White Paper, sponsored by AMD
OCTOBER 2025

AI Infrastructure for Business Impact: Enabling Agentic Intelligence with Scalable Compute

Peter Rutten
Research Vice-President, Performance Intensive Computing, Worldwide Infrastructure Research, IDC

Madhumitha Sathish
Research Manager, High Performance Computing, IDC

Heather West, PhD
Research Manager, Performance Intensive Computing, Worldwide Infrastructure Research, IDC

Table of Contents

CLICK BELOW TO NAVIGATE TO EACH SECTION IN THIS DOCUMENT.

In This White Paper

AI has evolved from experimental use cases to integral enterprise capabilities. In 2025, the advent of agentic AI signals a new era, one that redefines the standards by which enterprises measure progress and potential.

The rise of agentic AI is reshaping enterprise expectations. As customer demands accelerate and competitive differentiation pivots on intelligent automation, organizations face mounting pressure to move beyond traditional AI use cases. Agentic AI offers the promise of dynamic, context-aware systems that can act independently across datacenter, cloud, and edge environments — for instance, reallocating compute resources in datacenters to meet SLAs and reduce energy use, orchestrating cloud deployments based on latency and compliance, and adapting at the edge to sensor failures or bandwidth limits in smart manufacturing, logistics, and urban infrastructure.

Yet enterprise readiness is uneven. Insights from IT/AI infrastructure decision-makers highlight both promising opportunities and critical gaps in enterprise AI adoption. As per an IDC survey conducted for this purpose, survey data from over 400 enterprises reveal that more than 75% of organizations lack clarity on agentic AI use cases (see Figure 1).

FIGURE 1:

Enterprise Uncertainty Around Agentic AI Use Cases

Q. In your opinion, does your organization have a clear understanding of agentic AI use cases?

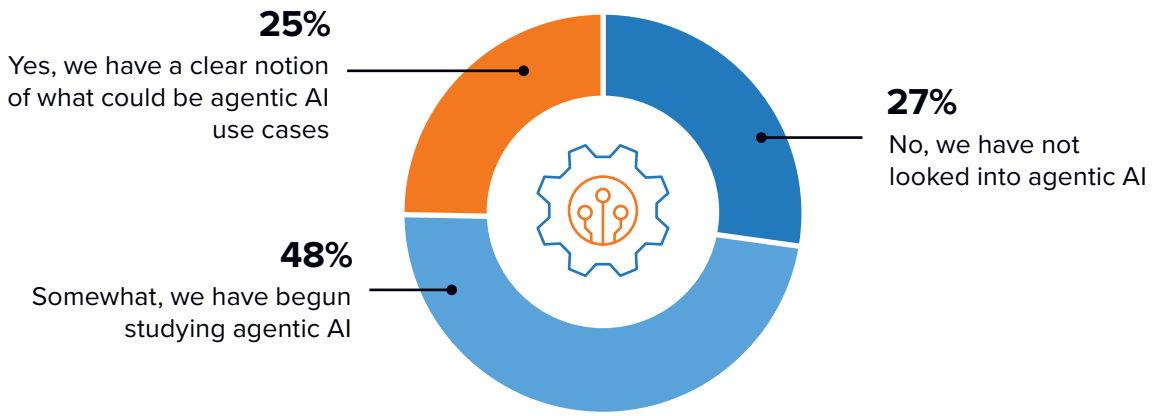

n = 410
Source: IDC's *AMD AI Survey*, July 2025

This gap underscores critical urgency. While AI adoption is on the rise, many enterprises struggle to define agentic use cases, navigate governance complexities, and manage infrastructure costs for scalable deployment.

To capitalize on the agentic AI wave, enterprises must align infrastructure strategy with business objectives, ensuring scalable compute, secure data pipelines, and low-latency execution environments. This white paper distills insights from IT and AI infrastructure decision-makers to offer strategic guidance for building AI-ready infrastructure that supports the next generation of intelligent systems.

Methodology

For this white paper, IDC conducted a quantitative survey and a qualitative study called an IDC QualBoard.

The study had the following characteristics:

- **Survey:** 410 respondents; QualBoard: 11 interviewees
- **Mode of survey and study:** Web based
- **Field time:** July 2025
- **Topics:**
  - AI strategy and use case selection
  - Data preparation and governance
  - Cost assessment and financial planning
  - AI infrastructure for model development and deployment (training and inferencing)
- **Geographies:**
  - United States
  - Canada
- **Minimum size of the organization:**
  - 1,000 employees

Industries:

- Financial services — 18%
- Healthcare and life sciences — 14%
- Retail — 11%
- Telecommunications — 10%
- Technology vendor — 6%
- Government — 6%
- Others – 35%

Respondents are IT directors, executives, or C-level who are:

- The primary decision-maker or a part of the team that makes decisions regarding AI server infrastructure purchases and management within their datacenter
- Knowledgeable about their organization's AI server infrastructure strategy and the servers that are used in their AI development and deployment environment
- Representing organizations that develop and deploy various types of AI models including generative AI and some that develop and deploy only generative AI models within their datacenter

Note that a datacenter can refer to the enterprise or academic own datacenter, a public or specialty cloud provider, a colocation datacenter, or a managed services provider.

Situation Overview

Need for Assessing the Business Impact of Identified

AI is driving transformation across sectors, including finance, healthcare, manufacturing, and retail. However, scaling from pilot to production presents significant challenges. According to IDC survey data, 83% of enterprises launch fewer than 10 AI use cases simultaneously, indicating a cautious and fragmented approach to scaling (see Figure 2).

