

# SRE Report

# 2025



Catchpoint is proud to continue dedicating this independent research to advancing reliability and resilience practices with the hopes of simply making it better.

The SRE Report is built upon insights from our annual SRE survey. It delves into recurring themes such as time spent and whether varying levels of managerial responsibilities influence perceptions about reliability and resilience. This year, we also explore new research such as production pressures and whether [digital] performance has come of age. True to form, we haven't set out to provide prescriptive advice but rather to present the data as objectively as possible, empowering you to derive insights most relevant to your unique challenges.

Site Reliability Engineering (SRE) has the power to transform how organizations manage and ensure system reliability. Embracing SRE is more than a technical enhancement; it's a strategic journey toward future success. However, the journey towards organizational transformation and success must commence with introspective individual observations. Only by understanding that our work is the foundation for organizational outcomes can we align on the significant opportunities ahead and the path for achieving them.

Whether this report data confirms existing beliefs or sparks new ways of thinking, we sincerely hope you enjoy reading it.

**Mehdi Daoudi**  
CEO, Catchpoint

# Introduction

Welcome to the seventh edition  
of The SRE Report.

# Table of Contents

Introduction	2
Insight I: Slow is Officially the New Down	4
Insight II: Toil Levels Rise for First Time Ever (So Much for AI)	10
Insight III: The Danger of Unstable Organizational Priorities	16
Insight IV: Single Panes or Multiple Pains?	23
Insight V: The Learning Path to Mastery... or Misery?	30
Insight VI: It's Not a Matter of If; It's a Matter of When	39
Insight VII: Acknowledge the Gap to Fix the Gap	46
The SRE Report 2025: Until Next Year	56
Demographics and Meta	57
About Catchpoint	58

# Insight I: Slow is Officially the New Down

It's no longer just about whether services are up or down—it's about whether they perform.

# Slow is the new down

We open this year's report with an investigation into the validity of an increasingly popular catchphrase—something we often hear bandied about in corridors: “Slow is the new down.”

The phrase “Slow is the new down” means that poor performance is as bad as complete downtime or unavailability. It illustrates an evolution in performance mindset, highlighting its importance as a critical dimension, beyond just uptime. Therefore, we set out to determine whether this expression holds true in practice. The majority (53%) of organizations agree with this expression even though only 21% say they have heard it before.

What makes this piece of the research so interesting is the high frequency of selected answer choices having nothing to do with the asked question. For example, note the second most popular statement, ‘Performance should be tracked against a service level objective’ (44%).

**“Slow is the new down” is an expression meaning that having bad performance is as bad as being down or unavailable. For this expression, which of these statements are true?**

I generally agree with this expression

53%

Performance should be tracked against a service level objective

44%

Performance should be measured as an indicator

35%

I have heard this expression, or some form of it, before

21%

I generally do not agree with this expression

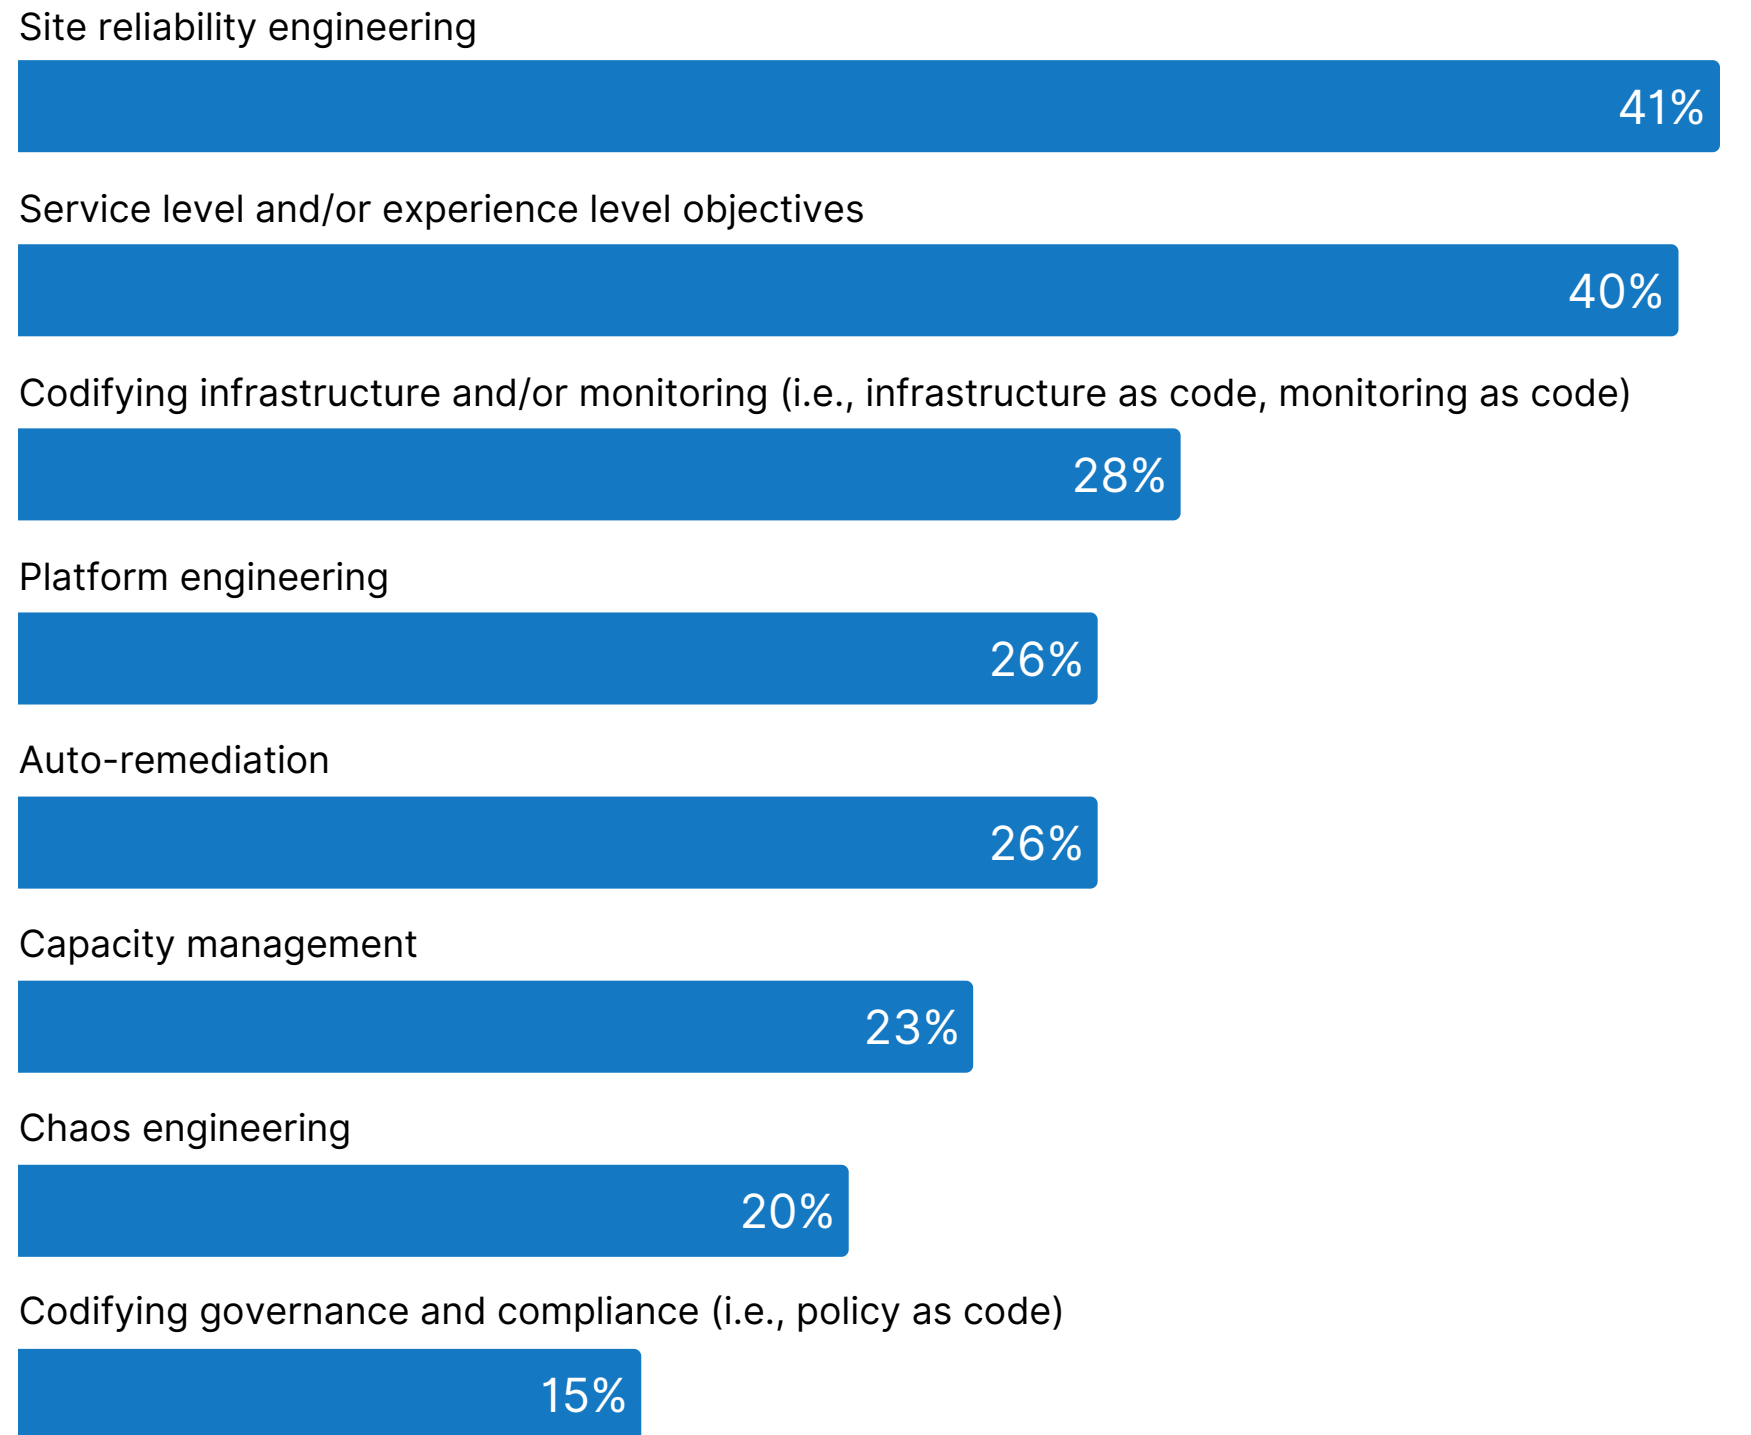
19%

# Prioritizing Service Level and Experience Level Objectives

Given the previous data on tracking performance indicators against service level objectives, we decided to put the results from this priority question here. For this report, Site reliability engineering should not surprise as the top priority result (41%). The choice of prioritizing service or experience level objectives (SLOs or XLOs) as the second highest choice (40%) only emphasizes the need to track performance indicators against objectives.

