



2026 State of Production Reliability and AI Adoption Report



Table of Contents

Executive Summary 02

Key Findings 03

Reactive Incident Management Is Consuming Engineering Capacity 04

Alert Fatigue Is Now a Reliability Problem 06

Current Incident Management Approaches Are Burning Profits And Teams 08

Stale Incident Management Practices Are Holding Teams Back 10

The AI Divide Between Leadership and Practitioners 12

Survey Methodology 17



Executive Summary

For years, platform and IT engineers have been asked to do two things at once: build reliable systems that power business growth and keep production running when things break. We surveyed more than 1,000 SRE, DevOps and IT operations professionals to get their insights on what happens when the second job starts winning. What the data describes is a system working exactly as it was built, just not in anyone's favor.

Incident management is no longer a support function operating in the background. It has become a major consumer of engineering time, a driver of team burnout and a source of direct financial exposure most organizations are significantly underestimating.

The numbers are unambiguous. The majority of engineering teams spend 40% or more of their time on incident management rather than building. Nearly half of organizations (44%) experienced an outage in the past year tied directly to suppressed or ignored alerts, and **78% report at least one incident where no alert fired at all, leaving engineers to learn about failures from customers rather than their own systems.**

Alert systems generate too much noise for engineers to treat every alert as a signal, tools are fragmented across platforms that don't communicate and runbooks exist but are out of date. These conditions compound: the more time teams spend on reactive incident work, the less time they have to address the underlying root causes that generate incidents in the first place.

AI is broadly recognized as part of the solution, and investment is flowing. Yet the data reveals a significant disconnect between executive perception and operational reality. Seventy-four percent of C-suite respondents say their organizations actively use AI for incident management, while only 39% of practitioners say the same.

What has been purchased and what is running in production are two different things, and the engineers who would benefit most from AI-assisted triage, root cause analysis and alert correlation are the least likely to have access to it. This report examines that divide in detail, along with where AI is already delivering results and what is blocking broader adoption.

The financial exposure from a single high-severity event can routinely reach six figures with 34% of organizations reporting downtime costs of \$100,000 or more per hour.

Key Findings

01 The top challenges SRE/incident management teams face:

- Alert fatigue and noise
- Insufficient automation
- Knowledge silos and documentation gaps
- Difficulty identifying root causes

02 Most engineering teams spend **40% or more of their time on incident management** versus innovation.

03 More than **70% of alerts received** are not actionable according to **57%** of organizations.

04 **More than 1 in 3** organizations (**35%**) report their **engineers occasionally ignore or dismiss alerts** due to alert fatigue and **nearly half (44%) experienced an outage linked to ignored or suppressed alerts** in the past year.

05 **78%** of companies report they've **experienced an outage or incident where no alert was issued.**

06 **More than quarter of on-call engineers report burnout symptoms** related to incident management according 39& of companies

07 **61%** of organizations estimate **downtime costs \$50,000 or more per hour; 34%** put that figure at **\$100,000 or more.**

08 **MTTR** (Mean Time to Resolve) is **the number one KPI** when it comes to incident response.

09 **74%** of C-suite respondents say their organization **actively uses AI for incident management,** but **only 39% of practitioners** agree.

10 **The top barriers to AI adoption are:**

- Budget constraints
- Data quality or availability issues
- Security and compliance concerns



Reactive Incident Management Is Consuming Engineering Capacity

Most organizations think of incident management as a cost of doing business with complex infrastructure. The data suggests it has become a structural tax on engineering capacity that is compounding faster than most organizations are accounting for.

The majority (53%) of engineering teams spend 40% or more of their time on incident management. This figure is consistent across all seniority levels, from practitioners to the C-suite, and company sizes. It is not a small-team problem or practitioner problem. It is an industry-wide business challenge.

In an environment where product delivery speed defines competitive advantage, this becomes a direct tax on the engineering capacity available for building. It's the difference between shipping and maintaining, between building new capability and keeping existing systems stable enough to stay in production. Current incident management practices are no longer sustainable.

The cost of an individual incident extends well beyond the first responder. Nearly all organizations (93%) involve at least three engineers when a SEV-0 or P0 incident occurs and close to 40% pull in six to ten

people. This represents a compounding cost: engineers pulled from other work and velocity lost.

On average, teams spend 28 hours a week on troubleshooting incidents and root cause analysis. That's nearly 3 full working days every week.

