dynatrace

RESEARCH REPORT 2026

The pulse of agentic AI

Building trust in autonomous operations from pilot to production

What's inside

INTRODUCTION

Agentic AI for autonomous operations in 2026

Autonomous operations are moving from experimentation into production as agentic AI enables systems to reason, learn, and execute complex tasks with minimal human intervention.

The goal of autonomous systems is to boost resilience, scalability, and customer experiences. But enterprises are discovering that autonomy without control can create more risk and reliability concern than value.

In this inaugural report on the pulse of agentic AI, we focus on project maturity, challenges, and scale to clarify:

- How organizations are **adopting, prioritizing, and deploying** agentic AI
- Where teams are **making progress and hitting barriers**
- When and why organizations **keep humans in the loop**
- How organizations **measure AI success**
- Why **observability is becoming a critical control plane** for operationalization

Observability, the real-time understanding of system behavior through connected data, has always played a role in system performance. But with the variability of generative models, observability must shift to a foundational control layer as organizations scale projects from pilot to production.

Why it matters

With a clear understanding of how organizations are advancing agentic AI deployments, the challenges they face, and where they're finding value, this report provides a data-driven resource for leaders looking to make strategic decisions, enabling them to advance and scale trustworthy autonomous operations using real-time observability.

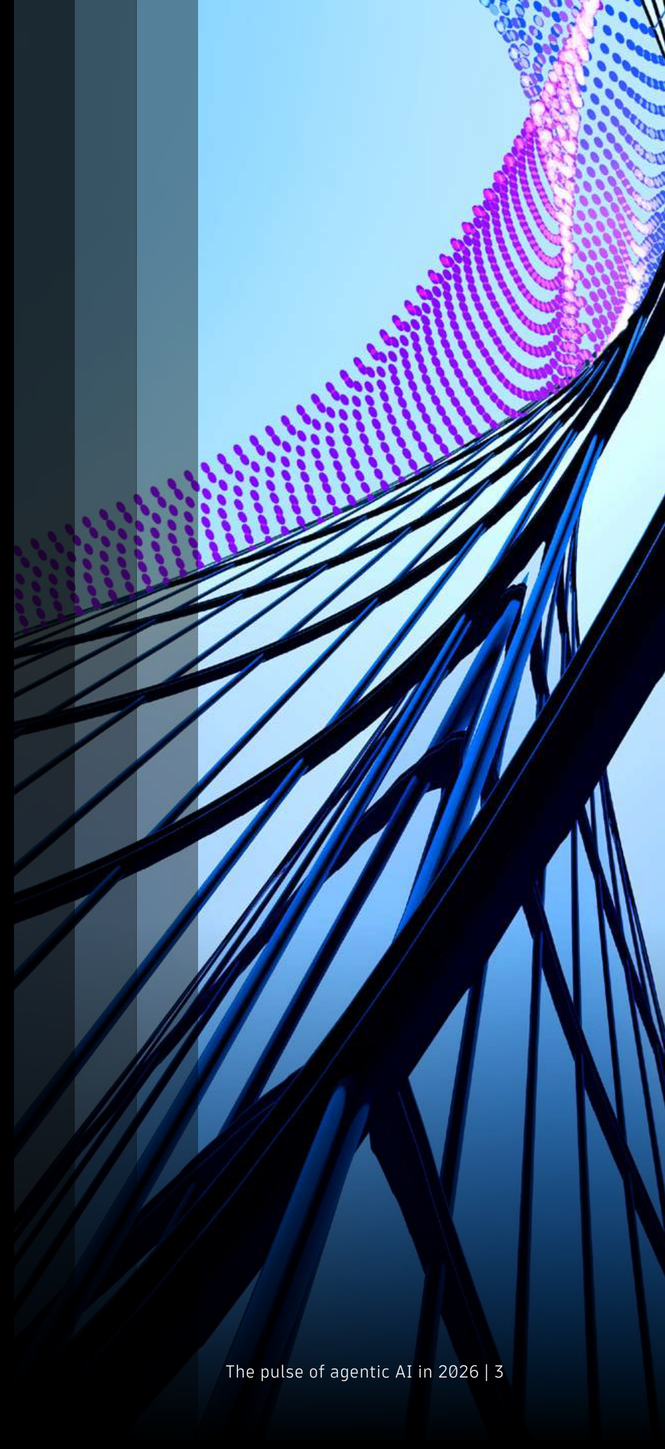

72%

use AI agents
for IT operations
and DevOps

50%

use agentic AI for both
internal and external
use cases

74%

expect agentic AI budgets
to increase in the next year

dynatrace

CHAPTER ONE

Agentic AI is rapidly expanding beyond IT operations

Agentic AI has gained the most traction in mission-critical proving grounds such as IT operations, software engineering, cybersecurity, and support—applications that can least tolerate failure. This importance emphasizes the need for foundational control, governance, and comprehensive observability of agentic systems.

KEY INSIGHTS

- Infrastructural **IT-related** use cases currently dominate, but **customer-facing** use cases are growing rapidly, demonstrating their importance and value.
- Agentic AI is attracting **priority funding,** with three-quarters expecting to increase agentic AI budgets in the next year.

An operational best practice is to prioritize the most measurable, repeatable use cases with the greatest expected ROI. The findings show that organizations are prioritizing foundational internal functions like IT operations, data processing and reporting, and cybersecurity.

Repeatable customer-related functions, like services and personalized engagement, rank as the top external applications. Nearly half (49%) use agentic AI for both, although 33% currently emphasize internal use cases.

