TATA COMMUNICATIONS

TATA

Built for AI readiness

Evolving your network for the next networking era

#InnovatewithDigitalFabric

Table of contents

Executive summary

The AI readiness illusion

AI is one of the top priorities for global CEOs, who are increasingly mandating their IT teams to prepare the business for AI's transformational impact *(BCG, 2026)*[1]. But the reality is that even after tackling their data silos and legacy tech, enterprises still struggle to support AI adoption: only 5% have achieved AI-readiness and the vast majority remain stuck between pilots and partial scale *(McKinsey's State of AI, 2025)*[2].

One important – and often overlooked – factor why enterprises fail to scale AI is not model accuracy, insufficient data or compute constraints. It's networks that cannot guarantee predictable performance across distributed systems, creating hidden operational risks as AI environments become more distributed and dynamic.

Enterprise networks must now enter a third era – following the physical network period (1990s-2000s) and the software-defined phase (2005-2025) – to leverage the new possibilities of AI models and agents.

From plumbing to platform

AI creates new demands for bandwidth, speed, east-west traffic and global connectivity – necessitating a new kind of cognitive network. Rather than passive and content-based "plumbing", the network must become a predictive, self-healing and context-based platform. Simply increasing capacity is not sufficient. It requires a deeper re-imagining of the connectivity fabric to ensure AI can scale without introducing greater complexity, inefficiency and security exposure.

Enterprises need a different architectural approach built on three core principles of intelligence, scalability and security. Their systems must understand workload characteristics and adjust performance dynamically. They need automation, unified visibility and multi-cloud abstraction to eliminate network fragmentation. And security and sovereignty controls must be embedded by design.

Making it happen

Fortunately, this network evolution can be managed pragmatically and doesn't involve a rip-and-replace approach. Enterprises should begin with assessment, followed by targeted pilots that prove performance, and then systematic scaling and integration into long-term architecture. The goal is to reduce complexity while enhancing coherence and control, helping organisations avoid the escalating operational and financial costs associated with fragmented infrastructure.

In the AI era, IT leaders should view the network as a strategic asset, where it becomes the control plane helping AI to scale efficiently. The stakes are already evident: 88% of AI initiatives fail to reach production[3], often due to gaps in infrastructure readiness *(IDC, 2025)*[4]. Businesses that modernise their network foundations are better positioned to accelerate deployment and realise value faster, while those that delay risk slower time-to-market *(Bloomberg, 2026)*[5], rising cloud costs *(Flexera, 2025)*[6] and increased exposure to security threats as AI environments expand.

The enterprise mandate is clear: treat the network as a strategic enabler of AI, before the cost of inaction begins to escalate.

1 BCG, 2026 2 McKinsey's State of AI, 2025 3, 4 IDC, 2025 5 Bloomberg's Building Durable AI Advantage Report, 2026 6 Flexera, 2025

Key considerations for defining your AI-ready network

As AI scales up, IT leaders are being tasked with ensuring their overall infrastructure is ready – with the network playing a critical role as a strategic enabler. While every organisation is at a different stage in this journey, here are six factors that enterprises must consider carefully as they define their future roadmap.

Tata Communications can help IT leaders assess how these considerations translate into network requirements, and how to design and optimise their network to support AI workloads effectively.

AI strategy	The ambition and maturity of your organisation's overall AI roadmap will influence network requirements, from data movement to performance and scale.
Cloud strategy	A single-cloud, multi-cloud or hybrid stance shapes how connectivity is architected across on-prem, edge and hyperscale environments, and how tightly performance, cost and control are managed.
Regional focus	Operating geographies influence how networks must handle data residency, latency and resilience requirements.
Data protection and regulations	Regulatory obligations dictate how data is segmented, encrypted, audited and routed, directly shaping network design and governance.
Industry vertical	Sectors such as banking, healthcare or logistics will have specific workloads, compliance standards and risk tolerances, all of which determine how the network must be designed to support AI use cases.
Security framework	Security models increasingly need to be embedded within the network itself, shaping how identity, access and traffic are managed across distributed environments.

The cost of inaction

If organisations view their network as a cost centre and hold back investment for fear of an uncertain return, they risk incurring serious costs – holding them back in the short term and hurt them strategically in the long term. Likely consequences of inaction include:

Failed AI initiatives

88% of AI projects are scrapped before reaching production scale, with a key reason being "the low level of organisational readiness in terms of data, processes and IT infrastructure" *(IDC, 2025)*[1].