Tracking these indicators against objectives also includes the concept of [budgets](#)—predetermined allowances for acceptable errors or deviations in system performance. These budgets ensure that resources are allocated to uphold service standards, which in turn reduces the risk of performance degradation.

## Which of these should your organization prioritize to adopt over the next 12 months?





# Let's Bridge the Gap

These performance indicators should **not** be relegated to exclusively internal or first-party perspectives. Instead, also consider external and [multi-party](#) perspectives. Managing performance should not be done in a silo. Performance practices should look holistically at the overall system.

We were pleased to see the delta between all applicable attributes is smallest when it comes to performance optimization. Unfortunately, the rating for the applicable performance optimization attribute ranked lowest in terms of frequency when compared to the other applicable attributes of instrumentation or stack maps/service maps. In this vein, we highlight the opportunity before us all to increase the importance of continual performance optimization because client's digital experience expectations will also continually rise.

**For the primary application or service you work on, which of the following statements describe its front end or back end?**

Front end		Back end
<b>41%</b>	It is instrumented to monitor its performance or behavior	<b>51%</b>
<b>39%</b>	We have documentation and understanding of interconnected components, APIs, or other dependencies	<b>45%</b>
<b>30%</b>	Someone continually optimizes our performance (e.g., full page load time or core web vitals for front-end web performance, or e.g., response times or latency for back-end performance)	<b>33%</b>
<b>24%</b>	None of these describe our ends	<b>15%</b>





# View from the field

The distinction between whether a service is just incredibly slow or completely unavailable has never been more irrelevant. From an online shopper abandoning a full cart to a consensus algorithm declaring a partition after a timeout, the impact is the same.

This matters in how we build services:

- Decoupling synchronous and asynchronous components so the latter does not block the former
- Graceful degradation when a lower level of service is better than none

This matters in how we handle data:

- Measures to speed up access like precomputation, caching and indexing
- A focus on IO at every layer from disk through the memory and CPU through to the network

This also matters in how we operate the services. It's been a long time since binary pass/fail monitoring was state-of-the-art, today it has been almost completely replaced with performance focused observability.

The end results are complex, distributed systems that don't just get the work done. Rather they can get it done as fast as possible while providing the level of service expected by the people and other systems that interact with them.

It is in this environment that awareness of, and use of, SLOs/XLOs has grown. They have gained traction, not because they are a "new hotness fad", but because they capture in a nuanced fashion what it means to provide quality of service. They help a business define what it means to be performant and track how well they meet those goals. Uptime is no longer a meaningful measure of success, performance is the current gold standard. Slow is the new down.

**Martin Barry**

Team Lead, Network Operations

# **Insight II: Toil Levels Rise for First Time Ever (So Much for AI)**

For most teams, it seems the burden of operational tasks has grown.

# Toil through the years

The survey used to generate insights for this report has employed a consistent methodology for the last seven years. This ensures the data remains trustworthy and comparable, leading to provocative conversations about industry trends. When toil levels rise for **the first time in five years** (according to this report), it certainly calls for a closer look.

In previous years, we theorized that the drop in toil was due to widespread work-from-home policies, with an expected rise as offices reopened. However, despite ongoing efforts to bring employees back, the predicted large-scale return to the office has not materialized.

Around what percent of your work, on average, is toil?

	p25	p50 (median)	p75
2025	10%	20%	30%
2024	7%	14%	26%
2023	10%	20%	30%
2021	15%	25%	40%
2020	25%	40%	60%

# So what could be behind this increase in toil?

Google defines toil as "the kind of work tied to running a production service that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows." Hence, the need to "automate all the things" is a key SRE tenet.

It's hard to ignore the AI-shaped elephant in the room—the major shift in how work has evolved over the last five years. The general expectation was that AI would reduce toil, not exacerbate it. The jury is still out on whether AI is adding toilsome activities in unexpected ways. However, the [2024 DORA Report](#) suggests a more nuanced reality may be at play, stating, "AI does not steal value from respondents' work; it expedites its realization." Paradoxically, the free time created by expediting valuable activities may end up being filled with toilsome tasks. This is at least one hypothesis for the rise in toil levels observed in this SRE report. As ever, this will be another key area of focus in next year's report.

# Operations' nefarious relationship with toil

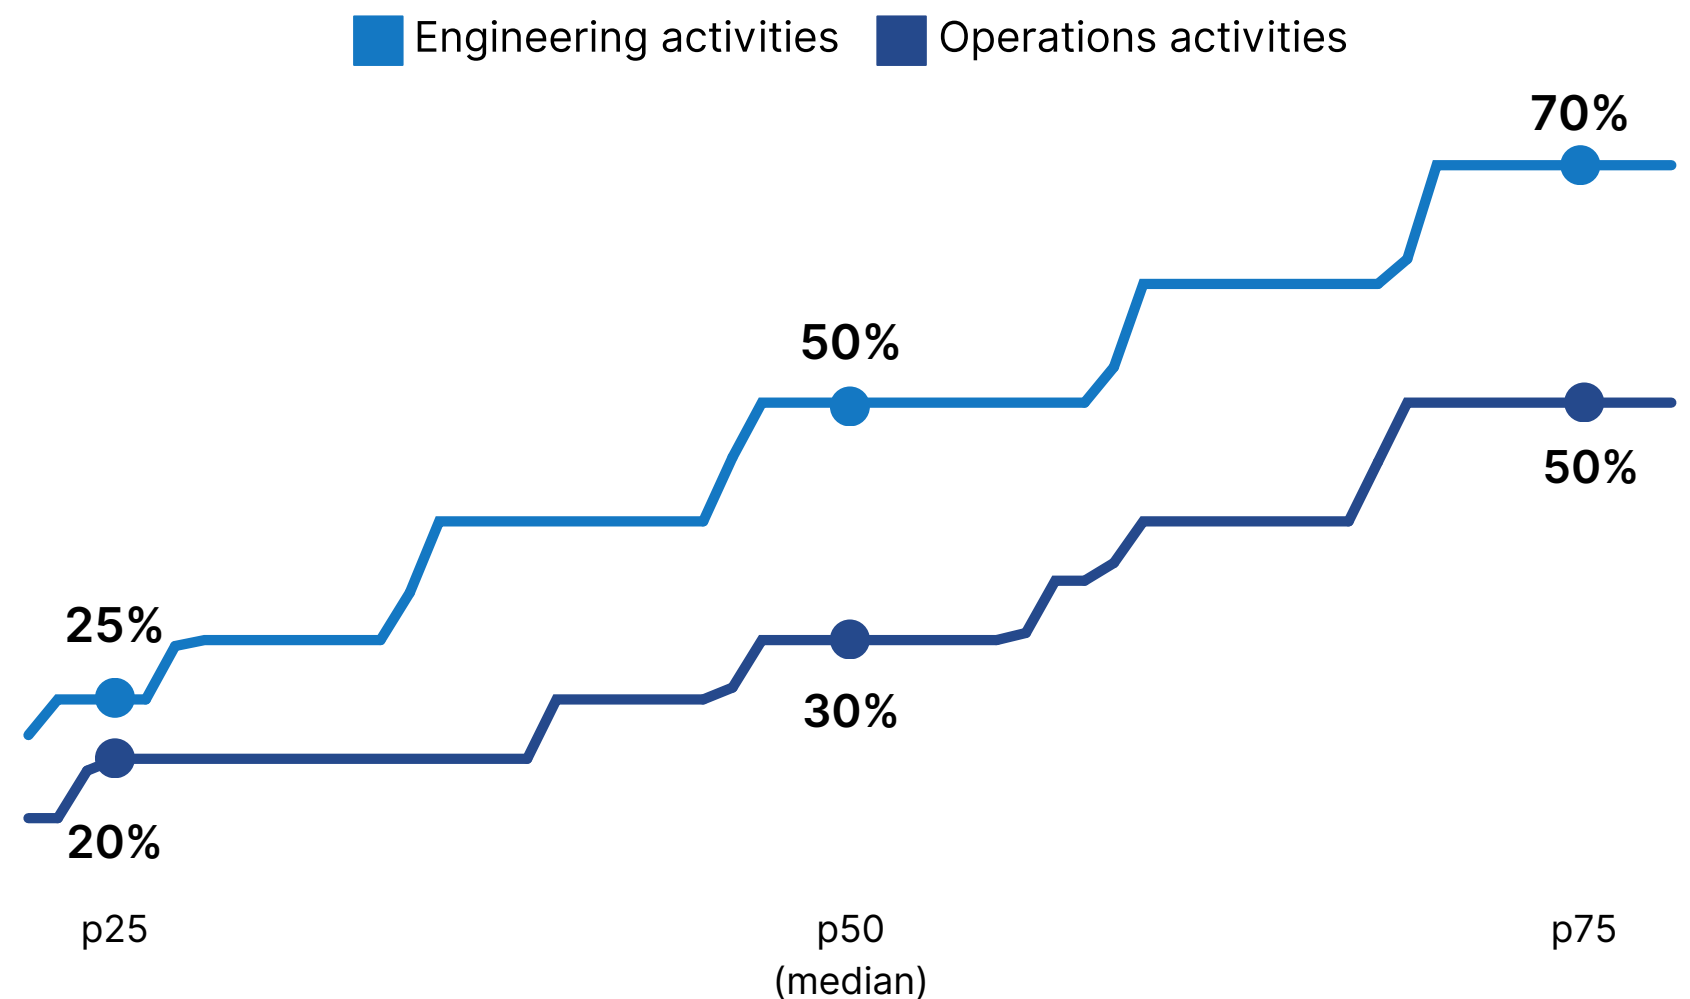
There are two critical takeaways on this page.

First, these time-spent questions are part of our recurring research. They are offered as industry indexes for organizations to benchmark against. For example, they may be considered against other data in this report, other research, or a given organization's own indicators.

Second, the p25, median, and p75 values for engineering activities are uncannily identical to those in [The SRE Report 2024](#). Yet, those values for operations activities have risen.

Last year, the median value for operations activities was 25%, but this year it has risen to 30%. For most teams, it seems the burden of operational tasks—often synonymous with toil—has grown, encroaching on time that could otherwise be invested in proactive engineering efforts. This uptick in operational load signals a potential red flag, where organizations may find themselves weighed down by routine, repetitive tasks, thereby limiting their capacity for innovation and strategic development.