FIGURE 2:
Enterprise AI Adoption

Q. How many AI use cases does your organization typically launch simultaneously?

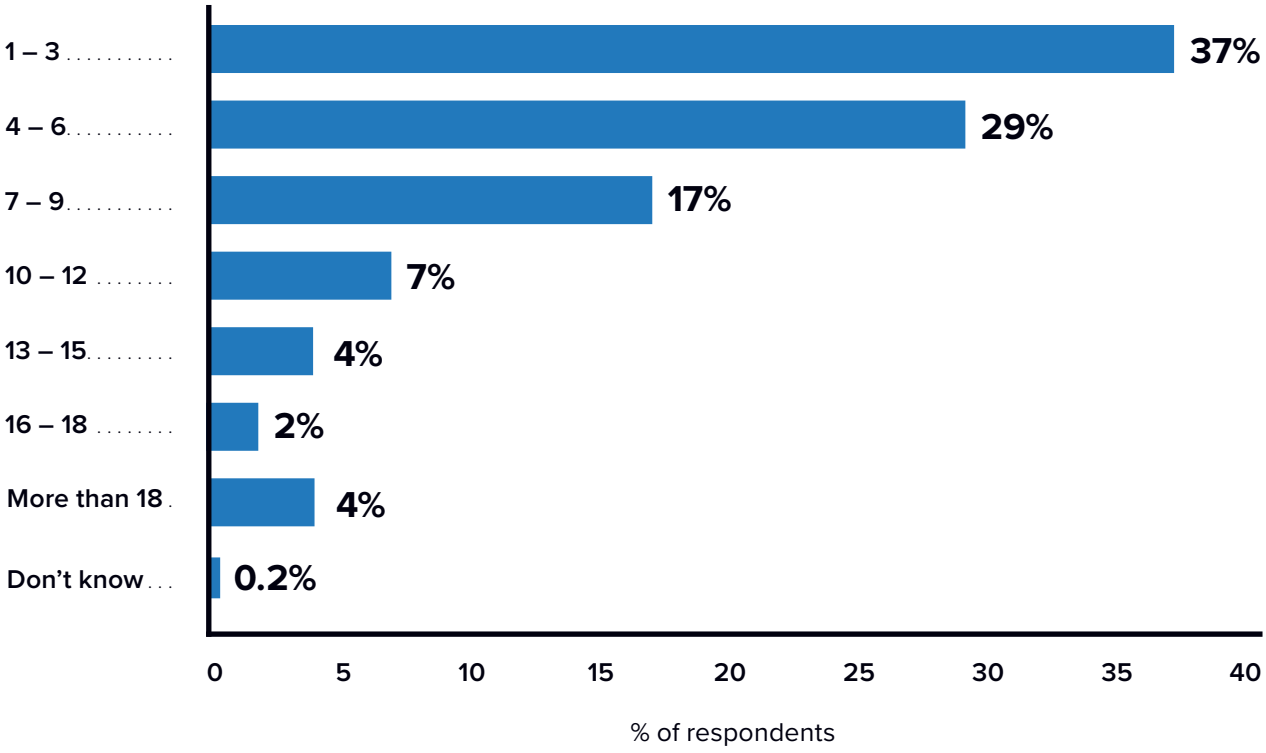

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Many organizations lack a structured framework for identifying relevant use cases and assessing their business impact, and this can lead to missed opportunities and inefficiencies. When AI initiatives are launched without clear alignment to strategic goals or measurable outcomes, they risk stalling in the proof-of-concept phase. Poorly chosen use cases can lead to wasted investment, delayed time to value, and increased compliance risk, especially as generative and agentic AI models introduce new governance complexities. For example, deploying agentic AI without robust oversight can result in autonomous actions that violate internal policies or regulatory standards, exposing enterprises to reputational and legal risks.

Realizing meaningful business impact from AI requires more than just adoption. It demands alignment with strategic objectives, robust infrastructure to support advanced workloads, and rigorous attention to governance and security. These pillars are essential to unlocking the full potential of AI initiatives.

Survey data reveals that only 21.7% of organizations conduct a full return-on investment (ROI) analysis, while 22.2% ensure alignment with business objectives

(see Figure 3). Others rely on rudimentary assessments or assumptions of value, which can lead to underperformance and missed opportunities. Without a disciplined evaluation process, enterprises may invest in models that fail to scale, lack operational relevance, or generate inconsistent outcomes.

FIGURE 3:

Limited Strategic Evaluation of AI Initiatives

Q. How does your organization primarily assess the business value of AI use cases that have been proposed?

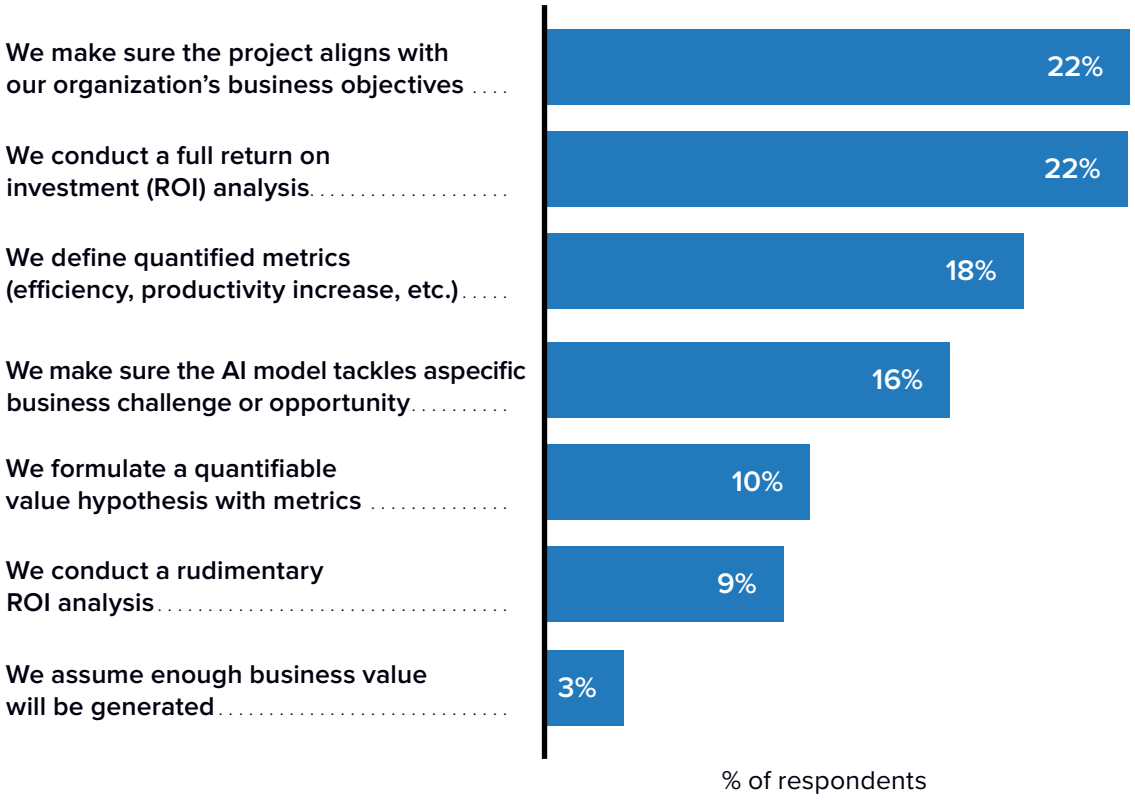

n = 410
Source: IDC's *AMD AI Survey*, July 2025

To avoid these pitfalls, organizations must adopt a strategic approach to AI use case selection that combines quantifiable value hypotheses, alignment with KPIs, and robust infrastructure planning. This includes evaluating compute, storage, and networking requirements for both training and inferencing, as well as ensuring that deployment strategies support scalability across datacenter, cloud, and edge environments. The most mature business value assessments result in ensuring that AI initiatives are not only aligned with strategic priorities but also deliver tangible, measurable outcomes that justify continued investment and support enterprisewide scalability.

However, even the most well-defined use cases can falter without a strong foundation of high-quality, well-governed data. As enterprises scale AI initiatives,

particularly generative AI and agentic AI, the focus must shift toward building resilient data pipelines, enforcing governance policies, and securing sensitive information. These elements form the backbone of responsible AI development and are critical to sustaining performance, trust, and compliance across the AI life cycle.

Data Foundations, Security and Governance for Generative AI and Agentic AI

Data serves as the input for any AI project. Ensuring high-quality, diverse, and well-governed data sets is essential for effective model training and fine-tuning. However, enterprises often face challenges such as fragmented data environments, inconsistent data quality, and evolving regulatory requirements. Overcoming these barriers requires robust data management practices, cross-functional collaboration, and infrastructure capable of supporting scalable, secure data operations. The rise of generative and agentic AI intensifies these demands, requiring substantial compute power, curated data, and robust safeguards.

Survey findings reveal that enterprises are adopting a multilayered approach to data governance in AI training and inferencing environments (see Figure 4).

