dynatrace

RESEARCH REPORT 2026

The State of Observability

Making AI reliable: How observability must establish the evidence for trusted outcomes

What’s inside

INTRODUCTION

AI ambition is outpacing operational proof 3

CHAPTER ONE

AI is resetting the observability agenda 5

CHAPTER TWO

Generative and agentic AI compound complexity and dominate visibility gaps 15

CHAPTER THREE

AI observability needs are outpacing execution 25

CHAPTER FOUR

Autonomous operations are moving from ambition to scale 31

CHAPTER FIVE

Runtime context is reshaping protection against AI-accelerated security threats 41

CHAPTER SIX

Open standards create the foundation for AI at scale 50

CONCLUSION

Reliable AI requires more from observability 57

INDUSTRY PERSPECTIVES

Observations and outlooks 58

APPENDIX A

Firmographics and demographics 61

APPENDIX B

Regional signals 63

APPENDIX C

The influence of company size 65

APPENDIX D

Industry-specific priorities 66

AI ambition is outpacing operational proof

Observability must close the gap.

AI is fast becoming the most mission-critical software many organizations operate. But ambition is moving ahead of operational proof: organizations must demonstrate that AI behaves as intended and delivers reliable, secure, governed, cost-effective results. Observability must evolve accordingly, providing the comprehensive visibility, context, intelligence, and control required to operate AI confidently at enterprise scale.

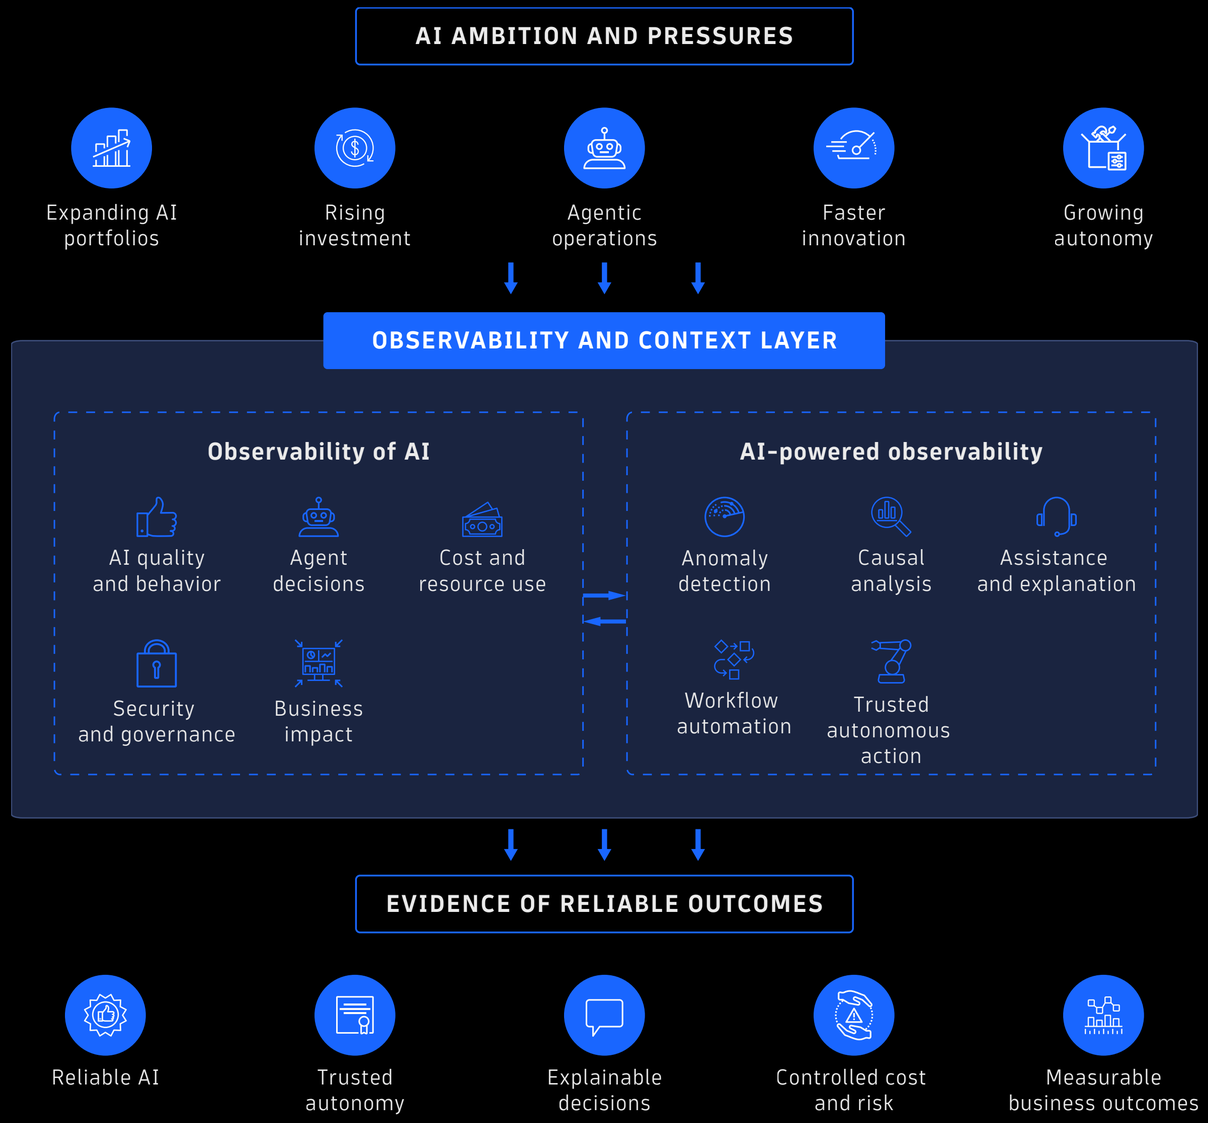
Observability must establish the evidence AI ambition requires

Observability should connect signals across the technology environment in real-time so organizations can understand, govern, and act on system behavior.

As AI gains traction, that mandate takes on two dimensions.

- **Observability of AI:** Uses continuous evaluation and operational evidence to verify AI behavior, quality, security, cost, and business impact.
- **AI-powered observability:** Uses causal analysis, assistance, automation, and agents to improve operations and advance autonomy.

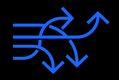
Together, these capabilities provide the operational proof required to run AI more reliably at scale.



Based on an August 2026 survey of 1,202 technology leaders at enterprise companies across 12 countries, the research explores what stands between AI ambition and operational proof:



Investment
How are priorities and budgets keeping pace with AI ambitions?



Complexity
How is AI complicating already fragmented technology environments?



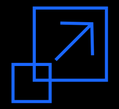
AI observability
What evidence proves AI is working as intended?



Autonomous operations
What must organizations trust and validate before AI takes greater responsibility?



Security
How are AI-accelerated threats changing detection, prioritization, and response?



Scale
How are open standards providing continuity across technologies and lifecycle stages?

The survey explores where AI ambition is advancing, where proof falls short, and what observability must deliver next.

METHODOLOGY

Survey conducted: August 10 – August 24, 2026

Total respondents: 1,202

Countries: United States, Brazil, Mexico, United Kingdom, Germany, France, Italy, UAE, Saudi Arabia, India, Japan, Australia

Qualifying criteria: C-suite (CIO, COO, CTO, CISO, CAIO, etc.) high-level exec (VP, SVP, EVP, etc.), mid-level-exec (Director, Senior Director, Department head, etc.) for organizations with annual revenue of $100M or more

Responsible for: Observability and/or AI observability, application security, resource optimization and sustainability, AI governance, software engineering, AI engineering, data-driven rules-based automation, or AI-driven autonomous operations

Question bases vary by topic: All 1,202 respondents completed the core survey, while specialized questions on autonomous operations, security, and OpenTelemetry were presented only to respondents who met relevant screening or routing criteria.

Research partner: Qualtrics with partner Y2 Analytics

Margin of error: ±2.8 with 95% confidence

CHAPTER ONE

AI is resetting the observability agenda

AI is now embedded across enterprise technology, resetting what organizations expect from observability.

Priorities and buying criteria are shifting, budgets are rising, and ambitions are advancing toward autonomous operations. Observability now has a broader mandate: support growing AI adoption while demonstrating the reliability and business outcomes organizations depend on.

49%

run 6-25 production AI projects (up from 42% a year ago)

78%

expect observability budgets to increase over the next year

57%

selected AI-powered capabilities as important when choosing an observability solution

44%

target fully autonomous operations within five years, while only 12% are there now

Proof points



AI adoption has become the enterprise baseline

99% use AI in their operations and are expanding production portfolios while predominantly building on existing models, APIs, and AI capabilities.



AI is resetting observability priorities and investment

AI-powered capabilities lead solution selection criteria, while AI adoption drives changes to observability budgets.



Observability supports critical business outcomes

Reliability and resolving issues lead, while reported improvements span all nine measured business outcomes.

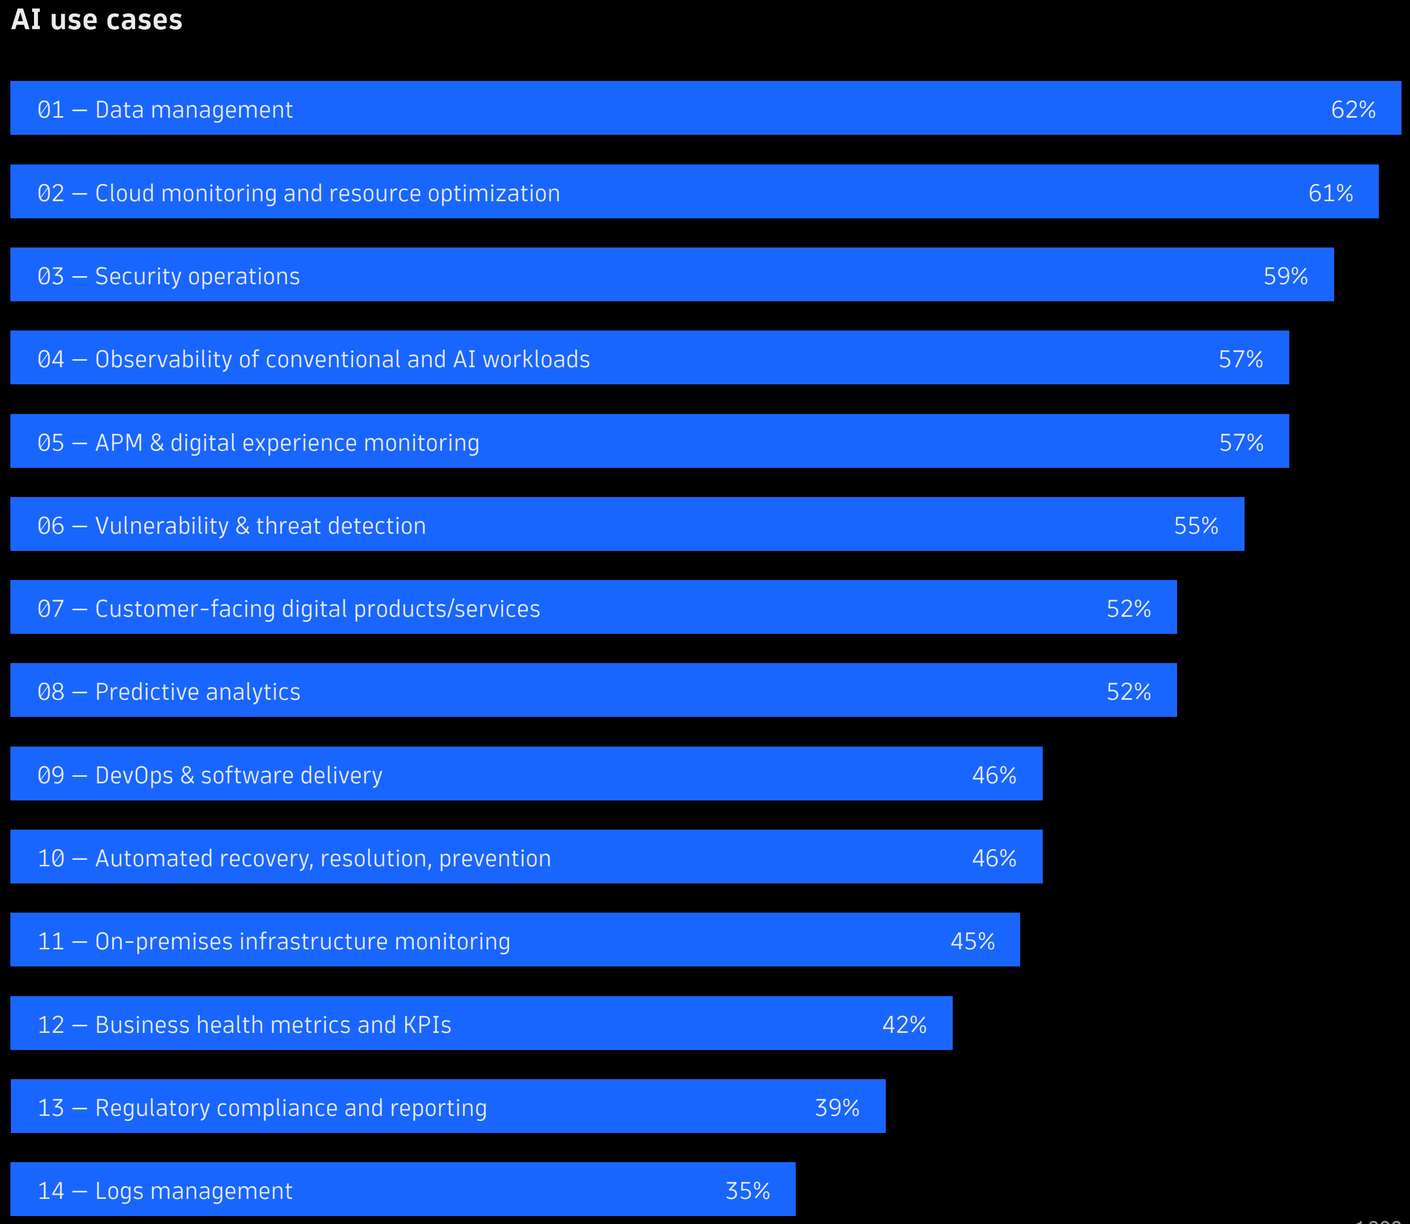


Autonomy is the next maturity frontier

Supervised autonomy leads today, but organizations are setting substantially higher targets for fully autonomous operations.

Agentic and generative AI are embedded across enterprise technology

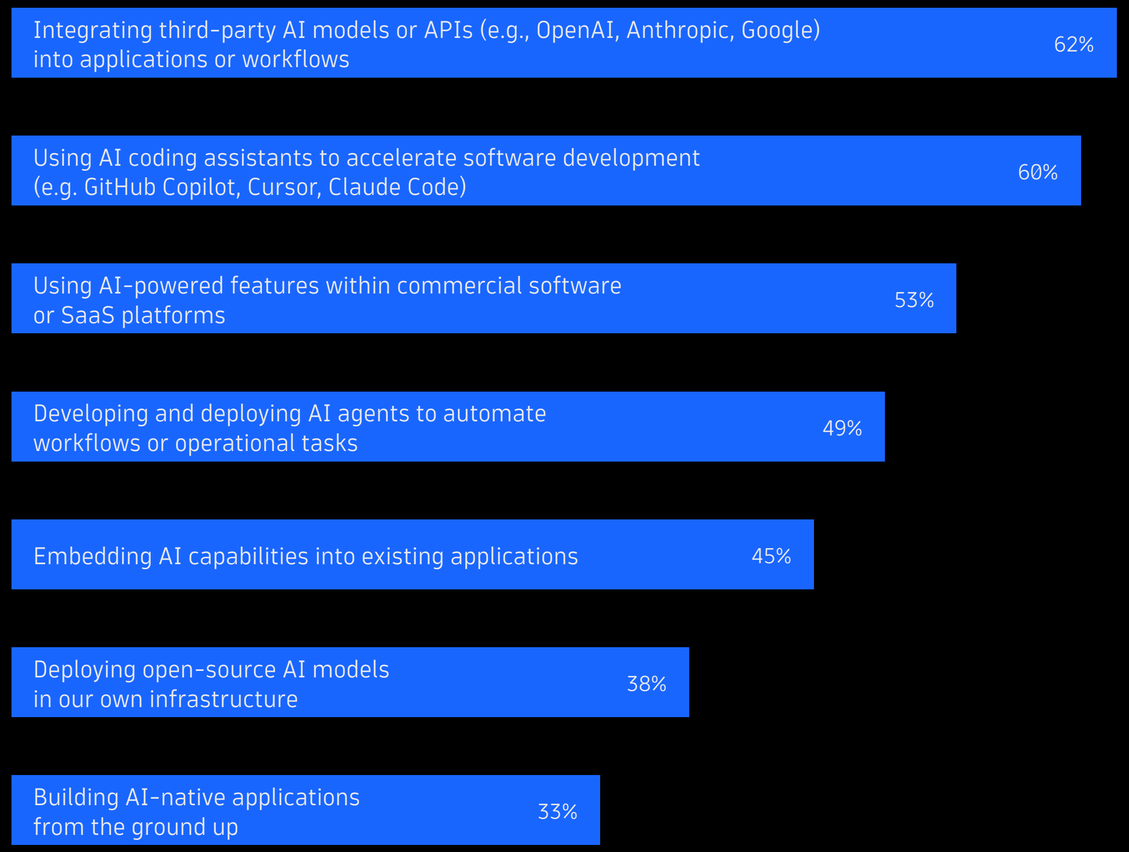
Respondents report near universal AI adoption, spanning data management, cloud operations, security, observability, and other functions.



n=1,202

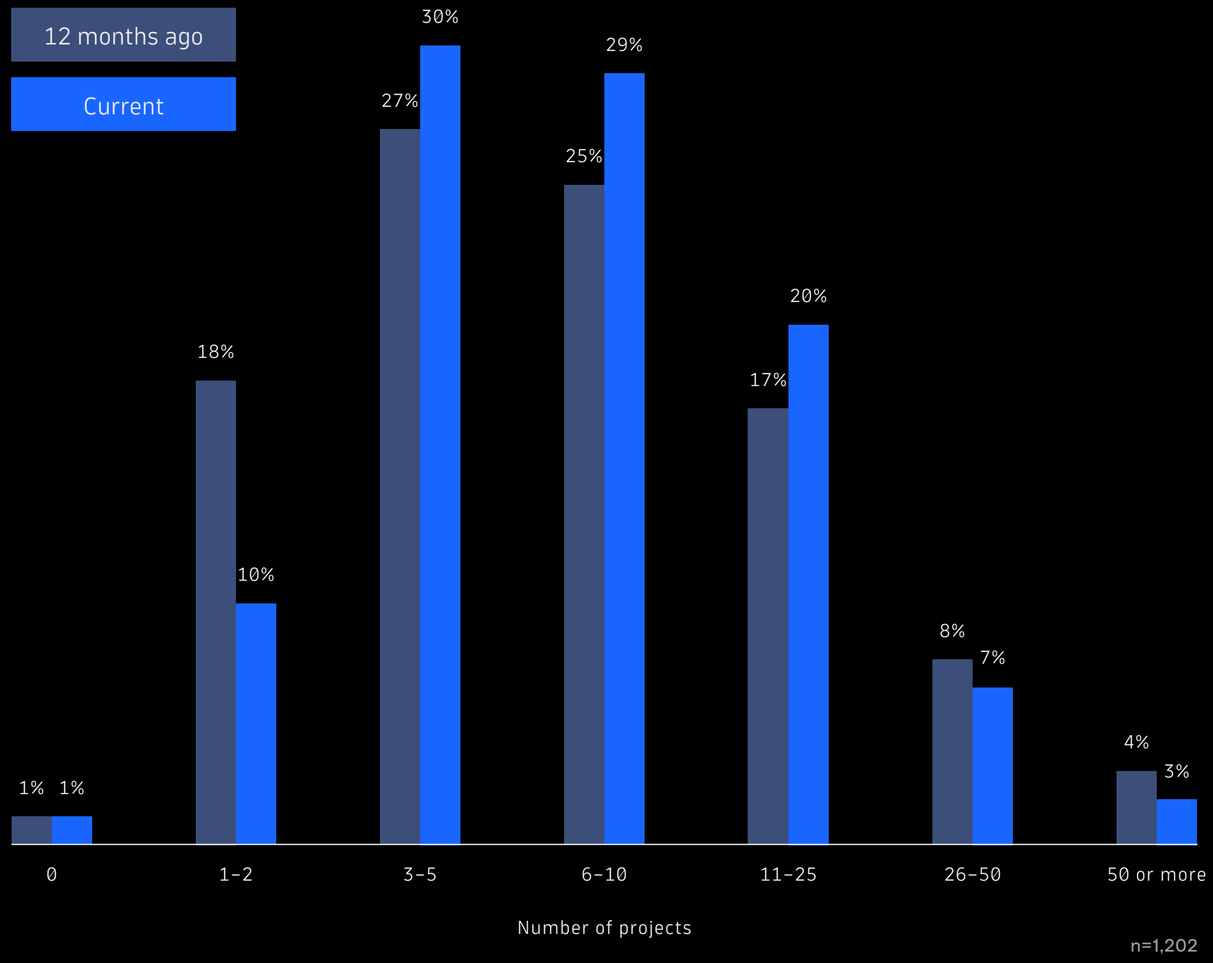
Most organizations are building AI portfolios on existing models, APIs, and development capabilities, and agent deployments rather than starting from scratch.

AI deployment models



And these AI portfolios are expanding: 49% now report running 6–25 projects, up from 42% a year ago.

AI projects in production



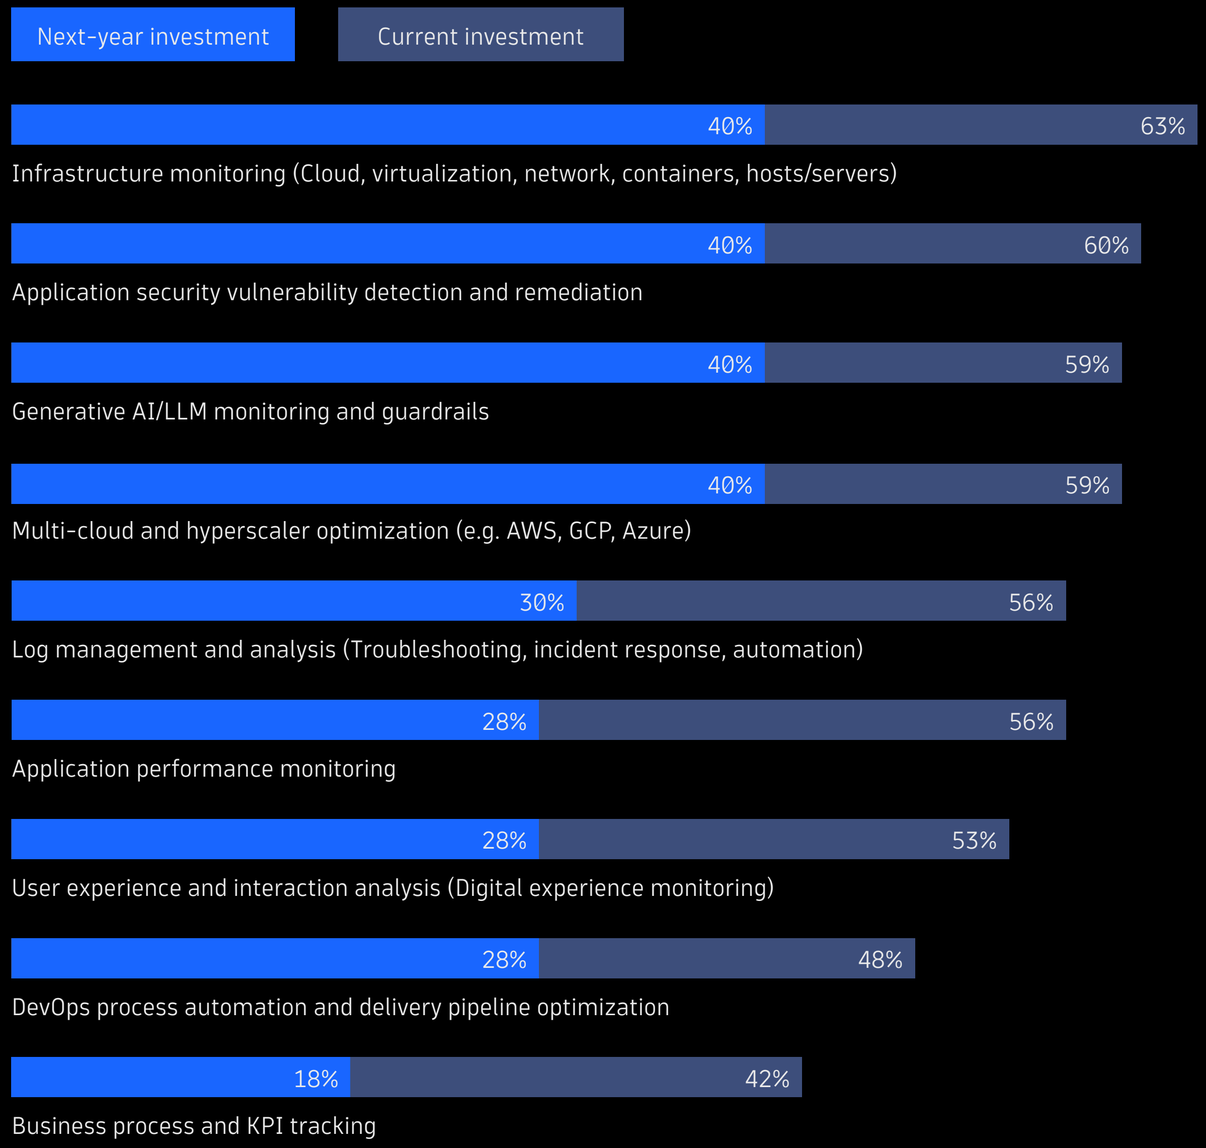
n=1,202

AI workloads and agentic operations are reshaping observability priorities and buying criteria

Organizations report that observability priorities now span conventional infrastructure, application security, multi-cloud environments, and generative AI.

