


2026 REPORT

OBSERVABILITY FORECAST



INTRODUCTION

For the second year running, AI is the top force driving demand for observability

Organizations are racing to monitor their AI systems, and engineering teams have found no bigger lever for running software faster without breaking it. Every other priority in this report, from tool consolidation to outage cost to agentic observability, now sits downstream of these facts.


The story pulls in two directions. AI helps teams ship more code, resolve incidents faster, and reclaim engineering hours that used to disappear into firefighting. But it also generates code faster than humans can review it, puts autonomous agents into production, and adds layers to the stack that are harder to see into.

This dynamic gives rise to *phantom velocity*, the illusion of rapid development where upfront productivity gains are quietly canceled out by downstream downtime, firefighting, and emergency fixes. The organizations that win the next phase are the ones that can capture true speed while keeping the system observable enough to eliminate hidden threats to runtime reliability.

This year's forecast surveyed 2,575 engineering and IT leaders and practitioners, up from 1,700 last year. Their answers point to the same conclusion. AI made software faster to build and harder to run. Observability is the layer that closes the gap.

METHODOLOGY

New Relic partnered with Enterprise Technology Research (ETR) to survey 2,575 IT and engineering leaders and practitioners across a wide range of industries, organization sizes, and regions. Respondents span executive, management, and practitioner roles. All figures reflect all respondents unless a segment is named.




KEY FINDINGS

AI is the number one observability driver


AI applications, agentic systems, and security risks define observability demand—expanding the focus of monitoring from uptime to AI reliability and governance.

TOP THREE DRIVERS OF OBSERVABILITY



AI observability is now standard practice


AI observability has quickly moved from experiment to baseline requirement, with adoption projected to hit 80% within 12 months.



Outage costs fall, but risks remain high

Faster detection and resolution times are bending the financial damage curve down per incident, even as overall outage frequency remains high.

YEARLY HIGH-IMPACT OUTAGE COSTS



One in four AI agents run completely unmonitored

Deploying autonomous software that alters code, configs, or infrastructure without live telemetry is the report’s most urgent, unmitigated operational risk.



AI-generated code demands observability


An overwhelming (82.6%) agree that the rise of AI-generated code makes reliable observability more critical, with 35.2% strongly agreeing, and virtually no opposition (5.2%).



Tool sprawl creeps back

New AI workloads have raised demand for AI-specific observability tools, reversing years of hard-won platform consolidation.

NUMBER OF OBSERVABILITY TOOLS



**AI is the
new center
of gravity**




AI pulls observability forward

AI strategy has replaced traditional priorities like cost management as the main driver of observability investments, forcing organizations to adapt.

Scaling AI deployment exposes systems to new attack vectors and unpredictable metered workloads. Observability gives organizations leverage over AI security risks and costs.


DRIVERS OF
OBSERVABILITY DEMAND



Beyond AI monitoring

Our survey has evolved to cover five categories of AI observability, each demanding dedicated deployment strategies as AI becomes normalized.


FIVE TYPES OF AI OBSERVABILITY



The agentic-AI visibility gap

Although autonomous agents are widely used in production, a quarter of organizations face major operational risks deploying this technology without monitoring or runtime visibility.

STATE OF AGENTIC MONITORING




You cannot run a workforce you cannot see

While 93% of organizations are adopting AI agents, 24.6% run them unmonitored. An agent altering code, configs, or infrastructure without observability is an incident waiting to be discovered by a customer.

AI incident response reaches production


Organizations have actively deployed the AI incident tools they wished for a year ago, turning 2025 aspirations into 2026 production baselines.

AI troubleshooting leads active adoption at 44.9%, while conversational assistants surged from 24% to 35%. Only 6.1% of engineering teams operate without AI incident tools.



AI pays off most in the investigation phase

AI delivers its highest operational value during investigation (35.8%), replacing tedious manual detective work with rapid, directed root-cause analysis.



PLATFORM TELEMETRY



New Relic platform usage data confirms AI compresses the discovery phase by surfacing failure patterns earlier, shifting engineering teams from detective work to guided remediation.

[Read the AI Impact Report →](#)


The AI-generated code reckoning



Most code is now written by AI, changing the role of observability

As our recent [State of AI Coding](#) found, two-thirds of organizations now generate over half their code with AI. While 94% of leaders rate AI code higher quality during review, shipped code tells a different story: 78% report more production incidents, and 86% see senior engineers spending more time fixing code.

