




451 Research Discovery Report

May 2025

Best practices for delivering AI at scale

Lessons from industry leaders

S&P Global
Market Intelligence

©Copyright 2025 S&P Global. All Rights Reserved.

in partnership with
verizon
business

Table of contents

Introduction	3
Methodology	3
Figure 1: Demographic breakdown of participants	3
Executive summary	4
Insight 1: Prioritize scalability and recognize that the cloud is not a solve-all	6
Figure 2: Vast majority of organizations leveraging the cloud for AI workloads	7
Insight 2: Data strategy and security must be addressed at the outset of AI projects	9
Figure 3: Most organizations say infrastructure could better support data pipeline management	10
Insight 3: Effective project prioritization is key to controlling costs	11
Figure 4: Required infrastructure upgrades for AI span security, storage, networking and compute	12
Insight 4: AI initiatives do not operate in isolation; organizations must emphasize both integration and collaboration	14
Figure 5: Wide array of stakeholders involved in infrastructure decisions to support AI/ML workloads	15
Implications: Laying the foundation for AI at scale	16
AI strategy	16
Networking	16
Security	16
About the author	17

Introduction

Organizations are actively investing in AI as the technology rapidly advances, but not all are reaping the benefits. According to [451 Research’s Voice of the Enterprise: AI & Machine Learning, Use Cases 2025 survey](#), 84% of organizations reported active investment in generative AI. However, 46% of these respondents did not see significant positive impacts from those investments against any organizational objectives assessed in the survey.

All are working to build the scaffolding that they hope will bring them AI value, but a clear divide is emerging between organizations effectively leveraging AI and those that are not. This new report on best practices delves into the practices of higher-performing organizations — those that have successfully implemented AI projects at scale. It explores how they handled the challenges of managing both the availability and security of the data that drives their applications, and the infrastructure and networks that enable them to flourish. We’ll explore strategies and recommendations from these leaders.

This report, the final installment in a three-part series, explores best practices. It highlights the most significant insights from AI leaders and their advice on how to avoid obstacles they encountered in their journeys. The first report in the series discussed networking strategies, and the second focused on security practices. This report takes a macro view to showcase overall project planning and AI strategy, to understand how these pieces fit together.

Methodology

The insights presented in this report are derived from comprehensive interviews and peer group discussions with senior executives responsible for AI initiatives at organizations that have successfully implemented AI at scale. These interactions included 15 in-depth interviews with AI and IT infrastructure decision-makers, as well as an executive discussion board with 30 participants. All participants in the study were specifically involved in managing and implementing their organizations’ IT infrastructure for AI workloads. Participants were based in the US, UK, Singapore, Australia, Sweden, Germany, Denmark and Japan. This report was developed by S&P Global Market Intelligence and commissioned by Verizon.

Figure 1: Demographic breakdown of participants

Executive discussion board				One-to-one interviews			
Region	#	Industry	#	Region	#	Industry	#
North America	10	Healthcare/life sciences	8	North America	5	Retail/wholesale	3
Europe	10	Financial services	8	Europe	5	Healthcare/life sciences	4
Asia-Pacific	10	Manufacturing	6	Asia-Pacific	5	Financial services	5
		Retail/wholesale	3			Manufacturing	2
		Utilities	1			Utilities	1
		Other	4				

Source: S&P Global Market Intelligence 451 Research AI at Scale study commissioned by Verizon.

Executive summary

Companies are eager to engage with AI, and many executives recognize its potential to address a wide array of objectives. They are building infrastructure to enable AI innovation while dealing with challenges related to data security, data management and interconnection in the context of increasingly hybrid IT environments. Study participants reflected significant differences in the specific value they seek from AI and the areas where they aim to apply the technology, with a frequent emphasis on industry-specific processes or goals. Nonetheless, the study identified four common areas where AI is creating opportunity: enhancing efficiency and cutting costs, improving customer experience and engagement, enabling data-driven decision-making, and addressing compliance and security needs.