Current agentic AI uses

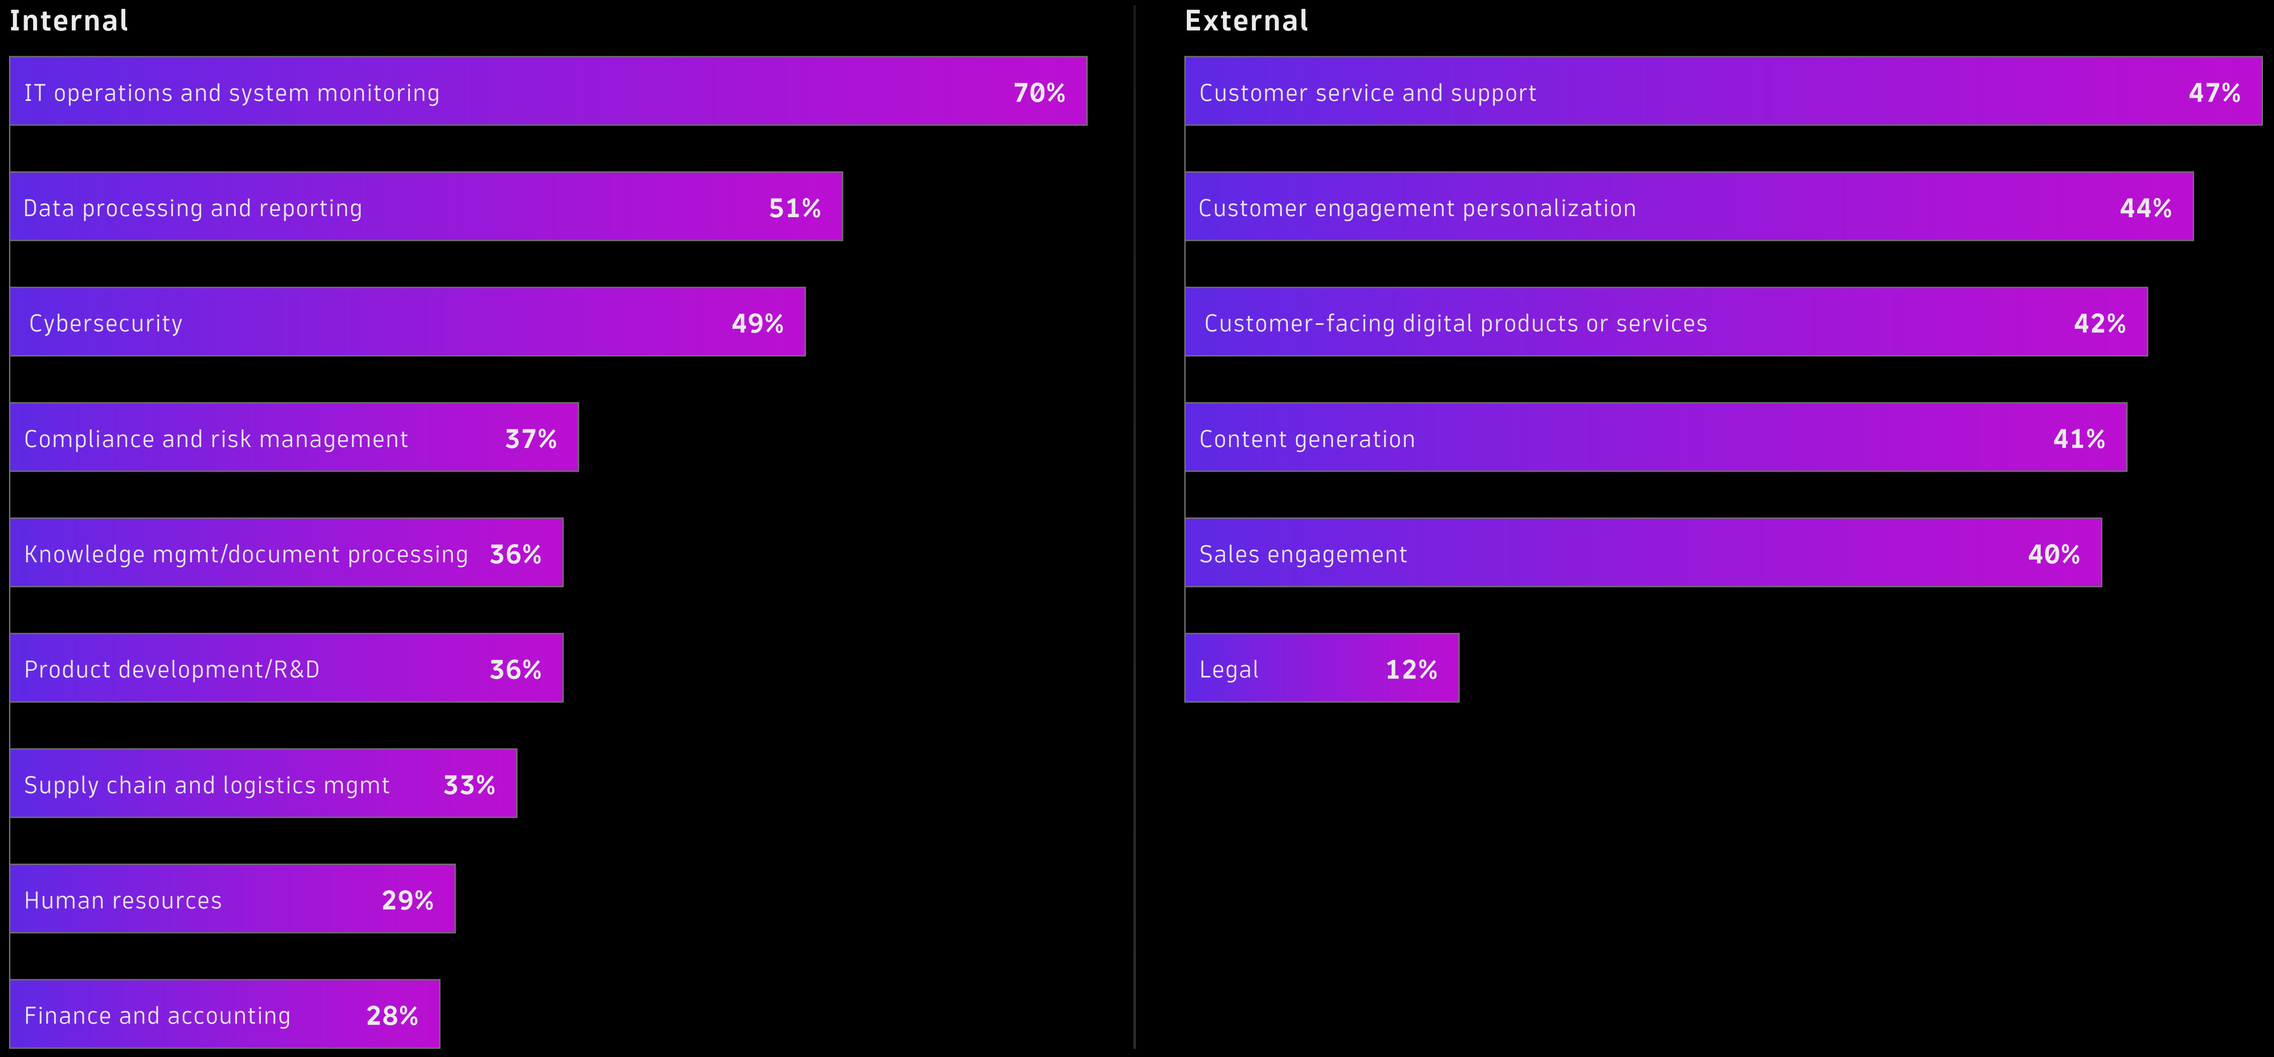

Customer-facing applications and budgets are rapidly rising

While internal use cases currently dominate, the fastest-growing implementations over the next five years are externally exposed, demonstrating an imperative to realize agentic AI value in customer-facing applications.

AI automation for legal, which is currently the least common use case, is projected to triple in the next several years, indicating strong growth potential for non-IT use cases.

Agentic AI use cases timeline

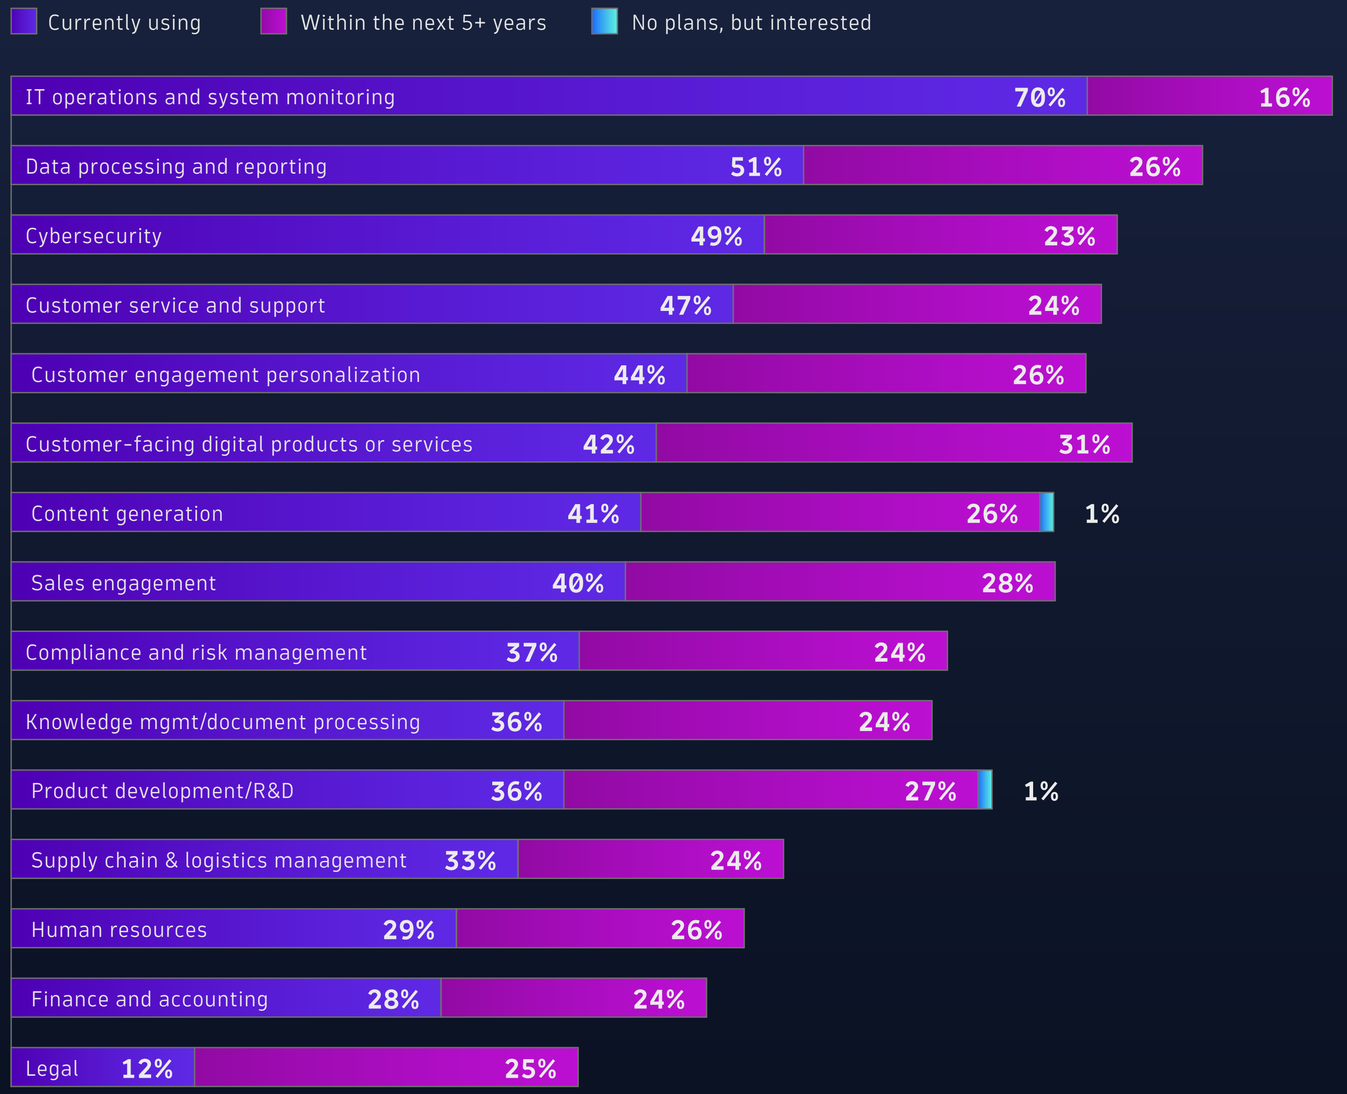

Agentic AI budgets remain robust

With current spending averaging $2-5M, the majority expect to increase their budgets, and nearly half expect spending to increase by another $2-5M in the coming year. This is a clear signal of sustained priority.

Current agentic AI budget

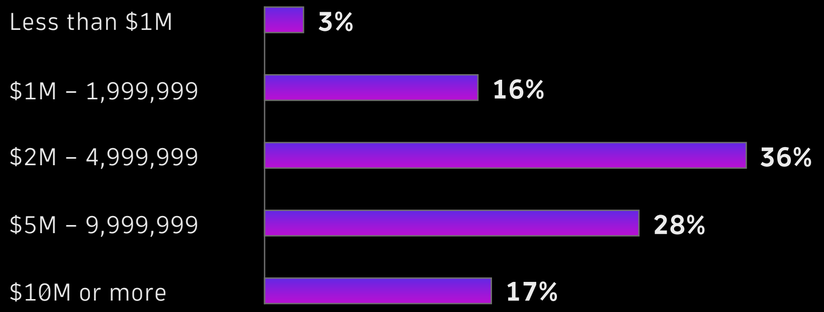

Projected agentic AI budget increase

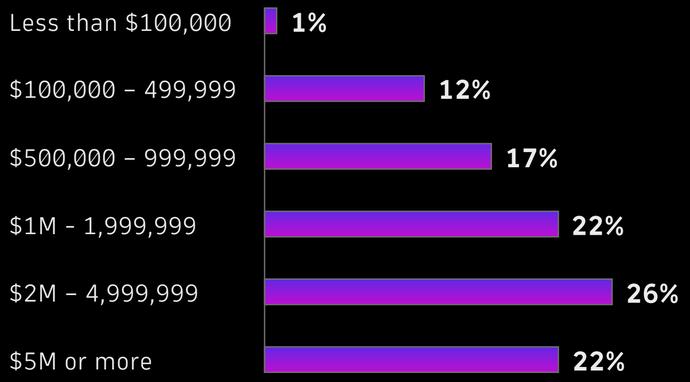

Greatest ROI expected

44%

IT operations/system monitoring

27%

Cybersecurity

25%

Data processing & reporting

Control starts with goal-based automation

Agentic systems represent a quantum leap in productivity, efficiency, and scale, but also complexity and often unpredictability. Gaining visibility into agent interactions and decisions helps organizations maintain control.

- **Develop clear goals**, guardrails, and oversight for agentic projects to avoid disorder and maintain reliability.
- **Design systems** so agent-to-agent communications remain accountable to humans.
- **Manage the nexus** of autonomous, real-time interactions with observability.

50%

of respondents have projects in production for limited use cases

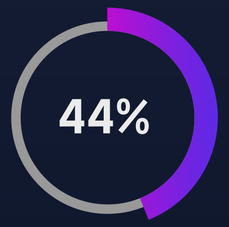

have projects in broad adoption in select departments

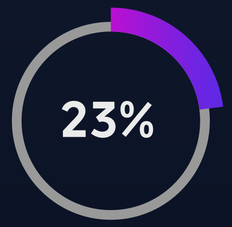

have projects in mature, enterprise-wide integration

dynatrace

CHAPTER TWO

Autonomous operations is making progress—and hitting barriers

Although agentic AI is still an emerging technology and adoption is relatively shallow, organizations are quickly moving toward maturity. But organizations face technical barriers to governing autonomous operations and monitoring agentic AI behavior, such as limited ability to set automation rules and limited visibility, which are hindering progress.

KEY INSIGHTS

- **50%** have projects in production for limited use cases; **44%** have projects in broad adoption in select departments.
- Nearly three-quarters (**72%**) have 2-10 agentic AI projects, indicating that agentic AI is still growing.
- **Trust and lack of observability** are top technical barriers to production.

Agentic AI adoption is still growing as organizations evolve small automation projects into autonomous operations. With adoption still in early phases, nearly three-quarters (72%) report having between 2-10 active projects, and about a quarter have 11-21 or more projects.

Most of these projects are in some phase of pilot or limited production, suggesting that organizations are feeling the benefits of agentic AI and looking to expand and perfect its implementation.

