NOKIA

Build first, lead forever.

Why AI infrastructure is America's next advantage.

Executive summary

The second wave of AI has arrived, and it is already reshaping how America works, builds, and competes. Nokia surveyed over 1,000 U.S. technology decision-makers across high-growth industries to understand where AI adoption stands, how fast deployment is accelerating, and what bottlenecks threaten the next stage of America's AI advantage.

What emerged is a nation moving quickly toward AI-powered productivity—but on foundations that were not designed for the for the internet era, not the connected intelligence era. Latency, energy scarcity, and fragmented data environments are now slowing the country's ability to deploy AI at scale and capture its full value.

Our research reveals the central tension shaping America's AI trajectory:

- **Adoption is racing ahead.** 97% of companies say AI is already transforming their business or society at large, and 74% report they are piloting or running AI in production today.
- **Infrastructure is not keeping pace.** 87% of firms say constraints in connectivity, power, or compute are limiting their AI efforts today, and 90% expect these constraints to hinder future scaling.
- **Security challenges are intensifying.** 90%of respondents believe AI introduces new forms of digital vulnerability, and 93% say digital sovereignty is "very" or "extremely" important to maintaining trust.
- **Competitiveness is at stake.** 1/3 companies says it may relocate AI workloads to regions with more reliable power, better networks, or stronger data environments, shifting the geography of innovation and threatening local economies.

The conclusion is clear: If America wants to lead the AI century, it must put infrastructure first to enable four foundational unlocks:

- **Connect Every Node:** Densify fiber, edge, and cloud interlinks. Guarantee sub-30-millisecond latency for enterprise and civic AI alike.
- **Empower Every Region:** Align state incentives to match local assets: power, talent, and land. Encourage balanced development rather than cluster congestion.
- **Secure Every Layer:** Make network integrity and data sovereignty non-negotiable. Trusted infrastructure must be treated as a strategic asset, not a vulnerability.
- **Include Every Citizen:** Use AI investments to create jobs, apprenticeships, and ownership models that keep benefits local.

Every technological revolution has depended on infrastructure to carry it forward. The United States now faces that same crossroads. To seize the economic and strategic potential of AI—not just adopt it but shape it—the nation must build first. Only then can it lead the global AI race.

The second wave: From experimentation to execution

America thrives on possibility. Its pioneering spirit, the drive to explore, experiment, and innovate, has powered human progress from the first flight to the microchip. Today, that same spirit hums inside the AI factories of a new industrial age. The whir of servers, the pulse of tokens, the loop between prompt and response: these are the rhythms of the second wave of AI, the new operating system of American ambition.

Between late 2022 and 2024 ChatGPT's launch pushed large language models from research into everyday use, marking the first popular wave of AI. That wave was a cascading discovery built on transformers and foundation models, where scale, more data, parameters, and compute became Silicon Valley's lingua franca and emergent abilities started to feel almost magical. But LLMs were only the most visible front: in the same window, diffusion models (e.g., DALL·E 2, Stable Diffusion, Midjourney) reinvented image generation; speech and audio systems (e.g., Whisper and modern voice/music models) leapt forward; and breakthroughs like AlphaFold and new robotics showed AI's reach beyond text. LLMs gave the broad public and businesses a usable interface to all this progress, turning AI from science fiction into practical possibility.

Now comes wave two, the era of return on investment. This phase is defined not by demos or hype, but by deployment: real business impact, measurable productivity gains, and sector-wide adoption. AI and data center companies like Microsoft, Google, OpenAI, Anthropic, and Meta are embedding AI copilots into productivity suites, customer service platforms, and developer tools. Healthcare systems are testing AI-assisted diagnostics and personalized treatment plans; manufacturers and logistics firms are piloting autonomous supply networks; utilities and energy providers are using AI to optimize grids and manage peak demand.

At the same time, the policy environment is catching up. The U.S. Executive Order on AI, the NIST AI Risk Management Framework, and new sector-specific oversight, from FDA pathways for medical AI to FTC guidance on algorithmic transparency, signal a shift from experimentation to structured, accountable deployment.[1] It marks the true beginning of the AI supercycle: the moment when the technology starts to pay its own way, reshaping productivity, governance, and growth for years to come.

The conviction around this shift is unmistakable. In our national survey of technology decision-makers within the leading growth industries in the US today, 97% describe AI as having a clear impact on their business and American society as a whole.

This transformation is not hypothetical. AI is no longer an experiment, but a new tool embedded in day-to-day business. 74% of all companies report they're running AI pilots or integrating AI into live products and operations, while virtually all have included it into their growth plans. Among the respondents, 51% already have private AI clouds, 50% use programs and natural language interfaces to help human-AI collaboration, and 47% use AI to aid compliance.

The conviction around this shift is unmistakable. In our national survey of technology decision-makers within the leading growth industries in the US today, 97% describe AI as an “incremental” or “transformative” force.

This transformation is not hypothetical. AI is no longer an experiment, but a new tool embedded in day-to-day business. 74% of all companies report they’re running AI pilots or integrating AI into live products and operations, while virtually all have included it into their growth plans. Among the respondents, 51% already have private AI clouds, 50% use programs and natural language interfaces to help human-AI collaboration, and 47% use AI to aid compliance.

