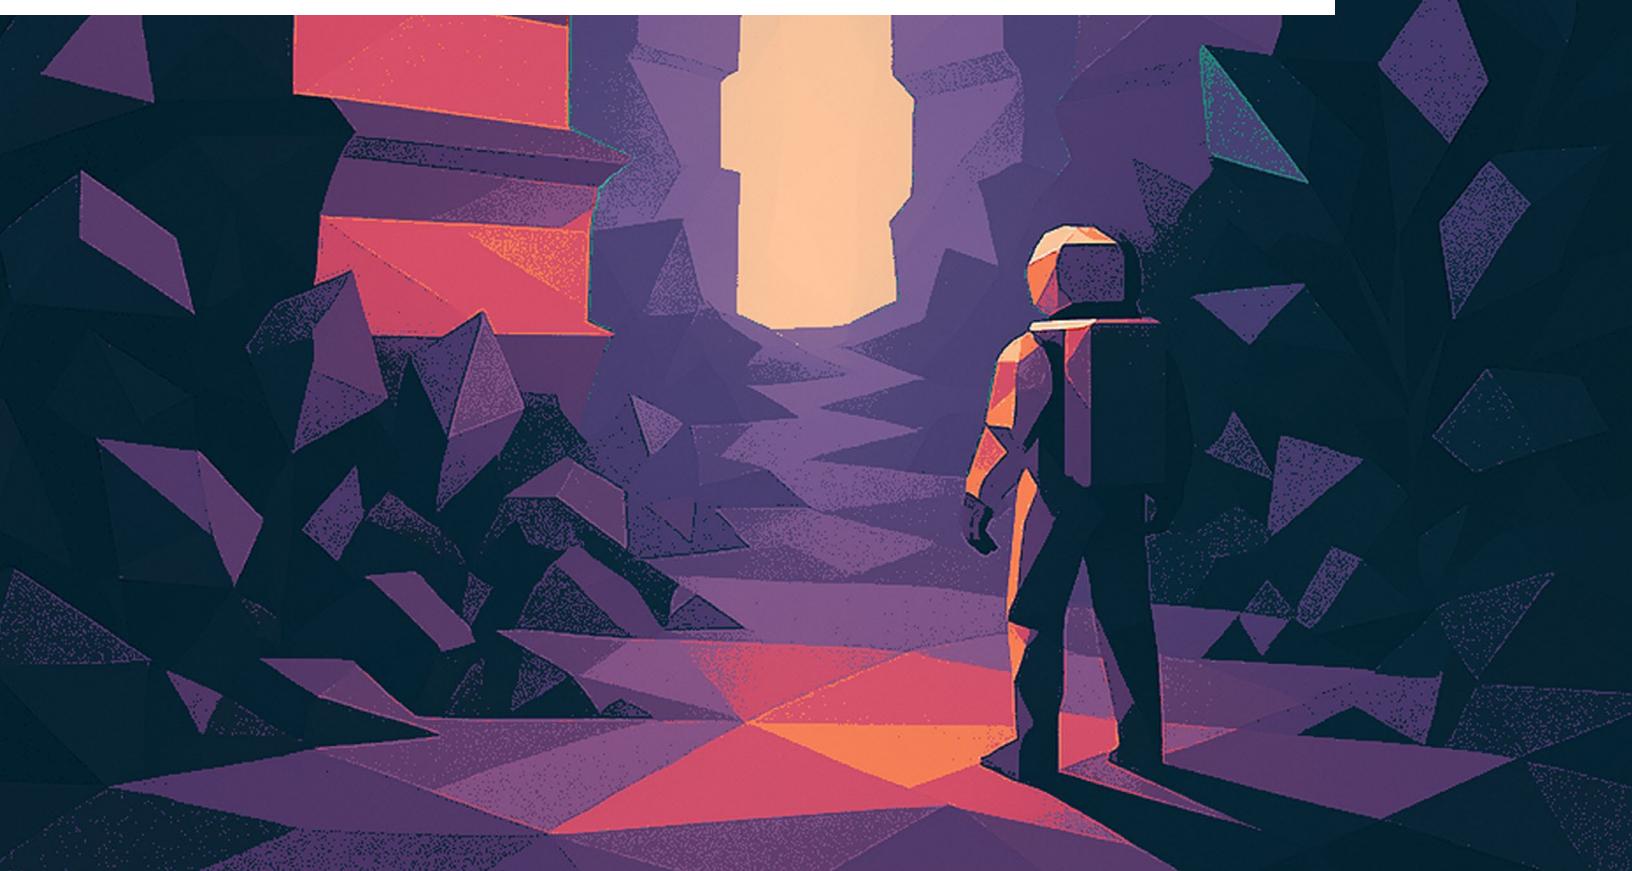


2026 AI IMPACT REPORT

NEW RELIC PLATFORM INSIGHTS



INTRODUCTION

Modern software systems operate continuously across distributed architectures, with changes deployed regularly.

Incidents are frequent. Potential issues surface at every layer of the stack and it's often difficult to identify the signal from the noise. Information about system performance grows faster than teams can keep up with. In fact, 33% of engineering time is spent firefighting or addressing these disruptions.

In 2025, New Relic customers created 2.2 billion alert events across monitored environments. 821 million issues, or potential problems, originated in production systems alone. This activity reflects the steady operational load of software at scale.

Under these conditions, operational effectiveness depends on speed of understanding. Teams must interpret signals, establish context, and decide on action in compressed timeframes. Alert volume rises quickly. Manual correlation of alert events to potential issues becomes increasingly difficult. Teams need to decide which of those alert events and potential issues are worth their time and which are just noise.

AI-strengthened observability addresses these constraints by organizing signals, surfacing context, and supporting faster decision making for live systems.