Agentic AI adoption and maturity

Active agentic AI projects

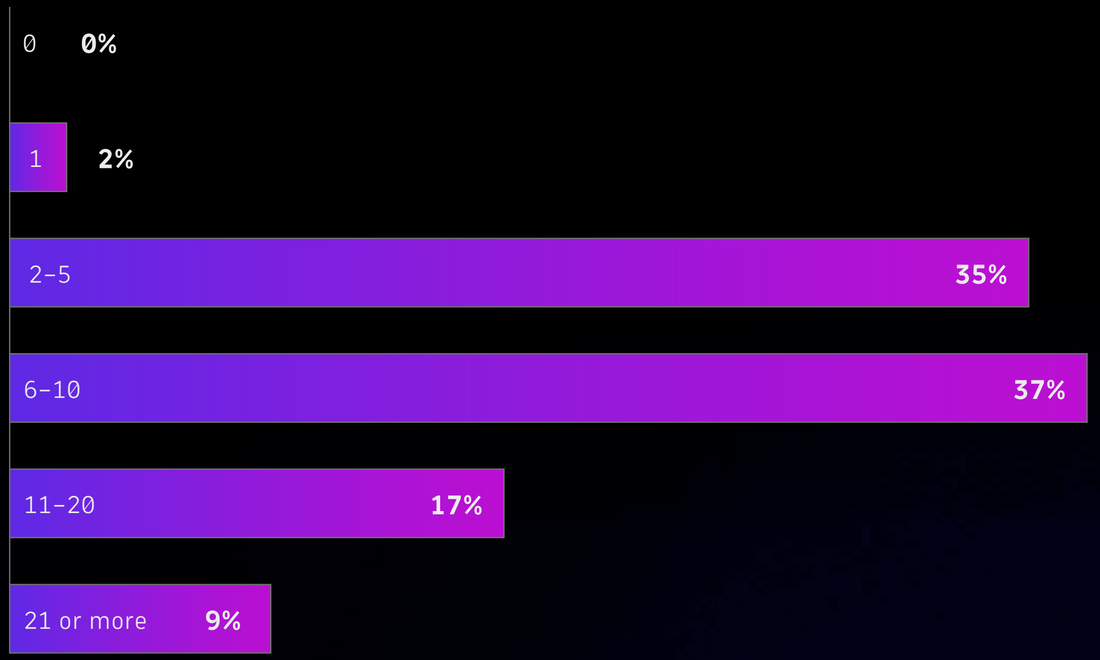

Agentic AI project stages of development

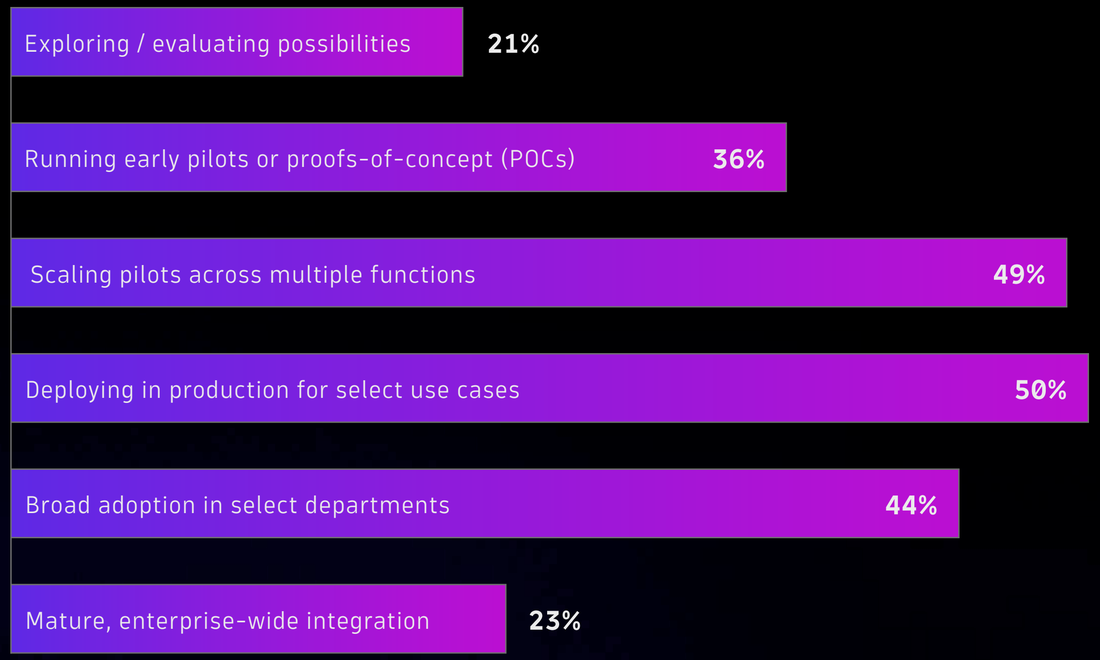

Production readiness is growing fast across all use cases

Consistent with current investment patterns and expected ROI, IT operations and system monitoring use cases show the largest share of production-level readiness. For most functions, about half of those projects are in production or being operationalized.

Agentic AI production readiness

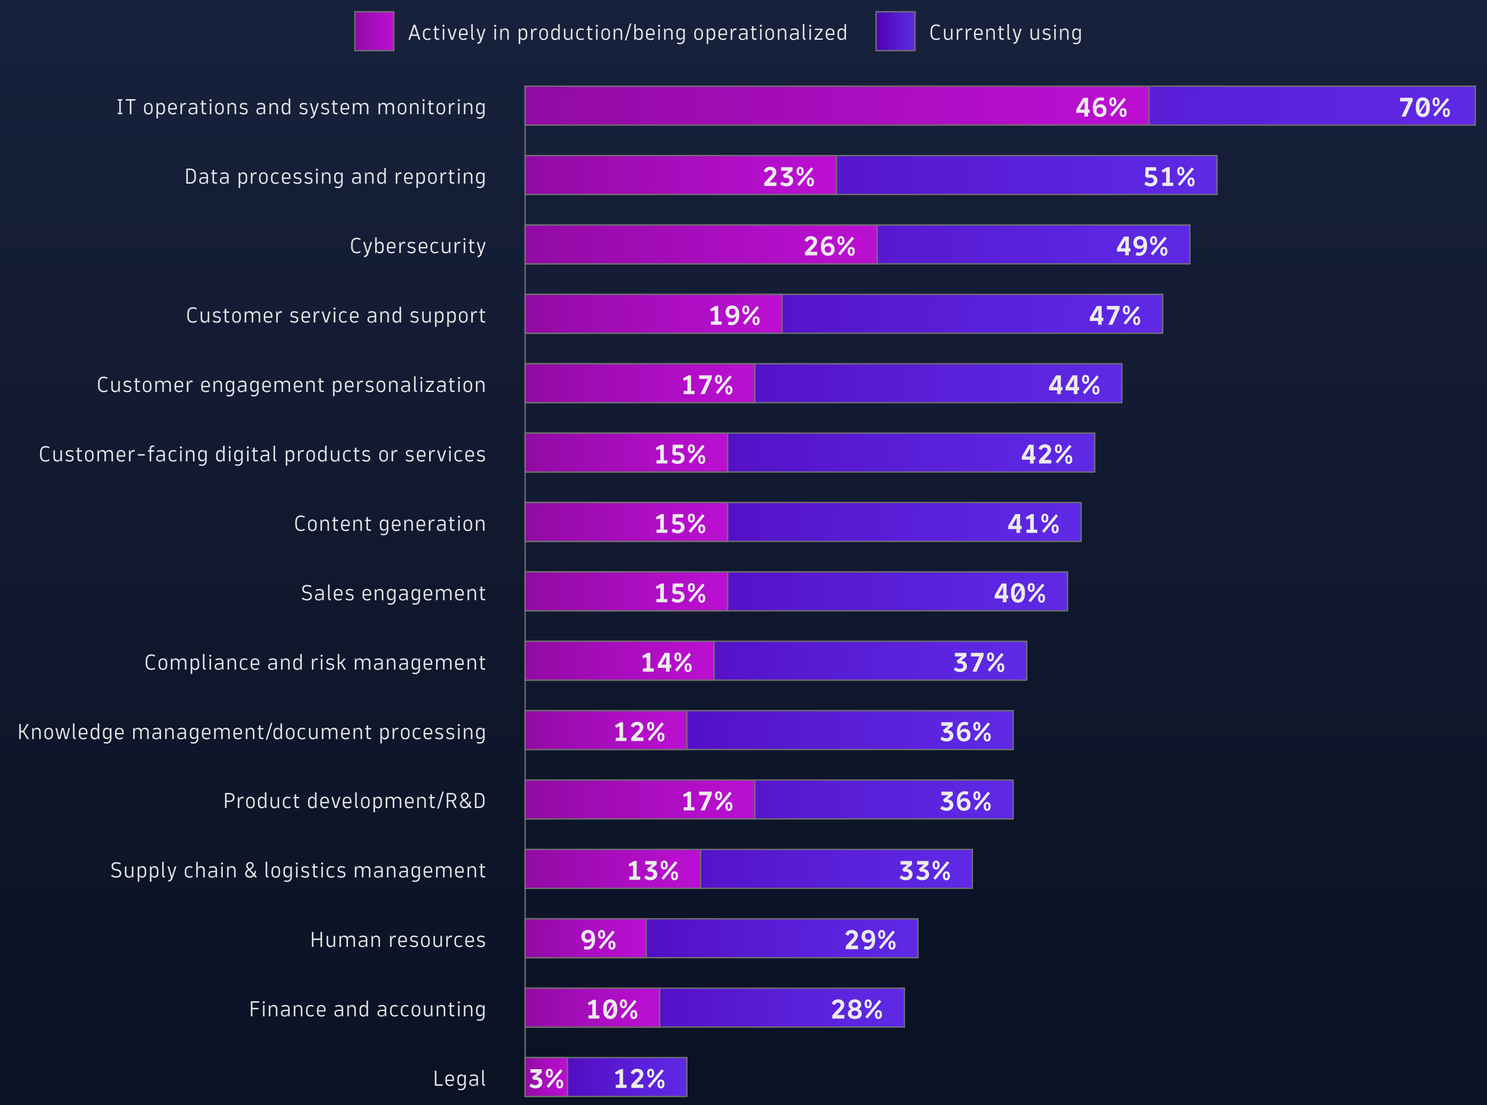

Production criteria prioritize security and reliability

Security and data privacy and AI accuracy are the top criteria for moving projects from pilot to production. Ethics and fairness rank lowest in priority, indicating that for many, technical accuracy currently outweighs ethical concerns.

Criteria for moving agentic AI from pilot to production

Technical barriers to production include observability concerns

The top technical barriers to advancing projects relate to establishing trust, such as setting clear rules for when an AI agent should act autonomously versus when it requires human approval.

Other technical challenges include limited visibility, difficulty monitoring, coordinating troubleshooting efforts, and difficulty tracing downstream effects, indicating that organizations are struggling with observability and traceability of agentic AI behavior.