“We are looking to deploy ML and AI solutions that will help to automate redundant workflows, increase operational efficiency, reduce total cost, [improve] quality, improve product design, help us to reduce the time it takes to deliver new product innovations.”

IT director

Auto manufacturing, >50,000 employees, US

“In the clinical space, a main driver is clinician burnout. We are trying to leverage AI tools to assist with clinical workflows to lessen the pyjama time for physicians as well as APPs [advanced practice providers].”

Senior infrastructure manager

Healthcare, 20,000-50,000 employees, US

“We are looking at leveraging AI to improve our customer experience. Using AI chatbots to get answers to products, services, technical support, billing”

Senior director of IT infrastructure

Consumer goods, 10,000-20,000 employees, US

These objectives remain elusive as organizations struggle to move AI initiatives from proof-of-concept trials into production environments. This is especially true for data-intensive projects and those that serve a wide user base. The rapid rise of generative AI has exacerbated this issue: Organizations are eager to leverage GenAI advancements and are exploring a broad range of opportunities, but the technology’s nascence and a lack of organizational experience make these projects particularly difficult to deliver.

According to 451 Research’s Voice of the Enterprise: AI and Machine Learning, Use Cases 2025 study, the proportion of companies that discontinued most of their AI initiatives between proof of concept and production increased from 17% in the fourth quarter of 2023 to 42% in the fourth quarter of 2024. On average, surveyed organizations ended an estimated 46% of their AI projects during the transition from a successful proof-of-concept to production. The same trend was also reflected in this “AI at scale” study, with multiple participants identifying issues that derailed projects.

Study participants across our interviews and executive discussion board shared a range of architectural insights critical to the successful delivery of AI projects:

- As organizations move beyond experimenting with AI, infrastructure rapidly emerges as a bottleneck. Without forward planning and modernization, projects can fail or face unforeseen costs and delays.
- Compute is not the sole limiter to AI ambitions. Without an early focus on networking and security requirements, AI projects risk stalling or veering off course.
- AI strategies rarely remain confined to the public cloud; failure to account for multi-venue complexity in networking and security planning will prove short-sighted. Without adequate planning, moving data between environments can prove costly and expose organizations to risk.
- Latency, bandwidth and availability should guide network strategy. These factors are critical for data-intensive AI applications, which often represent the most impactful use cases. Network segmentation can play a key role in enabling greater performance and control.
- Securing AI workloads demands layered defenses. A robust approach includes zero-trust principles, strong encryption and enhanced monitoring to protect sensitive data and models.

These infrastructure considerations are foundational to how successful organizations run AI workloads in production. Study participants also surfaced key insights related to AI strategies that inform these infrastructure approaches:

- **Insight 1:** Prioritize scalability and recognize that the cloud is not a solve-all.
- **Insight 2:** Data strategy and security must be addressed at the outset of AI projects.
- **Insight 3:** Effective project prioritization is key to controlling costs.
- **Insight 4:** AI initiatives do not operate in isolation; organizations must emphasize both integration and collaboration.

We examine each of these insights in turn and explore how they frame the networking and security considerations that participants raised.

“The following are very important to build robust: 1. Scalability, 2. Cost management, 3. Collaboration, 4. Resilience, 5. Compliance and security, 6. Automation.”

IT director

Life sciences, >50,000 employees, US

“I would have a different architecture design, bring in security integration much earlier, better data pipeline planning.”

Senior director of IT

Healthcare/life sciences, 10,000-20,000 employees, US

“I would surely make sure my architecture foundation is built to scale, which keeps in mind the security protocol that we need to follow in order to make that happen.”