This report examines how engineering teams operate when AI is used to monitor daily workflows. These insights are drawn from de-identified and aggregated usage data surrounding 6.6 million active users on New Relic over the course of 2025.

The following analysis highlights consistent improvements in signal clarity, resolution speed, and deployment behavior among teams leveraging AI.

HOW TO READ THIS DATA

The outcomes described in this report reflect real-world operating environments. Operational performance is influenced by many variables: architecture, team maturity, incident processes, deployment practices, and tooling. The results in this report reflect that reality.

That said, the consistency and magnitude of the differences observed between New Relic AI (NRAI) enabled accounts and non-NRAI accounts suggest that there are notable, quantifiable benefits to be gained from the use of AI-enhanced observability capabilities observed across months, metrics, and environments.

For the purposes of this report, New Relic AI or NRAI refers to the observability platform's generative AI features and AIOps features, which use a mixture of ML, generative, and deterministic methodologies to create value.

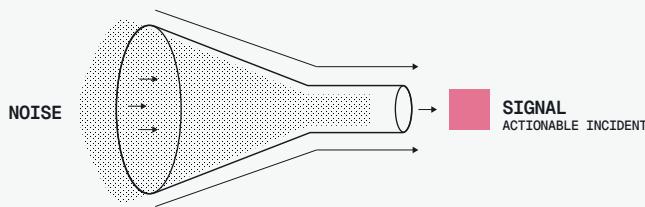
This report focuses on what the data suggests and how these outcomes plausibly connect.

KEY FINDINGS

AI-strengthened observability can decrease noise and increase productivity

IN 2025, NEW RELIC HAD 2.2 BILLION INCIDENTS CREATED IN PRODUCTION ENVIRONMENTS

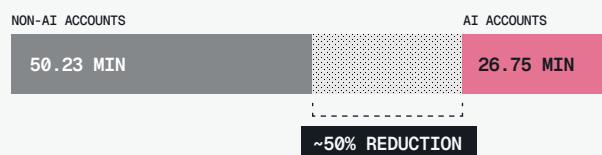
This massive volume of data highlights the alert fatigue crisis facing modern engineering teams and underscores the critical need for platform-level intelligence to manage the sheer scale of modern software operations.



25% FASTER RESOLUTION VELOCITY WITH AI

Across 2025, accounts using New Relic AI maintained a consistently lower Mean Time to Close (MTTC) than their peers, resolving issues roughly 25% faster on average. During peak performance months, the gap widened further. In May 2025, AI-enabled accounts averaged 26.75 MTTC, compared to 50.23 minutes for non-AI accounts.

MEAN TIME TO CLOSE IN MINUTES (MAY 2025)



REDUCING NOISE: A 2X HIGHER CORRELATION RATE

To combat operational noise, those using New Relic AI achieved a 2X higher correlation rate compared to non-AI configurations. A correlation rate is a measure of how effectively the software takes a massive pile of repetitive error messages and groups them into a single, actionable “incident.” By intelligently grouping related events, the platform can help prevent “noise” from overwhelming engineers, helping teams focus on root causes rather than symptoms — and spend more time shipping features.

2X  **HIGHER CORRELATION RATE**

AI MAKES TEAMS MORE PRODUCTIVE: THE 5X DEPLOYMENT MULTIPLIER

AI makes teams more productive. AI adoption has created a widening gap in deployment velocity. While non-AI users averaged 87 deployments per day during peak periods, AI-empowered teams achieved up to 453 deployments per day (a 5X increase in the ability to ship new features and updates to market).

5X  **DEPLOYMENT MULTIPLIER**

THE COMPOUNDING IMPACT

of AI-strengthened observability

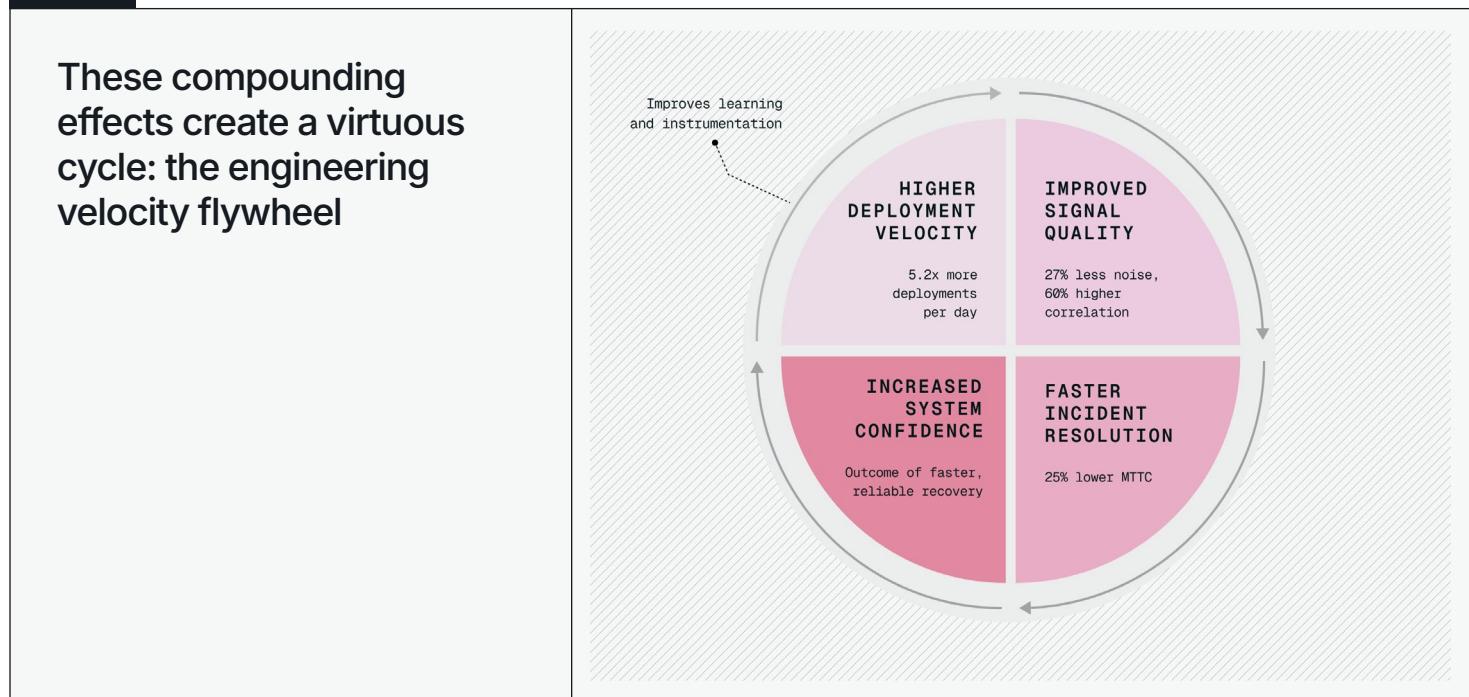
Engineering productivity can be hindered by a deluge of telemetry data. Distributed systems generate thousands (sometimes millions) of signals per minute, and most of it is noise. This report tracks how AI-strengthened observability fundamentally cuts through noise by altering the process of gleaning valuable insights from telemetry and leveraging them in the software development lifecycle (SDLC) process as measured against de-identified and aggregated metrics of New Relic usage data.