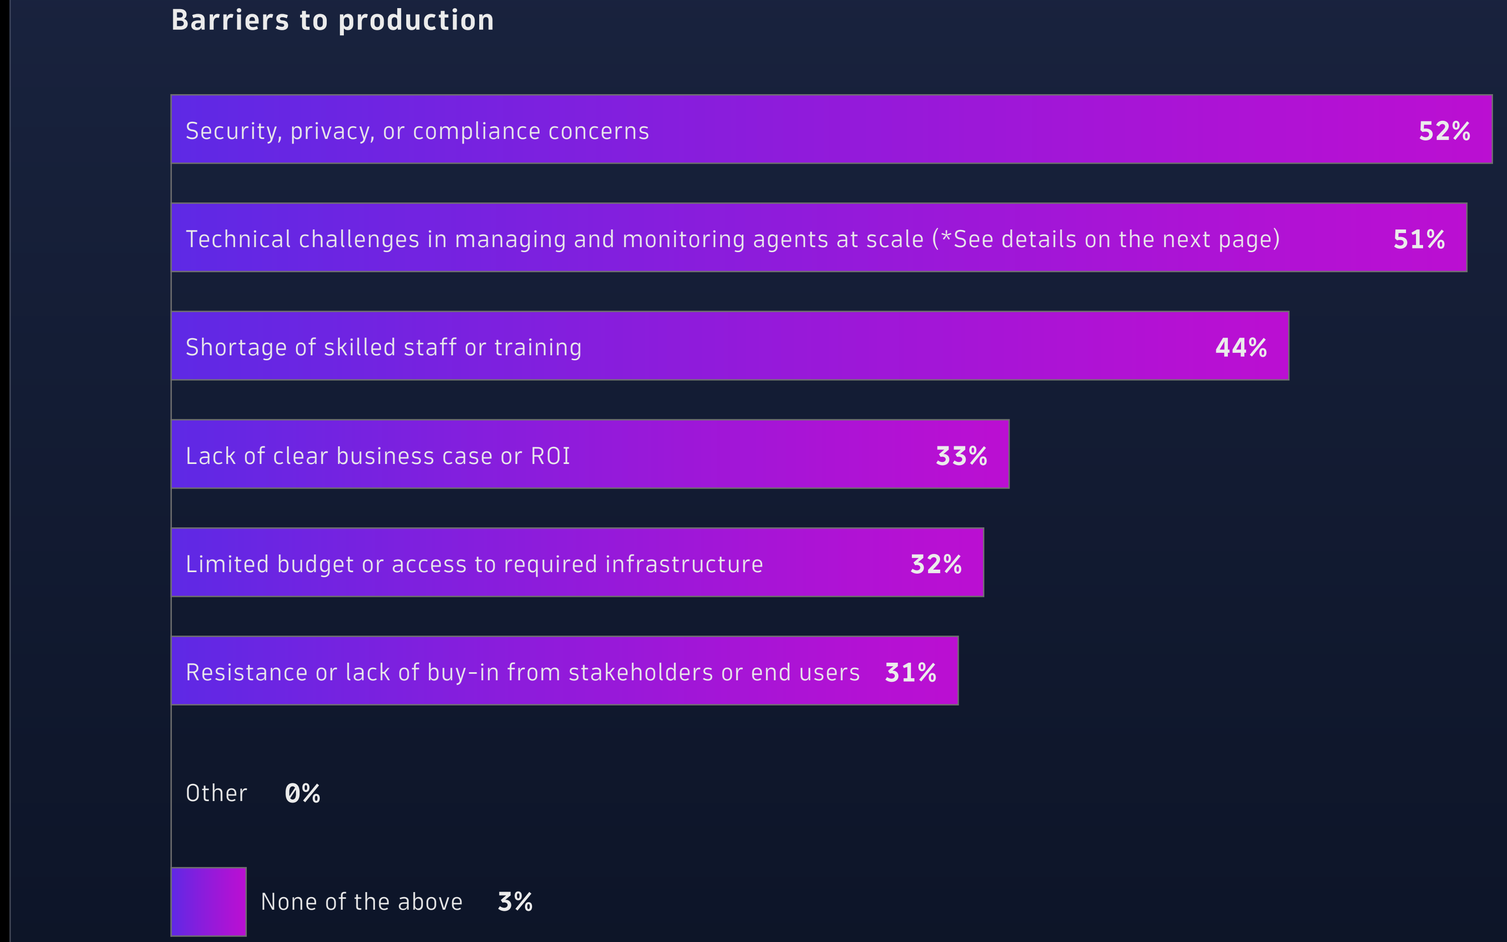

Technical barriers to production — detail

Build trust in autonomous operations incrementally

While expectations ride high for agentic AI projects, building trust in autonomous operations is best done incrementally. By implementing functionality in phases, organizations can reliably prepare the environment for fully autonomous operations.

- **Start with preventive operations** and recommendation-driven workflows.
- Using a human-in-the-loop model, **progress toward supervised autonomy** by hardening services, data sources, and contextual signals.
- **Graduate to full autonomy** using advanced observability data to understand each component in context and verify its performance in real time.

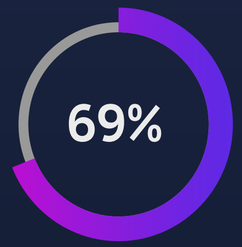

of agentic AI decisions are currently verified by a human

Top agentic AI validation measures

50% Data quality checks

47% Human review of agentic AI outputs

41% Monitoring AI outputs for drift or anomalies

dynatrace

CHAPTER THREE

Agentic AI trust and automation: Human oversight is intentional and enduring

AI-powered automation is the clear direction of travel, which raises the stakes for trust, governance, and control. But organizations are intentional about human involvement across the board and foresee a future of human-AI partnership.

KEY INSIGHTS

- More than **two-thirds** of agentic AI decisions are currently **verified by a human.**
- Data quality checks, human review, and monitoring for drift are the **top validation measures.**
- The future of AI-human interaction is **collaborative.**

Agentic AI powered by generative models is fundamentally probabilistic in nature and can produce hallucinations in the process. For more than two-thirds of respondents, humans validate agentic AI decisions using sophisticated systems—data quality checks, AI drift reviews, intelligent logs and traces—creating the primary trust mechanism that enables scale. These findings are consistent for organizations of all types, regardless of company size, region, and revenue.

Validation measures vary

Lower-revenue organizations ($100-$499M) are more likely to use security validations to verify agentic AI decisions (55%), whereas higher-revenue organizations ($10B or more) are more likely to validate these decisions by recording comprehensive logs and traces (31%).

Agentic AI validation measures

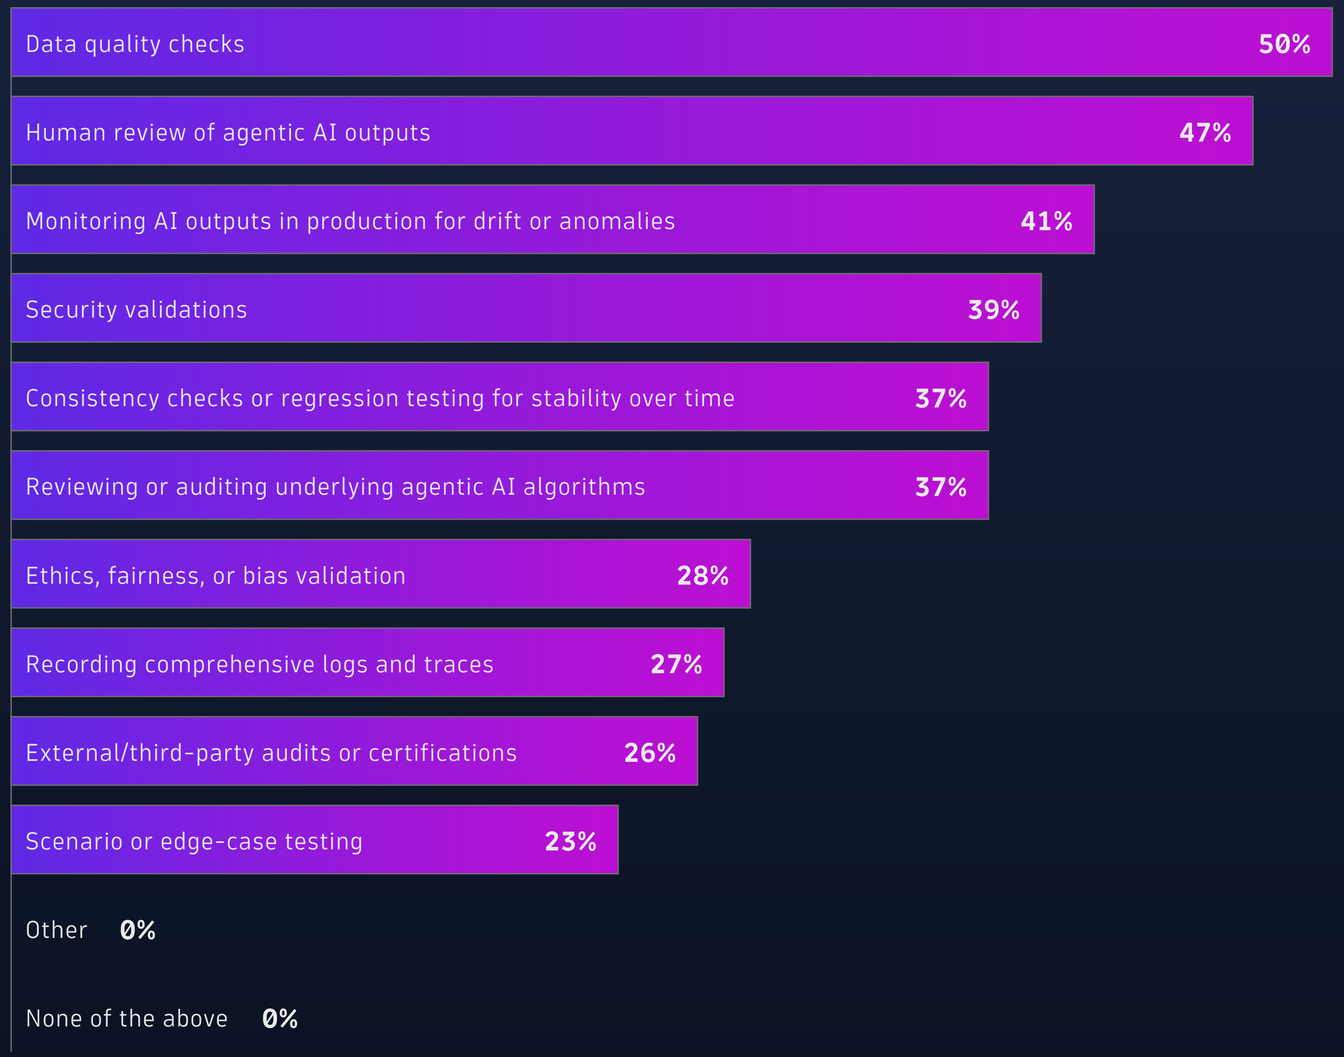

A mix of autonomy and human supervision

Nearly two-thirds of respondents (64%) are currently building both fully autonomous and semi-autonomous agents supervised by humans.

Only 13% are building solely autonomous AI agents, and about a quarter (23%) are building only human-supervised agents.

64% are building both **autonomous and human-supervised** agents

23% are building only **human-supervised** agents

13% are building solely **autonomous** agents

Automation goals favor humans-in-the-loop

Data processing, cybersecurity, and customer service/support functions currently show the highest rates of AI autonomy, but organizations foresee building greater autonomy for knowledge management and customer-facing digital products in the future.