**In a typical week, when you are not on call, around what percent of your time is spent on these activities?**




# Time is invested, not spent

The consistent survey methodology we use year after year helps us draw meaningful comparisons. We were curious if the increase in toil and ops work trended with time on-call. It turns out it doesn't. Much like last year's results, on-call rotation practices have remained fairly consistent.

And remember, time in these rotations isn't simply spent—it's invested. By effectively handling issues during on-call rotations, teams are investing in learning and development, building resilience, and proactive optimization.

## Do you participate in any on-call rotation?





Our systems are under constant change, and we as engineers are continually finding new ways to add toil to our day to day lives.

“Paying down” toil is a process that requires constant investment; as the industry has focused on immediate cost savings, long-term work that takes time to pay off in saved time and toil can easily be deprioritized. Engineering teams with reduced funding and under pressure to deliver product features are less able to find time to improve their operations practices for the long term.

We should also address the AI elephant in the room. AI systems are themselves a new source of operations we as an industry have yet to master: maintaining and updating models and running massive GPU clusters are both new problems for most teams. For teams not running those AI systems, AI proponents are keen to tell us that its rollout will reduce toil, but the evidence may suggest that AI is actually a source of increased toil. Manual supervision of AI systems that are mostly right, or make subtle and hard-to-predict errors, can easily raise the operational load of a team for both day to day work and incidents. We all know that a co-worker you can’t trust is a constant source of extra work... and AI is at best “a co-worker you can’t trust”. Teams and companies working to push AI adoption should look clearly at how much time and effort is being spent on supervising machine output.

**Laura de Vesine**  
Senior Staff Engineer

# View from the field

# Insight III: The Danger of Unstable Organizational Priorities

The higher the production pressure to perform, the less stable an organization's priorities appear to be.

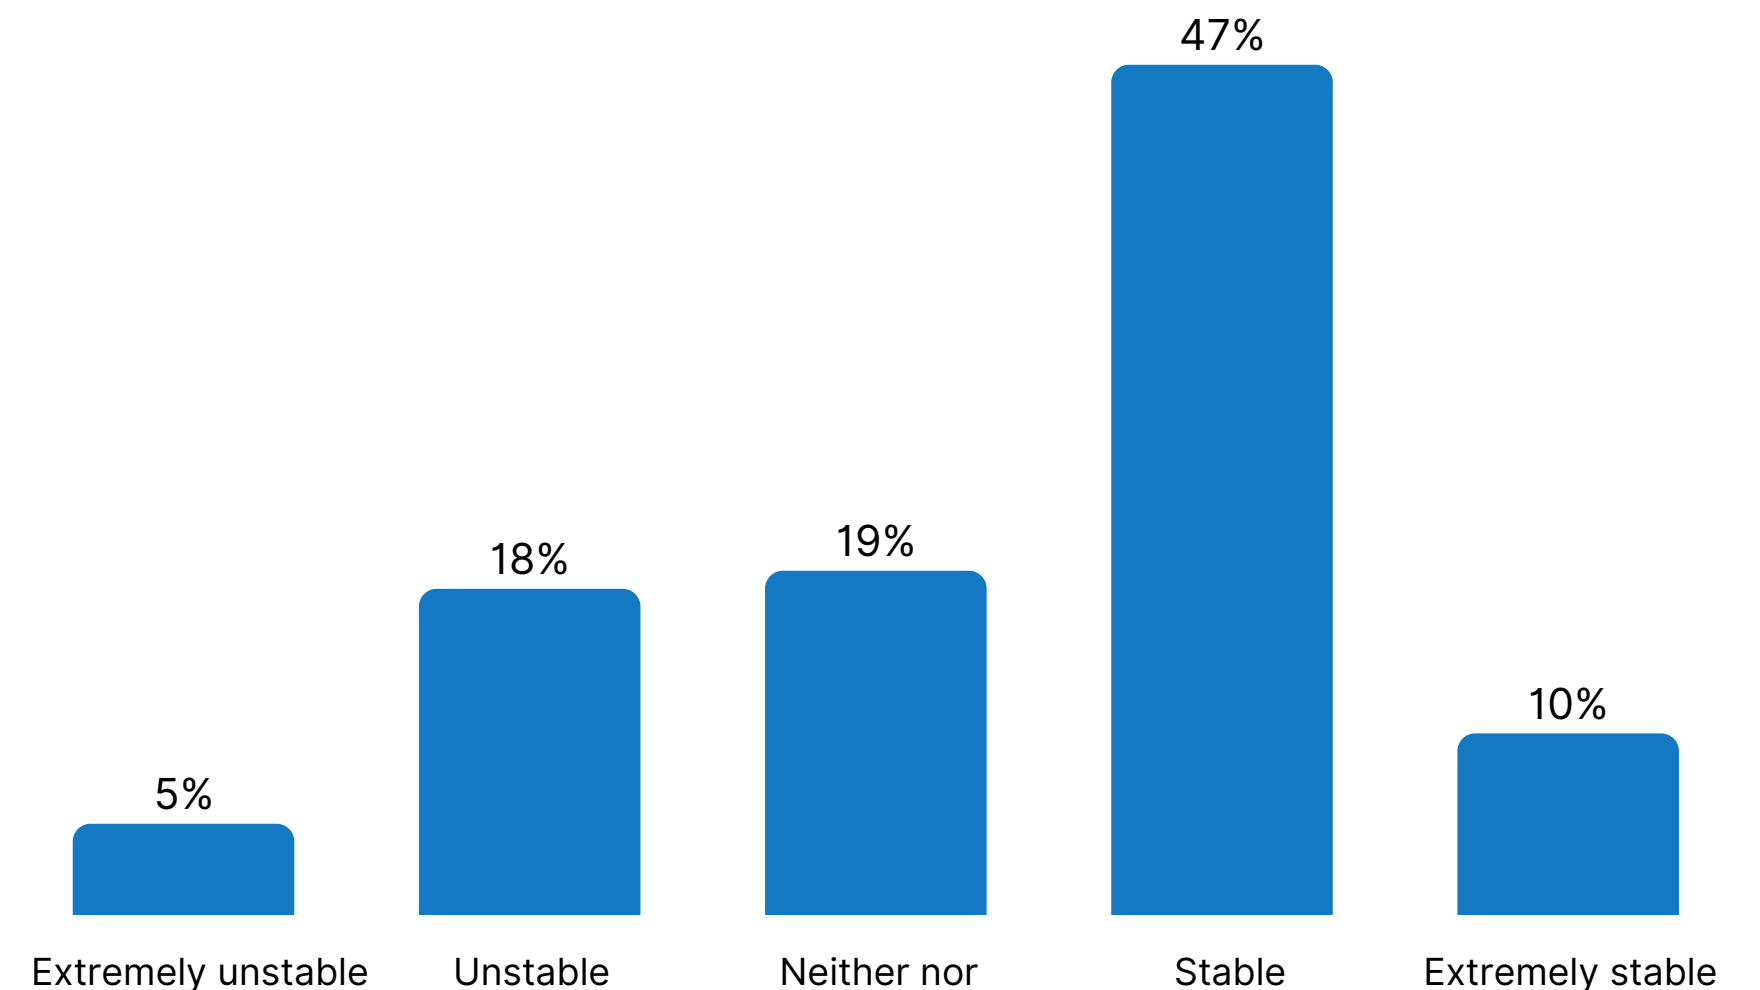


# Setting the stage: The initial illusion of stability

In this section of the report, **we explore the perceived stability of organizational priorities** against whether there is production pressure to perform. We also discuss the use of objectives and key results (OKRs) and whether reliability practitioners feel their problems and challenges are addressed.

Stability is the predominant experience for most teams (57% with a stable sentiment). This is encouraging, as it creates an environment where teams can better allocate resources, plan ahead, and minimize the disruptions that can come from shifting priorities.

How stable or unstable are the priorities of your organization?




# The need to communicate objectives and key results

The majority of organizations believe OKRs are clearly communicated (58% agree sentiment). Communicating OKRs is crucial for aligning team efforts with business goals, fostering transparency, and driving accountability. Clear OKRs ensure everyone understands priorities and enhances reliability practices. However, striving to meet OKRs can reveal challenges such as resource constraints or misaligned objectives. Addressing these issues promptly is essential to maintain morale, optimize performance, and ensure strategic alignment.

By ensuring two-way communications between the business and reliability practitioners, strategies can adapt, resources can be properly allocated, and both agility with stability can be had. This proactive approach not only mitigates risks but also reinforces a culture of continuous improvement and resilience. After all, change is inevitable in today's highly competitive marketplace.

**Business objectives and key business results are clearly communicated through your entire organization. Do you agree or disagree?**




# Address problems through culture

The majority of organizations believe reliability challenges and problems are addressed (53% agree sentiment). This cannot happen without a supporting culture.

A culture of transparency is vital for addressing reliability challenges. When teams openly share information, issues are identified and resolved more quickly, preventing minor problems from escalating. Transparency fosters trust and collaboration, encouraging team members to communicate openly about potential risks and failures. This approach ensures that reliability practices are continuously improved, as everyone is aware of the current state and can contribute to solutions. Additionally, transparent environments promote accountability, as individuals understand the impact of their actions on overall reliability. Ultimately, this culture leads to more robust and resilient systems, enhancing overall performance and reliability.

The challenges and problems of product reliability are addressed by the organization. Do you agree or disagree?



