



State of infrastructure in the agentic AI era

Key pivots for efficiency
and innovation from
1,402 global leaders.

2026

Executive summary

We're in the middle of an infrastructure renaissance. For decades, organizations built foundations that excelled in a world where applications moved at the speed of humans—as fast as a user could type or click. But agentic AI has fundamentally changed the speed of the game.

Multi-agent systems are rapidly collaborating, reasoning, and acting. Today, inference—the actual 'doing' of AI—accounts for the majority of all AI workloads, officially eclipsing training.

Traditional architectures are cracking under this pressure. According to our data, **62% of leaders** are seeing a costly 'inference tax' due to their legacy stacks—specifically data egress, storage bloat, and idle specialized hardware. The shift to agentic AI requires a departure from traditional foundations and a move toward dynamic, AI-optimized architecture.

To navigate this technical shift, our second annual report draws on the direct experiences of **1,402 global IT leaders**. Our findings show that innovative organizations are doing more than just upgrading hardware—they're reimagining their entire architectural approach. By shifting to distributed environments, securing workloads at the source, and embedding governance directly into the orchestration layer, leaders are ensuring their infrastructure is as ambitious as the agents it supports.

This report isn't just a survey of the landscape; it's a roadmap for establishing the new standard for production-grade autonomous systems.



Nirav Mehta
VP, Product Management
Google Cloud





Key findings

01

Agentic AI requires a new infrastructure standard

83% of organizations

require infrastructure upgrades to support production-grade agentic AI systems.

02

Governance is the prerequisite for scaling inference

4 out of 5 leaders

cite security, governance, and MLOps as the most significant challenges to scaling inference.

03

Hybrid multicloud is the new standard for AI deployment

52% of organizations

now use hybrid multicloud architecture, with 90% ranking edge deployment as important.

04

Efficiency is now a mission-critical requirement

91% of leaders

now factor power consumption into hardware selection.

The infrastructure mandate for agentic AI



Across the board, agentic AI is a growing global priority.

According to our research, **93% of technology leaders** recognize the value of AI tools and solutions, like agentic AI, and use them within their organizations. However, there is an increasing gap between the ambition of agentic AI and the infrastructural capability to achieve it.

A staggering 83% of organizations require infrastructure upgrades to support production-grade agentic AI systems.

Agentic workloads introduce a new set of infrastructural demands. A successful foundation for the agentic era needs to maintain context across complex workflows that span disparate data sources.

It has to be secure, providing granular data residency and security across on-premises, edge, and hybrid multicloud environments to meet the requirements of a fragmented global regulatory landscape.

And it should optimize for inference—delivering high-concurrency performance for always-on workloads while maintaining the cost-efficiency required for sustainable ROI.

Success in the agentic era requires more than just a cloud migration. To move from pilot to production, enterprises must optimize their infrastructure holistically, balancing the correct strategy across hardware, deployment models, security, and governance.

This report reveals how leading organizations are using their infrastructure as a catalyst for growth and innovation. From the practical challenges to strategic opportunities, we show how you can build an AI foundation as ambitious as the agents running on them—optimized for efficiency, security, and global scale.



Agentic AI requires a new infrastructure standard

Autonomous agents are transforming entire industries and demand a new approach. Success requires a foundation built for continuous intelligence—with cost-efficiency at the forefront.

The agentic readiness gap

AI has evolved from a tool that summarizes the world into an engine that acts upon it. In retail, it curates personalized styles and sources products instantly. In marketing, it builds and scales thousands of targeted campaigns at speed.


93% of senior IT decision makers recognize the value of AI and use it across their organization.

Yet, as AI workloads move from pilots to production, they are colliding with a functional ceiling—legacy infrastructure was not built for the demands of autonomous agentic systems.

For the 83% of organizations whose architecture requires upgrades, agents hit a speed ceiling—they can only move at the pace of the legacy systems they are confined within.

How prepared is your current infrastructure for agentic AI systems?

83% believe their infrastructure requires upgrades to support agentic AI



Survey question: How prepared is your current infrastructure to support the unique demands of agentic AI systems (e.g., persistent, stateful, long-running processes)?




Agentic AI demands a new scale of execution

The agentic era has fundamentally changed AI workloads. A single agentic prompt can trigger hundreds of downstream actions—independently browsing, querying, and executing across multiple systems.