In general, respondents expect a higher percentage of human-involved automation (60/40 on average) for business-related applications than IT and routine customer-support-related functions (50/50 on average).

Goals for AI agents revolve around more automation

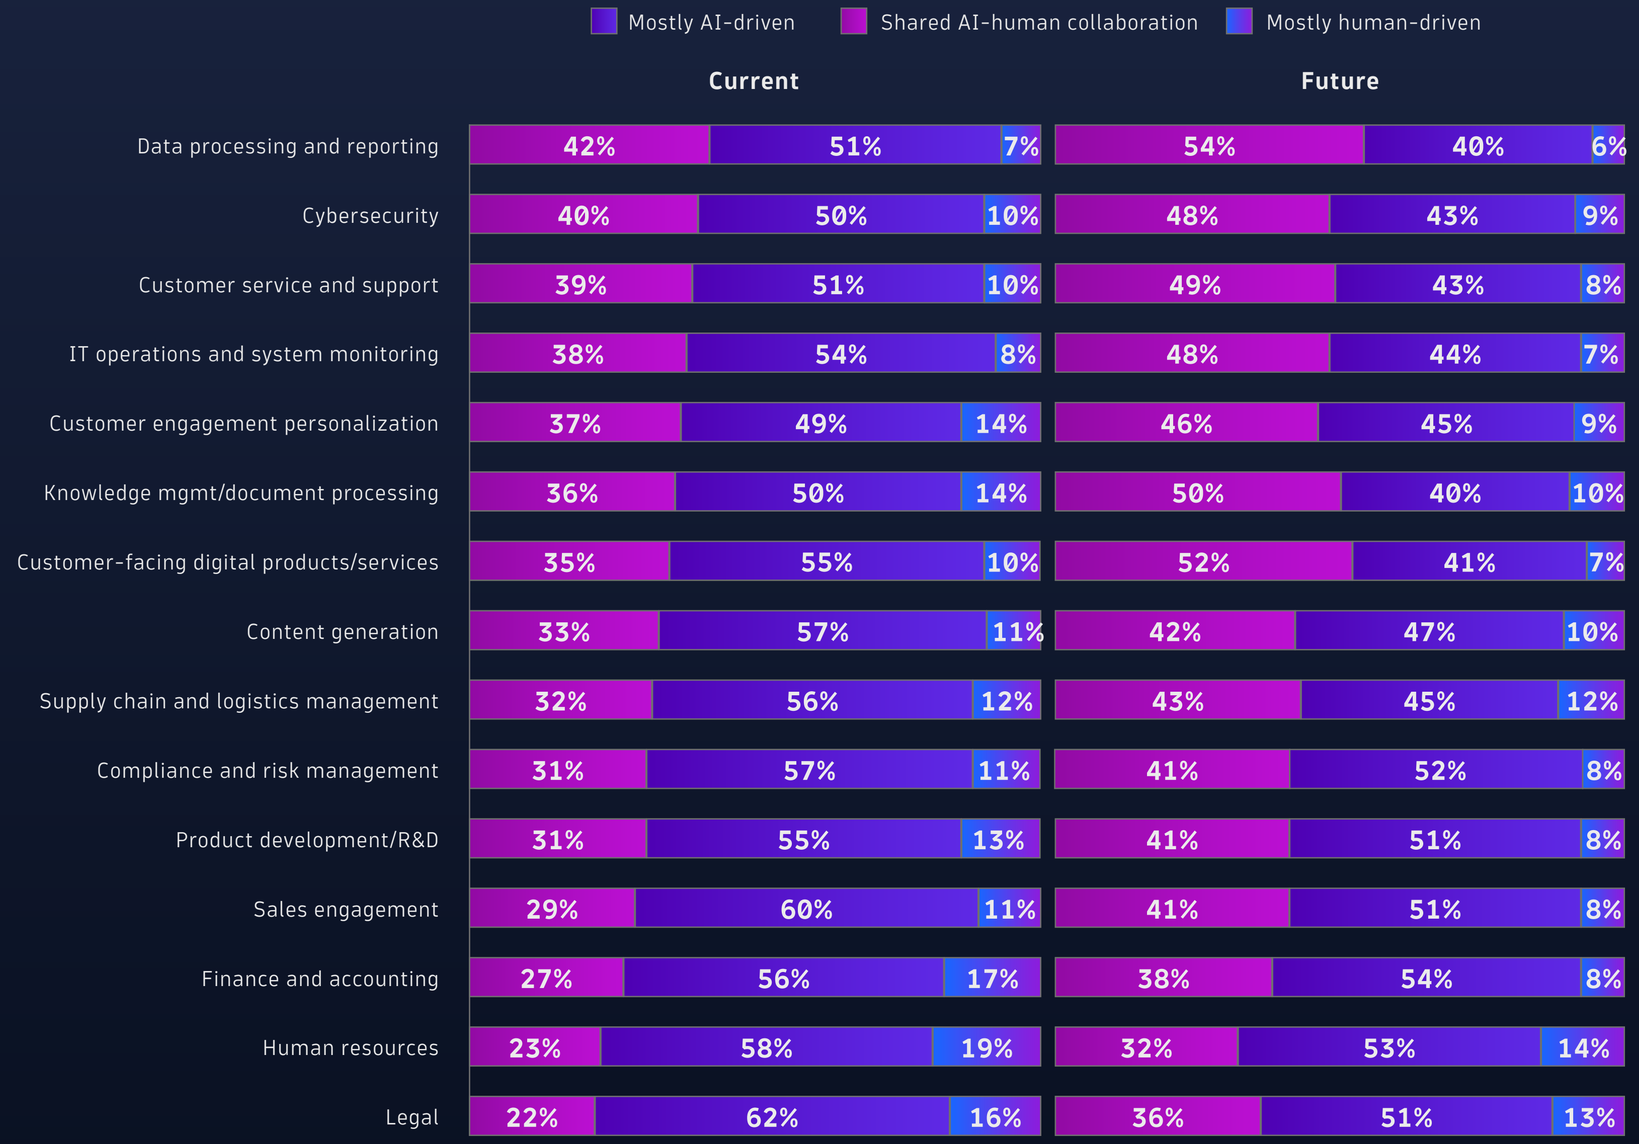

*Totals may not add to 100% due to rounding

dynatrace

The future is human-AI collaboration

Future automation goals indicate that organizations consider agentic AI as a crucial partner and enabler of human productivity, not a replacement.

Observability data is the crucial intelligence layer for human architects to govern and control the collaboration required to produce safe and reliable autonomous operations.

Redesign roles to optimize the human-AI partnership

As organizations adapt to the shifting priorities and demands of integrating agentic AI into workflows, leaders can be strategic about how AI facilitates human productivity.

- **Human judgment** guides the system by setting goals, defining boundaries, and ensuring accountability.
- **AI performs execution** for repeatable or time-sensitive tasks.
- An **observability-based control plane** helps manage the human/AI interface.

60%

use technical performance as their #1 agentic AI success measurement

52%

rely on logging, metrics, and traceability tools

44%

use manual methods to review communication flows among AI agents

dynatrace

CHAPTER FOUR

Reliability and resilience are key benchmarks of success

Organizations measure agentic AI success in terms of reliability and resilience. Real-time decision-making is a top priority, but monitoring methods still rely heavily on manual reviews, which limits the ability to scale.

KEY INSIGHTS

- Real-time decision-making, internal efficiencies, and system performance are the **top agentic AI goals**.
- **Technical performance** is the top success measurement as priorities and key performance indicators converge.
- **44%** are still reviewing agentic communication flows manually, revealing the need for accurate, real-time, and context-aware automation inputs.

Agentic systems are composed of a variety of specialized AI technologies (see appendix A). Because these systems are distributed and interconnected, a small flaw can quickly propagate through applications and environments, making reliability, resilience, and trust a driving concern. As a result, leaders are prioritizing real-time decision-making, internal efficiencies, and system performance above regulatory compliance, addressing labor shortages, and reducing mundane tasks.

