The AI reality check

The network essentials for successful deployment and monetization of enterprise AI.

verizon business

Explore the promises and challenges of AI in business

1. **Embracing AI calls for a reality check of your network**
2. **Going with the flow:** How will we deal with increased AI data flows?
3. **Industry in the moment:** How can we enable real-time AI use cases?
4. **Remaining in charge:** How do we stay in control in an AI world?
5. **Securing the system:** How do we stay secure in an AI world?
6. **Network preparedness checklist:** Is your infrastructure ready for AI?
7. **Your platform for the next era of enterprise connectivity:** Verizon Business unlocks network success with unparalleled insight and experience

verizon business

Embracing AI calls for a reality check of your network

The artificial intelligence (AI) revolution has accelerated many aspects of work, but it has also added pressure. How much time do you have to read this article, for example? Probably not a lot. Not while you're expected to keep pace with your rapidly evolving industry.

Enterprises are currently making the biggest investment in technology since the 2000s. The world's largest tech firms have spent around $400 billion in 2025 alone on their AI bet,[1] enabling a host of new applications—everything from robotic taxis to automated fraud detection.

For example, one major European industrial brand is using an engineering agent to reportedly drive up to 50% efficiency gains.[2] An insurer, at the same time, is employing a chatbot to support over a million customer interactions annually.[3]

There's plenty of hype and optimism surrounding AI, but there are also significant risks and technical issues to resolve. From our work helping firms adapt to this era, we've seen time and again that underdeveloped network and connectivity capabilities are posing significant barriers to progress.

The infrastructure that enables businesses to benefit from advanced AI, large language models (LLMs) and machine learning (ML) often isn't where it needs to be. While many companies have an AI strategy, few have the networks in place to execute it at scale.

That's exactly what the paper you're reading will help you address.

We know that networking essentials can make or break AI initiatives. Whether you're looking to embrace agentic AI to support internal teams, power computer vision or deliver AI-enhanced solutions to your customers, you will certainly be facing an avalanche of data traffic. Capable networks underpin all these innovations. After all, they are the invisible engine powering AI.

This paper details the implications of AI on different aspects of your networking infrastructure, providing a checklist to help you identify and overcome challenges before they have the chance to slow you down. You'll also hear directly from the experts on how businesses across the globe are seizing the opportunities AI offers.

[1]The Wall Street Journal, Big Tech's $400 Billion AI Spending Spree Just Got Wall Street's Blessing, 2025
[2]Siemens, Siemens brings AI to the physical world with Eigen Engineering Agent, 2026
[3]Matthieu Caillat for AXA Group Operations, AI at AXA: Leading the future of insurance with responsibility and innovation, 2026

Going with the flow

How will we deal with increased AI data flows?

Data is the lifeblood of AI. Without it, LLMs would have no generative ability, agents would be unable to find and action the next logical step in a process and mass personalization would still be a painfully slow, if not impossible, process. The flow of data today is enormous.

The question you may be asking is: "How will we deal with the increase in data flows that AI-powered operations demand?" To answer that, we can look at how some retailers and manufacturers have successfully been putting data to work.

AI-enhanced operations are data intensive

AI is speeding up the pace of development in all kinds of business processes—from production to customer service and beyond.

Manufacturers, for example, are growing more comfortable with and capable of utilizing AI-enhanced digital twins—digital models of real-world, physical products or systems that can be placed into simulations, tweaked and tested without impacting the real thing. These tools are helping them to achieve greater business intelligence and improve predictive maintenance, while making disclosure, compliance and risk management less costly and disruptive.

> "For manufacturers, using machine learning to understand their data is nothing new. But now people have realized that, with LLMs, they can chat with their data and get an answer."
>
> **Christoffer Sundgren**
> Senior Principal of Business Strategy, Manufacturing at Verizon Business

With the help of an AI-enhanced digital twin, you might be able to ask a model to identify otherwise invisible issues. While getting the twin of your manufacturing equipment to diagnose its own problems would save considerable time, this creates a huge data demand.

One European automotive manufacturer, for example, uses high definition video to capture images of products on the assembly line, with an AI system monitoring the entire process.[4] While this is apparently improving quality, the additional data load is surely significant, seeing as an hour of 4K video may require as much as 45 GB of storage space.

