



Scaling AI across the UK: why regional connectivity is becoming critical to the next phase of AI growth

 enquiries@neosnetworks.com

 neosnetworks.com



Contents

Executive summary	3
The UK's AI opportunity is becoming regional	5
Regional infrastructure is becoming critical to AI performance	6
Network conditions are already reshaping AI plans	7
Rising capacity demands will intensify the pressure	9
The service-provider view: pressure across the network	10
What we need to become regionally AI ready	12
Building the connected foundations for regional AI growth	13
Conclusion	14
Methodology	15



Executive summary

The UK is moving from AI ambition to infrastructure delivery. But the networks beneath that growth are already shaping where AI can scale.


Ninety-two per cent of UK enterprises say local connectivity influences where they deploy or expand AI workloads, while 38% have already kept workloads more centralised than planned because of network conditions.

Across financial services, manufacturing, life sciences, media and energy, AI is already being implemented in live operations. The experimentation stage seems to be over, and discussions are now about building the infrastructure to scale it.

The big question now is whether every part of the UK is geared up to support that expansion. Or whether the infrastructure reality risks repeating a familiar story in the UK: one about regional disparity.

Having surveyed large enterprises and communication service providers across the UK's key hubs, Neos Networks' research suggests the UK has credible foundations for regional AI growth. Enterprise respondents identify multiple regions as strategically important, while 78% believe their current connectivity is mostly or fully prepared for a two-to-three-times increase in AI workloads over the next year.

However, the survey also reveals a regional readiness paradox: organisations may feel prepared for near-term growth, but network conditions are already causing them to change deployment models, centralise workloads, move activity between regions, reduce project scale or delay implementation.



Nearly half (47%) of enterprise respondents say the AI applications they run need to be deployed in the same UK region as the end users. This includes private AI applications built on corporate data, model training environments and latency-sensitive inference services, where distance between workloads and users can affect the latency and performance required. Indeed, around a fifth (22%) of enterprises said deployment is needed within the same metro area. Only 22% of enterprises said workloads can operate anywhere in the UK.

Local connectivity is consequently becoming a key determinant of where AI lands.

When asked how a region's network infrastructure has affected their AI plans over the past 24 months, organisations reported a range of compromises. Thirty-seven per cent have changed provider, 36% have used cloud instead of regional or edge deployment, 34% have moved workloads to another UK region, 32% have reduced deployment scale and 29% have delayed deployment.

These pressures seem likely to intensify as AI becomes more distributed. Seventy-four per cent of enterprises expect a shift towards inferencing, every respondent expects AI to increase network capacity requirements, and 61% of service providers believe their most demanding workloads will require connectivity of 400Gbps or more.

The new digital divide is no longer between the organisations that do and don't use AI. It's now between the regions that can and can't support production-scale AI in the locations where enterprises need it, and those where infrastructure forces workloads to be centralised, relocated, reduced or delayed.



The UK's AI opportunity is becoming regional

The UK's AI opportunity is no longer confined to one national hub. Unsurprisingly, London will play a key role, as 42% of enterprise respondents picked the capital when asked which region is most relevant to their organisation's AI strategy over the next two to three years.

Yet the data points to a more dispersed opportunity for the UK. North West England was selected by 27%, the Midlands by 26%, East of England (including Cambridge) by 24%, and Scotland by 24%.

Respondents were permitted to select multiple regions, so these figures should not be read as a closed regional split. Instead, they emphasise that multiple regions carry strategic relevance for AI activity and raise the question: can regional infrastructure support AI growth at scale?


The regional spread also varies by sector. Media organisations are the most London-centric, with 54% identifying the capital as most strategically important. Manufacturing organisations show a different geographic pattern, recording their strongest interest in Scotland (34%) and North East England (32%).

This points to the idea that the UK's AI opportunity will not develop around a single national cluster. Different industries are more likely to drive distinct regional AI growth corridors, reflecting the locations of businesses, operational sites, skills and sector expertise.

That reinforces the case for more geographically distributed infrastructure investment. But identifying where AI demand could emerge is only part of the picture. We also need to ask whether those regions are currently equipped to support it as workloads move into production and begin to scale.