Slower time-to-market

Enterprises operating with fragmented infrastructure or legacy systems are hindered from moving as swiftly as those with more modern architectures, who differentiate by bringing AI's impact to bear in their operations *(McKinsey, 2025)*[2].

Soaring cloud costs

Data egress fees are material for many organisations, and managing spend is the top cloud- related challenge for 84% of IT decision-makers *(Flexera, 2025)*[3]. As AI usage grows, these costs will continue to skyrocket unless enterprises make deeper changes to their network architecture.

IP theft

Distributed AI estates expand enterprises' attack surface, and increase the risk of intellectual property and data theft. According to a recent data security report, "ungoverned AI systems are more likely to be breached and more costly when they are".

Competitive disruption

IDC[4] reports that AI infrastructure spending more than doubled in 2025 to **$318 billion**, driven by investments in accelerated compute and high-performance networking required to operationalise AI at scale.

1 IDC, 2025 2 McKinsey, 2025 3 Flexera, 2025 4 IDC

Chapter 1

Why networks must evolve

The triggers of the next networking era

Ambition is everywhere, but readiness is rare

AI has been on business leaders' radars ever since ChatGPT burst onto the scene at the end of 2022. The technology promises to revolutionise what companies do and how they work. Today, two-thirds of CEOs (65%) describe accelerating AI as one of their top three priorities *(BCG, 2026)*[1].

CEOs are tasking their IT teams to ensure they are ready for this brave new AI world, putting them under major pressure to deliver. But organisations are still struggling to make progress. Shockingly, only 5% of enterprises are AI-ready at scale *(McKinsey's State of AI, 2025)*[2].

On the surface, application-level innovation is accelerating and demos are common. IT leaders are busy trying to capture some of the huge economic value forecast from generative and agentic AI. Yet, 62% of enterprises, according to McKinsey, are still trapped in the AI foothills, either running early trials or testing initial use cases.

So while there's clearly plenty of appetite to harness AI's incredible productivity and efficiency gains, companies are still falling short. Ambition is all around us – but it's not sufficient on its own.

The challenge lies in the foundations. The key technology layer that enables successful AI adoption – the network – is simply not fit for purpose in most enterprises today. This structural barrier is what too often holds back AI readiness.

AI's failure to launch

Only **5%** of enterprises are AI-ready at scale, while **62%** of enterprises are still in the AI pilot or experiment phase

Source - McKinsey[3]

Enterprise AI maturity fell to **35%** in 2025, from **44%** in 2024, as businesses struggle to turn AI hype into real value

Source - ServiceNow[4]

Over **40%** of agentic AI projects will be cancelled by the end of 2027, due to unclear business value, escalating costs, or inadequate risk controls

Source - Gartner[5]

72% of CEOs in 2026 say they are now the main decision maker in their organisation, twice as many as in 2025

Source - BCG[6]

$2.6tn-4.4tn of economic value will be derived from generative AI annually *(McKinsey, 2023)*[7], with **$450bn** of value from agentic AI forecast by 2028 *(Capgemini, 2025)*[8].

1 BCG, 2026
2 McKinsey's State of AI, 2025
3 McKinsey
4 ServiceNow
5 Gartner
6 BCG
7 McKinsey, 2023
8 Capgemini, 2025

A brief history of enterprise networks

To understand why AI is revolutionising the enterprise network, we need to look back at how digitisation has developed over the past few decades. Broadly speaking, enterprise networks have evolved across three eras:

1990s-2000s ▶ The physical network era

This era saw applications hosted on server stacks in an organisation's data centres (either leased or wholly owned). Most users would come to the office to work, as hybrid or remote work was rare. Enterprises would typically have three 'blocks' in their network: a Local Area Network (LAN), a Wide Area Network (WAN) and a data centre network.

The LAN connected the office through a physical network of Ethernet cables. The WAN was primarily a private network based on leased lines, Ethernet or Multi-Protocol Label Switching (MPLS), and connected the office to data centres, where the security layer sat. Lastly, there was the network within the data centre, connecting servers and storage, and the network connecting to other data centres and disaster recovery.

