2024

State of AI Application Strategy Report

A New AI Stack Emerges

(But Few Organizations Are Positioned for Success)

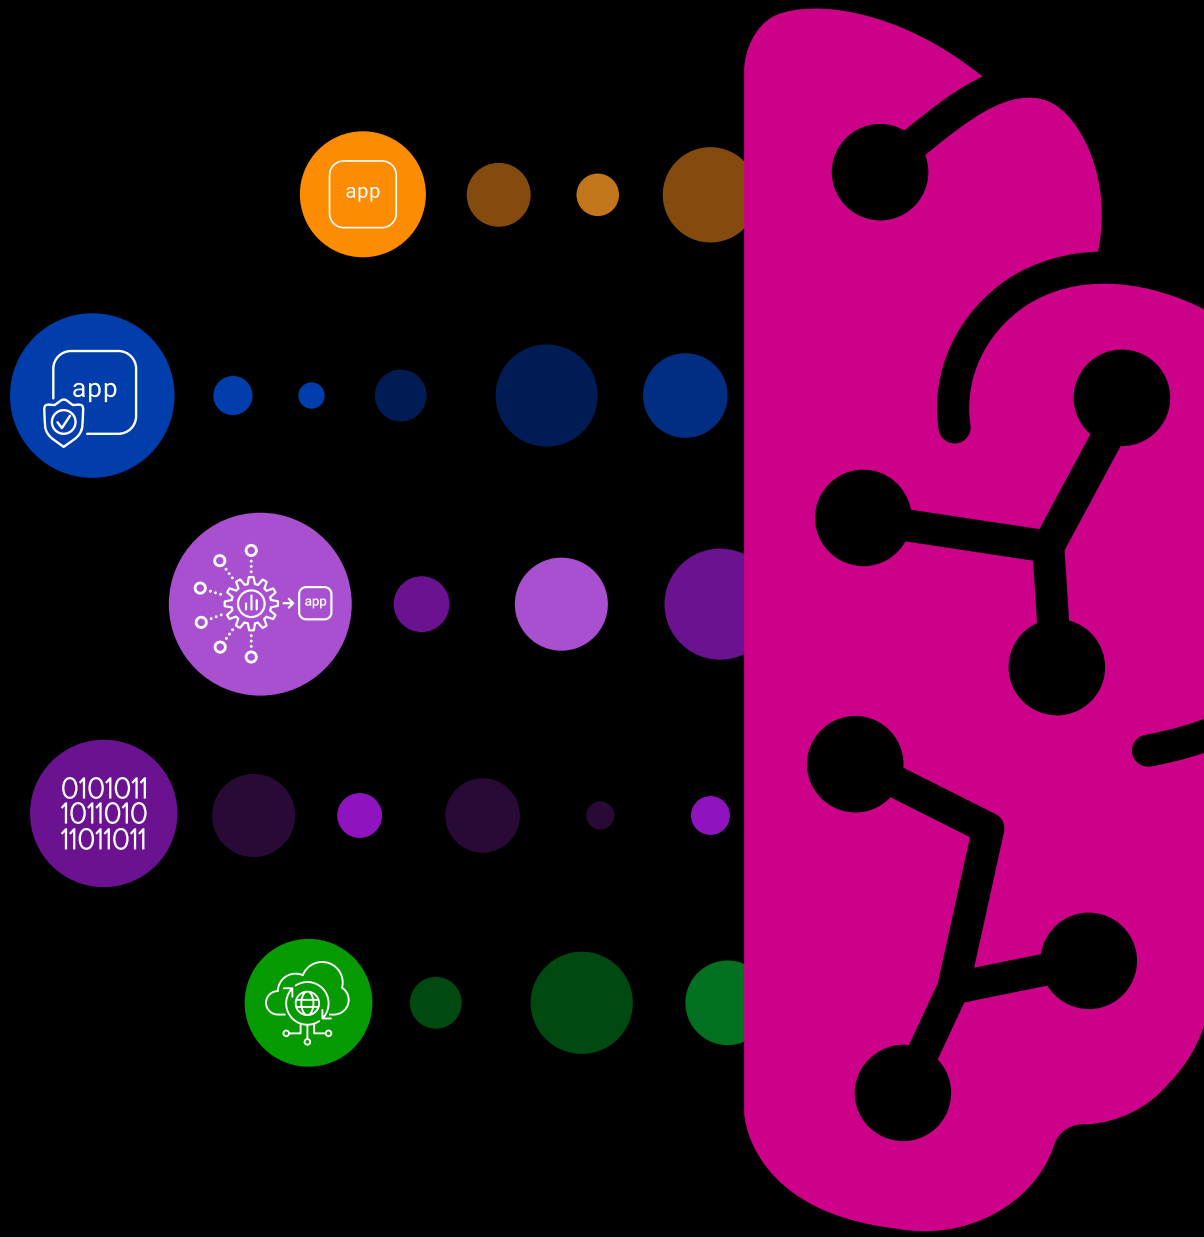

Contents

Introduction

AI Deployment Has Leapt Forward Despite Strategic Gaps

The digital transformation journey occupying organizations in recent years leads to new opportunities, new customers, and digital business success. Frequently, business use of artificial intelligence (AI) assistance is used as a marker for advanced progress on the journey.

Such progress is evident in the results of the tenth annual F5 State of Application Strategy survey, which reveals massive growth in the share of organizations operating in the third phase of digital transformation, AI assistance. (See terminology on page 7 for more about the three phases.)

Today, 75% of organizations are implementing AI assistance, and it's safe to say that virtually every senior IT leader is being asked to help determine where—or where else—AI might fit into strategic plans. This leap in activity in just a few years no doubt hinges on the introduction and accessibility of generative AI.

That's the good news. It comes alongside more sobering data because the journey is only a metaphor. Digital transformation is ongoing and with no fixed destination; technology and the digital economy will continue to evolve. In the meantime, it's worthwhile to ask what it really means to harness AI. What AI use cases truly serve a strategic purpose? And which are mostly low-risk experimentation or even public relations opportunities that imply sophistication but don't add much business value?

"[We are] using generative AI to build new revenue streams with novel applications."

Director, financial services organization (more than $10 billion in annual revenue)

Use of AI Has Skyrocketed

% of respondents with AI as a core business focus

Unfortunately, the results of our research suggest that AI deployments serving less-strategic functions may be common today. That's probably because the majority of organizations are still struggling to put into place everything needed. The danger is that they'll proceed without key prerequisites for long-term success.

Particularly concerning are the three-quarters of organizations lacking the data strategy or data quality required for the most effective AI implementations. Without a data foundation in place, too many organizations are in danger of succumbing to the significant challenges of AI while their better-prepared competitors surge ahead. And even those with both a formal data strategy and reliable, integrated data face a number of risks and potential threats that lie in the path of succeeding with this powerful technology.

Exactly what are those risks? The results of our annual global survey plus focused, subsequent research clarify the most pressing answers. The information we gleaned specifically from AI decision makers also suggests that a new AI stack is forming, at least for inference purposes. (The stack will look somewhat different for AI model training.)

The Next Stack Is Forming Around AI

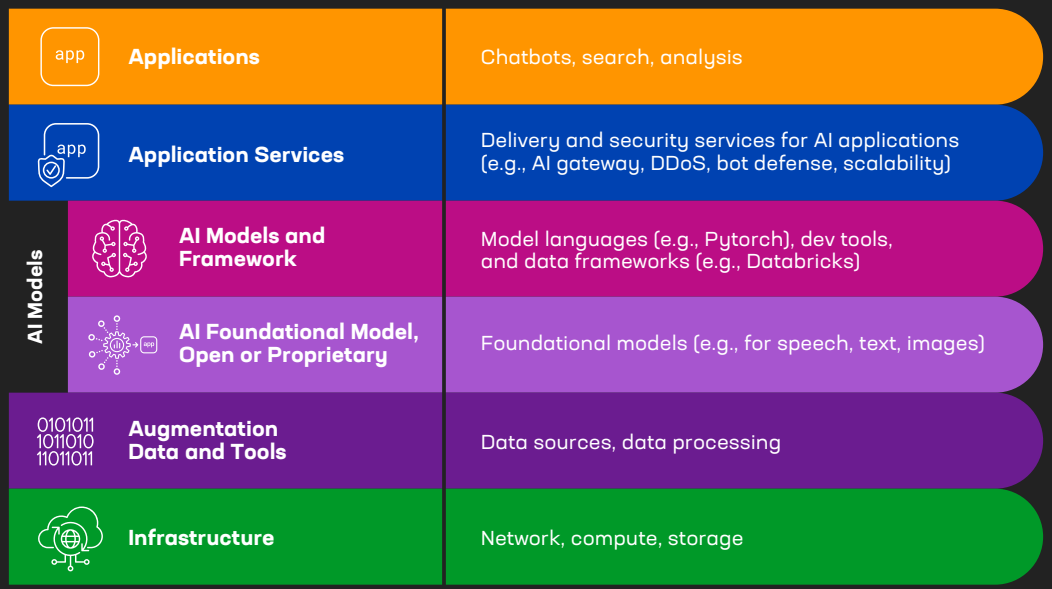

As implementations proceed, organizations must consider all tiers of this AI stack. AI models tend to attract the most attention, and many are available today, with model selection depending on the specific use case. But that's only one of the key decisions to make. Other considerations, some of them more fundamental, include what infrastructure will be needed to run the model, what data it will be trained on and ingest, how both model and data will be managed, which apps will rely on it, and how to deploy and protect those apps as well as the rest of the stack.

These questions tend to be interrelated. They often must be answered iteratively, if not simultaneously, with the combined answers guiding implementations. To prevent missteps, each tier of the stack needs attention. That specifically includes app services and technologies that can secure AI models and protect and deliver the apps relying on them.

Given perennial resource constraints, not every tier gets the thought it deserves. Opportunities for ill-advised detours abound. Fortunately, paths for leapfrogging competitors are becoming apparent, too. Read on to explore the state of AI today, which traits the leaders have in common, what decision makers anticipate in the near future, and how careful attention to this stack can help wise organizations overcome common challenges and capture the full benefits of AI.

"Our organization leverages AI to enhance supply chain efficiency and accuracy by implementing predictive analytics for demand forecasting, optimizing inventory levels, and automating procurement processes. AI-enabled logistics solutions improve route planning, reduce delivery times... [and help] identify and rectify inefficiencies in real time."

Senior leader, cloud service provider (more than $10 billion in annual revenue)

Terminology Used in This Report

The phases of digital transformation
The digital transformation journey can be viewed as a three-phase transition to digital business. With task automation as phase 1 and digital integration as phase 2, this journey is iterative rather than strictly linear. Phase 3, the use of AI to assist with business decisions and operations, indicates digital sophistication, at least in affected parts of the organization.

Artificial intelligence (AI)
A computational system that uses algorithms, other programming, and machine learning (ML) models to analyze large volumes of data, identify patterns, and use them as the basis for autonomous decisions or actions in response to new data or programming prompts.

Generative AI
A type of AI that can produce new content— such as text, images, or video—by following the patterns identified in training data sets to create new outputs with similar statistical properties.

Predictive AI
A type of artificial intelligence that can apply computational models or algorithms to historical data to predict future trends or events.

