



Five things we learned at Datacloud Global Congress 2025

We were at Data Cloud Congress 2025 in Cannes last week to evaluate the key narratives among data centre operators, investors and vendors. Here are our five key takeaways from the event.

Jonas Topp-Mugglestone, Consultant

Last week, 4,000+ attendees from 65 countries gathered in Cannes for **Datacloud Global Congress 2025**. While the agenda featured no shortage of optimism around AI-driven growth, the conversations among data centre operators, hyperscaler tenants, vendors and policymakers revealed a more complex picture. Below the headline figures and bullish growth forecasts, there is a growing recognition that the industry is approaching a critical inflection point: scale is accelerating, but questions remain around energy availability, efficiency, and a dynamic customer base.

1. Infrastructure buildout accelerates, future utilisation remains unpredictable

AI dominated Datacloud 2025 in every corner of the Palais des Festivals. The focus was not only on the scale of current deployment, but on the growing inevitability that data centres are fast becoming the infrastructure skeleton of the AI age – supporting what many now see as the beating heart of innovation across sectors, from healthcare and science to logistics and finance.

One hyperscaler shared that capacity had doubled annually globally for the past three years, with a 40% uplift planned in Europe alone by 2026. The industry is building fast, but demand continues to outpace delivery. As one panellist noted, “there’s no celebration when a project is complete - there’s usually a customer screaming and a schedule slipping”. The quote reflected a common theme: a growing pipeline of capacity under construction is being matched by relentless customer demand.

Billions are pouring into land, energy, and supply chain resilience, with data centre developers racing to keep up. One hyperscaler executive noted that “no one can build a data centre faster than anyone else, you just start earlier”. Even with significant capital deployed, non-monetary factors continue to be key blockers to facility delivery, with constraints on power transmission, skilled labour, and permitting continue to throttle delivery.

However, when there’s a will, there’s a way, and data centre operators should be cognisant of competing against the proprietary builds of the likes of **x.AI’s Colossus facility**, which was turned from a brownfield shell to ready for service data centre in just 122 days. Combined with Groq’s 51 day contract-to-live turnaround of their **LPU-based IaaS solution** for Saudi Arabia’s new AI leader, Humain, there are real routes to market in under 6 months for new entrants, eroding the perceived temporal barriers to greenfield market entry.

Much of today’s investment is predicated on the long-anticipated shift from AI training to inference. One global AI facility leader, for instance, has reoriented around GPU-heavy inference workloads, betting that this more scalable, transactional phase of AI will dominate compute demand for years to come, and one keynote visual captured this sentiment by forecasting the inference/training split at over 90% for all segments of the enterprise market. However, despite some early proof points, this is largely speculative, and widespread AI adoption faces significant cultural barriers across industries. Data centre operators need to look no further than their own facility operations when it comes to cultural blockers to the dynamic cross-domain facility optimisation which AI could enable.

With this uncertainty comes a cautionary term used at the opening keynote: “Braggawatts”. The phrase describes inflated, often unsubstantiated capacity announcements that grab headlines but may never materialise – hscale’s recent public launch is just one example of this – their **gigawatt ambitions matched by 100MW under construction and just 6MW operational**. With land, labour, and power all in tight supply, and several speakers arguing that energy transmission, not generation, is increasingly proving the real constraint, some questioned whether these bold announcements were grounded in reality.

Despite surging demand and investment, no one has a crystal ball when it comes to AI adoption trends. It’s far easier to sell promised returns to investors than to predict the precise shape of future demand, and markets that misjudge their demand forecasts may end up overbuilt and underutilised.

2. Power gets the headlines – but the interconnect angle is key

Power generation and distribution dominated the narrative as a key limiting factor, but several forward-looking panels honed in on a different, underappreciated bottleneck: connectivity. As data centre buildouts accelerate and AI workloads diversify, the assumption that power alone is the primary constraint is beginning to crack.

A CTO of a leading global colocation provider articulated this shift clearly: “Agentic functions and reasoning elements that are emerging are resulting in lots of hype around traditional data centre focus areas like power – but the interconnect is the one key element which is being overlooked.”

A case study in shifting network expectations came from the mobility sector. With the gradual rollout of autonomous fleets, mobility & logistics providers anticipate an uplift in data generation across a dispersed and dynamic footprint. Fleets of vehicles will generate video, mapping and sensor data continuously – pushing new volumes of traffic from previously unconnected or lightly connected geographies. Even a limited shift towards driverless cars could require significant upgrades in how traffic is aggregated and routed to centralised infrastructure, even accounting for a significant portion of workloads handled through compute built into these autonomous fleets.

Throughout the event, there was speculation about this trend playing out more broadly. AI workloads are no longer confined to bespoke hyperscale facilities designed for AI training – instead they are increasingly decentralised through a network of interdependent facilities with varying degrees of AI readiness. As one panellist remarked, “AI is now your customer”, a reminder that infrastructure must now serve not just users, but also systems acting independently of them. This transformation is accelerating with the rise of AI agents – systems that act with interconnected autonomy, rather than relying solely on a centralised architecture. As compute becomes more distributed and interdependent, traditional assumptions about how and where data flows are being upended. The result is a connectivity landscape that’s more complex, more fragmented, and without the weight of industry analysis going into power generation and distribution.