A supply chain agent, for example, can identify low levels of stock ahead of a peak sales season. It proactively checks your real-time sales against historic patterns and live inventory, coordinates with procurement agents to source stock, and negotiates with carrier agents to manage shipping—all before preparing a purchase order for approval.

This shift toward autonomy has dramatically changed AI workloads. Inference now accounts for **47% of all workloads**¹—**surpassing training (28%)**—as organizations move toward production-grade agentic systems that require continuous reasoning and autonomous decision-making at scale.

Percentage of budget consumed by type¹



PUMA has enabled over 90,000 fans to co-create more than 350,000 custom football kits through its AI Creator platform—turning product design into a fan-driven experience. Across two activations with major European football clubs, PUMA partnered with Google Cloud to implement dynamic workload scheduling—scaling compute in real time to deliver fast, seamless creation at global scale while optimizing cost and performance.

¹IDC, 2026, [The AI Efficiency Gap: From Total Cost of Ownership Crisis to Optimized Cost and Performance](#)



Building a foundation for continuous intelligence

Agentic systems operate through continuous, high-speed, secure access to multiple data sources.

Agents need to connect to ERPs and CRMs that were not designed for machine-to-machine conversations. They need long-term memory to remember a user's preference from three weeks ago while executing a task today. And they must sustain this chain of thought at scale.

Without a foundation built for continuous intelligence, agents remain forgetful and limited. These hurdles are where the innovation gap is widest.

What are your biggest agentic AI infrastructure gaps?

43%

Difficulty integrating with legacy APIs and data sources

36%

Lack of specialized, high-throughput vector databases

35%

Insufficient security for multi-system access

Survey question: What are the most significant gaps in your current infrastructure when it comes to supporting agentic AI?

Reigniting cloud migrations

Enterprises' fragmented data landscapes and legacy infrastructure create significant barriers to successful AI adoption. These challenges starve AI of proper context and slow response times, which produces incomplete results and poor user experiences.

To succeed with AI, enterprises need to migrate critical workloads to AI-ready cloud platforms that provide unified data lakes, fluid compute, millisecond latency, and integrated AI development tools. This closes the fragmentation gaps and builds the foundation that organizations need to succeed with AI.

Advanced AI adopters, those with AI embedded in most workflows that actively seek out and adopt new AI solutions, are now **50% more likely** to use fully or primarily managed services than organizations that are still in the earlier stages of AI maturity.

By offloading lower-level infrastructure management tasks, such as cluster provisioning, scaling, and hardware abstraction to cloud providers, organizations reduce overhead costs and free up valuable time for internal teams to focus on innovation.

☑ Centralizing control and observability

To solve the fragmentation problem, organizations are choosing cloud platforms with **orchestration and observability** tools. These platforms provide a single point of control to manage the flow of data between agents and legacy systems, coordinate complex, multi-step workflows, and allow humans to watch an agent's thought process and intervene if it veers off course.



Sabre shifted its infrastructure from 19 expensive data centers to Google Cloud, leveraging [Cloud Spanner](#) for transactional consistency and [BigQuery](#) to manage 50 petabytes of data. Beyond the cost savings, the partnership enabled Sabre to maintain 99.999% SLA for its global marketplace.

Driven primarily through this migration to the cloud, the global travel technology leader has saved \$150 million annually—which doesn't include any other savings resulting from development efficiency (AI, standards, tools, etc.)—and significantly accelerated its pace of innovation. Sabre now has more than 70 AI solutions in production, including agentic APIs built using Gemini's ADK.

Cost efficiency is a top priority

With agents driving an inference boom, enterprises are shifting their focus from the possibility of AI to its long-term profitability. **96% of respondents** say that cost efficiency is important to their AI infrastructure decision-making process.

While compute is the most visible expense, other factors beneath the surface are driving up the total cost of ownership (TCO).

How important is cost-efficiency when making AI infrastructure decisions?



Survey question: How important is cost-efficiency in your decision-making process for generative AI infrastructure solutions?

What are the top hidden costs when scaling AI?

81%
Operational complexity and engineering overhead

This includes the countless hours engineers spend on manual “glue work,” patching together AI agents across disparate systems and databases, as well as the operational cost of maintaining security and compliance.