The toil doesn't stop at resolution. More than a third of teams (36%) spend five to ten hours every week on incident reports and post-mortems. That is time after the incident is closed, time that does not show up in MTTR metrics but still comes out of engineering capacity.

When asked, "Who is primarily responsible for incident management at your organization?" Twenty percent of companies indicated it varies by incident, meaning no single team owns the problem. This ambiguity further contributes to resolution delays as organizations scramble to assemble the right team for every outage.

40% of engineering **time spent on incident management** on average

93% pull **three or more engineers** into a major incident


40% involve **six to ten people** to resolve a single SEV-0/P0 incident



Tool Sprawl and the Cost of Context Switching

MTTR is the top KPI cited by 61% of respondents. Yet the tooling environment engineers operate in actively works against resolution speed.

Eighty-three percent of teams work across four or more tools when handling a live incident, including observability platforms, monitoring systems and incident management software, and 41% sort through seven or more separate tools.



of teams work across four or more tools when handling a live incident

Every tool switch during an active incident is a delay. Engineers are not just diagnosing and fixing a problem; they are also translating data between systems, re-establishing context after each switch and reassembling a picture of what is happening across platforms that were never designed to work together. In an environment where minutes translate directly to dollars, this fragmentation is a material contributor to extended downtime costs.

What are your company's key KPIs when it comes to incident response?

01 Mean Time To Resolution (MTTR)

02 Number of incidents per month

03 Customer-impacting incidents

04 Mean Time to Detect

05 Mean Time to Acknowledge



Alert Fatigue Is Now a Reliability Problem

Alert fatigue was identified as the top operational challenge for SRE and incident management teams in this survey, ranked above insufficient automation, knowledge silos and difficulty identifying root causes. This has been a known problem for years, but what this data makes clear is that alert fatigue has moved from being a morale and productivity issue to being a direct contributor to infrastructure failures.

The chain of causation is short and well-documented in the responses: too many alerts, most of them not actionable, train engineers to discount the signal. When engineers discount the signal, real events get missed. When events get missed, outages happen that the monitoring system was theoretically designed to prevent.

The volume numbers alone are significant, but they become more consequential when set against signal quality. On-call teams are receiving a high volume of alerts that do not require action, and they know it.

The top-ranked challenges SRE/incident management teams face:


01 Alert fatigue and noise

02 Insufficient automation


03 Knowledge silos and documentation gaps

04 Difficulty identifying root causes


05 Integration challenges between tools



of on-call teams report getting at least **10 alerts per day**



of on-call teams report getting up to **50 alerts per day**



More than 70% of alerts received **are not actionable** according to 57% of organizations

Alert fatigue persists, even as 73% of teams report reviewing and tuning alert thresholds at least monthly. However, 21% of practitioners say they only review thresholds when problems arise, never, or annually (compared to just 3.5% of C-suite who believe this to be the case). This helps explain why alert noise doesn't self-correct; the people closest to it have the least systematic process for managing it.

When Engineers Stop Responding, Outages Follow

The behavioral consequence of sustained alert noise shows up clearly across the data. Over half (53%) of organizations report their engineers occasionally ignore or dismiss alerts due to alert fatigue and nearly half of organizations experienced an outage or incident in the past year tied directly to an ignored or suppressed alert.

When engineers learn that most alerts do not require a response, they start acting accordingly, and when engineers stop responding, outages follow:

83% of organizations report their teams are ignoring alerts

35% of organizations report their engineers occasionally ignore or dismiss alerts due to alert fatigue (10-30% of the time)

18% of organizations say engineers ignore alerts frequently or very frequently (more than 31% of the time)

44% of companies report their organization experienced an outage or incident that was linked to ignored or suppressed alerts over the past year



Current Incident Management Approaches Are Burning Profits and Teams

The human cost of incident management is visible in burnout rates, alert fatigue and engineering time lost to reactive work. The financial costs are also significant when analyzing the volume of issues handled monthly and the per-incident financial exposure.

What an Hour of Downtime Costs

Sixty-one percent of organizations estimate that one hour of infrastructure downtime costs \$50,000 or more. Considering the average MTTR is 1-2 hours and organizations experience an average of 20 incidents per month, the direct financial exposure of infrastructure downtime is considerable, even before accounting for engineering hours spent on diagnosis, root cause analysis and post-mortem documentation.