AI model
A computational framework comprising one or more algorithms trained on a curated data set to detect patterns and statistical relationships. These are then used as the basis for performing specific tasks or generating outputs.

Training model
An AI model currently undergoing training, during which it processes curated data to refine its algorithms and improve its ability to accurately identify patterns and statistical relationships relevant to the desired tasks.

Inference model
An AI model that has completed its training and is now used to interpret new data. It applies the knowledge gained during training to deduce appropriate actions or provide relevant information in response to the new input.

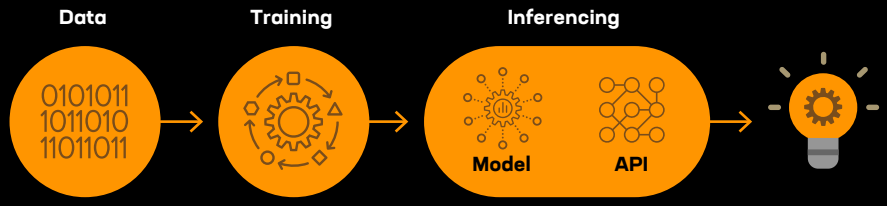

Section 1

Why AI Investment Is Growing

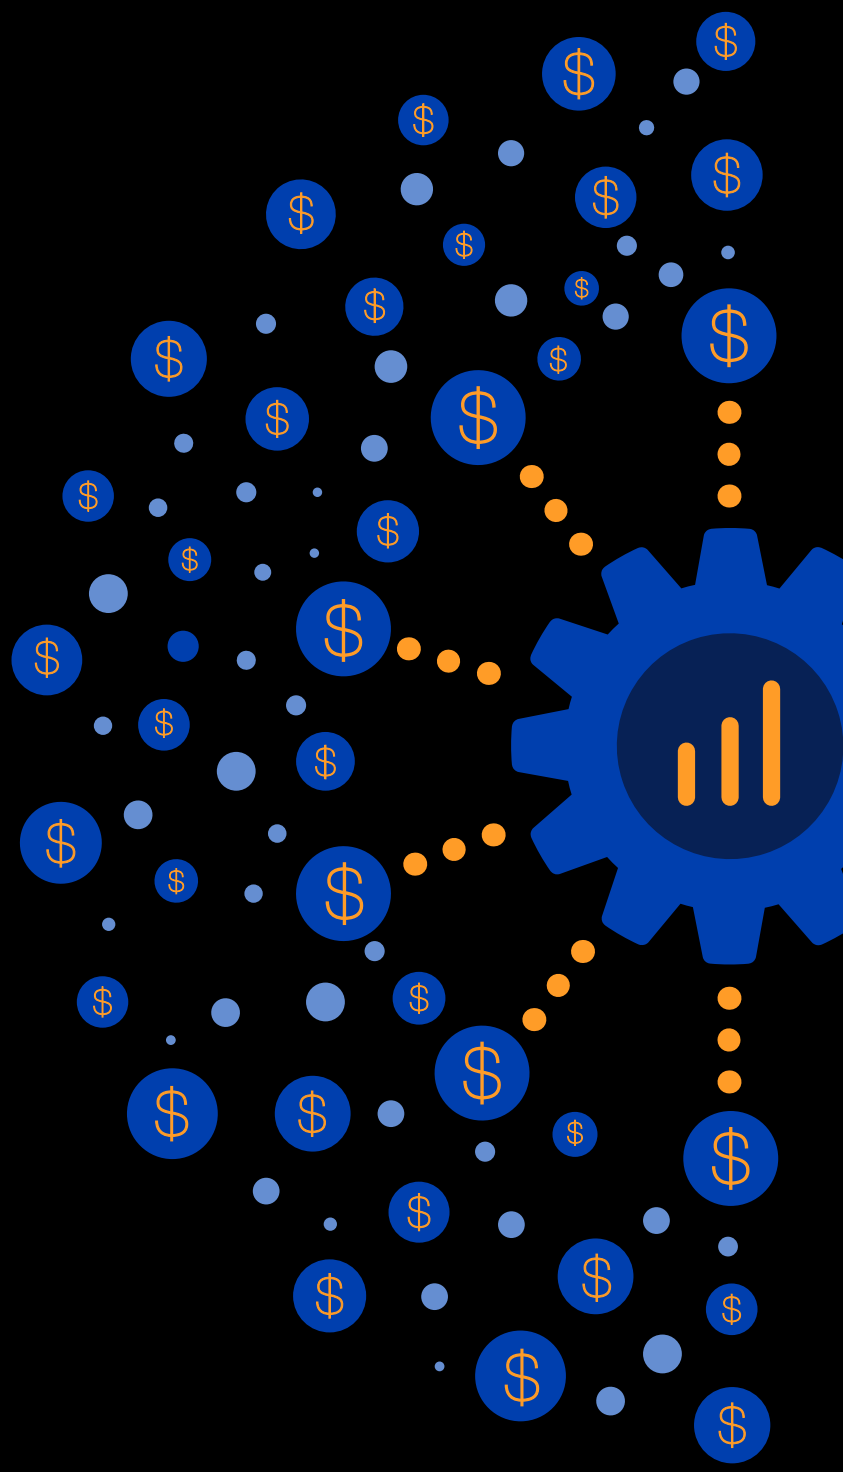

Currently, the use of AI as a strategic or core focus of the business varies significantly by industry. Cloud service providers and manufacturers are furthest ahead, deploying AI at rates of 84% and 82%, respectively.

On the other hand, the lowest use reported occurs in the government sector, at 64%, with healthcare and energy/utilties—both at 65%—right behind. The need in these industries to prioritize user privacy and data security likely accounts for their trailing positions.

That's a reason, but it's not an excuse—or at least not one that will long be sufficient, for either the industries or individual organizations that trail the leaders. While citizens may shudder at the idea of governments using AI, for instance, no one can ignore its vast efficiency benefits. Governance, energy, and healthcare are all particularly expensive, most of us would like more service at a lower cost, and AI may well become part of the answer.

> "[We are] utilizing Microsoft 365 CoPilot for internal/external unstructured data summarization, data location, and data creation. Looking at additional gains with ServiceNow and other product sets."

CIO, healthcare organization (more than $1 billion in annual revenue)

At any rate, that hope for increased efficiency is what motivates AI leaders. Decision makers expect healthy efficiency boosts from each of the most common AI use cases. For instance, they expect content generation (such as the creation of marketing copy and images) to boost productivity 42%. That amounts to the difference between working six hours versus 10 hours to meet a financial reporting deadline, for example. Customer service chatbots are expected to deliver a 33% increase—the difference between resolving a customer support case in 20 minutes rather than 30. Similarly, summarization of meetings or work products such as reports warrants hope for a 31% increase in efficiency. And employee satisfaction can only improve when taking a lunch hour is no longer a luxury because tasks that once took an hour can be completed in 40 minutes.

With such benefits on the horizon, it's no wonder investment in AI is rising across industries. On average, AI expenditures represent nearly one-fifth of organizations' overall IT budgets in 2024. Furthermore, nearly everyone (94%) expects that budget bite to grow, rising more than one-quarter (26%), on average, by 2026. Nobody expects it to drop.

The various tiers of the AI stack, starting with compute capacity in the infrastructure, certainly could consume spending on this order. Nonetheless, we suspect these investment predictions may reflect a momentary exaggeration brought about by the flashy promise of generative AI. Survey respondents named it the most exciting trend of 2024, and enthusiasm may be inflating their budget estimates.

Still, the energy with which organizations are pursuing AI is evident. As they source new technology solutions, it's increasingly likely those solutions will rely on AI. An explosion of AI-based tools will result. Of course, whether the IT team will buy and control all those tools or business units will spend separate budgets to inadvertently create a new flavor of "shadow IT" remains to be seen.

AI Budgets Will Increase

18%

of IT budgets to be spent on AI in 2024

94%

of organizations expect to increase that budget within two years

Section 2

The Use Cases Driving AI Spending

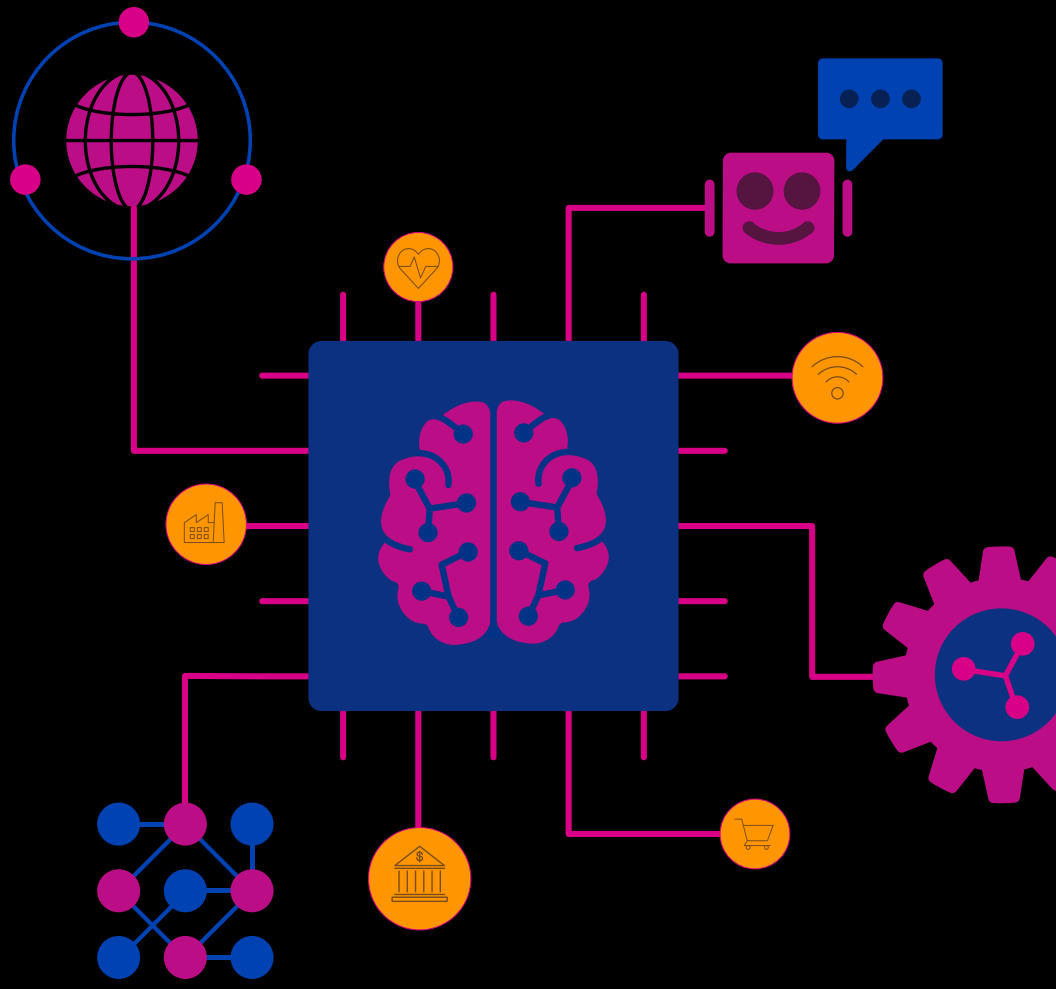

Survey respondents told us they plan to leverage AI for multiple purposes, particularly securing applications and data, managing more efficiently with AI operations (AIOps), and improving business processes.