The data reveals a compounding effect. The journey begins with signal optimization, where AI filters out the environmental noise that can lead to evaluation burnout. This clarity directly fuels faster and more precise incident resolution, enabling teams to minimize negative business impacts from disruptions while reducing engineer toil.

The final result is a considerable surplus of engineering hours reinvested into R&D, allowing organizations to ship code at a much higher frequency, approximately 80% higher on average.

Our 2025 telemetry data shows that AI-strengthened issue correlation (which groups related alerts, events, and signals into manageable incidents) improves the signal-to-noise ratio and drives the reinforcing gains shown below.

FIG 1.A



SIGNAL OPTIMIZATION

Eliminating the operational noise

Before you can move fast, you need to see clearly. In practice, engineering organizations are overwhelmed by noise: thousands of alerts with little context. The New Relic observability platform addresses this by structuring how operational data is evaluated, progressively reducing noise while preserving what matters.

The process begins with raw alerts generated across the system. Related alerts are grouped into an issue that reflects notable or abnormal behavior, which could be severe enough to warrant an incident. New Relic AI or the platform's AIOps features analyze the potential impact of the issue to help gauge severity based on what performance thresholds are important to or flagged by the user.

The system then identifies relationships between issues across services, dependencies, and infrastructure and correlates these with changes (such as deployments and anomalies) in key signals to determine the most likely cause. All of this information is displayed on a New Relic page where a responder lands after being paged, allowing them to quickly analyze the critical data related to a single underlying problem. By consolidating incidents into an issue, redundant notifications are reduced and engineers are presented with a clear work item tied to a suspected root cause.

This approach allows teams to maintain visibility without being overwhelmed and to hyper-focus on the suspected issue and root cause. Rather than investigating every symptom, engineers can focus on resolving the issue that explains the behavior they are seeing across the stack.

THE HIERARCHY OF ACTIONABLE DATA

- **Alerts:** Raw threshold violations generated by monitoring conditions. High volume, low context.
- **Issues:** The result of correlation across alerts. Issues represent a potential problem and serve as the primary work item for assessing the potential impact of a disruption before declaring an Incident.
- **Escalations:** The result of performing an escalation on an issue, which often involves paging in additional engineers or SREs to help remediate a problem. This is the primary work item while investigators identify and apply available mitigating actions that will return a system to a good state.

These terms describe the New Relic internal operational hierarchy and reflect common information technology service management (ITSM practices) employed by New Relic users.