AI coding tools read the source, but remain blind to the trace. When readability no longer guarantees reliability, runtime evidence becomes the only trusted signal. Consequently, 82.6% of leaders now agree that robust observability is critical.



LEADERS SAY OBSERVABILITY
IS CRITICAL FOR AI CODE

82.6%

Agree or strongly agree that
AI-generated code makes robust
observability more critical


DEVELOPER BENCHMARK

The Production Paradox: Research from New Relic's State of AI Coding report reveals why production breaks down: 62% of teams routinely ship AI-generated code without line-by-line manual verification. Unvetted code passes initial static checks, only to hit unexpected runtime edge cases under live traffic.

Read the [State of AI Coding Report](#) →

More coding tools create more need for observability

To understand the scale of AI-generated code, look at how saturated the toolchain is. We asked which AI-coding or code-generation tools organizations currently use, and the list is long.




GitHub Copilot leads at nearly half of all organizations, with Gemini Code Assist and Codex close behind. Only 2.6% of organizations use no AI coding tools at all. Most organizations run several of these at once, which is the detail that matters most for observability.

Every one of these tools generates code in its own way, prompted by different engineers using different conventions. The output lands in the same production environment, under the same service-level agreements, on the same outage clock as everything else. A codebase assembled from Copilot, Cursor, Codex, and Claude Code is a codebase with no single author and no single style. Reading the source is no longer a reliable way to understand the system, because there is too much of it, generated too fast, by too many different prompts. Runtime evidence becomes the more dependable signal, and observability is what captures it.

AI is reshaping stack complexity in both directions

AI does not simplify operations across the board. We asked how AI has affected the complexity of managing the software stack, and the answers split.

That split is the operations-side view of the trade-off the coding report named. Speed got cheaper. Reliability did not. AI cuts toil for most teams while adding opacity for a meaningful minority, and which side an organization lands on depends largely on whether it can see what its AI is doing.



OpenTelemetry makes AI-built code observable




OpenTelemetry standardizes a fragmented AI toolchain

When code is generated across Copilot, Cursor, Codex, Claude Code, and custom prompts, the resulting telemetry often fragments into incompatible formats. Without a shared standard, diagnosing incidents across disparate AI tools creates severe operational friction. OpenTelemetry provides the universal framework required to correlate metrics, events, logs, and traces across multi-engine codebases.

OTel adoption is accelerating

Momentum around OpenTelemetry is clear across production environments. Beyond internal standardization, open-source compatibility has become a decisive buying criterion, as engineering leaders demand vendor portability across their observability stack.



TELEMETRY MOVES

UPSTREAM INTO THE PROMPT

The New Relic [State of AI Coding Report](#) reveals 78% of teams now prompt AI tools to write specific telemetry directly into generated code. Standardizing on OpenTelemetry semantic conventions and structured logging ensures AI output remains readable before reaching production.

[Read the report](#) →


The business case for AI-strengthened observability



More engineer time is spent firefighting—instead of innovation


Over a third of engineering capacity is now consumed by system disruptions, pulling engineers away from shipping product. AI-strengthened observability exists to reverse this upward trend, converting wasted troubleshooting hours into an **innovation dividend**.

Despite the rapid adoption of AI coding tools, engineering teams are spending more time firefighting, not less. High AI code velocity without integrated runtime visibility is expanding the reactive toil tax.



Where reclaimed time would go

Reclaimed engineering hours do not vanish into administrative meetings. When troubleshooting hours are recovered, leaders reinvest them directly back into platform resilience and product velocity.



SYSTEM ARCHITECTURE & SCALABILITY SECURITY & COMPLIANCE NEW FEATURES TECHNICAL DEBT TRAINING

EXECUTIVE TAKE: THE INNOVATION DIVIDEND


New Relic platform usage data from the [AI Impact Report](#) confirms AI-assisted observability compresses the discovery phase of troubleshooting. By surfacing context and failure patterns earlier, teams shift from manual detective work to directed resolution, reclaiming the hours needed to fund architectural improvements.

[Read the report](#) →

Monitored AI agents drive measurable observability returns

Observability continues to yield strong financial returns, with three-quarters of organizations reporting positive ROI and planning to increase their budget over the next 12 months.


Holding steady from last year’s 75% benchmark, this ongoing performance shows observability is a rare, highly defensible investment that consistently pays for itself.



The agent visibility dividend

Financial returns diverge sharply based on agent monitoring. Organizations actively watching AI agents in production achieve significantly higher returns on their observability investment while mitigating critical operational risks.

