

KEYNOTE INTERVIEW

AI data centres go local



The rapid acceleration of AI is reshaping the architecture of digital infrastructure, say Morgan Stanley Infrastructure Partners' John Ghirardelli and Yacine Saidji

A fundamental shift in how compute is built, deployed and connected is now underway. What was once a highly centralised model – dominated by the concentration of infrastructure in a handful of core markets – is rapidly evolving into a more decentralised, multi-layered system of interconnected mission-critical compute capacity.

As AI workloads scale, the industry is moving towards a framework in which large, centralised facilities are complemented by a growing network of distributed data centres supporting inference, enterprise workloads and latency-sensitive applications.

This shift towards decentralised compute reflects a broader, longstanding dynamic – a natural trade-off between centralisation and distribution seen across infrastructure

SPONSOR

**MORGAN STANLEY
INFRASTRUCTURE PARTNERS**

sectors, according to Morgan Stanley Infrastructure Partners' head of digital for the Americas, John Ghirardelli, and Yacine Saidji, co-head of Europe.

Q How would you characterise the current state of the global data centre and AI infrastructure market?

John Ghirardelli: When you pare it back, we're in the early innings of a multi-year AI infrastructure supercycle where demand is expanding faster than the ecosystem can deliver power, grid interconnections and high-density-capable data centre capacity.

This is creating a "scarcity market" in the most power-constrained hubs. At the same time, the workload mix is shifting: training remains the headline, but inference – especially agentic, orchestration-heavy inference – is increasingly the volume driver, creating constraints in more of the stack (CPU, memory, storage, networking and metro fibre).

The scale at which infrastructure is being built – and the demand behind it – is surpassing most expectations. Capital expenditure from a handful of global players has quickly moved from hundreds of millions annually to recent estimates approaching a trillion dollars.

At the same time, adoption is annually accelerating far faster than in previous technology cycles. The personal computer took many years to



Q How critical is high-capacity fibre infrastructure in enabling the decentralised compute transition?

JG: Fibre is no longer just a supporting utility – it’s a core enabling layer of the digital infrastructure stack. High-capacity fibre networks determine and now enable where compute can be located, how efficiently workloads can be distributed and how effectively data centres operate as part of a wider system rather than as isolated assets. The traditional boundary between data centres and connectivity infrastructure is increasingly blurring, with value creation now sitting in the integration of compute, interconnection and network density. We expect this convergence to intensify further over time.

Against this backdrop, as data centre deployments become more decentralised and distributed – driven in part by ongoing power and supply constraints – the importance of purpose-built, high-capacity fibre infrastructure will intensify, as they will be essential to support this next phase of scaled, distributed compute architecture.

We’ve seen this first hand with our US portfolio company, Lightpath, which is actively building new metro and long-haul fibre infrastructure networks in new regions outside our original footprint to better enable connectivity across data-centre-rich ecosystems. Explosive AI workloads are fundamentally more bandwidth-intensive, therefore creating a structural shift in how compute is deployed and interconnected. AI and hyperscale cloud customers are turning fibre from a “connectivity utility” into a “mission-critical compute backbone”, requiring massive new builds across metro, middle-mile and long-haul simultaneously. This is a key thematic we’re actively pursuing across the globe for new investments.

YS: Our investments in UltraEdge in France and Flexential in the US also reflect our conviction in this thesis. In both cases, we’re backing platforms designed as connected infrastructure ecosystems, where data centres and fibre networks are structurally interdependent rather than separate investment themes.

reach maturity, and even smartphones followed a more gradual curve. AI tools, in contrast, have been adopted far faster – and that’s driving urgency around how quickly compute and transport infrastructure needs to be deployed.

Q What are the key demand drivers and how is agentic AI reshaping the landscape?

Yacine Saidji: As John said, AI workloads can broadly be split into training and inference. Training requires very large amounts of compute and is today concentrated in large-scale data centres, particularly in the US. Once models are trained, however, they’re deployed globally for inference, which is driving a more distributed infrastructure model.

Increasingly, enterprises want to run AI closer to their own data, particularly in sectors where control and confidentiality are critical. In industries such as pharmaceuticals and banking, data sovereignty is paramount. These organisations are often reluctant to place sensitive data in public cloud environments.

Instead, many are building localised AI environments – effectively ‘AI factories’ with dedicated compute pods and supporting infrastructure. In Europe, these typically range from 5MW to 15MW and are designed for single-tenant use. While modest by hyperscale standards, this represents a meaningful shift towards more distributed compute infrastructure.

On the demand side, agentic AI is fundamentally changing usage patterns too. AI is no longer primarily defined by human interaction with tools such as ChatGPT, but increasingly by autonomous systems operating 24/7, with swarms of agents continuously running in parallel.

In some corporate settings, primarily in the tech industry, employees leverage this agentic capability to maximise AI token usage and this ‘token-maxing’ behaviour has transformed AI

from intermittent, human-driven usage to persistent, machine-driven workloads. The result is a structural step change in demand intensity. Inference volumes have risen sharply in a very short period, compounding supply constraints. In my view, this points to a market that will remain structurally tight for years.