> “There are thousands, if not hundreds of thousands, of sensors in a factory and all the data is coming in at the same time.”
>
> **Colin Wilson**
> Enterprise Architect, Verizon Business

Retailers are also facing data surges as they embrace AI. Some have been able to improve inventory management and ensure the right products are on shelves. But the balance can easily be thrown off by the ebbs and flows of the industry.

Verizon Business’ James Hughes: Retail Chief Technology Officer (CTO) has seen first-hand how seasonal surges and complex needs can impact retailers’ network infrastructure.

Both retailers and manufacturers need a new baseline capability to handle vast amounts of data, but they must also be able to flex. Training proprietary models means provisioning even more data bandwidth—potential tens of gigabits every few days. And that needs to come with consistent uptime.

While a few hours of downtime might be acceptable if people can step in, for AI systems (like Agentic operations running at machine speed) those hours could prove devastating. Identity and access management systems also need to stay consistently online to see off AI-enhanced cyber threats.

As agentic AI becomes more common in workflows, these needs will be greater still.

> “AI tools won't be the only thing on the network. There'll be all the other things that retailers want to run as well, whether it's virtual assistant or payments. All of these are going to be taking bandwidth.”
>
> **James Hughes**
> Retail CTO, Verizon Business

[4]BMW Group, This is how DIGITAL the BMW iFACTORY is, 2023

Agents process more data, in different ways

Agentic AI is set to change how we think about network infrastructure, not just how we work. These are intelligent systems driven by AI that use tools and APIs to execute end-to-end workstreams, making real-time decisions to achieve specific goals with minimal human intervention.

As Christoffer Sundgren: Senior Principal (SVP) at Verizon Business puts it, agentic workflows essentially boil down to "machines talking to each other." Enterprises are now stringing multiple agents together to build strong, effective workflows. It is important to note, however, that machines do not 'speak' the same way we do. In fact, they send information at machine speed, generating high levels of traffic and more data per second than we ever could.

This change means moving from a traditional "north-south" way of thinking about networks to one much more horizontal. As machines begin sending information across data centers and your systems before it is delivered, we must start thinking about "east-west" capabilities and capacity.

Agents require persistent access to data and often need to work as close to the network edge as possible to minimize latency. This calls for a massive increase in data flows, token usage and network bandwidth, as well as an infrastructure to support real-time access to data.

The enterprises of today should prioritize developing a network architecture that can handle this vast amount of data.

> "You've got spider webs of data in one place, users in other locations, workloads and LLMs running a third location. That's driving huge network demand."
>
> **Colin Wilson**
> Enterprise Architect, Verizon Business

Self-assessment: How scalable is your network?

1. Do we have visibility into how much additional data our current and planned AI operations will generate – including from sensors, video feeds and agentic tools?
2. Can our network infrastructure flex to handle both peak demand and the baseline load of running multiple AI tools simultaneously?
3. Have we mapped where our data, users and AI workloads are located, and assessed whether our network can efficiently move data between these points?
4. Do we have a network architecture strategy that supports the east-west data flows that agentic AI agents require?
5. Do we have a good understanding of what is connected, and what it should have access to, thereby minimizing our attack surface?
6. Have we learned the lessons of 'technical debt'? Are we confident that all of our devices and systems are properly owned and managed such that when we want to decommission them, we can do so safely and with confidence that business systems will continue to operate?

Industry in the moment

How can we enable real-time AI use cases?

Today, AI moves faster than ever. Instant solutions and real-time actions are promised, but AI can only move as fast as the networks behind it.

Take autonomous vehicles, for example. Driverless cars must operate safely in a complex, 3D world. To do so, each of the 40 or so sensors that enable automated driving must absorb up to 50 different data points every millisecond, then inform split-section decisions in real-time.[5]

Considering that these vehicles may have up to 200 or so sensors for full functionality, something as simple as waiting for a traffic light becomes an intensive computing challenge. There's simply no time for signals to go up to the cloud and return. All that thinking needs to happen in the vehicle itself, right at the edge.