Existing AI-related service capabilities*

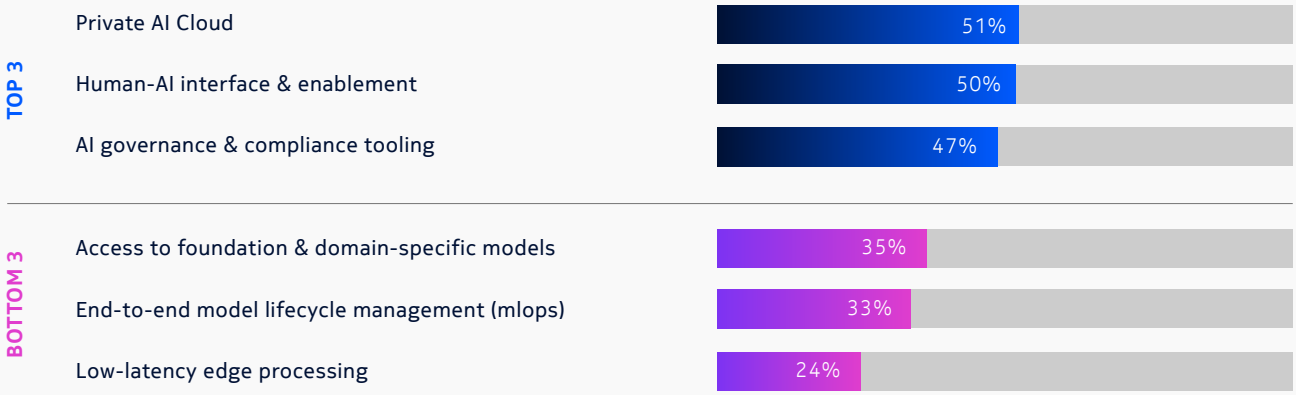

*Q8A. Which AI-related services or capabilities do you currently have?
Gemic for Nokia, 2025 Survey of American executives (N=1000)

As Nvidia’s Jensen Huang put it at GTC 2025, “It started with perception AI…Then generative AI…Now we’re entering the era of physical AI—AI that can proceed, reason, plan and act”.2 Physical AI has an almost infinite number of applications, with its only constraints found in the availability (or otherwise) or a few key assets: specialized compute, real-time connectivity, energy, and trusted networks. The companies that can provide an abundance of those assets will define the second wave, and so will the countries that house them. In other words, America may have invented the models upon which this grand undertaking rests, but it has no guaranteed monopoly on the value that will eventually spring from them.

Therefore, the next few years really matter. OpenAI’s Sam Altman has described the buildout of AI infrastructure as the biggest joint industrial project in human history, and if America can make AI abundant and useful for every sector, from manufacturing and finance to medicine and media, America will not just assert leadership in a critical industry and a once-in-a-generation Supercycle, but prove its continuing technological supremacy over competitors who are fiercer than ever.[3]

America must first make sure it can carry the load. Connectivity is essential to realizing this second wave of AI, and the underlying performance of networks will face increasing demands. Our study shows that companies' expectations are rising and becoming more specific, as they look for the second wave of AI to deliver real value. The delay—known as latency, the time it takes to transfer data—will be critical in determining the effectiveness of AI applications. An example: When asked about the network performance they'll require over the next few years, 72% say their AI applications will demand latency below 30 milliseconds. At that speed, every millisecond between data and decision can make the difference between advantage and obsolescence. and decision can make the difference between advantage and obsolescence.

Latency needs in next 2-3 years*

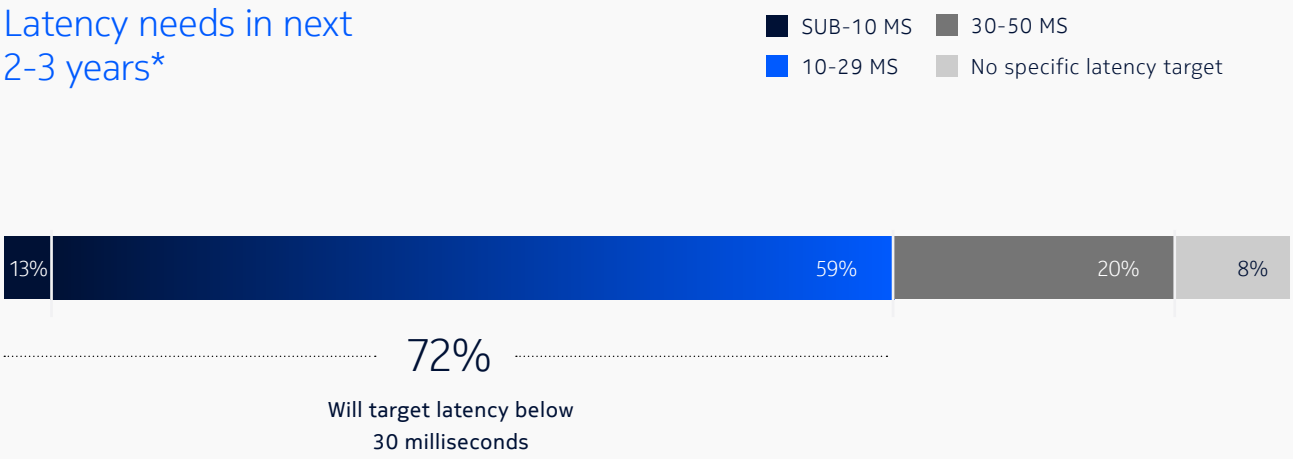

*Q7A. What level of latency will you need for your AI and data-intensive processes/products over the next 2-3 years? Gemic for Nokia, 2025 survey of American executives (n=1000)

This is a question of infrastructure: how quickly data can move, how much energy our systems consume, and whether the underlying networks can keep up. If the United States wants global AI leadership in the second wave, it needs to shore up its foundations. Otherwise, its plan for leadership will remain just that—a plan.