Next-year investment is distributed evenly across those leading use cases at 40%, suggesting AI is expanding the observability mandate rather than displacing established priorities.

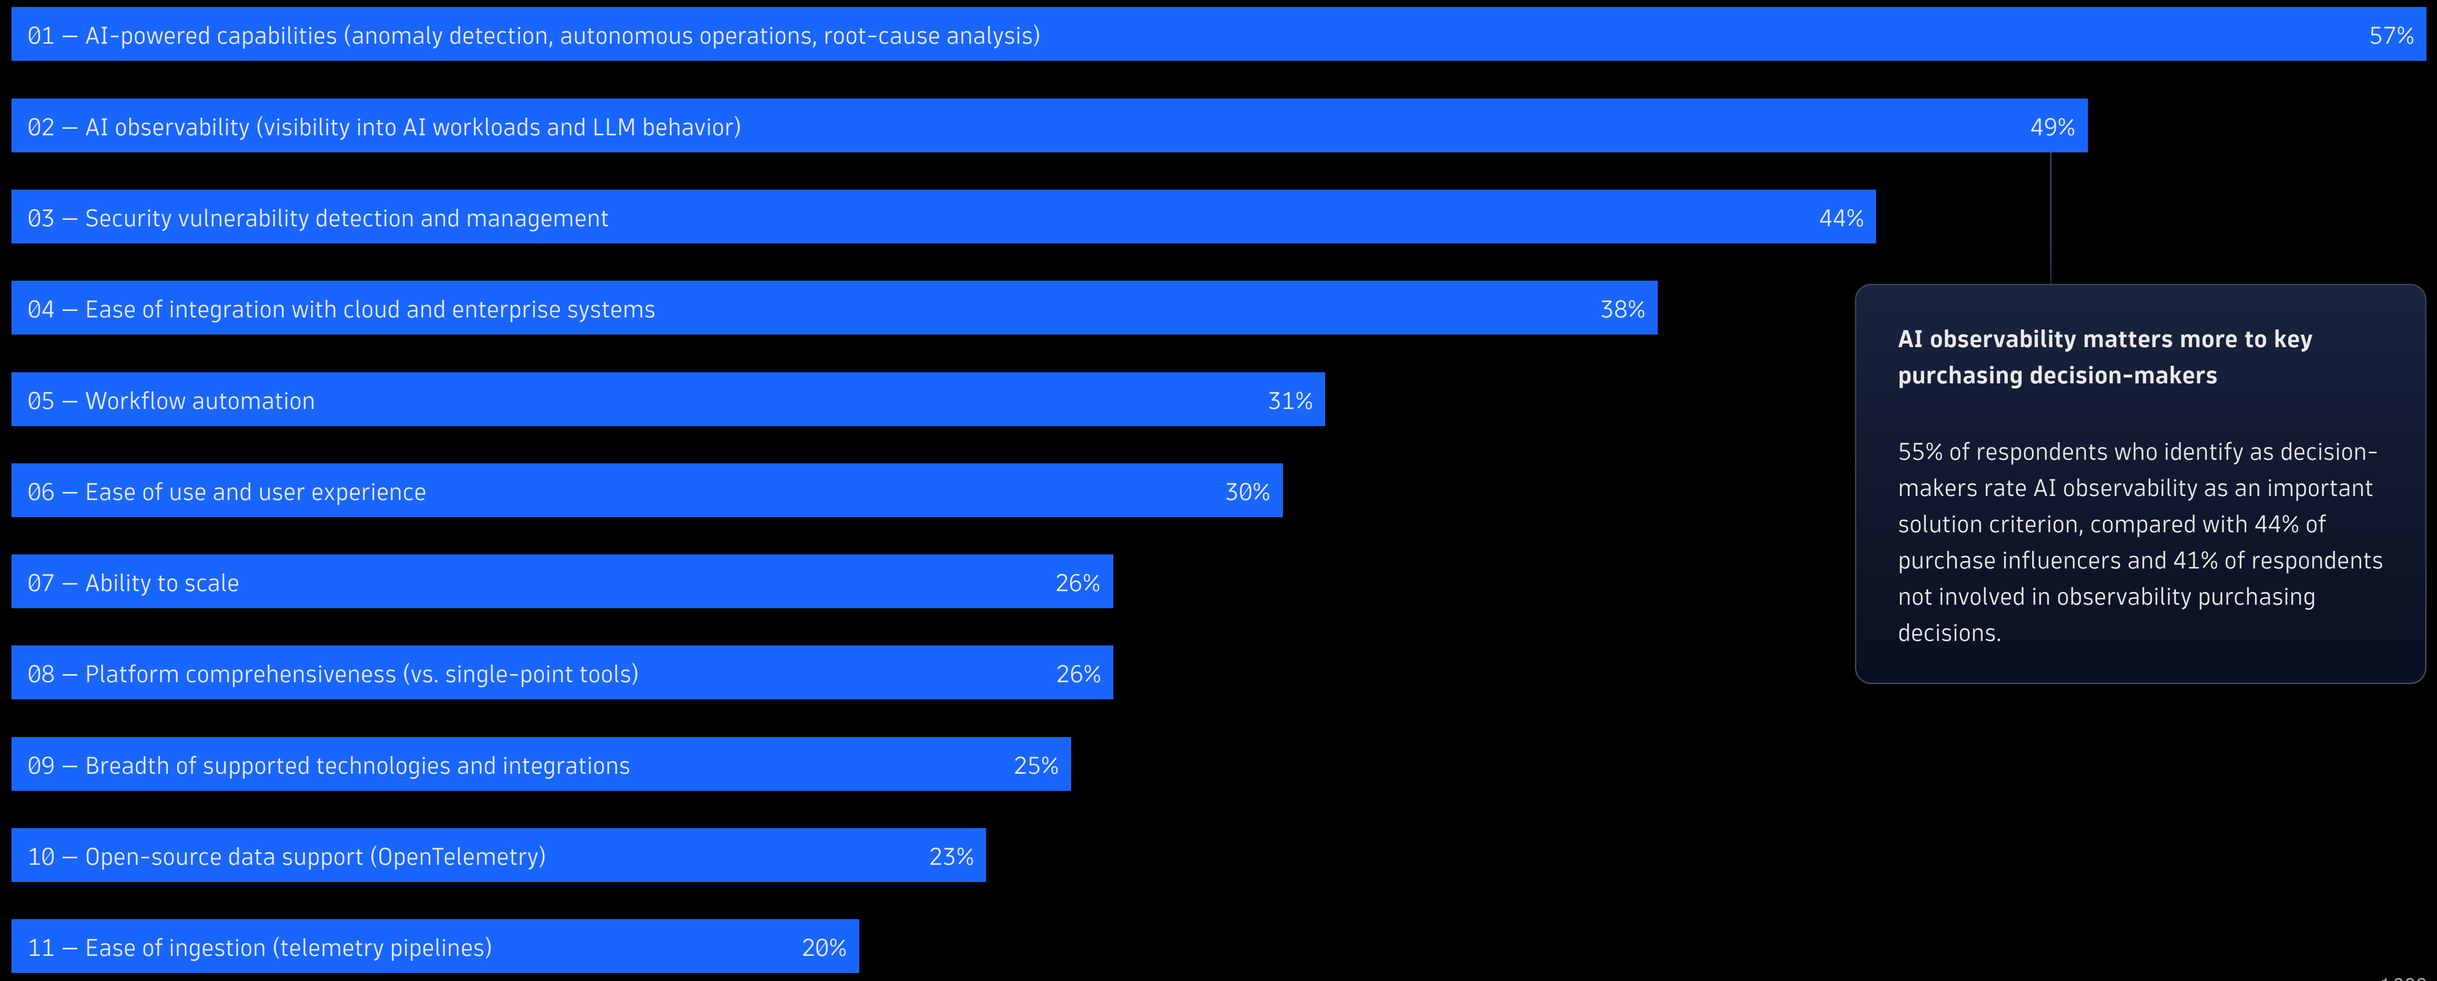
Observability uses and investment priorities



n=1,202

AI also leads solution criteria: 57% prioritize AI-powered capabilities and 49% AI observability, followed by security vulnerability management (44%). Observability priorities increasingly reflect the need to understand and manage conventional and AI workloads together.

Observability solution criteria



AI observability matters more to key purchasing decision-makers

55% of respondents who identify as decision-makers rate AI observability as an important solution criterion, compared with 44% of purchase influencers and 41% of respondents not involved in observability purchasing decisions.

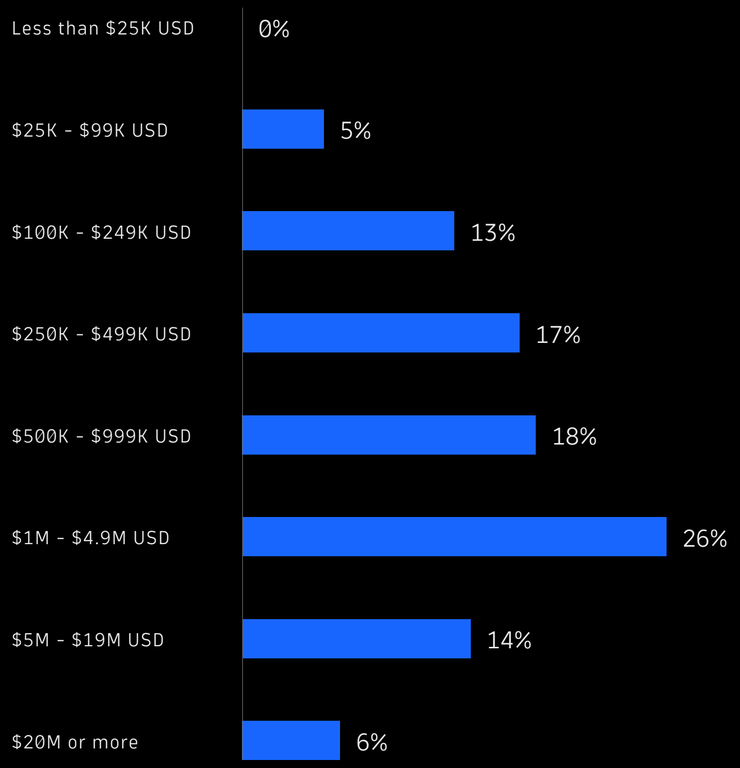
n= 1,202

AI adoption is driving the next wave of observability investment

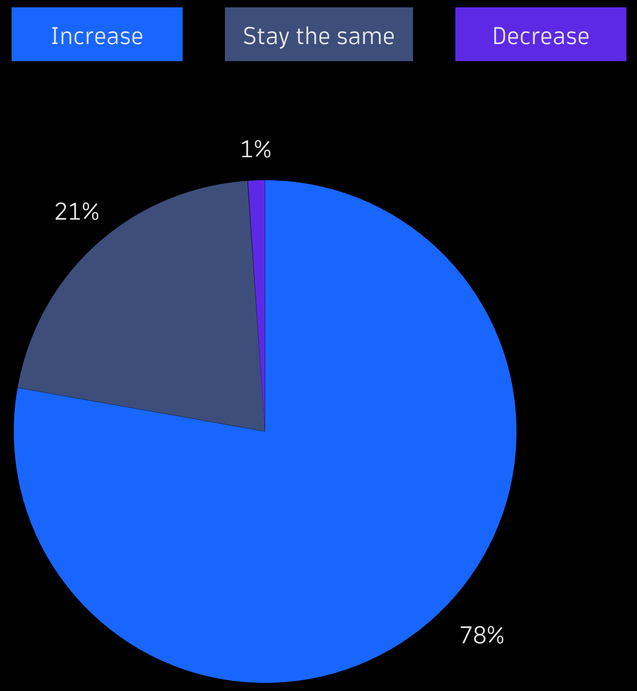
AI adoption is expanding the observability mandate and reshaping the investment required to support it.

Nearly half of organizations (46%) spend $1 million or more annually on observability. Investment is set to rise further: 78% expect their observability budgets to increase next year, and among them, 96% say AI adoption is at least a moderate driver of that increase.

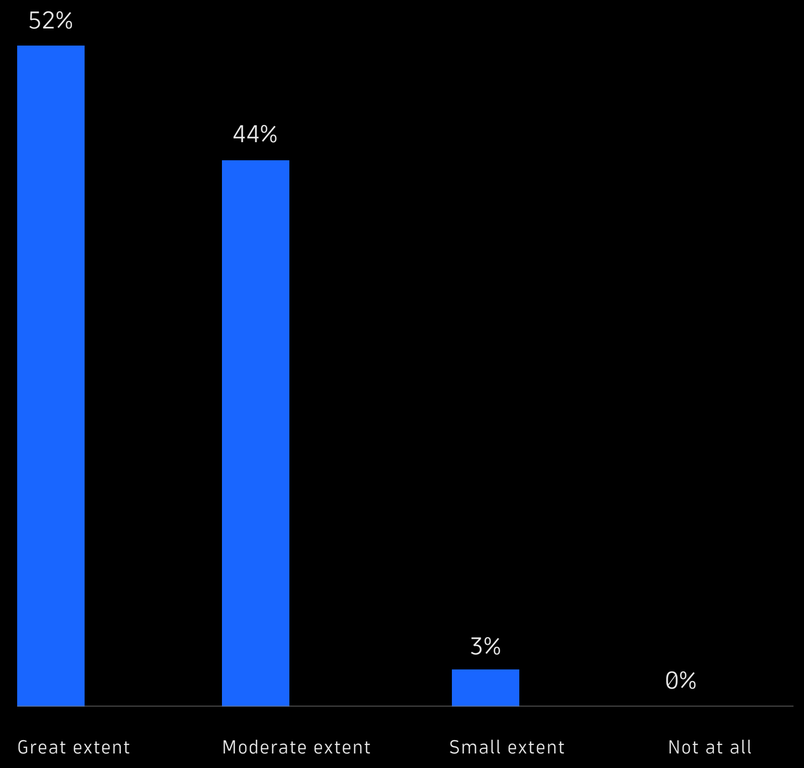
Current annual observability budgets



Observability budget change in the next year



AI adoption driving observability budget



n=1,202

Organizations report broad business value from observability, led by operational outcomes

Organizations report improvements from observability across every business outcome measured, with each cited by at least 42% of organizations.

Faster issue detection and resolution and improved reliability rise to the top. The breadth of reported outcomes demonstrates the reach of observability beyond IT operations, and reinforces its core value: keeping complex systems reliable and problems from becoming business disruptions.

Most improved business outcomes by observability



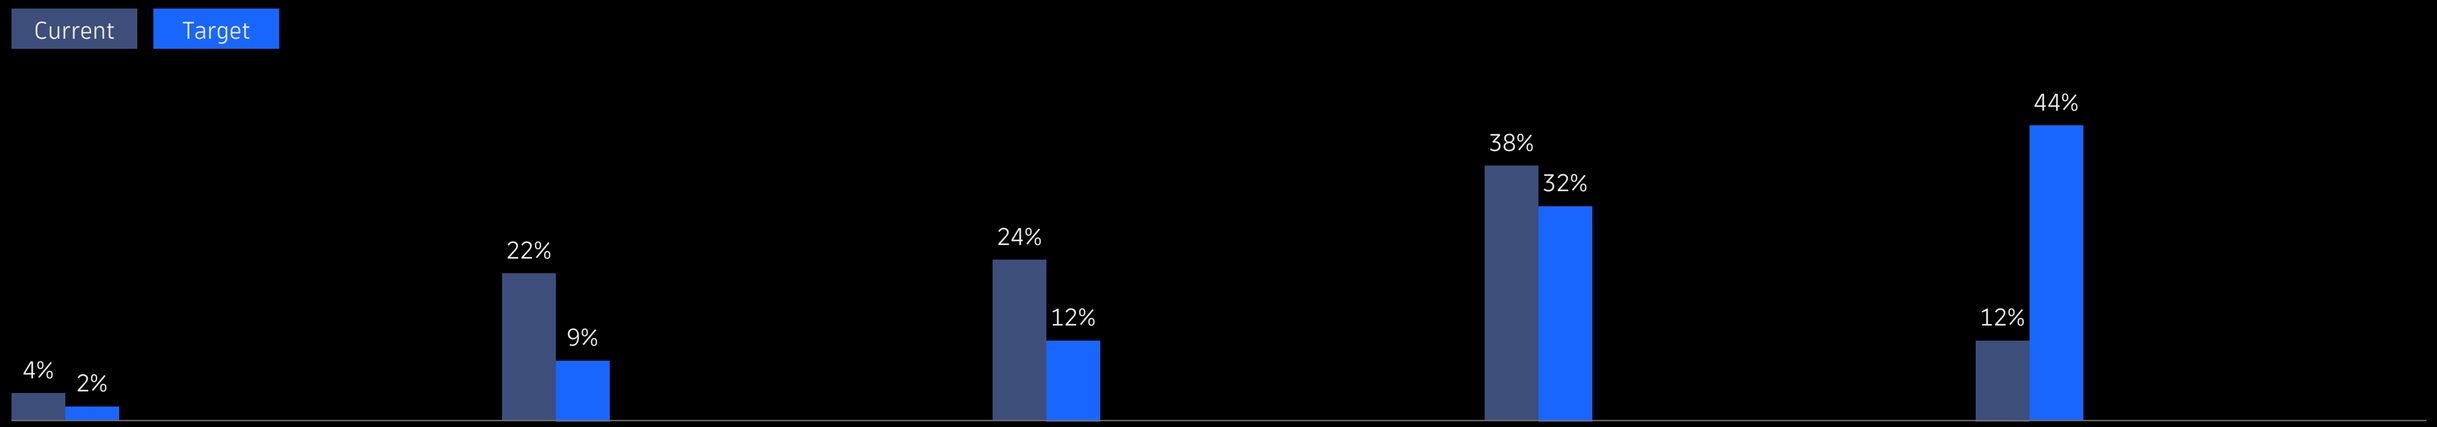
n=1,202

Supervised autonomy leads today, but full autonomy leads five-year ambitions

Thirty-eight percent of organizations report reaching supervised autonomy, where AI executes within human-defined guardrails, compared with 12% at fully autonomous operations, where agents can act, investigate, predict, and optimize on their own or on human approval. Another 24% operate predictively and 22% use AI-powered observability, while just 4% remain primarily reactive.

Yet 44% aim for fully autonomous operations within five years. The finding reflects organizational ambition, not which workloads respondents consider appropriate for full autonomy. This ambition rises with organizational scale: half of organizations with annual revenues of $10 billion or more target full autonomy, compared with 29% of those with revenues of $100–499 million.

Observability maturity



STAGE 1

Reactive monitoring

Manual dashboards and basic metrics; troubleshooting is reactive and human-driven

STAGE 2

Observability with AI-enabled capabilities

AI correlates telemetry across sources to identify root cause and track SLOs; humans still make all decisions

STAGE 3

Predictive operations

Machine learning predicts failures before users are impacted and recommends fixes, cutting manual triage and MTTR

STAGE 4

Supervised autonomy

AI executes within human-defined guardrails and approval gates; humans set goals, not steps

STAGE 5

Fully autonomous operations

Agents can act, investigate, predict, and optimize on their own or on human approval

n=1,202

THE TAKEAWAY

Build observability around business outcomes

As AI adoption grows, observability should connect technology and AI performance with evidence of business outcomes leaders need their investments to deliver.

- **Measure what AI changes.** Connect technology performance to revenue, customer experience, reliability, and operational resilience.
- **Invest in shared context.** Connect AI behavior, applications, infrastructure, security, and business data to move from “it’s up,” to “it’s working,” to “it’s worth it.”
- **Manage AI economics.** Balance performance, resource costs, and sustainability as usage expands.
- **Build for autonomy selectively.** Increase AI decision authority on appropriate workloads as the system demonstrates reliability, explainability, and adherence to governance.

AI adoption is resetting what organizations expect from observability. But turning those ambitions into reliable results means confronting the complexity AI compounds across already fragmented environments.

CHAPTER TWO

Generative and agentic AI compound complexity and dominate visibility gaps

AI is adding new complexity to technology environments already stretched across multiple clouds and modern application architectures.

It's also increasing skills requirements, creating visibility gaps, and driving resource costs higher. Organizations now face a dual challenge: managing AI-driven complexity while gaining the visibility and expertise needed to control its operational and business impact.

43%

name AI expansion as the leading complexity driver

47%

say AI has increased the skills teams need to manage observability

41–45%

cite AI-related concerns as top visibility gaps

81%

say AI workloads have increased compute and infrastructure spending

Proof points



AI compounds complexity on two fronts

Technology environments are more complex and require greater observability expertise to manage them effectively.



Complexity builds on an already fragmented foundation

AI workloads further complicate multi-cloud environments and numerous observability tools.



AI-related concerns now dominate visibility gaps

The top five observability gaps are clustered between 41% and 45%, while non-AI-specific gaps rank at 36% or below.



More data can’t solve the visibility problem

AI and conventional systems create data management challenges around integration, instrumentation, quality, and governance that make reliable context harder to establish.



AI complexity has a growing resource cost

Higher compute and energy demands collide with limited visibility into consumption and attribution.

AI is compounding complexity across the operating environment

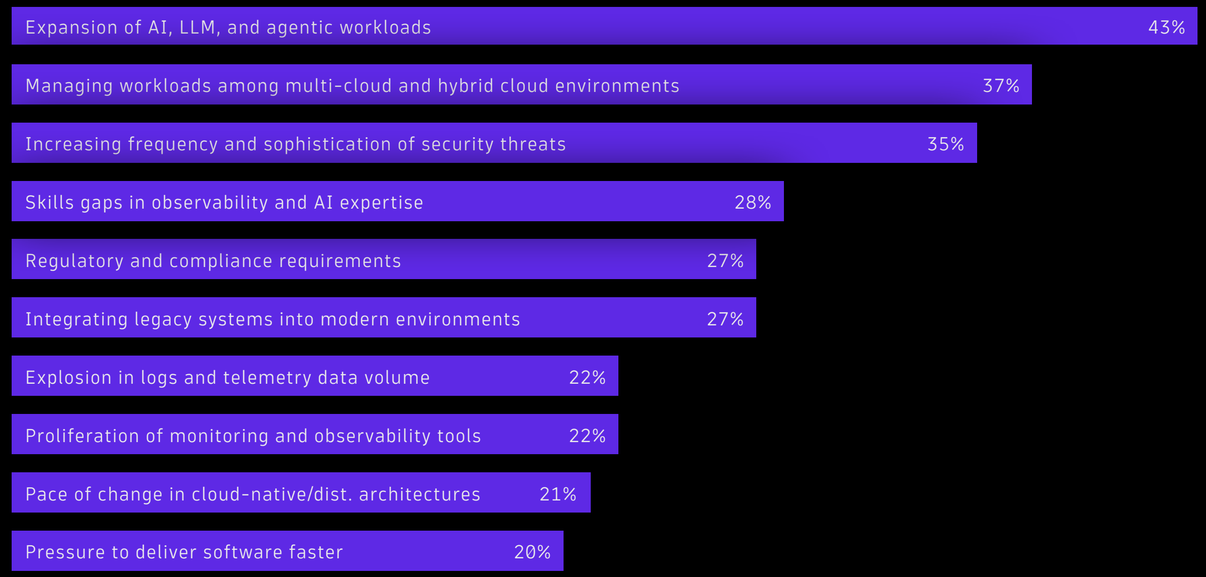
AI, LLM, and agentic workload expansion now leads reported sources of technology complexity, ahead of multi-cloud environments and security threats.

The pattern suggests AI is adding to the complexity, cloud, and security pressures organizations are already managing.

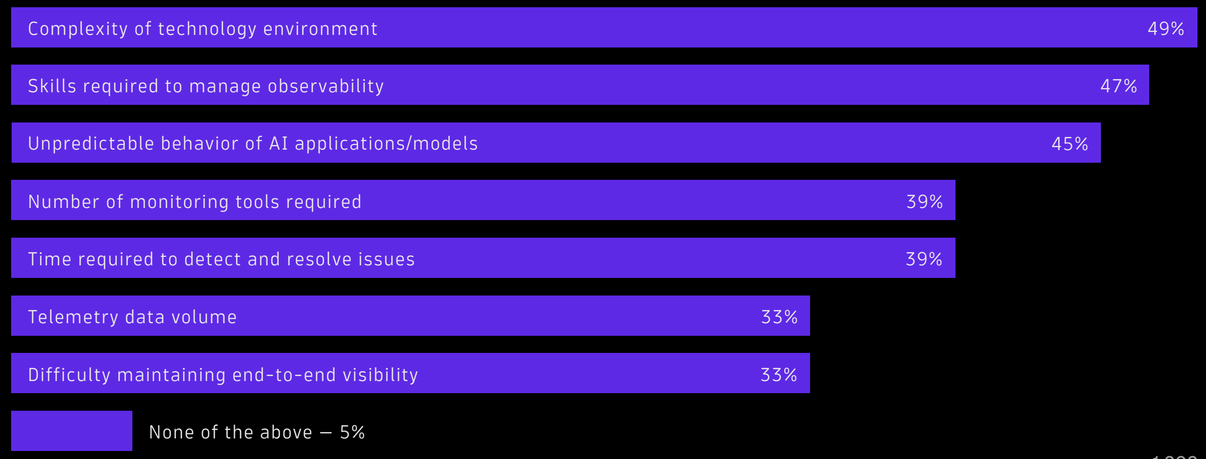
The effects are felt across operations. Nearly half say AI workloads have increased the complexity of their technology environment and the skills required to manage observability, while 45% cite unpredictable AI behavior.

The findings point to a dual challenge: AI is making technology environments more complex while increasing the expertise required to understand and manage them.