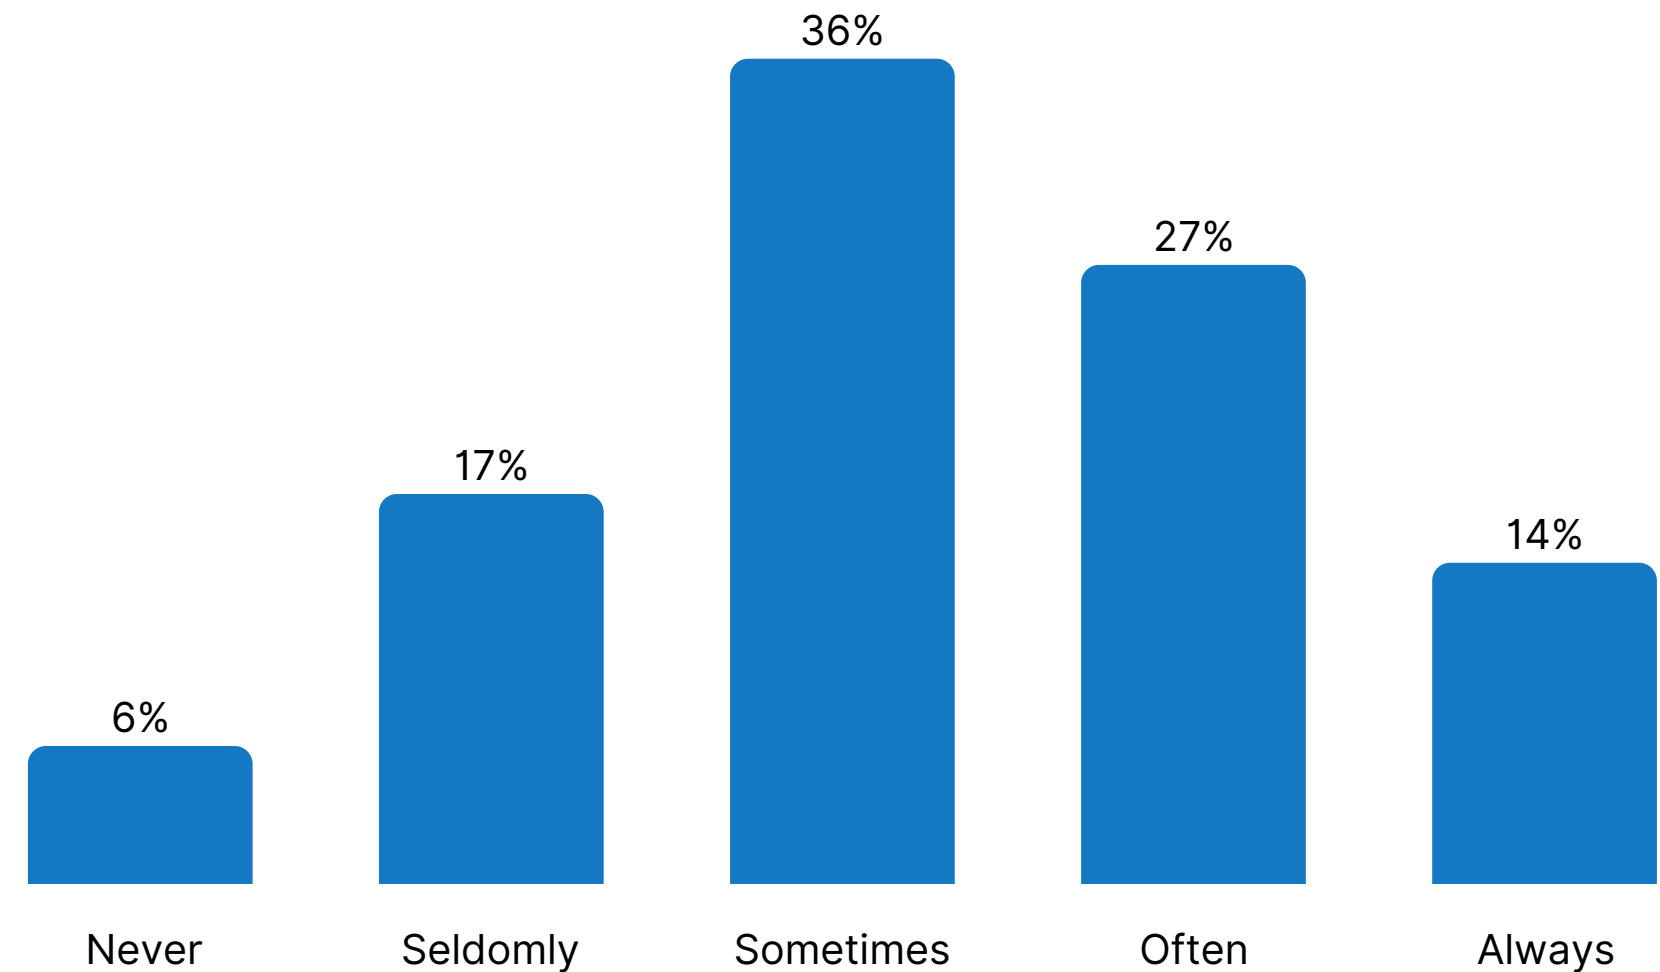
# Where it gets complicated

This is where the story gets a little complicated. Despite the majority of organizations believing that OKRs are clearly communicated, reliability challenges are being addressed, and stability the predominant posture for most teams, the majority still indicated feeling pressured to prioritize release schedules over reliability.

It's a classic case of agility versus stability: businesses want updates, new features, and revenue growth, whereas practitioners prioritize reliability and resilience. More than half of organizations are frequently caught in a tug-of-war between meeting tight deadlines and maintaining reliability. The tension is evident in these results.

The high number of respondents choosing "often" or "always" (41%) is especially concerning. It underscores the challenge of balancing speed with stability in fast-paced environments.

**Do you feel pressured to place release schedules or deadlines as the priority - instead of reliability as the priority?**



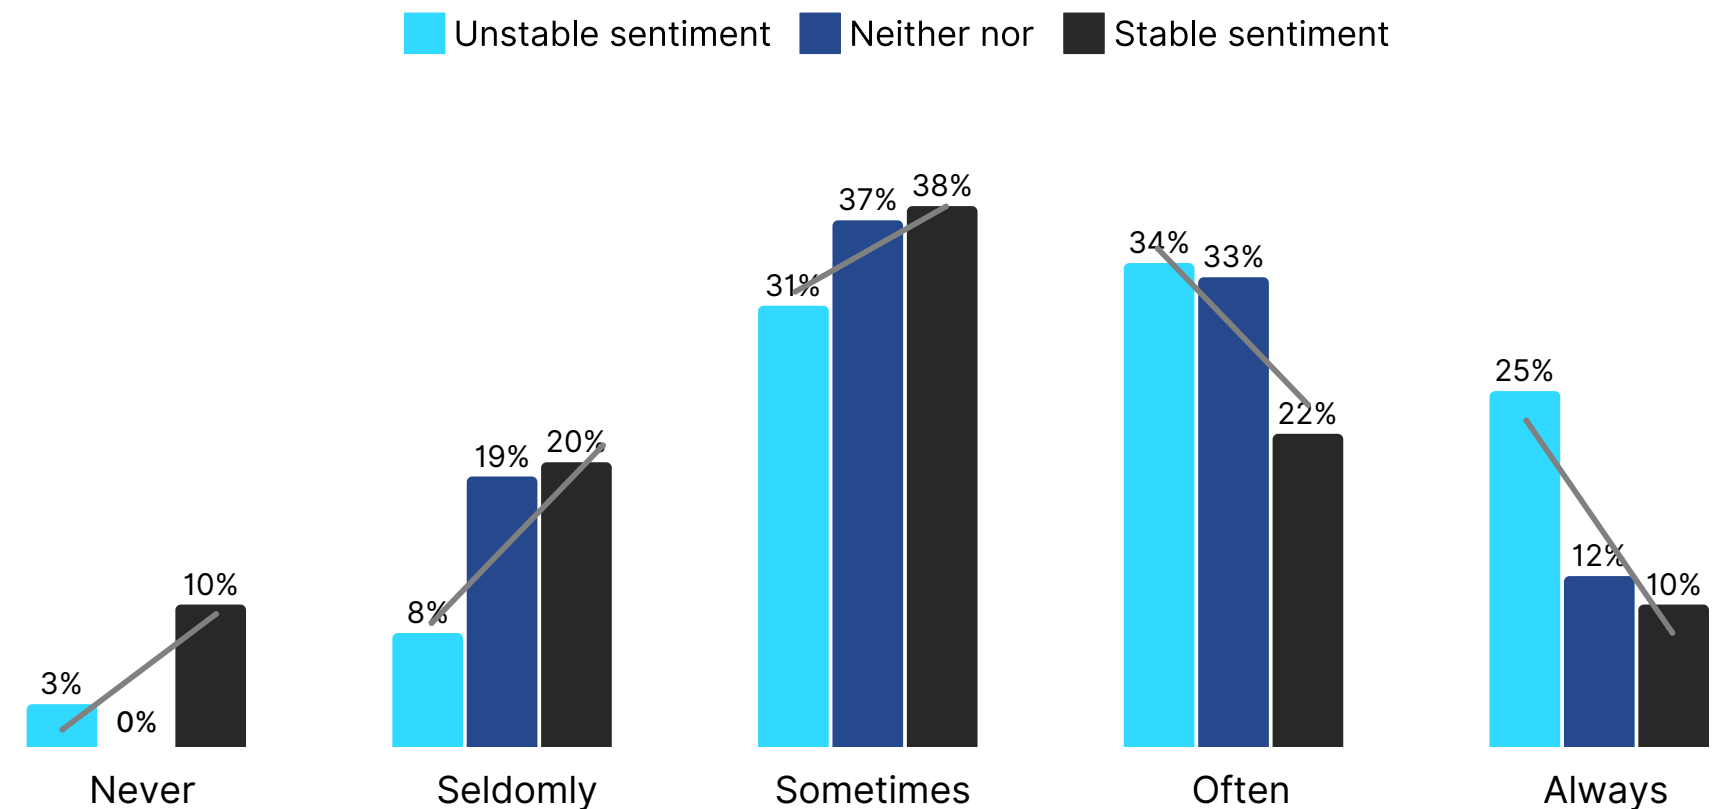
# Have resilient capabilities to manage instability

When production pressure to perform is more frequent, the less stable an organization's priorities appear to be. In other words, the agility versus stability war rages on.

If we accept business priorities will inevitably change, then priorities are stable until they aren't. Evaluate whether existing capabilities can be reused to account for new or different priorities. If they cannot be reused, and new capabilities are needed, this is a wonderful opportunity for reliability practitioners to build resilience and reusability into them. That is, the adaptive capacity of capabilities themselves should also be considered when discussing priorities.

## Do you feel pressured to place release schedules or deadlines as the priority - instead of reliability as the priority?

(by whether organization priorities are stable)



*Note the negative relationship trend of unstable organizational priorities was seen - to a lesser extent - when investigated across OKR communication and reliability problems being addressed by the organization. In this vein, we reinforce the idea of ensuring reusable, adaptable capabilities even if organizational priorities change.*



The “Classic Battle” between features and reliability rages on. If we’re honest, this is likely a struggle that isn’t going away in any organisation without a very well-developed reliability attitude capable of surviving a crisis (spoiler: this is almost all organisations!).

We’re seeing a number of trends that affect this ongoing tension:

- **Lore is Mobile** - The taste-makers who initially either sponsored or created an SRE group are increasingly likely to have moved on, for various - possibly unfortunate - reasons. This leaves org-wide priorities in flux, and it’s entirely possible for SRE groups to not be aware that the fundamental attitude to Reliability in their organisation has either shifted - or could shift - at any moment.
- **Investments are more Principled** - Product organisations are more likely to have developed a calm, holistic picture of their need for investment in Reliability. This is a good thing! However, there does tend to be less of an appetite to do the hard work of adjusting the organisation’s capabilities to suit. This can leave folks on either (or both!) sides of the features/reliability aisle frustrated that they seem to care a lot more about their domain than the organisation does.