The hidden bottlenecks of the second wave of AI

I. Connecting America for next generation of intelligence

Speed has always been a defining American advantage, but not just speed of production, speed of possibility. In the last two decades, the U.S. became the birthplace of on-demand life. Uber turned a phone tap into transportation. Airbnb turned spare rooms into global travel infrastructure. Amazon normalized same-day delivery. Americans didn't just want things faster: they reorganized entire industries around immediacy, convenience, and scale.

That same spirit was driving Silicon Valley, where innovation became synonymous with speed itself. AI now extends that logic even further: it doesn't just make workflows faster, it accelerates the rate of innovation. In our survey, 35% of executives identified "speed of innovation" as one of AI's most meaningful advantages, and for good reason. Speed compounds, with every rapid cycle generating more data, sharper intuition, and better tooling, widening the gap with slower competitors ultimately providing more breakthrough potential.

Yet 59% of respondents identified connectivity as their single biggest barrier to scaling the technology: more than power, more than cost, and more than the algorithms themselves. You can see this in how AI has been taken up by companies thus far. Out of 15 AI-related capacities that firms already have, low-latency edge processing was the least common, only 24%, even though 43% think they'll need it in the future. These are stark statistics. The future is being throttled by the country's old wiring.

Consider what that means for the American economy. Modern mobility, industrial control, robotics, and defense systems all depend on near-instant feedback loops, cycles so tight they mimic reflex. Such precision demands compute, memory, and networks that share physical space: edge nodes next to factories; inference servers near ports; fiber lines cutting latency to almost nothing. But the reality is patchwork, with too many U.S. regions still relying on legacy backbones or shared clouds located hundreds of miles away.

When we asked companies how infrastructure limits affect their current AI efforts, 87% said they were “somewhat limited” to “entirely halted”. That’s a number you would expect in emerging markets, not in the country that invented both the microchip and the internet. Future expectations don’t look too good either, with 90% saying that infrastructure limits will constrain their plans to scale up AI. That kind of pessimism will do no good for overall business confidence and, at a more profound level, America’s tenacious leadership in the second wave of AI.

Infrastructure limitations impact on AI efforts*

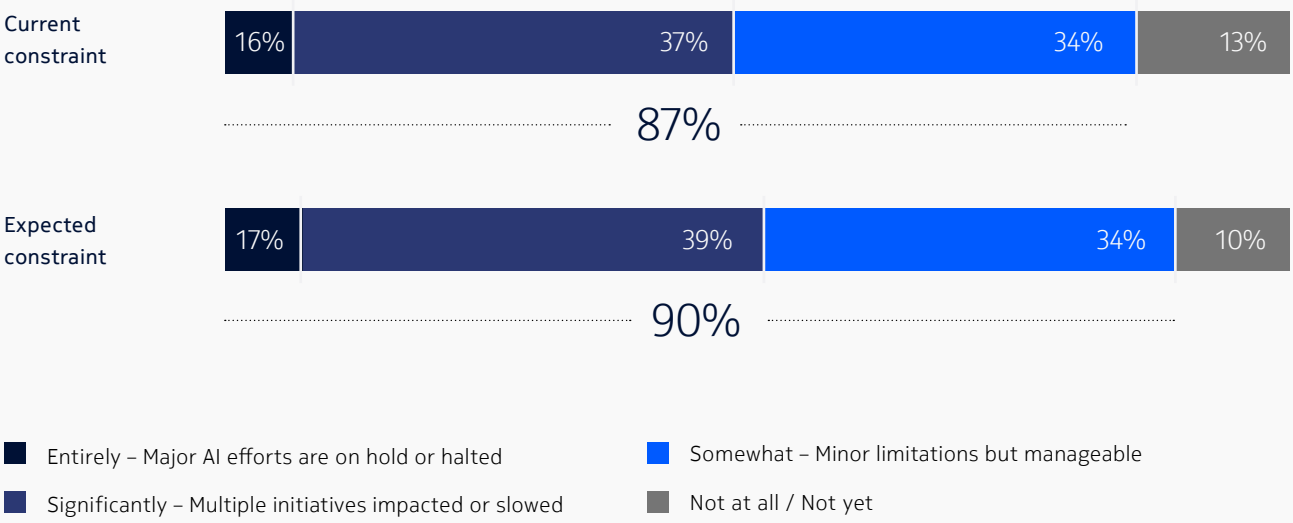

*Q12A. To what extent are your company’s current AI efforts currently constrained by infrastructure limitations (e.g. compute, bandwidth, latency, mobile/fixed networks, cloud, data center, or energy)?
Q12B. As your organization moves from pilots to scaled AI deployment, how much do you expect infrastructure limitations (e.g. compute, bandwidth, latency, mobile/fixed networks, cloud, data center, energy) to constrain your ambitions?
Gemic for Nokia, 2025 survey of American executives (N=1000)

II. Power: The great scarcity

If AI is the brain of our next digital world, then compute and networks are the neurons and synapses that let that brain think, learn, and act. We have spent decades building the nervous system that connects intelligence but the system only functions if it has a constant flow of energy. And right now, that lifeblood is running thin. In his “Abundant Intelligence” blog, OpenAI’s Sam Altman imagines factories capable of producing a gigawatt of AI infrastructure every week, but even he concedes that such scale will require breakthroughs across the entire stack: chips, robotics, and, critically, power.