62%
Infrastructure and direct usage costs

These expenses stem from a combination of mounting data egress fees, ballooning storage needs, and budget wasted on over-provisioning specialized hardware that often sits idle.

57%
Workforce and talent

The premium for specialized AI talent is high, and upskilling existing teams can drain time and budget without the right approach.



Key takeaways

The most successful organizations are moving toward a managed foundation that cuts down on operational complexity.

Modernize your foundation

Advanced AI adopters are pivoting to modern cloud platforms, with **24% now using fully or primarily managed services**. They can build, deploy, and scale agents with purpose-built databases and integrated AI development tools, bypassing the bottlenecks of a patchwork system.

Build for scale

Orchestration platforms provide the repeatable scaffolding to connect agents and systems, automate workflows, and monitor everything in one place.



LiveX.AI reduced operational costs by 66% and freed up crucial time for their engineers. By offloading cluster maintenance to Google Kubernetes Engine (GKE) Autopilot, this startup reduced overheads and gave its engineers more hours to focus on agent behavior.



To efficiently support the massive scale and daily traffic of KakaoTalk, Kakao optimized infrastructure costs and performance by adopting Google Cloud TPUs and the JAX framework. Through this transition, they successfully developed the Kanana model family—specialized for the Korean language—and achieved a 2.7x increase in model throughput, particularly when utilizing the latest Trillium TPUs. Furthermore, by applying a Mixture-of-Experts (MoE) architecture to enhance inference efficiency, Kakao has built a foundation to flexibly deploy advanced AI technologies across various services used by most South Korean populations.



Governance is the prerequisite for scaling inference

AI agents have fundamentally redefined the security landscape. Leading enterprises are placing security and governance at the heart of their infrastructure strategy.



Autonomous AI has changed how we think about security


In the agentic era, the challenge is no longer about stopping threats at the perimeter—it's about getting security and governance right so that scaling AI to production becomes a seamless, repeatable process.

79% of tech leaders cite security, governance, and operations as the most significant challenges to scaling inference.

Mature governance has become a strategic prerequisite. Organizations with mature governance are accelerating AI adoption with more confidence,² while organizations without these crucial structures fall behind, creating an innovation gap as the market moves toward autonomous execution.

Robust governance provides the freedom to innovate at speed. Instead of acting as a hurdle, a solid security and governance foundation acts as a launchpad, determining whether an organization can confidently lead the market or if it will slow down to address emerging risks.

What are your top challenges when scaling AI inference?



Survey question: What are your top challenges when it comes to scaling generative AI inference?




The agentic security paradox

Agents are designed to act—to read emails, query databases, and execute API calls. The question is, how do you give agents enough freedom to be useful with enough guardrails to be safe and secure? This is the defining governance challenge of 2026.

35% of senior IT decision makers cite “insufficient security for multi-system access” as a primary gap preventing agentic deployment.

Legacy security is built for humans, not autonomous agents. Agents require more access to fulfill user queries, but are vulnerable to new vectors. Examples include indirect prompt injection and tool poisoning, where attackers hijack an agent’s logic through the data it processes or the APIs it calls—turning an agent’s ability to act into a liability.



Securing the chain of thought is paramount

Traditional cybersecurity focused on the perimeter. AI security must focus on the model and its provenance—its origin, data sources, and training history.

Securing the AI/ML supply chain (48%) is the top security concern for leaders. Enterprises are focused on protecting their models from compromised training data or vulnerabilities introduced via open-source weights. The risk is that a compromised component deep in the dependency tree could introduce latent behaviors that only trigger under specific conditions.

As models run on shared infrastructure, the fear of **data leakage in multi-tenant environments (41%)** is also acute. This is particularly critical for GSIs and SaaS providers who must guarantee that one customer's prompts do not influence the weights or context of another customer's model.

Top security concerns around AI infrastructure

48%

Securing the AI/ML supply chain

41%

Protecting data in multi-tenant environments

39%

Unauthorized model access and extraction

Survey question: From an infrastructure perspective, what are your top security concerns related to AI?



Governance is driving a move to full-stack cloud platforms

As AI matures from pilot to production, organizations are struggling to manage governance on patchwork stacks. Increasingly, leading enterprises are pivoting to integrated cloud platforms for simplicity and centralized control.