2005-2025 ▶ The software-defined network era

This era was when cloud computing became mainstream. This period saw an explosion of data and new applications, as mobile devices boomed and working from home became commonplace. Enterprises moved applications from the data centre to public clouds, making the internet the network of choice.

When this phase reached maturity, the network had primarily five 'blocks'. The data centre network remained largely the same as before, while the office LAN was still important but shifted from wired to WiFi.

In addition, a new "Work From Anywhere" network emerged, connecting remote employees to the corporate network over an Internet Protocol Virtual Private Network (IPVPN) tunnel and latterly via Security Service Edge (SSE). Before Covid, this only handled a small percentage of enterprise users, but had to manage almost 100% of employees working remotely during the pandemic, before settling down to the roughly 30% level we see today.

The WAN pivoted from a private network (relying on private lines or MPLS) to a public network using the internet, as most key applications were increasingly hosted on public cloud. For enterprises who had shifted to the internet, Software-Defined WAN (SDWAN) became mainstream for traffic steering and centralised policy management. This is now evolving into Secure Access Service Edge (SASE) for stronger security and AI-driven automation for proactive journey routing and even self-healing.

The cloud network started with one link to EC2 and S3 in AWS, but over the years has sprawled to encompass multiple links to multiple clouds. Indeed, enterprises must now think about the network within as well as to the cloud. Poor network design is leading many enterprises to pay a growing "complexity tax" to their cloud provider, in terms of operational or network cost.

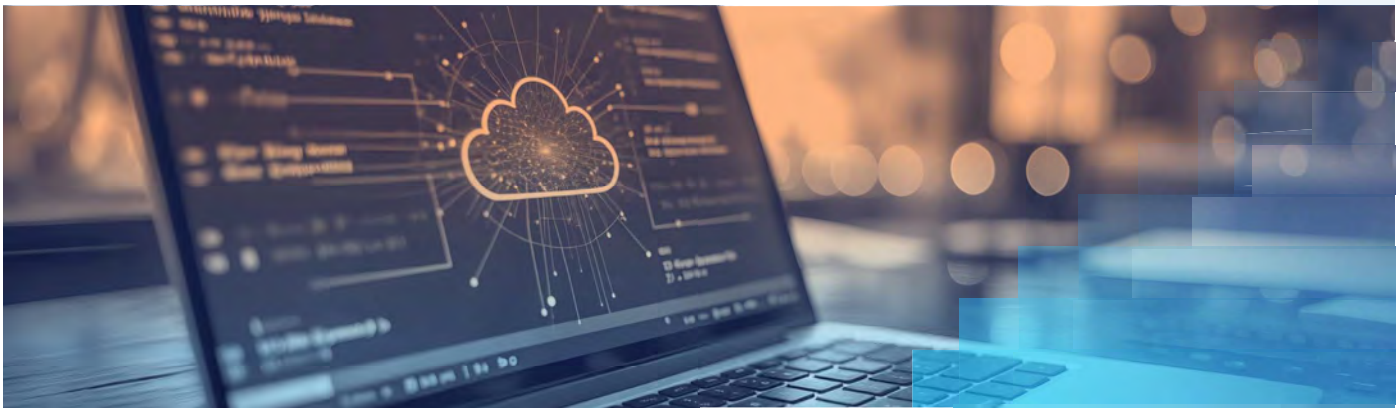

2025- The contextual network era

Now we're entering this era where AI – and specifically AI agents – will drive significant change. AI agents can reason, plan and complete multi-step tasks across different systems autonomously. As they boom across today's enterprises, they create a new set of network users that will complement human employees (and potentially even out-number them).

This "agentic network" will cause novel traffic patterns that existing architectures cannot cope with. Ultimately this will lead to the fundamental transformation of the network as we know it.

94% of enterprise IT leaders say the network is a bottleneck to AI adoption *(IDC and Expereo, 2025)*[1].

1 IDC and Expereo, 2025

What's different in AI-dominated networks

AI is a far more complex technology than the cloud and internet that dominated the previous era. The rise of AI creates two specific data flows that require enterprise networks to adapt in this third phase of networking history.

The first data flow challenge comes from training AI models within an enterprise's data centre. This is a massively demanding activity that can potentially overwhelm existing network capacity, which were built for human-generated "elephant and mouse" flows. AI workloads, in contrast, are persistent and continuous, with constant data ingestion or training loops, while AI datasets are massive and highly resource-intensive.