Q Which regions across the US and Europe are best positioned to benefit from the surge in decentralised compute demand?

JG: In the US, data centre development is increasingly dictated by where land and power are available and can be brought online quickly. As demand continues to outpace supply, development activity is shifting beyond traditional core hubs into tertiary locations.

Historically, infrastructure and supply chains have been concentrated in established hubs, but supply constraints are pushing development further out, extending timelines and reinforcing the structural tightness in the market.

Similar dynamics have already been seen in Europe, where the Frankfurt, London, Amsterdam, Paris and Dublin (FLAPD) markets are reaching saturation. In the US, activity is also moving away from legacy hubs such as Silicon Valley and Northern Virginia towards emerging markets including Phoenix and Dallas. Dallas, in particular, is set to overtake Northern Virginia after a long period of dominance as the largest data centre market in the country.

YS: In Europe, I would frame the opportunity in two phases. In the first phase, there are currently very interesting opportunities in the Nordic region. Certain pockets still have available power and capacity that can be absorbed by new data centre development. However, this is a finite window – these opportunities will likely be exhausted relatively quickly. Beyond

that, the longer-term phase favours France, which I believe will become a structurally important European data centre market once the initial low-hanging fruit in other regions have been absorbed.

France benefits from a large nuclear fleet, which provides stable, low-carbon baseload power, as well as a central geographical position in Europe. More broadly, factors such as energy security and geopolitical stability are also becoming increasingly relevant considerations.

When deploying billions in critical infrastructure, it's resilience and security that sit alongside power availability as key location drivers.

“High-capacity fibre networks determine and now enable where compute can be located”

JOHN GHIRARDELLI

Q What are the biggest bottlenecks that are constraining supply today?

YS: There are multiple bottlenecks, particularly in Europe. The first is grid connectivity. Most European markets rely on grid connection rather than behind-the-meter solutions. Connecting at scale can take three to four years, creating a significant delay in bringing capacity online.

Even once grid access is secured, there are long lead times for critical

equipment such as transformers, diesel generators and components for liquid cooling systems.

Construction permitting is another major constraint and can add a year or more in some jurisdictions. Together, these factors significantly slow down supply delivery.

JG: Similarly, in the US, the dominant constraint is power availability or time to power. In many markets it can take four to 10 years to secure sufficient power capacity – far longer than the current pace of demand growth.

This is driving experimentation with alternative power sources such as gas generation, solar and battery storage, often co-located with data centres. However, data centres require continuous, highly reliable power. Today, around 95 percent of operating US data centre capacity is grid-connected in some form and that will take time to change, unfortunately.

Additional constraints include electrical equipment shortages, the shift from air to liquid cooling and labour shortages for large-scale builds. Permitting and regulatory approvals are also becoming more complex, as stakeholder engagement becomes more nuanced.

Many companies now have teams working directly with municipalities and local stakeholders – roles that didn't exist a few years ago. This reflects a clear increase in complexity, as well as costs.

Competitive advantage is increasingly shifting towards operators and investors that can secure long-duration power access, navigate local stakeholder environments and deliver integrated infrastructure ecosystems at scale. In many respects, the bottleneck is no longer capital availability – it is execution capability.

Those able to combine technical expertise, local relationships and speed-to-market will be best positioned to capture the next phase of AI-driven infrastructure demand. ■

Disclosures: Morgan Stanley has paid for this sponsorship. The material contained in this article is current as of the publication date of the article, is intended for informational purposes only and does not purport to address the financial objectives, situation or specific needs of any individual reader. It has been obtained from sources believed to be reliable, but Morgan Stanley cannot guarantee its accuracy or completeness. The use of this article is not a solicitation, or an offer to buy or sell any security or investment product. Moreover, the opinions expressed in this article are not necessarily those of Morgan Stanley or its employees. Past performance is no guarantee of future results, and there is no guarantee that any of the investments listed above will perform well in the future. The views and opinions and/or analysis expressed are those of the author or the investment team as of the date of preparation of this material and are subject to change at any time without notice due to market or economic conditions and may not necessarily come to pass. Furthermore, the views will not be updated or otherwise revised to reflect information that subsequently becomes available or circumstances existing, or changes occurring, after the date of publication. The views expressed do not reflect the opinions of all investment personnel at Morgan Stanley Investment Management (MSIM) and its subsidiaries and affiliates (collectively "the Firm"), and may not be reflected in all the strategies and products that the Firm offers. This material is a general communication, which is not impartial and all information provided has been prepared solely for informational and educational purposes and does not constitute an offer or a recommendation to buy or sell any particular security or to adopt any specific investment strategy. The information herein has not been based on a consideration of any individual investor circumstances and is not investment advice, nor should it be construed in any way as tax, accounting, legal or regulatory advice. To that end, investors should seek independent legal and financial advice, including advice as to tax consequences, before making any investment decision. RedOak 5531322, Exp. June 5, 2027 © 2026 Morgan Stanley