Enterprise assessments of current readiness are also broadly positive. Northern Ireland, the Midlands, North East England and London were all rated somewhat or very well equipped by between 58% and 60% of respondents, while 78% believe their current connectivity could support a two-to-three-times increase in AI workloads over the next year. The tension is that this confidence exists alongside clear evidence of infrastructure-driven compromises – a paradox the remaining findings explore.

Organisations feel prepared for near-term growth but are already making compromises



*Between 58% and 60%

Regional infrastructure is becoming critical to AI performance

The importance of regional infrastructure becomes clearer when looking at the AI workloads enterprises are operating or planning themselves, including private AI models built on corporate data, model training environments and latency-sensitive inference services.

Nearly half of enterprise respondents (47%) say these workloads typically need to be deployed within the same UK region as their end users to achieve the required latency and performance. A further 22% say deployment is needed within the same metro area, and only 22% say workloads can operate anywhere in the UK.

This doesn't mean every AI application needs to be hosted locally. Indeed, public AI services and AI applications may be delivered from infrastructure outside an organisation's direct control. The finding instead reflects the growing range of AI workloads where organisations have more control over deployment architecture and where being close to users, operational sites, corporate data and compute infrastructure matters.

Only 5% say these workloads need to sit entirely on-site, while just 2% say they're comfortable with moving them offshore. Instead, for many UK enterprises, the required level of proximity sits somewhere between the two: at regional or metro level.

These findings once again draw out that AI readiness is shaped not only by the availability of compute, but by the infrastructure connecting users, operational sites, cloud platforms, data centres and edge environments within and between UK regions.

In short, location matters. Enterprises need networks that allow AI workloads to operate sufficiently close to the users, data and applications they support, while connecting effectively into wider cloud, data centre and national infrastructure.



Network conditions are already reshaping AI plans


One of the clearest findings from the research is that network infrastructure is already affecting enterprise AI plans.

Respondents were asked how network infrastructure in a UK region has affected their organisation's AI plans in the past 24 months. Only 3% say it had no impact.

For organisations looking to distribute AI workloads closer to users, data or operational sites, the availability of suitable regional infrastructure can limit the choices available. A location might have the right land, planning permissions and power required for new AI capacity, but still lack the high capacity optical connectivity or necessary network resilience to make that location viable.

For enterprises, the consequence can be fewer suitable regional colocation or GPU-as-a-service options, leaving them with little choice but to keep workloads in larger, more centralised environments.

That pattern is visible in the survey findings. Thirty-eight per cent say network constraints have kept workloads more centralised than originally planned. Meanwhile, 37% have changed their infrastructure provider, 36% opted for the cloud instead of regional or edge deployments, and 34% moved workloads to another UK region. Another 32% reduced the scale of their deployment, while 29% delayed it entirely.



92%

of enterprise respondents say local connectivity infrastructure availability is influential when their organisation makes decisions about where to deploy or expand AI workloads



This doesn't mean that UK regions are failing. Instead, it shows how connectivity has become a critical factor in enabling AI adoption at scale.


Network infrastructure conditions are forcing organisations to rethink where and how workloads are deployed. As a result, they're centralising workloads, shifting deployment models, changing providers, moving workloads between regions, reducing scale or delaying projects.

The effect is particularly pronounced in financial services, where half (50%) of respondents say network constraints have kept AI workloads more centralised than planned. Given the sector's demanding performance, resilience and governance requirements, connectivity has fast become the deciding factor in deployment.

This demonstrates how the digital divide could calcify. If certain regions cannot support AI locally and at production scale, organisations within them may be pushed into operational compromise.

Considering that 92% of enterprise respondents say local connectivity infrastructure affects decisions about where to deploy or expand AI workloads, and 36% say it's very influential, it's clear that companies are now choosing their AI locations based on connectivity rather than treating it as an afterthought.

Organisations reporting how regional network infrastructure has affected their AI plans



Rising capacity demands will intensify the pressure

Operational pressure is already heightened in enterprise AI plans, and it's only likely to intensify as workloads become more demanding and dispersed.

Almost three-quarters (74%) of enterprises expect a significant or moderate shift towards inferencing in their organisation's AI workloads across the network over the next two to three years. And service providers are in near-unanimous agreement: 96% expect a shift toward inferencing or edge deployment.

This is critical because inferencing often moves processing closer to the action. It puts processing power directly alongside users, applications, devices, operational sites and regional infrastructure.


With proximity already a key factor and network infrastructure determining deployment decisions, this predicted inference shift makes regional readiness hugely important.