Real-time AI is also utilized in the finance sector to detect fraud as it happens, responding to the rising threat of voice-cloning and AI-enhanced phishing attacks. The models these companies are using can be trained to recognize suspicious activity and risks as they occur, utilizing advanced pattern recognition and impressive scalability to catch things humans might miss.[6]

Some retailers are even embracing real-time AI to give customers a chance to try on items from the comfort of their own mobile device. But that brings its own network challenges, as Hughes explains.

> "Most of this retail activity is going to be consumed on a mobile device and a lot of that will be AI driven or generated. But these devices need to be connected in the retail environment. The infrastructure needs to be good and the applications need to be prioritized in a way that makes it resonate. And then the workloads need to be done in a way that makes it real time"
>
> **James Hughes**
> Retail CTO, Verizon Business

The AI arms race is happening now. Businesses need to think about how they will respond to the challenges posed by AI, as well as the opportunities.

[5] Automotive World, The power behind autonomous cars: time series data and AI, 2024
[6] IBM, AI fraud detection in banking, 2025

Why some AI needs to operate at the edge

Each of these use cases requires AI to operate where the action happens. Cloud-based architecture might not be the right approach for this as, sometimes, users can't wait for the round-trip to a distant data center.

A network that enables edge computing could be the bridge between simply requesting output from LLMs and unlocking full autonomous operations.

Verizon Business' Sundgren is helping to build these networks for industry clients today. He notes that manufacturers looking for real-time AI support need incredibly low latency, as well as traditional cloud uplinks for the purposes of model training and updates.

Uptime is always a key consideration—operations can't slow down just because one network component fails, nor can safety be put at risk. AI at the edge should help to alleviate some of this risk, dispersing processing so that, should one part of the chain go down, the whole system doesn't collapse.

Self-assessment: Is your network enabling real-time applications?

1. Which of our AI use cases require real-time responses? Could high latency make them unworkable or unsafe?
2. How resilient is our network infrastructure if one component should fail? Could our AI-driven operations continue safely?
3. What would our potential restore times look like should a component fail?
4. Are our applications and workloads being prioritized in a way that guarantees real-time performance, particularly on mobile or end-user devices?

Remaining in charge

How do we stay in control in an AI world?

> “If you're building a model for, say, pharmaceutical research, then you don't want to be sharing that data with anybody. You absolutely want it locked down, even if you do want to use all the clever capability that AI has.”
>
> **Colin Wilson**
> Enterprise Architect, Verizon Business

Businesses in all sectors need visibility and strong controls over their data. This obvious fact bears repeating, especially as we move to a world in which AI tools are being entrusted with sensitive information. If data access isn’t controlled, however, businesses risk not just giving up a competitive edge but also falling afoul of laws and data regulations.

Sovereignty is another increasingly discussed topic in the AI age. Microsoft CEO Satya Nadella suggests that the location of physical data centers is actually the least important factor (as we might assume if we think purely of “data sovereignty”). Instead, he suggests that the control over the models trained on proprietary knowledge should determine what he refers to as “corporate AI sovereignty”.[8]

Another definition comes from The World Economic Forum. While this organization might be focused on the ambitions of nation states, their definition can broadly apply to enterprises as well: “the ability of economies to shape, deploy, and govern AI ecosystems in accordance with their own values, whilst ensuring strategic and operational control, flexibility and ultimately, resilience.”[7]

AI sovereignty encompasses the entire AI stack, from infrastructure like Graphics Processing Unit (GPUs) and power grids to governance, like legal and ethical frameworks. If a nation (or at a smaller level, an organization) can maintain full operational control over this stack, it has achieved AI sovereignty.

Control and clarity over the ownership of data are essential to moving forward in this new age.

[7] European Data Protection Supervisor, TechDispatch #2/2023 - Explainable Artificial Intelligence, 2024
[8] The Register, Microsoft CEO: AI sovereignty isn’t where it runs, it’s who controls it, 2026

Proprietary models could be the answer to data panic

For businesses that aren't comfortable sharing information with public LLMs, the obvious answer is building their own.

Programmers and tech teams across industries can now build effective models from the ground up, within ring-fenced systems. Some businesses are choosing to use forks of public LLMs, with contracts preventing information from leaving the walled-off system or being used to train the parent model. Others are running AIs entirely within local networks, with no outward paths to prevent data leakage.