Top AI Use Categories Have Stayed Consistent

% of respondents planning to use AI

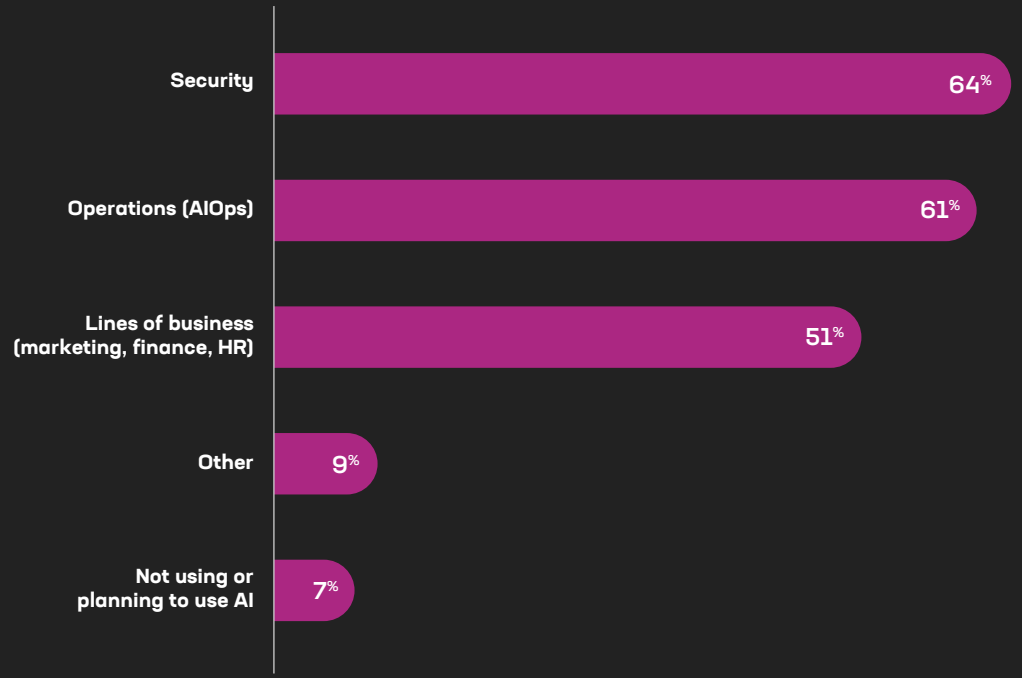

But security, AIOps, and line-of-business enhancement are very broad categories, with multiple use cases in each and some that are easier to address than others. Organizations generally are moving most quickly to implement AI solutions that represent quick victories, promise the highest efficiency benefits, or both. As noted, that typically means content generation, chatbots, and productivity copilots. Often, the speed is made possible because the appropriate model is available through Software as a Service (SaaS) or can be integrated into existing infrastructures via APIs. Indeed, generative AI solutions like these are often characterized as the "low hanging fruit."

For example, customer service chatbots have already been deployed by 36% of survey respondents. Chatbots are followed closely by office productivity solutions such as copilots, which more than one-quarter (27%) of organizations are currently using as one way to supercharge employee productivity. When summarization tools are added, productivity solutions represent the leading AI use category.

The trouble with low-hanging fruit is it's not always worth even minimal effort. Accordingly, some organizations may be wasting resources on use cases that won't deliver the benefits promised. For instance, we all encounter chatbots today that do little more than rely on a large language model (LLM) to spit out answers from an existing FAQ that customers could more efficiently access directly. Substantial questions stymie such chatbots, often simply delaying human support and annoying customers in the process.

Productivity and Chatbots Lead AI Implementation to Date

% of respondents with AI solutions already implemented

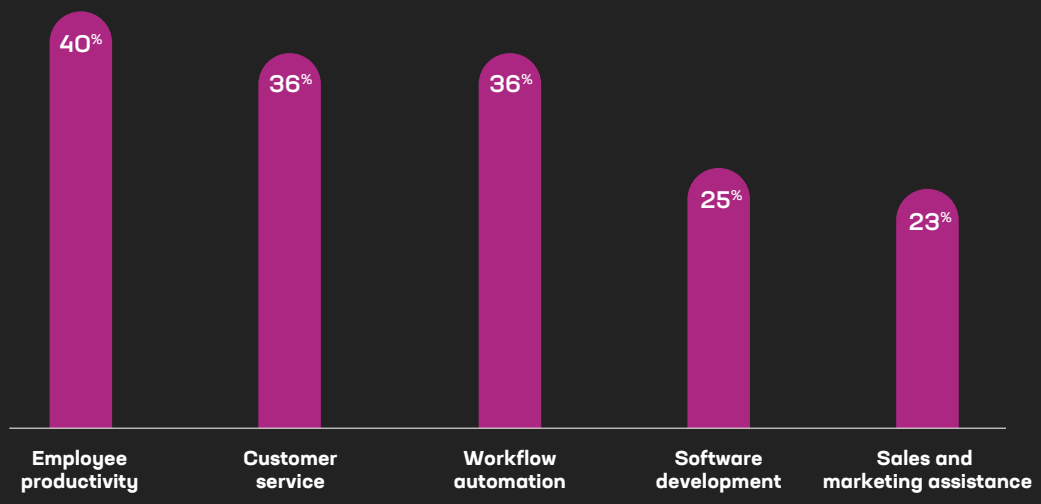

“With a target of improving efficiency each year, AI is key to improve the accuracy and speed of as many processes as possible.”

Senior leader, financial services organization (more than $1 billion in annual revenue)

Furthermore, only 24% of organizations say they’ve implemented generative AI at scale. More than two-thirds (69%) are still researching AI use cases, and nearly as many (63%) are working now on completing proofs of concept. Value assessments and use case priorities could change significantly as a result.

Scaled Implementation Are Still Mostly Pending

% of organizations at various stages of AI implementation

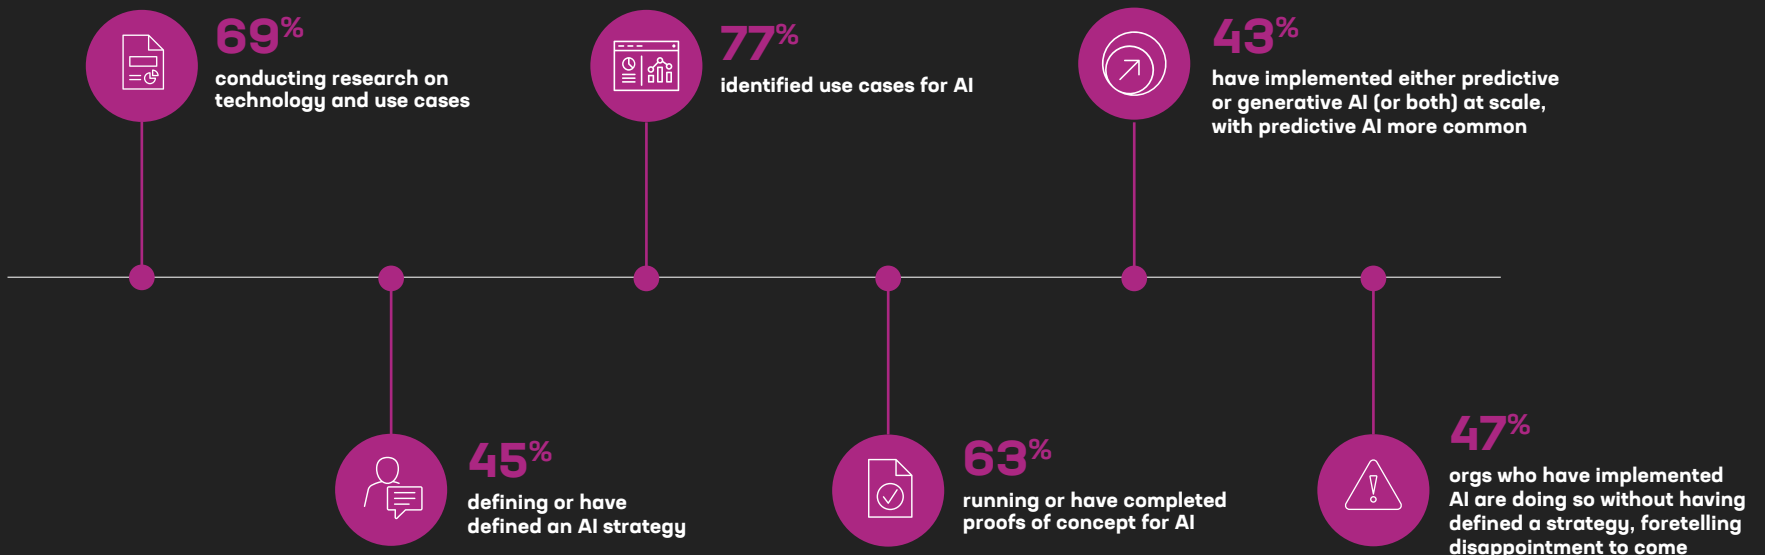

That may be why none of the use cases expected to yield the highest efficiency benefits rank as the top priority for AI in 2024. Instead, decision makers in most industries are focusing on deploying AI for workflow automation.

The exceptions are cloud service providers, manufacturers, and technology organizations. In these three verticals, other use cases edged out workflow automation as the highest priority.