Sources of technology environment complexity



Workloads increasing in complexity due to AI operations



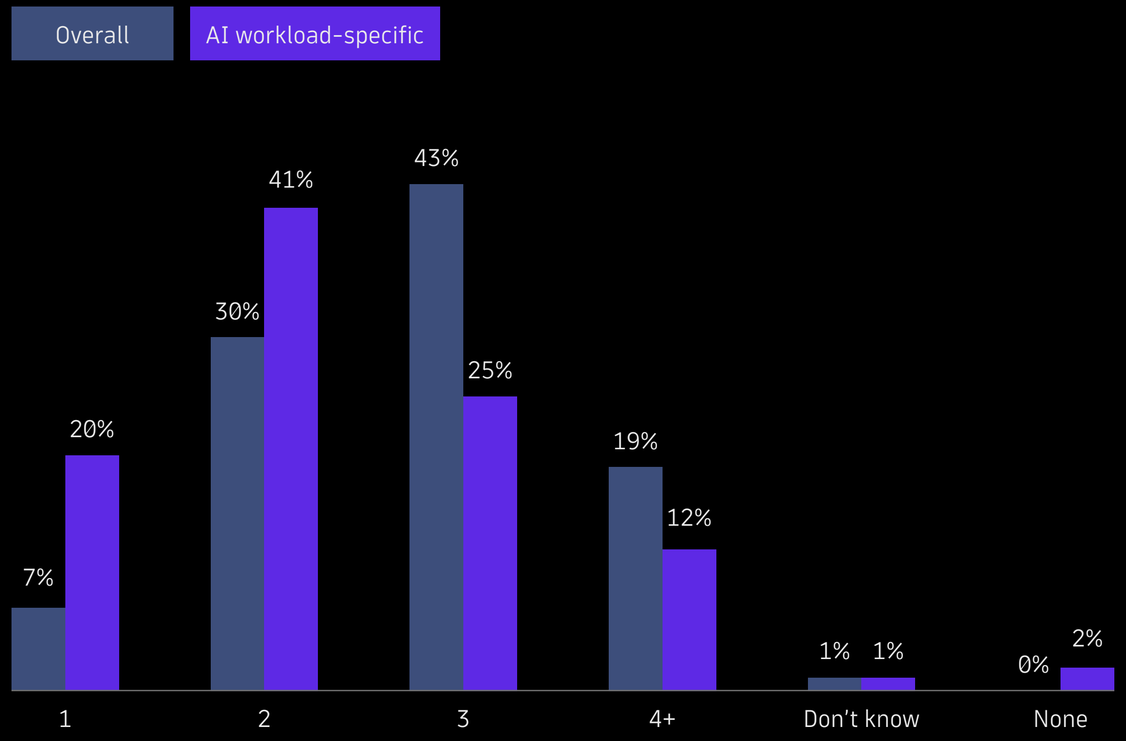
n=1,202

AI initiatives add to an already fragmented technology estate

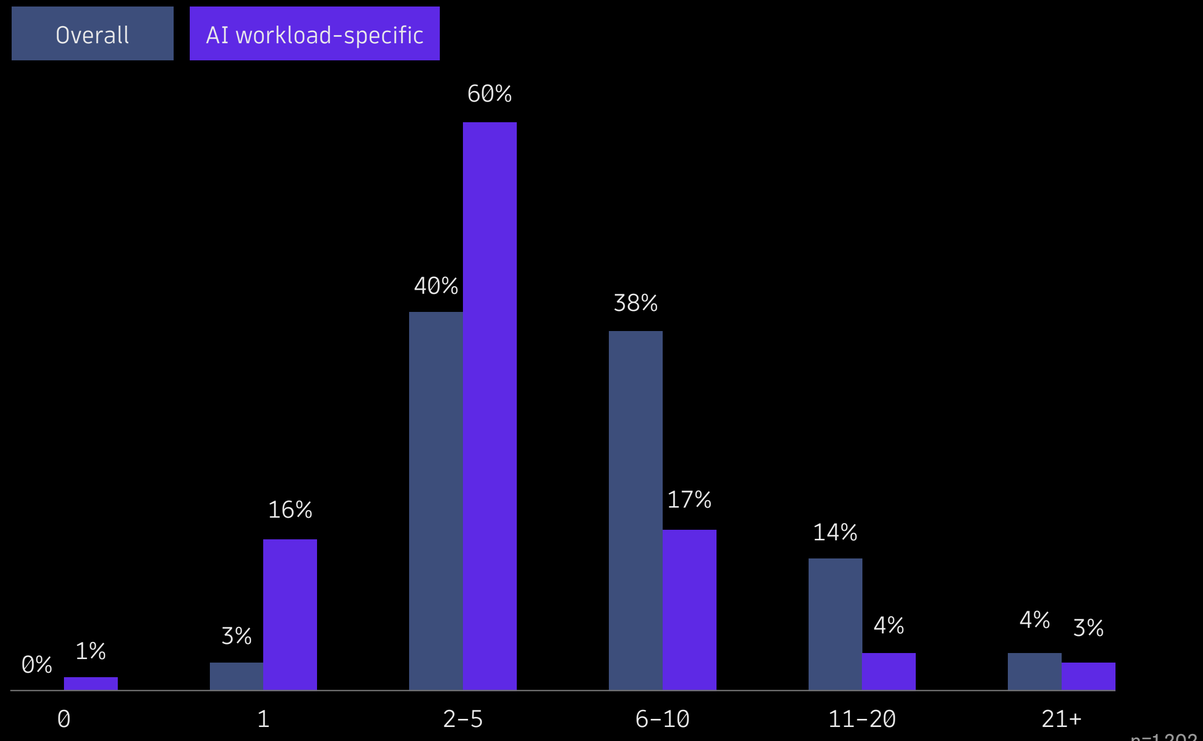
Organizations already operate across highly distributed environments: 62% use three or more cloud providers, while 56% use six or more observability tools.

AI workloads have a narrower footprint, with 37% using three or more cloud providers and 24% using six or more observability tools. As AI adoption expands, organizations must integrate these workloads into an already fragmented cloud and observability landscape.

Number of cloud providers



Number of observability tools or platforms



n=1,202

Rapid technology changes are testing observability readiness

Keeping pace with technology change emerges as the leading reported observability challenge, but the pressures span integration, multi-cloud visibility, and skills.

Together, they show that observability readiness depends on both the technology estate and the teams responsible for it.

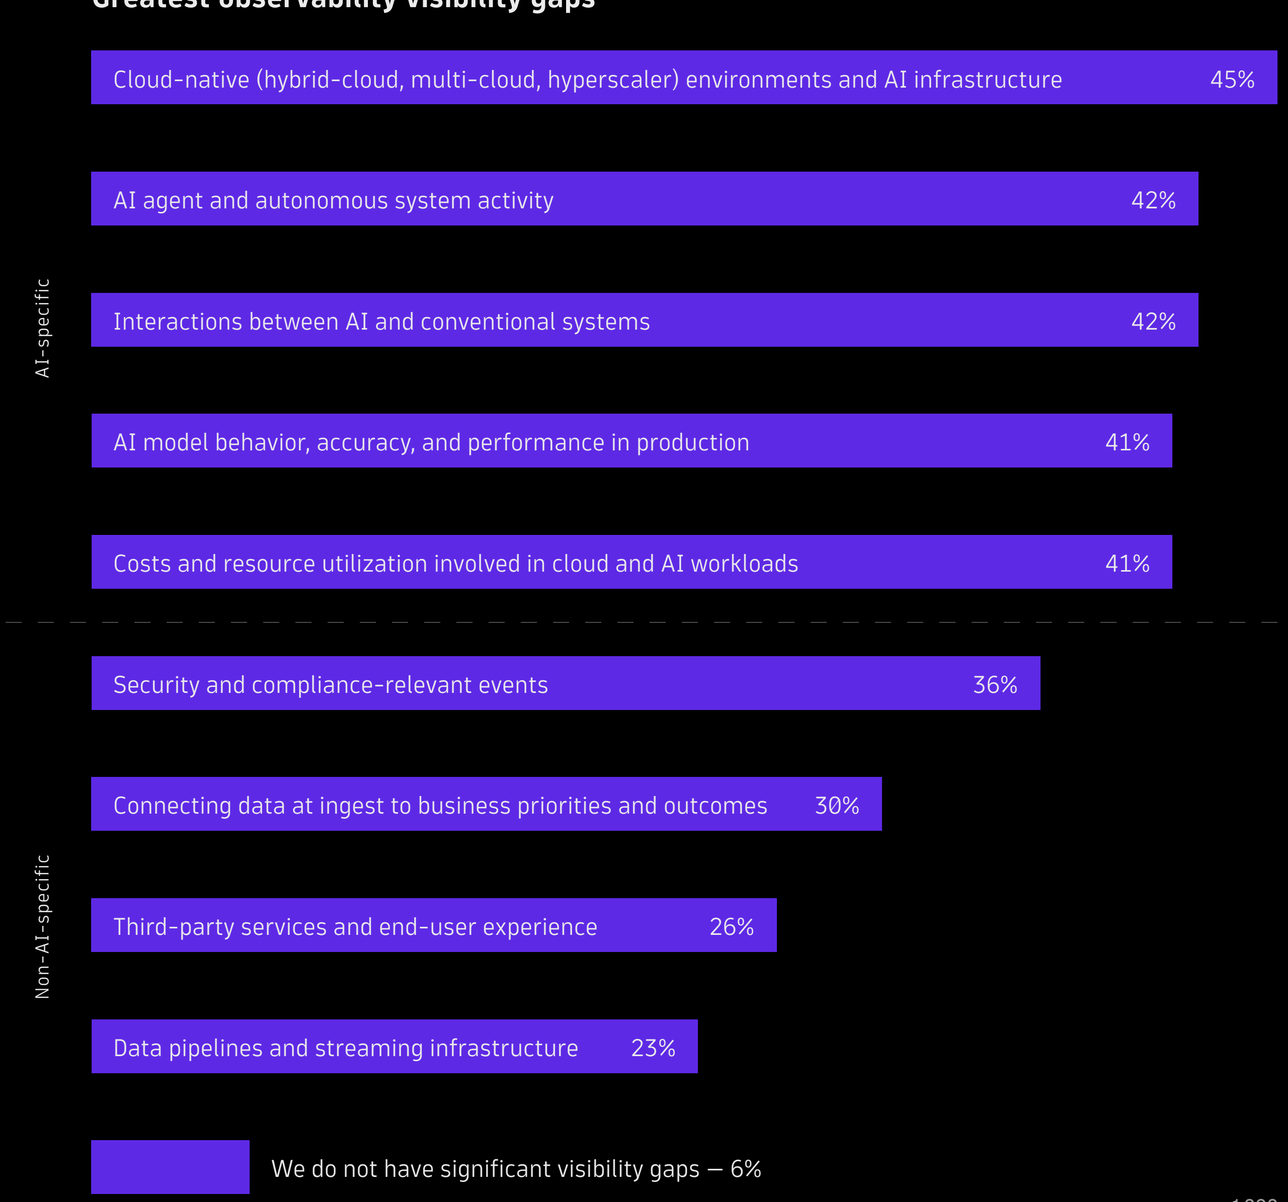


n=1,202

Visibility gaps are dominated by AI workloads

That growing complexity is showing up as blind spots. Only 6% report no significant visibility gaps, and AI-related areas dominate the top five deficits.

All AI-related visibility gaps rank above 40%, clustered tightly between 41% and 45%. By comparison, non-AI-specific gaps rank at 36% or below, pointing to AI infrastructure, agents, interactions, and costs as the new center of visibility challenges, making their behavior harder to verify across systems.



The data challenge is broader than volume

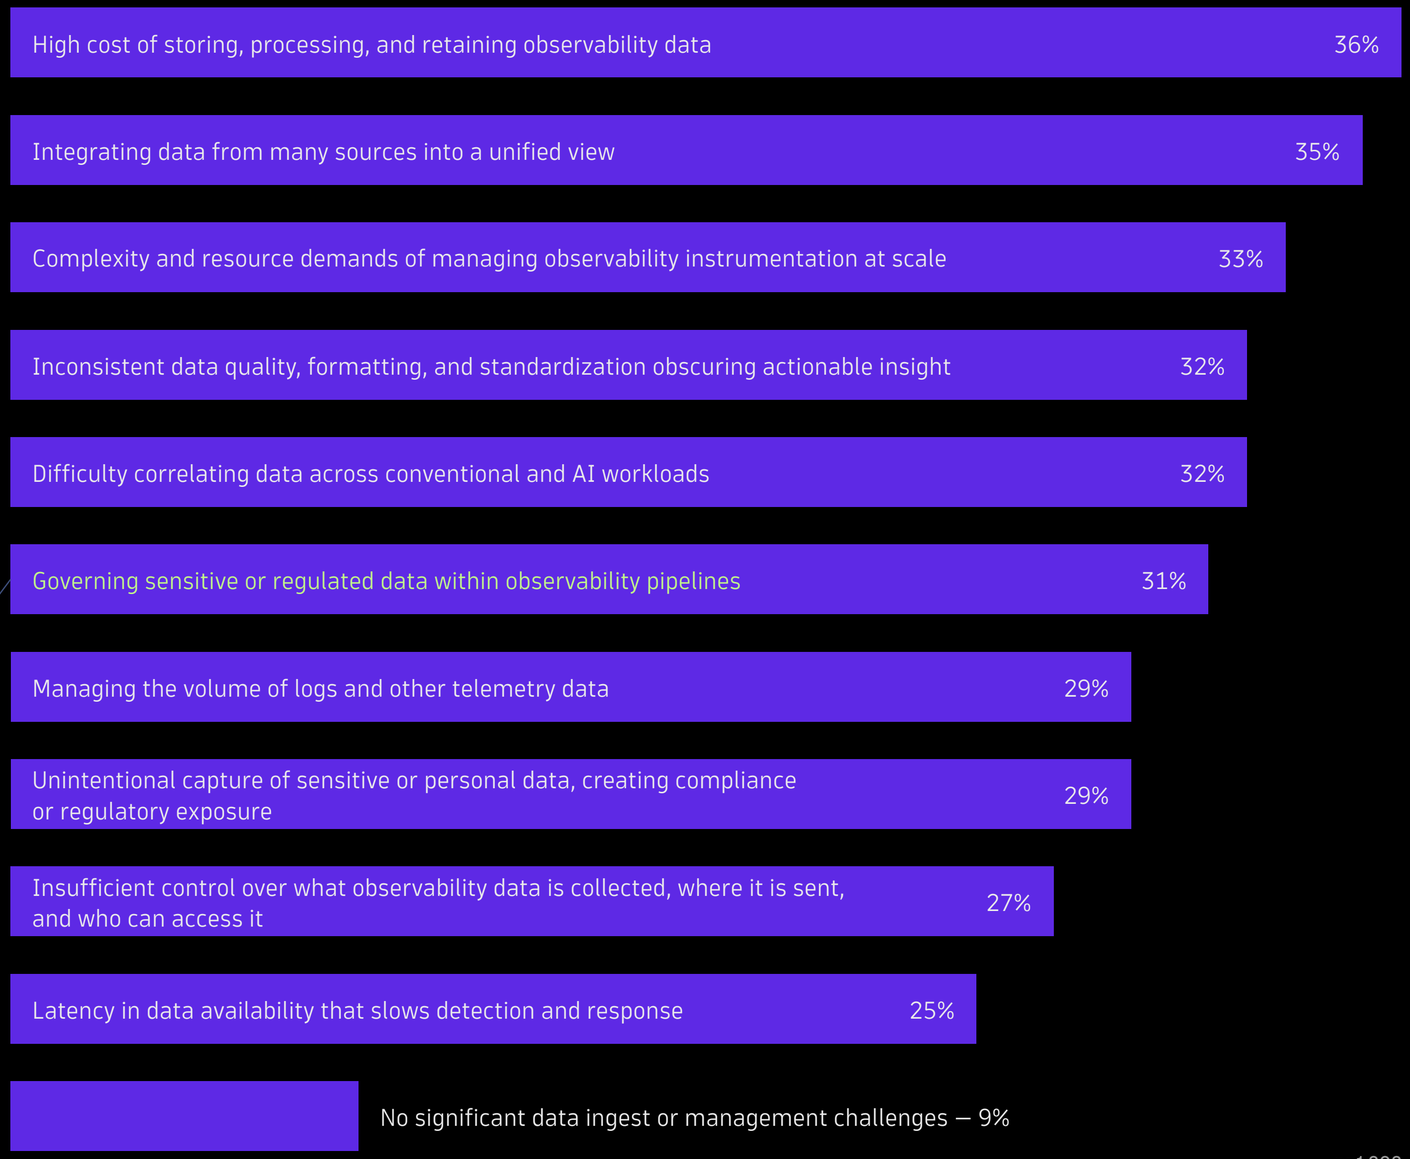
Organizations report challenges with data ingest and management issues that span cost, integration, instrumentation, quality, correlation, and governance.

To establish the evidence required to verify AI behavior and manage risk, data ingest and management must be consistent, connected, governed, and reliable across the environment.

Observability decision-makers put greater emphasis on governing sensitive data

33% of decision-makers cite governing sensitive data within observability pipelines as a challenge, compared with 19% of purchase influencers.

Data ingest/management challenges to observability progress



n=1,202

AI workloads are driving resource costs up without the visibility to manage them

Complexity is also raising resource demands.

AI workloads represent increased compute and infrastructure spending for 81% of organizations, energy costs for 78%, and the need for cost-allocation visibility for 77%. How organizations see these costs varies by deployment model, from token-based consumption charges for managed AI services to infrastructure utilization for AI workloads.

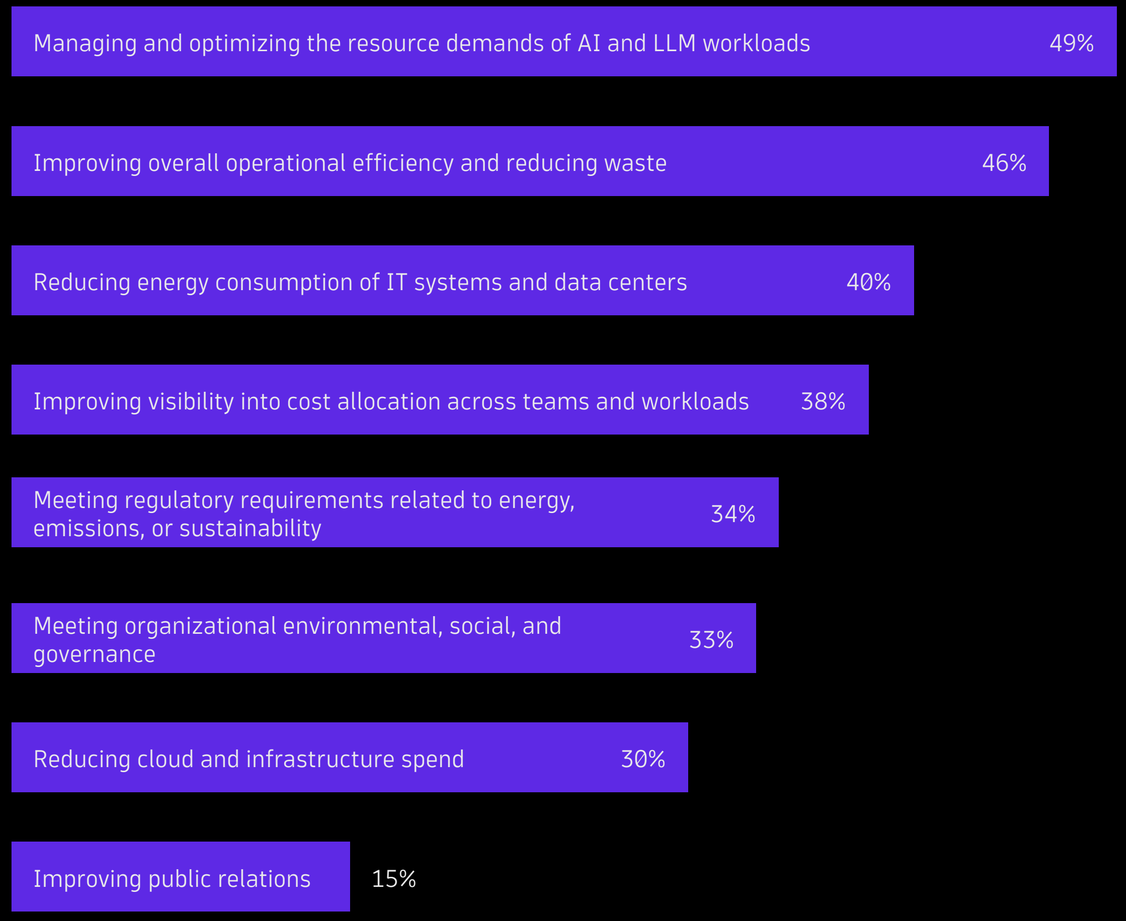
AI resource impact



n=1,202

Managing AI and LLM resource demands now leads sustainability priorities at 49%. Yet visibility remains the obstacle: 39% lack insight into energy consumption and emissions, while 39% can't attribute them to specific workloads. Organizations can't optimize AI resource consumption they can't measure and attribute.

ITOps sustainability priorities



ITOps sustainability barriers



n=1,202

THE TAKEAWAY

Use observability to simplify AI complexity

As AI compounds technology complexity, observability should reduce fragmentation and provide the shared context needed to understand and manage reliability, cost, and risk.

- **Work from one shared picture.** Connect cloud, applications, AI, security, and business services so context doesn't have to be rebuilt for every question.
- **Use AI to accelerate action.** Summarize findings, explain context in plain language, and orchestrate multi-step remediation.
- **Connect signals to consequences.** Link telemetry with dependencies, ownership, resource consumption, and business impact.
- **Govern data consistently.** Standardize telemetry, access, quality, and sensitive data controls.

More data and tools won't solve AI-driven complexity. Closing the visibility gaps requires AI observability that can show how AI actually behaves in production.

CHAPTER THREE

AI observability needs are outpacing execution

An AI system can be up and responding but still be wrong. That changes what organizations need from observability: visibility into AI quality, behavior, governance, and business impact.

Organizations recognize these requirements, but heavy reliance on provider-led approaches, fragmented visibility, and skills gaps make it harder to achieve unified AI observability across the enterprise.

66%

rely on AI-provider tools for AI observability

50%

say real-time AI output quality and accuracy is an increasingly important observability requirement

37%

cite skills gaps as the leading AI observability barrier

43%

consider agent decision visibility an important observability requirement; only 27% currently address it

Proof points

AI observability is widespread but provider-led

Over reliance on AI- and LLM-provider tools can fragment the evidence needed to understand AI across interconnected environments.

AI observability requirements are outpacing execution

Across nine capabilities, current execution trails perceived importance by an average 14 percentage points.

The hardest gaps concern what AI does and what it delivers

Execution trails requirements around agent decisions and business impact.

Organizations know where to focus next, but skills and fragmentation constrain progress

Real-time output quality is the top priority, but skills gaps and fragmented visibility are the leading barriers.

Provider-native tools lead AI observability today

Nearly all respondents use some sort of AI observability tool.

AI provider tools and platforms dominate: 66% use observability tools built into AI provider platforms and 57% use features built into LLM-provider monitoring.

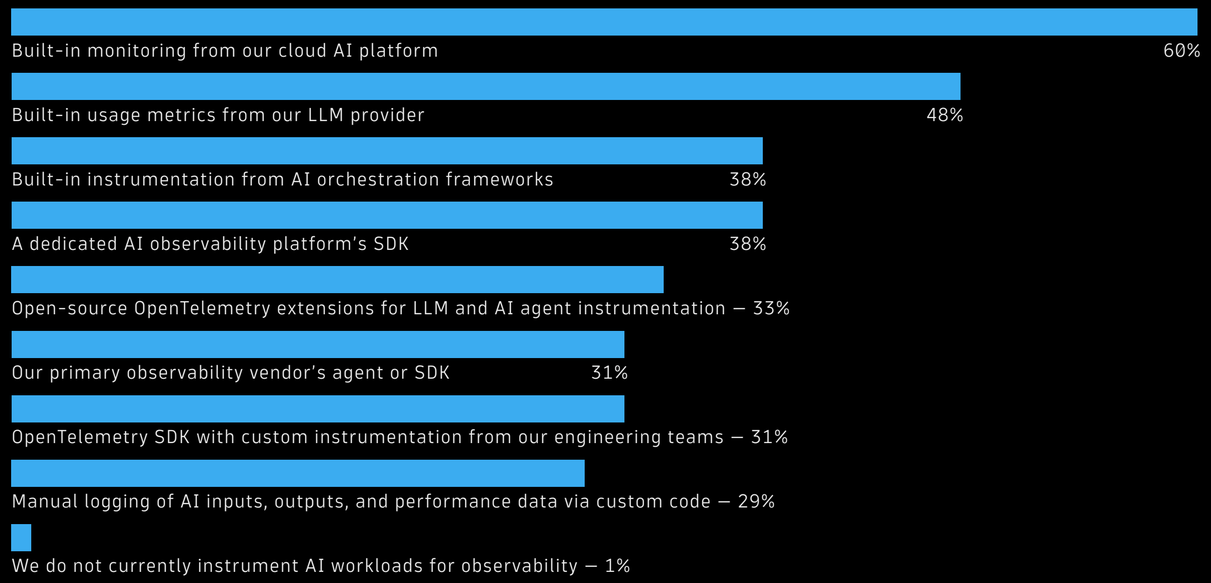
These tools report token counts, latency, and API errors, but leave it to organizations to figure out how to connect AI behavior and quality across providers and the broader technology environment.

Instrumentation follows the same provider-led pattern, while open and observability-platform approaches are less prevalent. That fragmentation can make end-to-end evidence harder to establish across providers and enterprise systems.