Globally, the math is sobering. Research published in arXiv this spring (2025) shows that AI supercomputers double their performance roughly every nine months, while hardware costs and power use double every year.4 Meta’s new El Paso data-centre project, is designed to scale to 1 gigawatt (1 GW) of capacity.5 That is enough electricity to power roughly 800,000 to 900,000 U.S. homes for a year (based on the average U.S. household consumption of ~10,500 kWh/year).6 In the United States, much of the grid was built for industrial demand and domestic consumption, not the unparalleled and rising load of AI operations.

Our survey echoes this. When companies are asked to rank infrastructure pain points, energy availability, price, and diversity of supply consistently appear at the top, with 54% citing the unavailability of reliable and sustainable energy as a limiting factor in their ability to scale AI, and up to 86% stating that the energy costs will impact their AI expansion efforts. A pattern is emerging: firms are relocating workloads to states or countries with cheaper, cleaner, or more reliable power. Which is why power is not just a resource, but a strategic enabler, one necessary for America’s transition from the information age to the intelligence age to not end up patchy and lopsided.

Impact of energy costs on AI expansion*

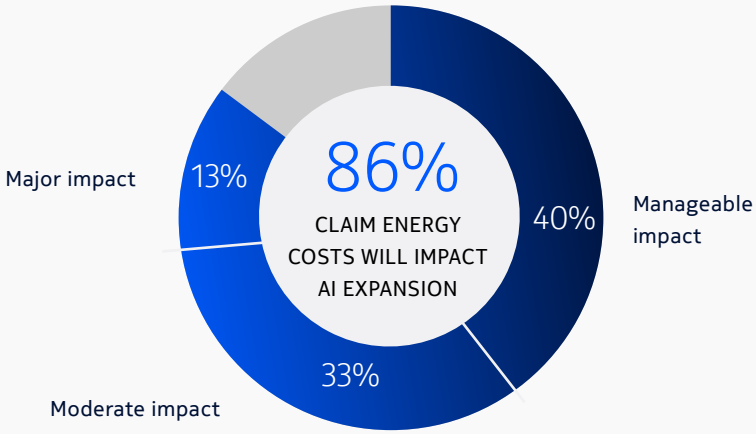

*Q43. How do you expect energy costs to impact your digital or AI expansion plans over the next 2-3 years? Gemic for Nokia, 2025 survey of American executives (N=1000)

III. The data dilemma

Even with enough power and bandwidth, intelligence is useless without data that's clean, current, and accessible. Our study revealed that 44% of companies cited data availability or quality as their top obstacle to scaling AI internally. That figure outranked every other internal issue. In practical terms, it means models are often starved not only of compute but of credible context. Bad or biased data doesn't just slow AI down, it corrupts it. Bad AI is not only bad for business, but for the reputation of the technology as a whole.

Scaling barriers according to AI decision makers*

TOP 3 internal barriers limiting AI adoption at scale

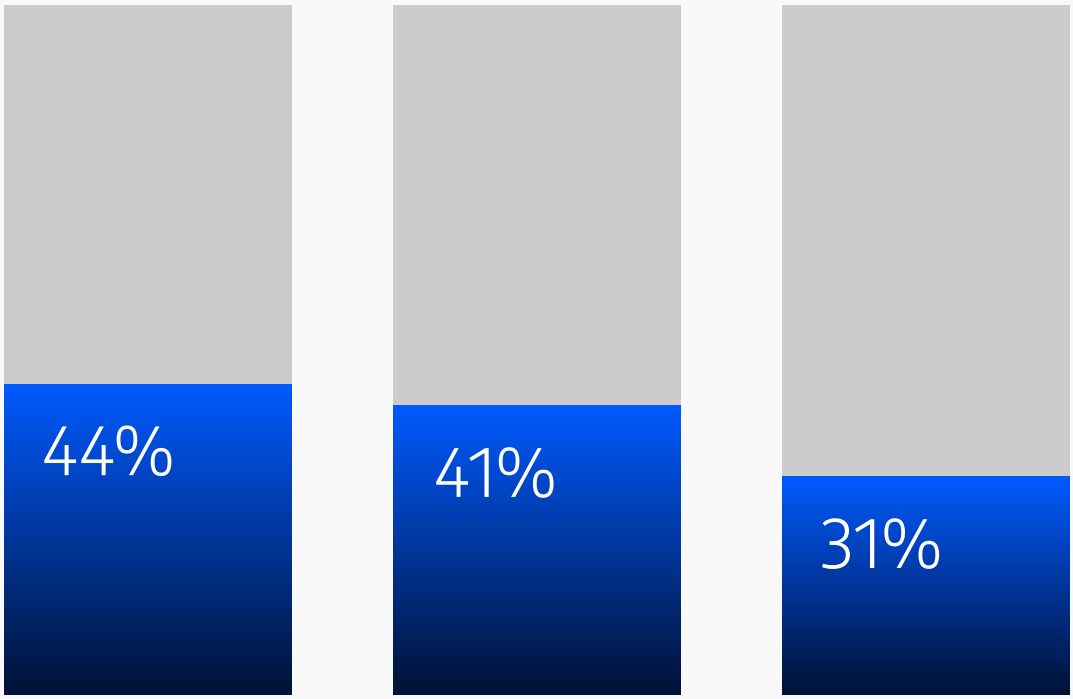

Data availability or quality

Lack of AI-skilled talent

Inadequate infrastructure (network, compute, storage)