By moving from disconnected tools to a central control plane, organizations gain a single system of record for agent permissions and workflows. This helps them achieve something legacy setups can't deliver—straightforward governance at scale.


69% of surveyed executives now rate a full-stack platform as a critical requirement.

Governance is the primary driver of this shift, with **80% of leaders** identifying data compliance as the number one factor dictating their infrastructure choice.

For the modern enterprise, the cloud provider is no longer just a source of compute—it is the primary anchor for a secure AI strategy.

This reliance on major cloud providers has surged nearly 30 points in a single year, with **78% of organizations** now sourcing their gen AI solutions directly from their primary cloud partner—up from just 48% in 2025.³

³ Google, 2025, [State of AI infrastructure](#)





Key takeaways

Effective governance and security in the agentic era requires clear frameworks at the heart of development and unified architecture behind your AI.

Make AI “secure-by-default” with SAIF

Instead of reacting to problems after they happen, use the [Secure AI Framework](#) (SAIF) to make security a natural, embedded part of building AI. This framework helps mitigate AI-specific risks like prompt injection or compromised data, deploy AI faster, and scale with greater confidence.

Centralize oversight on a unified platform

Managing agents across disparate systems creates risk. Use an integrated platform to act as a central control plane. These platforms unify permission management and workflow logic, allowing for human-in-the-loop oversight to flag when an agent requires approval before moving forward.



“

We are using [Model Armor](#) not only because it provides robust protection against prompt injections, jailbreaks, and sensitive data leaks, but also because we’re getting a unified security posture from [Security Command Center](#). We can quickly identify, prioritize, and respond to potential vulnerabilities—without impacting the experience of our development teams or the apps themselves. We view Model Armor as critical to safeguarding our AI applications, and being able to centralize the monitoring of AI security threats alongside our other security findings within SCC is a game-changer.”

Jay DePaul,
Chief Cybersecurity and Technology Risk Officer,
Dun & Bradstreet



Hybrid multicloud is the new standard for AI deployment

Agentic readiness means moving toward a hybrid architecture that integrates the power of the public cloud, edge performance, and sovereign governance into a single, cohesive standard.




The shift to hybrid multicloud architecture

In the agentic era, the debate between public cloud and local computing has settled into a clear consensus: hybrid is the destination.

Hybrid multicloud adoption surged 26.8% YoY, moving from 41% in our 2025 report⁴ to 52% of organizations this year.

This 11-point jump signals a fundamental pivot. Enterprises no longer view hybrid as a transitional state—it's the strategic choice that enables them to balance the raw power of the cloud for model training with the speed of local environments for inference.

Cloud infrastructure approaches for generative AI



Survey question: Which cloud infrastructure approach does your organization primarily use for generative AI workloads?



Recursion's hybrid architecture helped reduce costs by 50%. Using both its on-premises supercomputer and [Google Cloud infrastructure](#), Recursion can manage over 50 petabytes of biological data across a distributed stack, utilize TPUs to reduce costs by 50%, and maintain the massive data gravity required for its models.

Distributing AI to the edge

Organizations are rethinking where AI physically processes data to deliver unparalleled speed and performance. In addition to centralized cloud environments, organizations are running models closer to the source of data generation, including directly on mobile and IoT devices.

For years, the public cloud was the default laboratory for AI, providing the power needed to train massive models. However, as AI workloads expand from periodic training to always-on inference, relying solely on centralized infrastructure can introduce physical and financial limits.

To optimize performance, organizations are embracing a hybrid AI architecture, using the cloud where necessary while handling specific, high-volume inference at the edge for three core reasons:

Latency

For agentic workflows, the round-trip time to a central data center is a massive speed limiter, particularly for those involving voice or video. Low-latency, real-time inference must happen locally.


The inference tax

By running inference at the edge, organizations can handle the majority of their compute locally, reserving central cloud infrastructure for training and complex, cross-functional tasks. This dramatically reduces variable per-token costs.

Operational resilience

Edge deployments allow AI to continue functioning even when the primary network connection is interrupted.

How important is deploying AI at the edge for your organization?