FIGURE 4:

Approach to Data Governance

Q. How does your organization ensure data governance for the data used in AI training/fine-tuning?

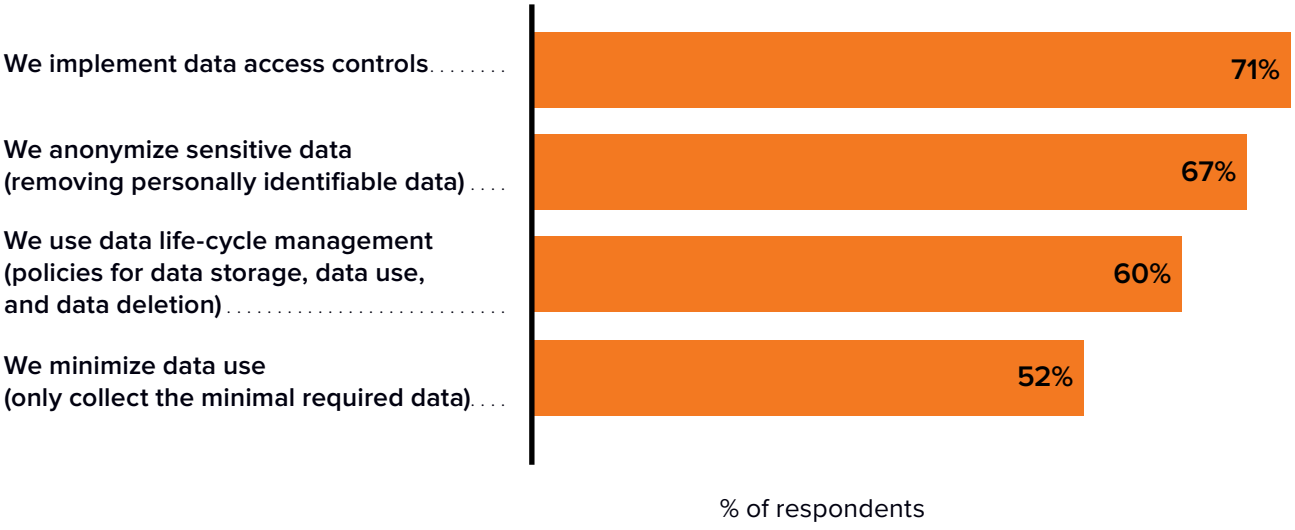

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Common practices include implementing data access controls to restrict sensitive data to authorized users; anonymizing personally identifiable information to mitigate privacy risks; applying data life-cycle management policies for storage, usage, and deletion; and minimizing data collection to only what is necessary for model development. These measures reflect a growing emphasis on responsible AI development and the need to align data governance with regulatory compliance and enterprise risk management.

In addition to this, AI models are among the most valuable digital assets enterprises can develop, requiring robust protection throughout their life cycle. Security begins at the development stage, where safeguarding training data, proprietary architectures, and access controls is essential to prevent unauthorized use and ensure integrity (see Figure 5).

FIGURE 5:
Security Testing of AI Models

Q. How does your organization test an AI model for security?

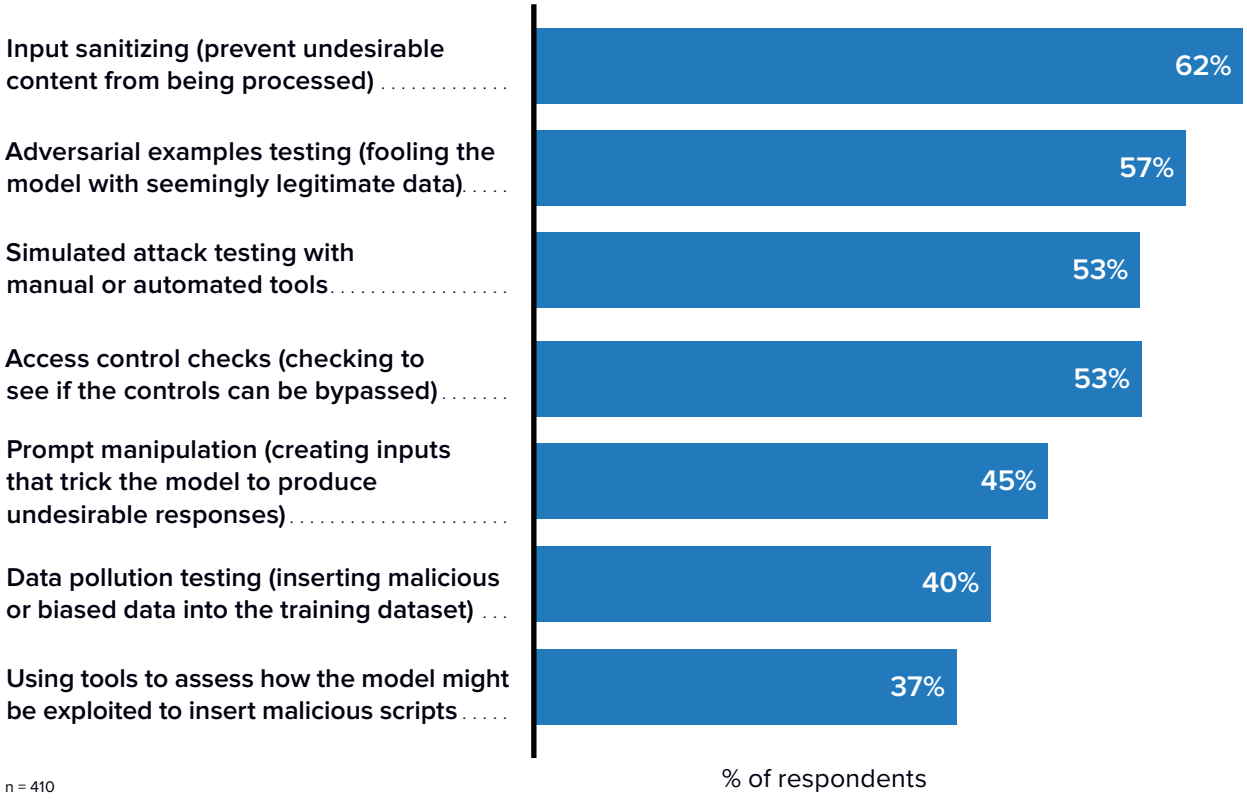

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Enterprises are increasingly embracing diverse methods to test their models against real-world vulnerabilities. These include simulating adversarial attacks to expose weaknesses in model behavior, testing for data pollution to identify risks from biased or malicious inputs, and manipulating prompts to evaluate how models respond to edge case scenarios. Input sanitization is widely practiced to

prevent undesirable content from being processed, while access control checks and simulated attacks help uncover potential bypasses and exploitation paths. Some organizations also deploy specialized tools to assess how models might be leveraged to inject malicious scripts or compromise downstream systems.

Together, these practices reflect a growing awareness that AI security must be embedded throughout the development and deployment pipeline — not only to protect intellectual property but also to ensure safe, compliant, and trustworthy AI operations.

Security, however, is only one dimension of enterprise AI readiness. As organizations scale AI initiatives, they must also confront the economic realities of deployment. The next section outlines the need for a comprehensive cost evaluation framework that enables informed decision-making by accounting for infrastructure, licensing, labor, and long-term scalability.