One benefit of completely owning an AI model is being able to train it on proprietary information and tailor its output to your specific needs. While the huge data sets of popular LLMs are useful for everyday tasks, certain industries may require specific knowledge and expertise that would be difficult for these AI tools to get right.

Of course, running and training these models is incredibly data intensive. Doing so on top of typical network activity is going to place strain on even the most prepared of infrastructures.

It is vital that any business looking to develop or integrate a ring-fenced AI into their workflows evaluates the suitability of its network infrastructure—and invests in upgrades as appropriate.

Self-assessment: Are you really in control of your data?

1. Is our network infrastructure capable of handling the additional load of running and training a proprietary or ring-fenced AI model on top of existing operations?
2. Do we have network controls in place, such as air-gapped environments or private connectivity, to prevent sensitive data travelling outside approved boundaries?
3. How do we keep data protected as it moves between distributed compute nodes?
4. How do we keep latency low enough to avoid 'pathway rerouting,' which could result in data moving across borders into areas we don't have sovereignty?
5. Can our infrastructure enforce data sovereignty requirements, ensuring data is processed and stored only within the correct geographic regions?
6. If part of our network were to be compromised, do we have the safeguards in place to prevent unauthorized access to our AI systems and the data they use?
7. How do we ensure we remain compliant with frameworks like the General Data Protection Regulation (GDPR) or the European Union Artificial Intelligence (EU AI) Act ?

Securing the system

How do we stay secure in an AI world?

LLMs and AI agents may be incredibly useful for automating many tasks, but they also expand the attack surface of your operations. As these tools autonomously call on data from outside of local networks and transfer it from machine to machine, the need to secure your business becomes far greater.

Consider, for example, the fact that AI agents are increasingly being given more autonomy to call APIs, query databases and even execute code. If a bad actor were to be able to pull off a successful prompt injection (in other words, manipulate the AI into proceeding with a harmful action despite guardrails), the consequences could be enormous—especially in industries such as finance and logistics.

Malicious intent isn't always involved, either. As well-meaning employees begin experimenting with interconnected agents and automation, there is always the possibility that these tools may act in an unpredictable and potentially damaging way. Strict oversight and mitigations are necessary at all stages to prevent this.

Keeping a lid on sensitive data is a big responsibility

Every business needs to ask itself a few key questions around security:

- Is our customer and payment data fully protected against breaches?
- Are we compliant with global privacy laws?
- Will we remain secure and compliant with AI operating independently across our networks?
- If not, how can we better secure and control it?

Colin Wilson: Enterprise Architect (EA) at Verizon Business suggests organizations also ask themselves: "Is my data being compromised and sent somewhere it shouldn't? Is it accessible to external threat actors?"

"How are we making sure that someone else has not intercepted and replaced the data that's been sent ... with data that is incorrect, biased or nefarious?"

While AI and agentic security is just beginning to evolve, it is almost certain that network visibility and controls will be the first line of defense that alerts your higher order systems to a possible threat.

The network sees all. It is the invisible engine behind it all, fueled by data and enabling essential processes. Monitoring the engine to identify and address unusual patterns should be at the heart of any AI security strategy.

There are three critical elements to network security in the AI age:

- Secure Access Service Edge (SASE) ensures your customers and employees can use a variety of wireless and broadband access options to securely connect from anywhere within your wide-area network.
- Zero-trust policies that assume users or devices, whether inside or outside the network perimeter, should not automatically be trusted.
- Visibility will become increasingly important as regulators demand AI be able to explain itself and make its workings more transparent.

> “You don't solve security with a black box. You solve it by having a culture of security in your organization. The challenge with AI is there are no people. So, you've got to kind of build that culture into the organization or into the AI models from the outset.”
>
> **Colin Wilson**
> Enterprise Architect, Verizon Business