Survey question: How important is deploying generative AI models at the edge (e.g., on IoT devices, mobile devices) for your organization?

“

With Google Distributed Cloud, we have been able to reimagine the possibilities of building and deploying innovative technology securely across NAPA's distributed locations. GDC is central to our strategy to simplify infrastructure and accelerate software delivery. It provides the agility to bring new features to our stores with greater speed and security while optimizing our long-term costs.”

Navin Anandaraj
Technology Fellow, Genuine Parts Company





Digital sovereignty and governance in a global landscape

The defining boundary for agentic AI is no longer just the network, but the physical jurisdiction of the data itself. In the agentic era, security, privacy, and international laws have reshaped the approach to cloud computing.

The push for **digital sovereignty** is driven by three concerns:


- The jurisdictional risk posed by foreign data requests and shifting regulations.
- The economic dependency risk created by a lack of domestic cloud infrastructure providers.
- The need to safeguard critical local services against global disruptions.

Together, these pressures are driving enterprises and governments to reclaim control over their data, infrastructure, and digital future.

48% of leaders are prioritizing infrastructure with data residency controls, supporting compliance with local data security laws.

75% of all enterprises outside of the US will have a digital sovereignty strategy by 2030.⁵

⁵ Gartner, 2025, [Gartner CIO Survey](#)





Key takeaways

Building for AI impact in 2026 requires a fundamental architectural pivot: moving intelligence closer to the data source and governance closer to the local jurisdiction.

Adopt a hybrid architecture strategy

Balance cloud hardware with edge computing. By moving high-volume inference to the edge, enterprises can reduce their hardware costs and increase speed, while reserving the central cloud for heavy training and complex reasoning.

Prioritize sovereignty-first design

Architect for local jurisdiction from day one to avoid legal bottlenecks and costly innovation gaps. Use managed cloud platforms that offer granular data residency and access controls to satisfy local laws, while maintaining a single, centralized control plane for global operations.



Deutsche Telekom proved that you can move massive amounts of data to the cloud while balancing strict local laws. By partnering with T-Systems and [Google Cloud](#), they built a platform that keeps their data strictly within European borders and under their own digital lock and key. This setup didn't just satisfy the sovereign regulations; it helped the company turbocharge the idea-to-insight journey on sensitive data much faster—proving that high security can be a speed booster rather than a bottleneck.



Efficiency is now a mission-critical requirement

Effective infrastructure optimization balances speed with energy efficiency, using the right hardware to improve the sustainability and productivity of compute workloads.



The energy equation:

Cost, capacity, and compliance mandates


In the agentic era, energy has moved from a technology concern to a boardroom priority.

91% of leaders now factor power consumption into hardware selection, with 61% rating energy efficiency as a primary or significant factor.

A stark new reality drives this shift—power consumption now presents a two-fold barrier to growth, with grid scarcity and regulatory compliance acting as equal constraints. Energy efficiency is no longer a secondary metric; it dictates both physical scale and market access. While the limited availability of power limits how much compute infrastructure can be provisioned and deployed globally, a rapidly tightening web of regional legislation is simultaneously turning energy management into a strict legal prerequisite for operation.

In hubs like Germany, new data centers must achieve a PUE (Power Usage Effectiveness) of 1.2 or lower and reuse 10% of waste heat. Ireland now mandates that large centers provide 100% on-site dispatchable generation to match their grid draw. These new mandates push enterprises to adopt energy efficiency strategies to remain compliant.

To what extent does energy efficiency factor into hardware choice?



Survey question: To what extent are power consumption and energy efficiency a factor in your selection of AI hardware and platforms for inference?



Key takeaways


For your organization to be a leader in the energy efficiency race, hybrid architecture and a diverse portfolio of accelerators are the way forward.

Balance the load with hybrid architecture

The most successful organizations are embracing a hybrid multicloud strategy. This allows for more efficient handling of high-volume workloads on local hardware, while reserving more costly queries for centralized cloud architecture.

Use the right accelerators for the job

To maximize energy efficiency, innovative organizations are using a diverse portfolio of accelerators. By matching the right hardware to the specific demands of each workload, they can better balance performance and energy efficiency.