VP of infrastructure, ITops & security

Retail, >50,000 employees, US

Insight 1: Prioritize scalability and recognize that the cloud is not a solve-all

Executives and AI practitioners from leading organizations noted significant benefits from starting small, experimenting and building up. In many instances, participants reflected that project teams derived valuable insights from their first wave of initiatives, particularly regarding effective areas of AI application. One interview participant, a US-based vice president of infrastructure and security at a healthcare organization, cited the opportunity to identify low-hanging fruit — in this instance, using pattern-recognition technology to identify under-utilized or redundant assets — and using smaller projects and quick wins to accelerate AI strategy. A healthcare and life sciences participant from Japan spoke about the need to “use AI to help and do parts of tasks in the beginning, rather than trying for full solutions.”

“All these questions that you have around security, compliance, AI governance. Those are best learned and experienced by just doing it in smaller environments, that are in controllable environments.”

Director of IT for EMEA

Manufacturing, 10,000-20,000 employees, Germany

“Start small, make it safe and experiment.”

Director of digital health infrastructure

Healthcare, 5,001-10,000 employees, Australia

Starting small and experimenting does introduce the risk that organizations will focus excessively on the immediate needs of pilot initiatives or draw incorrect conclusions from small-scale samples. Some participants mentioned that, without proper perspective, experiments can be misleading.

“I believe the key point is to recognize that a PoC is not simply a scaled-down version of the production environment. The primary purpose of a PoC is to validate the technical feasibility, but it’s important to remember that the configuration of the production environment will differ. It is essential to plan with operations in mind and to clarify the non-functional requirements of the production environment (such as performance, availability and security) before starting the PoC.”

Director of technology

Professional services, 1,000-5,000 employees, Japan

Experienced organizations viewed flexibility to scale as a critical infrastructure enabler for transitioning from small beginnings to enterprise-level execution. Several study participants highlighted the importance of intentional forward planning to accommodate scalability. When respondents were asked what they would do differently if they could start over, planning for greater scalability emerged as a common theme.

“Planning, designing, strong and scalable architecture based on functions we want to perform.”

IT director

Manufacturing, >50,000 employees, US


“[I would] redo my architecture to ensure it has a future scalable solution in place.”

Head of infrastructure services

Banking, 20,000-50,000 employees, Australia

Study participants identified cloud infrastructure as an enabler for greater elasticity. While several participants suggested that the compute requirements of generative AI drove them to the cloud out of necessity, many viewed the cloud as a powerful tool for addressing the pressures of scaling and managing uncertainty. As Figure 2 highlights, the majority of participants said they are executing AI workloads in cloud environments.

Figure 2: Vast majority of organizations leveraging the cloud for AI workloads



Q. How would you characterize your organization's current approach to infrastructure development for your AI/ML workloads?

Base: Executive discussion board participants (n=30).

Source: S&P Global Market Intelligence 451 Research AI at Scale study commissioned by Verizon.

"We probably don't have enough chips to do that on-prem. So, it's just a matter of fact. We won't get them ... So, there will — and has to be — a cloud solution for that."

Director of IT for EMEA

Manufacturing, 10,000-20,000 employees, Germany

"If I were to start from scratch, I would start on cloud because I will avoid the risk of investing into an infrastructure where I'm not completely sure of how much computing power I would need."

VP of IT

Utilities, 10,000-20,000 employees, France

While the cloud is part of the solution to scalability, it is not the complete answer. For many use cases, hyperscale public cloud may be poorly suited due to cost and privacy concerns. Participants also noted that the cloud can present challenges to flexibility, with several respondents citing vendor lock-in.

"Ensure that every single component I use in my solution was cloud agnostic ... I would do that, obviously, because then I know that even if I have problems with my cloud provider, I could switch easily to another infrastructure [provider], even my own ... And just use mainstream tools ... But again, avoid vendor lock-in."

VP of IT

Utilities, 10,000-20,000 employees, France

To create greater capacity for growth, participants recommended increasing modularity in AI stack architectures. Some noted that using flexible, self-contained components can enhance companies' ability to swap out models and software as needed, as well as scale capacity more quickly.