Training an AI model means enterprises now use hundreds of thousands or even a million GPUs, compared to a few thousands GPUs in years past. This requires data centres to be connected with high bandwidth, minimal latency and zero packet loss (because one packet loss may lead to the entire training cycle starting again, wasting expensive GPU investment). Networks that were previously used to handling 100Gb of data might now need capacity for 400Gb or even 800Gb to ensure lossless communication and no jitter.

Crucially, performance is no longer defined by average latency alone – tail latency becomes a key factor. In distributed AI workloads, where GPUs must synchronise continuously, even small delays in a fraction of traffic can bottleneck the entire system.

The second data flow challenge is from conducting inference, where human users talk to the AI and where agents talk to other agents. This creates multiple issues for existing network architectures. In many enterprise environments, this traffic still traverses the public internet when moving between clouds or data centres, introducing unpredictable routing behaviour that is misaligned with AI performance requirements.

Most obviously, inference changes the direction of network traffic, from mainly north-south flows between human users and applications, to much more east-west traffic between machines. As agent-driven traffic flows increase, a network might need to be ready to handle 50x or 100x of its capacity for human users.

Furthermore, the speed at which agents can operate is much faster than that of humans, posing a further set of network challenges. Devices like phones and laptops will be increasingly used both by humans and agents, so office LANs and Work From Anywhere networks will need to be able to operate at machine speed. Failure to do so risks damaging the performance of AI systems and missing out on AI's productivity gains.

Lastly, the importance of inference also reinforces the value of fixed data centres, as in the first phase of network history mentioned above. In the cloud era, hyperscaler regions meant that 90% of enterprise network traffic was actually regional in nature. But in the AI era, enterprises will operate with a mix of data centres for certain applications, plus hyperscalers, neoclouds, and specialised AI platforms for other applications. The network has to allow you to reach the application – wherever it is hosted – in the fastest way possible, returning us to a network that truly supports global traffic between regions.

Millisecond-level delays unnoticeable for humans are critical bottlenecks (or worse) for AI agents and other real-time systems. Re-sending a packet 200ms later, for instance, is useless for an algorithmic trading or autonomous driving tool, because the moment for decision will have already passed.

Chapter 2

What the AI-ready network looks like

The key changes that IT leaders must implement

Three principles for the future

To get AI-ready, networks must evolve from being passive and content-based to predictive and context-based, in which it can detect, for example, a heavily synchronous AI training workload and automatically provision a dedicated, low-latency optical path, bypassing standard internet routing. This is enabled through programmable network layers and dynamic provisioning across private backbone and multi-cloud interconnects, rather than relying on static, internet-based routing paths.

This future network will be built on a triad of principles, making it natively:

Intelligent and autonomous

Unlike today's enterprise networks, AI-ready networks will be cognitive, with deep observability built in as standard. They need to be able to understand whether traffic is from a human or agentic user, deal with it differently, and potentially send it to different locations (such as a normal internet network or edge network). In practice, this requires unified visibility across hybrid environments and the ability to classify and steer traffic dynamically across cloud, edge and private backbone infrastructure.

They will also be able to use AI-powered orchestration and deterministic routing to move beyond reactive firefighting of network issues, and instead apply predictive and preventative measures that, for example, stop congestion before it happens. This includes leveraging intent-based networking and automated traffic engineering to re-route workloads in real-time across optimal network paths.

Elastic and scalable

Given enterprises' hunger for data, IT leaders will need to architect their future networks to handle exponential scalability. This involves leveraging a converging set of technologies, like 5G and WiFi for wireless density, software-defined networking for dynamic routing, and a core Ethernet network to handle massive IoT- and AI-driven data flows of 400Gb and beyond.

Crucially, this scalability must extend across multi-cloud environments, requiring high-capacity interconnects and transit layers that can move large volumes of east-west traffic efficiently between clouds, data centres and edge locations.

However, scalability is not simply about increasing link capacity. It requires reviewing congestion domain boundaries and establishing controlled interconnect paths between clouds. IT leaders must evaluate whether internet-based optimisation is sufficient, or whether policy-controlled backbone interconnect is needed for critical AI workloads.