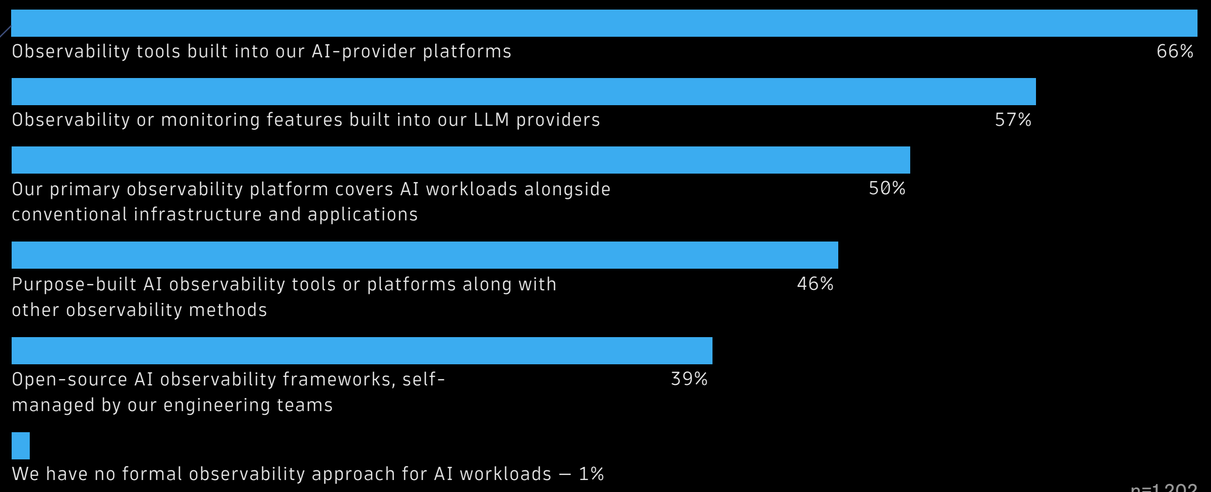
C-suite leaders report greater use of AI-provider observability tools

71% of C-suite leaders use observability tools built into their AI provider platforms, compared with 61% of high-level executives and 63% of mid-level executives.

AI observability instrumentation



AI observability approaches



n=1,202

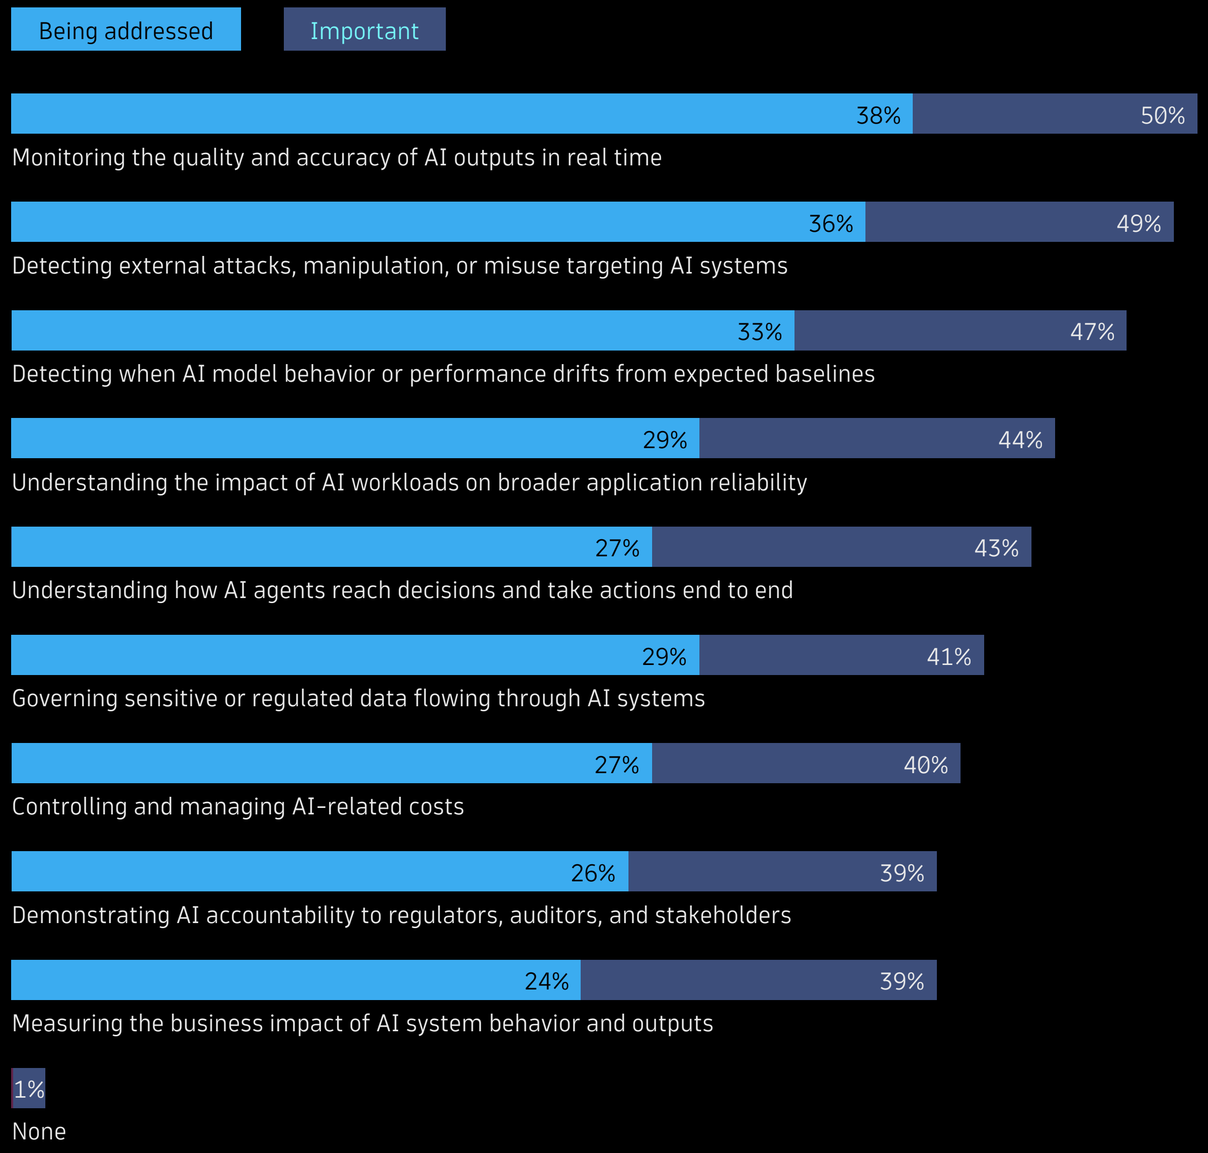
Organizations are not yet meeting their own AI observability requirements

AI changes what organizations need observability to verify.

Real-time output quality and accuracy leads requirements at 50%, followed by detecting attacks or misuse (49%) and model drift (47%). Across all nine requirements, an average 14-point gap separates importance from current execution.

The gaps are especially pronounced around agent behavior and business impact: 43% prioritize understanding how AI agents reach decisions and actions but only 27% currently address it, while only 24% currently measure business impact compared to 39% who consider it important.

Observability requirements for AI workloads



n=1,202

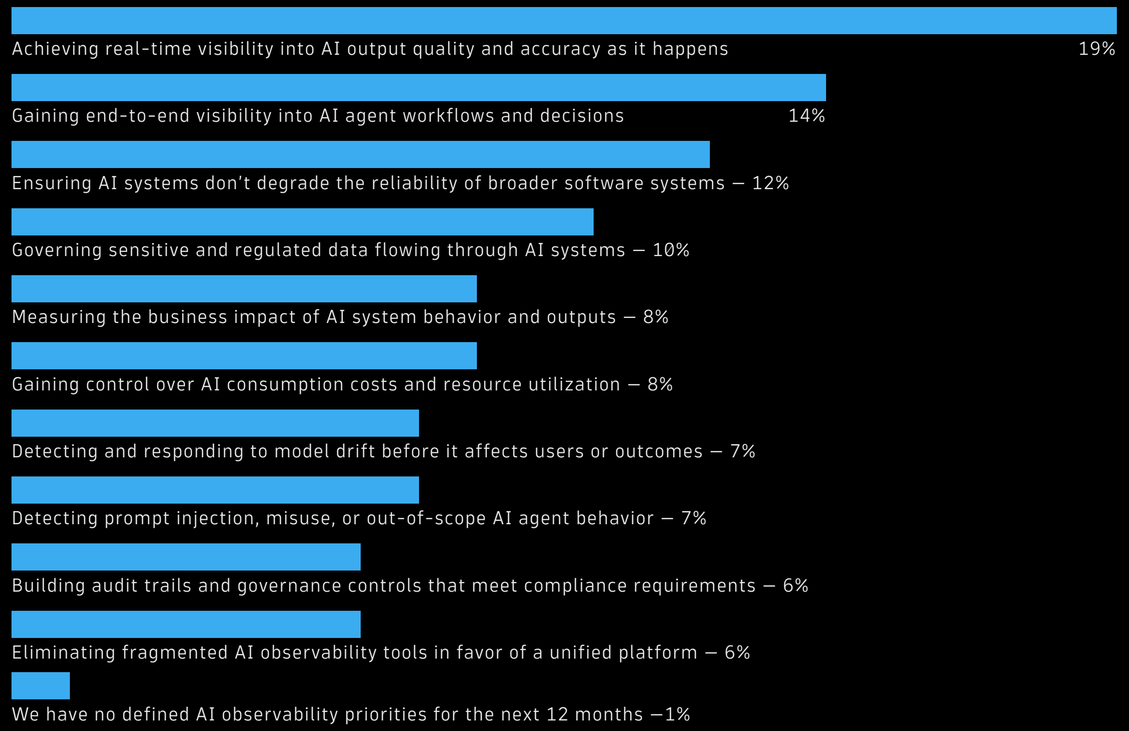
Organizations are prioritizing visibility into AI behavior, but skills and fragmentation constrain progress

Organizations report concentrating near-term attention on AI output quality, agent decisions, and broader system reliability.

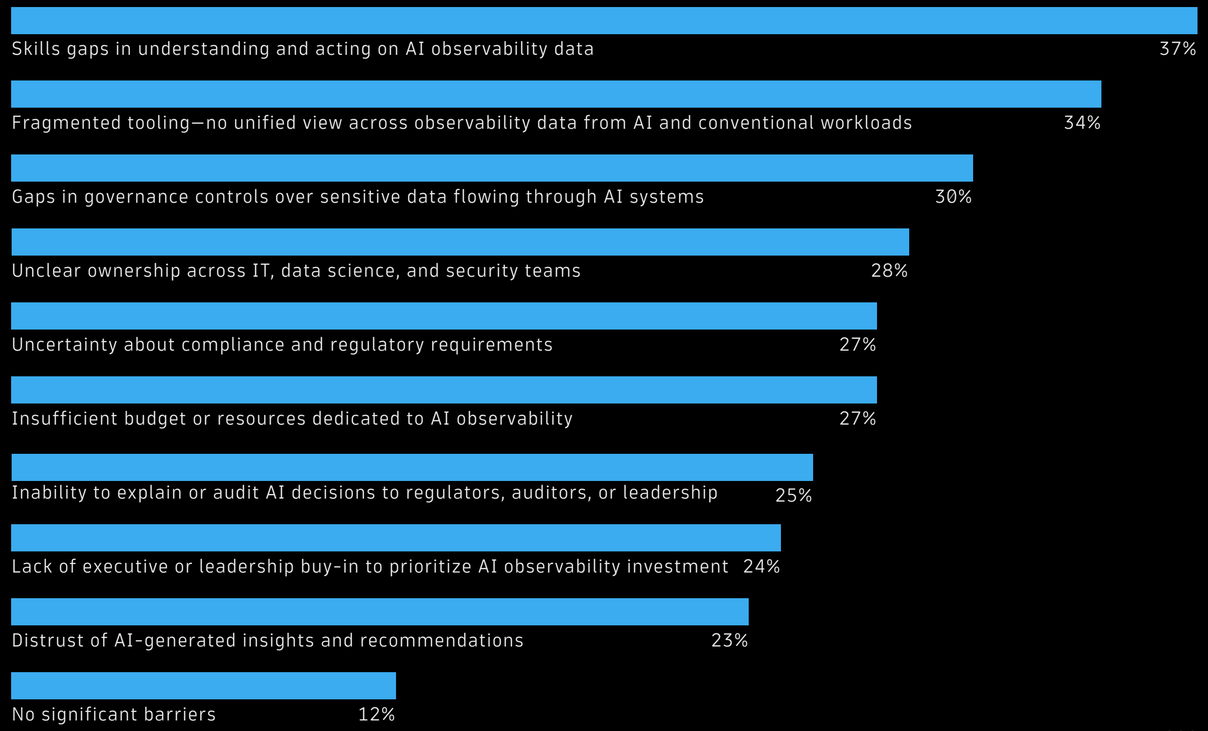
The leading barriers point beyond monitoring itself, particularly to skills gaps and fragmented tooling across observability data from AI and conventional workloads.

Closing that gap requires early instrumentation, continuous production evaluation, and context that connects AI behavior with infrastructure, applications, and business outcomes.

Highest AI observability priority for the next 12 months



Barriers to implementing AI observability



n=1,202

THE TAKEAWAY

Build the AI observability foundation that keeps AI accountable

Closing AI observability gaps requires design decisions. Leaders need shared context that connects production signals with engineering, so teams can fix causes, not just symptoms.

- **Instrument early.** Capture prompts, model versions, agent handoffs, tool calls, tokens, and evaluation scores.
- **Set multidimensional service levels.** Define ranges for relevance, faithfulness, hallucination, and cost.
- **Evaluate continuously.** Score production traffic for quality and drift.
- **Unify AI and operational data.** Connect AI telemetry to application, infrastructure, and business context to trace degraded output to its cause.
- **Make AI traceable.** Embed access controls, data lineage, and auditability to substantiate what AI did and why.

Understanding and verifying AI behavior is only the first step. As organizations give AI greater decision authority, the challenge shifts from understanding what AI does to determining when it can safely act.

CHAPTER FOUR

Autonomous operations are moving from ambition to scale

AI is accelerating autonomous operations.

Although 44% have set full autonomy as their five-year goal, only 12% have achieved it. Expanding AI decision authority requires evidence that its actions are reliable and governed, supported by causal intelligence, validation, and appropriate human oversight.[1]

87%

use AI agents for observability either moderately or extensively

66%

of AI-powered decisions are verified by humans

45%

verify AI outputs against trusted source data

31%

cite integration complexity as the top automation barrier

1. Findings in this chapter are based primarily on 989 respondents responsible for data-driven automation or AI-driven autonomous operations.

Proof points

Agentic operations are already taking shape

87% use AI agents for observability and ITOps moderately or extensively, while agent frameworks lead technologies organizations consider important for autonomy.

Organizations report broad AI value throughout the operational lifecycle

Reported impacts span predictive detection, software delivery, root-cause analysis, security, optimization, toil reduction, and compliance.

Agentic AI is moving toward broader orchestration

Cross-team workflows show the largest projected increase as organizations expand AI beyond individual operational tasks.

Human oversight remains integral as AI responsibility expands

Organizations combine human verification with data validation, continuous evaluation, audit trails, and behavioral monitoring to validate AI outputs and actions.

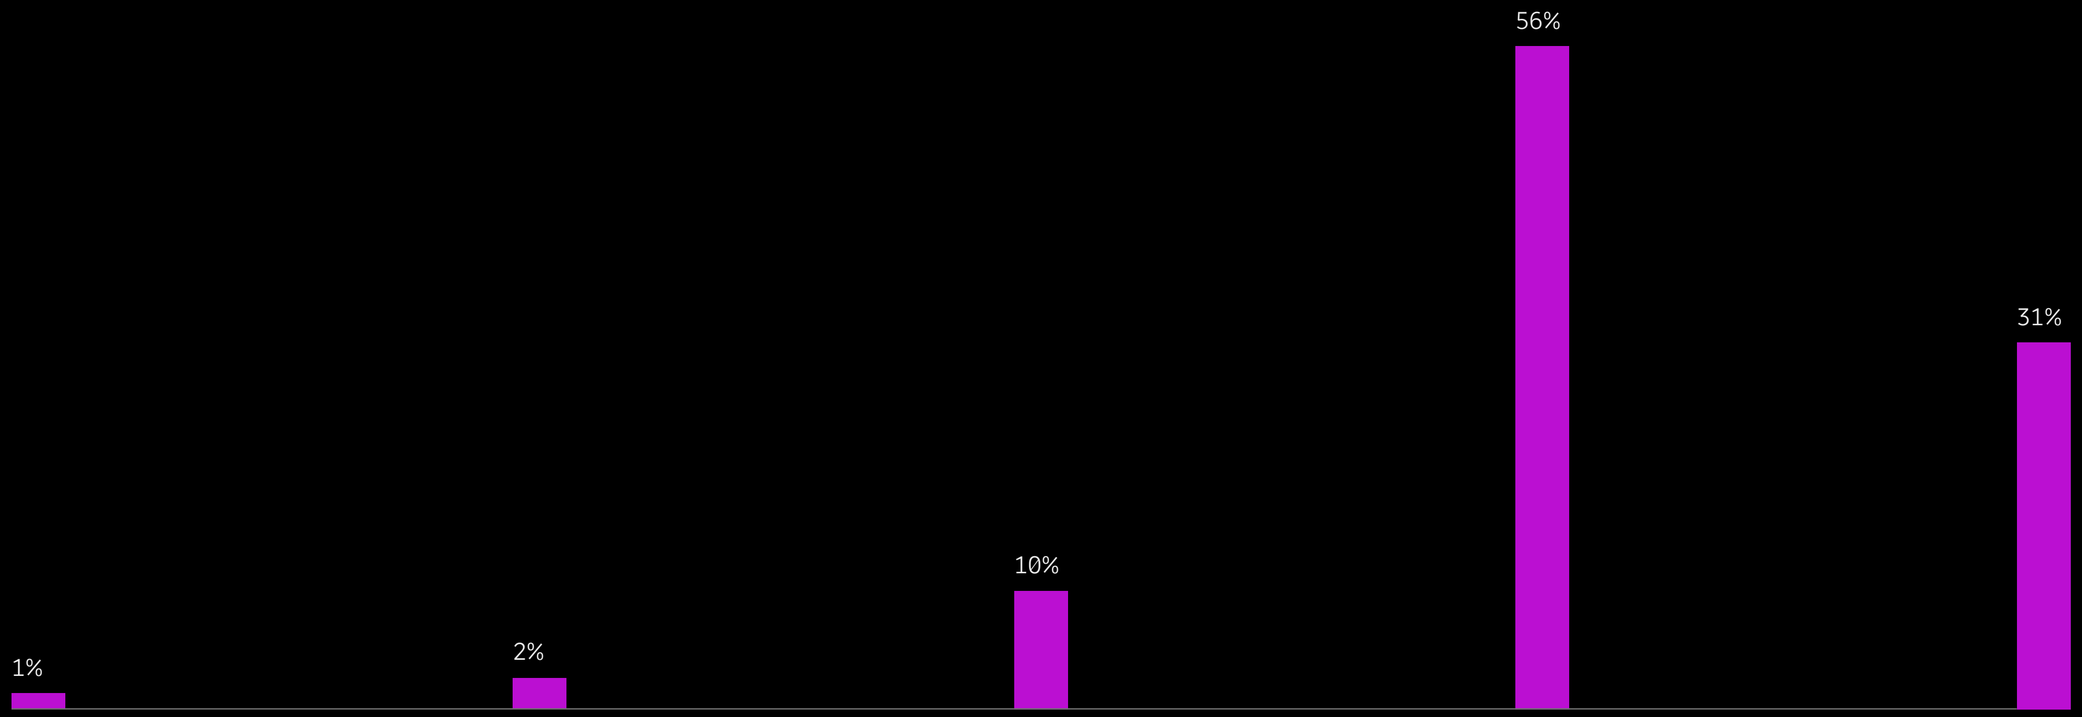
Broader AI concerns become constraints on autonomy

Security, privacy, and data quality are top concerns, and integration, skills, and trust are top barriers to further progress.

AI agents are laying the foundation for autonomous operations

AI agents are already widely used to support observability and IT operations:
87% report moderate or extensive use.

Current use of AI agents to support observability and ITOps



None
No use and no plans in the near term.

None today
No use today, but planned within 12 months.

Limited use
AI agents are used in pilots or a small number of workflows.

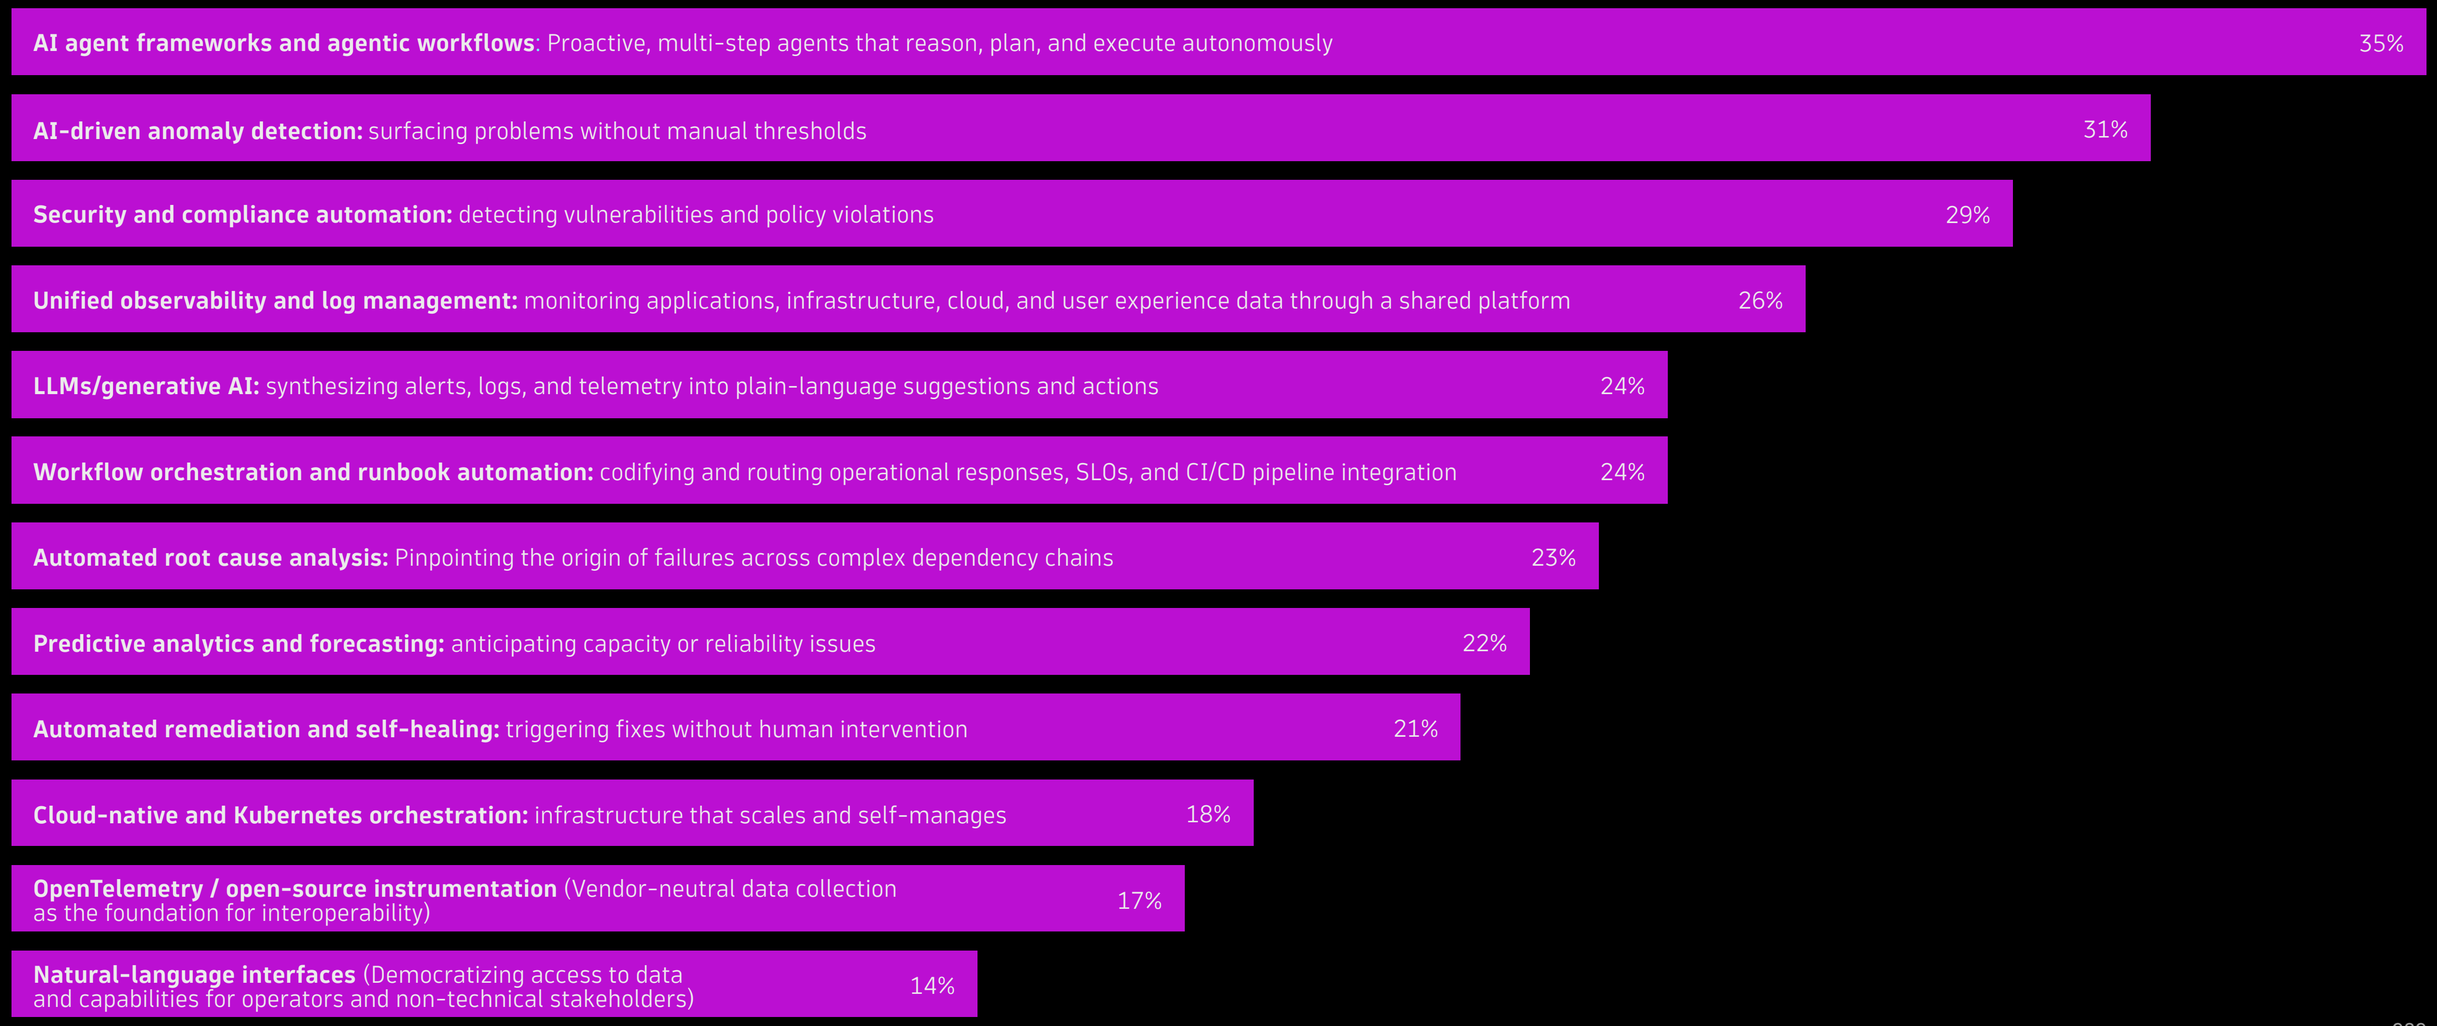
Moderate use
AI agents assist with several observability and IT operations tasks, but many tasks are still done conventionally or manually.