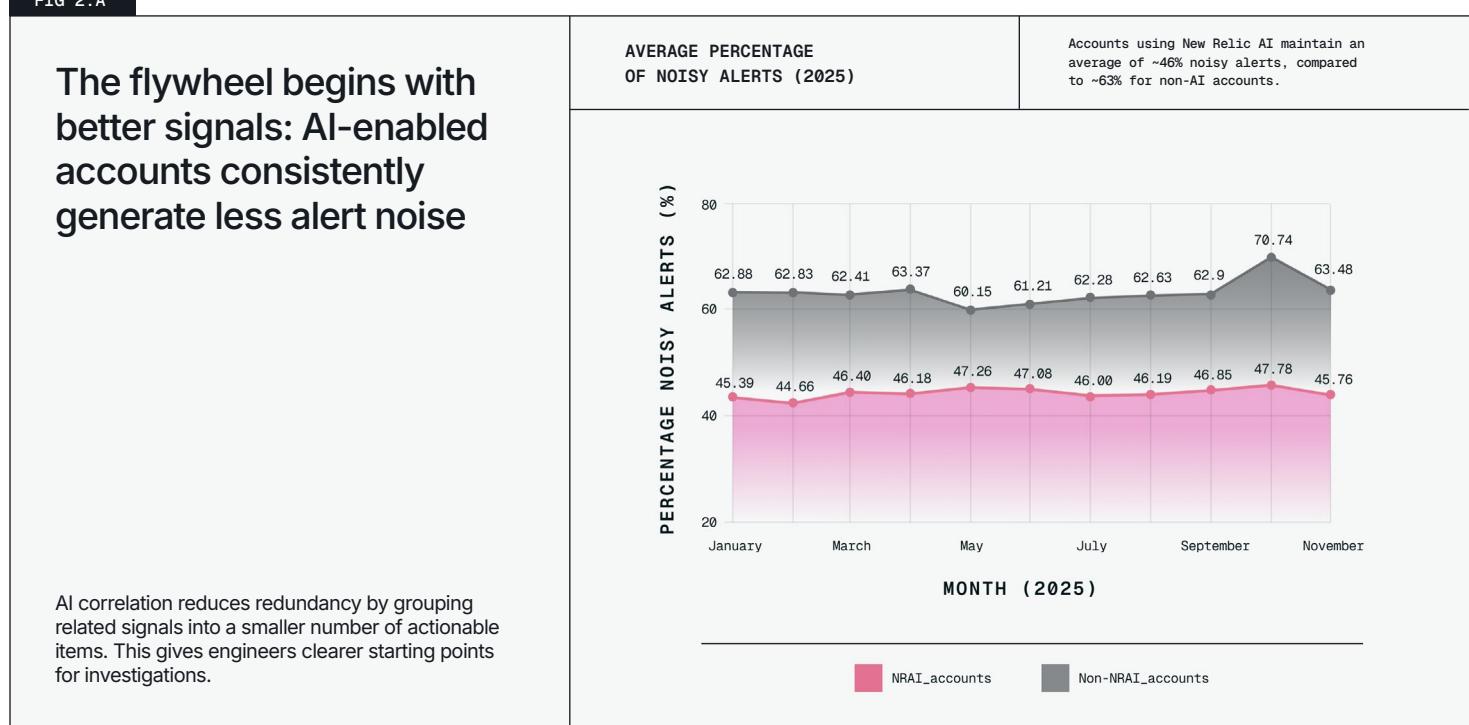
AI-CORRELATION

The analysis layer that evaluates relationships between alerts to surface a single issue and relates that issue to potential causes, such as deployments, anomalies, and errors. Notification volume is reduced as a downstream effect of issue correlation and resolution times drop as a result of discovery of relationships between issues and inciting actions, such as a deployment.

AI-ENABLED ACCOUNTS

consistently generate less alert noise

FIG 2.A



In practice, alert noise doesn't come from "too much telemetry." It comes from poorly scoped alert policies.

We define a noisy policy as an alerting rule with more than six triggering conditions. These policies attempt to encode every possible failure mode into a single alert. The result is predictable: frequent firing, ambiguous context, and low confidence in whether an alert represents a real incident.

Over time, teams learn to treat these alerts as background noise. Engineers delay investigation, mute alerts entirely, or rely on tribal knowledge to decide what's safe to ignore. None of this shows up as an error rate, but it directly degrades response quality and speed.

AI-enabled accounts behave differently.

By correlating related signals into a smaller number of actionable incidents, AI-assisted workflows reduce the need for overly complex alert definitions in the first place. Instead of compensating for uncertainty with more conditions, teams rely on correlation to surface meaningful failures.

Across 2025, accounts using NRAI maintained an average noisy-alert rate of approximately 46% compared to about 63% for non-AI accounts, with non-AI environments frequently exceeding 70%. The gap is stable month over month, suggesting a structural difference in how an alert is designed and maintained, not a temporary tuning effect.

HIGHER CORRELATION RATES

may explain why alert noise drops

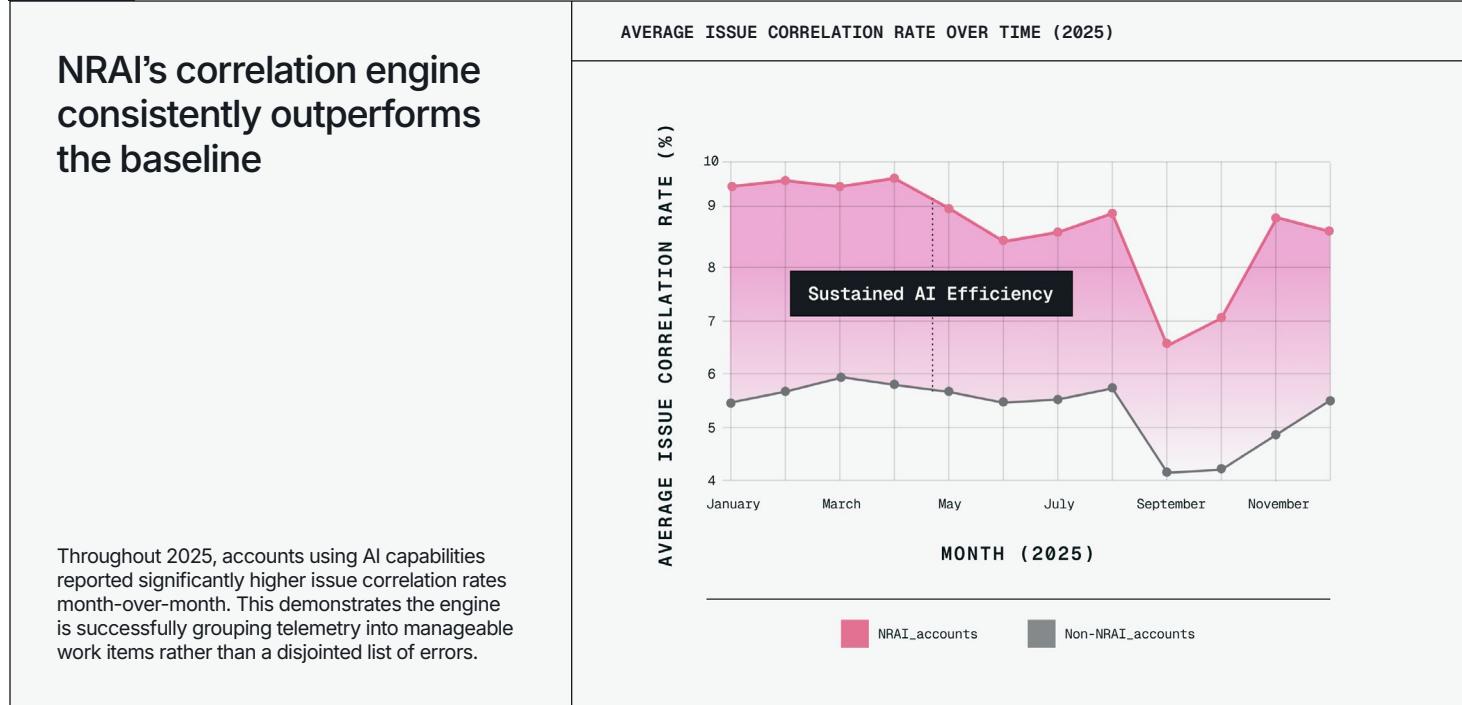
Alert noise and correlation are two sides of the same system. When correlation is low, teams compensate by adding conditions to alert policies. When correlation is high, alerts can remain simple because related signals are grouped automatically.