Unlike traditional internet routing, which relies on best-effort protocols such as Border Gateway Protocol (BGP) to determine the shortest available path, AI workloads increasingly require deterministic routing – where traffic follows predefined, performance-assured paths that guarantee bandwidth, latency and reliability.

In many cases, this means leveraging private, deterministic network backbones and direct cloud interconnects to ensure consistent performance and avoid the unpredictability of public internet routing.

Secure and controlled

The old perimeter-based security model is defunct in an era of AI-powered threats, when data moves continuously across domains. Security architecture must therefore be embedded into routing logic, not bolted on. Sovereignty-aware routing and identity-based segmentation become architectural design requirements. This is typically implemented through integrated network and security layers, where policies are enforced dynamically as traffic moves across clouds, regions and users.

Resilient networks must be built on zero-trust principles with a philosophy of “never trust, always verify”, and only grant users least-privilege access. The best defence against AI is AI itself, and modern secure networks must deploy this technology as a key part of threat detection and prevention. This includes embedding security controls such as SASE and Zero Trust Network Access (ZTNA) directly into the network fabric to ensure consistent protection regardless of where users or workloads are located.

AI infrastructure networking architecture

Enterprise multi-cloud AI workload connectivity

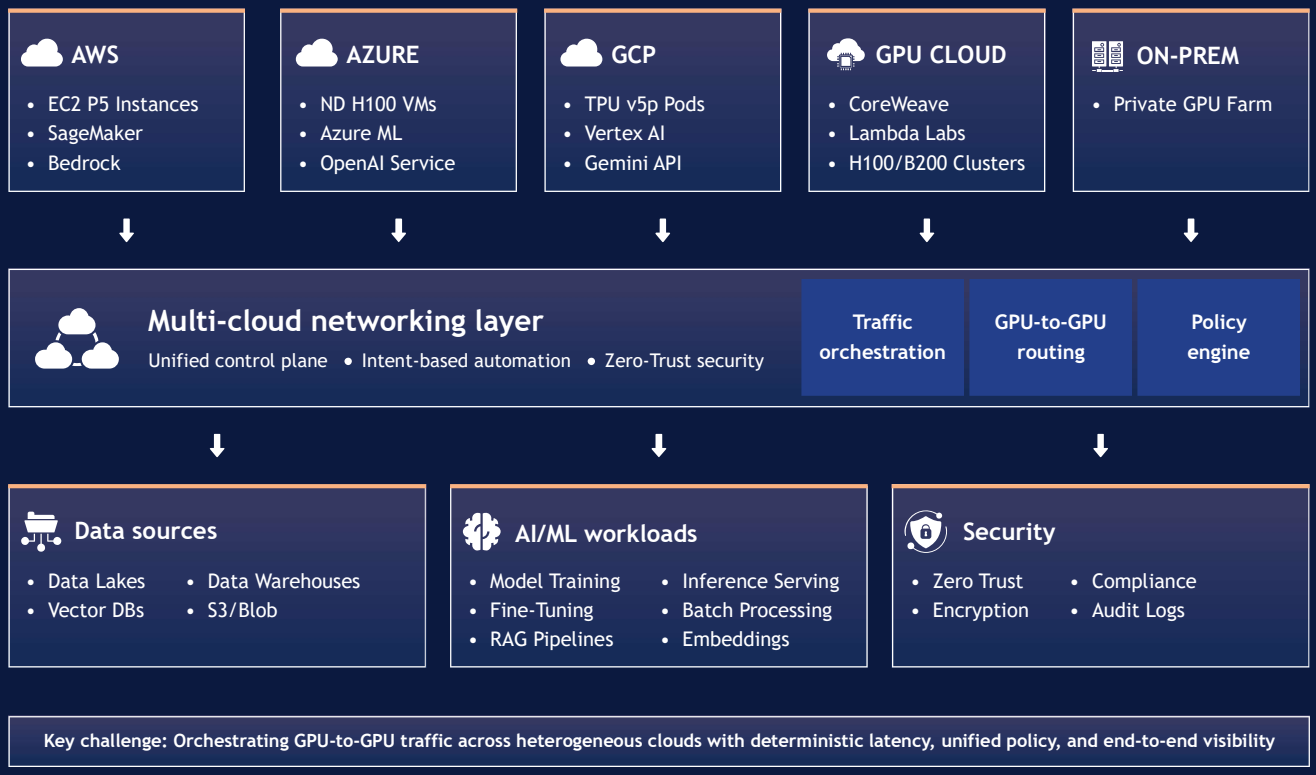

Why sovereignty matters

Just a few years ago, data sovereignty was primarily a storage problem. But now the constant transfer of data across regions and clouds, driven by AI, makes it a network problem. Geopolitical regulation of data flows is growing across the EU, US, China and India, leaving enterprises increasingly unable to follow this complex patchwork of rules without network-level controls.