Agentic AI priorities

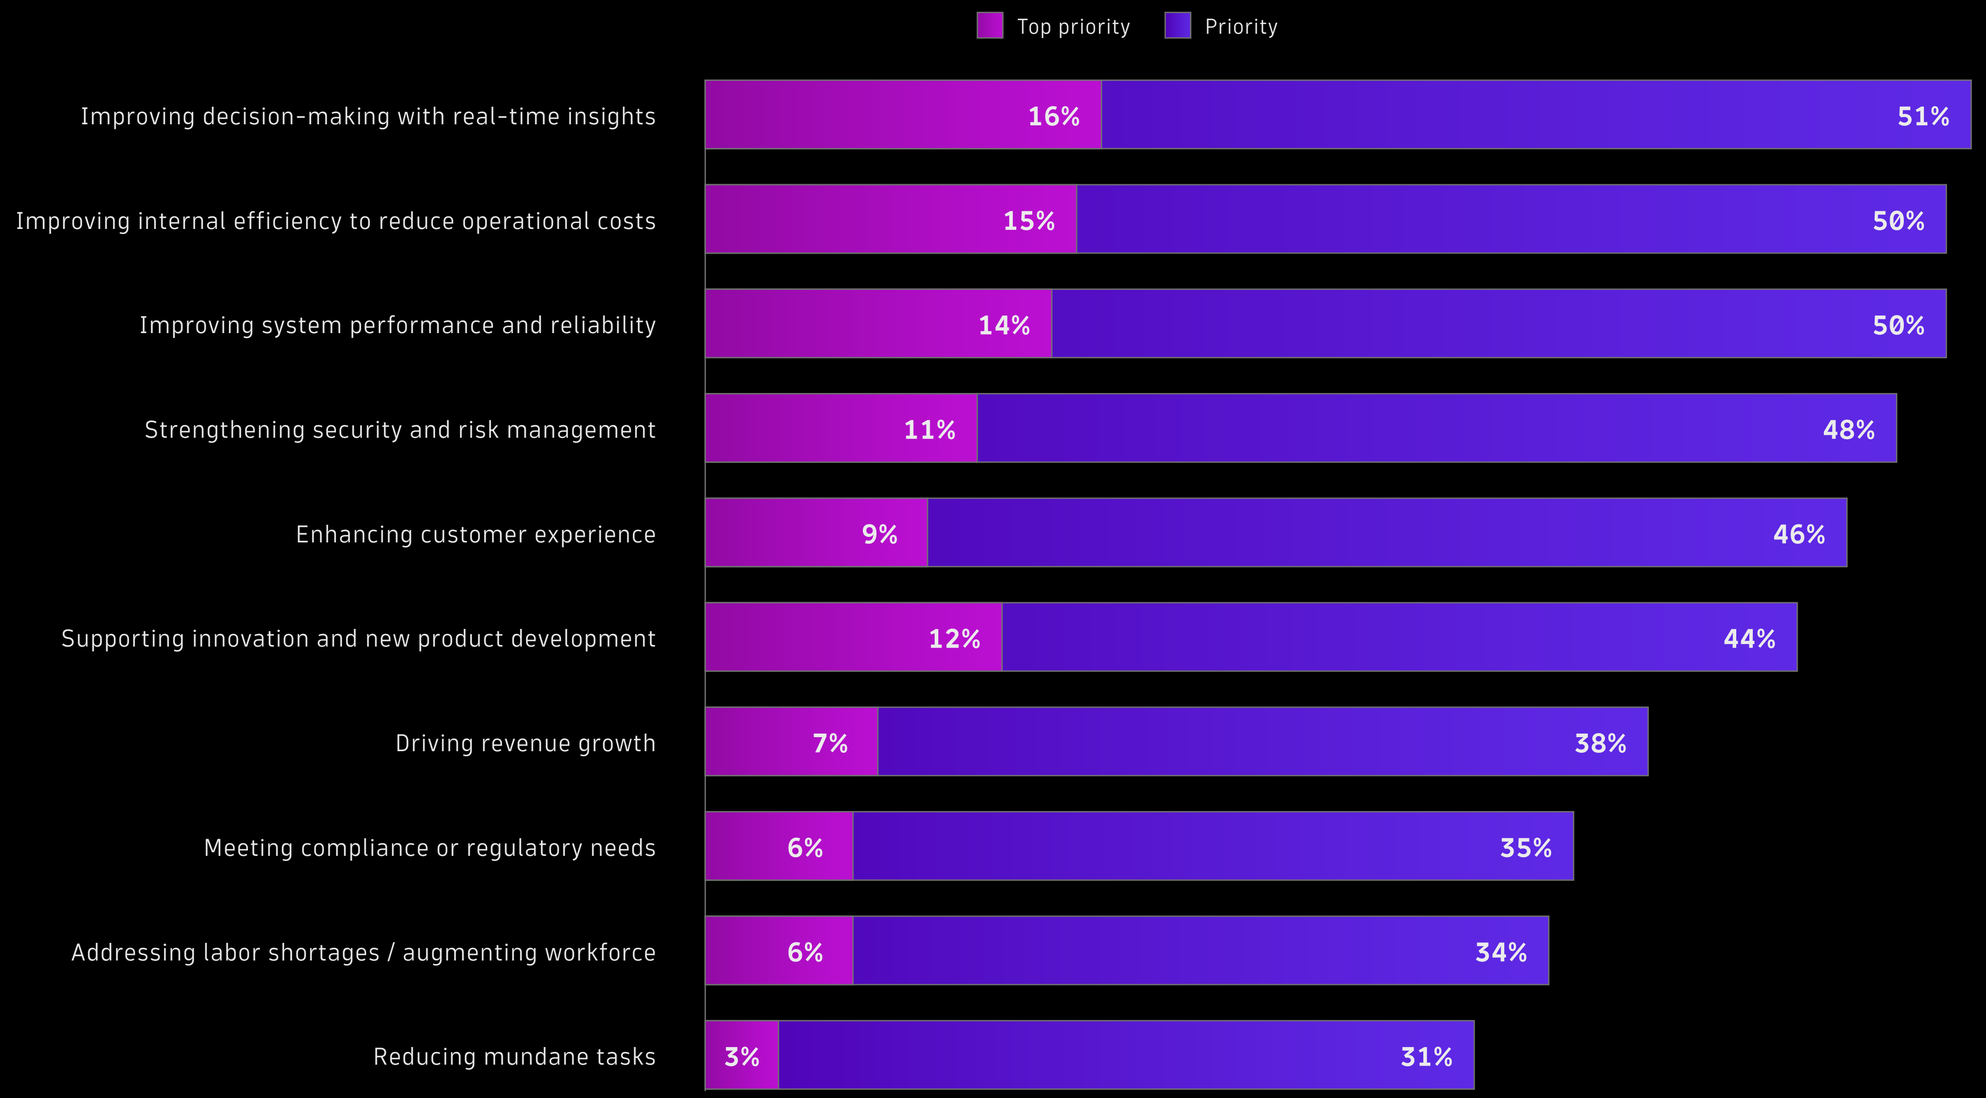

Priorities and key performance indicators are converging

In terms of how they measure agentic AI success and ROI, organizations are focusing less on daily use rates and sustainability than on technical performance, operational and developer efficiency, and customer satisfaction. Reducing risk is a fast-follow, suggesting that for systems to remain available and secure, reliability and security are becoming inseparable.

The larger an organization's size and revenue, the greater the likelihood it uses compliance and security benchmarks for measuring agentic AI success. Conversely, lower revenue organizations ($100M-$999M USD) focus more on increased productivity/efficiency and time-to-value.

How agentic AI success is measured

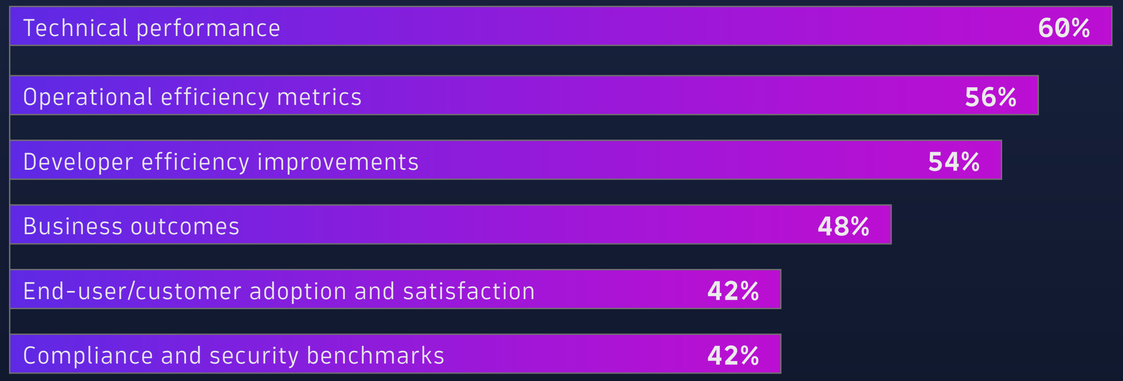

KPIs used to measure agentic AI ROI

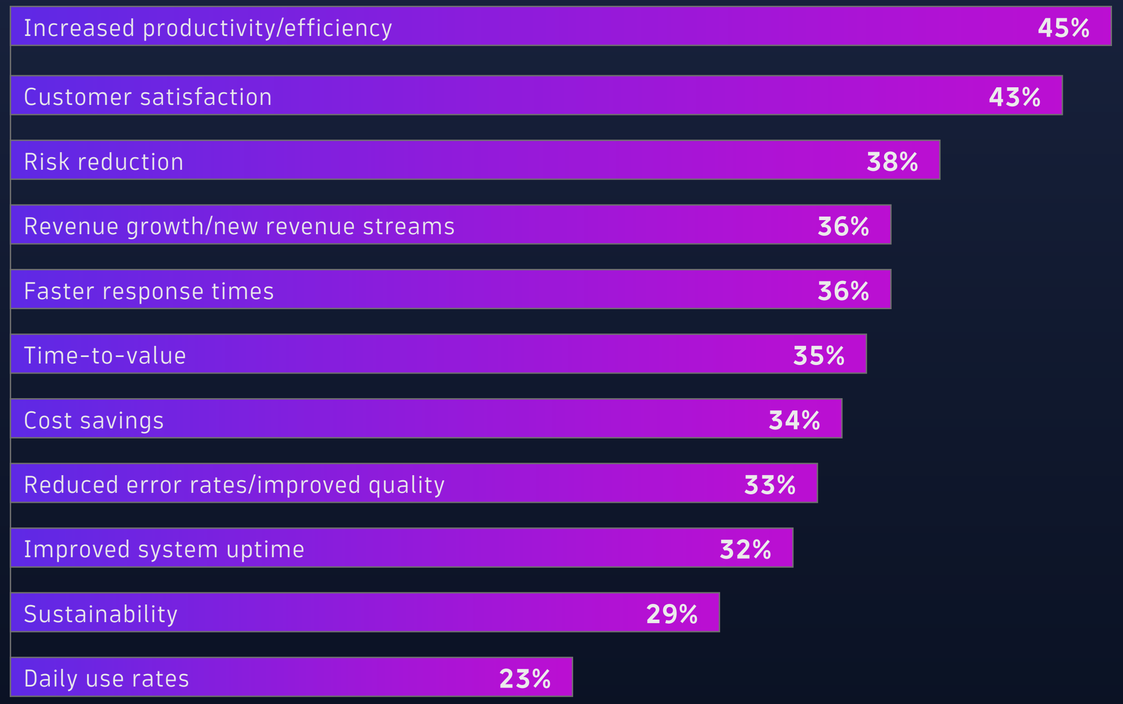

Monitoring methods focus on outcomes

Methods organizations use to monitor and analyze agentic AI systems are largely outcome-driven (about 60%), with a heavy reliance on logging, metrics, and traceability tools. While they're trying to automate anomaly and drift detection, 44% are still reviewing communication flows manually, indicating a need for automation inputs that are accurate, real-time, and context-aware.

Methods for monitoring/analyzing agentic AI systems

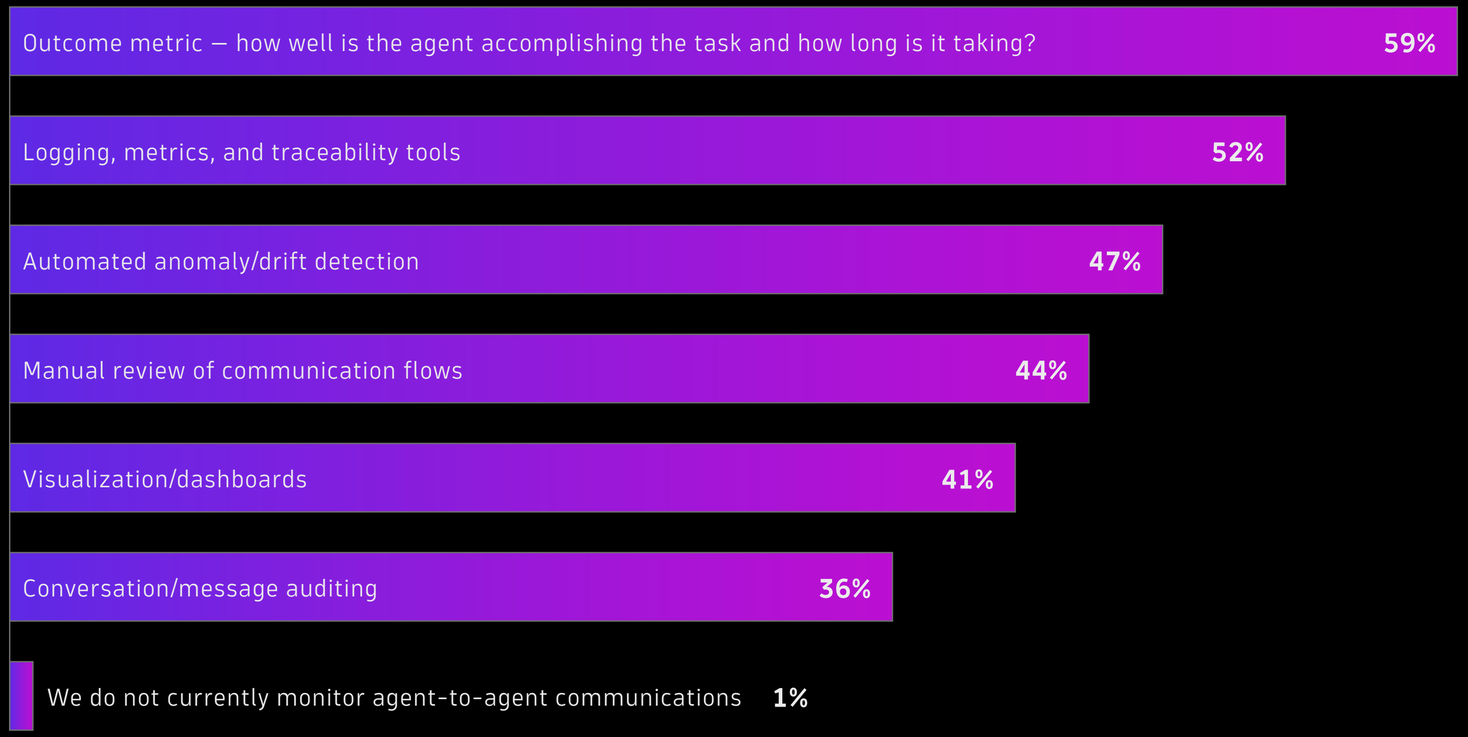

Measure success in resilience and reliability

Successful agentic systems must be able to absorb disruption, recover quickly, and maintain consistent performance under peak loads. Detecting problems early and rapid recovery are essential, because faults spread faster through interconnected agentic systems.

