

# State of SRE Operations 2026

The operational reality for SRE, platform, and infrastructure teams.

# The Operational Reality for SRE, Platform, and Infrastructure Teams

OpsWerks surveyed SRE, platform, and infrastructure engineering practitioners\* to understand what they deal with operationally: not in theory, but in production, on call, every week.

The data reflects the field as it actually operates: reactive-heavy, under-resourced on coverage, and increasingly clear on what good operational partnership requires.

Respondents spanned a cross-section of industries, including financial services, technology and software, travel and hospitality, automotive, cybersecurity, pharmaceuticals, and entertainment.

The questions covered team structure, operational challenges, coverage models, alert quality, vendor selection, and investment priorities.

Use these findings to benchmark your team, sharpen your operational roadmap, or anchor a conversation with leadership about where the real capacity constraints are.

\* Based on responses from 32 SRE, platform, and infrastructure engineering practitioners surveyed in early 2026. Directional; reflects practitioner experience, not a statistically representative sample. Percentages are rounded to the nearest whole number.

# Key Takeaways

- 01 Alert fatigue and platform bottlenecks are the top operational challenge.**

Alert fatigue and incident overload (25%) is the single biggest operational challenge; platform support bottlenecks follow close behind (22%). Combined, 47% of teams name them the core capacity drain.

---
- 02 Most teams are reactive at least half the time.**

Only 38% report spending more time on proactive work than reactive. The majority operate at a 50/50 split or worse, which makes it structurally difficult to reduce future incident load.

---
- 03 24/7 coverage runs almost entirely on internal on-call rotations.**

72% rely on internal on-call. Only 9% hand coverage fully to an outsourced provider. Counting hybrid setups, just 13% involve an outside vendor at all.

---
- 04 Alert noise is the operational baseline, not an outlier.**

63% say fewer than half their alerts are actionable. Alert noise is the operational baseline for most teams, and it compounds every other capacity problem.

---
- 05 Technical depth matters most when evaluating a vendor.**

Technical depth in their specific stack was the most-cited vendor selection factor, named by 66%. Speed of onboarding, 24/7 SLAs, and cultural fit are what get vendors shortlisted.


---
- 06 Observability and technical debt lead investment priorities.**

Improving observability (31%) tops the investment list. Reducing technical debt (25%) is second. Both reflect the same constraint: teams spend too much time firefighting to fix the systems that cause fires.

# What best describes your current role?

# 56%

of respondents are practicing site reliability engineers



Percentages are rounded to the nearest whole number.

SREs and practitioners who own reliability outcomes day-to-day make up the majority. Platform engineers and engineering managers round out the next tier.

# How large is your SRE or platform engineering team?

# 38%

of teams run full production on fewer than 10 engineers, the single largest group




Most teams represented here are lean. Teams under 30 engineers are the majority, and under 10 is the single largest group. Small teams holding full production responsibility is standard.

# What is your team's biggest operational challenge right now?

47%

cite alert fatigue or platform support bottlenecks as their primary challenge




Alert fatigue is the single biggest challenge, with platform bottlenecks close behind. Legacy systems and high-stakes migrations follow. On-call burnout ranks lower than it should: not because it is absent, but because teams have normalized it.

# How would you describe your team's current reactive vs. proactive work split?

63%

are reactive half the time or more, leaving scarce capacity for the fixes that stop incidents recurring




Percentages are rounded to the nearest whole number.

Only 38% operate with more proactive than reactive work. Most are split evenly or skewed reactive. This makes it structurally difficult to reduce future incident load.

# How does your team handle after-hours and weekend coverage?

# 72%

carry 24/7 coverage entirely on internal on-call, absorbing the full human cost




Percentages are rounded to the nearest whole number.

Internal on-call is the overwhelming model. Very few teams have moved to an outsourced or hybrid provider. This is a sustainability question as much as a coverage question: 72% of teams absorb all of 24/7, including nights, weekends, and the compounding fatigue.

# Where does your team spend the most engineering effort?

66%

cite observability and monitoring as the top area consuming engineering time



Multi-select. Percentages reflect share of respondents selecting each option.

Observability and monitoring lead by a wide margin. Kubernetes orchestration and CI/CD follow, tied at 50%. These are the core operational surface areas for platform and SRE teams in 2026.

# What percentage of your alerts are actually actionable?

# 63%

say fewer than half their alerts are worth acting on: the rest is noise




Percentages are rounded to the nearest whole number.

Alert noise is the operational baseline, not an outlier problem. Only 16% report more than three-quarters of alerts are actionable. The rest manage a signal-to-noise problem that burns time, compounds burnout, and buries the alerts that actually matter.

# How would you describe your team's staffing and budget situation over the past year?

# 38%

are stable on paper with scope that keeps growing around them




Stable on paper with growing scope is the most common situation. It means the team absorbs more operational surface area without headcount to match. The gap between what is covered and what needs to be covered widens quietly.

# When evaluating an external operations partner, what factors matter most?

66%

put technical depth in their specific stack first, well ahead of price




Multi-select. Percentages reflect share of respondents selecting each option.

Technical depth leads. 24/7 SLAs, cultural fit, and speed of onboarding follow. Teams lead with capability and reliability, not cost: they are looking for outcomes, not bodies.

# If your team had more capacity tomorrow, where would you invest it first?

# 31%

want to invest more time in observability and monitoring



Improving observability leads investment priorities. Reducing technical debt and AI/ML tooling follow. These are the areas that reduce incident load; reactive work crowds them out.

## About the Data

This survey was conducted by OpsWerks among SRE, platform, and infrastructure engineering practitioners. Respondents represent organizations of varying scale, primarily in technology and software sectors.

The data is directional. It reflects practitioner experience, not a statistically representative sample of the broader market. Percentages are rounded to the nearest whole number. All findings are published in aggregate.

## About OpsWerks

For over a decade, OpsWerks has embedded with the world's most demanding platform and infrastructure engineering teams. We deliver outcome-based managed services across Cloud and Infrastructure, Platform Operations, Monitoring and Incident Response, AI and Data Platforms, and Security and Compliance. Fixed pricing. Guaranteed SLAs.

We do not bill hours. We own results. Teams that start with a single operational problem expand because the model works: fewer incidents, faster recovery, engineers freed to ship.

[Schedule a discovery call →](#)

# You define the outcomes. We own the delivery.

Managed services for leading enterprise platform, SRE,  
and DevOps teams. 24/7 operations. Fixed pricing.  
Guaranteed SLAs.

[Schedule a discovery call](#)