As data centre growth extends into new geographies, driven by factors like access to power and local digital transformation, the geographical distance between where data is collected, analysed and outputted is widening, as well as the complexity in the service provider ecosystem delivering such insights to enterprises. Addressing this will require forward planning across the value chain, from network service providers (NSPs) to data centre operators, and a more collaborative approach to network planning. For further insight, see [STL Partners’ article on the implications of AI for NSPs](#).

3. Public perception is shaping delivery and timelines

The conversation around siting and permitting has shifted from a regulatory formality to a material delay risk. Across Europe, attendees acknowledged that local opposition could delay projects by months or even years – threatening timelines, driving up costs, and hitting IRRs. As STL Partners has shown, even [modest delays can erode margins and derail investment cases](#), to the tune of \$14.2 million per month of delay for our indicative 60MW build in the US. Building community trust early is now a financial imperative to ensure operators do not have to pay the price of such delays.

This was underscored by [new research presented by CyrusOne](#) surveying 13,000 people across the UK, Ireland, Netherlands, Spain, Italy, France and Germany. While 93% of respondents across Europe reported neutral or positive views on data centres, just 52% could accurately define what one is. The disconnect was most pronounced in the UK, where only 38% made the correct definitional choice from a selection of options. Among 16–24-year-olds, fewer than a third recognised the role data centres played in their day-to-day activities.

One speaker captured this public sentiment in a particularly poignant way, stating that “people think the internet runs on magic”. But this lack of awareness is no longer benign – it’s becoming a material risk to growth.

The sector has been pushed to the fore following years of relative public anonymity, and it faces a public perception problem. With growing media attention, environmental scrutiny, and land-use tensions, community sentiment can make or break a project. Nimbyism remains a challenge, but several presenters argued that it can be mitigated, if not reversed, by genuinely engaging local communities.

One data centre operator shared an example of how biodiversity concerns had emerged as a core issue during local consultations. In another case, job creation became the linchpin for support. In both, meaningful dialogue early in the planning cycle helped avoid costly resistance later, through simple (and cost-effective) action such as changing a specific type of brick and more effectively communicating job creation to the local community through engagement programmes.

A panellist’s mantra, “Ask, Listen, Act”, perfectly captured the need for a broader shift in tone from compliance to collaboration. While regulation, such as the UK’s **designation of data centres as critical national infrastructure**, can help streamline permitting and design stages, proactive community engagement remains one of the most cost-effective ways to prevent local resistance and avoid costly delays.

4. The industry is committed to sustainability – but questioning the policy playbook

Sustainability was ever-present in Cannes, but this year’s conversation struck a more mature and, at times, combative tone.

On the one hand, the industry remains confident in its trajectory. Hyperscalers continue to lead the charge with carbon-free sourcing strategies – one hyperscaler reaffirmed its pursuit of “additionality” in clean energy procurement,¹ and others pointed to the sector’s broader track record in catalysing investment into renewables. Indeed, a recent state of the market report from the **European Data Centre Association** indicated that 94% of the European data centre sector is powered by renewable energy sources. As one speaker put it, “No industry has done more to facilitate renewable energy than the data centre industry”.

Despite political shifts in the US and evolving global priorities, panellists agreed that market forces still overwhelmingly favour sustainable energy solutions. Corporate customers and investors continue to treat sustainability credentials as key – carbon footprint and power sourcing are now baseline requirements from prospective tenants, not differentiators.

Bullish attitudes to the sustainability imperative was tempered by growing frustration with regulatory dynamics – particularly in the EU. Several attendees criticised the bloc’s sustainability reporting framework as burdensome, poorly targeted, and at odds with real environmental performance. Sub-500kW facilities, often the least efficient, with average PUEs of 5 and peaks above 11 according to one panellist, are exempt from the most stringent regulations, while larger operators face stringent efficiency targets including, but not limited to, PUE – a much-maligned metric in its usage and comparability which many argue is outdated.

“There’s no harmony across member states”, one speaker noted. “The implementation was rushed, and the burden pulls resources away from where they could drive real change.” Another from the panel offered a more colourful summary: “It’s easy to shoot elephants. The rats are the problem.”

¹ “Additionality” is carbon crediting terminology that refers to the fact that a project’s emission reductions or removals would not have happened without the carbon finance provided by the credits

The underlying message was clear: the industry supports robust sustainability goals – but it wants smarter, more targeted regulation that incentivises outcomes rather than administrative box-ticking.

Still, the consensus is that the green imperative is here to stay in the eyes of private investors and many regulators, and operators should continue to invest in sourcing sustainable energy and optimising facility efficiency across a diverse array of legacy and new metrics, such as PUE, WUE and ERF.

5. Infrastructure design norms are in flux – and flexibility is the only constraint

Perhaps the clearest takeaway from Cannes: data centre infrastructure is evolving faster than the frameworks used to design and operate it – and few feel confident they're building to the right long-term assumptions.