- **Measure resilience** by how systems respond in real time under stress.
- Treat **reliability and security as equal requirements** to keep systems available and secure.
- **Leverage factual observability** signals to uplift all KPIs and build trust for automation.

69%

use observability in the implementation stage of the agentic AI lifecycle

54%

use observability in development

57%

use observability in operationalization

dynatrace

CHAPTER FIVE

Observability is a control plane for autonomous operations

Organizations are already using observability across the AI delivery lifecycle, especially during implementation. But existing tools that rely solely on deterministic signals are not designed for probability-based, goal-driven systems.

KEY INSIGHTS

- Nearly **70%** use observability during agentic AI implementation and for monitoring training data quality.
- Existing monitoring and analysis tools leave **gaps in agentic AI behavior** that's hampering trust building and scale.
- To build trust in AI, observability needs to function like a **networking control plane** that orchestrates probabilistic AI with deterministic facts.

With conventional software, organizations can test quality during development. However, generative AI models introduce variability that can't be fully verified in pre-production. This variability requires real-time systems to monitor performance in production, identify unexpected behaviors, and implement corrective actions quickly.

Most respondents report using observability during at least one phase of the agentic AI life cycle. Significantly, more than two-thirds use observability during implementation for integrating with existing systems, detecting anomalies, and validating outputs. These indicators reflect the value of observability data at every stage, from identifying pre-deployment risks to unexpected AI outcomes in production.

Nearly all organizations use observability at some point during the agentic AI lifecycle

Development

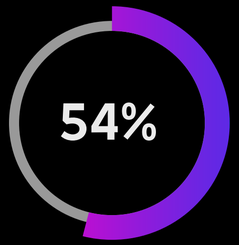

use observability in the development stage

- Monitoring training data quality **67%**
- Ensuring compliance and security **57%**
- Tracking prototype or POC performance **57%**
- Identifying risks before deployment **57%**
- Debugging and testing agent behavior **53%**

Implementation

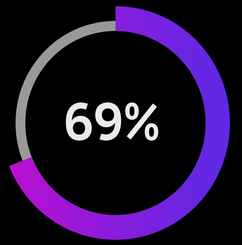

use observability in the implementation stage

- Monitoring integration with existing systems **59%**
- Detecting anomalies in real time **59%**
- Ensuring compliance and audit readiness **57%**
- Validating outputs during rollout or pilots **54%**
- Tracking adoption and user feedback **53%**

Operationalization

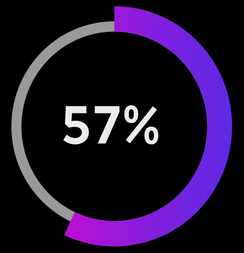

use observability in the operationalization stage

- Optimizing resource usage and efficiency **60%**
- Continuous monitoring of agent performance **59%**
- Improving transparency & explainability **57%**
- Detecting anomalies or unexpected outcomes **52%**
- Supporting governance, risk, and compliance **52%**
- Tracking A2A communication and MCP **40%**

Traditional monitoring and analysis tools leave gaps in understanding

However, respondents also report clear gaps in their current methods for understanding how agentic AI systems behave, the risks they expose, and drawing a throughline to business outcomes. This recognition suggests that current practices lack sufficiently structured inputs to ensure trustworthy outputs over time.

Respondents describe the following types of blind spots:

- Black-box behavior
- Lack of real-time insights
- Fragmented monitoring
- Weak linkage between technical signals and business outcomes

"There's a need for improved capabilities to assess the **ethical implications and unintended behaviors** of agentic AI systems."

—C-SUITE, USA

"The lack of a transparent, real-time view of the agent's decision-making logic makes its behavior resemble a **'black box'.**"

—HIGH-LEVEL EXECUTIVE, MEXICO

"Current tools lack **real-time insights** and proactive anomaly detection."

—MID-LEVEL EXECUTIVE, UK

"Existing tools lack sufficient **monitoring capabilities** for the long-term behavior of AI agents."

—MID-LEVEL EXECUTIVE, USA

"It's difficult to **directly link** technical monitoring metrics to business results."

—MID-LEVEL EXECUTIVE, JAPAN

Observability as a control-plane approach

Using observability data for monitoring is a good start, but to build trust in AI, observability needs to function like a control plane that can orchestrate deterministic data and context with generative models.

An intelligent observability framework needs to detect and take real-time actions to detect hallucinations, understand downstream implications, and take immediate action to limit impacts. A precise fact basis for reliable decision-making produces autonomous operations that can scale.

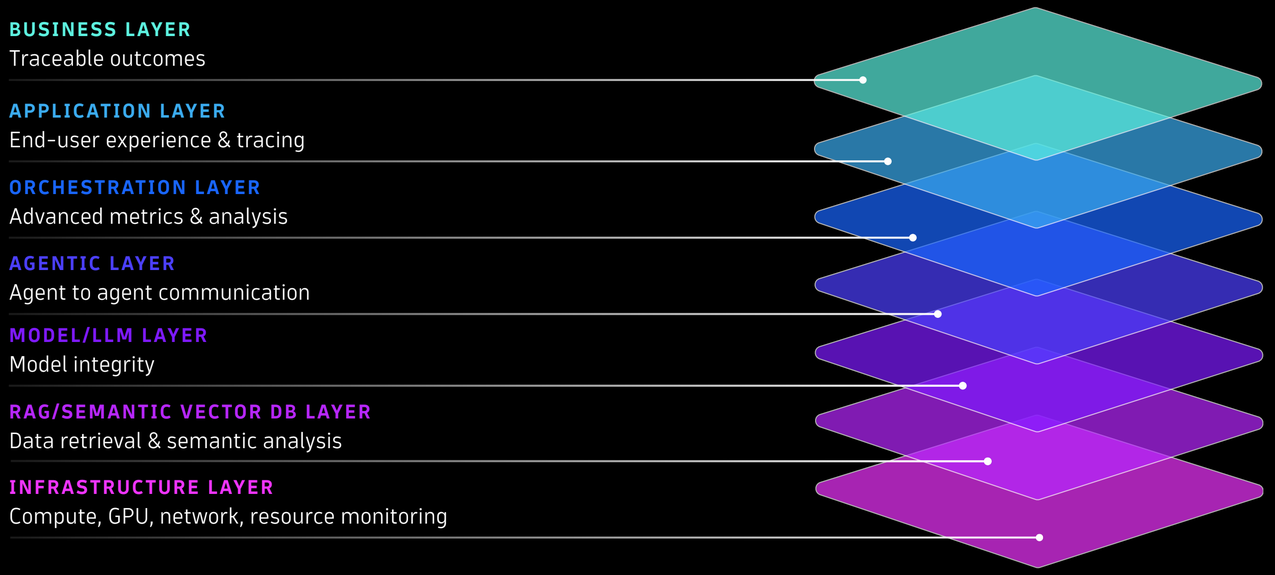

Establish observability as the control plane for autonomous operations

Because generative AI models are based on a probabilistic approach to decision making, building trust in automation requires that AI agent actions be grounded in deterministic, real-time facts.

- Ensure agents act on the same **factual signals** and remain accountable to the human operator.
- **Establish deterministic guardrails** that enable agentic systems to behave safely, act predictably, and collaborate using reliable real-world signals.
- **Instrument thoroughly** and standardize common semantic conventions, metrics, and trace attributes.

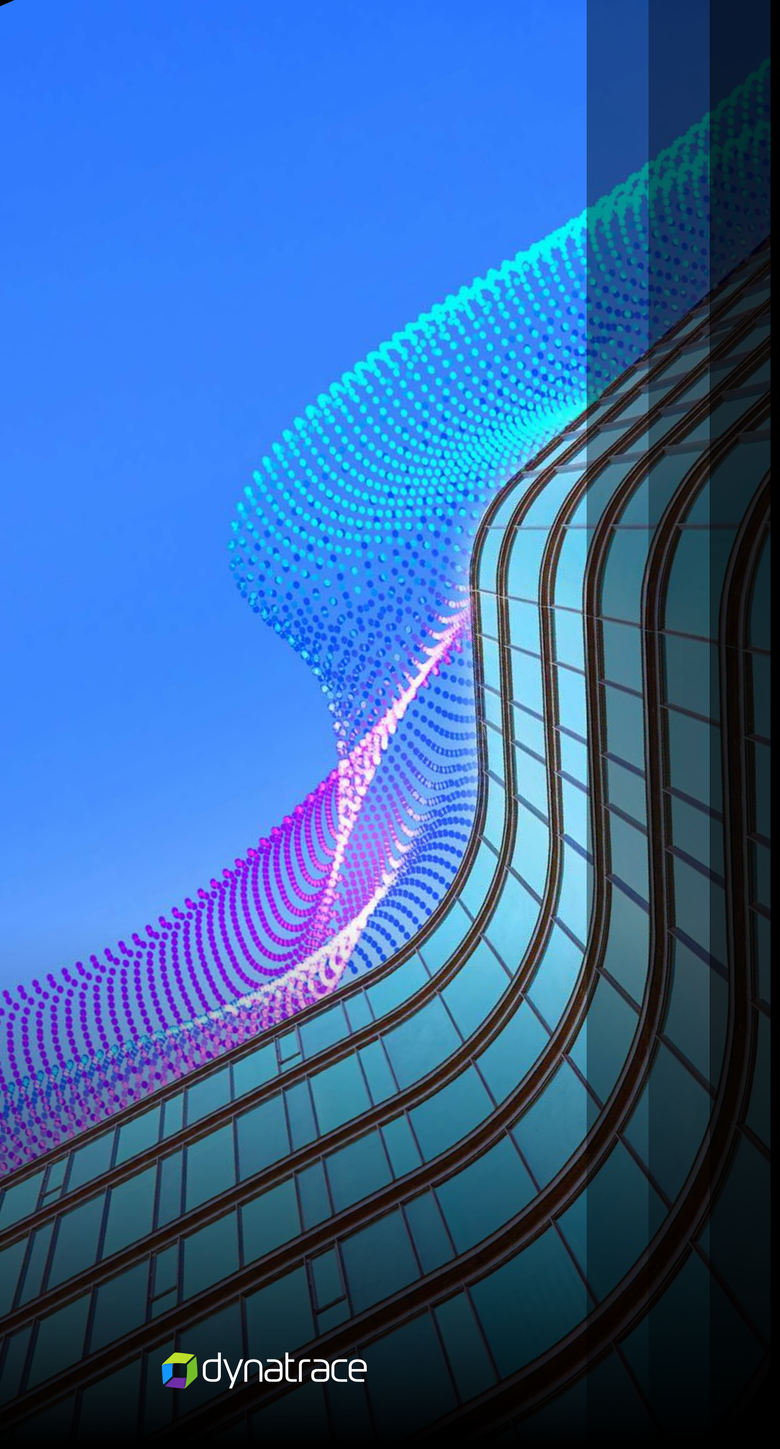

WHAT'S NEXT

The formula for operationalizing agentic AI

AI transformation is no longer coming—it's fully here, funded, and operationalizing fast. The challenge now is scaling agentic AI into reliable, resilient, and safe autonomous operations with informed human control and oversight.