Extensive use
AI agents autonomously handle multiple observability and IT operations tasks.

n= 989

Their importance is also evident in future technology priorities, where agent frameworks and workflows rank highest for advancing autonomous operations.

Most important technologies for the journey toward autonomous operations



n=989

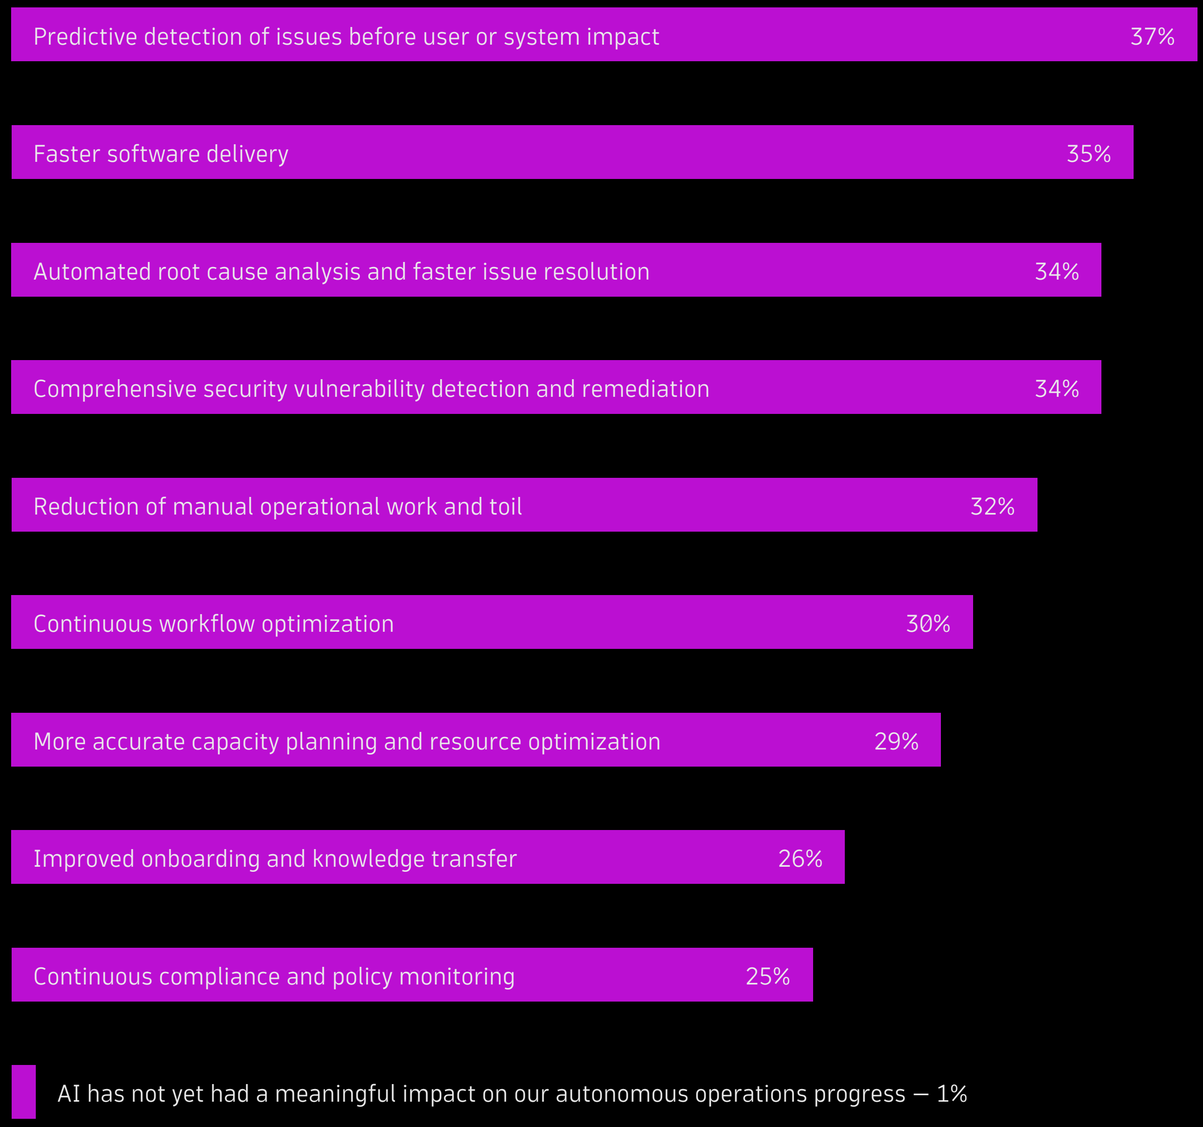
Organizations report broad AI value across the operational lifecycle

AI's reported contribution to autonomous operations is broad rather than concentrated in one benefit.

All nine measured impacts fall within a 12-point range, spanning predictive detection, delivery, root-cause analysis, security, optimization, toil reduction, and compliance. Only 1% report no meaningful impact.

Broad AI value by itself doesn't establish readiness for greater autonomy. As AI influences more of the operational lifecycle, observability must provide context for its decisions and consequences, giving teams the evidence they need to determine where greater AI decision authority is warranted.

Most meaningful impacts towards autonomous operations



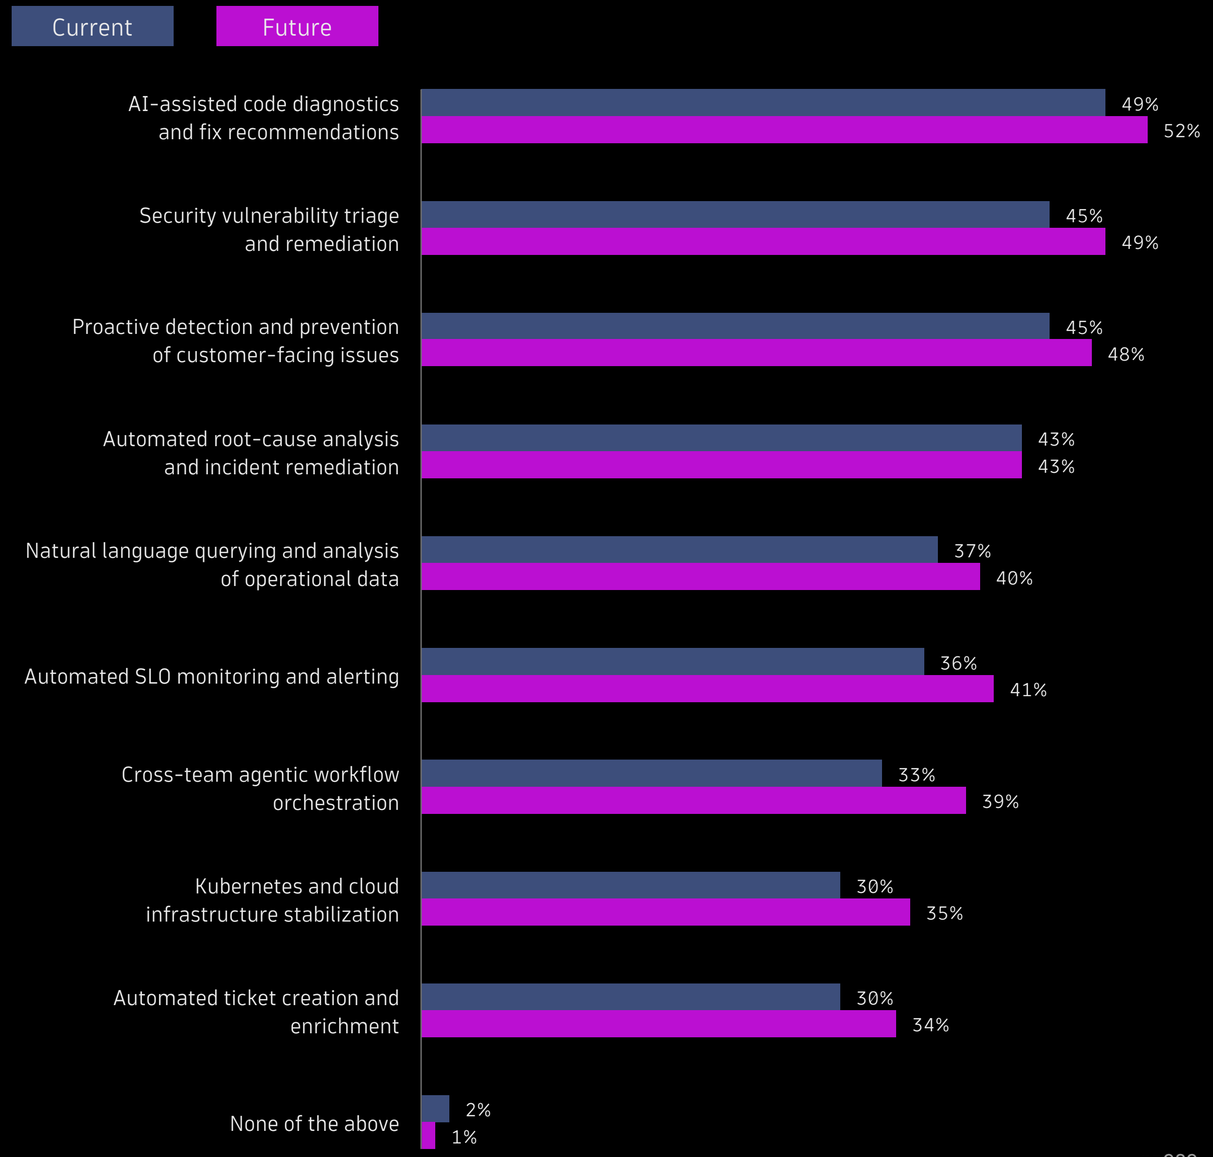
n=989

Agentic AI is expanding from individual tasks to broader workflows

Organizations report current agentic AI use spans development, security, incident response, and customer-facing operations.

Their future plans broaden that scope further: cross-team orchestration shows the largest expected increase, followed by Kubernetes stabilization and automated SLO monitoring.

Current and future agentic AI solutions



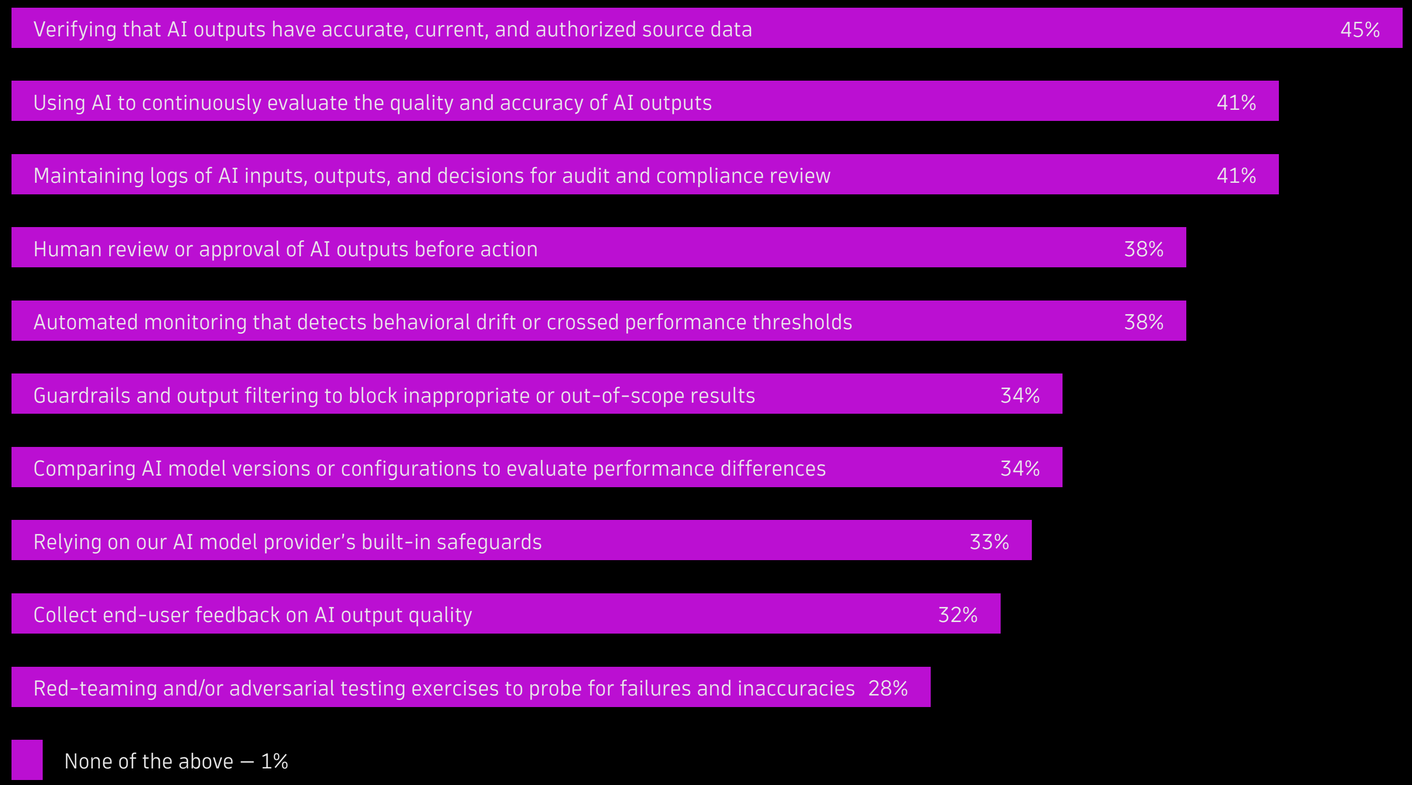
n=989

Organizations use layered safeguards to govern autonomous action

Organizations report using nearly four methods on average to validate AI outputs in production.

These methods include source-data verification, continuous evaluation, auditability, human review, and automated drift monitoring.

Current AI output validation measures



None of the above — 1%

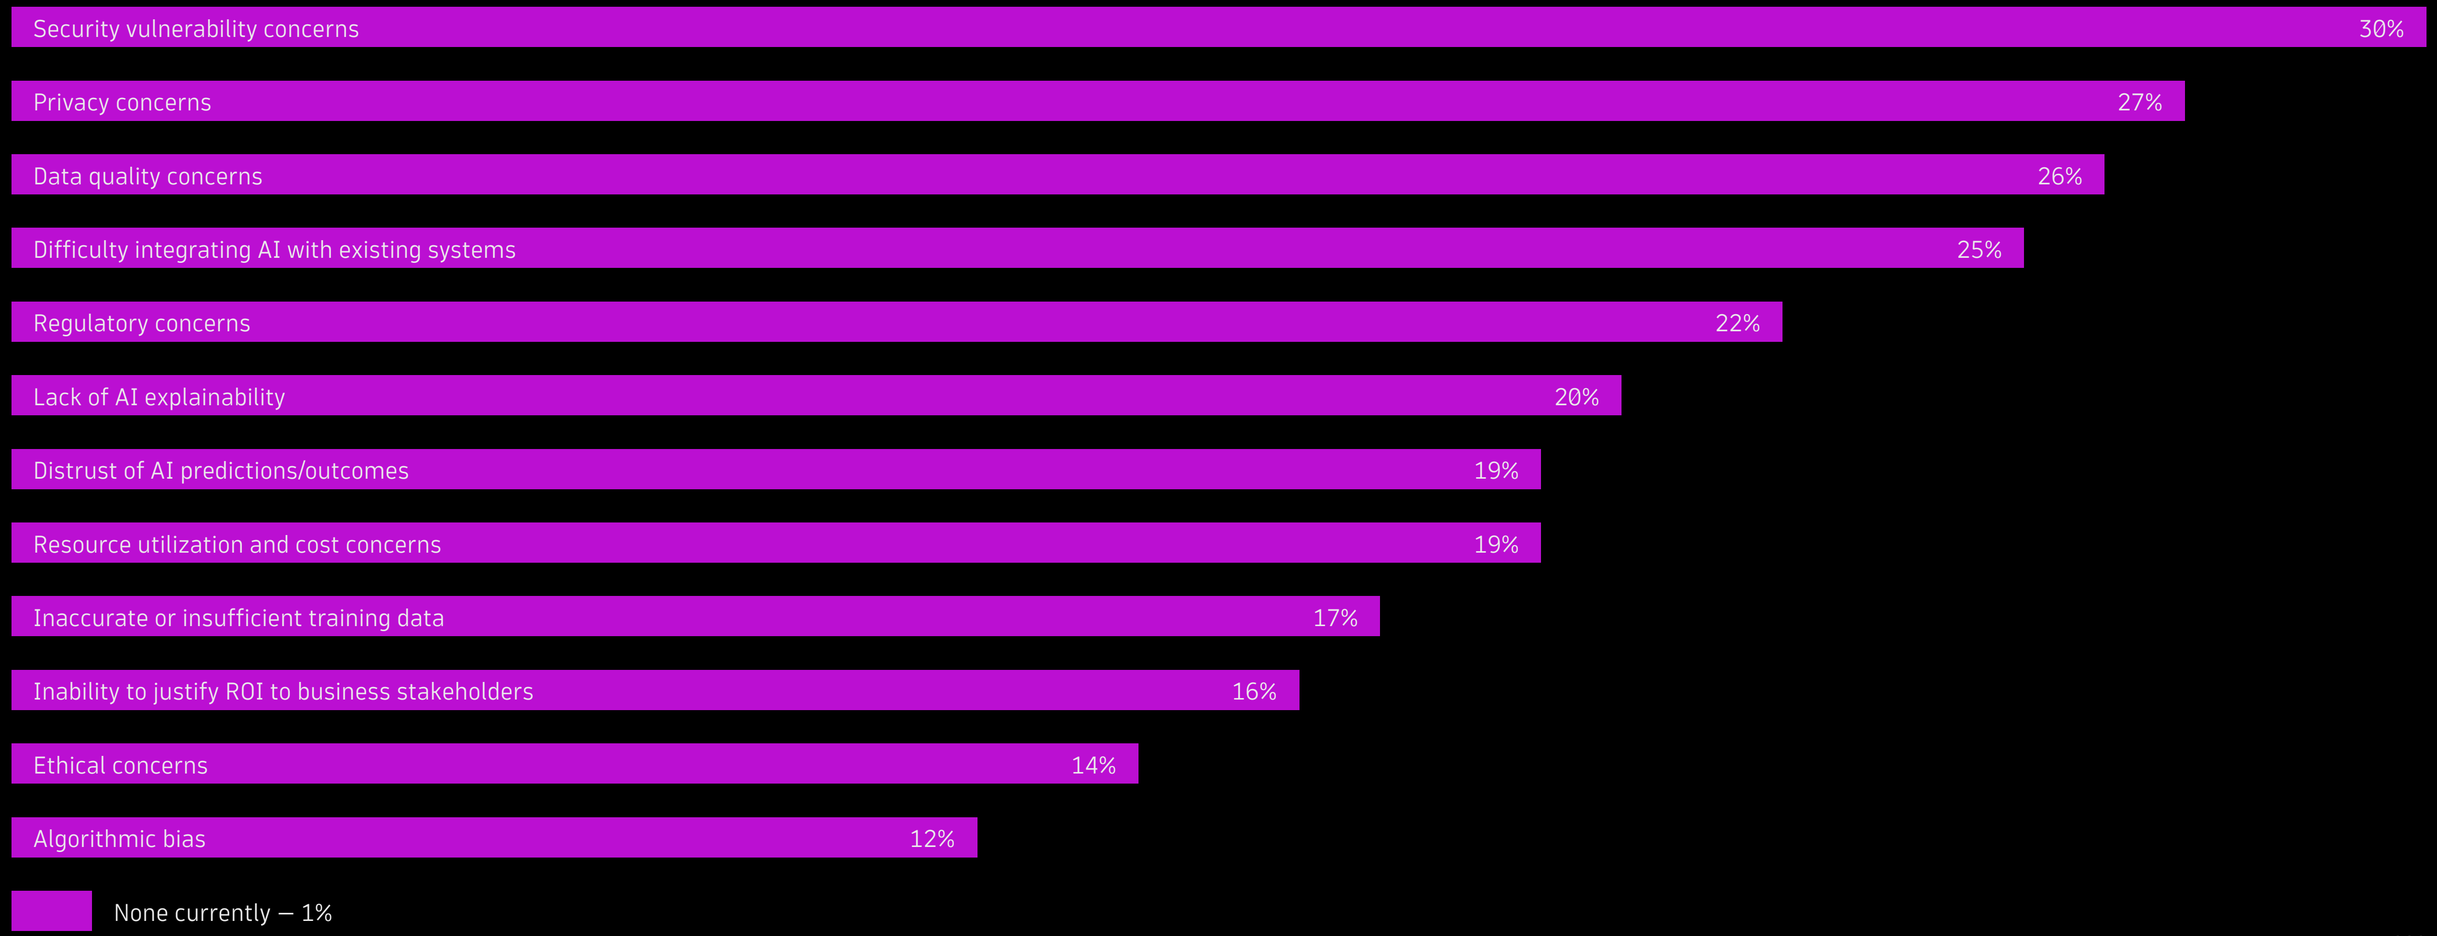


Human verification is also applied more broadly: respondents estimate that humans verify an average of 66% of AI-powered decisions.

n=989

Security, privacy, and data quality lead concerns about using AI for autonomous operations, reinforcing why organizations are retaining substantial oversight as AI assumes greater decision authority.

Concerns using AI for autonomous operations



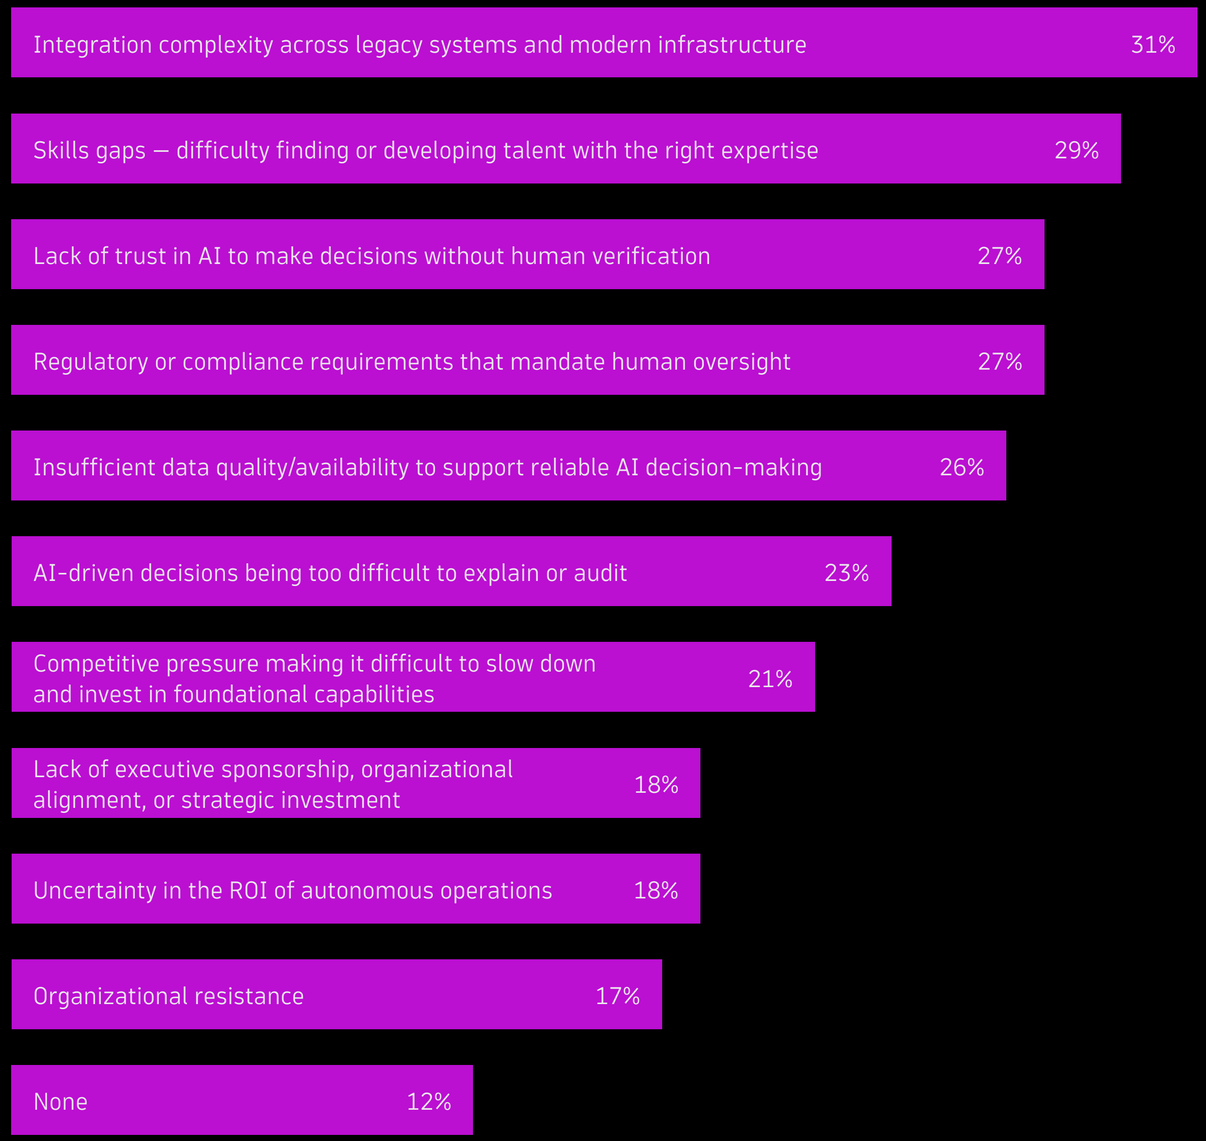
n=989

Broader AI concerns become barriers as autonomy expands

The concerns organizations have about AI reappear as practical constraints on autonomous operations. Integration complexity leads at 31%.