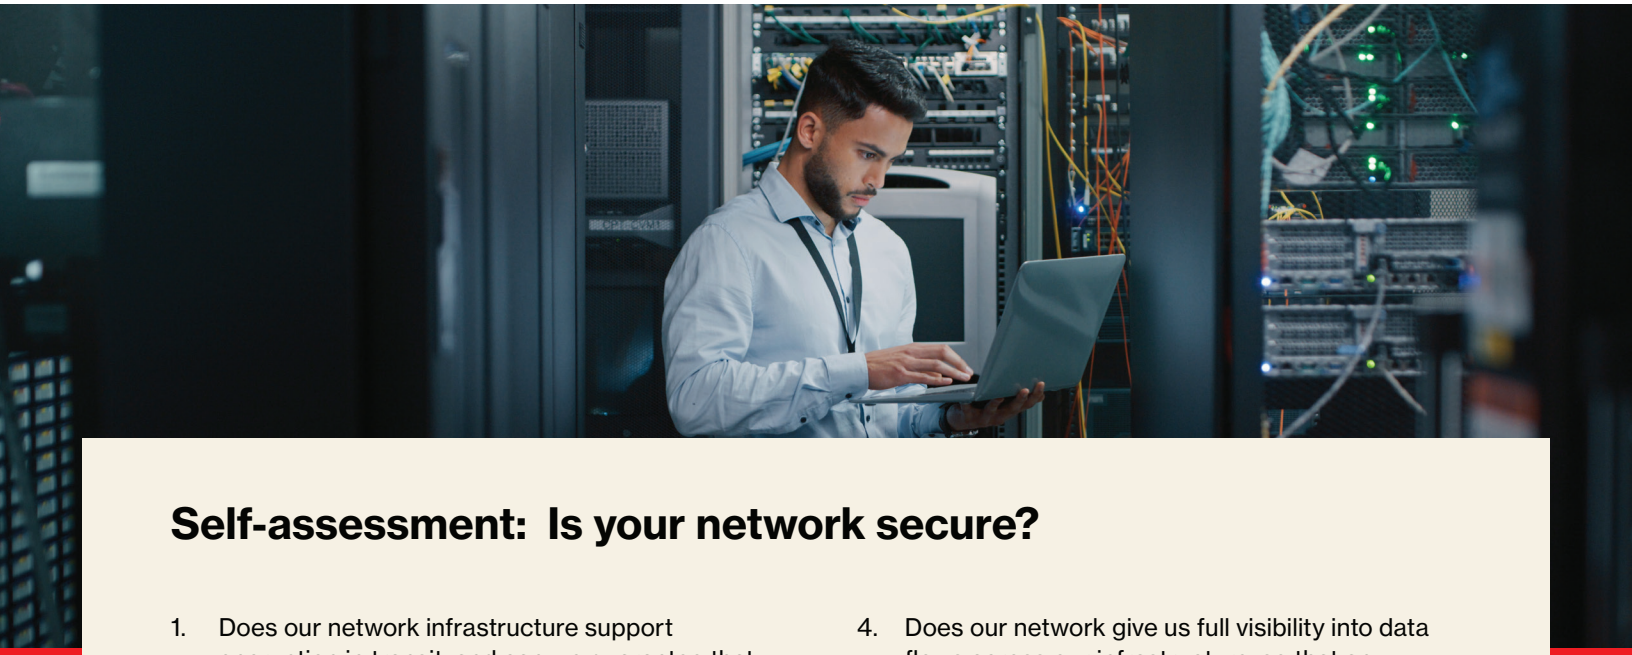

Self-assessment: Is your network secure?

1. Does our network infrastructure support encryption in transit, and can we guarantee that data cannot be intercepted or manipulated as it moves?
2. Do we have the architecture in place to control and monitor where data flows, ensuring it only crosses jurisdictions that meet our compliance obligations?
3. Is our security built on zero-trust principles, ensuring that no device, user or AI agent can move across it without authorization at every step?
4. Does our network give us full visibility into data flows across our infrastructure, so that any unusual or unauthorized movement of data is detected and flagged in real time?
5. Do we understand the capabilities of each agent we build and run? What data can it access and what is the destination of its outputs?
6. How are we modelling and mapping the interactions (both internal and external) between agentic systems? Can we be sure the limits of capability and risk are understood?

The invisible engine of AI

The right network is fundamental to AI success

The success of any AI initiative hinges on a robust, secure and high-performance network. Businesses need a network and connectivity infrastructure with scalable bandwidth and incredibly low latency.

While network infrastructure may typically be viewed as a cost, especially as data volumes increase, it is really a key strategic asset. After all, AI innovation can be effectively monetized[9], as long as critical loads do not impact business as usual.

Before you can reap the benefits of AI, you need to be sure your network can handle the weight of high-bandwidth data and have low enough latency to ensure fast reaction to service and safety needs.