- **The future is autonomous.** While agentic AI is still in its early stages, the trajectory is clear: Autonomous agents are quickly evolving to bolster human productivity, decision-making, efficiency, and scale. Although adoption is still early, projects are rapidly progressing to human-governed autonomous operations.
- **Limited visibility is jeopardizing reliability and control.** But limited real-time visibility into model behavior and regulating human-in-the-loop processes are hindering progress and leaving systems vulnerable to security, privacy, and compliance concerns.
- **Scaling agentic AI requires a new control plane.** To operationalize AI safely, organizations need observability that functions like a control plane, providing a deterministic, context-based foundation for the decisions of generative AI agents.

The next phase of autonomous operations belongs to organizations that can build trust, reliability, and resilience into their AI operating models with accurate, contextual observability.

dynatrace

Methodology overview

This report is based on a global survey of 919 senior leaders and decision makers directly involved in or responsible for agentic AI development and implementation in large enterprises with annual revenues of $100 million or more. It was conducted and analyzed by Qualtrics partner Y2 on behalf of Dynatrace during November and December 2025.

Industry verticals: Financial Services, Retail, Government, Utilities, Transportation & Fulfillment, Software & Digital Services

Respondent roles: Vice Presidents and directors of engineering leading AI projects; platform engineering leaders scaling AI systems; executive and IT decision-makers responsible for AI investments and risk; AI/ML and automation Leads; SRE and observability stakeholders

Respondent countries: United States (150), Canada (50), Brazil (35), Mexico (50), United Kingdom (100), France (50), Germany (100), Italy (50), Austria (30), Netherlands (27), UAE (46), Saudi Arabia (35), Australia (26), Japan (100), India (70).

All findings were analyzed to uncover success patterns, competitive advantages, and enterprise strategies for scaling agentic AI through intelligent observability across global markets.

The margin of error is ±3.2% at a 95% confidence level.

APPENDIX A

Additional findings

Currently used AI technologies

The agentic systems that fuel autonomous operations are made up of many AI technologies that perform specialized functions. The most used types are conversational and generative AI, such as chatbots and virtual assistants, as well as predictive and analytical AI. These findings support the value organizations expect from AI technologies and the priorities to which organizations are actively applying them.

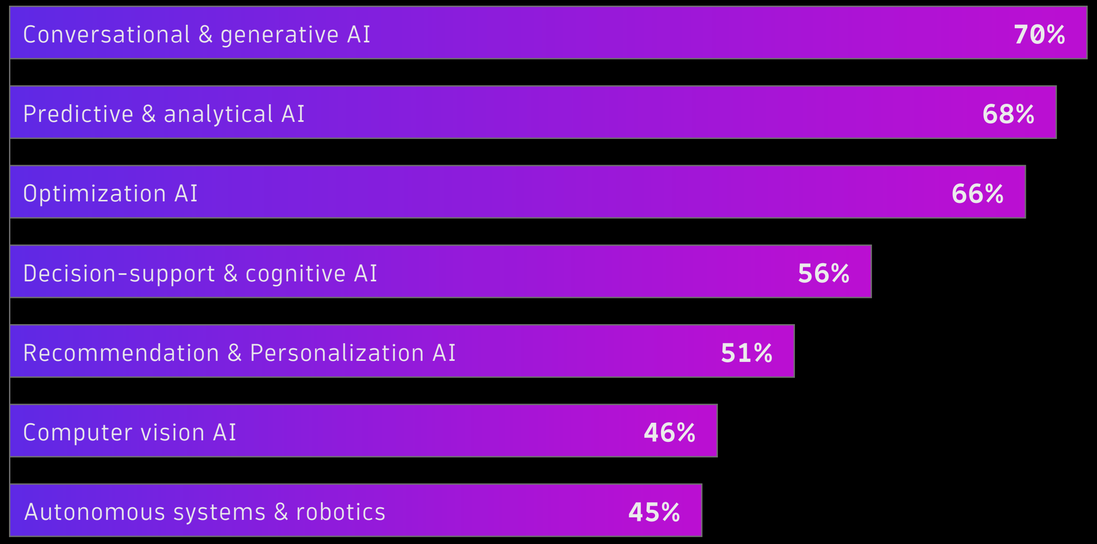

Agentic AI infrastructure

Hyperscaler platforms, such as Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure, host 82% of agentic AI infrastructure, followed by neocloud GPU-as-a-service platforms. 42% of respondents host their agentic AI infrastructure using on-premises assets.

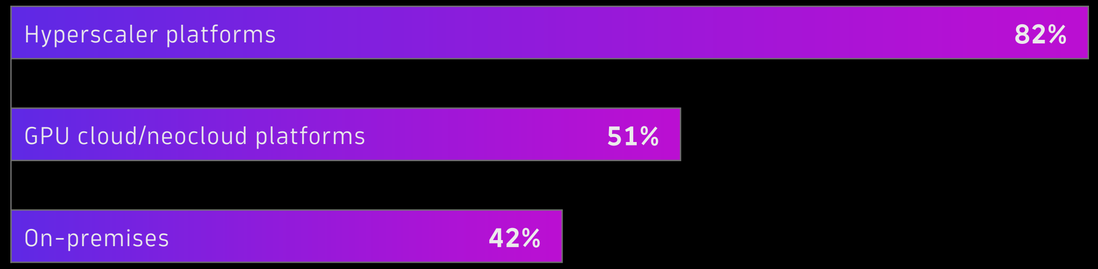

APPENDIX B

Global data summary

The Americas

Canada (50), US (150), Brazil (35), Mexico (50)

- Respondents in the Americas expect ROI for agentic AI to be greater in supply chain and logistics management than their counterparts in the Asia-Pacific region.
- **68%** of leaders in the Americas say that security and data privacy are their top criteria for moving a project from pilot to production than their counterparts in EMEA (**53%**).
- Top barriers to production in the Americas include technical challenges in managing and monitoring agents at scale (**45%**), and shortage of skilled staff (**51%**).
- Priority AI validation measures among respondents in the Americas are data quality checks (**51%**) and human review of agentic AI outputs (**48%**). They also place more emphasis on recording comprehensive logs and traces (**33%**) than their EMEA (**24%**) and APAC (**24%**) counterparts.
- Organizations based in the Americas place a higher priority for their agentic AI projects on driving revenue growth (**44%**) than their counterparts in APAC (**36%**) and EMEA (**36%**).
- When measuring agentic AI success, developer efficiency improvements are a higher priority in the Americas (**62%**) than respondents from EMEA (**52%**) and APAC (**49%**).
- The top observability uses among respondents in the Americas are ensuring compliance and security from the start during development (**70%**), ensuring compliance and audit readiness during implementation (**63%**), and continuous monitoring of agent performance during operations (**62%**).
- Respondents in the Americas are more likely to use recommendation and personalization AI (**60%**) than their counterparts in other regions.
- Organizations based in the Americas are more likely to host agentic AI workloads using on-premises infrastructure (**47%**) than their other regional counterparts.

Europe, Middle East, Africa (EMEA)

UK (100), France (50), Germany (100), Italy (50), Austria (30), Netherlands (27), UAE (46), Saudi Arabia (35)

- In EMEA, respondents expect that agentic AI will yield greater returns on investment in supply chain and logistics management than their counterparts in the Asia-Pacific region.
- **53%** of leaders in EMEA cite security and data privacy as their top criteria for moving agentic AI projects from pilot to production.
- **50%** of EMEA organizations report technical challenges in managing and monitoring agents at scale and **39%** cite shortage of skilled staff as barriers to production.
- EMEA respondents prioritize data quality checks and human review of agentic AI outputs as top AI validation measures equally (**47%**) and rely more on consistency checks and regression testing (**38%**) than their counterparts in the Americas (**34%**).
- When measuring agentic AI success, respondents in EMEA place greater emphasis on end-user/customer adoption and satisfaction (**45%**) than their counterparts in APAC (**43%**) and the Americas (**39%**).
- The top observability uses in the EMEA region include monitoring training and data quality in development (**66%**), detecting real-time anomalies during implementation (**60%**), and optimizing resource usage and efficiency in operations (**57%**).
- Organizations in EMEA host **39%** of their agentic AI infrastructure on-premises compared to APAC (**41%**) and the Americas (**47%**).

Asia-Pacific (APAC)

Australia (30), Japan (100), India (70)

- Cybersecurity and data processing/reporting carry greater expectations of return on investment among APAC respondents.
- **60%** of Asia-Pacific-based organizations report technical challenges in managing and monitoring agents at scale as a top barrier to production, and **46%** cite a shortage of skilled staff.
- For validating agentic AI, respondents in Asia-Pacific prioritize data quality checks at a higher percentage (**56%**) than respondents in the Americas (**51%**) and EMEA (**47%**), as well as security validations (**45%**) versus **34%** for EMEA and **41%** for the Americas.
- When validating agentic AI, APAC respondents place ethics, fairness, or bias validation at a higher priority (**32%**) than respondents in the Americas (**28%**) and EMEA (**27%**).
- For organizations in APAC, improving internal efficiency is a greater priority for their agentic AI projects (**56%**) than for their counterparts in the Americas (**49%**) and EMEA (**48%**).
- When measuring agentic AI success, respondents in APAC place greater emphasis on end-user/customer adoption and satisfaction (**43%**) than their counterparts in the Americas (**39%**).
- The top observability uses in the Asia-Pacific region include monitoring and training data quality during development (**74%**), monitoring integration with existing systems in implementation (**64%**), and continuous monitoring of agent performance in operations (**65%**).
- Respondents in the Asia-Pacific region are more likely to use conversational and generative AI (**77%**) than their counterparts in other regions.
- Organizations in APAC host **41%** of their agentic AI infrastructure on-premises.

Dynatrace: The observability control plane for advancing agentic AI from pilot to production

We hope this report has provided valuable insights about how agentic AI projects are advancing to autonomous operations and how observability provides vital intelligence at every stage.

Dynatrace is committed to providing enterprises the data and intelligence they need to be successful with their enterprise cloud and digital transformation initiatives, no matter how complex.

Learn more

To experience the Dynatrace difference, visit www.dynatrace.com/platform for assets, resources, and a free 15-day trial.

ABOUT 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

dynatrace

01.23.26 bae22697-ebk-cs