OBSERVABILITY ROI BY AGENTIC MONITORING MATURITY



Budget trend



Technology leaders are expanding budgets to match the proven ROI of AI-strengthened observability.

Connecting technical visibility to AI and business outcomes

Observability has evolved from a backend debugging utility into a central driver of business performance. While quantifying incident costs remains a baseline need (51%), runtime visibility for AI execution (45%) has rapidly surged to near-parity with traditional priorities.

Because nearly half of organizations already rely on observability to safely roll out and scale AI applications, AI strategy and observability strategy have officially merged.

OBSERVABILITY INTO BUSINESS OUTCOMES

Quantify the business impact of incidents	51.1%
Improve customer experience	48.8%
Mitigate service disruptions and business risk	48.5%
Shift from reactive to proactive work	45.7%
AI application development and deployment	45.3%


FASTEST  RISING

Value for teams and their leaders


The benefits of AI-strengthened observability span multiple tiers, making it easier for both engineers and their leaders to achieve their goals.



PRACTITIONERS



EXECUTIVES & MANAGERS



The cost of
outages is falling,
but the risk is not



Annualized outage losses

Lower downtime costs stem directly from faster response times rather than fewer disruptions. Shrinking detection and resolution windows reduces the financial damage of every event: outage costs fell not because incidents became rarer, but because teams catch and clear them faster.

Per hour, high-impact business outages cost our respondents an average of \$1.85 million. That's over \$30,000 a minute in financial exposure.


RESPONSE TIME IMPROVEMENTS




70.3%
REPORT FASTER
MEAN TIME TO
DETECT (MTTD)



69.3%
REPORT FASTER
MEAN TIME TO
RESOLVE (MTTR)




14%+
REPORT 50%+
IMPROVEMENTS ON
BOTH METRICS




Outages are still frequent

Faster recovery does not mean the underlying threat is receding. High-impact outages remain a regular operational event, with more than a quarter of organizations hitting major disruptions weekly or more often—including a 200% increase in multiple daily outages.



Detection and resolution still take time

For high-impact outages, mean time to detect averages 41.3 minutes and mean time to resolve averages 53.6 minutes. The medians are close behind, at 35 and 51 minutes. Every minute in those windows is revenue, customer trust, and engineering attention spent on recovery instead of building. That is the exposure the \$73.5 million figure measures, and it is the exposure observability is built to compress.



Why observability matters more in 2026

The past year made the stakes plain. A failure in a single AWS region on October 20, 2025 took 15 hours to fully resolve, generated more than 17 million outage reports across 60-plus countries, and took down Snapchat, Reddit, Coinbase, Venmo, and airline and banking systems ([CNBC](#), [Ookla](#)). The same event disrupted observability and incident-response vendors themselves, delaying the alerts teams rely on to respond ([IncidentHub](#)). Nine days later, a configuration change inside Microsoft's Azure Front Door caused a nine-hour global outage that hit Microsoft 365, Xbox, and the Azure portal, with loss estimates ranging from \$4.8 billion to \$16 billion ([Network World](#), [ThousandEyes](#)). Three weeks after that, a Cloudflare failure took down X, ChatGPT, Spotify, and Discord ([NYT](#)).


The pattern is consistent. The biggest outages of the past year started inside the shared infrastructure that nearly every business depends on, and they cascaded fast. When the failure is upstream, the only lever a team controls is how fast it sees the problem and routes around it. The organizations that cut their detection and resolution times are the ones bending the cost curve down while everyone else watches it climb.

The operational reality




What still causes outages

While upstream cloud failures are leading causes, deployment and environment changes represent 60% combined—a clear warning for the AI era. As more code ships from more tools at higher frequency, these controllable failure modes become the most important to watch.



Nearly half of teams still detect incidents reactively

Despite growing observability investments, many organizations still learn of outages through manual checks, customer reports, or internal complaints. Outdated practices leave a sharp dividing line between resilient organizations and reactive ones.



The consolidation imperative



Tool sprawl is creeping back


For two years the story was consolidation. Organizations had cut the average number of observability tools from 6 in 2024 to 4.4 in 2025. This year the trend reversed. The average organization now runs ~5 observability tools.

The likely cause is the AI boom itself. As organizations rush to monitor AI applications, agents, and models, they reach for new AI-specific point tools that fragment the stack again. The same force driving observability demand is driving observability sprawl. The cure and the disease share a root.