“Simplicity and scalability often beat complexity and over-optimization. Early on, I spent a lot of time tuning everything for theoretical performance, but in practice, what mattered more was a stable, modular setup that we can monitor, debug and extend. Keeping things lean with well-documented components, clear logging and just enough automation made a world of difference.”

VP of IT architecture and infrastructure

Healthcare, 20,000-50,000 employees, US

“Would have prioritized modularity earlier, using microservices, making scaling and maintenance easier.”

Principal for AI and cloud

Professional services, >50,000 employees, UK

Several participants also recommended implementing strong infrastructure monitoring practices to preempt bottlenecks. Few organizations felt they were monitoring infrastructure KPIs as closely as they should, and many saw this as a gap in their approach.

“More redundancy, more monitoring, more overall visibility into what we plan to do in the future to ensure we are more agile.”

Head of cybersecurity and infrastructure

Machinery manufacturing, 1,000-5,000 employees, Germany

“It is hard to predict what the loads from the various models will be, so we need to benchmark and establish metrics.”

Senior manager for solutions architecture

Banking/insurance, >50,000 employees, UK

Just 33% of study participants said their organization was “very closely” monitoring infrastructure efficiency performance for AI workloads. This measurement gap becomes even more evident when evaluating specific components such as networking, where just 23% said their organization was “very closely” monitoring performance.

Insight 2: Data strategy and security must be addressed at the outset of AI projects

Initially, many organizations sought to use large pre-trained generative AI models to quickly harness AI capabilities without needing to clean and curate their own data. However, as they advanced in their AI journeys, study participants recognized the growing importance of data quality and management practices, even for generative AI workloads. To support high-impact use cases, such as using proprietary data to enhance the accuracy and relevance of AI responses, or using generative AI to access and summarize company data, many participants wished they had improved their data practices before embarking on large projects.

“You cannot compete with existing models like foundation models ... so we would use those, and you would combine that with your own data and your own tuning.”

VP of IT

Utilities, 10,000-20,000 employees, France

“There have been a bit of experimentation with maybe fine-tuning some of the models. But we say right now, we are not in the business of building models from scratch ourselves ... we either enrich [third-party models] with retrieval augmented generation techniques or with our own content ...”

Global head of data and AI architecture

Healthcare/life sciences, >50,000 employees, Germany

“Many of the assumptions was that AI would minimize or even eliminate the need to normalize your data, and it does to a degree, but not completely.”

Senior director of IT infrastructure

Consumer goods, 10,000-20,000 employees, US

Many survey participants encountered significant data-related challenges as AI initiatives transitioned from pilot to production. Multiple study participants noted that poor data strategies derailed entire projects, and inadequately planned data pipelines drove unforeseen levels of expense.

“We need to reconsider the design of our data pipeline. By processing streaming data in batch mode, we compromised real-time capabilities. Additionally, placing frequently accessed data in the cool storage tier instead of the hot storage tier resulted in increased latency.”

Director of technology

Professional services, 1,000-5,000 employees, Japan

“So, the data is fragmented. Like any other big corporate organization, especially with generative AI, nobody developed these documents thinking that generative AI is going to come. Second, all these documents are not curated for generative AI. So, a lot of times, what happens is people just try to use whatever data they have ... it leads to hallucination. It leads to exposing some confidential data, and things like that ... that is a huge challenge, and that’s a multiyear initiative.”

Senior director of AI products and platforms

Financial services, 20,000-50,000 employees, US


“Storing PB+ data on a hot tier of storage within the cloud is expensive. Making thousands of calls to retrieve that data within the cloud, also expensive. A hybrid setup would add latency to these ‘low latency’ calls and also incur egress/ingress charges.”