Correlation rate measures how effectively the platform groups related errors, events, and signals into a single actionable incident. Higher correlation rates indicate fewer fragmented alerts and more coherent investigative starting points for engineers.

Throughout 2025, accounts utilizing AI capabilities reported higher issue correlation rates month-over-month. This indicates not just better grouping in isolated incidents, but sustained operational efficiency as environments scale and change. For example, instead of separate CPU, memory, and latency alerts firing independently, correlated telemetry is surfaced as one actionable issue with related signals attached.

By reducing the number of events an engineer must process, New Relic AI gives engineers more time to focus on shipping features rather than managing alerts.

FIG 3.A



Throughout 2025, accounts using AI capabilities reported significantly higher issue correlation rates month-over-month. This demonstrates the engine is successfully grouping telemetry into manageable work items rather than a disjointed list of errors.

RESOLUTION VELOCITY

Closing the feedback loop

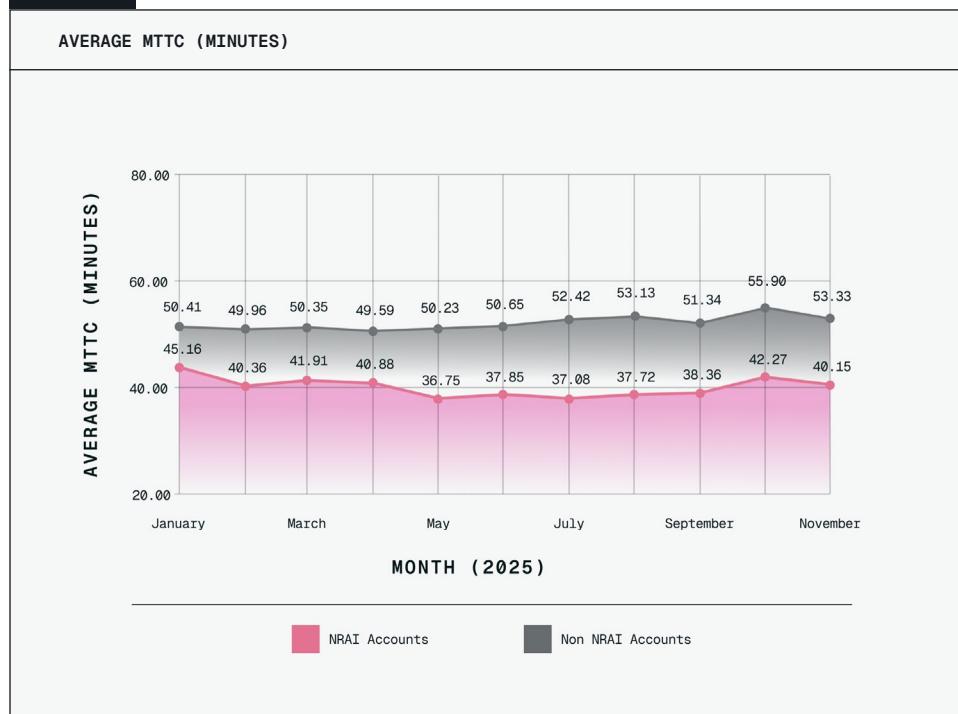
Once noise is suppressed, the primary metric of success shifts to how quickly an engineer can review and close issues and any resulting incidents. For issues (correlated alert events representing a potential problem) this is measured via Mean Time to Close (MTTC), which tracks the average duration issues remain open. MTTC bears a strong correlation to how quickly a system recovers from disruption.

THE MTTC ADVANTAGE

Across 2025, accounts using New Relic AI maintained a consistently lower Mean Time to Close (MTTC) than their peers, resolving issues roughly 25% faster on average. During peak performance months, the gap widened further. In May 2025, AI-enabled accounts averaged 26.75 MTTC, compared to 50.23 minutes for non-AI accounts.

This roughly 23-minute difference per issue compounds quickly. It reduces the wasted time that stalls investigations and preserves engineer momentum during high pressure events.

FIG 4.A



While resolution speed is influenced by many factors, the durability of this gap over time suggests that AI-strengthened capabilities can meaningfully reduce the manual “discover” phase of troubleshooting by surfacing context, related signals, and likely failure patterns earlier in the ITSM lifecycle.