It’s good to have an organisation that can adapt to shifting priorities, as long as this is the exception. The baseline as we move into a leaner, more principled, and more evolved attitude to Reliability is to be closer to the needs of the business; step one toward that is knowing what those needs are, and making (potentially difficult) adjustments for a longer-term attitude that makes sense. Increasingly, the choice is between doing this yourself, or having it done to you.

**Dave O'Connor**

Google SRE 2004-2021, Reliability Consultant

# View from the field

# Insight IV: Single Panes or Multiple Pains?

Most organizations are using multiple monitoring tools, and that's okay.

# Is there really a tool sprawl problem?

All of us, not just SREs, are under pressure to cut costs. However, this drive to reduce spending should never come at the risk of losing value.

We've explored this topic in multiple SRE Reports. Though the questions may have been expressed differently, they were meant to research the same idea: Is there a tool sprawl problem? This year, we decided to ask yet another form of this question: Is the received value from your monitoring/observability tooling greater than the cost?

Cost takes many forms, e.g., in the form of hard dollars or in the form of time spent implementing and maintaining. Either way, it's crucial to understand that tool sprawl isn't simply about "How many tools are in the stack?" Rather, if the received value from a toolchain is net positive when compared to cost, then there is no tool sprawl problem. In fact, organizations answered with the same directional value trend regardless of how many tools they have in use.

**The received value from our monitoring/observability tooling is greater than their cost. Do you agree or disagree?**

Strongly disagree



Disagree



Neither nor



Agree



Strongly agree






# Multiple tools, multiple needs

Most organizations use between 2 - 10 monitoring or observability tools, with little change from last year. With so many organizations relying on multiple tools, we suggest self-evident value in doing so. It's unrealistic to think that everything can be effectively monitored with just one tool. Different technology stacks, such as application versus Internet Stacks, require different tools.

Many tools are designed to address specific scenarios. Practitioners may need to rely on a variety of different tools from either single or multiple vendors for comprehensive coverage - and this is okay.

Rather than fixating on the idea that 'there are too many tools in the stack,' the focus should shift to evaluating if the value gained from these tools justifies their total cost of ownership (or operation)—understanding that cost can take various forms.

## How many monitoring or observability tools does your organization use?




# Mind the observability gap

Observability is undeniably critical, but are we truly getting the right level from our tools? The majority of organizations believe their level of observability falls short (51% 'less than' sentiment). Interestingly, 16% with a 'more than' sentiment suggests another side of the challenge.

Striking the right balance in observability is essential for effective system monitoring. Too little observability can leave critical issues undetected, leading to prolonged downtime and difficult troubleshooting. On the other hand, too much observability can flood teams with more noise than signal, making it hard to surface and explore meaningful insights.

The key is to implement an approach focusing on capturing relevant telemetry that provides actionable insights. In this context, it's worth noting the data we previously mentioned: 41% of organizations said their front end is instrumented, while 51% said their back end is instrumented. We mention this again because conversations about relevance and value - as compared to cost - will most undoubtedly include digital experiences and business outcomes. It is in that vein we stress the importance of instrumentation to monitor performance or behavior of both front and back end, and having the ability to tie those metrics to what's important to the business.

## What level of observability has been instrumented in your organization?

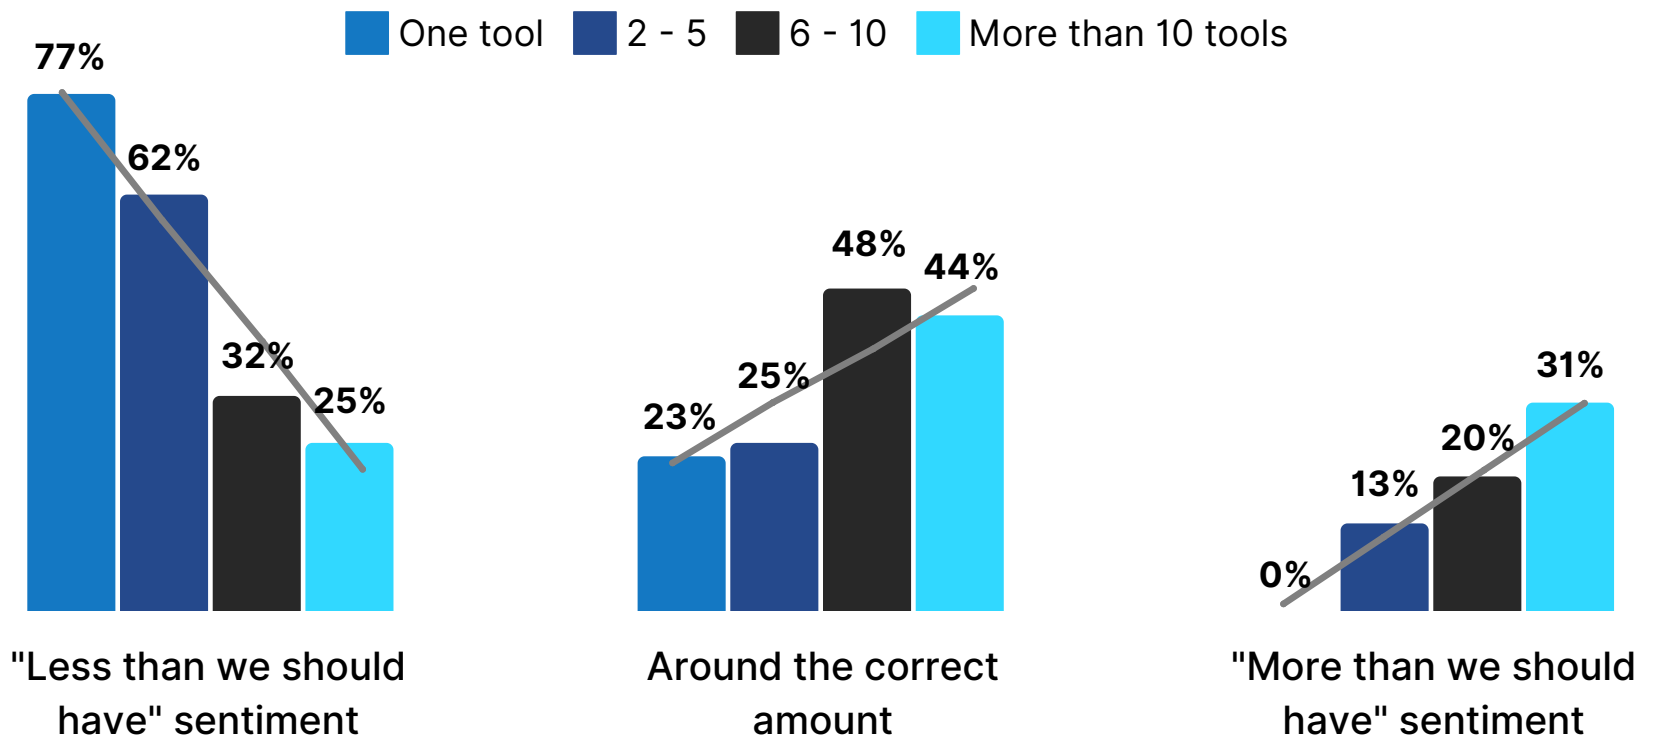


# The allure of one pane to observe them all

Consolidating tools might seem like a solution to manage complexity, but this data shows that less tools trend with less perceived observability instrumentation. If observability truly is a critical business function - not just a “nice to have” - then be leery of trying to consolidate or minify tools only for the sake of, well, trying to consolidate or minify tools. One pane of glass isn’t enough.

Organizations using ‘more’ tools trended with much better outcomes in terms of feeling they have the “correct amount” of observability. But remember, it’s not just simply a count of the tools in your profit and loss statement. It all comes down to received value. If the value outweighs the cost, then by all means, go and get the tool for the thing you need. A best-of-breed approach allows you to cover different stacks—whether it’s application or Internet—each requiring specialized visibility.

What level of observability has been instrumented in your organization? (by number of monitoring tools)




# The building blocks of observability

We close this section with a breakdown of perceived value for observability components: logs, metrics, events, and traces. The good news is the value for all listed components has 'high' as the majority vote. Also keep in mind the question was asked, 'How much value do these provide to **you**?' This contrasts with, for example, if we had asked, 'How much value do these provide to your organization?'

Combining insights from logs, metrics, and traces allows for a more holistic view of the system, enabling better decision-making and more effective problem resolution. So we felt it important enough to present this breakdown and discuss it as a 'greater than the sum of their parts' mindset. In the context of the 'cost versus value' research, we suggest this will be part of the conversation when purchasing (or building) the correct tooling.

## How much value do these provide to you?





# View from the field

65M USD. Now, that's a telemetry bill! ... even for a company with \$3 billion in revenue.

But is it too much?

The answer, in Coinbase's case, was Yes. In 2023, they successfully renegotiated, and got a much better deal. But the broader question remains: How much is too much for telemetry? How much value are we getting from these tools? How does this translate to Business impact?

Quantification of the value of observability is hard since it is highly multi-dimensional. Bottom-line contributions by increased reliability ("Reduce MTTR") are only part of the picture. Here are a few things to consider:

- Telemetry is a Tier-0 system. If Telemetry is down you are flying blind.
- Telemetry tools are relied on by every engineer who is deploying to production in the most stressful situations.
- Telemetry vendors are hard to change. Migrations take 6 months+ and require effort from all teams.

You are not just buying a telemetry product, you are engaging in a long term relationship with a trusted partner to deliver a central piece of your developer platform to your internal audience. Polished UI, good documentation and support directly translate into productivity of your developers.

What to do now? Here is some tactical advice: Compare total-cost-of-ownership for Telemetry against your current cloud provider's bill. If it's more than 20%, it's definitely time to re-evaluate. Watch out for pay-per-query costs, this can lead to surprises! Don't send junk data - if you're paying for volume. When vendor costs are a major concern consider best-of-breed solutions: They can give you lower cost and better service at the cost of increased integration efforts and reduced polish of developer experience.

Happy monitoring!

**Heinrich Hartmann**  
Principal SRE

# Insight V: The Learning Path to Mastery... or Misery?

The desire for technical training on artificial intelligence (AI) is universal.

# Training widely supported, but not universally so

Technical training should be a key investment area for any organization to increase velocity and modernize implementations.