Director of engineering infrastructure

Financial services, >50,000 employees, US

When asked about changes they would make to previous AI initiatives, study participants consistently highlighted the need for greater focus on integration and data readiness. Notably, several study respondents wished that data and integration requirements had been more adequately considered not just in infrastructure planning, but when selecting AI projects to prioritize. “Making sure that our data was ready for our AI model” and “starting with an integrated platform in mind” were two representative reflections. As Figure 3 illustrates, even experienced organizations consider data pipeline management an in-progress area of development.

Figure 3: Most organizations say infrastructure could better support data pipeline management



Q. How well does your current cloud/on-prem infrastructure (compute/storage/networking/network security, etc.) support the following?
Base: Executive discussion board participants (n=30).
Source: S&P Global Market Intelligence 451 Research AI at Scale study commissioned by Verizon.

The importance of data as a driver of AI project success is also reflected in the results of 451 Research’s Voice of the Enterprise: AI & Machine Learning, Use Cases 2025 study: 54% of organizations with a below-average AI project failure rate said they considered data availability when prioritizing AI use cases, versus just 42% of those with above-average project failure rates. Interview and executive discussion board participants in this study echoed those findings and attributed data pipeline bottlenecks not only to technology limitations, but also to process and governance issues.

“For us, what drives the infrastructure is the data. What’s the data need? What’s the data requirement? Well, to be fair, it’s the business. We start with the business requirements, then you go into the data requirements that would address the business needs, and then you go into the AI infrastructure that would address the data. So, it’s a triangular approach, if you like.”

Director of data strategy
Life sciences, 50,000-100,000 employees, UK

“It starts from data labelling, to purging the data ... we already have a very robust life cycle ... to train a model, you may need some immediate data, but we have this huge governance life cycle where you might have to go through multiple processes and steps to get access to some data. So, we are trying to simplify that right now, as we speak.”

Senior director of AI products and platforms
Financial services, 20,000-50,000 employees, US

“Usually, what takes most time is, if you want to onboard new data sources, is to find someone in the company who wants to act as a data owner so that we have the governance straight ... this can take a little bit of a while. We are now putting a data governance community in place with the goal to get more data stewards and data owners on board.”

Global head of data and AI architecture

Healthcare/life sciences, 50,000-100,000 employees, Germany

Financial services organizations were notably more likely to describe their data management practices as highly mature. Half of the respondents from the financial services sector rated their practices at four or five on a five-point maturity scale, compared to 29% of respondents across all sectors. Among manufacturing participants, this figure was just 17%.

Insight 3: Effective project prioritization is key to controlling costs

There is widespread enthusiasm for adopting AI across organizations, with many senior leaders recognizing the technology’s potential benefits for their teams and objectives. However, in many organizations, the rush to harness AI has resulted in a fragmented, scattergun approach to investment. Numerous projects are being launched, but not all are strategically planned, and many are failing to deliver the desired results.

“The biggest pain point is, we’re going to try and do too many things too fast. And that’s going to create a bottleneck on resources. Those are internal resources as well as infrastructure resources. The more you add, the more you need to manage, right?”

Senior director of IT

Wholesale/retail, 1,000-5,000 employees, US

“Pushback to the business/investor that they cannot make the decision to implement the solutions before giving us enough time to evaluate them.”

Senior director of ITOps and strategy

Chemicals manufacturing, 1,000-5,000 employees, US

The breadth and rapidity of this investment is also evidenced in 451 Research’s Voice of the Enterprise: AI & Machine Learning, Use Cases 2025 study: Over 40% of surveyed departments indicated they expect to purchase or implement a new AI solution in the coming year. In many organizations, these new AI products and services will sit alongside longer-standing AI investments.

A key challenge amid rapid growth and scaling is the sharp increase in costs associated with AI initiatives. Experienced organizations in the “AI at scale” study repeatedly cited this concern. In some instances, ballooning expenses have threatened to undermine projects as they scale or forced organizations to narrow the scope of their projects.

“It will cost a lot more than you think! Be careful on presenting budgets for funding!”