An AI-ready network helps ensure compliance by providing enterprises with visibility into data paths as well as metadata. It can dynamically identify, say, that a specific packet contains German citizen data and must be physically routed through Frankfurt, not London, even if the London path is faster.

On top of data sovereignty concerns, the need for cost optimisation is leading many enterprises to bring selected workloads back from public clouds into their own data centres or private environments. For instance, some are training workloads in neoclouds or private clouds to save on egress costs and reduce the amount of inference conducted on hyperscalers. That's driving more hybrid strategies and greater fragmentation of data, creating a requirement for more resilient network backbones to move it around at speed.

Hybrid architectures

As AI workloads scale, enterprises may also need to consider whether existing network architectures are sufficient to support their unique performance characteristics. That's especially important given that network performance will no longer be defined by average latency alone. Tail latency also becomes critical, as it can constrain the efficiency of entire AI workloads.

This raises an important architectural consideration: should AI workloads co-utilise existing enterprise networks, or operate within more dedicated, segmented environments? In some cases, organisations may choose to isolate certain AI data flows within separate logical or physical network domains. While this will not be necessary for all use cases, it highlights the need for enterprises to assess whether their current network architecture can support the scale, performance and consistency requirements of AI, or whether a more hybrid approach may be required over time. This may involve combining shared enterprise networks with dedicated high-performance paths for critical AI workloads.

Key building blocks of an AI-ready network

Infrastructure built for the demands of AI display core characteristics that set them up for success as this technology scales:

AI-native routing
Optimising east-west traffic deterministically, to maximise GPU usage, managing costs and reducing waste

Policy automation
Establishing self-healing networks and ensuring infrastructure-as-code

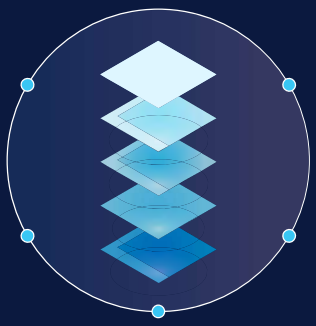

Unified visibility
Simplifying management through a single pane across all clouds and edge

Embedded security
Building in zero trust principles, micro-segmentation and encrypted transit as standard

Cost intelligence
Smartly managing egress arbitrage and data placement optimisation

A network ready for AI workloads

Traditional workloads

Latency	Data	Traffic
100-500s acceptable	KB-MB per request	North-South (client-server)

Scaling	Location
Gradual, predictable	Single region typical

Optimised for: Web apps, Databases, Microservices

AI/ML workloads

Latency	Data	Traffic
<10ms critical	GB-TB per training job	East-west (GPU-to-GPU)

Scaling	Location
Burst, dynamic	Multi-cloud, multi-region

Optimised for: LLM training, Inference, Agentic AI

The multi-tier architecture emerging

Internet infrastructure is evolving from regionally concentrated hyperscaler hubs to a distributed model:

Neoclouds	Hyperscalers	Edge
Training facilities	Orchestration & storage	Inference serving

Chapter 3

How enterprises should respond

Getting AI ready at pace

AI changes what networks must do and multi-cloud changes where networks must live

At AI speed and scale, digital environments simply become too dynamic and distributed for manual network management. Consequently, IT leaders need to unlearn the network approach of the public cloud era, leaving behind design principles that worked a decade ago but no longer apply.

Instead, the AI-ready network must function as a coherent fabric across clouds – not a collection of isolated domains – with an abstraction layer across providers to give IT leaders a single point of visibility and control. In an AI-first enterprise, the network is much more than just connective tissue. It's the control plane for performance, compliance and resilience.

Operational success is now inseparable from network readiness. The AI winners will be those who shift to deterministic infrastructure and gain structural competitive advantage.