*Q6. What are the biggest internal barriers limiting your company's ability to adopt AI at scale?
Gemic for Nokia, 2025 survey of American executives (N=1000)

While big tech has long defined, refined, and operationalized data as a core input to their business models, this is still an emerging capability for many SMEs. Yet it is now a prerequisite for capturing the value of the second wave of AI. As these organizations begin adopting AI, their needs for data collection, processing, storage, and secure communication will expand rapidly. And as they develop and deploy numerous smaller, specialized models tailored to their operations, their cumulative data demands will begin to rival those of much larger players, not in scale of company, but in aggregate volume.

Company responses if infrastructure limitations persist*

CURRENT AI CAPABILITIES

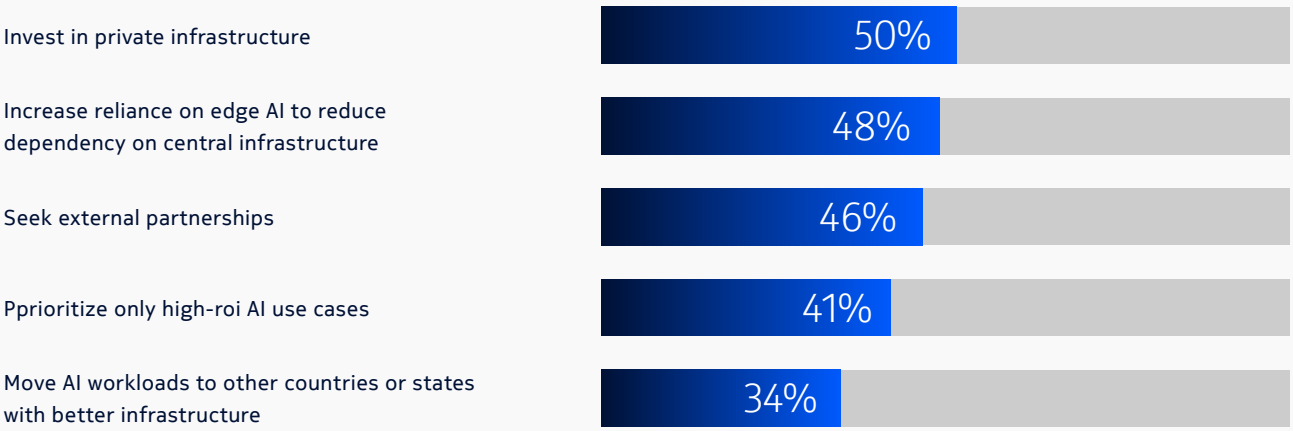

*Q20. If infrastructure limitations persist, what will your company most likely do?
Gemic for Nokia, 2025 survey of American executives (N=1000)

IV. The new geography of intelligence

What happens when companies can't get the power, connectivity, or clean data they need? They move. They relocate data centers across state lines. They shift inference loads to other regions, or other continents. They follow infrastructure like 20th-century manufacturers followed rail lines, with 1 in 3 companies saying they're likely to relocate AI workloads if local limits persist.

Respondents show clear gravitational pulls toward states perceived to offer better conditions: California (32%), Texas (23%), New York (31%), Florida (20%) and Washington (17%). When companies move, they don't just move servers; they take high-value technical jobs, construction and maintenance contracts, long-term service spend, and the future employers those assets attract.

Infrastructure limitations are prompting relocation of AI clusters*

Most popular us states to consider relocationg to (Among those would consider moving within the us)

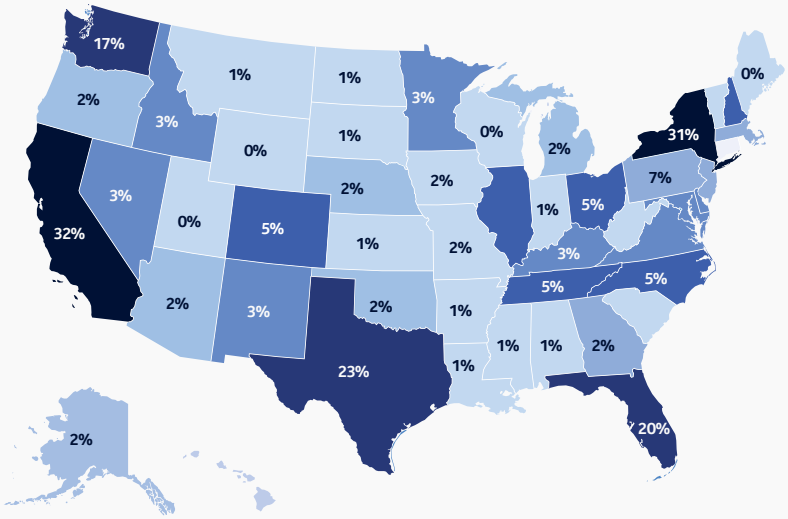

*Q20B. Which state(s) specifically would you consider relocating to?
Gemic for Nokia, 2025 survey of American executives (N=1000)

Regions that once hoped to build AI clusters instead face a slow unravelling: tax revenues shrink, universities lose industry partnerships, local talent migrates, and innovation pipelines wither. What begins as an infrastructure gap becomes a talent gap, then an economic gap, then a relevance gap.

