


BUILDING FOR AUTONOMY

Why Agentic AI Demands a
New Kind of Infrastructure



INTRODUCTION



Agentic AI isn't just another workload.

Unlike traditional AI models, autonomous agents generate unpredictable demand patterns, run sprawling multi-agent workflows, and evolve rapidly with new frameworks and silicon refreshes. As a result, data centre infrastructure is being pushed far beyond its traditional limits.

96% of enterprises plan to expand their use of AI agents within the next 12 months, with over half targeting organisation-wide deployment. This means infrastructure must become more flexible, responsive, and resilient than ever before.

This is why working with trusted partners matters more than ever. Operators don't just need added capacity. They need support to rethink modularity, connectivity, and lifecycle management. The right partners bring the expertise to help design infrastructure that is ready for agentic AI today and scalable for tomorrow.

In this viewpoint, we explore five critical shifts data centres must make to support agentic AI – from modularity and edge deployments to AI-optimised hardware, lifecycle planning, and coordinated operations.

1. FROM FIXED SYSTEMS TO INFRASTRUCTURE THAT ADAPTS

Agentic AI workloads don't scale in neat, predictable ways. A new agent, workflow, or model update can spike compute demand overnight.

Traditionally, “modularity” meant containerised builds that could be bolted together to expand capacity. That sped up deployment but didn't make sites flexible. Once workloads shifted, operators were stuck with static blocks that couldn't adjust.

Today, modularity needs to mean something different: interchangeable IT, power, and cooling blocks, all connected by a cabling foundation that allows plug-and-play upgrades. This new modularity gives operators the flexibility to keep pace with agent-driven demand, adding capacity in smaller, faster steps, and refreshing silicon without tearing down entire sites.



Modularity today can help operators overcome key challenges:

- 01.** Add capacity fast, in small increments, adding pods in weeks, not months, and adapting to grid delays.

- 02.** Isolate extreme rack densities, enabling racks that exceed 100 kW while avoiding overbuilding entire halls.

- 03.** Deploy multiple cooling approaches side by side, scaling liquid-cooled pods for extreme densities while retaining air-cooled zones for lighter loads.

- 04.** Support rapid hardware refreshes, swapping in new components without site-wide downtime.

- 05.** Enable low-latency network locality for multi-agent workflows.

Without this redefined modularity, data centres risk downtime, stranded capacity, and premature refreshes, all incompatible with the fast, iterative nature of agentic AI.

2. FROM CENTRALISED TO EDGE-FOCUSED DEPLOYMENTS

Speed is non-negotiable for autonomous systems.

Agents managing real-time operations – whether within IT operations or production lines – can't wait for data to travel halfway across the globe. This makes the edge an essential part of agentic AI infrastructure. Edge data centres bring compute closer to the source, cutting latency and enhancing security by keeping sensitive information local. Analysts estimate the global edge data centre market will reach \$317 billion by 2026, driven primarily by AI demand.

But success at the edge hinges on three key areas:



Stable power delivery in environments where local grids are often fragile.



Cooling systems that can absorb the unexpected heat loads of AI accelerators – like GPUs or TPUs – in facilities never designed for them.



Cabling designed as a foundation, not an afterthought. Too often, hardware gets budget priority, while the connectivity that ties systems together is squeezed into “what's left.” The result is costly, disruptive retrofits when refreshes arrive after 12 months.

When these three elements are aligned, the edge becomes a powerful AI accelerator for autonomous systems. When they're neglected, deployments stall before they start.

3. FROM GENERAL-PURPOSE TO HIGH-PERFORMANCE, AI-OPTIMISED HARDWARE

Agentic AI pushes infrastructure harder than traditional enterprise workloads.


A single multi-agent workflow may orchestrate thousands of GPUs, CPUs, and interconnects. General-purpose facilities can't cope with that scale.

Operators must support high-performance AI accelerators, low-latency networking, and advanced cooling. With hardware refresh cycles now as short as 6–12 months, sites need to integrate new hardware without costly rebuilds.