Need for a Comprehensive Cost Evaluation

AI initiatives are often derailed not by technical limitations but by a lack of clarity around the true costs involved. From initial development to long-term deployment, AI systems demand significant investment across multiple dimensions. These costs are not limited to compute resources or model licensing. They span infrastructure, labor, compliance, and ongoing operations. A robust cost assessment framework is essential to evaluate return on investment, ensure long-term feasibility, and guide strategic decision-making. About two-thirds of organizations surveyed report that their organization's GenAI initiatives begin with a comprehensive cost assessment to ensure the project is financially viable, covering infrastructure, licensing, labor, and scalability before moving beyond proof of concept.

Cost Dimensions to Consider

Strategic evaluation

- ROI analysis to measure impact and justify investment
- Cost-benefit analysis to assess long-term viability

Licensing and vendor ecosystem

- Licensing fees for large language models (LLMs) and proprietary tools
- Costs associated with third-party vendors, platforms, and consulting services

- **Infrastructure and hosting**
  - Cloud compute and storage pricing (e.g., GPU instances and data transfer)
  - On-premises hardware, virtual machines, and datacenter operations
  - Scalability of compute and storage to support growing workloads
- **Model development and operations**
  - Data acquisition, cleaning, and labeling for training
  - Expenses for model training, inference, fine-tuning, and red teaming
  - Software development, integration, and deployment pipelines
- **Labor and support**
  - Salaries and overhead for AI engineers, data scientists, and DevOps teams
  - Ongoing support, monitoring, and maintenance of AI systems
- **Compliance and risk management**
  - Security and privacy safeguards to meet regulatory standards
  - Governance frameworks and risk assessments for responsible AI use

Understanding and planning for these multifaceted financial factors is critical to ensuring the long-term viability and success of AI projects. Beyond hardware, software, and infrastructure, organizations must also account for licensing fees, vendor contracts, and the scalability of compute and storage resources. The costs of model development, including data preparation, training, inference, and ongoing integration, can be substantial. Labor expenses for skilled personnel, along with continuous support and maintenance, further add to the financial burden. Finally, compliance with security, privacy, and regulatory standards introduces additional layers of cost and risk.

A comprehensive cost evaluation that includes ROI analysis and long-term feasibility is essential for making informed, strategic decisions in AI adoption. As per the survey, among the factors considered, organizations place the greatest emphasis on costs related to specialized infrastructure for both AI training (60.7%) and inferencing (54.5%) (see Figure 6).

FIGURE 6:

Infrastructure Cost Assessment for AI Training and Inferencing

Q. Which of the following cost factors are the three most important to include in a full cost assessment before taking an AI model into production?

Factors considered for a training cost assessment (n = 275)

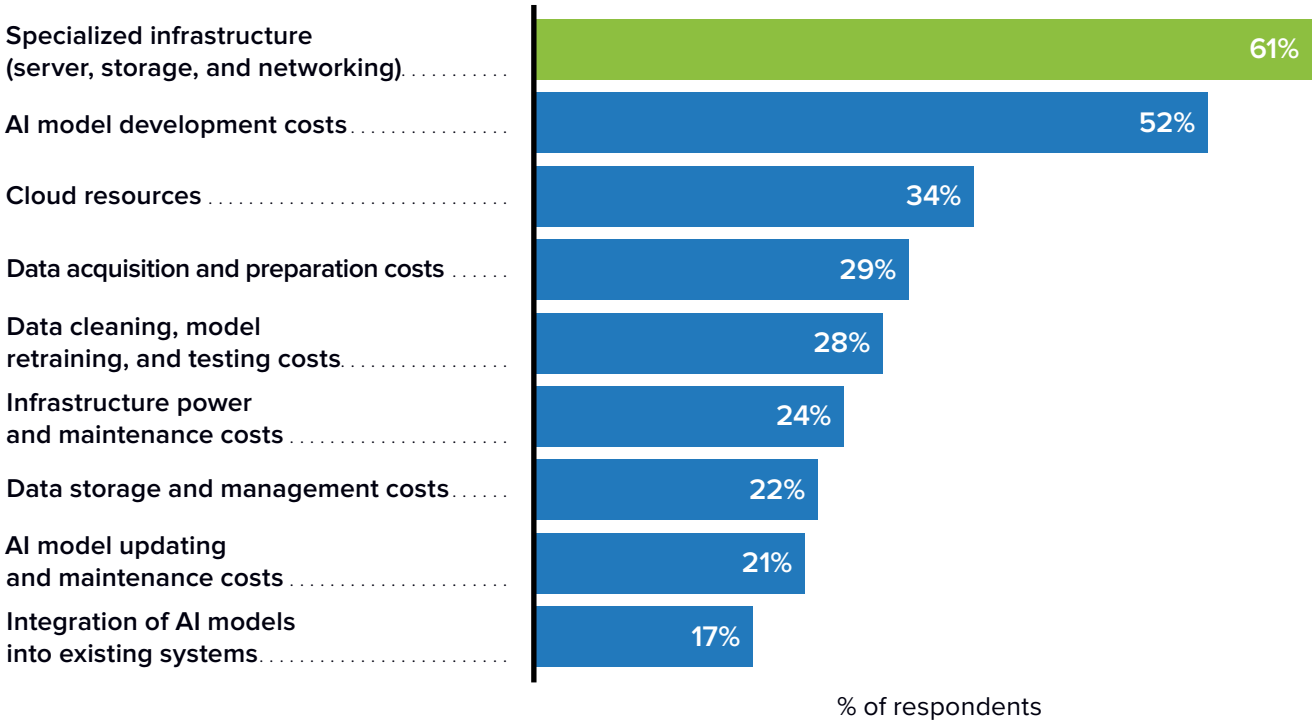

Factors considered for an inferencing cost assessment (n = 275)