No single issue dominates, reinforcing that greater autonomy depends on technical integration, expertise, verifiable AI decisions, quality data, and appropriate oversight advancing together.

Barriers to advancing toward autonomous IT operations



n= 989

THE TAKEAWAY

Build the guardrails for expanding autonomous operations

Scaling autonomous operations requires a shift from passive monitoring to intelligent decision making and remediation. Leaders need trusted data, operational context, governance, and evidence that AI actions are reliable.

- **Design for integration.** Standardize interfaces and context across legacy, cloud, and AI environments.
- **Make trusted data foundational.** Govern quality, relevance, and context before AI acts.
- **Equip teams to oversee AI action.** Give operators shared context to validate decisions and manage exceptions.
- **Expand autonomy progressively.** Increase AI decision authority as reliability, explainability, and governance are demonstrated.

Greater autonomy increases the speed and consequence of machine-driven decisions. AI-accelerated security threats raise the pressure further, making runtime evidence essential to identifying and managing risk.

CHAPTER FIVE

Runtime context is reshaping protection against AI-accelerated security threats

The window between vulnerability disclosure and weaponized exploitation has collapsed from months to hours, or even before public disclosure.

That compression raises the stakes for detecting and prioritizing known and unknown vulnerabilities. Despite maturing security practices, zero-day exposure persists, driving the need for continuous monitoring, runtime evidence of actual risk, AI-aware threat detection, and closer convergence between security and observability.[1]

55%

experienced a zero-day incident in the past year

45%

continuously prioritize vulnerabilities using runtime exploitability status

50%

extend security monitoring to AI workloads and agents

82%

describe their security and observability approach as significantly or fully converged

1. Findings in this chapter are based primarily on 741 respondents responsible for application security.

Proof points



Security maturity doesn’t eliminate zero-day exposure

Organizations need continuous awareness as exploitation windows shrink and AI expands the attack surface.



Security is shifting toward runtime risk

Real-time detection and exploitability context help teams focus remediation on vulnerabilities that present actual risk.



AI is reshaping security operations

AI-aware monitoring is expanding, yet skills and organizational silos constrain teams’ ability to respond.



Security and observability are converging

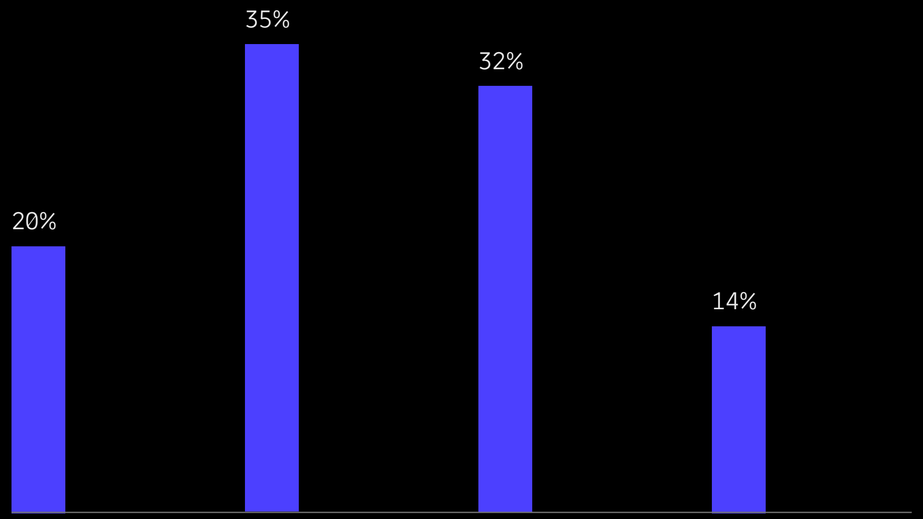
Shared context supports faster threat detection, incident response, and vulnerability prioritization based on business impact.

Zero-day incidents challenge organizations despite advancing security maturity

Two-thirds of respondents responsible for application security describe their organization's security posture as runtime-connected or fully integrated.

Yet regardless of maturity, 55% experienced a zero-day incident in the past year, including 20% that reported significant operational or business impact. Another 14% cannot confirm whether they experienced one. Reported security maturity has not eliminated exposure to vulnerabilities exploited before organizations can remediate them.

Zero day security incidents (in last 12 months)



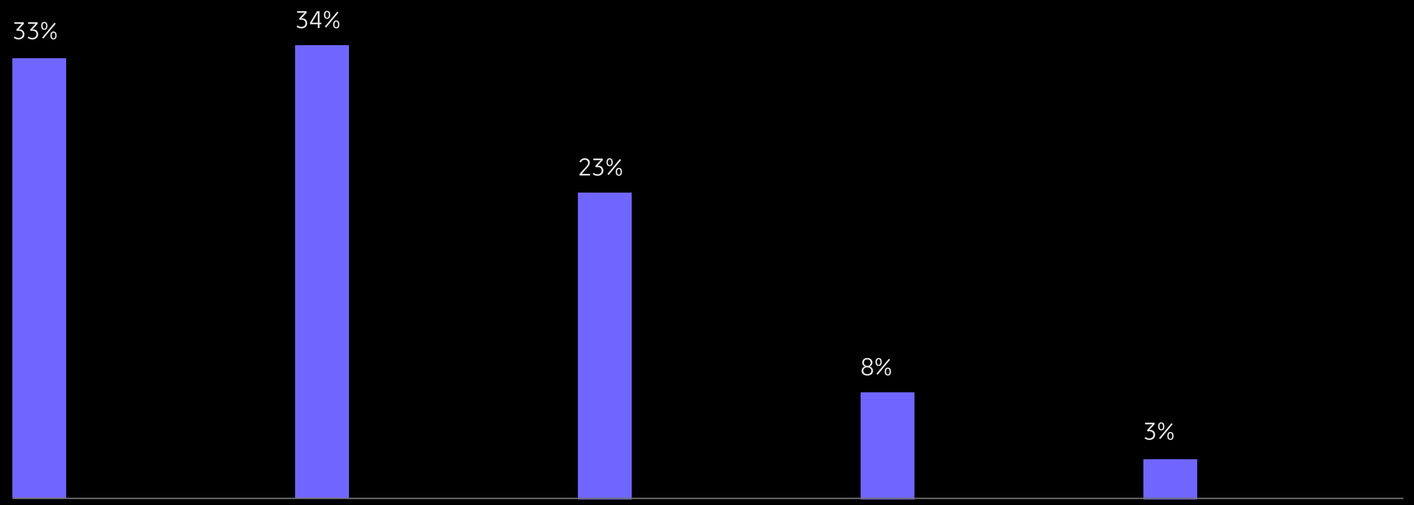
Yes
Resulted in significant operational or business impact.

Yes
Impact was contained.

No
Confirmed that we haven't experienced a zero-day attack.

Not aware of one
Cannot confirm whether we've experienced one.

Security maturity



Fully integrated
Vulnerability management, runtime detection, and incident response operate as one.

Runtime-connected
Runtime detection feeds directly into vulnerability prioritization and incident response.

Risk-informed
Prioritize known vulnerabilities based on active threat intelligence, not just severity.

Reactive
Formal vulnerability management program and patching of publicly disclosed vulnerabilities.

Ad hoc
No formal vulnerability management program; we address vulnerabilities as they come to our attention.

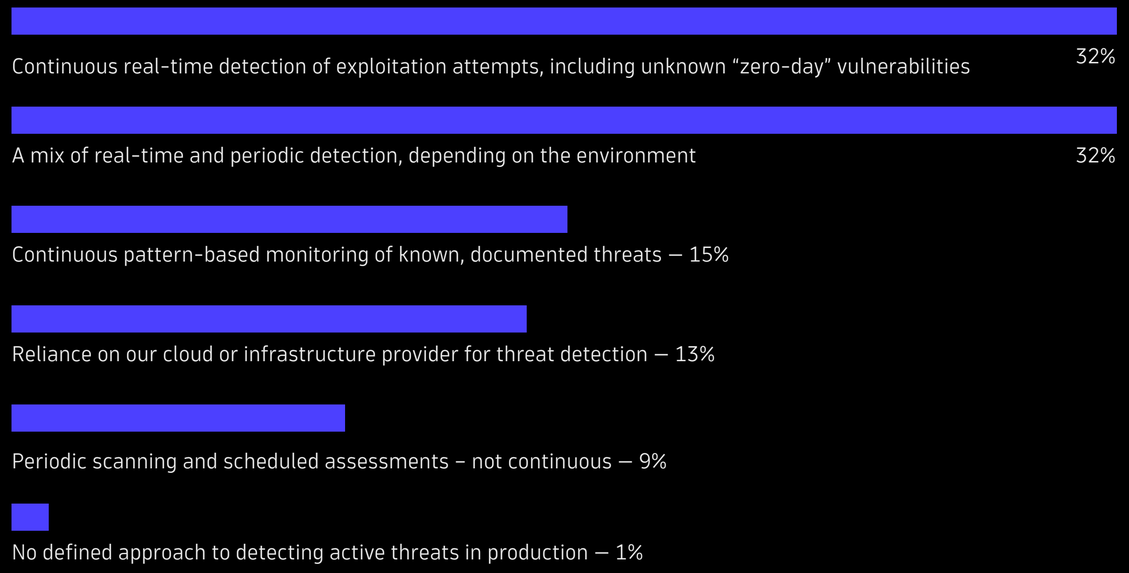
n=741

Detection is moving toward real-time, but prioritizing threats remains the challenge

Respondents report a growing emphasis on methods that incorporate real-time detection of known and unknown threats in production.

And while vulnerability remediation closely tracks organizational service level agreements—28% remediate critical vulnerabilities within 24 hours—faster detection doesn't resolve the harder question of prioritization: which vulnerabilities present actual risk in the runtime environment?

Security threat identification methods in production



Security vulnerabilities remediation time



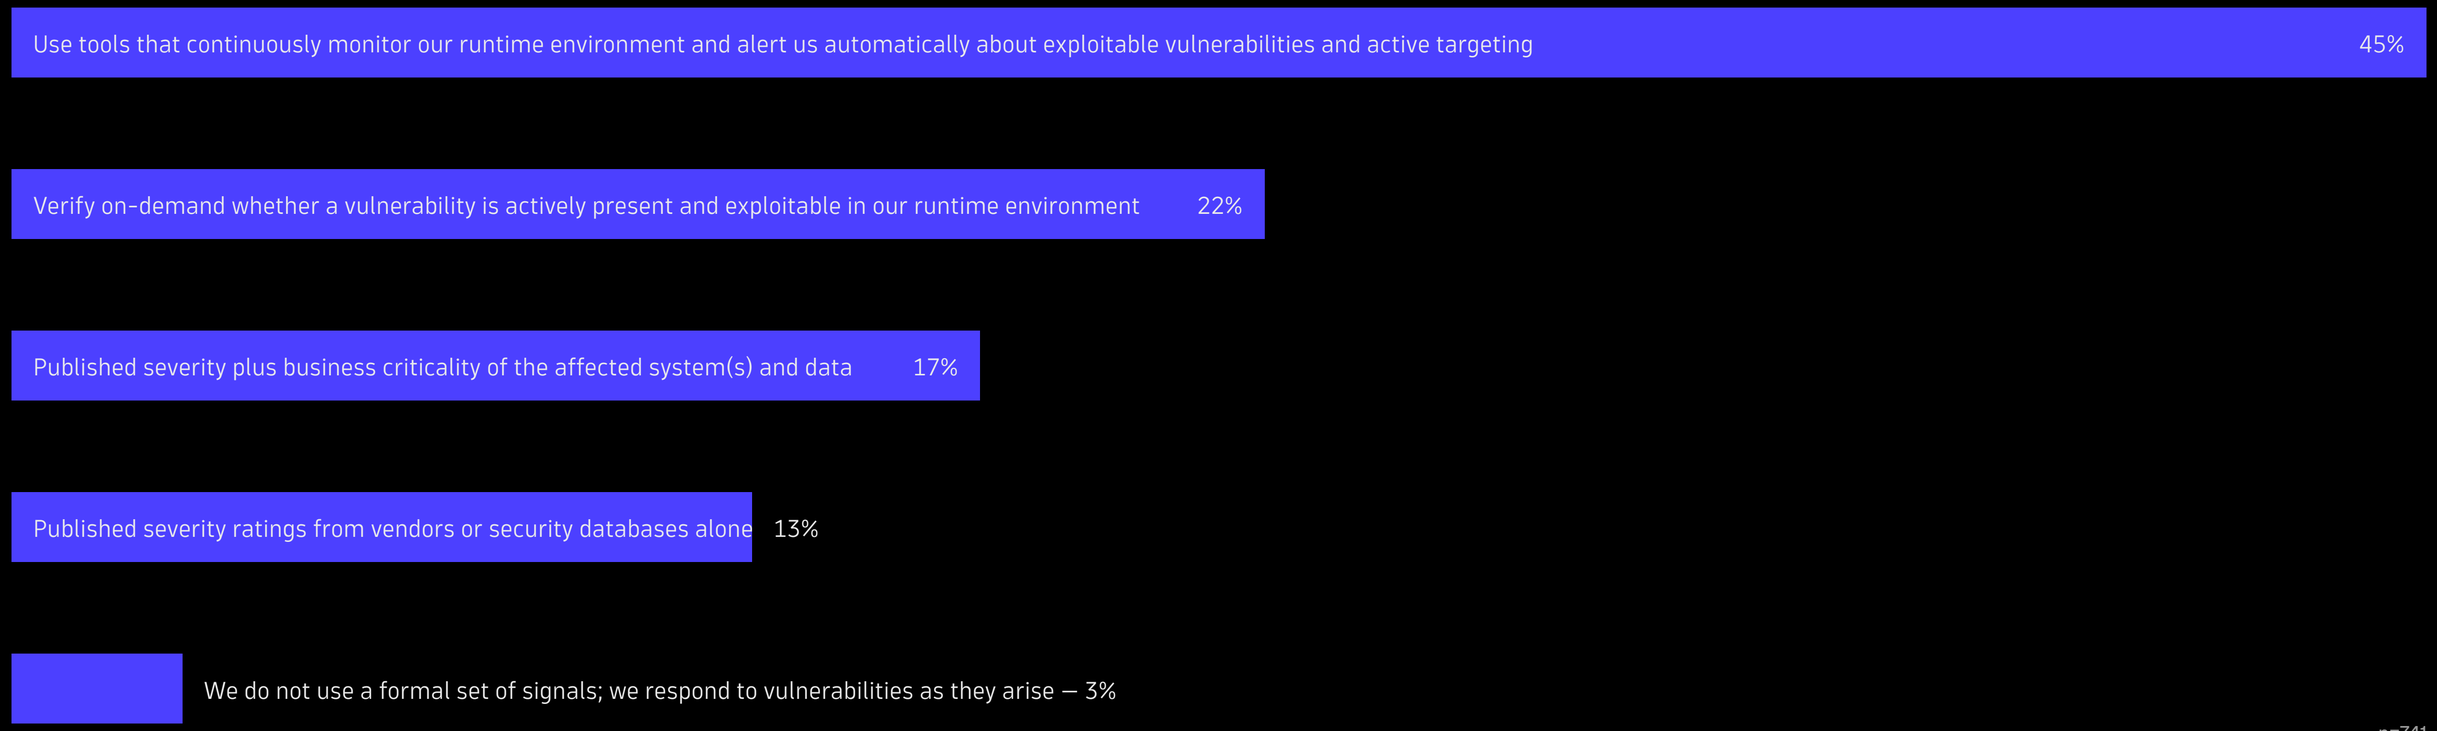
n=741

Runtime context is becoming the new prioritization standard

Nearly half of security respondents (45%) say their organizations continuously monitor runtime environments to identify vulnerabilities that are actually exploitable or actively targeted, while another 22% verify exploitability on demand.

Runtime evidence shifts vulnerability management from theoretical threats to prioritizing actual immediate risk.

Publicly disclosed patch priority



n=741

AI is reshaping both the attack surface and security response

Organizations report preparing for future exploits with greater AI-aware monitoring and AI-powered detection.

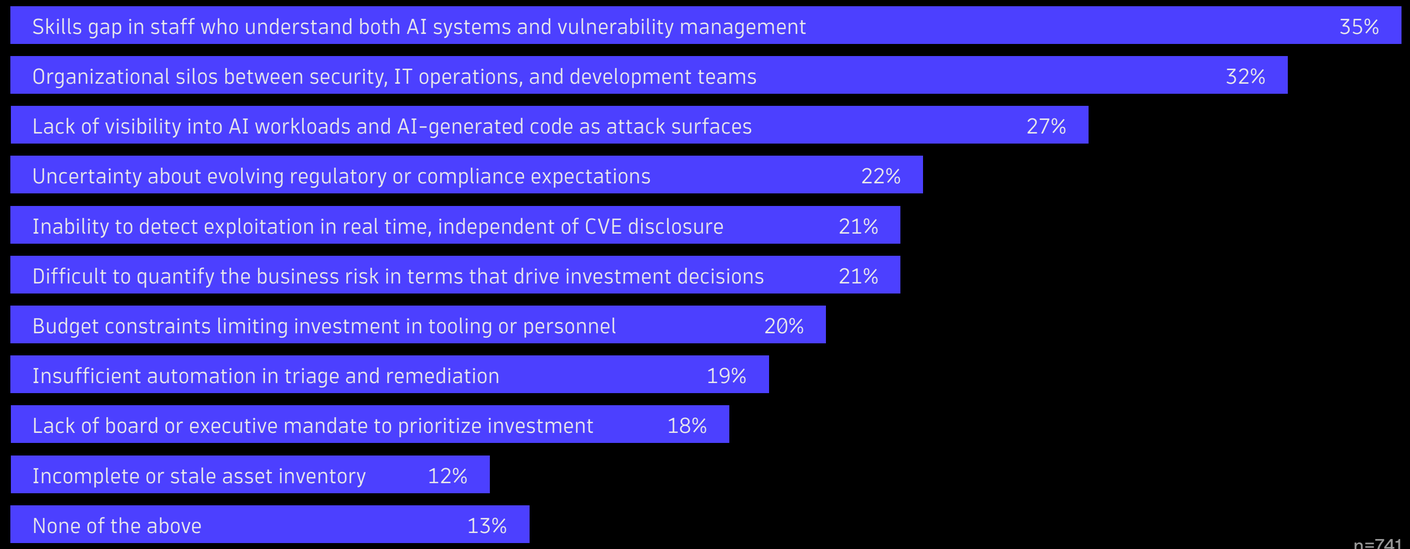
Yet vulnerability-management barriers remain rooted in expertise, organizational silos, and limited visibility into AI workloads.

Preparations for future security exploits



n=729

Top barriers to improving vulnerability management

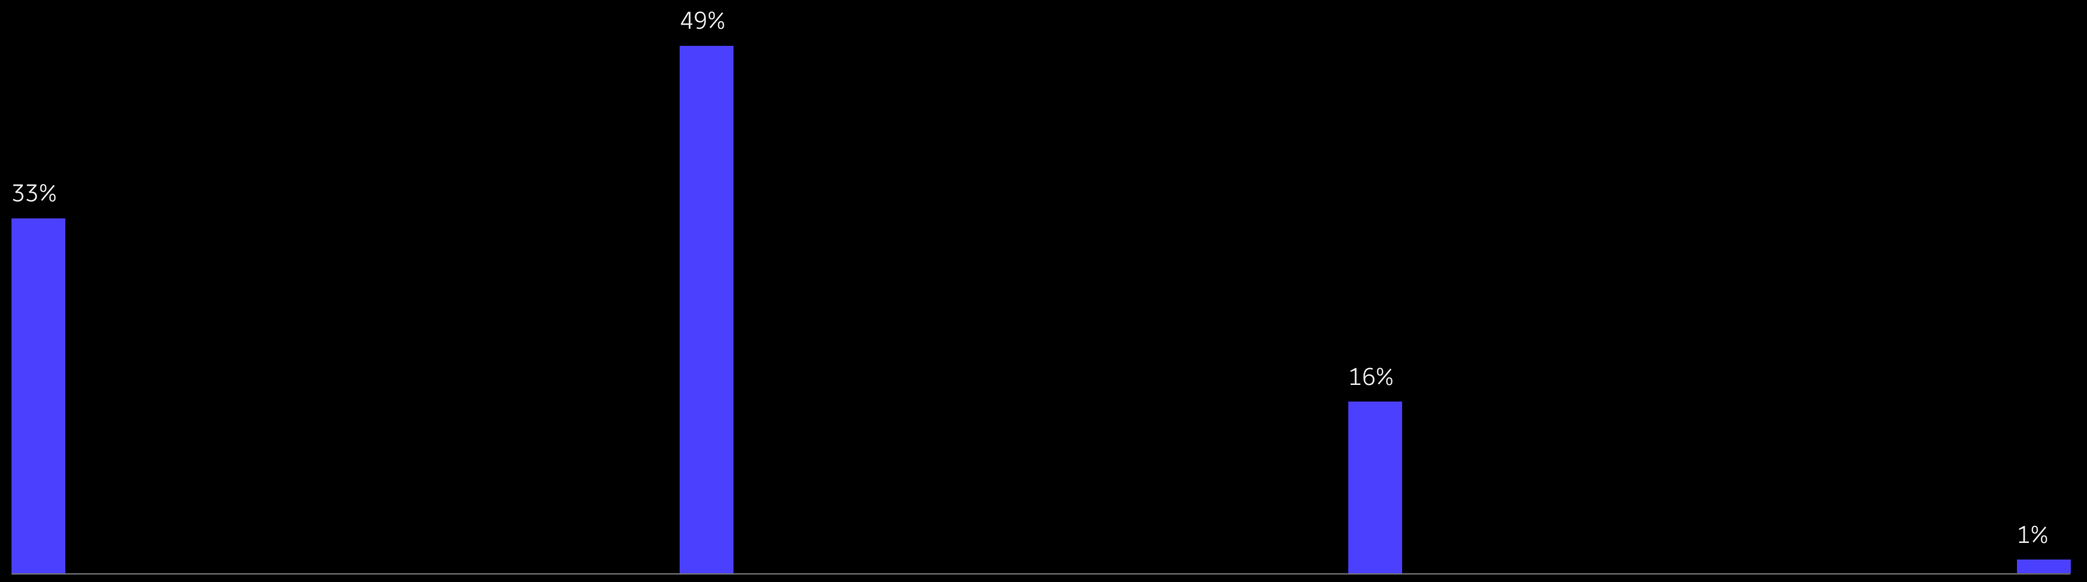


n=741

AI is strengthening the case for converging security and observability

Eighty-two percent describe their technical approach to security and observability as converged, either significantly (with some separation—49%) or fully (shared platforms, data, and workflows—33%).

Security and observability function convergence



Fully converged
Security and observability operate from a shared platform, data, and workflows.

Significantly converged
There is substantial integration, but some separation remains.

Partially converged
We have begun integrating security and observability but remain largely separate.