Senior infrastructure manager
Healthcare, 20,000-50,000 employees, US

“Lack of sufficient budget only allows us to use some not very advanced and elegant solutions.”

Head of IT infrastructure and security
Retail, 1,000-5,000 employees, Germany

In many departments, additional software licenses contribute to growing cost concerns. However, for more experienced organizations, infrastructure requirements are also driving up AI-related investment. As these organizations scale their AI efforts to support a larger user base, they recognize the need to invest across networking, storage, security and compute — an insight illustrated in Figure 4.

Figure 4: Required infrastructure upgrades for AI span security, storage, networking and compute



Q. Which infrastructure component requires the most urgent upgrade to support your AI/ML initiatives?
Base: Executive discussion board participants (n=30).
Source: S&P Global Market Intelligence 451 Research AI at Scale study commissioned by Verizon.

Cost optimization is a common infrastructure-related challenge. When asked to assess the challenge of cost optimization on a scale of 1-5, the majority, 63%, assigned it one of the two highest-severity tiers. The largest organizations — those with more than 50,000 employees — identified the cost-optimization challenge as particularly severe, with 38% assigning the highest rating of 5. This compares to 10% of organizations with a headcount below this threshold.

Several executives cautioned that to help manage costs, organizations must avoid rushing into operationalizing projects before they fully understand the scope and concept. This strategy enables organizations to more accurately forecast potential costs and avoid unexpected financial burdens as projects scale.

“Assess cost factors to infer possible major cost impacts.”

VP of architecture

Professional services, 1,000-5,000 employees, UK

“Wait till the technical concept is complete to the high degree to avoid numerous and costly reworks.”

Head of IT infrastructure and security

Retail, 1,000-5,000 employees, Germany

Many study participants also emphasized the importance of managing costs by effectively prioritizing projects. By aligning initiatives with strategic objectives and focusing on high-impact projects, organizations can avoid inefficiencies and cost overruns that can result from pursuing too many initiatives at once.

“We have larger initiatives coming out of the strategic pillars that directly tie into that ... we drive based on, ‘OK, what business capability do we want to have in a higher maturity level, and is that a strategic priority for us?’”

Director of IT for EMEA

Manufacturing, 10,000-20,000 employees, Germany

“By narrowing down the features that are best realized with AI, we should minimize pay-per-request charges.”

Digital transformation lead

Manufacturing, >50,000 employees, Japan

To enhance project prioritization and determine which initiatives merit scaling, participants noted the value of effective measurement practices. Nevertheless, many organizations lack a structured post-implementation review process. Such reviews are critical for capturing lessons learned, improving cost and effort estimations, and ultimately enabling more efficient and cost-effective project delivery.

“I never saw any review of one particular project and basically analyzing whether this has been able to achieve the value that they were promising two or three years ago.”

VP of IT

Utilities, 10,000-20,000 employees, France

“Costs are going to be in the infrastructure ... which we then chop up per department per GL [general ledger] code in the, in our accounting world ... once we ask them, is this working for you, then we could say, well, you're spending this, let's draw it back and see if it still is just as successful or making you just as happy if we reduce spending for your department just a little bit.”

VP of infrastructure and security

Healthcare, 5,000-10,000 employees, US

“We follow the metrics. So, when we start working on our initiative, we check what kind of value proposition this particular initiative will have ... whether this will help us saving some cost. Then how much. Whether this will help us in improving our department's efficiency. And then by what extent.”

Senior engineering manager

Financial services, 10,000-20,000 employees, UK

Insight 4: AI initiatives do not operate in isolation; organizations must emphasize both integration and collaboration

Study participants often pointed to legacy infrastructure and outdated processes as barriers to AI progress. As their AI initiatives matured, many recognized the need to modernize both governance and infrastructure. Several participants emphasized the importance of considering AI workloads in the context of overall infrastructure modernization efforts, as AI initiatives exert additional pressure on those broader efforts.