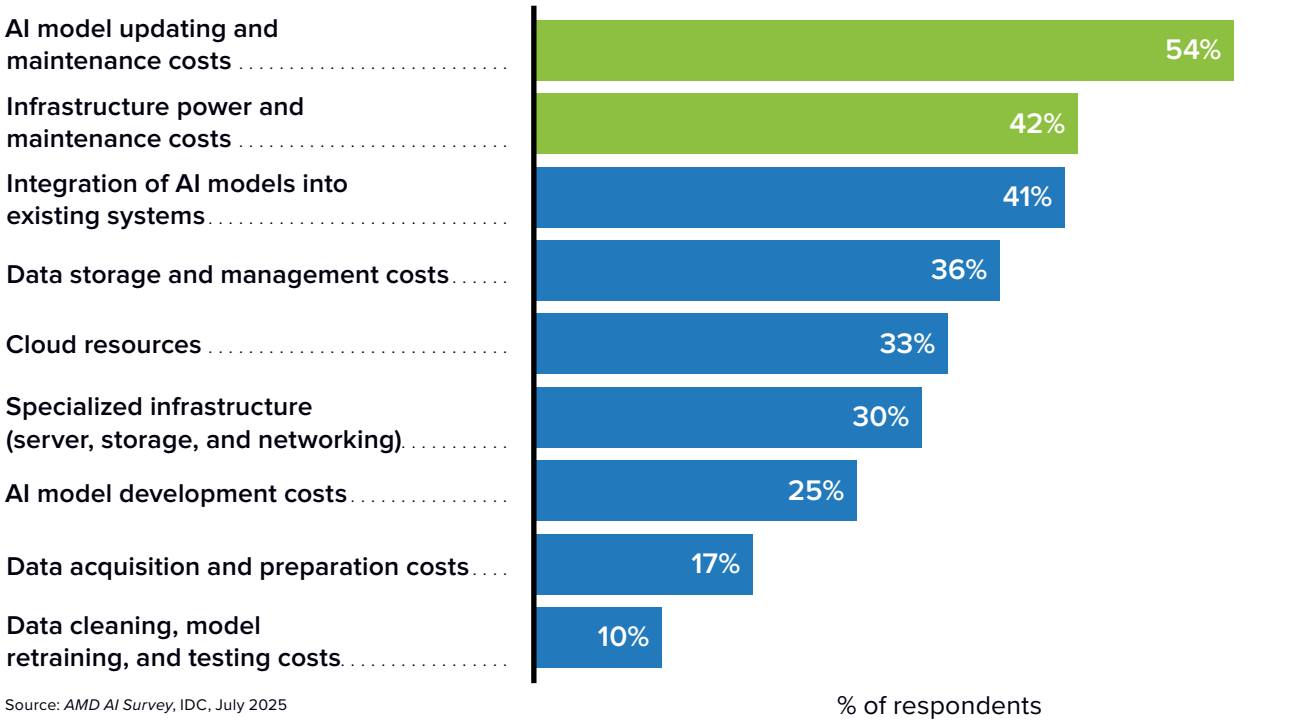

Source: *AMD AI Survey*, IDC, July 2025

Given the financial weight of AI initiatives, infrastructure decisions become pivotal. Survey data highlights that organizations are prioritizing factors such as ongoing model maintenance, infrastructure power and upkeep, and seamless integration with existing systems when assessing the full cost of taking models into production. Hardware choices also play a critical role, with GPUs, high-speed interconnects, and robust cluster networks emerging as key components for scalable, low-latency inferencing.

The type of infrastructure selected whether built in-house, procured as turnkey systems, or integrated via third-party providers directly influences cost efficiency, scalability, and time to value. As organizations evaluate their options, the role of performance-intensive computing infrastructure becomes increasingly critical. The sections that follow explores how enterprises are approaching infrastructure to meet the evolving demands of AI at scale.

Infrastructure Choices

AI Training Infrastructure

Survey findings underscore the importance of *specialized infrastructure* in achieving successful AI model training that is defined by accuracy, timeliness, and budget adherence. GPUs and GPU memory emerged as the most critical components, reflecting the computational intensity of modern AI workloads. Interconnects between GPUs and CPUs, high-bandwidth memory, and cluster-level networking and management software were also identified as essential for scaling performance and minimizing bottlenecks. These results highlight the need for a balanced infrastructure stack that supports parallel processing, fast data movement, and efficient resource orchestration (see Figure 7).

FIGURE 7:

Infrastructure Enablers for AI Model Training

Q. Which of the following are the five most critical infrastructure components for successful training of the AI model?

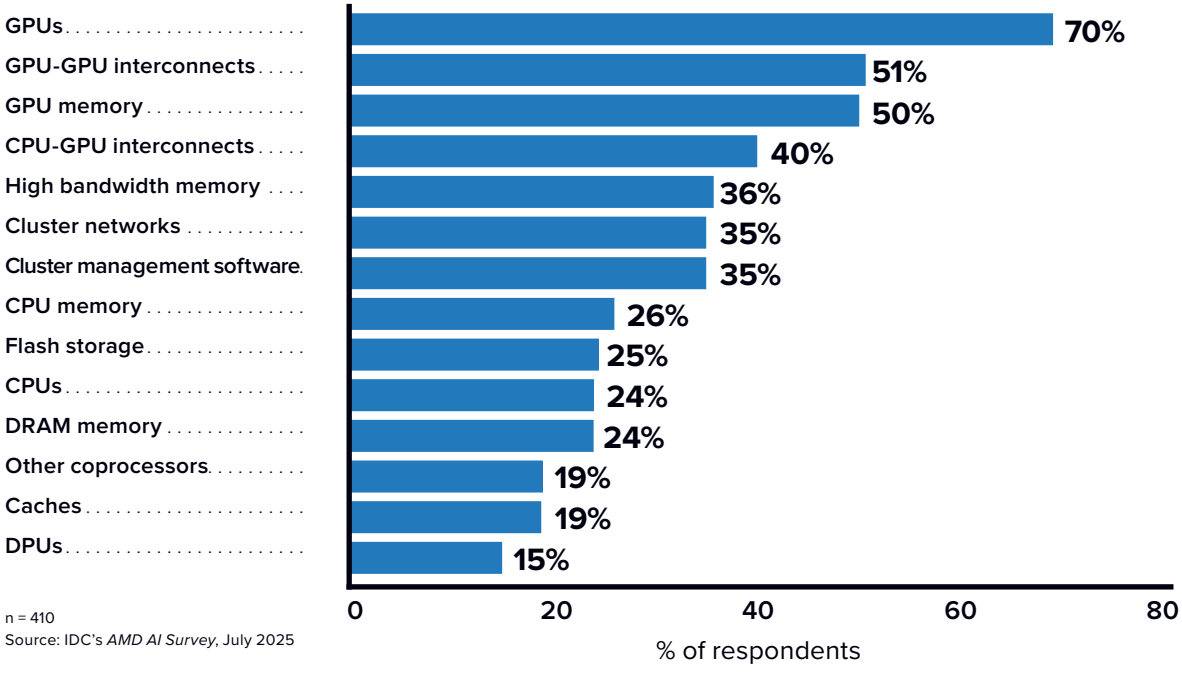

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Given these priorities, *infrastructure acquisition strategies* play a pivotal role in enabling AI success. Organizations are taking varied approaches; a significant 25% relied on value-added resellers (VARs) or systems integrators to design and install the infrastructure, while the rest were evenly split between building in-house solutions and purchasing fully configured turnkey systems from vendors. Infrastructure choices are also closely tied to *deployment preferences* (see Figure 8).

FIGURE 8:
Primary locations for AI Model Development

Q. In which locations does your organization develop (train or fine-tune) its AI models?

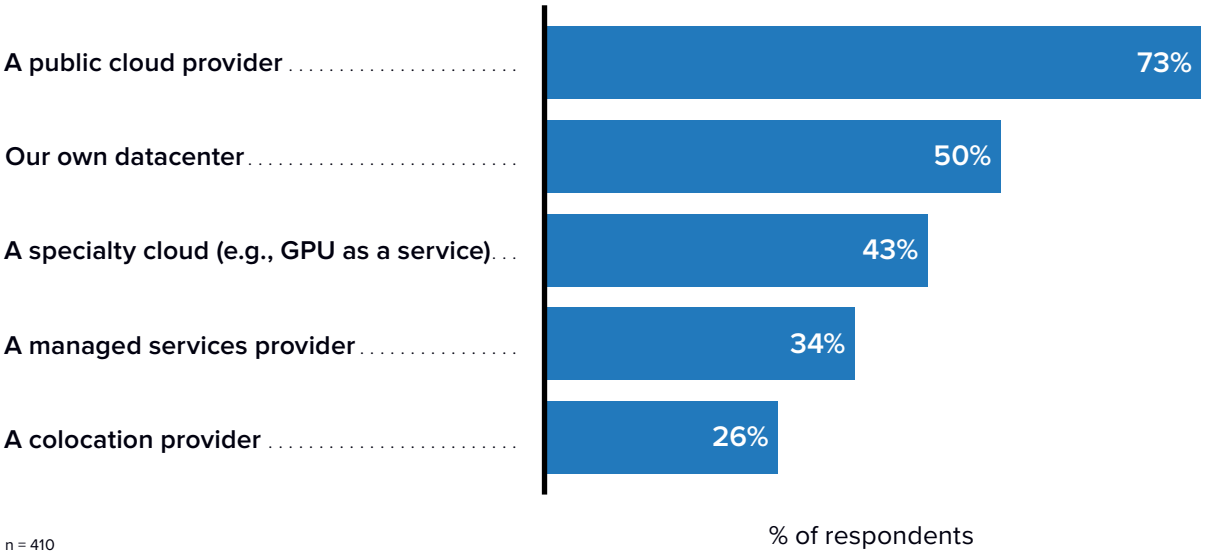

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Survey data indicates that though public cloud remains a popular choice for AI model development, a significant proportion of organizations favor on-premises or colocation environments for AI model development to maintain architectural flexibility, meet compliance requirements, and be able to tailor compute configurations to specific workloads.