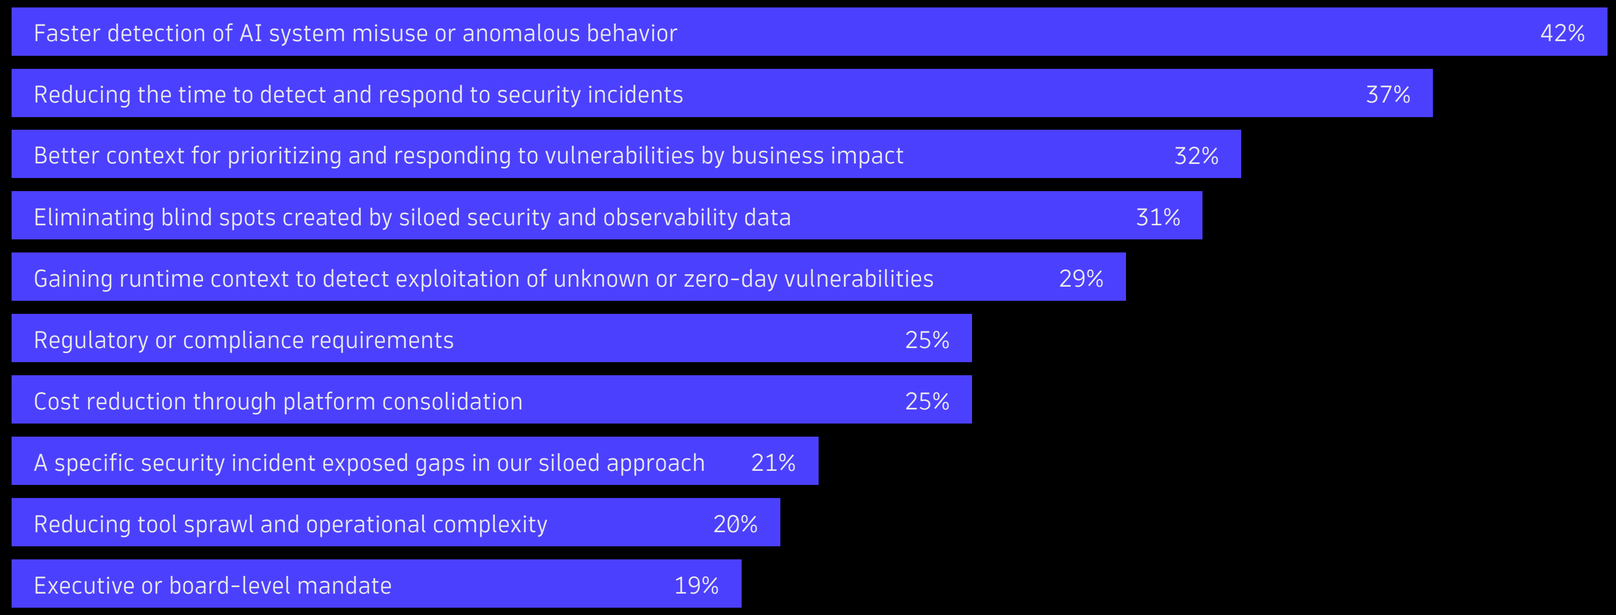
Not converged
Security and observability operate independently with separate tools and teams.

n=741

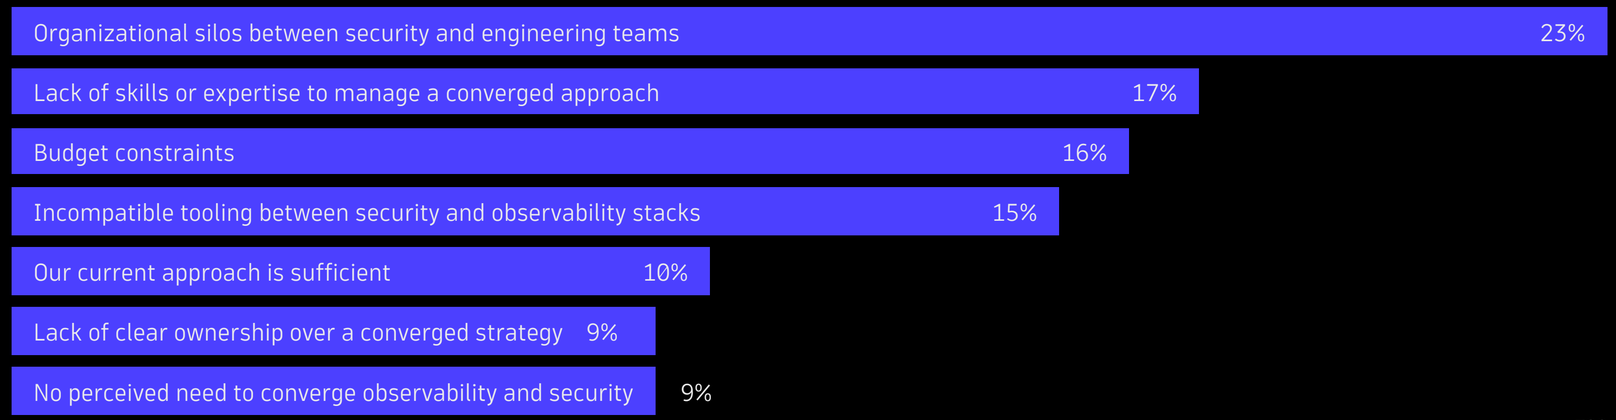
The strongest reported motivations for this convergence center on detecting AI misuse faster, accelerating incident response, and prioritizing according to business impact.

The leading reported barrier to greater convergence is organizational rather than technical, led by organizational silos, skills gaps, and budget constraints.

Primary drivers to converge security and observability



Biggest barriers to converging security and observability practices



n=729

THE TAKEAWAY

Runtime context turns security signals into remediation priorities

AI-accelerated threats require security strategies that test resilience early, continuously evaluate production behavior, and prioritize the risks that matter most.

- **Test before production.** Use AI debugging and red-teaming to expose weaknesses early.
- **Prioritize actual exposure.** Focus remediation on runtime exploitability, reachability, and business impact.
- **Maintain continuous runtime context.** Detect changing behavior and risk beyond pre-production testing.
- **Converge security and observability.** Give teams shared context to detect, understand, and respond faster.

Runtime context depends on telemetry that moves across technologies and teams. Open standards provide the interoperable foundation needed to extend that intelligence as AI scales across the business.

CHAPTER SIX

Open standards create the foundation for AI at scale

AI workloads need evidence and context that can move across technologies, teams, and lifecycle stages.

The Cloud Native Computing Foundation has recognized OpenTelemetry as a production-ready standard for vendor-neutral telemetry, enabling organizations to change analysis tools without reinstrumenting applications. As AI scales, that open foundation can provide continuity from telemetry collection through AI development, evaluation, and production observability.[1]

77%

have adopted OpenTelemetry selectively, broadly, or fully

20%

have made OpenTelemetry a strategic standard

37%

cite AI and LLM workload observability as a primary adoption driver

33%

cite skills gaps as their top OpenTelemetry implementation challenge

1. Findings in this chapter are based primarily on 1,096 respondents who have adopted or are evaluating OpenTelemetry in their environments.

Proof points



OpenTelemetry adoption outpaces strategic standardization

77% have adopted it beyond evaluation, but only 20% make it a strategic standard.



AI is strengthening the case for open standards

Standardization, AI observability, security, and avoiding vendor lock-in are top adoption drivers.



OpenTelemetry is extending deeper into AI observability

Most adopters use it as part of their AI observability approach, but strategic depth and consistency vary.



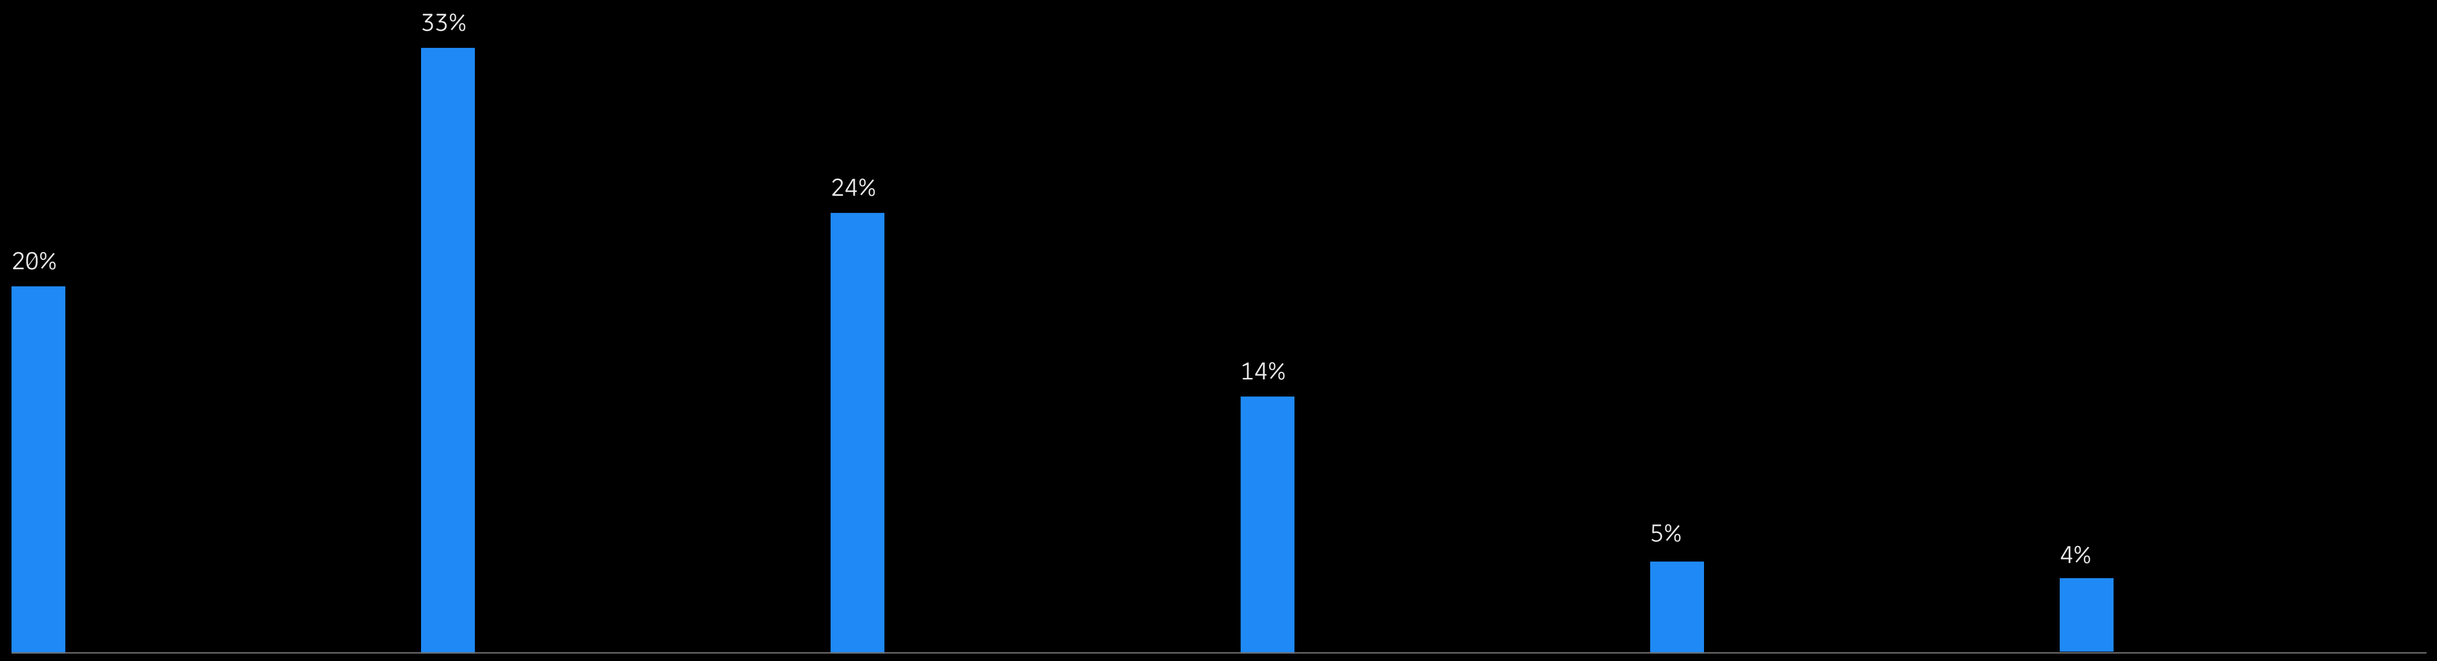
Value is clear, but scale introduces friction

Most adopters report meaningful value while skills, telemetry costs, engineering effort, and integration complicate expansion.

Most organizations use OpenTelemetry, but strategic standardization lags

More than three-quarters say they've moved OpenTelemetry beyond evaluation, but only 20% have made it a strategic standard, indicating that most organizations are still progressing from implementation toward an organization-wide standard.

Current adoption and strategic approach to OpenTelemetry



Strategic standard
OpenTelemetry is fully adopted and a core pillar of our long-term observability strategy.

Broadly adopted
OpenTelemetry is widely implemented but secondary to our observability platform.

Selectively adopted
OpenTelemetry is used by some teams or in select areas without a unified organizational approach.

In evaluation
We are piloting or assessing OpenTelemetry but have not committed to broader adoption.

Aware but not pursuing
We are familiar with OpenTelemetry but have no current plans to adopt it.

Unfamiliar
Not familiar with OpenTelemetry.

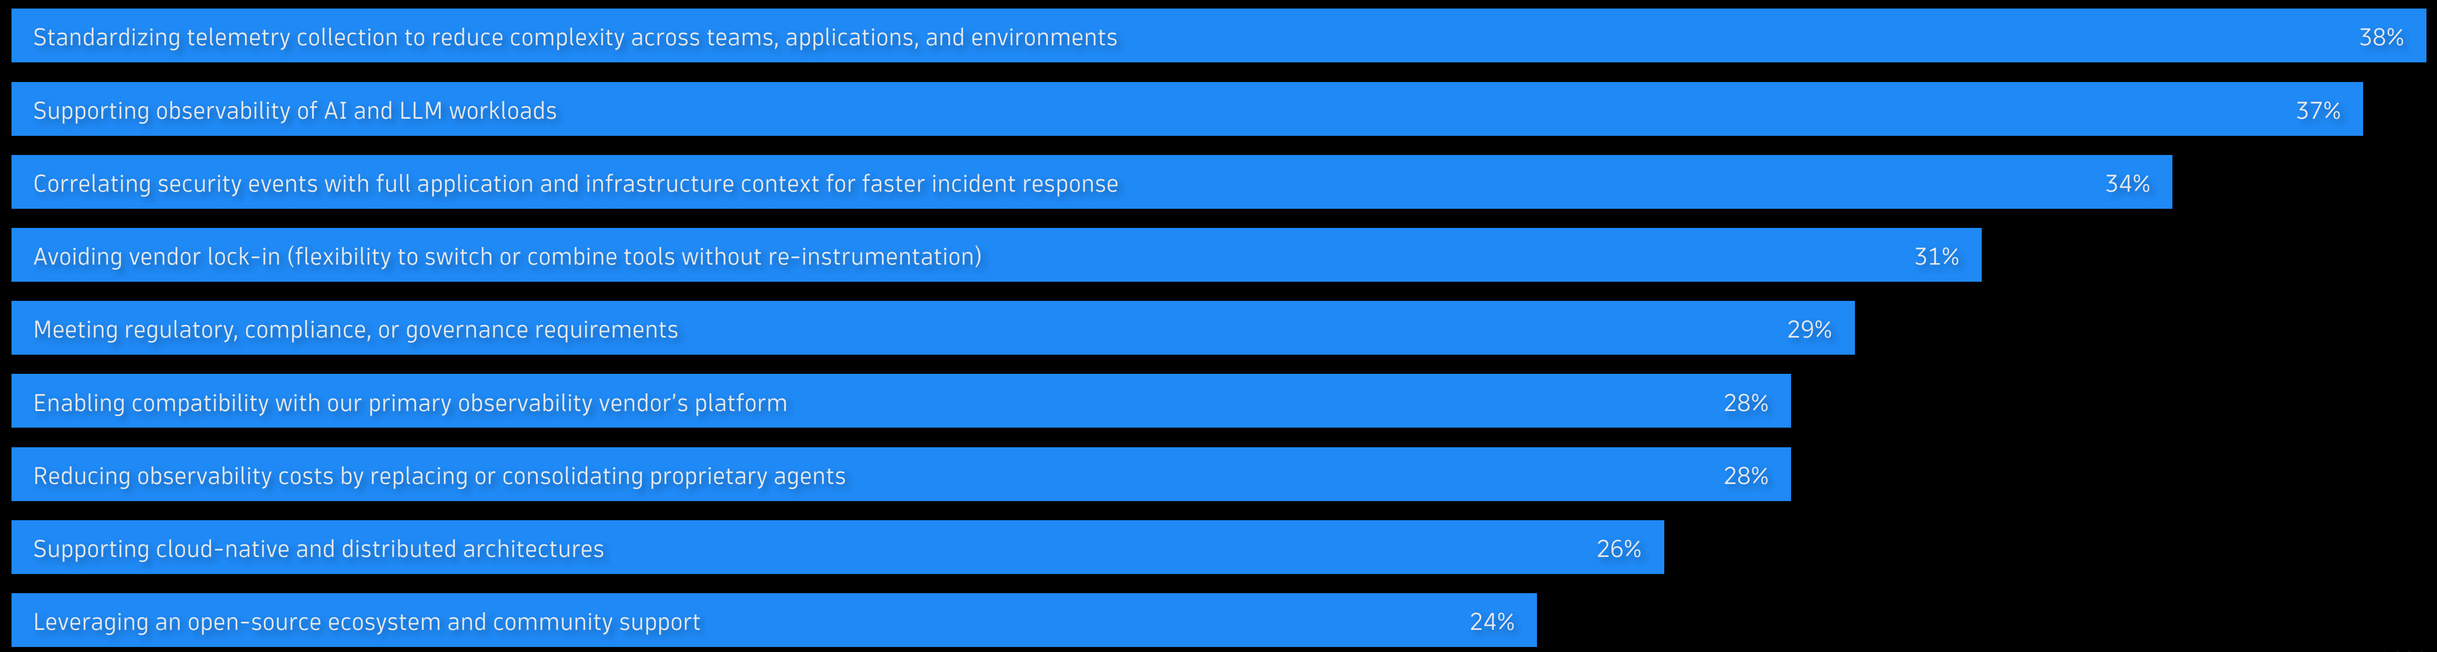
n= 1,202

Open standards are becoming foundational to AI observability

Respondents cite several reasons for adopting OpenTelemetry: standardizing telemetry, supporting AI observability, correlating security events, and avoiding vendor lock-in.

Together, these priorities point toward establishing an interoperable telemetry foundation that can span increasingly diverse cloud, AI, security, and observability environments.

Primary reasons for adopting/considering OpenTelemetry



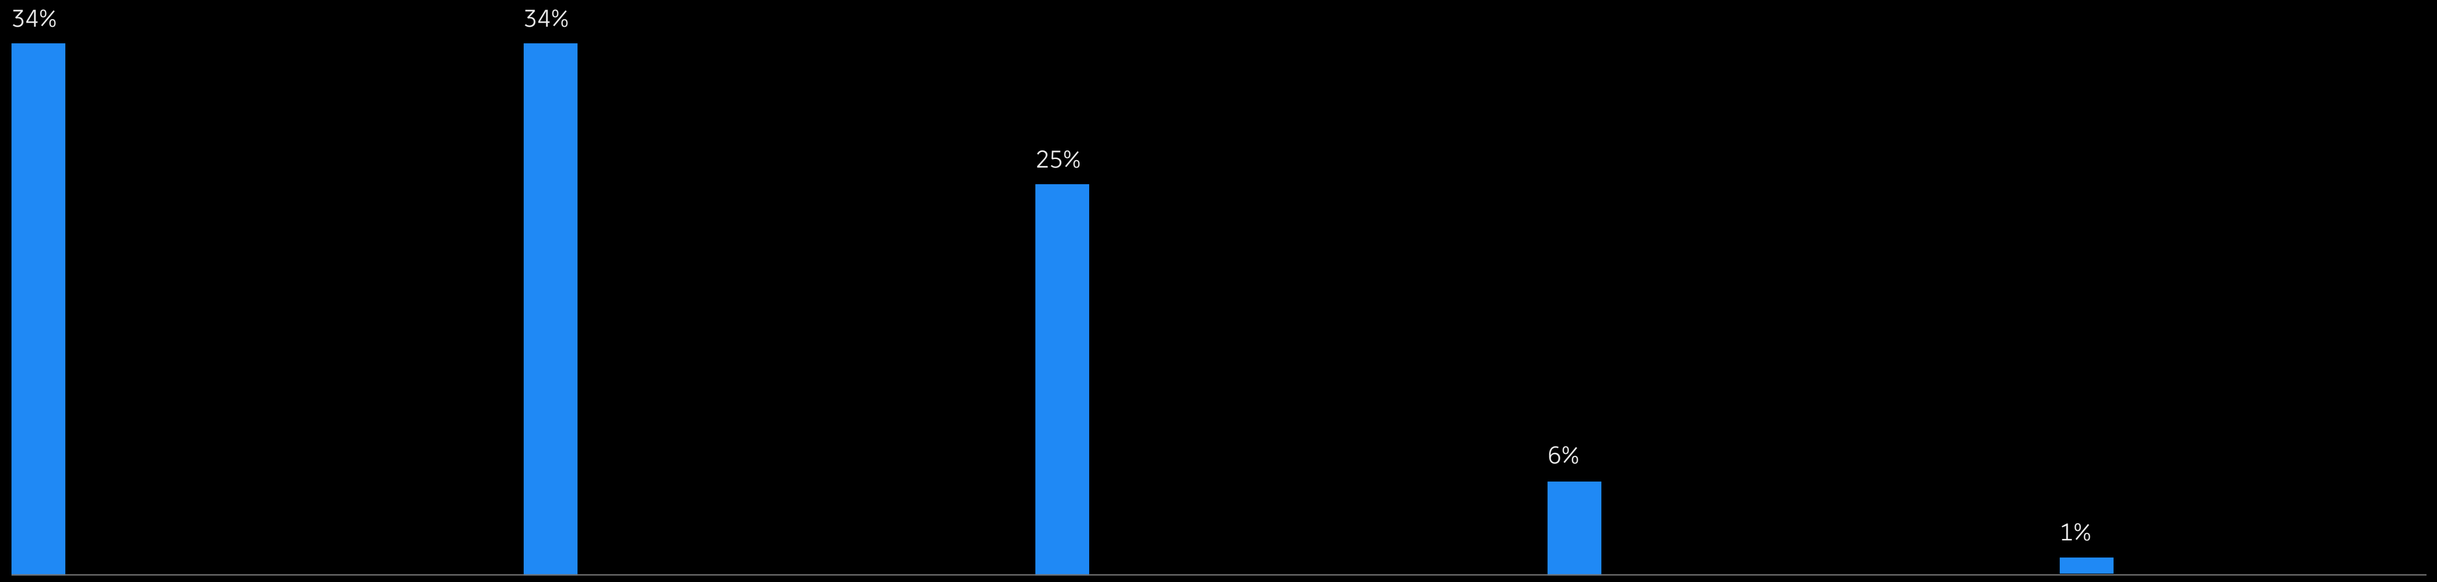
n=1,096

OpenTelemetry is extending deeper into the AI lifecycle

Among OpenTelemetry respondents, most now incorporate it into their AI observability strategy, but its strategic role varies from core pillar to secondary or selective use.

The opportunity is to extend open instrumentation beyond conventional telemetry into AI-specific behavior, maintaining continuity as AI moves from development and evaluation into production.

OpenTelemetry's role in observing AI workloads



Standard
OpenTelemetry is a core pillar of our AI workload observability strategy.

Secondary
OpenTelemetry is useful, but secondary to other approaches for observing AI workloads.

Unorganized
OpenTelemetry is adopted for some AI workloads or use cases, but there is no unified organizational approach.

Not used
OpenTelemetry is not currently part of our AI workload observability strategy.

Not applicable
We don't currently have AI workloads in production.

n=1,096

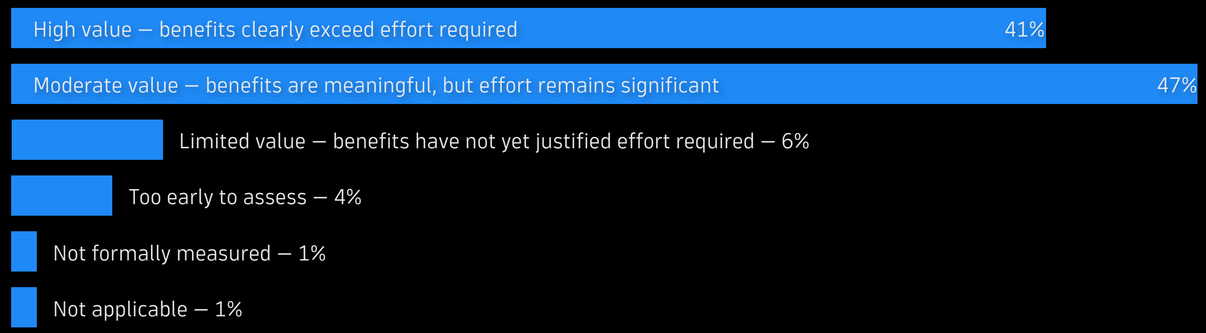
OpenTelemetry adopters report meaningful value, but scale introduces friction

Eighty-eight percent of adopters report high or moderate value from OpenTelemetry.

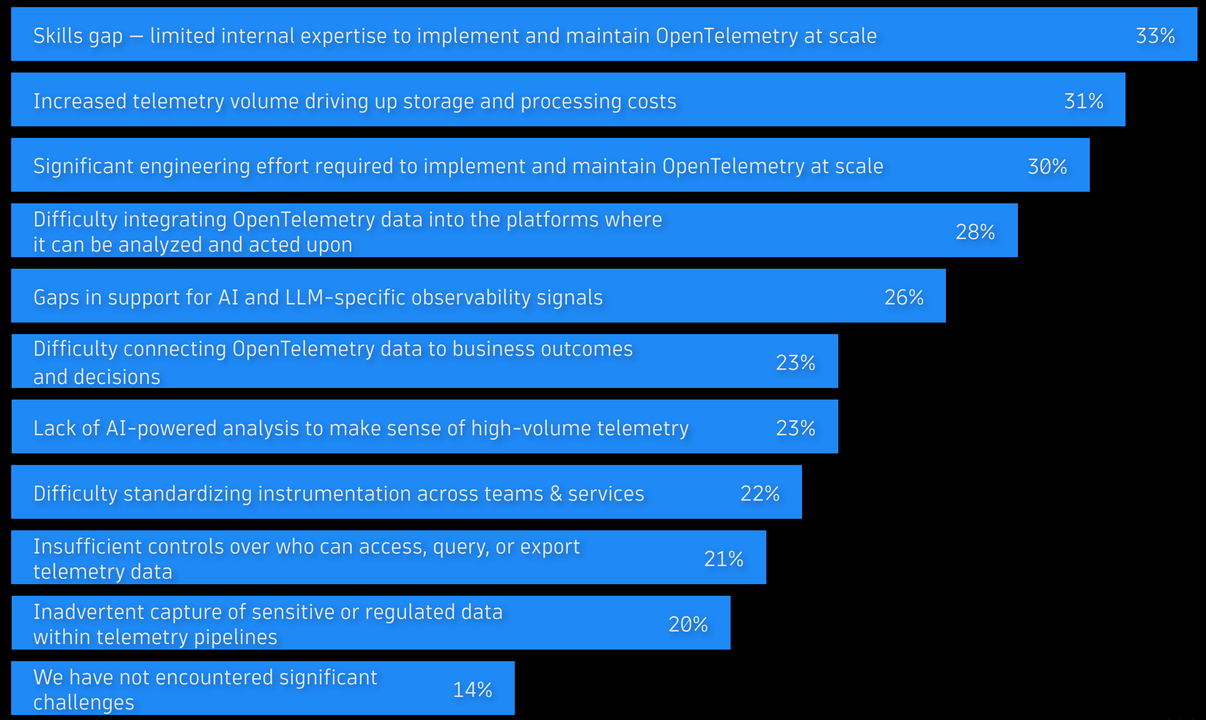
But friction emerges as implementations expand: respondents cite skills gaps, telemetry costs, and engineering effort as top barriers to scaling.