Key advantages for AI-ready enterprises

Accelerated AI deployment

Driving successful adoption of AI tools by ensuring the right systems and structures are in place, helping them to making the leap from AI pilots to implementation at scale.

Faster product launches

Enabling the roll-out of new capabilities thanks to fewer manual approvals, fewer integration delays and less firefighting.

Lower cloud egress fees

Correctly architecting enterprise networks for AI workload needs, for instance by eliminating unnecessary cross-region movement, can lead to a 30-50% reduction in egress charges for those with heavy data demands.

Resilient and future-proofed operations

Embedding automation, segmentation and deterministic routing to become better equipped to absorb growth and withstand shocks.

Building confidence fast

Transitioning to AI-ready infrastructure should not be an overwhelming task. Through a pragmatic three-step approach, enterprises can make incremental improvements that avoid a wholesale rip-and-replace.

Assess and align

You can't fix what you can't see. So the starting point for IT leaders must be to audit their current network against industry best practice and their AI roadmap. That helps identify infrastructure and policy gaps, and find the biggest points of friction.

In many organisations, it's beneficial to look at how the network handles latency, bandwidth and east-west traffic, particularly across clouds. For example, can your infrastructure handle a throughput of 400Gb or more? Can virtualised devices scale or are they bottlenecks? And is cross-cloud latency benchmarked under realistic workload conditions?

Pilot and prove

It's then easier to run a targeted proof of concept on a high-value AI project that will show ROI rapidly. Start small, with a tightly defined scope, and measure impact in terms of latency, cost and performance.

This stage is about stress-testing the network under real production conditions. By choosing a technically contained yet commercially impactful use case, IT leaders can demonstrate that improving determinism, visibility and traffic steering translates into efficiency and speed gains. Early evidence builds internal confidence and unlocks fresh funding.

Scale and integrate

Now enterprises can embed this modern backbone into the network fabric, rather than bolting it on. The lessons from the pilot can be built into long-range strategic planning.

Scaling is about institutionalising what worked – formalising multi-cloud abstraction, codifying sovereignty controls, expanding automation across regions, and ensuring observability spans edge, core and cloud. Integration also requires aligning network architecture with procurement, security and platform teams so that AI workloads are designed with data locality and performance in mind from day one.

The power of partnership

The Tata Communications approach

The three key challenges to AI-ready networks are explosive east-west traffic, fragmented multi-cloud estates, and the need for sovereignty-aware routing – all of which are difficult to solve through slow, incremental upgrades to existing infrastructure. As set out in this whitepaper, enterprises need a network that's architecturally coherent from the outset: one that unifies connectivity, cloud access and security under a single programmable fabric, rather than assembling them from disconnected point products.

That's precisely the design principle behind Tata Communications' Unified Network Fabric – a globally programmable platform built to address the demands of AI-era infrastructure. Its foundation is one of the world's most extensive private networks.

Built on the world's most extensive private network

Global fibre reach

500,000 km of wholly owned subsea & terrestrial fibre

Internet backbone

Direct access to **35%** of global internet routes

Worldwide presence

Operating across **190+** countries

Cloud connectivity

Connected to **80%** of the world's cloud providers

Critically, Tata Communications brings particular depth in high-growth markets for global enterprises – such as India, China, Brazil and the Middle East – where regulation and compliance challenges are often most complex.

At the core of this architecture is the combination of an intelligent overlay (IZO™+ MCN) with a high-performance private underlay (Interconnect). Together, these layers create a predictable, private transit path for AI workloads, designed specifically to handle the shift from traditional north-south traffic to high-volume, agent-to-agent east-west flows.

This isn't shared public internet infrastructure – it's a dedicated, deterministic backbone capable of supporting the low-latency, high-bandwidth, lossless data flows that AI training and inference workloads require. In practice, when an AI workload communicates between environments – for example, between cloud platforms or between cloud and data centre–traffic is directed from the source environment to a local IZO™+ MCN gateway, rather than routing via the public internet.

From there, it is carried onto private cloud on-ramps and transported across Tata Communications' Interconnect backbone, before exiting via a corresponding private connection into the destination environment, where it is delivered to the target workload.

By keeping traffic entirely off the public internet, this model avoids the unpredictability of best-effort routing, where protocols such as BGP prioritise shortest path over performance. Instead, traffic is pinned to dedicated, SLA-backed paths across a private optical and MPLS infrastructure, ensuring consistent latency, minimal packet loss and reliable performance for AI workloads.