Pressure on existing network infrastructure is also mounting. In fact, every enterprise respondent expects AI workloads to increase their network capacity requirements over the next three years.

And powering these heavy workloads requires serious network muscle. When asked what level of high capacity connectivity their most demanding AI workloads will require over the next two to three years, nearly all (97%) of enterprises said they expect connectivity of 100Gbps or more. This included 44% who expect 400-799Gbps and 4% who expect 800Gbps or more.

For service providers, the jump is even bigger. Asked the same question, 61% said they expected 400Gbps or more, including 20% who expect 800Gbps or more.

With infrastructure already forcing enterprises to centralise or delay deployments, the combination of rising capacity and shifting inference is placing immediate pressure on regional readiness.



The service provider view: pressure across the network

The enterprise findings show how connectivity is already influencing AI deployment decisions. The service provider research provides the other side of the picture: how networks will need to evolve as customer demand becomes more distributed and capacity-intensive.

Today, the delivery of AI-intensive services remains relatively concentrated. Private-cloud colocation is the most common primary environment, selected by 40% of service providers, while a further 15% primarily support workloads through private-cloud infrastructure in their own data centres. By comparison, only 6% identify edge or network-edge locations as the main environment in which they currently deliver AI-intensive services.

However, service providers expect that balance to change quickly. Almost all (96%) anticipate a shift towards inferencing or edge delivery over the next two to three years.


This won't create pressure in one isolated part of the network. When asked where they expect AI-driven workloads to generate the greatest traffic growth, service providers were divided across edge and access networks, selected by 30%; core and backbone networks, at 27%; data centre interconnect, at 25%; and metro networks, at 18%.

The relatively narrow spread is significant. It suggests AI demand will affect the network end to end. Edge and access infrastructure will support applications operating closer to users, devices and operational sites. Metro networks will connect those locations to regional cloud and data centre environments. Data centre interconnect will support the movement of data and workloads between compute facilities, while core and backbone networks will carry growing volumes of traffic between regions and national or international infrastructure.

Regional AI readiness cannot therefore be delivered through an edge strategy alone. Nor can it be solved simply by adding capacity between the country's largest data centres. It requires connectivity across the full network: from local access and metro infrastructure to national backbones and data centre interconnect.

The scale of expected demand reinforces that requirement. Sixty-one per cent of service providers expect their most demanding AI workloads to require connectivity of 400Gbps or more within the next two to three years, including 20% anticipating requirements of at least 800Gbps.

Where service providers expect to see the greatest AI-driven traffic growth in their networks



For network operators, the challenge is not only meeting a predictable rise in average traffic. It's preparing for concentrated and potentially uneven demand around regional data centres, cloud connections, enterprise locations and high-volume inference environments. This also means network capacity cannot be planned separately from the physical infrastructure supporting AI.

Ninety-six per cent of service providers say data centre power capacity and network capacity are closely linked when planning for AI infrastructure growth. More than half (51%) say the two are very closely linked. Power may determine how much computing infrastructure can be supported at a particular site. But without sufficient network capacity, that compute cannot interact effectively with other data centres, enterprise locations, cloud platforms or end users.

Data sovereignty and local infrastructure add another consideration. Ninety-seven per cent of service providers say these issues are important to customers deploying AI workloads on their networks, including 54% who say they're very important.

As AI becomes embedded within business operations, organisations need confidence not only in where systems and data are located, but in who operates the infrastructure, who has access to it and the governance arrangements surrounding it.

The service provider outlook therefore reinforces the central finding of the enterprise research. AI infrastructure is becoming more distributed, but its success will depend on the networks connecting those distributed environments into one resilient and scalable system.



What we need to become regionally AI ready

The findings suggest that regional AI readiness won't be delivered by one infrastructure characteristic alone.

For enterprises' most demanding AI workloads, high resilience and route diversity rank as the leading connectivity priority, selected by 33% of respondents. Guaranteed bandwidth follows at 31%, with cost efficiency at 29%, low and predictable latency at 28%, and direct connectivity to cloud or data centre environments at 27%.

Enterprises also expect to use different connectivity models for different workloads. Forty-one per cent prefer a hybrid approach combining shared and dedicated connectivity, compared with 26% favouring fully dedicated private connectivity and 24% preferring shared, best-efforts services.