The CEO of a European colocation provider summed up the moment: "More has happened in the last two years than the twenty before, especially around cooling and power." If Nvidia's forward guidance is stuck to, the transition to GPU-heavy rack configurations designed for AI workloads is expected to push rack densities to **600kW by H2 2027**. Legacy cooling, power and networking solutions are simply not up to the challenge of a 100x increase in facility power density, and the associated impacts on thermal optimisation and network configuration.

Liquid cooling is rapidly progressing from proof-of-concept to scaled deployment. However, as several speakers noted, standardisation is still lacking – and the air-to-liquid ride-through mismatch is a real operational challenge. Today's most advanced deployments are moving from rear-door heat exchangers to direct-to-chip (D2C) cooling, and in some test cases, full immersion. Yet each method comes with trade-offs in complexity, maintenance, and failover protection. Liquid typically provides just 30 seconds of ride-through during an outage, compared to up to two minutes with air. For operators, balancing both systems effectively is becoming non-negotiable.

Flexibility has become the new design principle. Developers are increasingly delaying key infrastructure decisions to keep pace with rapid shifts in silicon and reference architectures. As one panellist noted, many currently design around peak loads that "may only occur for 37 minutes per year". But this approach is prompting debate – especially for non-real-time AI workloads like model training. Does every site need to build component-level redundancy for a momentary peak capacity, if innovation in IT solutions can permit momentary downtime? For certain facilities hosting a select few workload types, overprovisioning may no longer be a sign of resilience, but a drag on efficiency.

One CTO of a greenfield operator serving hyperscaler customers went further: "AI agents must be embedded to succeed" – not just in customer use cases, but within the facility itself. Agentic AI systems monitoring and managing infrastructure in real time remain the stuff of dreams as facilities prioritise uptime and trust above marginal efficiency gains. However, innovation in this space is expected to achieve widespread adoption quickly once a single established operator can demonstrate that closed-loop, cross-domain automation can be balanced with trust and uptime assurance. Allaying existing fears that such a solution introduces a single point of failure and additional attack surface will be crucial to enabling widespread adoption.

And yet, as the industry scrambles to adapt, the gap is widening. Hyperscalers, armed with scale, purchasing power, and end-to-end visibility from facility to application, are setting a new bar for efficiency. Self-reported hyperscaler PUEs are just 1.15 according to the latest **EUDCA report**, a marked difference with figures of 1.39 for colocation facilities and 1.85 for proprietary enterprise facilities – all of which feel low in comparison to true industry averages. They're also shaping rack reference architectures and driving wholesale build strategies across the market. In contrast, colocation providers without deep ecosystem integration or long-term partnerships with anchor tenants often lack the same incentives, and ability, to optimise across both

facility and IT infrastructure. Without direct influence on hardware roadmaps or early visibility into hosting and enterprise IT trends, they risk falling behind.

Retaining flexibility to adapt to changing market pressures is paramount for data centre operators to future proof their business models for resilience to dynamic pressures coming from the enterprise IT industry. The only certainty is change, and those without a strategy designed for agility risk being locked out of the next wave of growth opportunities in data centres.

Final thoughts

The dawn of the AI adoption is undeniably an inflexion point for the data centre sector, but beneath the growth lies growing complexity. As Datacloud 2025 made clear, infrastructure is no longer commodity infrastructure to be built and leased – it's a dynamic, evolving building block enabling the AI ecosystem, and one that must flex with xPU innovation, unpredictable customer workloads, and regulatory scrutiny that shifts as fast as AI innovation further up the stack.

To navigate this new dawn, data centre operators need more than capital and construction velocity. They need modularity in design to accommodate new rack architectures. They need to shift from facility monitoring to dynamic optimisation in operations, embedding automation intelligently without compromising trust and uptime. They need to anticipate demand in emerging digital markets – and be willing to engage customers to stimulate this demand. And they need a commercial strategy that not only aligns with their facility specifications, but supports long-term investment returns in a world of volatile AI adoption curves.

Across the board, customer relationships will need to evolve. Data centre operators must be ready to take on a more consultative and advisory role, helping tenants optimise their IT infrastructure estate to optimise against a backdrop of performance, cost, and sovereignty constraints. That means guiding them on how your site(s) fits into a wider strategy, whether this is through commercial and technical redundancy, guaranteed GPU availability, or sustainable and cost-effective hosting models.

At STL Partners, we're working with stakeholders across the ecosystem to navigate these transitions – from evaluating shifting demand patterns for corporate development to rethinking infrastructure design, and aligning commercial strategy with tomorrow's customer requirements. For data centre operators, tenants and investors alike, the path forward will require sharper foresight, closer collaboration, and a willingness to rethink assumptions in an industry where the only constant is change.

Jonas Topp-Mugglestone is a Consultant at STL Partners, specialising in data centres.

Get in touch with the author to learn more

jonas.topp-mugglestone@stlpartners.com

Or visit STL Partners' Data Centres Hub

www.stlpartners.com/data-centres