Workflow Automation Ranks as the Top 2024 AI Priority Overall

% of respondents identifying each use case as their 2024 priority

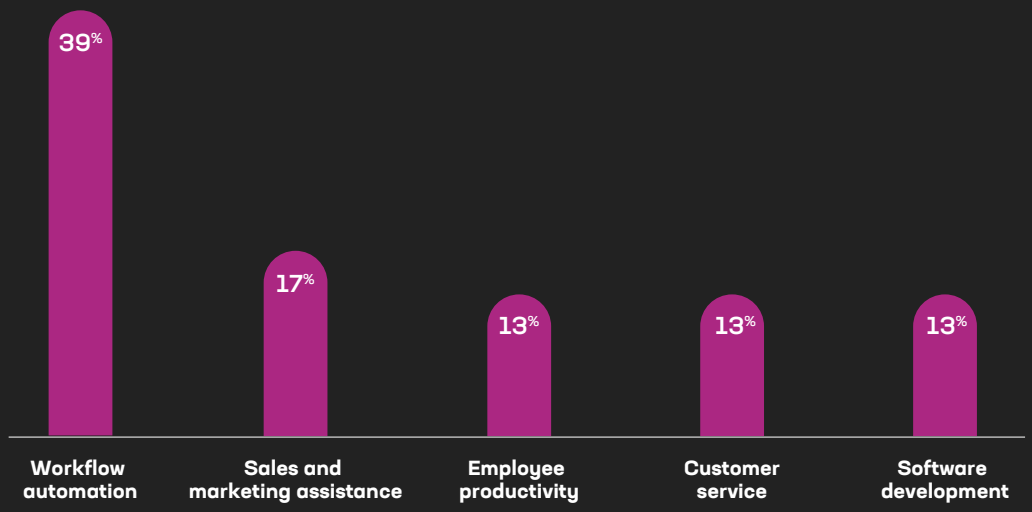

The Top 2024 AI Priority, by Industry

% of industry respondents calling this use case their top AI priority

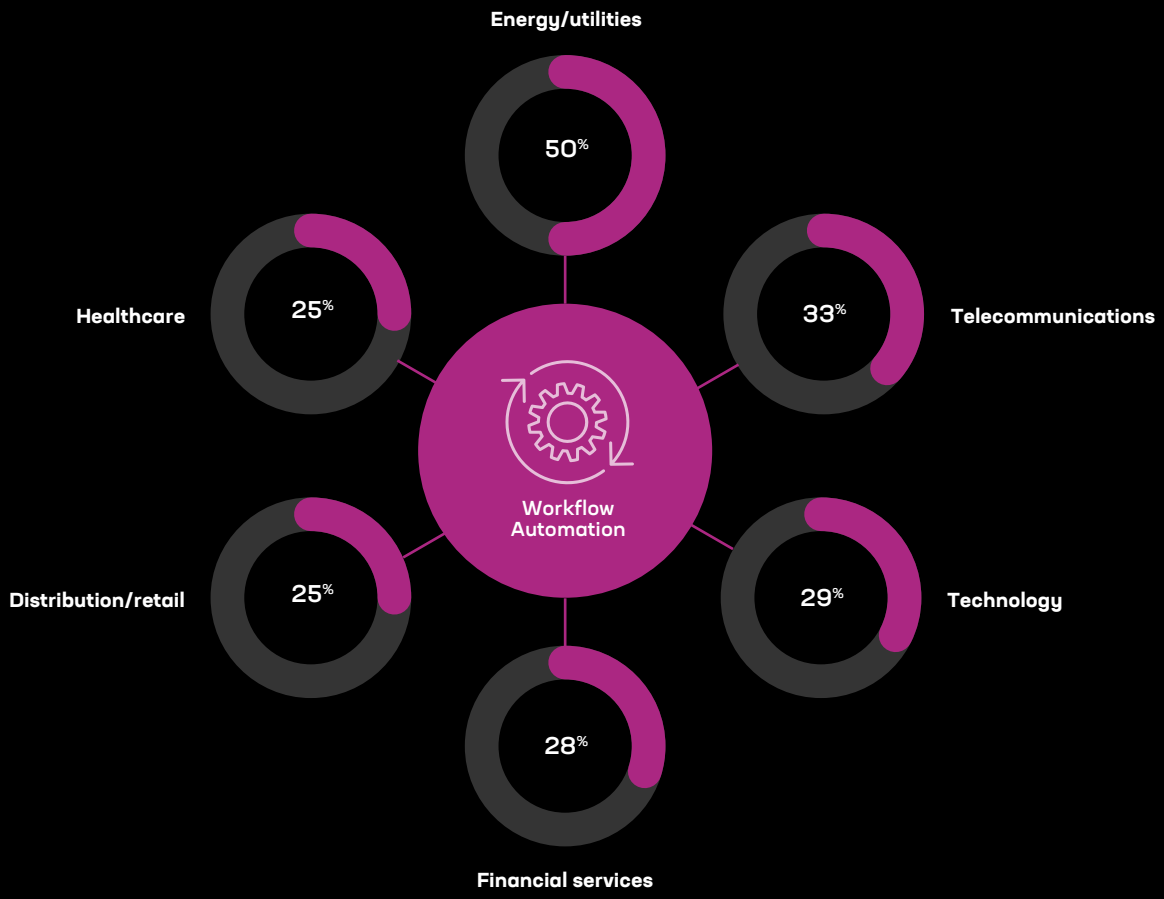

The availability and previous implementations of SaaS-based AI probably sway this prioritization. SaaS AI solutions such as copilots are already available and cost-effective for employee productivity and content generation use cases. Workflow automation is more challenging to obtain, but as existing or SaaS-based AI solutions increase employee throughput, the importance of optimizing those processes grows. It's the digital transformation journey all over again: First, improve individual processes, then integrate those processes, and finally use process data to optimize the whole system. Workflow automation is analogous to the third step. It can be the hardest because it demands the most integration and data integrity, but those who accomplish it find it worthwhile—and likely a competitive advantage.

Data from our annual research has long emphasized the use of generative AI to automate application security and delivery processes—AIOps—as a means of taming complexity. By comparison, plans to automate workflows in lines of business have lagged. That makes today's focus on workflow automation somewhat surprising. But many organizations need more time and effort to successfully integrate all the moving parts necessary to automate app security and delivery operations. That time and effort translate into dollars, and budget is a perennial—and the single largest—barrier to change.

Several Factors Slow the Automation Required for Effective AIOps

% of respondents naming each barrier

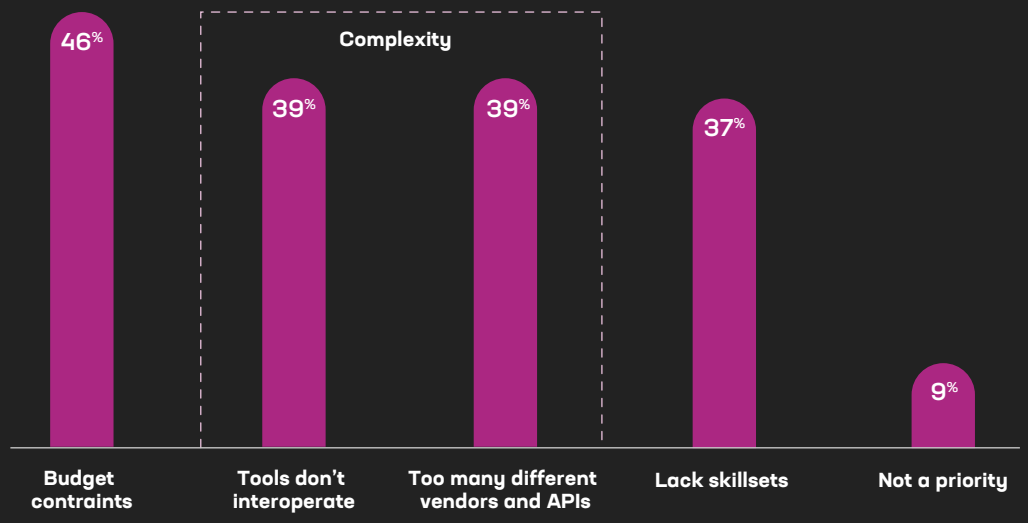

“[We are] implementing AI-powered recommendation systems, demand forecasting algorithms, and customer sentiment analysis to personalize offerings, optimize inventory management, and anticipate market trends for introducing innovative products.”

Executive IT leader, distribution/services/retail company ($1-5 billion in annual revenue)

The value currently being placed on workflow automation, plus the capabilities that generative and predictive AI offer to achieve it, will likely spur significant technology standardization. Moving forward, that standardization will ease automation and in turn facilitate the implementation of AIOps.