After all, new technologies like AI aren't going to implement themselves (not yet anyway). Employees need to be supported with the right training to meet evolving business objectives, but as the results show, support for technical training is not universal. One in five organizations is not making this investment, and the fact that 8% of respondents indicated "I do not know" regarding training support is concerning, pointing to a potential gap in communication or awareness about available resources. Training isn't just about keeping up; it's about staying ahead, ensuring the workforce is equipped to fully leverage new tools and frameworks, from traditional tech skills to newer AI capabilities.

## Does your organization provide support for paid technical learning?

Yes, I can spend some amount of dollars on third-party training (in-person, Coursera, etc.)



Yes, my organization has professional training staff



No



Yes, my organization brings in training specialists



I do not know



# Learning preferences

In a digital-first, Internet-centric landscape, we are not surprised that online training was the most popular option (55%) for valuable learning sources. We are also pleased that onsite, in-person training also ranked high on the list (45%).

We again stress that technical training should be a cornerstone for any organization. There will be many contributing factors shaping what that looks like, though. For example, on the next page, levels of managerial responsibility trend differently for which learning sources are most valuable. Before continuing, we also wanted to give an honorable mention to [some form of] labs or laboratories as a valuable learning source. This question had an open-ended option, and labs was a popular write-in.

## Which of these technical learning sources are the most valuable to you?

Purpose-built online training sites (Coursera, Khan Academy, etc)

55%

Onsite, in-person training

45%

Conferences

41%

Peers/colleagues

33%

Blogs (Medium, Substack, Blogger, etc)

31%

Social media (X, LinkedIn, etc)


15%



# Balancing learning in hybrid times

The skew between individual contributors and management learning preferences reveals a potential disconnect. While management favors onsite, in-person training, likely to encourage collaboration and maintain in-person culture, individual contributors clearly lean toward online training formats that offer flexibility. This disconnect might highlight differing values—management seeking control and face-to-face interaction, whereas individual contributors prioritize autonomy and adaptability. Or it could be that managers are looking to enhance their “people skills” which might be more effectively done with an in-person format.

Which of these technical learning sources are the most valuable to you? (by rank)




# Time-starved for learning


While technical training is universally seen as critical, most people expressed that they simply don't have the time. The gap between intent and action here speaks to the pressure organizations feel to prioritize other tasks over upskilling—especially considering that reliability engineering is, at its core, a deeply technical undertaking. When learning is deprioritized, the ability to adapt to new tools and methodologies suffers, which will become evident as we explore the role of AI.

Regarding time spent on technical learning, there is a differing trend between ranks. This is qualitative, though, as we theorize the underlying dimension is simply managers (1 mgmt level) and directors (2 mgmt levels) may feel like they need less tech training to begin with. So we are not surprised to see this.

## How much time do you spend on technical learning?



-- by rank --



# AI and the training paradox

Training is important...but on what? We asked about AI, but the sentiment around it remains varied. People are understandably hesitant about diving into AI implementations. Notably, the desire for technical training was the second most selected, following only the cautionary sentiment. While AI is undeniably the hot shiny thing right now, are there other crucial areas that warrant your attention? It remains to be seen whether training on AI is the path to mastery or misery.

## How do you personally feel about investing in, or implementing, AI technologies?

I would approach AI implementations or investments with caution

37%

I would like technical training to better know how to use them effectively

30%

I am excited or eager to implement or invest in them

24%

I would consider them only if our organization sets guidelines on their use

23%

They will increase our workforce productivity

22%

I feel pressured to implement or invest in them

16%

To be honest, I have not really thought about them

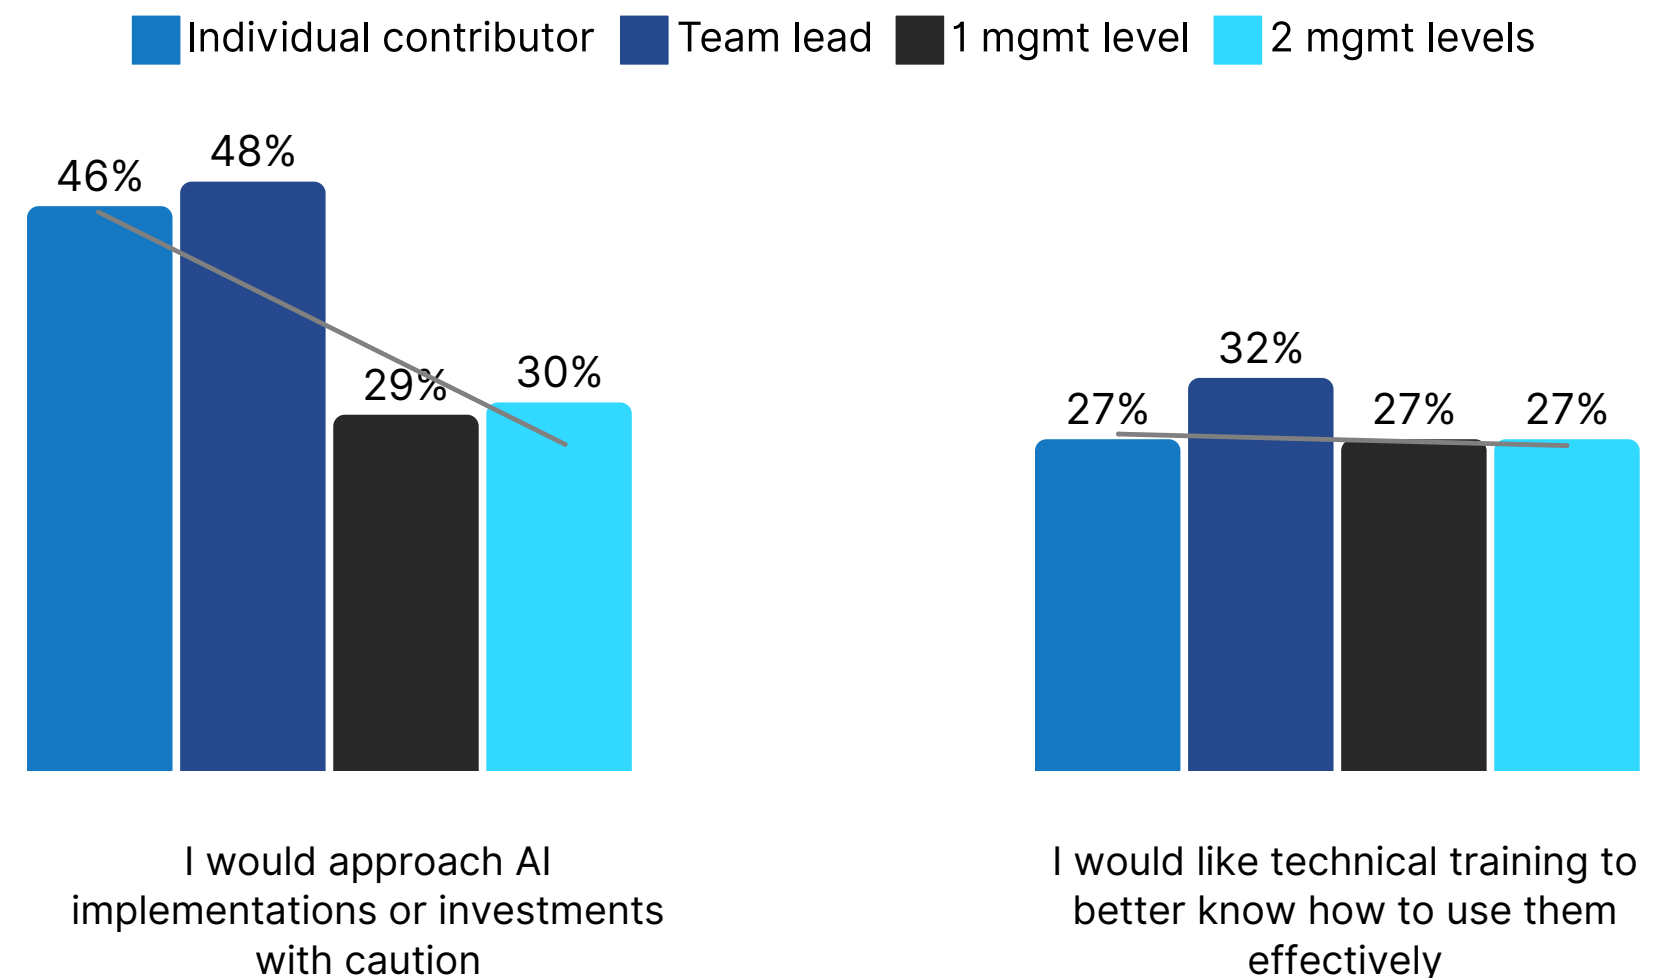
12%

# The desire for AI tech training is universal

Analyzing the top two choices from the AI implementations sentiment question by rank reveals alignment on the desire for technical training. That is, the trend line is nearly flat across ranks. We refer to these rare alignments by rank as 'universal opportunities', and suggest organizations take note.

However, when it comes to the caution sentiment, the skew between individual contributors (46%) and management levels (30%) is notable. Individual contributors seem to show greater caution, and this differing opinion matches the general pattern of previous SRE Reports.

## How do you personally feel about investing in, or implementing, AI technologies? (by rank)




## How useful will AI be in the following activities within the next 12 months?

None
  Low
  Moderate
  High

### Writing Code



### Incident Management



### Web Performance Optimization



### Infrastructure as code



### Service level management



### Capacity Management



### Release Management



# AI use cases

Similar to the [2024 SRE Report](#), 'Writing Code' saw the highest positive sentiment regarding AI's usefulness (39%). Writing code also emerged as the most common AI use case in the [2024 DORA Report](#). The optimism around AI for coding is understandable, given the universal need to "write more code". What remains less clear is whether this reliance will ultimately lead to significant software delivery or quality improvements.

Conversely, 'Release Management' garnered the lowest positive sentiment, at 17%, down from 27% last year.



# View from the field

Engineers create stuff; that's basically ingrained in our character. Learning, in contrast, is often a series of passive tasks where you consume content created by others. To be better learners then, we should keep in mind the ways to make learning a creative process too.

- Don't just read about a new concept or technology, have a lunch session where you discuss it with your teammates.
- Don't just attend a conference, try to submit a talk. Even if you feel that you're at the early stages of your learning journey, conferences are often open to "the view from the field" sessions.
- If you've completed a lab, see if you can adjust the examples to make them more relevant to your organization's environment - and then share your ideas with your colleagues so that the training will be even better for them.
- Start a blog series or post about your learning and your ideas.

Of course, some of these suggestions are easier said than done, especially for junior engineers. That's where teamwork comes in - we never have to do these things alone. Discussing how to convert your passive learning into active learning with your manager and mentors can lead to creative results at a larger scale than you can achieve on your own.

Active learning, or learning by creating & teaching, is a way to advance your career by getting noticed, helps fight the imposter syndrome which plagues so many engineers, and is a great way to both deepen your personal learning and help your colleagues - all at the same time!