Hosting AI workloads locally offers several advantages: it allows for greater control over data access and storage, ensures data remains within specific geographic regions to support data sovereignty, and enables customization of hardware and software configurations to suit early-stage experimentation. It also facilitates seamless integration with existing enterprise systems, supports model tailoring for specific business needs, and reduces latency by keeping data local. In addition, organizations benefit from enhanced security through encryption, access controls, and compliance certifications while avoiding ongoing cloud service contracts and potentially lowering total cost of ownership through hardware amortization.

AI Inferencing Infrastructure

Unlike training, which is often centralized, inferencing frequently occurs at the edge to support real-time applications, requiring infrastructure that balances latency, throughput, and cost efficiency. As AI models move into production, inferencing performance becomes a key determinant of user experience and operational efficiency. Survey data highlights that GPUs, GPU memory, and high-speed interconnects are considered the most critical components for successful inferencing (see Figure 9).

FIGURE 9:

Infrastructure Enablers for AI Model Inferencing

Q. Which of the following are the five most critical infrastructure components for successful inferencing of the AI model?

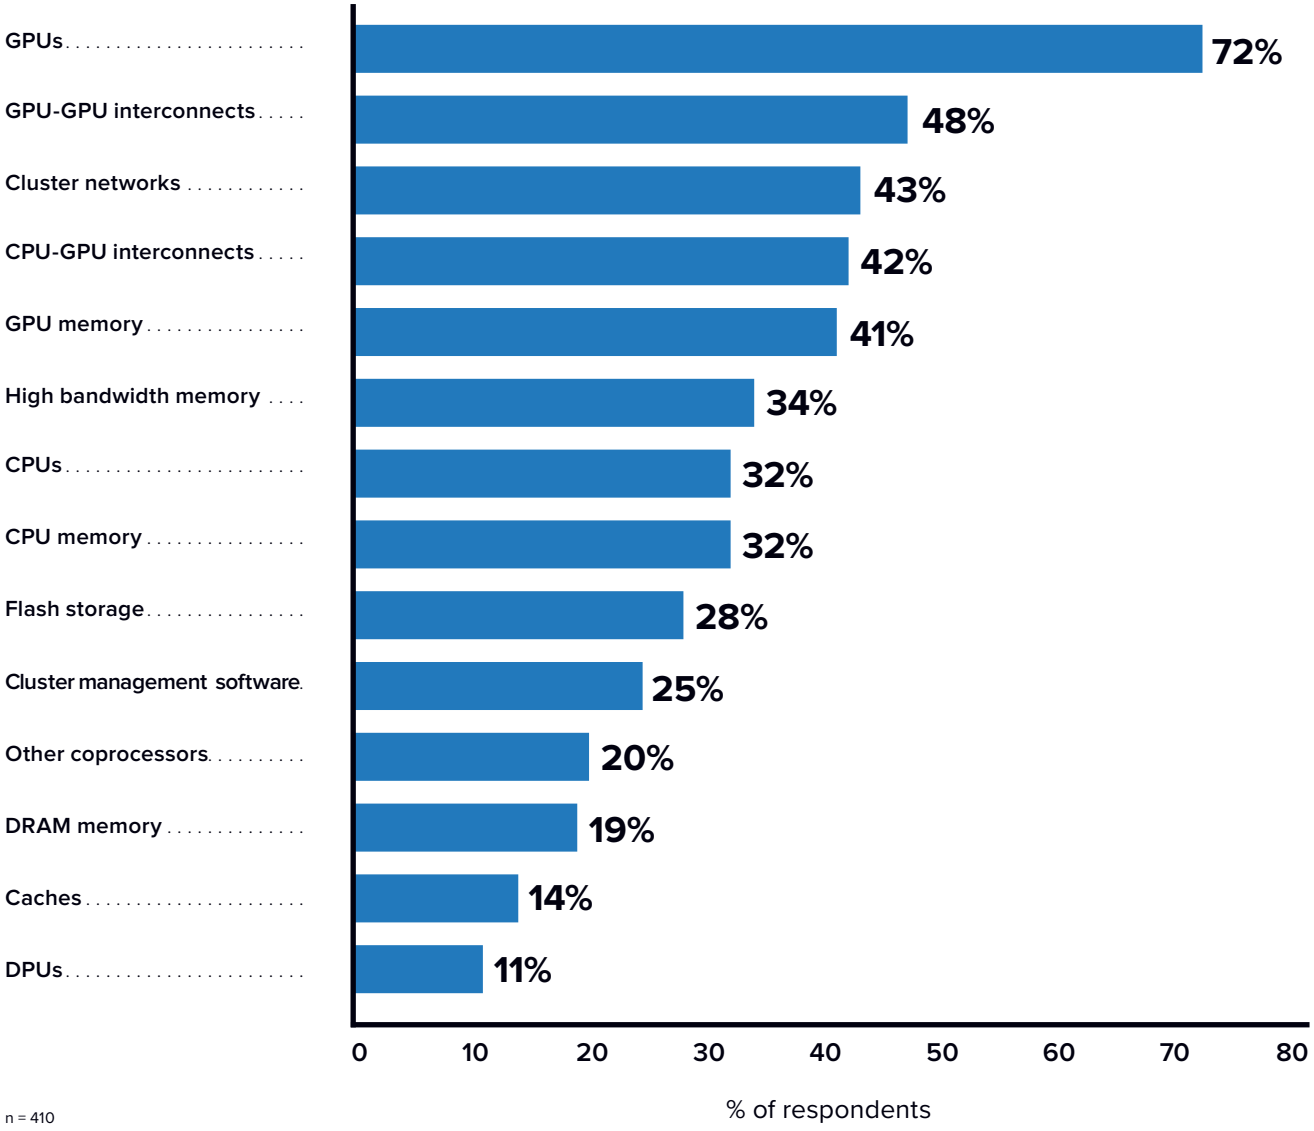

n = 410
Source: IDC's *AMD AI Survey*, July 2025

These findings highlight the need for robust, scalable systems that can deliver low-latency insights while maintaining operational and financial efficiency, and therefore, a varied choice of high-performance components is required. With respect to *deployment preferences*, inferencing is distributed across public clouds, colocation facilities, and edge environments to support latency-sensitive and real-time applications. This distributed model allows organizations to align infrastructure choices with workload requirements, regulatory constraints, and operational priorities.