Generative and agentic AI also stress infrastructure differently. Generative AI relies on massive, centralised GPU clusters with extreme cooling. Agentic AI tolerates higher latency but depends on dense, low-latency interconnects across many servers, often in smaller edge or colocation sites.

That makes cabling critical. High-bandwidth links are the backbone of multi-agent workflows, yet many facilities weren't built to scale fibre at this density. Without structured, forward-looking strategies, operators risk bottlenecks and stranded compute. Increasingly, sites will need parallel cabling designs – one optimised for centralised generative clusters, another for distributed agentic workloads.

For this reason, data centres can no longer be designed as general-purpose IT boxes. They must be specialised, high-performance, and built for autonomy.



4. FROM ONE-OFF CONSTRUCTION TO LIFECYCLE PLANNING

Agentic AI isn't a one-time build. It forces operators to think in decades, not years.

A data centre designed as a static asset will become obsolete as AI accelerators evolve, cooling standards change, and customer needs shift.

Lifecycle planning means approaching design holistically from the start. Power, cooling, and cabling can't be treated as silos. They must be engineered as interdependent systems that can evolve together. It also means planning for "day two" realities: what happens when tenants leave, when demand spikes, or when silicon refreshes come faster than expected.

Most facilities weren't designed with spare space to swap out cables at scale. That means upgrades can trigger massive rewiring, or even force new builds. Lifecycle planning must factor cabling in as a priority. Sustainability is another dimension. Agentic AI's hunger for GPU racks risks driving up energy use just as ESG scrutiny intensifies. Lifecycle planning ensures facilities can extend their lifespan, minimise waste, and remain compliant under tighter sustainability regulations.


5. FROM DISCONNECTED TEAMS AND TOOLS TO COORDINATED OPERATIONS

The complexity of agentic AI makes siloed operations unsustainable.

Power, cooling, networking, and IT teams can no longer work alone. Instead, operators need coordinated operations: integrated planning, unified monitoring, and aligned vendor strategies.

This shift is as much cultural as technical. Teams must blend infrastructure expertise with AI analytics to anticipate shifting workloads. Vendor fragmentation must give way to tighter partnerships, ensuring cabling standards, hardware refreshes, and cooling strategies stay aligned over time.

There's also a skills gap. Cabling, racking, and installation are increasingly undervalued compared to AI software roles. But when undertrained contractors treat fibre and power as plug-and-play, the result is poor workmanship, harder maintenance, and reliability issues at scale. Only coordinated operations can keep pace with the demands of agent-driven infrastructure.





SUMMARY

Building for Autonomy

Agentic AI is rewriting the data centre playbook. Fixed systems and siloed operations won't cut it. The new infrastructure must be modular, distributed, AI-optimised, lifecycle-aware, and operationally coordinated.

But operators can't go it alone. The complexity of modern infrastructure means they need global infrastructure partners that can act as extensions of their teams – experts in power, cooling, cabling, and local markets.

Partners can provide:

- Holistic design perspectives
- Expertise across diverse environments and regulatory regimes
- Continuity through rolling refreshes and unpredictable demand
- Risk mitigation through coordinated planning and execution

No operator can afford to navigate agentic AI's demands solo. Those who build infrastructure today – with global partners – will be best positioned to scale and deliver on demand.

ABOUT ONNEC

Onnec is a leading Infrastructure Solutions and Services company for tech and enterprise, specialising in structured cabling, managed services, and network solutions. Our team of experienced designers, project managers, and engineers, supported by world-class vendor partnerships, delivers top-tier services and solutions.

Onnec's expertise spans all environments and can support customers with:

- Structured cabling design and installation
- Installation of cabling, ODFs, PDUs and containment solutions
- Network hardware installations, changes and support
- Connectivity and equipment upgrades and changes
- Smart Hands support services



Contact us:

www.onnecgroup.com

info@onnecgroup.com

020 3929 9915