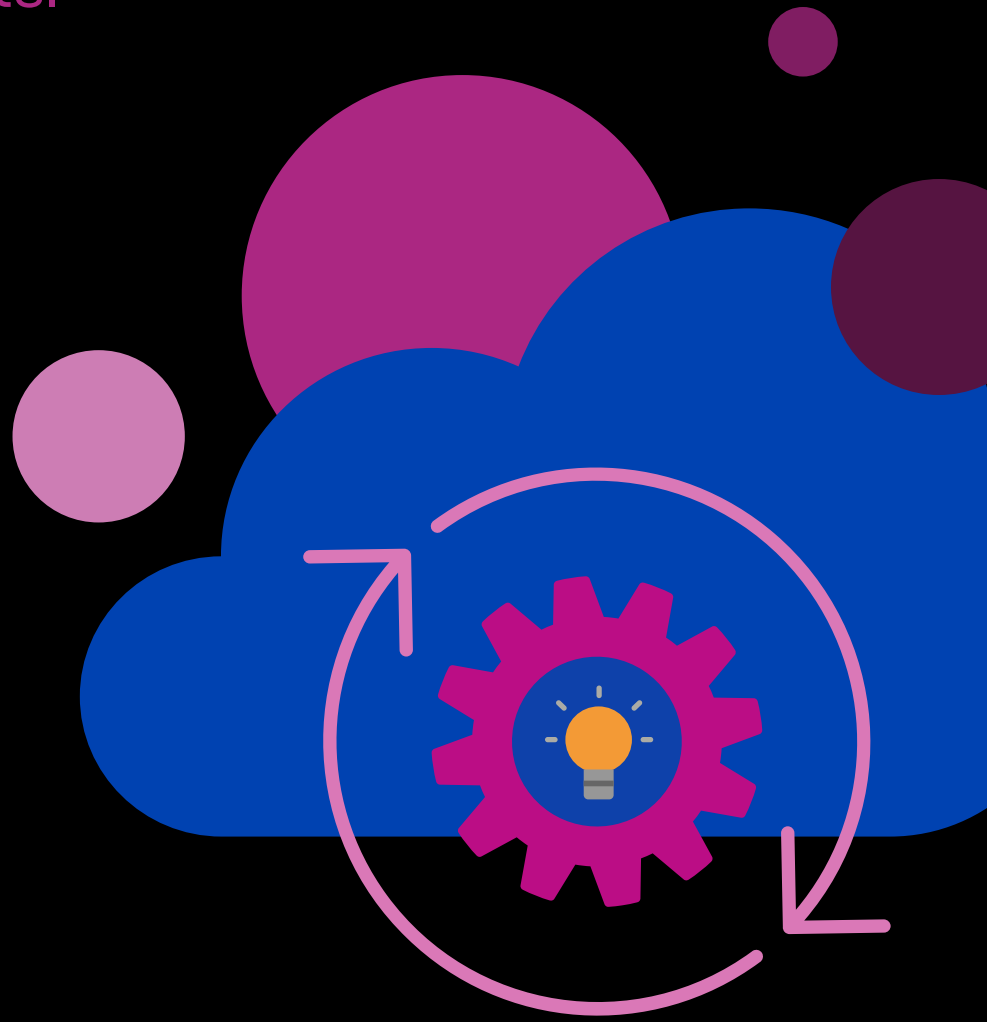

Section 3

The Emerging AI Stack and Its Implications

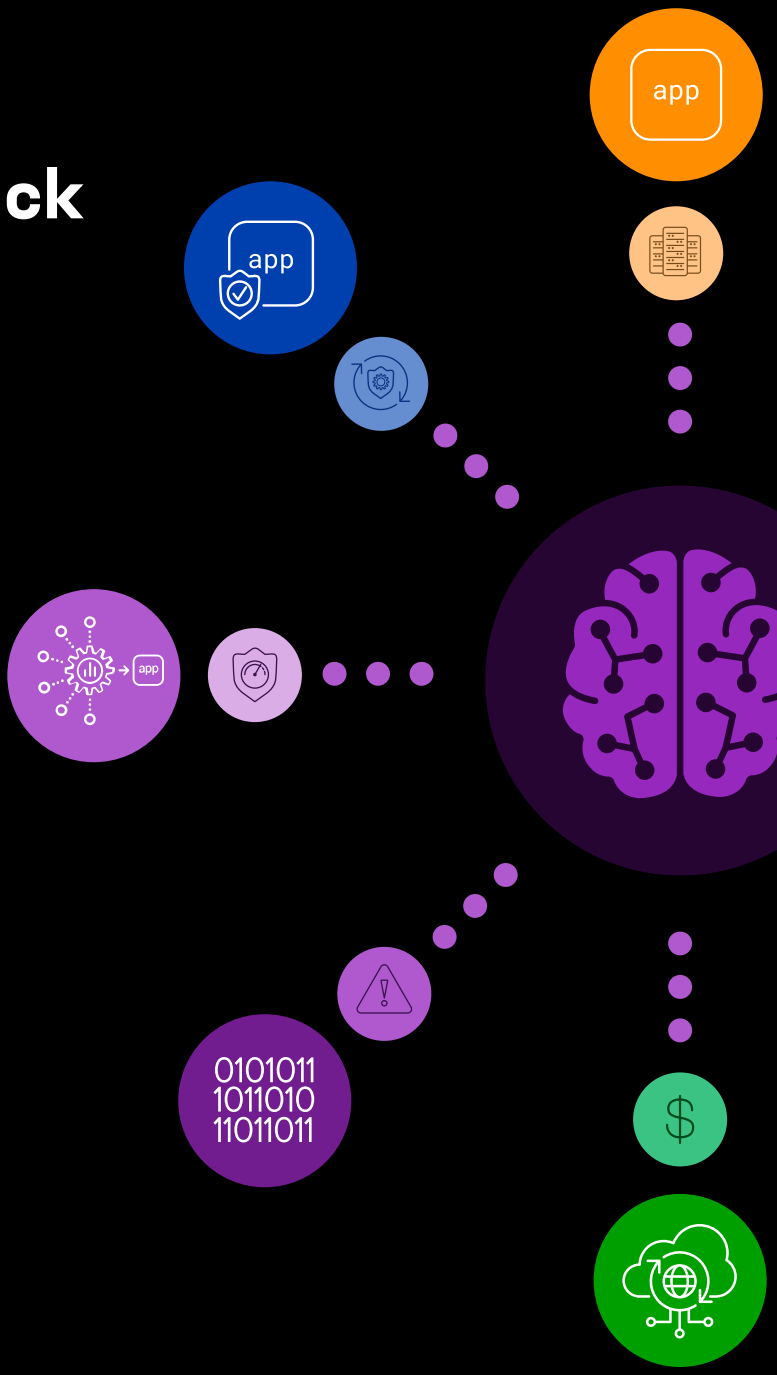

In addition to clarifying AI hopes and priorities, our research provides insight into how organizations are thinking about and putting resources behind each tier of the emerging AI stack. The occasionally alarming implications can guide decisions yet to be made, especially for organizations that hope to take leading positions (or at least not be left in the dust).

Infrastructure: Concerns about costs

Compute capacity sits at the foundation of the AI stack—and is one reason decision makers expect AI budgets to swell. In fact, when asked their biggest concerns about AI model development, training, and inference moving forward, nearly two-thirds of respondents (62%) in our annual survey told us they were worried about compute costs.

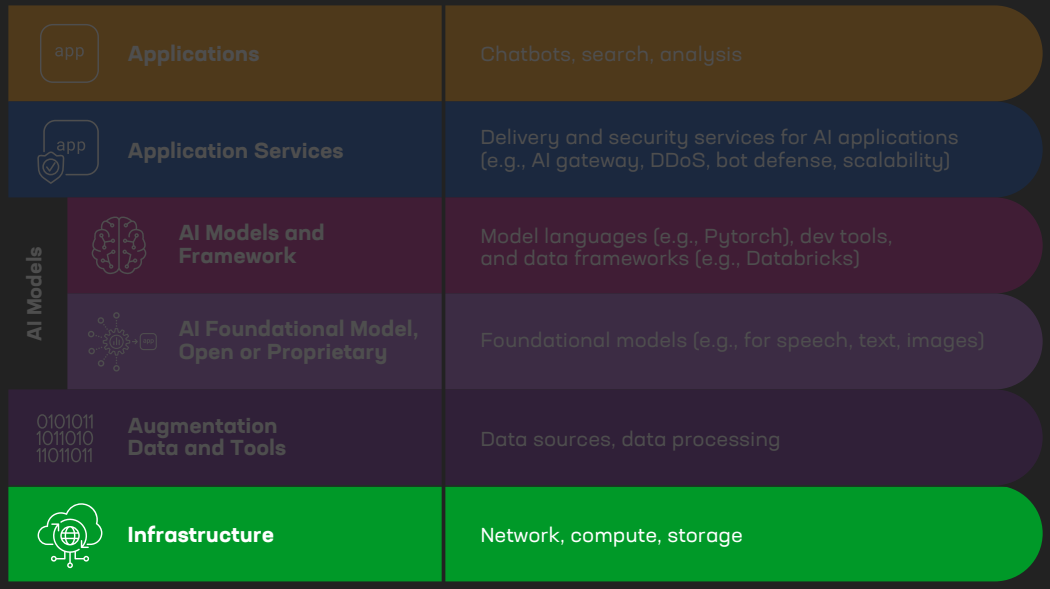

AI will most likely follow other, similar technology cycles, where costs at the infrastructure layer such as computing capacity initially represent the bulk of the investment. As infrastructure competition matures and the ecosystem expands, spending will shift to other layers of the stack.

Specifically, once compute capacity is sufficient, the number two concern—the security of AI models—will probably become the new top priority for spending. Survey respondents validate this imperative by saying they expect to spend 44% more on security within a few years as they scale deployments. Those organizations deploying AI at the core of the business especially can't afford to suffer a breach.

"[We are] using it to properly understand financial models and metrics, understand customers' pain points, and make it easier to provide better support for customers."

Cloud decision maker, financial services organization ($1-5 billion in annual revenue)

Compute Resources Are a Substantial Concern

% of respondents naming each concern

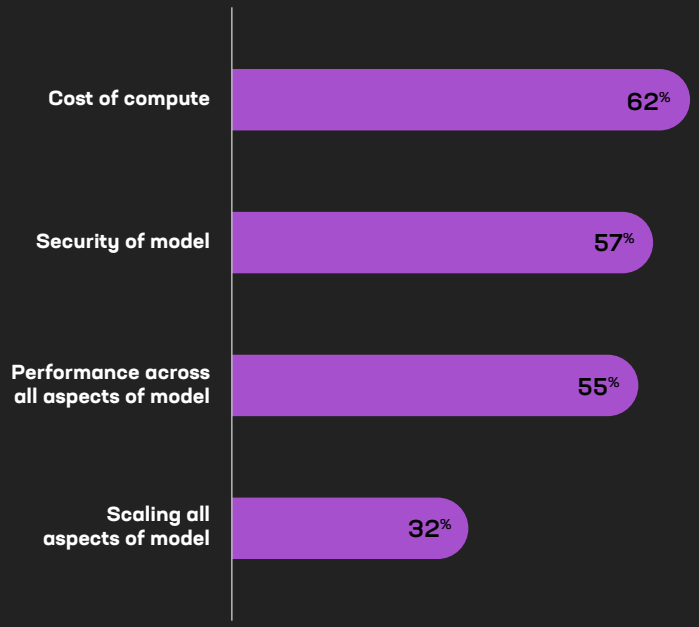

The data tier: In distress

For all the concern about compute costs and security, data maturity is a more immediate and potentially bigger challenge. Data—including what resides in AI models themselves—are an integral component of the stack, and weakness at this tier will hinder the full power of both predictive and generative AI.

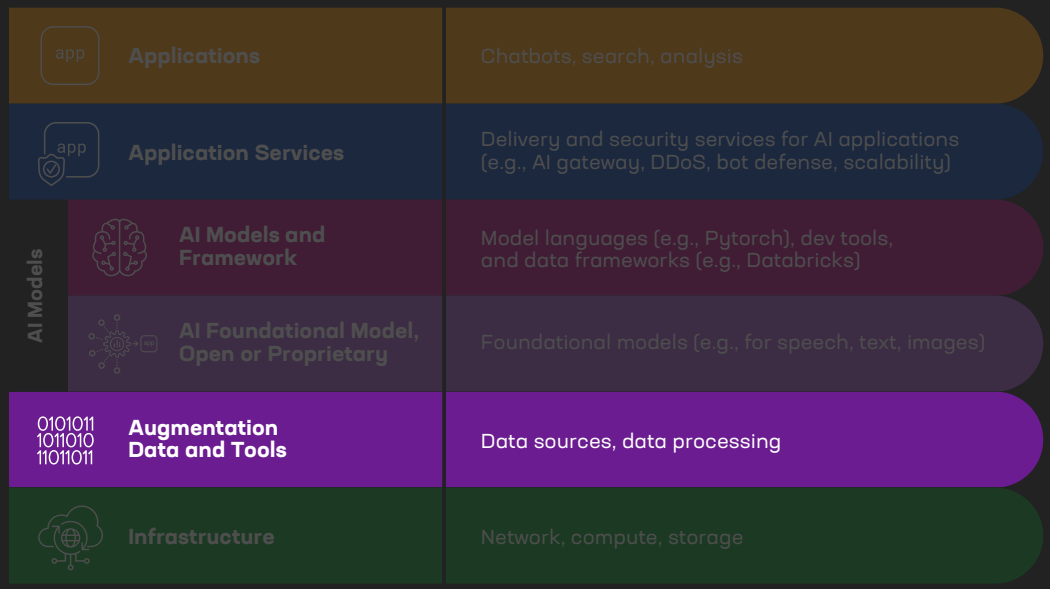

Appropriately, when naming overall challenges blocking AI adoption, nearly three-quarters of respondents in our annual survey—many of whom have already implemented AI—say data issues, including an inability to scale their data practices, are still their top hurdle.

Data Maturity Tops the Challenges to AI Adoption

% of respondents naming each challenge

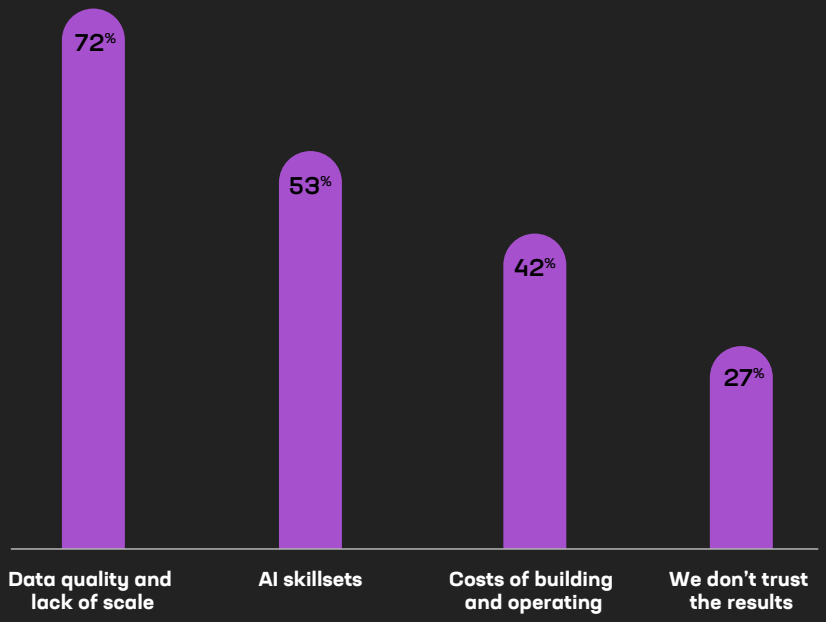

Organizational data strategies generally lack maturity. Just over half of organizations (53%) have defined a formal data strategy. Given the relative newness of data science and established best practices for operational telemetry and other data, organizations often fall back on established business data practices instead. Those simply may not be appropriate for AI and operational use cases.