"You can't run a 2026 tool—like a digital twin—on a 2005 network infrastructure," says Sundgren.

"You've got to think about how to size the network for all of that data that you can't yet take advantage of, but you will want to in the future," adds Wilson.

To help you understand your current and future networking needs, we've developed a handy reference, designed for operations of all shapes and sizes. It will highlight the network implications of AI initiatives, giving you an edge when building your infrastructure.

[9] https://www.verizon.com/business/en-gb/resources/deploying-ai-at-scale/

Network preparedness checklist

Is your infrastructure ready for AI?

Business need	AI need	Network impact
Achieve machine speed operations	AI can interact with other systems, requiring speed for real time responsiveness.	Extremely low latency in network communication is critical.
Enable real-time synchronous operations	Some automatic operations, such as autonomous vehicles, need real-time synched data.	Synchronous data exchange in network communication should happen at the edge.
Handle exponential data growth	Emerging AI tools are generating a massive increase in data flows and token usage.	An increase in overall bandwidth demand and capacity is required.
Unify internal workflows	Agentic interactions inherently require communication between systems within the data center.	Agentic network traffic patterns shift significantly from 'north-south' to 'east-west'.
Maintain robust cybersecurity	AI tools can significantly expand the attack surface, necessitating advanced threat detection.	Cybersecurity solutions must operate at machine speed to detect and neutralize threats.
Ensure continuous task performance and context	Today's AI tools require consistent access to memory data to maintain context, learn and perform continuous tasks.	Reliable, high-speed network access to distributed memory and storage is essential.
Achieve explainable AI	EU regulations will require AI output to be evidenced and explainable, turning the "black box" into a "glass box".	Networks will need to be able to cope with extra traffic and enable the secure storage of attributable data for compliance.
Maintain AI sovereignty	Full governance and control over AI data usage is critical to ensure data privacy, ethical alignment and prevent the leaking of proprietary data.	Secure, isolated network environments, robust identity and access management – as well as potentially new protocols for agent authentication – are required.
Optimize performance	The critical need for ultra-low latency and persistent memory access dictates optimal placement of AI tools.	Your network needs to bring AI agents and their data closer to where decisions are made, to the edge.

Your platform for the next era of enterprise connectivity

Verizon Business unlocks network success with unparalleled insight and experience

As AI moves from pilot to production, your network should no longer be seen as simply a utility, but rather as the strategic asset that it is.

The enterprises that will win in the AI era are not simply those with the best models or the most data. They are those with the infrastructure to move that data reliably, securely and at speed—across every location, every cloud, every device, every sensor and every user.

Yet, for a majority of executives, infrastructure remains the single biggest barrier to scaling AI effectively. Legacy environments built in silos over decades weren't designed for the demands of today's AI workloads, let alone tomorrow's.

With agentic AI on the horizon, bandwidth requirements are expected to increase tenfold by 2030. The network must be ready.

The tier 1 global network supporting you

Verizon Business is uniquely positioned to be the partner businesses need. As a tier 1 global provider operating across 180 countries, we deliver the reach, resilience and performance that modern AI demands—from the factory floor to the cloud edge. We also use AI in our own business for predictive network optimization, real-time customer service intelligence and data-driven decision making. This gives us a practical understanding of AI that goes far beyond theory.

A network, at your service

Most enterprises today are managing dozens of suppliers, contracts and networks across multiple regions. It's an enormous challenge and one that's holding businesses back from realizing the true potential of AI.

Verizon Business Network as a Service (NaaS) changes that. Our global, cloud-first platform brings together global connectivity, cloud links, managed SASE, AI-ready Wide Area Network (WAN), embedded security, consulting and digital experience management into a single, software-defined solution—all managed by Verizon Business.

One contract. One partner. One platform.

Built around survivability, zero trust, efficiency, visibility and AI readiness, this service delivers consistent, secure performance across every user, application and location.

Our experience is yours

We've spent decades managing and deploying global networking solutions, but we know today it's not just about infrastructure. It's about being the partner that's built around the challenges you're actually facing.

The AI opportunity is here. With Verizon, your network—the invisible engine behind it all—will be ready.

This content was created with the assistance of artificial intelligence and validated by humans.