

Systems integrators, cloud operators and data centre providers are likely to move closer to chip vendors, co-designing optical compute fabrics tailored to specific AI or networking workloads. Co-packaging - where photonic and electronic functions share a single substrate - may become standard practice.

Regions investing early in fabrication and R&D could gain a competitive edge. Europe's public funding for photonics pilot lines reflects a clear intent to secure leadership in this emerging domain. Similar initiatives are expected in Asia and North America as governments seek to balance semiconductor resilience with sustainability goals.

In the longer term, photonic computing offers a way to extend Moore's Law beyond electronic scaling limits and to address the environmental costs of digital growth. For sectors such as AI, high-frequency trading, and scientific simulation, the ability to move data faster and cooler could become a defining differentiator.

Adoption, however, will be evolutionary rather than revolutionary. The likely trajectory involves hybrid systems - where photonics accelerates specific data flows - before broader integration into general-purpose computing. Each stage will depend on reliability, cost reduction, and alignment with existing infrastructure.

Should those challenges be met, photonic computing could transition from being seen as an exotic research niche to industrial mainstay, underpinning the next generation of sustainable, high-performance data
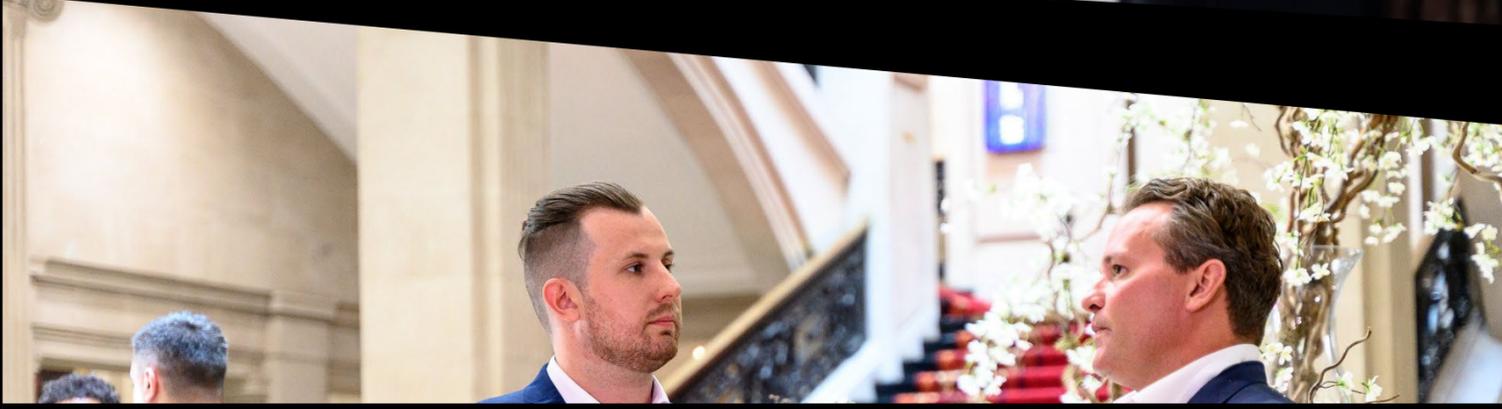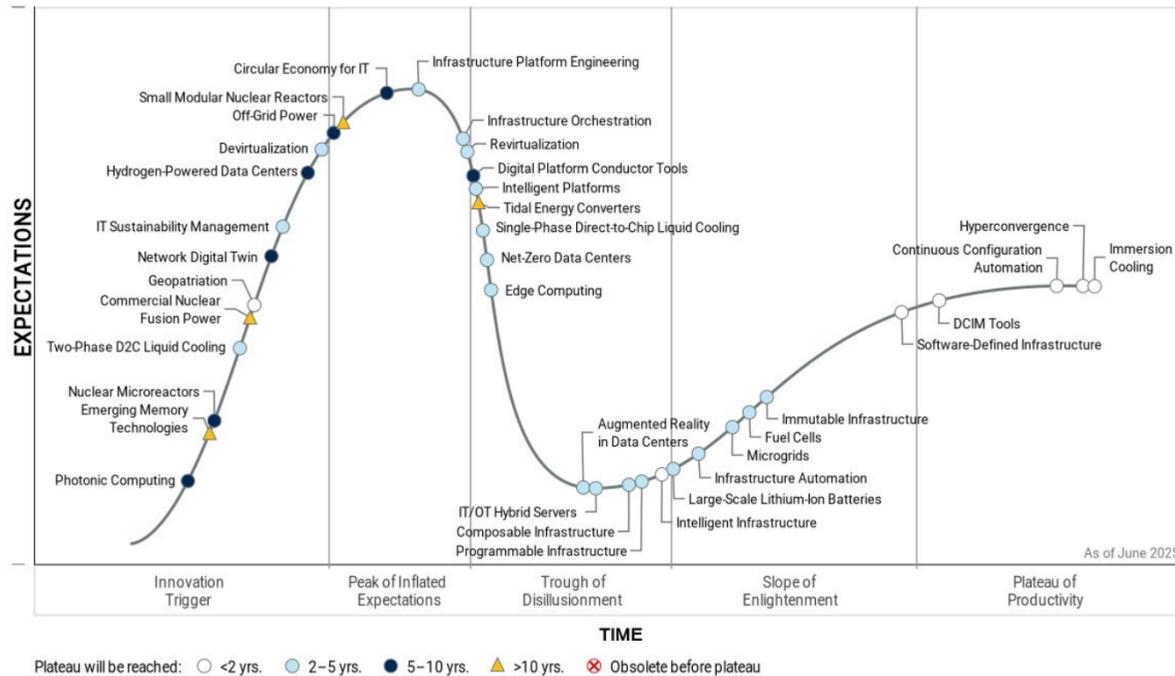
Hype Cycle for Data Center Infrastructure Technologies, 2025