61% estimate downtime costs \$50,000 or more per hour

34% estimate downtime costs \$100,000 or more per hour

58% report an MTTR between 30 minutes to two hours for critical incidents

26% report an MTTR over two hours for critical incidents

Engineering Burnout Adds a Compounding Cost

The financial exposure from downtime is measurable and immediate. The cost of burnout is slower to accumulate but ultimately more expensive. Nearly 40% of organizations report that more than a quarter of their on-call engineers show burnout symptoms tied to incident management, making reactive operations no longer sustainable at scale. The cost of replacing a burned-out senior engineer, including recruiting, onboarding and the lost institutional knowledge they carry out the door, exceeds the cost of most individual incidents.

26–50% of on-call engineers are reported as burned out by a quarter of organizations

51–75% of on-call engineers are reported as burned by 11% of organizations

75%+ of on-call engineers are reported as burned out by 3% of organizations



Customers Are Finding Failures Before Monitoring Tools

Organizations report that almost 40% of their incidents are discovered by customers. When a customer finds a failure before the monitoring stack does, the monitoring stack has not done its job.

Additionally, 78% of respondents report at least one incident where no alert fired at all, meaning the monitoring system contributed nothing to detection.

51%

of organizations have customers reporting more than 30% of incidents

70%

report they experienced up to 19 incidents in the last year where no alert fired.



Stale Incident Management Practices Are Holding Teams Back

Alert volume and MTTR are the visible metrics. Below them sit a set of structural conditions that make improvement slow: documentation that is out of date and there’s a significant divide between what observability practices leadership believes are in place and what engineers are working with.


Forty percent of organizations report they have runbooks and incident response playbooks for major systems, but those playbooks require updating. For practitioners, an outdated runbook is often worse than no runbook at all: it introduces false confidence, sends engineers down wrong paths under time pressure and creates post-incident disputes about whether the documented procedure was followed correctly.

Leadership is currently operating under the assumption that their company’s runbooks are comprehensive, regularly updated and widely used, which runs contrary to the reality practitioners are dealing with.

C-suite respondents were 23 points more likely than practitioners to describe their runbooks as comprehensive, regularly updated and widely used (57% vs. 34%).

The 23-point perception gap is not a communication failure. Executives evaluate documentation quality based on whether it exists and was recently reviewed. Practitioners evaluate it based on whether it works when they need it at 2am during a production incident.

Runbooks described as current and widely used





Observability Maturity: Another Perception Divide

When asked to describe the maturity of their organization's observability function, respondents split roughly into three tiers: 28% described it as advanced, 58% as moderate and 12% as basic. The distribution shifts meaningfully by role.

Nearly 20% of practitioners describe observability as basic, standard monitoring and alerting only, compared to 6.5% of C-suite. The disparity between what leadership believes the observability stack can do and what engineers are working with has direct consequences for detection speed, alert quality and incident resolution.

31.5%

of C-suite describe their organization's observability as "advanced," compared to 22.3% of practitioners.

20%

of practitioners say observability is "basic — standard monitoring and alerting only," compared to 6.5% of C-suite.

How would you describe the maturity of your organization's observability function/approach/system/practice?

- **28%** said: Advanced
Full observability across all systems with traces, metrics and logs
- **58%** said: Moderate
Good monitoring and logging with some distributed tracing
- **12%** said: Basic
Standard monitoring and alerting in place



The AI Divide Between Leadership and Practitioners


On AI in incident management, executives and practitioners are operating in two different realities. Executives are largely confident that AI deployment is underway and AI tools are being used effectively for incident management. The engineers on the ground, the people who would benefit most directly from AI-assisted triage, anomaly detection and automated root cause analysis, are reporting a very different reality.

Seventy-four percent of C-suite respondents say their organization actively uses AI for incident management, but only 39% say the same. This 35-point divide reflects the distance between a procurement decision and a production deployment. AI tools can be purchased, licensed and integrated at the platform level without being meaningfully available to the engineers running incidents day to day. Executives see the investment. Practitioners experience the tools, and many of them are not experiencing AI.


74% of C-suite say their organization actively uses AI for incident management, only 39% of practitioners say the same.

Practitioners are not resistant to AI, more than half (52%) say they are actively evaluating AI solutions, the highest evaluation rate of any group. They are waiting for it to exist in their workflows rather than in their organization's software inventory.

Actively using AI for incident management



Evaluating AI tool usage for incident management





The AI Impact Gap

The adoption gap is only part of the story. Among the respondents who report using AI, sentiment on its effectiveness varies dramatically by seniority. C-suite respondents are nearly three times as likely as practitioners to say AI has significantly reduced operational toil, defined as a 30% or greater reduction in workload.