In addition to that backbone, Network Fabric delivers on the three capabilities acknowledged to be non-negotiable for AI readiness:

Intelligence and autonomy

- Integrates AI-driven orchestration, predictive routing and self-healing automation — all managed through a single pane of glass
- Provides end-to-end visibility across LAN, WAN, cloud and edge, enabling the network to dynamically adapt to changing workloads
- ThreadSpan™ aggregates telemetry across the full stack — enabling teams to see the complete path of an AI workload in a single view

Elastic scalability

- Spans global VPN, private line, IP transit, SD-WAN and multi-cloud connectivity to AWS, Azure, GCP — delivered as-a-service
- High throughput AI traffic identified through orchestration-layer signals or real-time traffic analysis trigger dynamic provisioning of bandwidth and optimised paths across the network without human intervention
- Bandwidth provisioned in seconds to minutes using bandwidth-on-demand, vs. days or weeks with traditional methods to respond instantly to spikes in demand

Security and sovereignty

- Advanced network security with SASE and ZTNA embedded into the fabric — not bolted on
- 8,000+ threat intelligence operation centers enabling continuous threat detection
- Enables enterprises to enforce governance, compliance and data residency requirements directly at the network layer

The result is a network that actively manages how data moves across environments, enabling enterprises to optimise performance, control costs by reducing unnecessary data movement, and maintain compliance in increasingly complex regulatory landscapes. AI readiness is also a network readiness problem. Tata Communications' Network Fabric is built to solve it. Seamless, secure, trusted.

The time to act

We're moving from the software-defined era of enterprise networks, dominated by public clouds and internet connectivity, to a contextual era where AI models and agents are reshaping data flows and architectural requirements.

As a result, networks are becoming a decisive factor in whether organisations can achieve positive ROI on their AI investment. Undoubtedly, access to the latest applications, sufficiently upskilling the workforce, and setting direction from the top of the organisation are important too. But to bear fruit, they all rely on an AI-ready network being in place. Fundamentally, as AI agents spread, the network is emerging as a critical – yet frequently overlooked – enabler of AI success, alongside factors such as model accuracy, data and compute.

The benefit of investing in AI-ready infrastructure is not just theoretical. It's tangible in direct cost savings, accelerated revenue, and IP protection. Connected intelligence is increasingly essential to sustained competitive edge. Those who move to modernise their networks and embed automation and sovereignty controls will build a durable advantage. Those who delay risk locking themselves into fragile, expensive architectures.

The AI era is developing fast, and a new phase of network architecture is taking shape. We're at a pivotal moment, where tomorrow's winners and losers will soon be set. The time for enterprises to act is now.

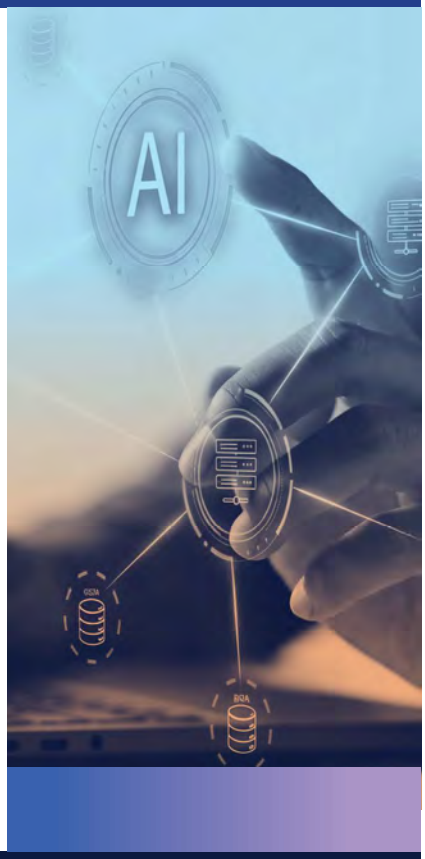

"Wherever you are in your network's AI readiness,
Tata Communications can help."

Contact us today to request a demo and find out more.

And check your enterprise's score on our Multi-Cloud Complexity Calculator to get a clearer view of where fragmentation, inefficiencies and risk may be holding back your AI performance.

For more information, visit us at **www.tatacommunications.com**

Contact us