FIG 5.A

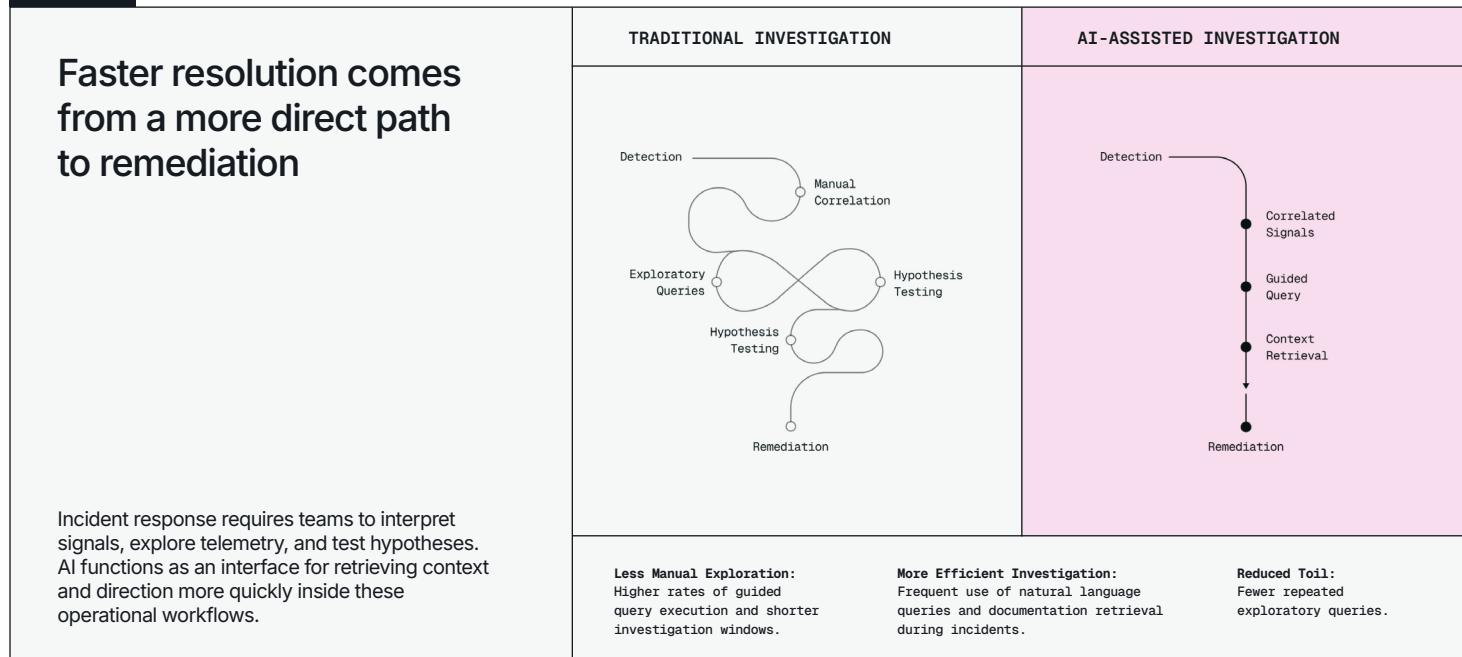
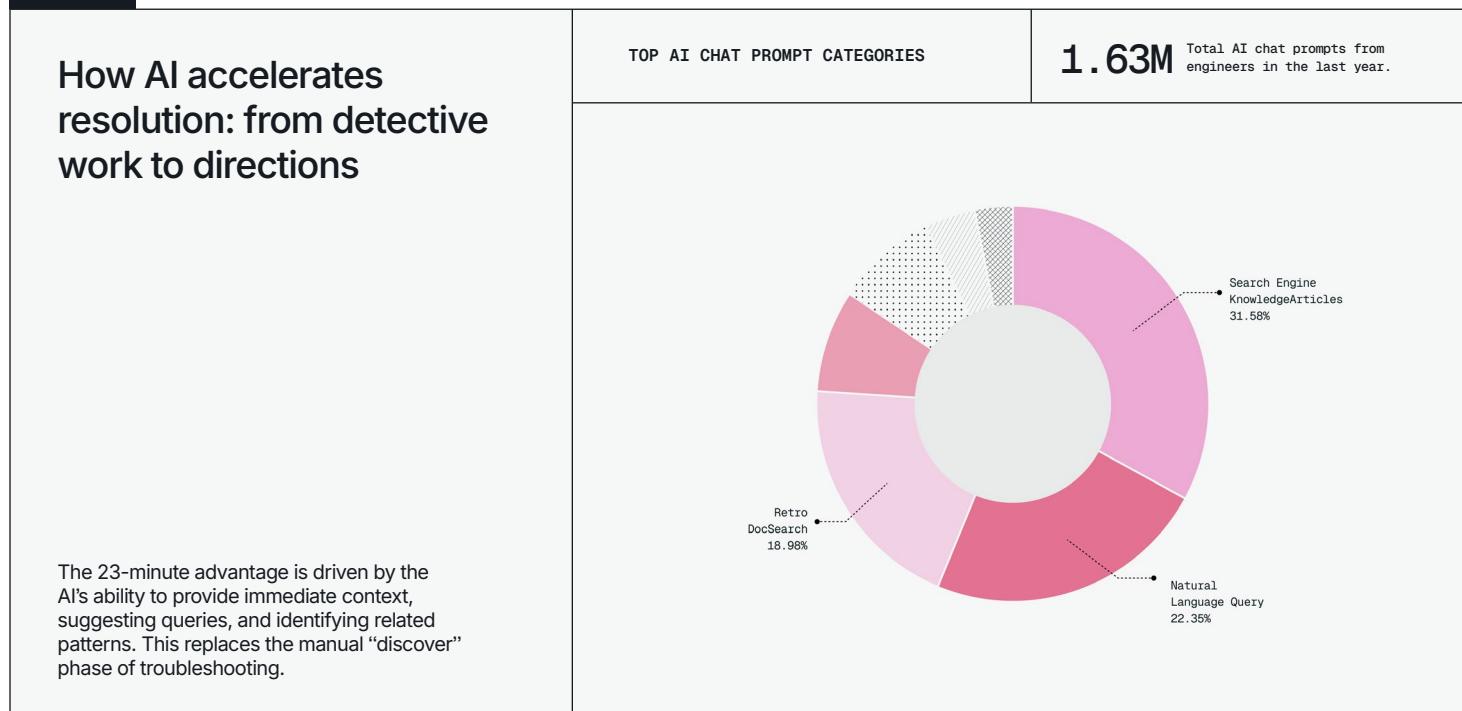


FIG 6.A



R&D THROUGHPUT

Reinvesting reclaimed hours into innovation

The final chapter of the 2025 findings is the most significant for technical leadership: R&D throughput. In many organizations, deployment frequency is capped not by developer speed, but by the “tax” of manual triage, troubleshooting, and reliability engineering tasks, plus the fear of breaking a fragile environment.