Among practitioners who use AI tools, 28% report the impact on their workload has been less than 10% and 5% report AI has increased their operational toil due to added complexity.

AI has significantly reduced operational toil (30%+)

36%

C-suite

27%

Leadership

13%

Practitioners

AI has had minimal impact (less than 10% reduction)

5%

C-suite

13%

Leadership

22%

Practitioners



Company Size and AI Adoption in Incident Management


AI adoption patterns by company size reveals a counterintuitive pattern. The largest organizations are not the furthest ahead.

Mid-market organizations with 500 to 4,999 employees, lead active AI adoption at 70% These companies are large enough to have mature SRE and DevOps functions with real incident volume justifying the investment, and agile enough to move through procurement and deployment cycles without the friction that slows enterprise organizations.

Larger enterprises with more than 5,000 employees report the lowest active adoption at 47%, a full 22 points behind mid-market leaders.


Smaller organizations with 100 to 499 employees report a 56% adoption rate, ahead of large enterprises. Smaller teams feel the burden of manual incident management acutely and typically have fewer legacy systems creating integration complexity.

56%




100 - 499 employees

69%




500 - 1,999 employees

68%



2,000 - 4,999 employees

47%



5,000+ employees

How Organizations Are Using AI in Incident Management

Among organizations that have deployed AI for incident management, the leading use cases directly address the problems ranked highest by respondents throughout this survey. Automated root cause analysis tops the list, followed by anomaly detection and prediction and alert correlation and noise reduction. All three target the same underlying problem: engineers spending time on work that machines can do faster and with more consistent accuracy, leaving human judgment for the decisions that require it.

AI use cases in incident management, among AI users:

- 01 Automated root cause analysis
- 02 Anomaly detection and prediction
- 03 Alert correlation and noise reduction
- 04 Automated incident triage and prioritization
- 05 Incident reports and post-mortem

Barriers to AI Adoption in Incident Management

The top barriers to AI adoption are practical rather than philosophical with budget constraints cited as number one, followed by data quality and availability issues and security and compliance concerns. These are procurement and infrastructure readiness problems, not signs that organizations are skeptical about AI's value.


The appetite is there. The 52% of practitioners who are actively evaluating AI solutions are not waiting to be convinced of the value; they are waiting for the conditions that allow responsible deployment.

Leadership and practitioners diverge on what is blocking AI adoption. C-suite respondents are significantly more likely than practitioners to cite "concerns about AI increasing system complexity" as a barrier (43% vs. 28%) and more likely to cite data quality issues (42% vs. 31%). Budget constraints are shared across groups. For executives, the blockers are largely architectural; for practitioners, security and complexity concerns are nearly as prominent.

Barriers to AI adoption in incident management:

- 01 Budget constraints
- 02 Data quality or availability issues
- 03 Security and compliance concerns
- 04 Concerns about AI increasing system complexity
- 05 Lack of technical expertise





With 45% of large enterprises actively evaluating AI, there is likely to be a significant near-term AI deployment wave. The barriers above are largely addressable: compliance requirements can be satisfied, ROI can be demonstrated from mid-market deployments, integration friction can be reduced with purpose-built tools. When those conditions are met, the evaluation stage converts to deployment and the practitioner adoption rate, currently 39%, will follow.



Survey Methodology

The 2026 State of Incident Resolution and AI Adoption Report is based on a survey of 1,039 SRE, DevOps and IT operations professionals at organizations with 100 or more employees, conducted in February 2026.

Respondents by role

C-suite executives (CTO, CIO, CDO, Chief AI Officer) → 209 respondents (20%)

IT and engineering leadership (Head/VP/ Director of IT, Engineering, SRE, ITSM, Observability) → 417 respondents (40%)

Practitioners (software engineers, system administrators, DevOps engineers, SREs, ITOps engineers, platform engineers) → 413 respondents (40%)


Respondents by company size

100 to 499 employees → 291 respondents (28%)

500 to 1,999 employees → 189 respondents (18%)

2,000 to 4,999 employees → 223 respondents (21%)

5,000 or more employees → 336 respondents (32%)





2026 State of Production Reliability and AI Adoption Report

About NeuBird

NeuBird is pioneering the use of agentic AI for IT operations to address the scarcity in skilled human talent in keeping up with an increasingly complex modern technology stack. NeuBird simplifies complex data analysis and offers actionable insights in real time, empowering companies to innovate faster and more effectively.

Visit [Neubird.ai](https://neubird.ai) to learn more.