Greater OpenTelemetry maturity depends not only on generating telemetry using open instrumentation standards, but on managing it efficiently and turning it into actionable context.

Organizational value realized from OpenTelemetry



Challenges in OpenTelemetry implementation or operation



n=1,096

THE TAKEAWAY

Build open standards into the AI lifecycle

Open standards create flexibility, but reliable AI at scale requires consistent telemetry management, intelligent lifecycle control, and context across AI and production environments.

- **Standardize at the source.** Apply consistent telemetry and policies across cloud, conventional, and AI environments.
- **Control and streamline telemetry before ingest.** Automate configuration, filtering, enrichment, routing, and ongoing optimization to control cost and improve signal quality.
- **Connect AI across its lifecycle.** Link development and evaluation with production behavior and context.
- **Invest for comprehensive intelligence.** Pair open instrumentation with a flexible backend that connects technology signals to operational and business outcomes.

Open standards provide the foundation, but telemetry alone can't prove that AI is working as intended. End-to-end observability turns that foundation into the intelligence required to operate AI reliably and deliver better business outcomes.

Reliable AI requires more from observability

AI is expanding what enterprises can automate, optimize, and achieve. But as organizations depend more on AI to deliver consequential outcomes, they need to demonstrate that it behaves as intended. Observability must evolve in response, connecting AI behavior with operational and business context so organizations can scale AI with greater reliability, control, and confidence.

The results point to a higher standard for observability:



Deliver reliable experiences

AI must work as intended, not merely remain available. Observability must continuously evaluate quality and behavior and connect AI with the applications and infrastructure it affects.



Move faster with greater autonomy

AI can accelerate decisions and action, but greater authority requires causal intelligence, trusted data, validation, governance, and appropriate oversight.



Control cost and risk

AI increases resource demands and expands the attack surface. Observability must make consumption attributable and expose actual runtime risk so teams can prioritize what matters.



Preserve flexibility as AI evolves

Open standards can maintain continuity across technologies and lifecycle stages, while end-to-end intelligence connects those signals to operational and business impact.

Together, these capabilities help organizations understand AI behavior, verify results, govern action, and connect AI performance to the outcomes the business expects.

Observations and outlooks

servicenow

Trust enables autonomy to scale

"Most organizations have already introduced AI into their operations. The challenge now is building the trust framework that allows AI to take on greater responsibility over time.

Observability provides the context, governance establishes accountability, and workflows turn intelligence into action. Together, they create the foundation for trusted autonomous operations at scale."

— **Adrian Phillips**, Senior Director Marketing, Technology Workflows, ServiceNow

Building trust in AI-driven decisions requires a framework that combines operational context, governance, and workflow orchestration. AI agents can only act reliably when they understand the systems and services around them, operate within clear guardrails, and produce transparent, accountable outcomes.

Observability must provide operational context and a real-time understanding of system behavior, dependencies, business impact, and risk. Connected to workflows and governance, real-time observability is a crucial component to scaling autonomy with trust.

Microsoft

Governance must be designed in

"As models and agents move into production, organizations need visibility into how these systems actually behave, alongside strong identity, policy, security, and evaluation controls.

Open, interoperable observability helps connect runtime behavior with those controls, giving enterprises better evidence to operate, evaluate, and improve AI systems responsibly at scale."

— **Uli Homman**, CVP, Distinguished Architect, Cloud + AI, Microsoft

AI governance cannot be retrofitted. When organizations move models and agents into business-critical operations, governance must evolve in parallel. Observability provides evidence of what actually happened across models, agents, tools, applications, infrastructure, and business processes.

Open, interoperable approaches connect runtime signals with security, policy, evaluation, and business context across the technology estate, while runtime enforcement helps ensure policies are applied when actions occur. Governance then becomes a continuous operational capability that helps organizations scale AI responsibly with greater confidence.

Red Hat

Open standards provide continuity

"As AI expands across hybrid and multi-cloud environments, open standards and open platforms become essential for maintaining visibility and control.

By standardizing telemetry on an open foundation, organizations can reduce fragmentation, strengthen governance, retain control of telemetry data, and evolve AI architectures without sacrificing observability or creating dependency on a single vendor ecosystem."

— **Joe Fernandes**, Vice President and General Manager, AI Business Unit, Red Hat

Vendor lock-in and tool fragmentation are no longer just efficiency headaches. In the AI era, they're governance and visibility risks. As enterprises scale AI across hybrid and multi-cloud environments, critical new blind spots emerge. You cannot govern, secure, or optimize what you cannot see.

Open standards such as OpenTelemetry and open platforms allow telemetry, context, and governance to move consistently across hybrid architectures. These standards give IT leaders greater control over their telemetry data, more consistent governance, and the flexibility to evolve their AI strategies without tying visibility to a particular platform or environment.

accenture

Operating models must evolve

"As AI dissolves the boundary between building and running systems, enterprises need an operating model that helps navigate that convergence. We see real-time runtime context as foundational: understanding how agents behave, what they depend on, and how their actions affect outcomes.

Combined with mature delivery pipelines and strong governance, that visibility gives organizations the confidence to extend greater decision-making authority to AI."

— **Scott Wintermute**, Global Observability Lead, Managing Director — Master Technology Architect, Accenture

AI is not only accelerating operations, it is dissolving the boundary between building systems and running them. The same agents that write and deploy code increasingly create the autonomation that operates it. Operating models that treat build and run as separate disciplines will create gaps in ownership, context, and control as those boundaries disappear.

From an enterprise transformation perspective, autonomous operations require real-time awareness of how AI behaves in production: what agents do, their dependencies, and their impact on performance, quality, cost, and business outcomes. This runtime context enables better decisions as build and run automation converge.

Deloitte.

Choosing the right role for AI

"The next phase of enterprise AI is not about deploying more agents, but scaling autonomy where it creates value.

Observability provides the evidence to guide that transition, validate business outcomes, and evolve the controls and operating model required for confident, AI-driven enterprise transformation."

— **Alvaro Martin**, Observability Strategy Lead, Technology & Transformation, Deloitte

As organizations modernize their operating models, progress should not be measured by the number of models, agents or automated processes deployed, but by whether AI is improving outcomes, operating within appropriate boundaries and creating sufficient value to justify greater autonomy.

Modernizing operations requires choosing the right role for AI, automation, human expertise, and deterministic methods. Observability provides the evidence to determine where autonomy should scale, where other approaches are more appropriate, and whether those choices are delivering the intended outcomes.

aws

AI economics must be measurable

"The challenge isn't simply scaling AI. It's scaling AI efficiently. As organizations expand AI workloads, they need visibility into the relationship between performance, infrastructure consumption, and cost.

We aim to support our customers to treat observability, cloud economics, and resource optimization as a single operational discipline rather than separate initiatives."

— **Mona Chadha**, Director, Worldwide Infrastructure Partners, AWS

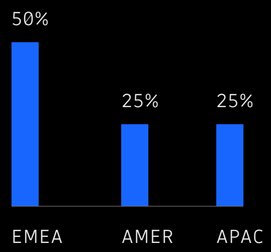
We are seeing organizations rapidly scaling AI from experimentation into production, but many are discovering that infrastructure costs rise faster than anticipated.

Organizations cannot optimize AI investments they cannot accurately measure or attribute. As AI adoption grows, leaders must understand not only whether AI systems deliver accurate outcomes, but also what those outcomes cost to produce. Cost allocation, workload attribution, capacity planning, and infrastructure efficiency should be treated as foundational architectural requirements rather than post-deployment exercises.

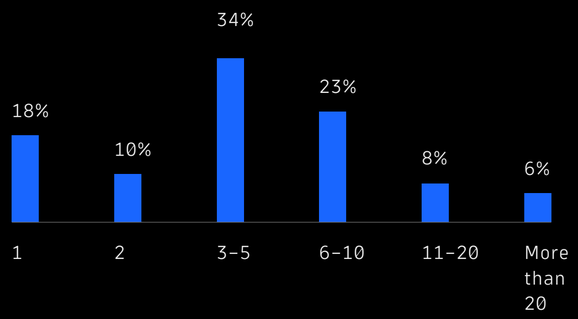
APPENDIX A

Firmographics and demographics

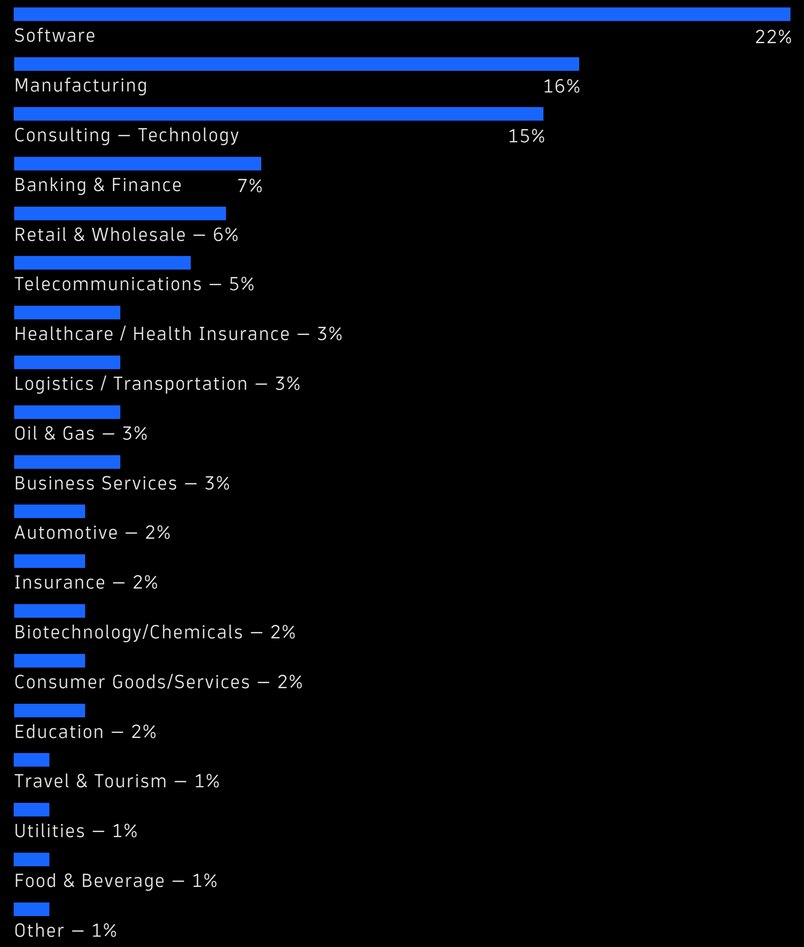
Region



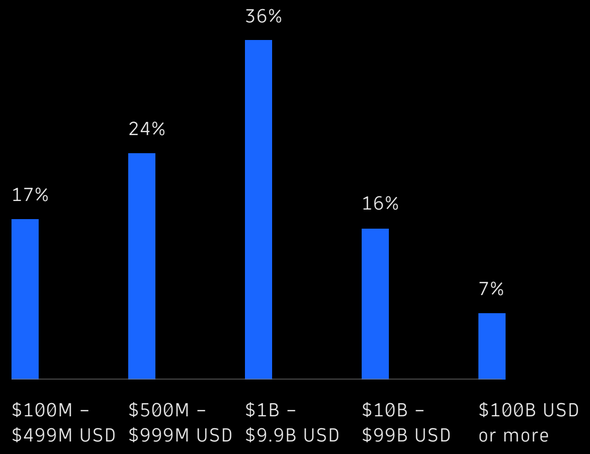
Number of countries operating in



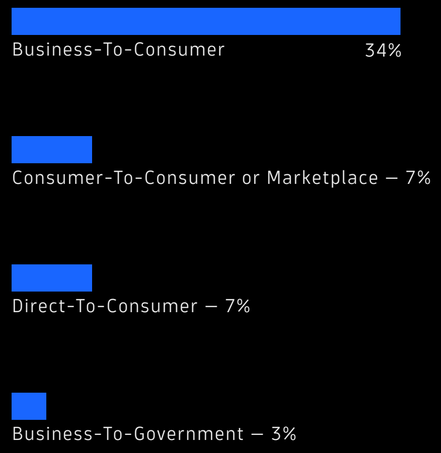
Industry



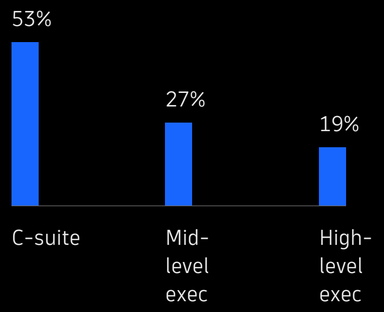
Annual organization revenue



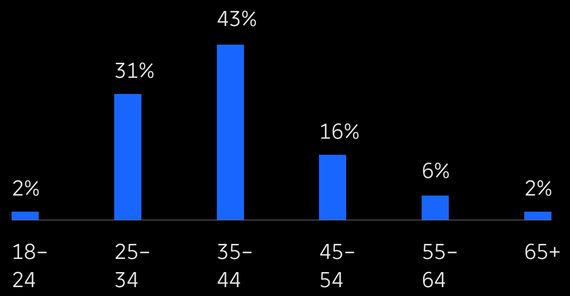
Business type



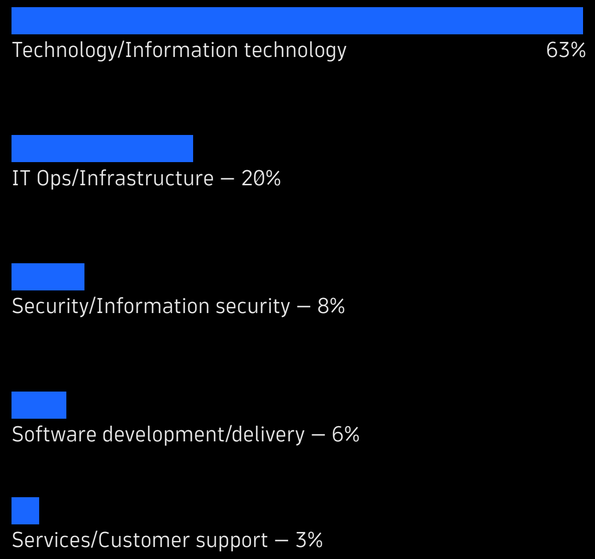
Role



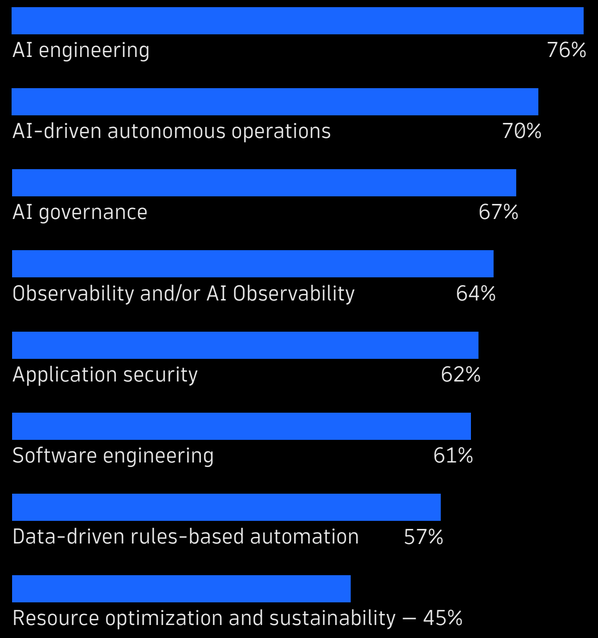
Age



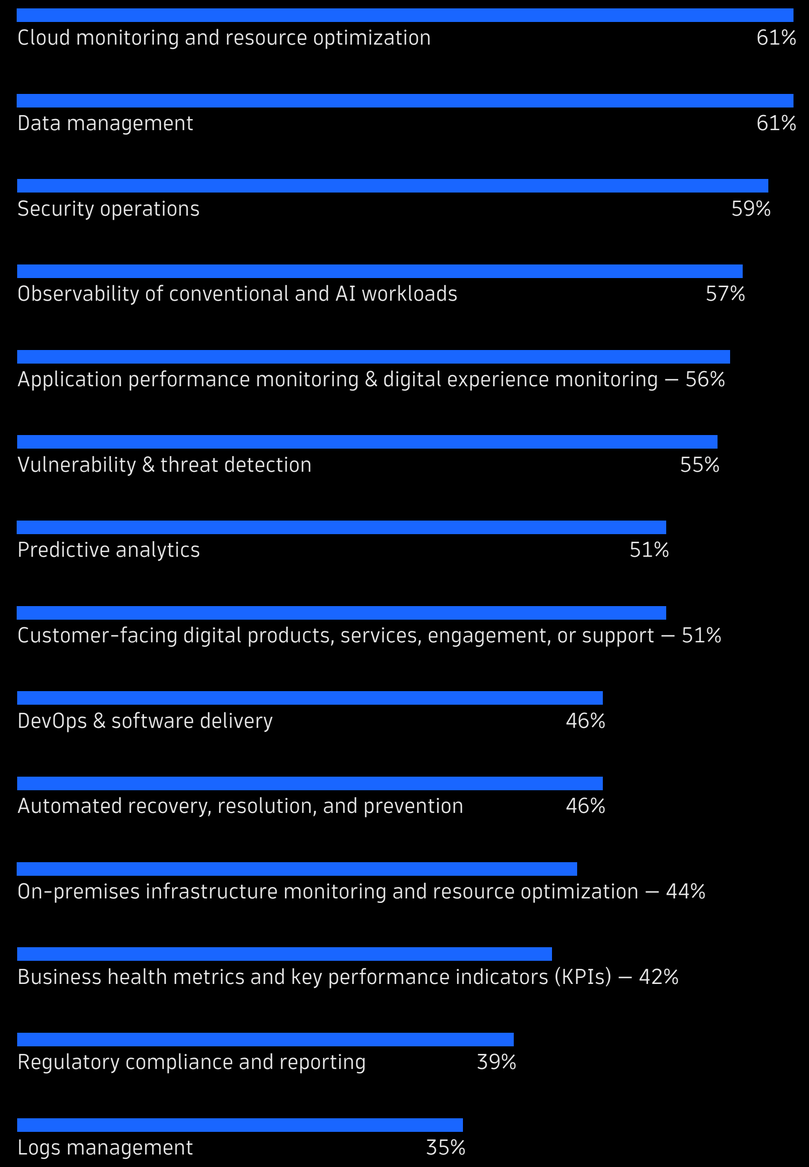
Department



Responsibility



AI-facilitated management



APPENDIX B

Regional signals

The research reveals broad consistency in observability priorities and maturity across regions, alongside several meaningful differences in the complexity organizations experience and the outcomes they derive from observability and AI.

THE AMERICAS

Skills gaps and faster problem resolution

US (100), Brazil (100), Mexico (100)

- Organizations in the Americas are more likely to report that AI presents the **need for greater skills to manage observability**. More than half of respondents in the Americas cite this challenge at 56%, compared with 43% in EMEA and 46% in APAC.
- Organizations in the Americas also report greater growth in the number of monitoring tools required, at 45%, versus 35% in EMEA and 42% in APAC.
- That pressure accompanies strong reported operational outcomes. Faster issue detection and resolution is the leading observability-enabled business outcome in the Americas, cited by 60%, compared with 53% in APAC and 48% in EMEA.
- Automated root-cause analysis and faster issue resolution also lead the ways AI has advanced autonomous operations in the Americas, at 39%.

EUROPE, MIDDLE EAST, AFRICA (EMEA)

Less pressure from security threats

United Kingdom, Germany, France, Italy, UAE, Saudi Arabia

- EMEA organizations are less likely to identify the growing frequency and sophistication of security threats as a source of increased technology complexity. Only 29% of EMEA respondents cite this factor, compared with 39% in the Americas and 43% in APAC.
- EMEA also reports fewer challenges related to the complexity and resource demands of managing observability instrumentation at scale. This challenge is cited by 29% in EMEA, compared with 36% in AMER and 38% in APAC.
- For autonomous operations, predictive detection before user or system impact is the leading reported AI contribution in EMEA, at 39%.
- Of the 741 global respondents responsible for application security, 33% in EMEA set more aggressive 24-hour remediation targets than other regions (29% in the Americas and 20% in APAC). But 37% in the Americas report actually meeting the sub-24-hour remediation window versus 28% in EMEA and 22% in APAC.

AI-driven complexity and security impact

India, Japan, Australia

- APAC respondents are more likely to identify expanding AI, LLM, and agentic workloads as a source of technology complexity. 47% cite AI expansion, compared with 41% in both the Americas and EMEA.
- APAC respondents are also more likely to cite the growing frequency and sophistication of security threats, at 43%, versus 39% in the Americas and 29% in EMEA.
- The region reports similar patterns when asked how AI affects operational complexity. Fifty-six percent say AI has increased the complexity of the technology environment, compared with 47% in the Americas and 46% in EMEA.
- More than half of APAC respondents (52%) cite increasingly unpredictable AI application and model behavior, compared with 43% in each of the other regions.
- APAC respondents also report distinctive outcomes. Improved reliability and uptime of business-critical systems leads observability-enabled business outcomes at 58%, compared with 51% in the Americas and 48% in EMEA. Comprehensive security vulnerability detection and remediation is the region's leading AI contribution to autonomous operations, at 37%.

APPENDIX C

The influence of company size

Company size differentiates investment and ambition much more than it differentiates the underlying complexity and data-management problems organizations face.



Larger organizations invest more in observability

Observability spending rises with company size. 17% of organizations with $10 billion or more in annual revenue spend at least $10 million annually on observability, compared with 6% overall.



Autonomy ambitions rise with organizational scale

Half of organizations with $10 billion or more in annual revenue target fully autonomous operations within five years, compared with 29% of organizations with $100 million–$499 million in revenue.



Complexity and data challenges cut across company sizes

Greater scale does not produce equally large differences in the challenges organizations report. Sources of technology complexity are similar across revenue levels, including AI expansion, multi-cloud environments, security threats, and skills gaps. Organizations also report similar data-ingest and management challenges regardless of annual revenue.

APPENDIX D

Industry-specific priorities

Industry differences emerge most clearly in how organizations apply and invest in observability, the complexity of their technology environments, and their progress toward operationalizing AI.

SOFTWARE

APM investment and AI execution

- 36% plan to prioritize APM over the next year, compared with 26% of respondents in other industries. Software respondents are less likely to prioritize log management and analysis, at 22% versus 32% elsewhere. (n=269)
- 35% currently address end-to-end understanding of AI agent decisions and actions, compared with 25% of respondents in other industries. 35% also address AI-related cost management, versus 25% elsewhere. (n=266)
- 32% cite difficulty explaining or auditing AI-driven decisions as a barrier to autonomous operations, compared with 20% of respondents in other industries. (n=212)

TECHNOLOGY CONSULTING

Autonomous operations and observability maturity

- 31% use at least four cloud providers across their technology environment, compared with 17% of respondents in other industries. For AI workloads specifically, 21% use at least four providers, versus 11% elsewhere. (n=175)
- 21% say their organizations have reached fully autonomous operations, compared with 11% of respondents in other industries. (n=175)

MANUFACTURING

Narrower multi-cloud footprints

- 11% use at least four cloud providers, compared with 20% of respondents in other industries. (n=188)

BANKING AND FINANCE

Business context

- 59% use observability to track business processes and KPIs, compared with 40% of respondents in other industries. (n=87)

TELECOMMUNICATIONS

Multi-cloud and hyperscaler optimization

- 63% plan to prioritize multi-cloud and hyperscaler optimization over the next year, compared with 38% of respondents in other industries. (n=56)

Turn AI ambition into reliable outcomes

Dynatrace and Arize, now part of Dynatrace, help organizations know what their AI is doing, what it costs, and whether it's accurate, from evaluation to business outcome.

Explore Dynatrace AI Observability

See how AI observability connects applications, LLMs, agents, and the supporting technology stack with the context needed to understand behavior, quality, reliability, cost, and impact.

Explore AI Observability

- **Go from insight to reliable action** with Dynatrace Intelligence, an agentic operations system that fuses deterministic insight with agentic action. Explore Dynatrace Intelligence.
- **Learn what your agents are actually doing using Arize**, the agent observability and evaluation platform. Explore Arize.
- Explore the Dynatrace platform | Start a free trial

dynatrace

Dynatrace is advancing observability for today's digital businesses, helping to transform the complexity of modern digital ecosystems into powerful business assets. By leveraging AI-powered insights, Dynatrace enables organizations to analyze, automate, and innovate faster to drive their business forward. Learn more at www.dynatrace.com.

blog @dynatrace linkedin.com/company/dynatrace instagram.com/dynatrace

11.05.26 bae28903-ebk-cs