“Outdated infrastructure governance, which is burdened with technical debt and not aligned with technological advancements, is hindering transformation. For example, there are excessive restrictions on data movement and accessible devices, and change management processes are rigid, requiring complex approval workflows for system updates. In the era of AI, it is crucial to apply more flexible policies and take technological advancements into account.”

Director of technology

Professional services, 1,000-5,000 employees, Japan

“Current policies and architecture are heavily optimized for traditional IT workloads, not data-intensive, dynamic AI workflows. String[ent] outbound traffic controls, limited access to cloud GPUs, and lengthy approval cycles for integrating new tools or APIs slow down experimentation.”

VP of IT architecture and infrastructure

Healthcare, 20,000-50,000 employees, US

“Organizational maturity to integrate AI within embedded solutions, due [to] data handling restrictions and legacy systems, [is a top-three challenge].”

Senior infrastructure/network security architect

Medical devices, 1,000-5,000 employees, Denmark

Given AI's strategic importance, a wide range of stakeholders is involved in both selecting AI projects and making infrastructure decisions. Participants acknowledged the need to proactively manage these stakeholders to ensure alignment across the organization.

“Decision-making often involves collaboration across departments, with a focus on aligning AI/ML initiatives with broader corporate strategies and ensuring ethical and security compliance.”

Global head of data and AI architecture

Healthcare/life sciences, >50,000 employees, Germany

“AI is never stand-alone. At the end of the day, it is embedded into a software application. You have the infrastructure teams, which are different people in charge of either on-premise[s], cloud or edge infrastructure.”

VP of IT

Utilities, 10,000-20,000 employees, France

“I would consider developing principles, standards, frameworks, and check basically the studies and references patterns everywhere Because today ... they want to go into the AI, they have no clue what it is all about ... they will constrain the whole organization to adopt this without basically having an understanding of the consequences and the implications of these constraints.”

Senior data architect

Financial services, 10,000-20,000 employees, Singapore

Broad inclusion can result in tensions due to misaligned visions. In some cases, the enthusiasm of less technical staff about AI meant that study participants felt they had to educate stakeholders about AI and where it can be most effectively applied.

“We spent some time educating the business ... because they wanted to run a model to predict whether there was a potential fraud, suspicious activity in this data set. The problem is that no model could detect anything abnormal because there was not even one single data point that we could label as abnormal. So, the model could learn nothing from that.”

Head of IT

Healthcare/life sciences, 5,000-10,000 employees, Japan


“I’d say, especially in 2024 and beyond, it’s much more of a strategic role where I’m talking to individual board members on how the AI is not going to steal our secret business, our trade secrets, and how [our primary vendor] is not in the business to have data exposed. So, a lot of that is educating C-suite, educating the lines of business about the possible technologies that we have, but also rightsizing them.”

Senior director of IT
Wholesale/retail, 1,000-5,000 employees, US

“Educate the business on the difference between deterministic and non-deterministic systems.”

Head of IT strategy and transformation
Banking, 500-1,000 employees, Germany

Figure 5: Wide array of stakeholders involved in infrastructure decisions to support AI/ML workloads



Q. Specifically, regarding your AI/ML workloads, which teams/functions are involved in your infrastructure decisions to support these AI/ML workloads? Please indicate their roles (Lead, Influence, Inform or Not involved).
Base: Executive discussion board participants (n=30).
Source: S&P Global Market Intelligence 451 Research AI at Scale study commissioned by Verizon.

European participants were most likely to report that non-technical business stakeholders are involved in AI project decision-making, with 80% saying non-technical stakeholders either lead, influence or have the final say in these decisions. This contrasts with 60% of US-based participants and 45% in Asia-Pacific.

Implications: Laying the foundation for AI at scale

Study participants shared a clear set of recommendations and insights that broadly fell into three key areas: AI strategy, networking and security.