As inferencing workloads become critical to enterprise operations, organizations are making strategic decisions about where and how to deploy AI models. According to survey data, 63.9% of organizations deploy AI inference workloads in public cloud environments, while 50.7% continue to leverage their own datacenters, reflecting a balanced approach to infrastructure planning (see Figure 10).

FIGURE 10:
Primary Locations for AI Model Deployment

Q. In which locations does your organization deploy (inference on) its AI models?

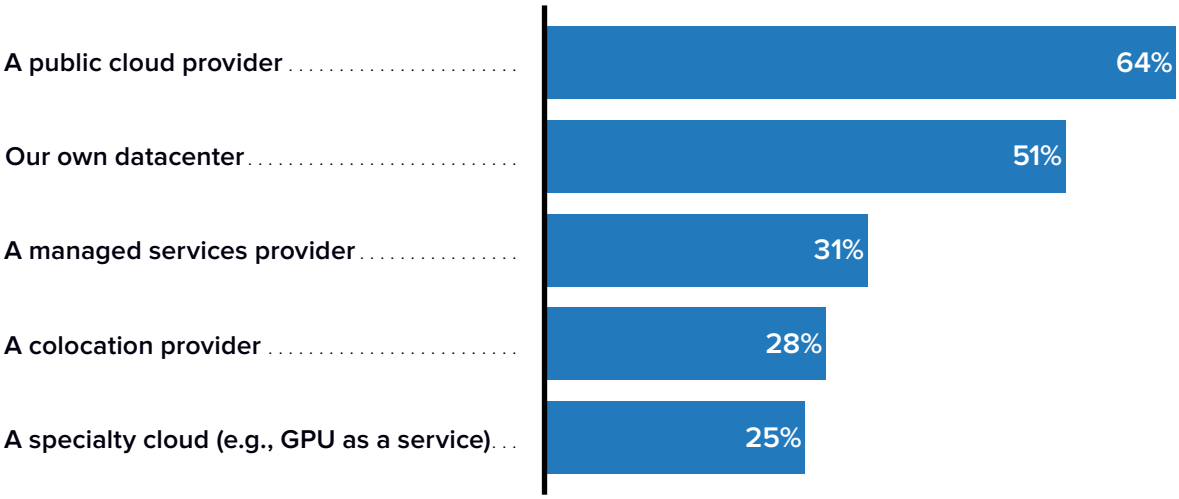

n = 410
Source: IDC's *AMD AI Survey*, July 2025

On-premises and colocation environments continue to play a vital role, particularly for use cases involving sensitive data, regulatory compliance, and the need for hardware and software customization. These environments offer deeper integration with existing systems and enable precise control over data access, storage, and model behavior — making them well suited for applications that demand privacy, precision, and long-term cost efficiency.

At the same time, public cloud platforms, managed services, and specialty AI clouds are gaining traction for their agility and operational simplicity. Elastic scaling, access to AI-optimized instances, and reduced infrastructure overhead

make these environments ideal for dynamic workloads and rapid experimentation. Organizations value the ability to deploy inferencing workloads quickly without up-front capital investment while leveraging specialized services to accelerate development cycles.

Infrastructure acquisition strategies reflect this. Some enterprises build in-house solutions to maintain control and customization, while others opt for turnkey systems or partner with systems integrators to streamline deployment. Regardless of the approach, reliability, availability, and serviceability (RAS) features are becoming standard requirements. Nearly 38% of organizations have installed RAS capabilities that match or exceed those in their broader IT landscape, underscoring the mission-critical nature of inferencing workloads. As AI becomes embedded in core business operations, infrastructure decisions must balance performance, resilience, and long-term sustainability.

Inferencing Infrastructure Preferences and Strategic Implications

As AI inferencing becomes embedded across enterprise environments, infrastructure strategies must evolve to support deployment across a diverse range of systems. Whether models are hosted in datacenters, cloud platforms, or edge devices, the need for interoperability and standardization is crucial. Success in inferencing is no longer defined solely by compute power. This is because it depends on the ability to deploy consistently across heterogeneous environments without compromising latency, accuracy, or security.

The research data reinforces this trend, with respondents identifying datacenter servers, edge systems, and personal computing devices as key targets for inferencing workloads (see Figure 11).

Figure 11 illustrates that the majority of inferencing is anticipated to occur on datacenter/cloud servers, followed by edge servers. The data reveals a clear organizational preference for centralized and edge-based infrastructure when deploying AI inferencing workloads. With the majority of inferencing expected to occur on datacenter or cloud servers, enterprises are prioritizing environments that offer high performance, scalability, and control. The sizable share of inferencing on edge servers underscores the growing need for low-latency, real-time processing closer to data sources, particularly relevant in sectors such as manufacturing and logistics. In contrast, the relatively low share of inferencing on end-user devices such as laptops, mobile phones, and robotics indicates that most organizations are not yet relying on distributed or consumer-grade hardware for production AI. This distribution reflects a strategic focus on reliability, cost-efficiency, and infrastructure consistency in AI deployment planning.

FIGURE 11:

Device-Level Distribution of AI Model Inferencing: Strategic Deployment Across Compute Environments

Q. When considering inferencing on your organization's trained AI model, to what extent do you expect this inferencing will take place on the following devices?

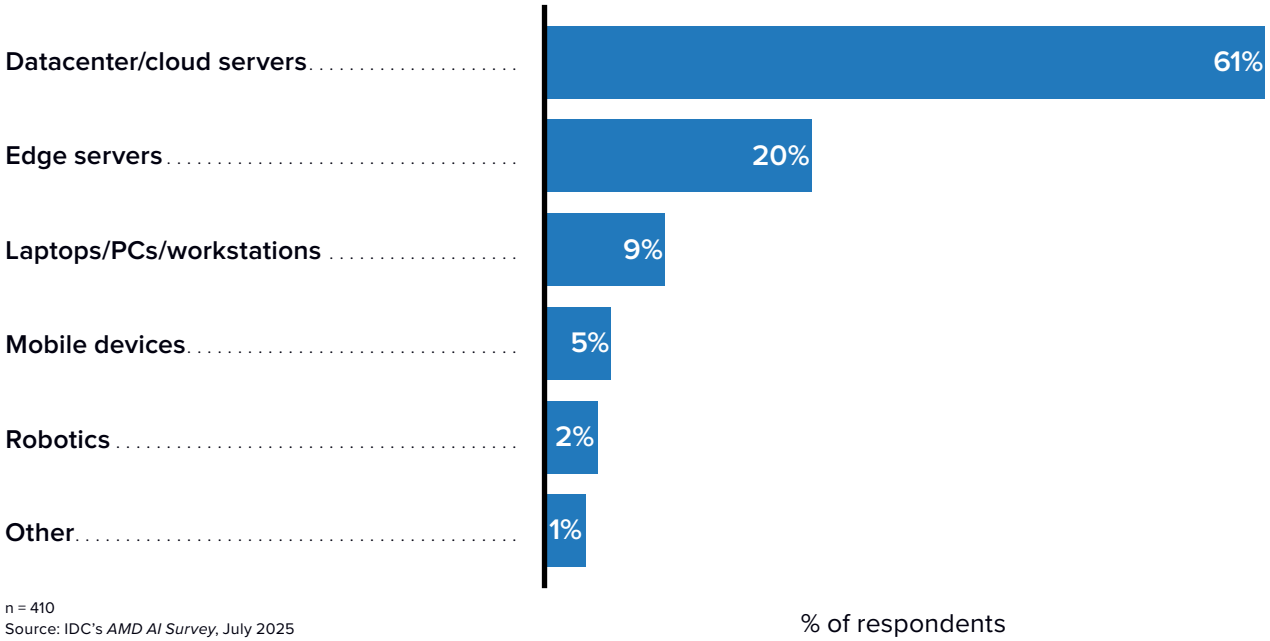

n = 410
Source: IDC's *AMD AI Survey*, July 2025

Supporting this trend, the survey also explored how important it is for organizations that servers, laptops/PCs/workstations, and edge devices processing the same AI model operate on consistent hardware and software platforms (e.g., same CPUs, GPUs, and DPUs). A strong majority of respondents (77%) rated this as either very important or extremely important, indicating that hardware and software uniformity is seen as a key enabler of performance predictability, model portability, and operational efficiency across diverse deployment environments.