Not surprisingly, organizations that consider themselves "data savvy" and use data in most processes are also the most likely to have implemented generative or predictive AI at scale (33% and 31% of those respondents, respectively). But they're also most likely to have defined a separate AI strategy and specific use cases to pursue, actively conduct research on technology and use cases, and have proofs of concept underway or complete. The existence of a data strategy may be seen as a marker of their overall AI commitment and leadership.

That said, most organizations (77%) still lack a single source of truth, and that includes more than half of those with data strategies in place.

Data Quality Lags Even When Data Strategies Exist

% of respondents at each stage of data maturity

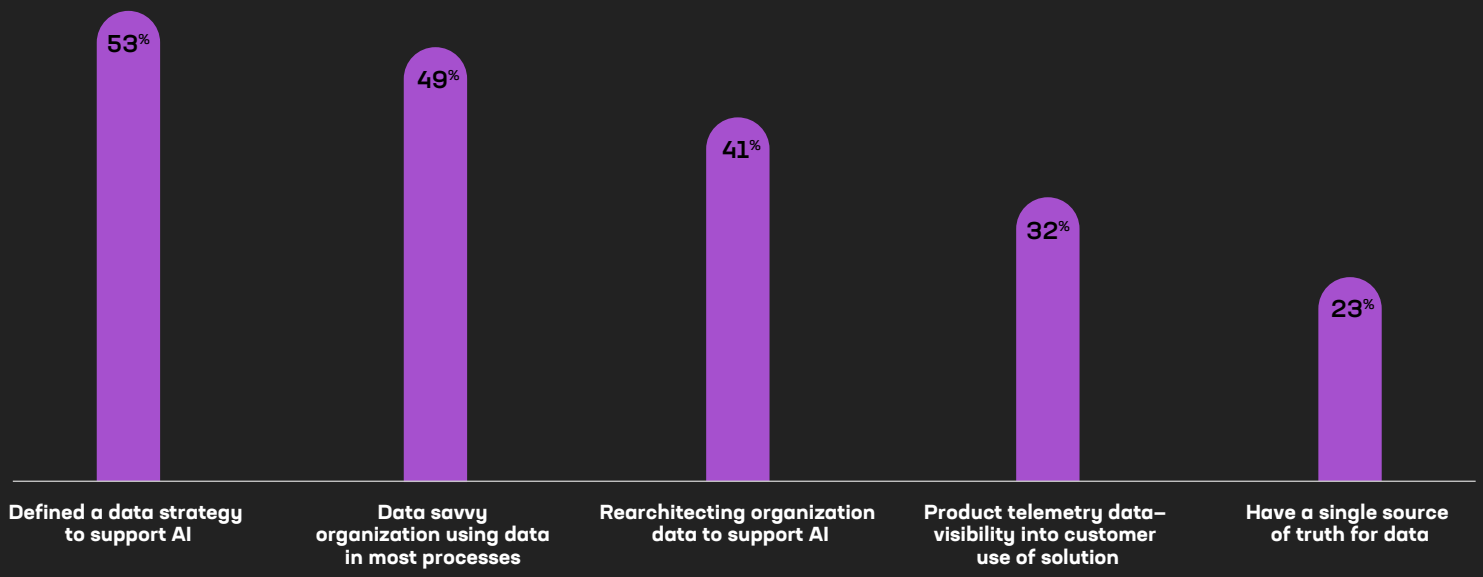

Such organizations face an uphill struggle to extract the insights they'll need to power strategic AI use cases such as workflow automation or product development. They'll first have to solve their data quality challenges by, for instance, using APIs to integrate disparate data sources. In fact, more than one-quarter (27%) of organizations are currently taking this approach. Another quarter (25%) are adopting cloud data platforms, while 17% are focused on new data schemas to support greater extensibility and usability of data companywide.

Those efforts are worthwhile, because one thing is certain: Data and the systems for obtaining and handling it—from security and management strategies to related tools, practices, and technologies—are key for successful adoption and optimization of AI assistance, no matter what enterprise domain it may empower.

AI models: The beating heart of the stack

AI models, which rely on patterns in curated data to assess new data and respond, are the beating heart of the AI stack, along with their associated languages, libraries, data frameworks, and development tools.

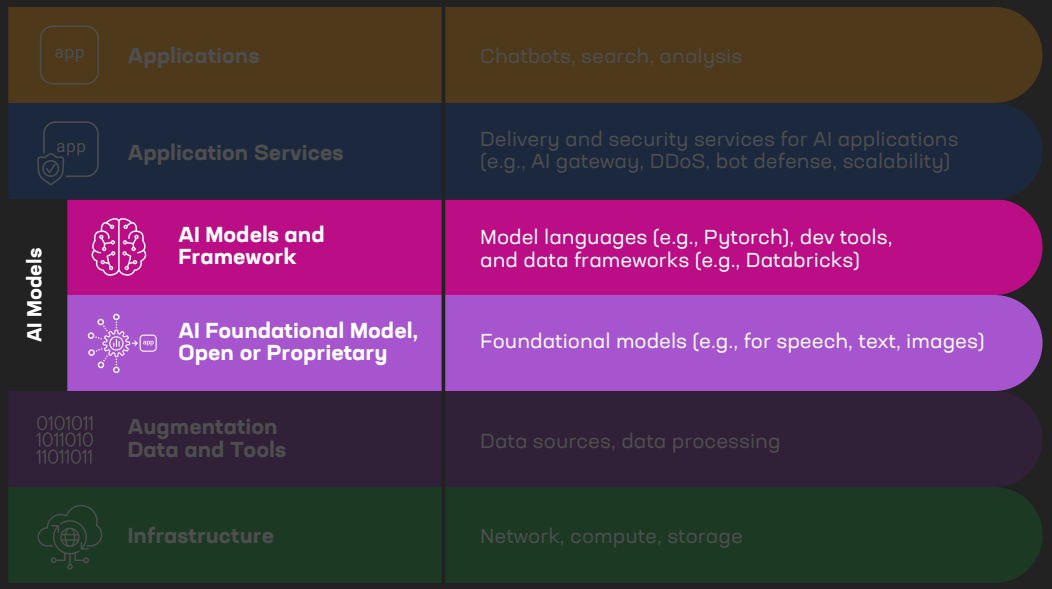

Though hundreds of models are available today, organizations already using AI rely on 2.9 different models, on average. The model selected depends on the desired use case and objectives, and organizations tend to gravitate toward a handful of popular open source and proprietary models. In general, proprietary models from well-known providers such as Microsoft, OpenAI, Google, and Amazon are favored.

But decision makers are not averse to open source models. Several, such as Hugging Face, Llama 2, and Mistral, enjoy healthy adoption rates. In fact, some use cases prompt a significant preference either for open source or proprietary models.

For instance, half of organizations (50%) use Llama 2 for office productivity use cases, while the Microsoft model (e.g., the power behind Microsoft Copilot) is preferred over OpenAI's GPT series (42.8 % to 21.43%). For customer chatbots, however, open source models don't seriously compete with Microsoft's Bing Chat model (42.11%), OpenAI GPT models (31.58%), or Google Gemini (15.79%).

Yes, Microsoft's AI is OpenAI's GPT-4; the difference is the services Microsoft makes available for integration with the LLM. These include authentication, security, and monitoring that are not as readily accessible with core OpenAI services. But the topic of services brings us back to workflow automation and the standardization that supports it. It's not only the model itself that's important; it's the availability of a robust ecosystem of services that support access to the model and integration with it.

This is why the most popular AI models are often offered by cloud providers. The providers already deliver a rich set of services that support applications of all kinds to greatly simplify adoption of their models.

"We utilize AI chat to compare our product and solution set to competitor offerings and identify gaps that can feed into our roadmap for product development."

Senior leader, financial services organization ($1-5 billion in annual revenue)

That's important, because when it comes to where those models will reside and how they'll be delivered and managed, the responses to our annual survey suggest they'll be deployed just like most applications today—that is, spread across hybrid IT estates both on-premises and in public clouds. Even looking two years into the future—a long time in a fast-changing world—AI implementation will only reinforce IT hybridization.

Within the context of hybrid IT and the complexities that result, we expect more standardization of enterprise architectures. This belief is supported by the popularity of SaaS-based AI models such as C3.ai and Shift Technology. As the primary means of integration with SaaS tools and software is via a well-defined API, decision makers expect this popularity to endure. Greater standardization of platforms and vendors will reduce the required number of APIs for most use cases. That's especially true for the AI use case most valued now—workflow automation.

It's worth noting that multicloud networking (MCN) also can ease the growing complexity of hybrid estates and ease management of the infrastructure tier required to support the AI stack. Multicloud networking can deliver connectivity while securing AI models, no matter where they reside.

Finally, recognize that current AI models—open source, proprietary, SaaS, and combinations—are not set in stone. Current rates of adoption and relative market shares could change drastically as organizations modernize and eventually shift from experimentation to scaled, strategic deployments supported by standardization and mature automation.

Deployment of AI Models Will Be Hybrid Too

% of organizations planning to use each deployment model

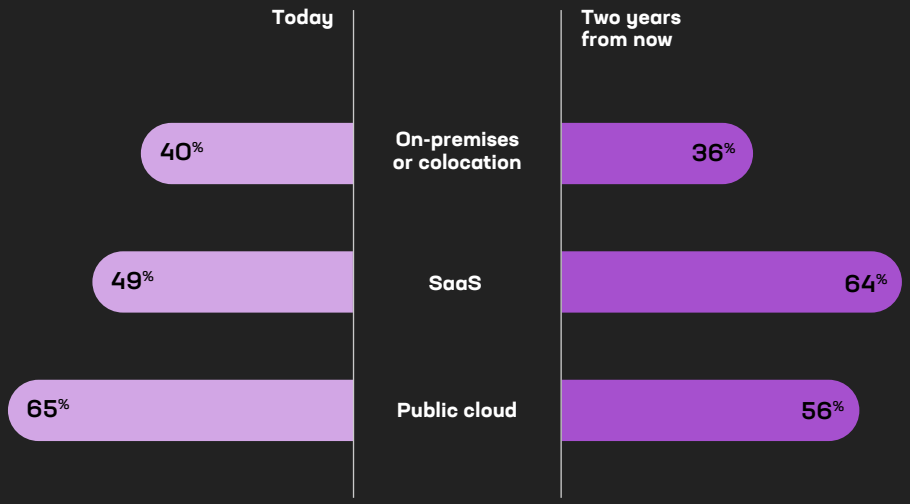

App services: Protection and management support

Regardless of where they're deployed, AI models will need support from technologies for accessing, managing, and optimizing them, and both the models and the data associated with them will need protection. Application services can help fill these roles.