AI strategy

Organizations benefit from a “start small and scale” approach to AI. However, this must be guided by a few critical considerations. Starting small should not equate to a lack of direction or proper prioritization. A forward-looking mindset is essential, and many participants wished they could start again with better planning practices. This includes considering the architectural impact of scaling successful pilots, understanding the ongoing data requirements of initiatives, and building with modularity and agility in mind. Additionally, project selection must account for existing systems and tools while also addressing implications for people, processes and governance frameworks.

Networking

Organizations intending to invest in AI must take a proactive approach to network planning. This involves looking beyond the immediate demands of pilot projects to anticipate full production needs, factoring in network performance, availability and security from the start.

To effectively support AI initiatives, networks should be built for low latency, high bandwidth and high availability, and they can benefit from investments in optimization techniques such as distributed storage and caching. Network segmentation can also prove to be a key enabler, with dedicated subnets and isolated routing for AI hardware helping manage traffic efficiently and prevent congestion. For more information on networking considerations, please see the companion report on [networking for AI at scale](#).

Security

Organizations investing in AI must develop a comprehensive security strategy that spans all IT environments to protect company data and intellectual property. Security measures must adapt to evolving AI capabilities, and the strategy must acknowledge the expanding attack surface created by AI operations. Many participants recommended adopting a zero-trust framework, and several suggested that organizations should integrate advanced security tools such as improved authentication mechanisms, API security, enhanced threat detection and monitoring systems to safeguard against unauthorized access and insider threats.

Many participants acknowledged the need to balance project performance with security and governance, ensuring that fast experimentation does not compromise safety and compliance standards. By fully embedding security into their governance frameworks, organizations can mitigate the risk of regulatory fines and reputational damage while fostering a secure environment for AI delivery. For more information on security considerations, please see our companion report on [security for AI at scale](#).

About the author



Alex Johnston

Senior Research Analyst

Alex Johnston is a senior research analyst on the 451 Research Data, AI & Analytics team at S&P Global Market Intelligence. He focuses on emerging technologies and how they can be applied in business contexts. Alex's primary coverage areas are artificial intelligence, distributed ledger technology, event stream processing and data marketplaces. Alex's recent areas of concentration include monitoring the emerging generative AI market, tracking the evolution in blockchain use cases and investigating real-time architectures.



About Verizon Business

Verizon Business is all about helping organizations like yours succeed in today's dynamic digital environment. We provide essential network solutions that support and enhance business operations and empower how millions of people live, work, and play every day. This research was commissioned to provide insights into the challenges of AI deployment, share peer experiences, and clarify the strategic choices facing IT leaders. Our aim is to offer a deeper understanding of these intricate dynamics, especially the critical importance of network and security infrastructure. We believe these insights can help your organization to scale AI confidently, maximize its potential, and consistently stay ahead in this rapidly evolving landscape.

About this report

A Discovery report is a study based on primary research survey data that assesses the market dynamics of a key enterprise technology segment through the lens of the “on the ground” experience and opinions of real practitioners — what they are doing, and why they are doing it.

About S&P Global Market Intelligence

At S&P Global Market Intelligence, we understand the importance of accurate, deep and insightful information. Our team of experts delivers unrivaled insights and leading data and technology solutions, partnering with customers to expand their perspective, operate with confidence, and make decisions with conviction.

S&P Global Market Intelligence is a division of S&P Global (NYSE: SPGI). S&P Global is the world's foremost provider of credit ratings, benchmarks, analytics and workflow solutions in the global capital, commodity and automotive markets. With every one of our offerings, we help many of the world's leading organizations navigate the economic landscape so they can plan for tomorrow, today. For more information, visit www.spglobal.com/marketintelligence.

CONTACTS

Americas: +1 800 447 2273

Japan: +81 3 6262 1887

Asia-Pacific: +60 4 291 3600

Europe, Middle East, Africa: +44 (0) 134 432 8300

www.spglobal.com/marketintelligence

www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2025 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.