But the reverse is also true. Consider the rise of Silicon Valley, where early investment in research infrastructure and digital connectivity created a self-reinforcing ecosystem of talent, capital, and innovation. Or look to Raleigh-Durham today, where sustained public–private commitment to high-speed networks, universities, and advanced manufacturing is attracting new AI ventures and reshaping local economies. With adequate investment, such regions don't just keep pace, they set the pace.

V. America needs its own blueprint for secure infrastructure and safe governance

All problems already outlined are urgent, and digital sovereignty is adding another layer of pressure. The companies surveyed had security top of mind: 90% are concerned that AI will introduce new, complex threats to America's digital infrastructure, with 68% actively increasing their investment to address the issue.

Security concerns are intensifying in the second wave of AI*

90%

are concerned AI will introduce new, complex security threats to America's digital infrastructure

*Q22. How concerned are you that AI will introduce new, complex security threats to American digital infrastructure? Gemic for Nokia, 2025 survey of American executives (N=1000)

Practices like using US-based cloud services, or on-site AI models for sensitive tasks, are widespread. This goes beyond immediate security concerns to maintaining a fundamental degree of societal trust in the entire AI revolution: 93% of companies say that digital sovereignty is "very" or "extremely" important to that specific end.

Sovereignty is no longer just about where data is stored, it's about ensuring that national networks are not exposed to new classes of AI-driven vulnerabilities. As AI fuels exponential growth in traffic, decision-making, and automation, network security becomes an extension of national security.
In an AI-enabled environment, both mobile and fixed networks are increasingly controlled,

optimized, and even defended by machine learning models that can themselves be manipulated. Attack surfaces expand as network functions become virtualized, inference workloads shift rapidly across distributed edge sites, and data centers serve as both energy-intensive processing hubs and strategic intelligence assets.

A compromised inference engine in a mobile core, or a poisoned dataset feeding a predictive maintenance system, could quietly degrade performance across national grids, logistics networks, or emergency communication systems. In many cases, this would happen long before human operators detect the threat.

Priorities to maintain data security & sovereignty*

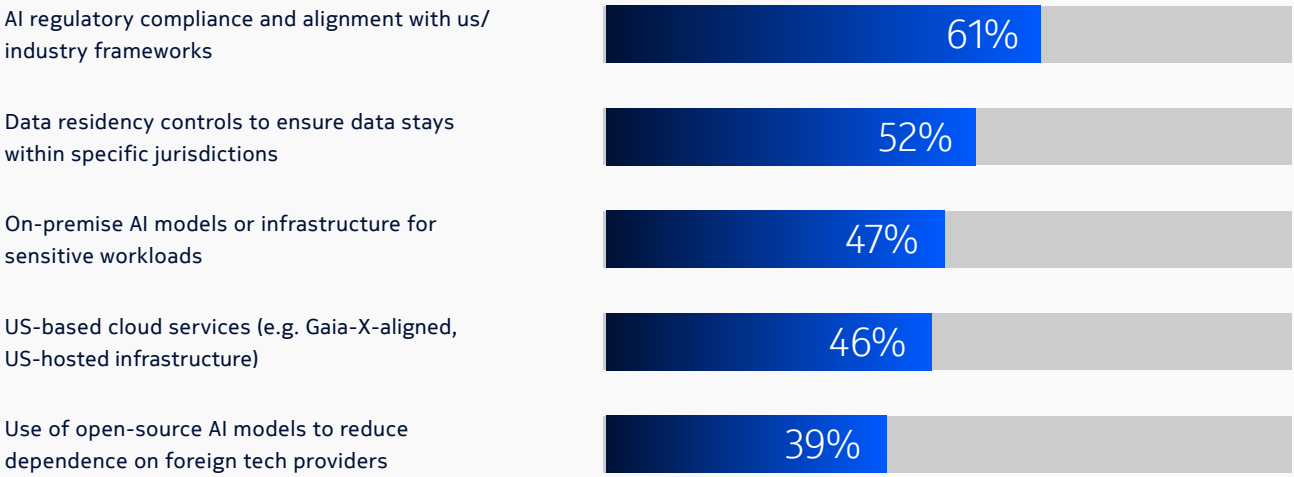

*Q46. Which of the following solutions does your organization prioritize to maintain data security and sovereignty? Gemic for Nokia, 2025 survey of American executives (N=1000)

To align with America's interests, AI must first be properly sovereign. So what does that look like? How do we build AI infrastructure the American way?

Building it the American way

For all America's inventive power, its capacity to build, to translate breakthrough into backbone, has historically been an even more formidable strength. The transistor became Silicon Valley. The affordable car became the highway, which became suburbia. The Apollo mission became the microchip economy.