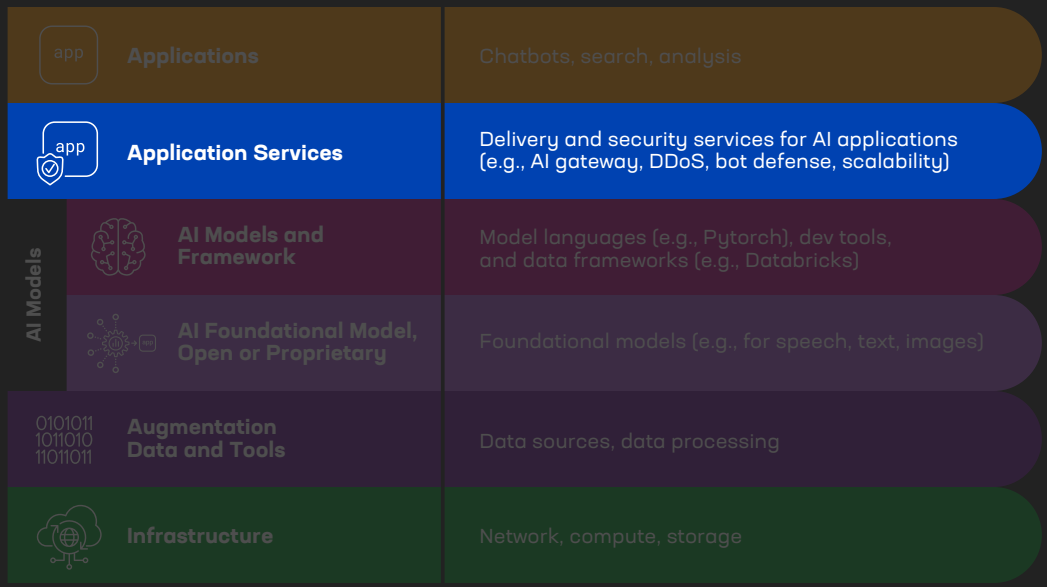

In today's sophisticated cybercrime landscape, most organizations are aware of the risks that threaten AI implementations. Indeed, decision makers rank attacks—themselves powered by AI—as their top AI security concern.

Their fears are warranted. The status quo of finding "zero days" with understanding of signatures or patterns to counter are probably over. Only extreme programmability of the data path, also fueled by AI, will be able to counter AI-driven attacks.

But that isn't all. Notably, all four concerns in the chart draw significant worry, and the ranking generally persists regardless of where the organization expects to train or deploy its AI models. One exception: Organizations deploying in the public cloud harbor somewhat greater concerns about liability (25.8% of respondents) than do those deploying on-premises (21.8%). Alternately, perhaps their ranking of security concerns helps drive their deployment choices; the cause/effect relationship is not yet clear.

Regardless, the close rankings suggest that all four security concerns will trouble the sleep of AI decision makers. The insomniacs may be less aware of how application services can address all four, along with other security and app delivery challenges.

AI Powered Attacks Lead AI Security Concerns

% of respondents concerned about each

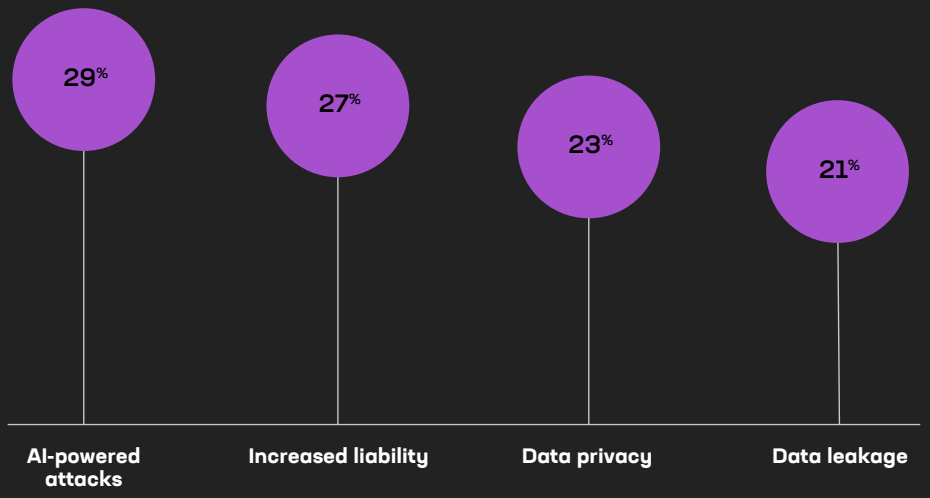

“We strongly leverage AI to accelerate our pace of bringing new products and offerings to the market.”

Operations team member, financial services organization ($1-5 billion in annual revenue)

Which app services could be most helpful? Decision makers asked how they plan to secure their AI implementations (or already are doing so) name traditional app services such as API security, DDoS protection, and bot protection. Data masking and exfiltration prevention are also high on the list.

A Variety of Protective Services Are Needed

% of respondents using or planning to use each type of protection

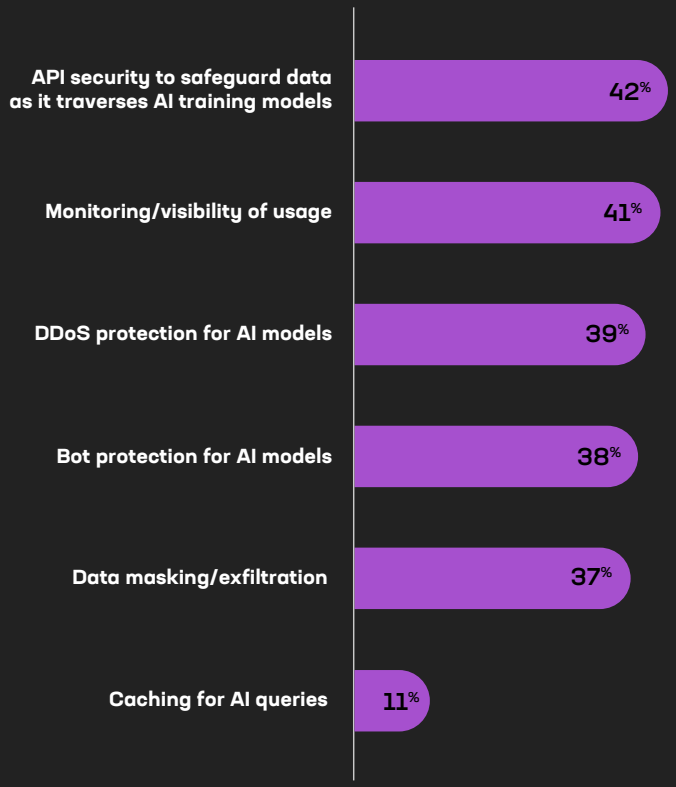

When asked to rank available AI security capabilities, they place preventing data exfiltration and leakage of personal identifiable information (PII) at the top of the list. But several other protections are considered nearly as important.

Note that two of the protections in the ranking—preventing bias and hallucinations—are vital for the responsible use of AI but are not strictly data security issues. Yet they're often included with security because they relate strongly to governance and guidelines. Perhaps a separate market for relevant solutions will arise—much as posture management solutions have developed alongside, but separate from, app security services.

Finally, app services could help address not only the security and performance of AI models but the data quality issues many organizations still face. In particular, there's almost certainly a role for technologies centered on the acquisition and analysis of telemetry from business operations, including the models, the apps that rely on them, and the defensive tools in place. The most prepared organizations can mine this telemetry for insights that can feed evolving AI models while helping IT teams ensure the ongoing suitability of the AI infrastructure. That's another reason application delivery and security services ought not be overlooked as AI deployment planning proceeds.

AI Protections Vie for Most Important

Average overall ranking (of 100 possible points assigned)

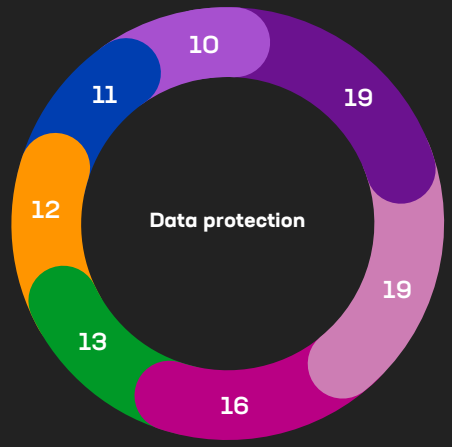

- Blocking data exfiltration
- Preventing PII leakage
- Preventing prompt injection attacks to LLMs
- Preventing LLM from presenting biased data
- Encrypting data in motion
- Preventing LLM from presenting hallucinations
- Encrypting data at rest

Applications: As vulnerable as any

The apps driven by the AI models and data—from well-known consumer assistants like Alexa and Siri to narrowly focused enterprise tools such as inventory optimization engines—represent the top tier of the AI stack.

As with the AI models themselves, AI apps are likely to be deployed both on-premises and in public clouds. In either case, they'll need the same kinds of delivery and security services as apps serving other use cases. That means everything from authentication and identity technologies to API security and bot protection.

A few of these support services might apply specifically to AI apps, but most won't. That doesn't mean the others are not important. At its heart, AI consumes and generates data—lots and lots of data. That data and the apps through which users interface with the models will be as vulnerable to performance issues or attackers as any, leaving business security and data privacy at risk.

In addition, understanding how those AI apps are performing, and for whom, is insight app services can deliver as part of comprehensive data collection and processing. Organizations that hope to become digital leaders can capitalize on such insight to assess the value they're gaining from their AI implementations, how to maximize it, and what use cases and AI-driven apps to implement next.