By reducing the time spent on toil (triaging and cutting through noise), organizations are successfully reinvesting that cognitive surplus into shipping code faster than their competitors.

QUANTIFYING THE SURGE IN SHIPPING

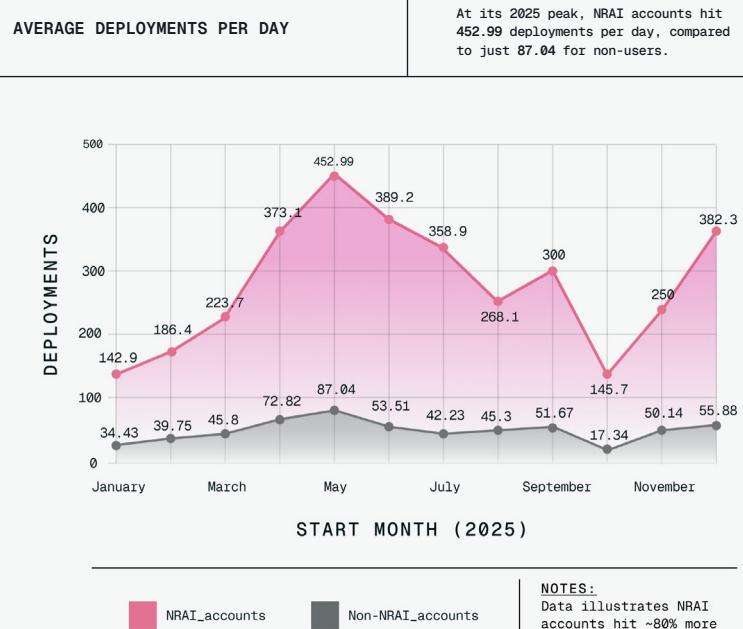
Our analysis of around 6.6 million users reveals a considerable disparity in output based on AI adoption: NRAI-enabled accounts maintain a daily deployment rate nearly double that of non-AI accounts.

At its 2025 peak, NRAI accounts hit an average of 452.99 deployments per day, compared to just 87.04 for non-users.

FIG 7.A

The innovation dividend: an 80% surge in deployment frequency

NRAI-enabled accounts maintain a daily deployment rate nearly double that of non-AI accounts.



THE INNOVATION DIVIDEND

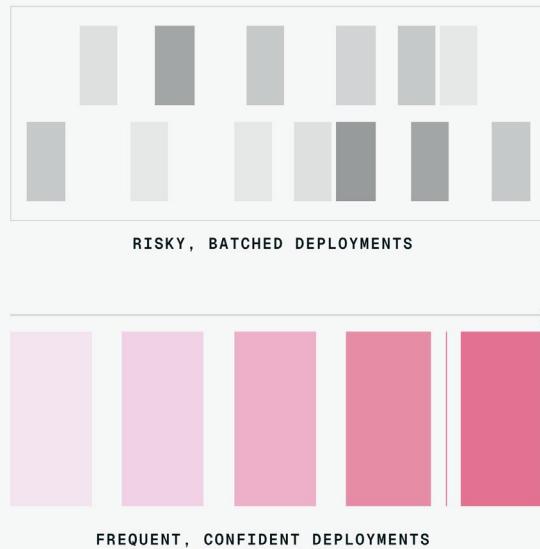
The story the data tells is one of a new operational baseline. By reducing noise and accelerating resolution, teams have unlocked the ability to focus approximately 80% more effort on building new features rather than maintaining existing ones.

This is a measure of an organization's ability to respond to the market. When the operational tax is minimized, the speed of the entire company accelerates. From the first correlated alert to the final deployment of the day, this report demonstrates that New Relic AI helps teams reach a new standard of high-velocity, high-stability engineering.

FIG 8.A

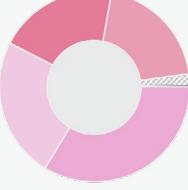
More deployments, more confidence

The data reveals a new operational baseline. With lower critical alert rates and faster resolution, teams are empowered to deploy smaller, more frequent updates rather than large, risky batches. When the operational tax is minimized, the speed of the entire company accelerates.