Consolidation still wins, but the margin is narrow

Organizations still prefer a unified approach, but less decisively than the tool-count reversal might suggest. Half (49.6%) prefer a single consolidated platform, while 40.2% prefer multiple point solutions and 9.7% have no preference. The pull toward point solutions is real, and again it is likely the AI tooling wave that explains it. The organizations that consolidate fastest will be the ones that resist adding a new dashboard for every new AI capability.



What buyers prioritize

When choosing an observability vendor, cost has emerged as the single most decisive factor as engineering teams face increasing budget scrutiny. Simultaneously, AI and machine learning capabilities have surged into the top three buying criteria, signaling that buyers expect intelligent automation to deliver both efficiency and cost control.



Strategic imperatives for 2026 and beyond

Instrument your AI before it ships, including your agents

Almost a quarter of organizations (24.6%) are running AI agents in production without monitoring them, while 68% already monitor them or plan to within 12 months. The dividing line in 2026 is whether AI systems are observable from the moment they go live. Waiting to add observability after an agent causes an incident is waiting too long.

Treat production telemetry as a first-class input to AI development

The most consequential observability practice of the next two years is closing the loop between runtime data and code generation. The forecast shows 82.6% of leaders already believe AI-generated code makes observability more critical. The next step is feeding what production reveals back into the next prompt, the next review, and the next deployment.

Standardize on OpenTelemetry before the toolchain fragments your telemetry

With code now coming from Copilot, Cursor, Codex, Claude Code, and more, OTel is what keeps the telemetry those tools emit consistent enough to correlate. Nearly three-quarters of organizations are already standardizing, migrating, or testing. The ones that prompt their AI tools for OTel conventions today will be able to read their AI-generated systems tomorrow.

Consolidate before AI tooling re-fragments the stack

The average tool count rose from ~4 to ~5 in a single year. The consolidation gains of the past few years are reversing under AI point-tool sprawl. A unified platform that absorbs new AI capabilities without adding new tools is the way to capture AI's benefits without paying its complexity tax.

Adopt AIOps to turn reactive work into innovation dividend

An average of 37% of engineering time still goes to disruptions. AIOps is deployed by only 35.6% of organizations today, with 34.8% planning to adopt it within a year. The teams that move first will be the ones that convert reclaimed hours into architecture, security, and new features.

Eliminate phantom velocity by closing the loop between runtime and development

AI allows teams to generate code faster than ever, but unchecked speed creates a hidden tax on runtime reliability. Phantom velocity occurs when rapid development gains are quietly erased by downtime, emergency fixes, and firefighting. For measurable results, teams must close the feedback loop between production telemetry and development—ensuring speed leads to durable innovation rather than costly outages.

ABOUT THIS REPORT

All data in this report derives from a survey conducted by New Relic in partnership with Enterprise Technology Research (ETR), reflecting responses from 2,575 IT and engineering leaders and practitioners. Supplementary data referenced from the 2026 New Relic AI Impact Report and the 2026 State of AI Coding report. The \$76 million annualized outage cost for 2025 reflects the prior year's Observability Forecast. Industry-wide outage cost figures are drawn from third-party sources cited inline.


ABOUT ETR

ETR is a technology market research firm that leverages proprietary data from its targeted ITDM community to deliver actionable insights about spending intentions and industry trends. Since 2010, ETR has worked diligently at achieving one goal: eliminating the need for opinions in enterprise research, which are typically formed from incomplete, biased, and statistically insignificant data. The ETR community of ITDMs is uniquely positioned to provide best-in-class customer/evaluator perspectives. Its proprietary data and insights from this community empower institutional investors, technology companies, and ITDMs to navigate the complex enterprise technology landscape amid an expanding marketplace.

ABOUT NEW RELIC

New Relic arms businesses with the trust and confidence required to thrive in the AI era. The New Relic Intelligent Observability Platform is the leading AI-strengthened platform designed to unify telemetry and business outcomes, bringing intelligence and automated actions to the most complex digital environments. The platform shifts teams from reactive firefighting to intelligent orchestration, leveraging AI-driven automation to optimize technology spend and protect revenue in real-time. That's why global leaders—Adidas Runtastic, Domino's, Ryanair, Swiggy, Topgolf, and William Hill—run on New Relic to drive innovation and deliver exceptional customer experiences.

© Copyright 2026, New Relic, Inc. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies. 09.2026



[Learn about the New Relic Platform](#)