This reflects a growing demand for standardized infrastructure that simplifies deployment, reduces operational friction, and ensures consistent performance throughout the AI life cycle. In parallel with these infrastructure preferences, organizations are also refining their approaches to planning and scaling inferencing workloads across environments.

Infrastructure Planning for Scalable and Predictable AI Deployment

As inferencing workloads expand across diverse environments from datacenter servers to edge devices, the need for consistent performance and seamless deployment becomes paramount. Organizations are increasingly prioritizing

standardized infrastructure to reduce latency, simplify integration, and mitigate security risks across heterogeneous systems. This emphasis on consistency sets the stage for a more strategic approach to infrastructure planning, one that accounts not only for technical requirements but also for operational realities.

To meet anticipated query volumes and data sizes, enterprises are adopting a layered planning methodology. It begins with forecasting and scoping that includes analyzing user types, estimating query loads, and accounting for business growth and seasonal fluctuations. Buffer capacity is built in to absorb demand spikes, while predictable workloads are managed through reserved resources and traditional capacity planning. For dynamic workloads, cloud elasticity, autoscaling, and modular APIs enable responsive scaling and concurrency. Monitoring tools and feedback loops refine forecasts and guide infrastructure adjustments, ensuring alignment with actual demand. Cost considerations are tightly integrated into this process, with decisions shaped by budget constraints, ROI expectations, and long-term strategic goals. Whether deployed in cloud or on-premises environments, successful inferencing infrastructure requires a forward-looking strategy that balances agility, performance, and fiscal discipline.

Future Outlook

Outlook: Preparing for 2026–2027

Enterprises will need to prepare for the next wave of AI innovation. Agentic AI — autonomous decision-making and task execution — will demand even greater infrastructure agility and governance sophistication. Energy efficiency and sustainability will become central to infrastructure planning, as organizations seek to reduce carbon footprints and optimize power consumption across AI workloads. By aligning infrastructure strategy with business value, governance standards, and future-readiness, enterprises can unlock the full potential of AI — transforming operations, accelerating innovation, and driving competitive advantage.

In addition, semiconductor supply chain resilience will be a strategic priority, with enterprises and partners working together to ensure availability, scalability, and innovation continuity in a rapidly evolving global landscape.

Conclusion

As AI has become foundational to enterprise strategy, infrastructure decisions must be guided by business outcomes, governance, and operational efficiency. Based on the findings and analysis presented in this white paper, three key imperatives emerge for organizations seeking AI successfully.

Align AI with Business Value

Successful AI initiatives begin with clearly defined use cases that are tied to measurable business outcomes. Enterprises are launching multiple AI use cases simultaneously, with most deploying from 1 to 6 at a time. However, only a minority conduct full ROI analyses or define quantifiable metrics upfront. To maximize impact, enterprises must prioritize ROI-driven planning, ensure strategic relevance, and continuously evaluate performance against evolving KPIs.

Build Governance Foundations Early

As generative and agentic AI models gain traction, governance becomes a critical success factor. Survey data shows widespread adoption of security testing methods, including adversarial testing, input sanitization, and access control checks, which highlights the need for robust safeguards. Organizations must embed governance into every stage of the AI life cycle to mitigate risk and ensure compliance.

Adopt Flexible Infrastructure Models

Infrastructure choices vary widely, with enterprises split between building in-house systems, purchasing turnkey solutions, and working with systems integrators. Inferencing workloads are deployed across datacenters, cloud platforms, and edge devices, with a strong preference for unified hardware and software environments to reduce latency and complexity. Flexibility in deployment and acquisition models is essential to meet performance, scalability, and cost-efficiency goals.

Message from the Sponsor

AMD

AMD is the trusted, open, and flexible partner that accelerates enterprise AI transformation. Our broad portfolio of compute engines provides choice, flexibility and proven performance enabling enterprises to optimize compute for every workload including training and inference. Our commitment to an open ecosystem empowers industry collaboration, providing enterprises with faster access to innovations. And our full-stack solutions include advanced orchestration capabilities that help enterprises easily deploy and scale AI initiatives. AMD is the customer-first AI leader that meets the evolving needs of enterprises. Together, we transform AI potential into measurable business outcomes and sustainable competitive advantage.

Learn more

About the IDC Analysts

Peter Rutten

Research Vice-President, Performance Intensive Computing, Worldwide Infrastructure Research, IDC

Peter Rutten is Research Vice-President within IDC's worldwide infrastructure research organization and global research lead for the performance-intensive computing (PIC) practice. IDC's PIC coverage includes research on High-Performance Computing (HPC), Artificial Intelligence (AI) and Generative AI (GenAI), Big Data and Analytics (BDA) and Quantum Computing (QC) infrastructure stacks, deployments, solutions, workloads and use cases. It includes coverage of classical and hybrid quantum-classical supercomputing, and institutional and mainstream HPC.

More about Peter Rutten

Madhumitha Sathish

Research Manager, High Performance Computing, IDC

Madhumitha Sathish is Research Manager within IDC's worldwide infrastructure research organization and part of the performance intensive computing (PIC) practice. She leads IDC's artificial intelligence (AI) infrastructure research and plays a supporting role in IDC's research on high-performance computing (HPC) infrastructure stacks and deployments. In her role, Ms. Sathish delivers syndicated and custom qualitative and quantitative insights, including market size and forecast studies, and custom market models.

More about Madhumitha Sathish

Heather West, PhD

Research Manager, Performance Intensive Computing, Worldwide Infrastructure Research, IDC

Heather West, PhD, is Research Manager within IDC's worldwide infrastructure research organization and part of the performance intensive computing (PIC) practice. She leads IDC's quantum, analog and neuromorphic computing research and plays a supporting role in IDC's research on artificial Intelligence (AI) and high-performance computing (HPC) infrastructure stacks and deployments. Dr. West is deeply engaged with her clients on their solutions and services, as well as on their business and technology strategies.

More about Heather West, Phd

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.)

IDC Custom Solutions

This publication was produced by IDC Custom Solutions. As a premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets, IDC's Custom Solutions group helps clients plan, market, sell, and succeed in the global marketplace. We create actionable market intelligence and influential content marketing programs that yield measurable results.

IDC Research, Inc.
140 Kendrick Street, Building B, Needham, MA 02494, USA
T +1 508 872 8200

@idc | @idc | idc.com