Personally, I can't define exactly what makes a great engineer, but the ones I most admired were the ones who taught the most.

**Robert Barron**  
SRE & Architect

# Insight VI: It's Not a Matter of If; It's a Matter of When

Incidents are not rare, nor are they isolated.

No discussion of reliability would be complete without considering its obverse - the outage. As we saw in the first section of this year's report, it isn't just full outages that matter anymore; it's also performance degradations that cause customers to rage-click, go off to another supplier, hate their online lives, or otherwise be unhappy with your service.

So how are companies faring this year in handling incidents?

# 40%

The number of respondents who said they have responded to between 1 and 5 incidents in the last 30 days.




# Incident response: par for the course

Responding to incidents is part of the job. Understanding the frequency and impact of these events is crucial for improving resilience.

40% of respondents reporting they handled between 1-5 incidents in the past 30 days may seem manageable, but it's important to recognize the cumulative effect of incidents on both individual well-being and team effectiveness. For those dealing with 6-10 incidents per month (23%), the load becomes even more challenging—especially when compounded by other responsibilities.

Curiously, 14% of respondents reported dealing with zero incidents in the last month. This may indicate that their systems are either highly resilient or they are not in the response loop for incidents, or, perhaps, that their monitoring and alerting practices haven't picked up critical signals—a potential concern.

How many incidents have you responded to in the last 30 days?

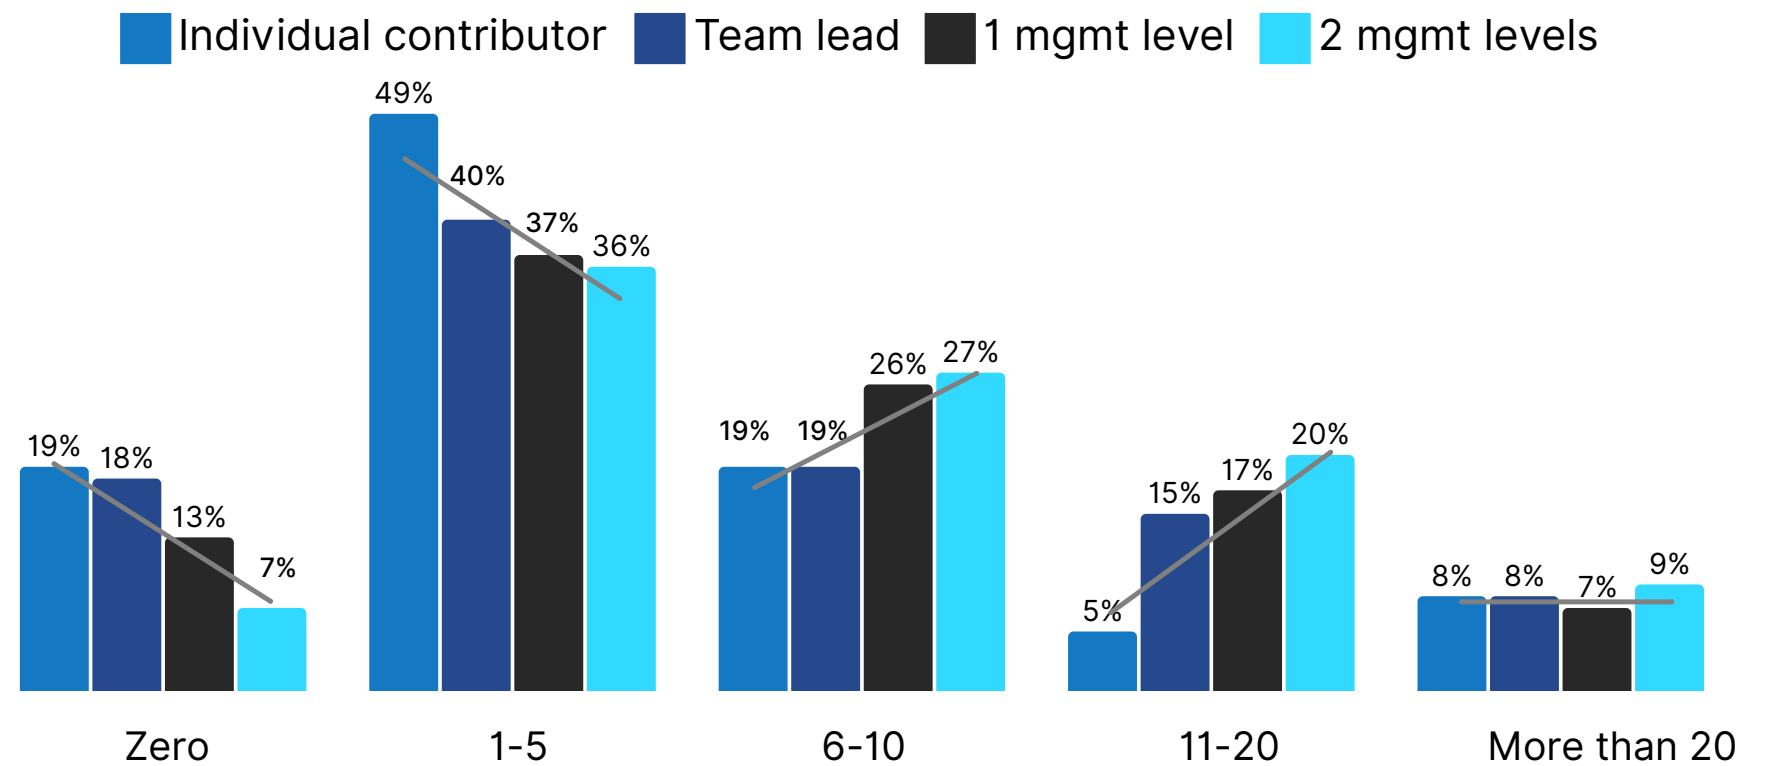


# Incident response isn't just for individual contributors

We hypothesized that being an individual contributor would translate into handling more incidents, but our findings tell a different story. Higher-level managers are just as involved, if not more, in incidents. Perhaps it's because managers get called into most calls and postmortems for each of their direct reports.

This reveals an interesting insight into the dynamics of incident response within teams: **managers aren't exempt from incidents.** In fact, they are deeply involved in both managing and understanding them. It could also indicate that **responsibility increases involvement** rather than decreasing it.

How many incidents have you responded to in the last 30 days?  
(by rank)



# Incidents don't end when they're over

Every incident, regardless of severity, contributes to the ongoing stress of maintaining reliability. The majority of respondents said they experienced higher levels of stress during incidents. However, we also looked at how many reported higher levels of stress after incidents. Fourteen percent of respondents said stress levels were higher after incidents (versus during incidents). For example, if they indicated seldomly getting stressed during incidents, then 14% selected either 'Sometimes,' 'Often,' or 'Always' for increased stress after incidents. This could indicate ongoing repercussions of incidents that were not fully resolved or perhaps the pressure of learning from failures in an environment lacking blameless post-incident practices.


## Do your stress levels increase during and/or after incidents?

■ Never ■ Seldomly ■ Sometimes ■ Often ■ Always

During incidents



After incidents



# Support during and after incidents

Support levels tend to be higher during incidents compared to afterward, matching the superficial expectation that once an incident is “resolved,” all the work is done. This highlights a potential gap in post-incident support. When your site (or feature) is obviously broken, it is easy to get people engaged to fix that obvious problem, but often the longer work to identify and address the contributing factors is much harder to explain to people and maintain attention or priority.


## How supported do you feel by your team during and/or after incidents?

■ None ■ Low ■ Moderate ■ High

During incidents



After incidents





This year, I received fewer incident pages. I'd love to say it's because there were fewer incidents, but that's not true. Incidents are part of life when running a globally distributed system with thousands of devices.

There will be incidents – encompassing factors both within and outside our control. However, we've built a resilient system. While incidents happen, they rarely impact our customers. There may be a regional outage requiring traffic to fail over to another provider, a bad deployment needing a rollback, or of course the occasional software bug that need patching.

War rooms, or as they exist today – [pick your favorite teleconferencing system] rooms – are never “fun,” but they are eye-opening. They reveal the team members who can take control of a potentially bad situation and show how committed everyone is to resolving issues. Accountability and ownership of the problem are evident as people work day and night until it's resolved.

That's why I've received fewer incident pages. We've changed the escalation tree to exclude me in most cases because the team's got it, and I have full confidence in that they do.

After the incident 'ends', my contribution 'begins'. I help make sense of what went wrong, prioritize action items, and sometimes talk to customers. We work as a team to understand what happened and minimize the chance of it happening again. So, yes, I agree: “Incidents don't end when they're over”.

**Sergey Katsev**  
VP, Engineering

# View from the field

# Insight VII: Acknowledge the Gap to Fix the Gap

This isn't a diatribe. It is a wonderful opportunity to address the IT-to-business communications gap.



Before we proceed with the final section of this year’s report, we would like to express our gratitude to our executive sponsors, the reliability community, our esteemed readers, and anyone who simply ‘wants to make it better’. They push us to pour our heart and soul into the writing of this report - to dive a bit deeper into the data to uncover true insights.

This last section is no different.

In it, we present data—without distracting text—from survey questions about the current state of reliability. The aggregate answers paint a positive picture of those practices and are presented in a table on the next page. However, when we broke down the questions by level of managerial responsibility, the sentiment trends varied—some substantially.

We present the data not to complain about the proverbial ‘they don’t understand what it takes’. Instead, this is an explicit declaration of the wonderful opportunity that lies before us all: making it better by acknowledging the IT-to-business communications gap exists.