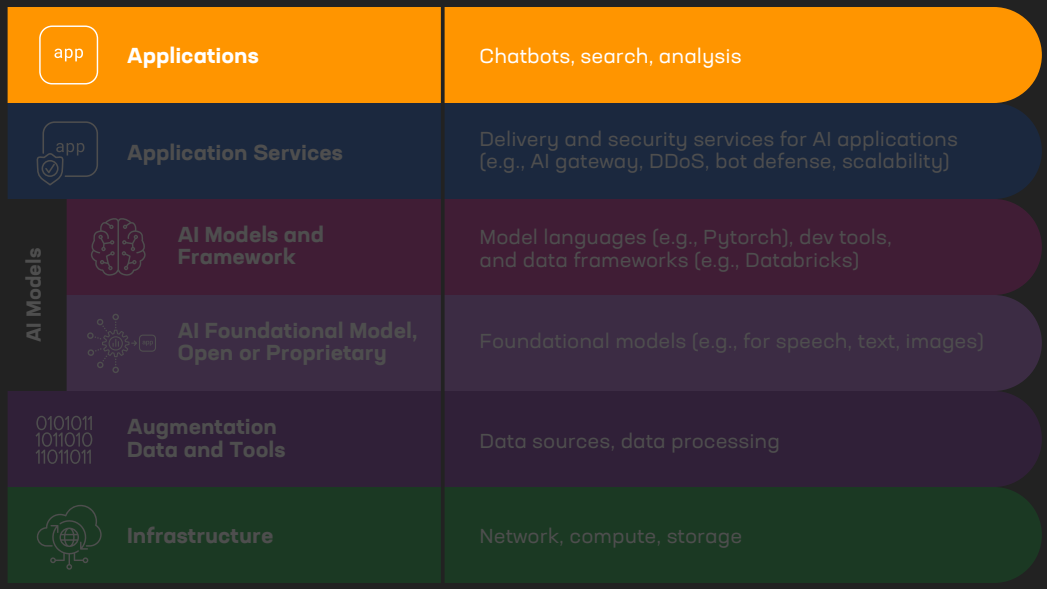

Conclusion

Mapping Out the Path Forward

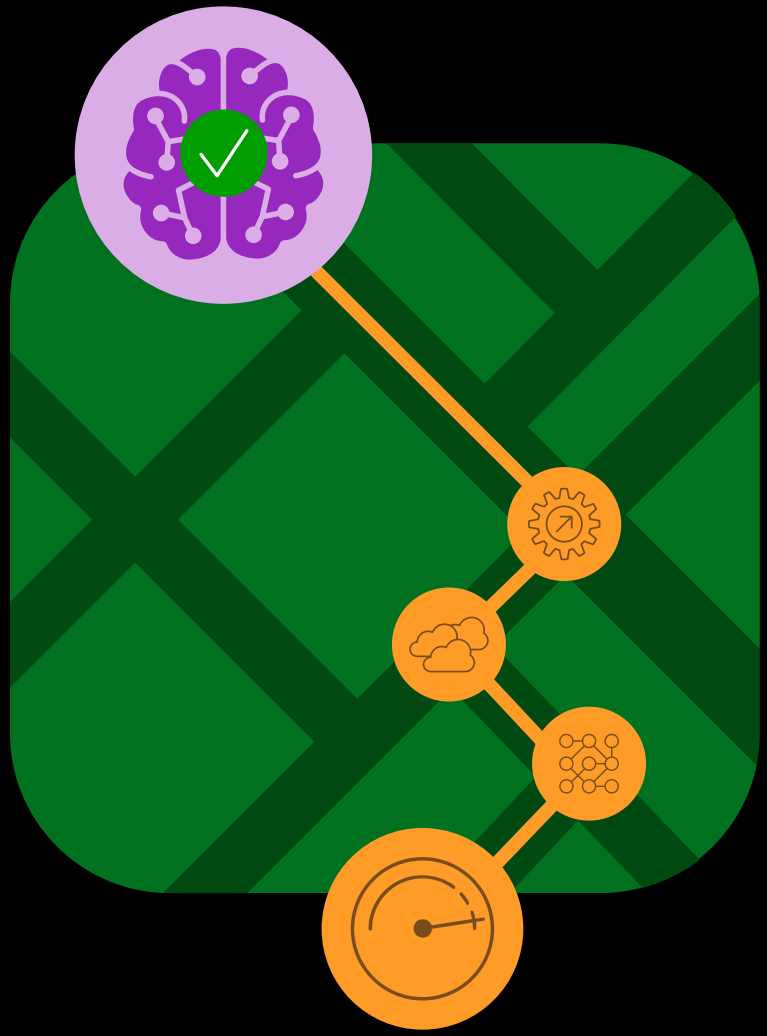

The promises of productivity and revolutionary business empowerment held forth by AI are rapidly moving organizations forward—sometimes faster than many can readily manage without skipping important steps. As a result, shiny chatbots or similar attractions can distract from a focus on more crucial tasks such as fundamental modernization, data integration, focused resource allocation, or less-sexy but higher priority measures like the automation required for more strategic AI implementations.

In particular, data strategies and quality need attention now. Organizations that lack a data strategy are already behind, and decision makers ignoring that inconvenience may be making the biggest mistake of their careers. In addition, a data strategy without data quality—or vice versa—amounts to a peanut butter sandwich without the jelly, rendering any AI achievements sticky and unsatisfying.

On the other hand, as AI deployments scale and grow more mature, app services and related technologies can support the whole stack in four obvious and not-so-obvious ways:

First, traffic management services such as high-performance load balancing can ensure the infrastructure tier operates efficiently as the AI models supported by that infrastructure ingest and process many billions of parameters.

> "Our organization prioritizes AI-driven product innovation as a core strategy to stay competitive and meet evolving market demands. We focus on leveraging AI technologies to enhance product features, automate processes, and deliver personalized experiences to our customers."

CIO, financial services organization ($200-500 million in annual revenue)

Next, in addition to protections such as API security for automatically discovering and securing AI training and inference endpoints, security services that prevent data leakage or theft can protect the data tier while defeating other threats to successful AI implementation.

Multicloud networking can deliver the high-throughput connectivity needed to feed, manage, and access the models, aiding data integration as well as operational ease. Such connectivity can help tame the hybrid deployments likely for AI models.

Finally, AI models running at the far edge also need to be integrated into the estates, which both app services and multicloud networking can help perform.

App Services Are Critical for Mature AI App Deployments

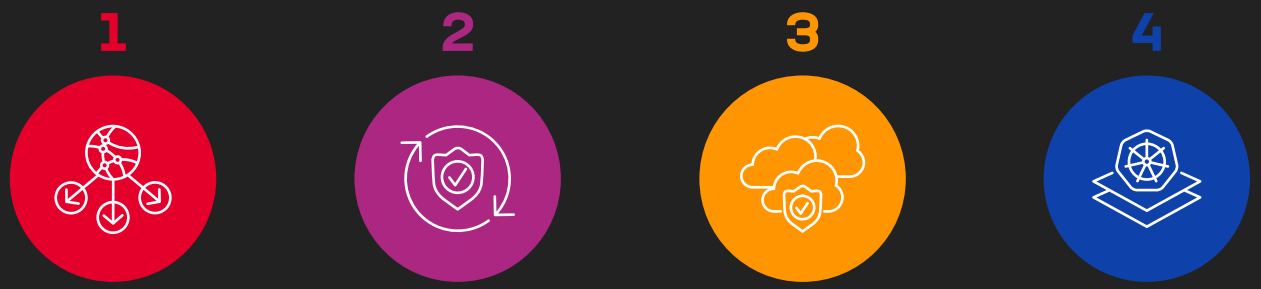

All these proven technologies are available today for organizations ready to take their next AI steps. Stragglers seeking to catch up, as well as digital leaders hoping to accelerate their progress, would do well to consider the emerging AI stack and the recommendations in this report as a map. Thus equipped, they can better plan how to capture the AI lightning that will drive the business forward while avoiding the insufficient compute capacity, data disasters, model mistakes, security failures, or other hazards that almost certainly lurk in the shadows.

The final phase of the digital transformation journey may be endless, but with a data strategy and reliable data in hand and each tier of the AI stack well plotted, the trip to the digital economy of the future can be as rewarding for customers and the business as any finite destination would be.

For more of the latest data about AI and many other trends in app strategy, read the full 2024 F5 State of Application Strategy Report.

About this report

This report compiles data and analysis from both the tenth annual F5 State of Application Strategy survey, which explores the current interests of more than 700 IT decision makers across industries worldwide, and more in-depth research specifically on AI technology and strategies. The in-depth research targeted 75 North American and United Kingdom decision makers who are specifically responsible for their organizations' AI strategies and implementations. As with the global survey, focus survey respondents were drawn from a variety of industries and companies of all sizes. Nearly half—45%—identified their organizations as providers of digital services, from ecommerce and payment processing to SaaS and digital gaming.

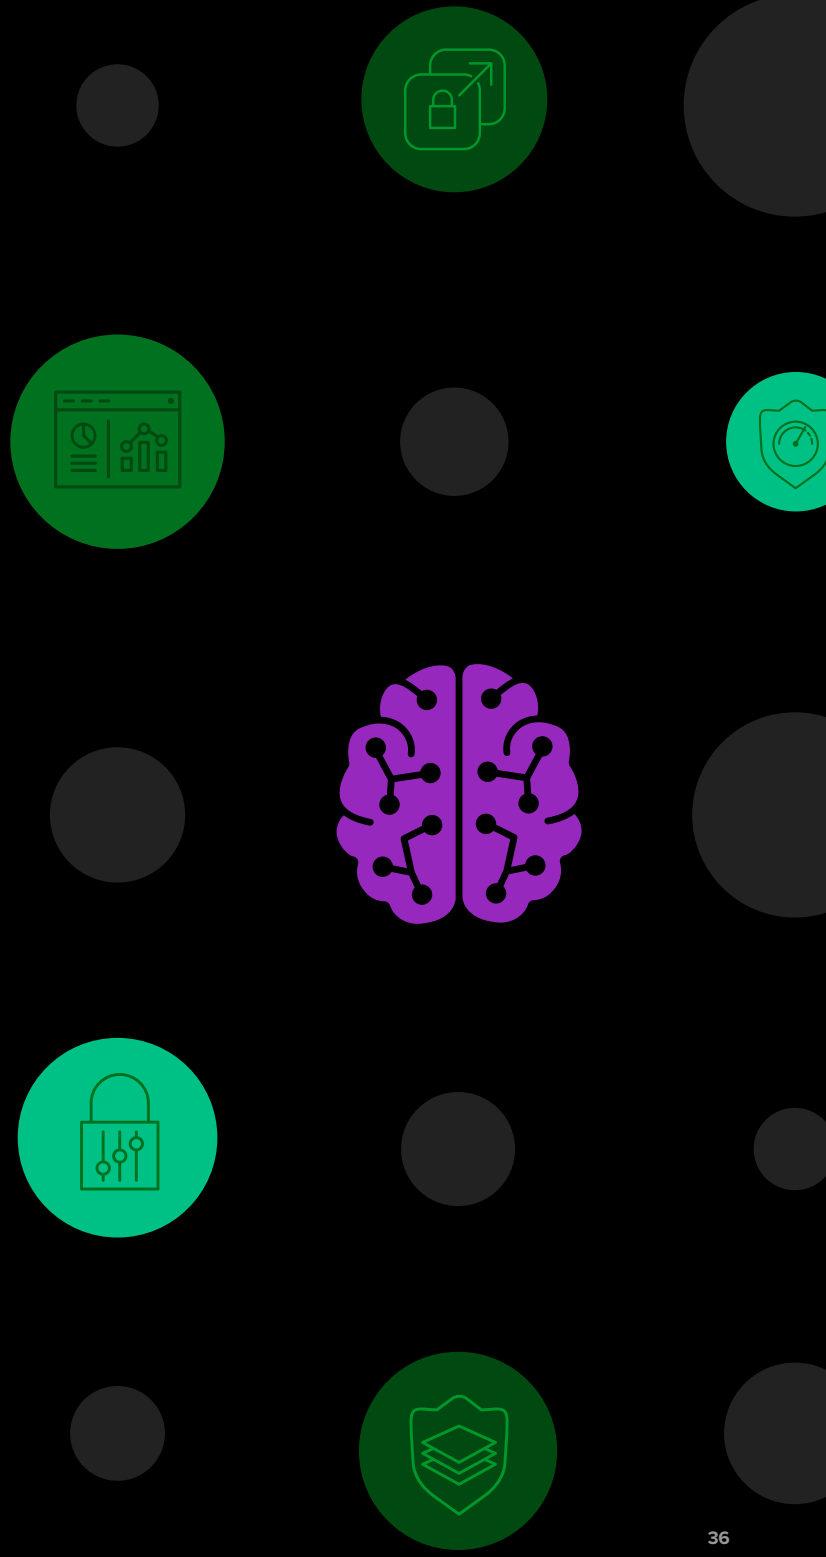

ABOUT F5

F5 is a multicloud application services and security company committed to bringing a better digital world to life. F5 partners with the world's largest, most advanced organizations to secure every app and API—on premises, in the cloud, or at the edge. F5 enables organizations to continuously stay ahead of threats while providing exceptional, secure digital experiences for their customers.

For more information, go to f5.com. (NASDAQ: FFIV).