The question isn't whether AI adds value.
**THE QUESTION IS WHETHER
 YOU CAN AFFORD TO OPERATE
 WITHOUT IT?**

APPENDIX: AI USAGE AND PERFORMANCE DETAILS

NRAI QUERY VOLUME BY %		AI TOOL LATENCY & RESPONSE PERFORMANCE																																																										
		<table> <thead> <tr> <th>Tool Name</th> <th>p50(sec) (50%)</th> <th>p95(sec) (95%)</th> <th>p99(sec) (99%)</th> </tr> </thead> <tbody> <tr><td>Root Cause Tool</td><td>3.573</td><td>3.573</td><td>3.573</td></tr> <tr><td>Retro DocSearch</td><td>3.359</td><td>6.469</td><td>8.313</td></tr> <tr><td>Search Engine NR_Docs</td><td>2.941</td><td>4.086</td><td>4.339</td></tr> <tr><td>Search Engine ForumTopics</td><td>2.934</td><td>3.863</td><td>4.14</td></tr> <tr><td>Insight Cause Interpreter</td><td>2.678</td><td>3.448</td><td>3.448</td></tr> <tr><td>Search Engine KnowledgeArticles</td><td>2.559</td><td>3.254</td><td>3.785</td></tr> <tr><td>Search Engine ForumBestAnswers</td><td>2.119</td><td>2.544</td><td>2.54</td></tr> <tr><td>SessionReplay Trail Analyzer</td><td>1.182</td><td>1.185</td><td>1.185</td></tr> <tr><td>Summarize NRQL Result</td><td>1.561</td><td>1.185</td><td>1.968</td></tr> <tr><td>Natural Language Query Indexer</td><td>1.144</td><td>1.185</td><td>1.968</td></tr> <tr><td>Natural Language Query</td><td>1.141</td><td>1.641</td><td>2.344</td></tr> <tr><td>Condition Coverage Gaps</td><td>0.617</td><td>0.617</td><td>0.617</td></tr> <tr><td>Issue Insights v2</td><td>0.259</td><td>0.306</td><td>0.306</td></tr> </tbody> </table>			Tool Name	p50(sec) (50%)	p95(sec) (95%)	p99(sec) (99%)	Root Cause Tool	3.573	3.573	3.573	Retro DocSearch	3.359	6.469	8.313	Search Engine NR_Docs	2.941	4.086	4.339	Search Engine ForumTopics	2.934	3.863	4.14	Insight Cause Interpreter	2.678	3.448	3.448	Search Engine KnowledgeArticles	2.559	3.254	3.785	Search Engine ForumBestAnswers	2.119	2.544	2.54	SessionReplay Trail Analyzer	1.182	1.185	1.185	Summarize NRQL Result	1.561	1.185	1.968	Natural Language Query Indexer	1.144	1.185	1.968	Natural Language Query	1.141	1.641	2.344	Condition Coverage Gaps	0.617	0.617	0.617	Issue Insights v2	0.259	0.306	0.306
Tool Name	p50(sec) (50%)	p95(sec) (95%)	p99(sec) (99%)																																																									
Root Cause Tool	3.573	3.573	3.573																																																									
Retro DocSearch	3.359	6.469	8.313																																																									
Search Engine NR_Docs	2.941	4.086	4.339																																																									
Search Engine ForumTopics	2.934	3.863	4.14																																																									
Insight Cause Interpreter	2.678	3.448	3.448																																																									
Search Engine KnowledgeArticles	2.559	3.254	3.785																																																									
Search Engine ForumBestAnswers	2.119	2.544	2.54																																																									
SessionReplay Trail Analyzer	1.182	1.185	1.185																																																									
Summarize NRQL Result	1.561	1.185	1.968																																																									
Natural Language Query Indexer	1.144	1.185	1.968																																																									
Natural Language Query	1.141	1.641	2.344																																																									
Condition Coverage Gaps	0.617	0.617	0.617																																																									
Issue Insights v2	0.259	0.306	0.306																																																									
<table> <tbody> <tr><td>Search Engine KnowledgeArticles</td><td>123 k</td><td>31.59%</td></tr> <tr><td>Natural Language Query</td><td>86.9k</td><td>22.35%</td></tr> <tr><td>Retro DocSearch</td><td>73.8k</td><td>18.99%</td></tr> <tr><td>Search Engine NR_Docs</td><td>72.8k</td><td>18.74%</td></tr> <tr><td>Search Engine ForumTopics</td><td>6.86k</td><td>1.76%</td></tr> <tr><td>Insight Cause Interpreter</td><td>578</td><td>0.15%</td></tr> <tr><td>Issue Insights v2</td><td>232</td><td>0.054%</td></tr> <tr><td>NL2NRQL</td><td>113</td><td>0.029%</td></tr> <tr><td>Summarize NRQL Result</td><td>105</td><td>0.027%</td></tr> <tr><td>Logs Parsing Pattern Generator</td><td>105</td><td>0.026%</td></tr> <tr><td>SessionReplay Trail Analyzer</td><td>89</td><td>0.023%</td></tr> <tr><td>Alert Condition Recommendation</td><td>58</td><td>0.011%</td></tr> </tbody> </table>		Search Engine KnowledgeArticles	123 k	31.59%	Natural Language Query	86.9k	22.35%	Retro DocSearch	73.8k	18.99%	Search Engine NR_Docs	72.8k	18.74%	Search Engine ForumTopics	6.86k	1.76%	Insight Cause Interpreter	578	0.15%	Issue Insights v2	232	0.054%	NL2NRQL	113	0.029%	Summarize NRQL Result	105	0.027%	Logs Parsing Pattern Generator	105	0.026%	SessionReplay Trail Analyzer	89	0.023%	Alert Condition Recommendation	58	0.011%																							
Search Engine KnowledgeArticles	123 k	31.59%																																																										
Natural Language Query	86.9k	22.35%																																																										
Retro DocSearch	73.8k	18.99%																																																										
Search Engine NR_Docs	72.8k	18.74%																																																										
Search Engine ForumTopics	6.86k	1.76%																																																										
Insight Cause Interpreter	578	0.15%																																																										
Issue Insights v2	232	0.054%																																																										
NL2NRQL	113	0.029%																																																										
Summarize NRQL Result	105	0.027%																																																										
Logs Parsing Pattern Generator	105	0.026%																																																										
SessionReplay Trail Analyzer	89	0.023%																																																										
Alert Condition Recommendation	58	0.011%																																																										

© Copyright 2026, New Relic, Inc. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies. 01.2026