AXIA Energia, Brazil's largest power company, helps prevent power outages for millions of customers by running advanced weather modeling on TPU clusters. By forecasting severe weather conditions up to 10 days in advance, the company can proactively plan and execute preventative measures before storms hit, significantly reducing service interruptions year over year.

- 60 million people covered.
- +99.9% transmission lines availability assurance.
- 15% YoY improvement in extreme weather accuracy.



Building for agentic production

The agentic era has fundamentally changed the infrastructure mandate. Success is no longer defined by who trains the largest models, but by who builds the most resilient, secure, and integrated systems to run them.



As organizations move toward agentic maturity, our research reveals critical lessons from leaders.

Bridge the gaps in your AI stack

Agentic workflows require context and memory across complex workflows that span disparate data sources. Leading enterprises are investing in infrastructure to connect legacy systems, automate complex workflows, and ensure human oversight is built in. It's the simplest way to keep AI projects moving in the right direction.

Robust governance accelerates AI innovation


To scale agents safely, security and compliance can't be afterthoughts. The right foundation provides granular data residency and security across on-premises, edge, and hybrid multicloud environments to meet local regulations. Robust governance doesn't slow down innovation, it provides the guardrails to experiment and deploy at speed amid emerging threats and evolving regulations.

Embrace a hybrid multicloud strategy

The AI footprint spans from data center to device, so leading organizations use a hybrid mix of cloud and managed services to stay flexible. This enables them to shift high-concurrency performance for real-time tasks to the edge for speed, while minimizing energy costs required for sustainable ROI.

Success in the agentic era is built from the ground up. Most leaders are planning on investing heavily in core infrastructure and operations—the foundation for further investments in high-performing data pipelines and a secure, resilient AI posture.

Most organizations plan to invest heavily in AI infrastructure this year



Survey question: Where do you plan to make the most significant investments in your AI infrastructure over the next 12–24 months?





Google Cloud is your partner in the agentic era

Success is defined by the **efficiency and reliability** of the agents you deploy.

Google Cloud turns infrastructure into an engine for growth, managing backend complexity so your teams can focus on building. With Google Cloud, you can:

1 Build on an AI-optimized stack

A full-stack platform combines performant hardware, AI-development tools, fully managed cloud storage, powerful analytics, cutting-edge models, and ready-to-deploy agents.

2 Scale on infrastructure optimized for the agentic era

Achieve exceptional performance and efficiency for demanding workloads with AI Hypercomputer. This integrated supercomputing system is engineered with co-design at every layer—from specialized compute, storage, and networking to software and consumption models. It gives you complete control over the latest TPUs, CPUs, and GPUs, providing the system-level advantage you need to confidently deploy agents at scale.

3 Accelerate innovation with decades of R&D

Launch grounded, governed, autonomous agents that are ready for impact. Get direct access to Google DeepMind's breakthroughs and use the same infrastructure powering Google's 2B+ user products.

4 Turn governance into a competitive advantage

With our Secure AI Framework and leading AI security products in Security Command Center, Model Armor, and Wiz, we turn security into an enabler that allows you to scale faster than the competition.

5 Adapt in a multcloud world

Whether on-premises, at the edge, or across clouds, we support your unique sovereignty, scalability, and connectivity demands with Google Distributed Cloud and Cross-Cloud Network.

Move beyond AI potential.
Lead the way in agentic production with Google Cloud.

Build agent-ready infrastructure today.

Contact us



Methodology

This report is based on quantitative research finalized in January 2026, drawing on a survey of leaders of global enterprises conducted by Google Cloud and GBK Collective.

1,402 senior IT decision makers surveyed

Exemplary titles included Infrastructure Architects and Technology Leads. Respondents held positions at the manager level or higher.

Across 12 countries

Countries included: United States, Canada, United Kingdom, France, Germany, Australia, Japan, Korea, Mexico, Brazil, China, and India.

GBK Collective is a decision-grade insights consultancy built to deliver clarity when it matters most. As the independent research partner for this report, GBK provided unbiased end-to-end execution, spanning initial scoping and survey design through to rigorous analysis and strategic recommendations.

The companies represented in this research had 1,000+ employees (500+ employees in Asia Pacific and Latin America).

Respondents were familiar with generative AI, considered AI important to their organization's future, and either had a workload utilizing generative AI at the time or planned to have one within the next 12 months.