Now, in the second wave of AI, the materials are new: grid, connectivity and compute, but the payoff could be just as great as, if not greater than, past industrial revolutions. If the United States builds smart, and builds fast, the next trillion-dollar industries may rise not on the coasts but in Kansas, Alabama, Ohio and Nevada. What could be a better affirmation of the American Dream than a technology pioneered by Silicon Valley finding full economic expression in the country's heartlands.

The second wave isn't about who can code the smartest model, but who can wire the most resilient nation, one that keeps intelligence close to its ideals, its innovators and its citizens.

To do that the task for America is to put infrastructure first.

I. AI-native connectivity and new network architectures

The first requirement for success here is clear. AI needs speed, and speed is enabled by high performance connectivity. In the first wave, networks were just delivery channels: data in, data out. In the second, networks become part of the product itself. Every AI-enabled service, from logistics optimization to real-time translation, depends on latency, reliability, and bandwidth as much as on algorithms. Connectivity is now, more, than ever, a strategic asset.

In the emerging era of AI applications, from consumer to mission critical contexts (such as healthcare), fast and reliable connectivity is then no longer a convenience. It is a must. Hospitals and rural clinics alike now depend on real-time, AI-augmented decisions with zero delay. The same networks that guide emergency triage and surgical imaging in cities must also interpret scans and link patients to specialists hundreds of miles away. And for AI assistants to follow us everywhere—at home, in the car, or on a hiking trail—coverage itself must be universal. AI cannot pause for signal loss when diagnosis, safety, and daily continuity rely on millisecond inference. For this future, high-quality network coverage is essential not optional.

To deliver that performance, America must not only build coverage of resilient, high-speed connectivity across the nation, but to also redesign its network configuration. Traditional networks were built for consumption: for fast downloads, streaming and content delivery. But AI depends on bi-directional flow, with continuous uploads of sensor data, real-time training feedback, and cloud-to-edge dialogue.

A recent Bell Labs report shows that in AI-intensive environments, traffic is becoming far more symmetric as uplink demand accelerates. AI-driven creation, machine-to-machine activity, and real-time interaction are reshaping network loads and eroding the long-standing download-heavy pattern of the internet. As Bell Labs concludes, networks now must shift from content-centric delivery to intelligence-centric connectivity, architectures optimized not just for consumption, but for high-speed interactivity and continuous AI-powered exchange.7

In this new order, networks must be treated as mission-critical infrastructure as integral to national security as power plants or ports. Reliable, sovereign, low-latency connectivity underpins AI command systems, healthcare delivery, manufacturing automation, and national resilience. To lead the world in AI, America must lead in networks: trusted, with built-in redundancy, and fast enough to think in real time.

Old vs. New Network Architecture*

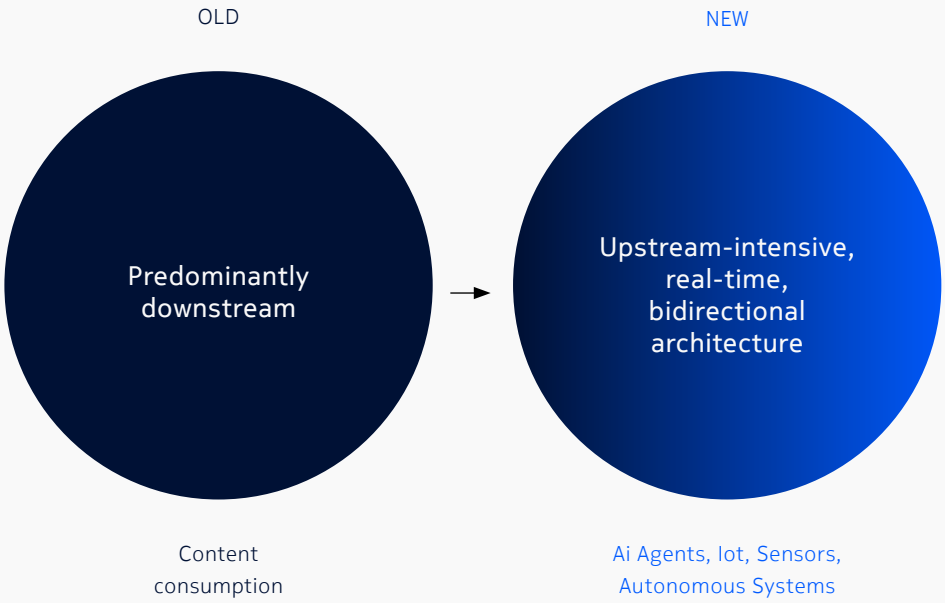

*Nokia Bell Labs. (2025). Global network traffic report: Understanding the growing impact of advancing technologies on future networks.

II. AI needs a new power mix and a grid able to support it

We know that energy scarcity threatens growth, but the solution is more than just more. Power must be smart as well as abundant: clean, distributed, resilient. The next generation of AI infrastructure must be designed not just to consume power efficiently but to reevaluate the social contract between business and communities.

That means balancing scale with legitimacy. A grid that fuels AI while driving up local electricity bills will not last. Communities that see no share of the upside will push back, and social license will evaporate. This is already a live issue: a recent Bloomberg investigation found that the wholesale price of electricity has leaped as much as 267% over the past five years in areas near American data centers.8

Energy sources may be different in the future but regardless of their origin or process, 89% say it's "extremely" or "very" important that AI infrastructure growth occurs in harmony with America's climate and social values.