The research points to five connected foundations



Capacity

Infrastructure that can scale as AI workloads grow.



Performance

Predictable latency and sufficient proximity to users, applications and data.



Resilience

Route diversity, guaranteed bandwidth and dependable operations.



Interconnection

Seamless links between enterprise sites, cloud platforms, data centres and regional infrastructure.



Control

Clear sovereignty, governance and operational assurance across the underlying network.

Regional AI readiness is the ability to bring these foundations together in the right combination for each workload, not to impose one architecture across every organisation or region.

Building the connected foundations for regional AI growth

The UK's AI opportunity is becoming more regional. However, the ability to scale AI will depend on whether regions can support demanding workloads locally, reliably, and at high capacity. With AI workloads becoming more distributed and latency-sensitive, the UK needs a highly resilient national fibre footprint.

At Neos Networks, we see it as our core mission to connect the places where AI capacity is concentrated with the regions where applications, businesses and users are located. Our national network, regional reach and UK-based operations support organisations that need high capacity connectivity alongside resilience and operational control.

Project Reach can help that ambition. Through our partnership with Network Rail, Neos is installing new 432-count fibre along railway corridors linking London with Cardiff, Manchester and Newcastle. These routes will add high-performance backbone infrastructure between major cities, data centre hubs and AI Growth Zones, connected via direct routes which are geographically diverse to alternatives.

Project Reach demonstrates how the UK can prepare for regional AI growth: by adding scalable capacity along nationally important corridors and connecting major centres of compute with the locations where AI will be deployed.





Conclusion: turning regional readiness into national opportunity

“ Artificial Intelligence will be the most important technology of our lifetimes. As a nation, the UK boasts world-class research, deep technical talent and a growing base of compute and data centre infrastructure. And with AI now firmly on the UK government's agenda, there's no shortage of ambition for the technology either.

But turning that ambition into economic growth that spreads right across the country will depend on the infrastructure connecting it all.

That was already becoming clear when our previous research found that 82% of UK data centre operators had delayed site builds or expansion because of fibre availability. This latest research took the question a step further: if AI demand continues to grow rapidly, could businesses deploy and scale it in the parts of the UK where they actually need it?

The findings give us reason for confidence, but also a clear warning.

Large enterprises across our key sectors – from financial services to manufacturing – and the communications service providers connecting the nation, believe we have credible foundations on which to build. They see AI opportunities emerging across multiple UK regions, not just London

and the South East, and many are broadly confident in the infrastructure available today.

At the same time, however, network conditions are already influencing where AI workloads are located, how they are deployed and, in some cases, whether projects go ahead at their intended scale.

That becomes even more important as AI moves towards inferencing, real-time applications and more distributed environments. The UK cannot treat regional AI readiness as a question of compute or data centre capacity alone. To avoid a new AI infrastructure divide, all the following need to be considered together from the outset: connectivity, power, data centre capacity, cloud access and local infrastructure.

If we get those foundations right, regional AI growth can become a national opportunity rather than another source of geographic disparity. That means building ahead of demand and making sure the networks are there wherever the next wave of AI activity develops.

”

Lee Myall, Chief Executive Officer, Neos Networks




Methodology

Neos Networks commissioned Censuswide to survey 250 senior IT infrastructure and digital transformation decision-makers at UK organisations with more than 500 employees. Respondents were evenly split across financial services, manufacturing, life sciences, media and energy. The research was conducted between 8 and 14 May 2026.

Censuswide also surveyed 100 senior decision-makers at UK service providers, carriers and telecoms operators involved in wholesale connectivity, network strategy, infrastructure planning or enterprise service delivery. This research was conducted between 8 and 15 May 2026.

Some questions allowed respondents to select more than one answer, meaning figures may total more than 100%.

Censuswide is a member of the Market Research Society and the British Polling Council, and a signatory of the Global Data Quality Pledge. The research adheres to the MRS Code of Conduct and ESOMAR principles.



In collaboration with





About Neos Networks

Neos Networks has the UK's largest business-dedicated network. With over 600 points of presence and 90 data centres nationwide, Neos provides high capacity, critical connectivity for businesses, from telecoms and energy to banking and emergency services.

Its nationwide infrastructure underpins critical services and the UK's digital economy, supporting the organisations and services Britain depends on every day.



enquiries@neosnetworks.com



neosnetworks.com