**Leo, Kurt, and Denton**

# Authors’ Personal Note

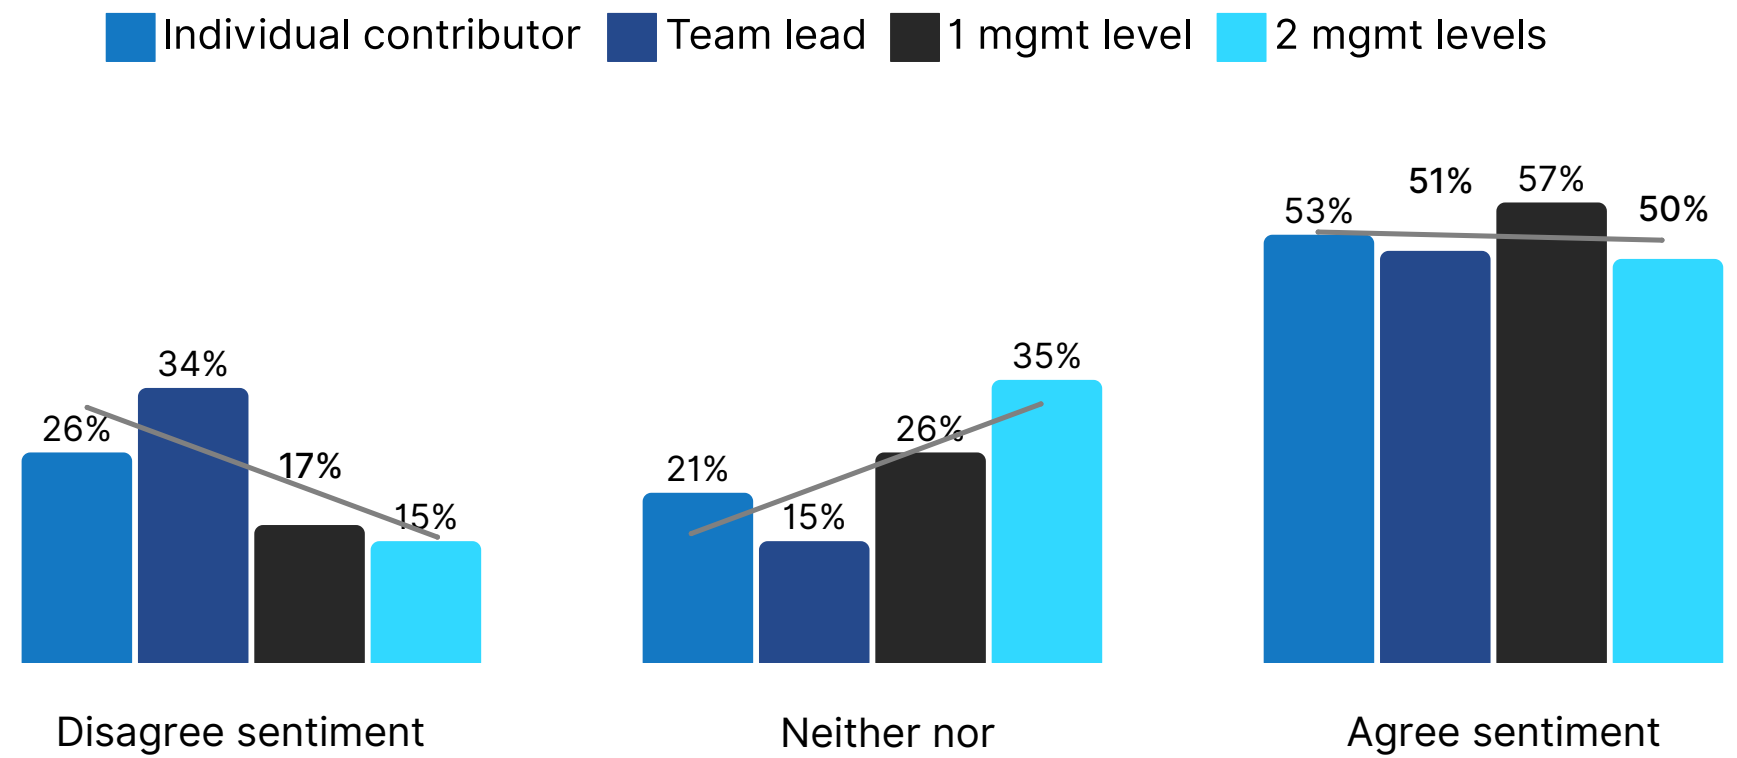
# Do you agree or disagree?

	Strongly disagree	Disagree	Neither nor	Agree	Strongly agree
For the primary application or service I work on, my team regularly reviews and revises reliability targets based on evidence.	7%	17%	22%	46%	7%
For the primary application or service I work on, when we miss our reliability targets, we perform improvement work, adjust our development work, and/or re-prioritize.	5%	14%	24%	46%	11%
For the primary application or service I work on, my team works to improve the reliability of the existing system throughout the lifetime of the product (not only during initial design or immediately after an outage).	3%	11%	18%	53%	14%
For the primary application or service I work on, we regularly test our reliability incident preparedness through simulated disruptions, failover exercises, table-top exercises, or etc.	10%	23%	21%	37%	9%
For the primary application or service I work on, my team has well-defined incident management procedures (e.g., responder roles, communications channels, playbooks, or etc).	4%	11%	20%	49%	16%
For the primary application or service I work on, we include third-party vendors or providers in our incident management procedures when applicable.	5%	12%	20%	51%	12%




**For the primary application or service I work on, my team regularly reviews and revises reliability targets based on evidence.**

**Do you agree or disagree? (by rank)**




**For the primary application or service I work on, when we miss our reliability targets, we perform improvement work, adjust our development work, and/or re-prioritize.**

**Do you agree or disagree? (by rank)**




**For the primary application or service I work on, my team works to improve the reliability of the existing system throughout the lifetime of the product (not only during initial design or immediately after an outage).**

**Do you agree or disagree? (by rank)**




**For the primary application or service I work on, we regularly test our reliability incident preparedness through simulated disruptions, failover exercises, table-top exercises, or etc.**

**Do you agree or disagree? (by rank)**




**For the primary application or service I work on, my team has well-defined incident management procedures.**

**Do you agree or disagree? (by rank)**



**For the primary application or service I work on, we include third-party vendors or providers in our incident management procedures when applicable.**

**Do you agree or disagree? (by rank)**





# View from the field

Aligning on priorities can be challenging when there are perceived gaps in the current state of ‘What do our current reliability and resilience practices look like?’ These gaps between work-as-imagined and work-as-done may lead to misunderstandings and miscommunications (if communications even happen at all), as stakeholders may have different perspectives on what is most important.

Without a clear, shared understanding of the current situation, it becomes difficult to set common goals and agree on the steps needed to achieve them. This misalignment can result in wasted resources, duplicated efforts, and missed opportunities. However, fear not. This isn’t a diatribe. Instead, it is an explicit declaration of the wonderful opportunity for advancing reliability-as-a-feature practices by acknowledging the existence of these gaps. To bridge them, it’s crucial to establish transparent communication channels, regularly perform and update assessments, and ensure all parties are informed and engaged in the decision-making process.

To this end, we encourage readers to conduct their own research on their current state. This research should include understanding current - or even non-existent - capabilities and how they map to achieving business outcomes. These capabilities are like gears in a clock. If even one gear is out of sync, the clock won’t function correctly, and you’ll end up with the wrong time (that is, business outcome). This is the situation we want to avoid, and we hope this provocative data set helps to start or catalyze independent research within the context of your own organization.

**Kurt Andersen**  
Software Architect



# The SRE Report 2025: Until Next Year

**“We can’t do it without you.  
Looking forward to next year’s  
research.”**

I was recently having a conversation with a VP of Software Engineering. During our banter, he said he was about to do a presentation on the importance and evolution of monitoring, and was going to use some of the data from last year’s report.

It is because of examples like this that the SRE Report will always hold a special place in my heart. I started writing it (along with [Kurt Andersen](#) and others) in the first year of COVID, even though Catchpoint had started this annual contribution a couple of years before.

But it’s not only about me. It is about all of our readers who give us feedback and suggestions. It is about our contributors who take time out of their day job to offer their views from the field. It is about all practitioners, leaders, or evangelists who share their stories and are sometimes the inspiration for what we choose to research.

For this, I, and the entire report production team, extend our deepest gratitude.

**Thank you.**

We can’t do it without you. Looking forward to next year’s research.

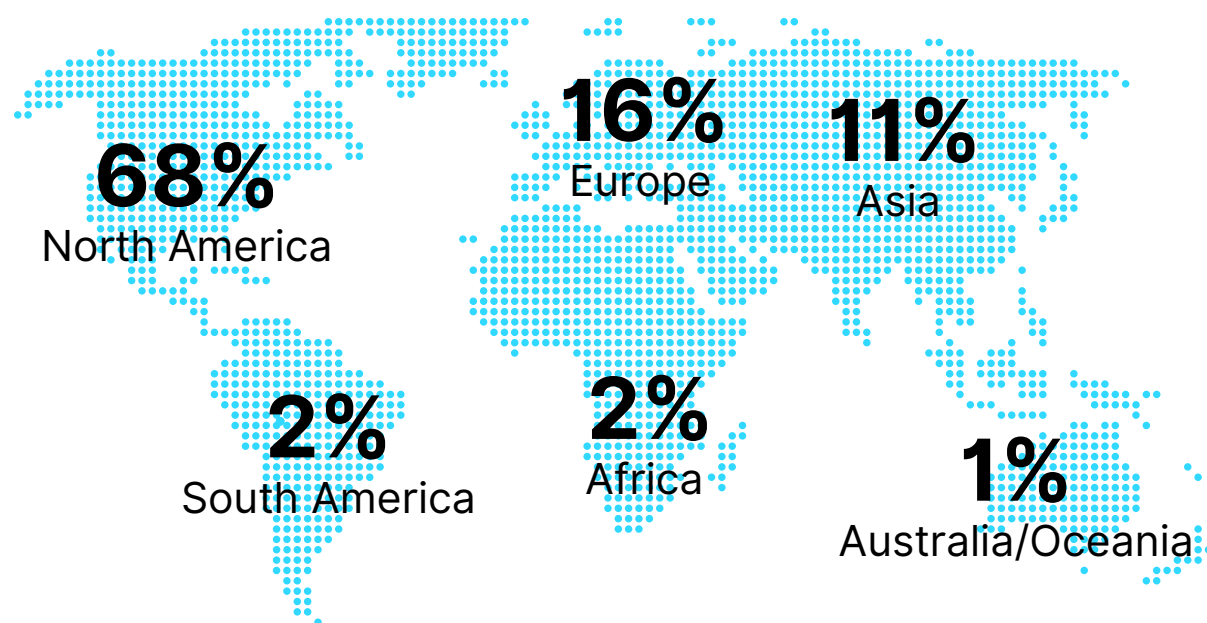
**Leo Vasiliou**




# Demographics and Meta

The SRE Survey, used to generate insights for this report, was open for six weeks during July and August 2024. The survey received 301 responses from all across the world, and from all types of reliability roles and ranks.

## Where are you personally located?




## Which is closest to your level of managerial responsibility?



## In what industry(ies) does your company operate?



## How many employees does your company have?





In today's exacting digital age, performance is paramount. The top online retailers, Global2000, CDNs, cloud service providers, and xSPs all rely on Catchpoint to ensure high performance and digital resilience by catching issues across the Internet Stack before they impact their customers, workforce or digital experiences.

Catchpoint's Internet Performance Monitoring (IPM) suite offers Internet Synthetics, RUM, BGP, Tracing, performance optimization, high fidelity data and flexible visualizations with advanced analytics derived from the world's largest, most detailed, active observability network.

Learn more at <https://www.catchpoint.com/>



# About Catchpoint