Importance of climate and social commitments in AI infrastructure build-out*

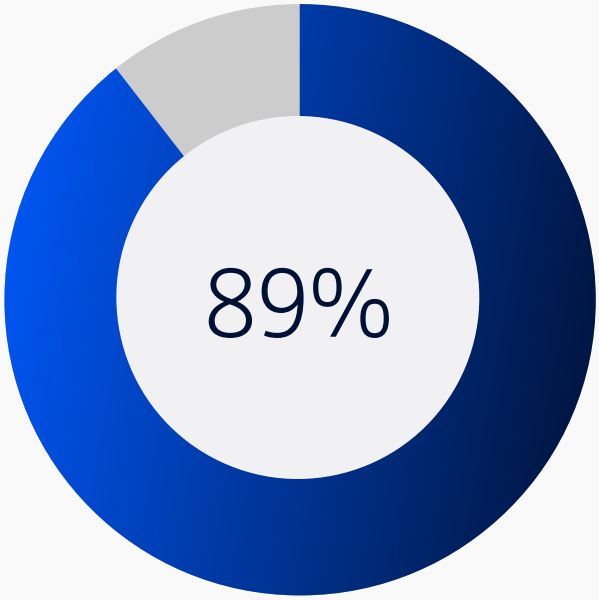

It is important that AI infrastructure aligns with America'S climate & social pledges

*Q18. How ambitious should the us be in developing its own AI-ready internet infrastructure?
Gemic for Nokia, 2025 survey of American executives (N=1000)

Building the American way is not just a decorative idea but is embedded in the ethos of the executives leading the industries in the second wave of AI. 91% of companies think America should be “ambitious” in developing its own AI-ready internet infrastructure, to the point where it is leading the world rather than just keeping pace with competitors.

Strong consensus: U.S. leaders call for ambition in AI infrastructure development*

91% Feel the US should be ambitious

*Q18. How ambitious should the us be in developing its own AI-ready internet infrastructure?
Gemic for Nokia, 2025 survey of American executives (N=1000)

The next wave of AI is an unparalleled opportunity to unleash American ingenuity by bringing together the country’s best ideas, boldest builders, and most ambitious communities. To make that possible, we need to modernize the American AI stack so that advanced compute, real-time intelligence, and low-latency performance is available to everyone in America.

To win the global AI race, America must put infrastructure first to enable four foundational unlocks:

1

Connect every node

Densify fiber, edge, and cloud interlinks. Guarantee sub-30-millisecond latency for enterprise and civic AI alike.

2

Empower every region

Align state incentives to match local assets: power, talent, and land.
Encourage balanced development rather than cluster congestion.

3

Secure every layer

Make network integrity and data sovereignty non-negotiable. Trusted infrastructure must be treated as a strategic asset, not a vulnerability.

4

Include every citizen

Use AI investments to create jobs, apprenticeships, and ownership models that keep benefits local.

Building AI the American way means developing and scaling these technologies intentionally and ensuring that innovation is not dictated by geography, capital concentration, or infrastructure proximity, but enabled by a network designed for access, dynamism, and national competitiveness. In doing so, it opens the next frontier of American ingenuity.

Sources

1. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). NIST. https://doi.org/10.6028/NIST.AI.100-1
2. Huang, J. (2025, January 6). NVIDIA CEO Keynote Livestream | CES 2025 [Video]. YouTube. https://www.youtube.com/watch?v=k82RwXqZHY8
3. Altman, S. (2025, September 23). Abundant Intelligence. Sam Altman blog. https://blog.samaltman.com/abundant-intelligence
4. Pilz, K. F., Sanders, J., Rahman, R., & Heim, L. (2025). Trends in AI Supercomputers (2504.16026). arXiv. https://arxiv.org/abs/2504.16026
5. Wang, E. (2025, October 15). Meta commits $1.5 billion for AI data center in Texas. Reuters. https://www.reuters.com/business/meta-commits-15-billion-ai-data-center-texas-2025-10-15/
6. U.S. Energy Information Administration. (2024, April 30). Use of energy explained: Use of energy in homes. EIA. https://www.eia.gov/energyexplained/use-of-energy/homes.php
7. Nokia Bell Labs. (2025). Global network traffic report: Understanding the growing impact of advancing technologies on future networks. Nokia.
8. Bloomberg. (2025). AI data centers are fueling America's energy bill crisis [Online Article]. https://www.bloomberg.com/graphics/2025-ai-data-centers-electricity-prices/

Methodology

This white paper is based on new research conducted by Nokia in collaboration with Gemic, a growth strategy firm. The foundation of the work is a quantitative survey of 1,000 senior decision-makers across the United States, carried out online by Dynata between September and October 2025.

Respondents were drawn from three key groups: telecommunication providers, executives in America's ten largest industries, and leaders from the ten fastest-growing sectors. All respondents held C-suite or director-level positions, ensuring the data reflects high-level strategic perspectives on AI, infrastructure,
and competitiveness.

In addition to survey data, the white paper is supported by secondary research conducted jointly by Nokia and Gemic, as well as expert interviews with Nokia leaders working across AI infrastructure, security, sustainability, and digital sovereignty.

About Nokia

Nokia is a global leader in connectivity for the AI era. With expertise across fixed, mobile, and transport networks, we're advancing connectivity to secure a brighter world.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

CID215238