2025

State of Observability

The Rise of a New Business Catalyst

splunk>
a CISCO company

Contents

Executive foreword

When I started in observability over a decade ago, our mission was straightforward: keep services and systems running. Understand what's going on and who's affected, isolate the issue, and fix it. Today, software doesn't just support the business, it *is* the business.

As digital experiences have become the primary vehicle of customer engagement, observability practices have an impact beyond server rooms and NOCs. Correlating telemetry data with business outcomes is powering major decisions, like how to improve customer satisfaction or even what products to build. And as AI triggers the next seismic shift, observability practices are taking on a new level of responsibility — monitoring complex and dynamic AI workloads to ensure performance, reliability, and trust. This evolution positions observability not just as a foundation for customer experiences, but as a key enabler of AI-driven innovation and business growth.

For the *State of Observability 2025*, we surveyed 1,855 ITOps and engineering professionals to better understand this transformation and to identify what sets high-performing teams apart. A standout group of respondents contributes to the bottom line more than their peers. They're doing it by collaborating more with security teams, handling incidents more strategically, and investing in more forward-looking technologies and newer practices.

Observability practitioners are now embedded in high-level planning and decision-making, bringing their expertise (and their data) to customer engagement strategies, product roadmaps, and the boardroom. They've got a piece of the action — now it's time to use it.

Patrick Lin
SVP and GM, Observability, Splunk

Chapter 1

Observability becomes a new business catalyst

Behind every bold business move or innovation is an essential spark — a catalyst that drives action and purpose. It could be market research that reveals a surprising truth, like a different demographic using your product in unexpected ways, opening doors to a new market. Or, it could be a minor, overlooked feature that turns out to be the key to customer engagement, and suddenly, the product roadmap shifts.

These moments of insight don't happen by accident. They're powered by something sometimes quiet but critical. Observability is emerging as that new force. And while observability and the insights it can surface may not make the final decisions, without it, nothing moves fast enough to matter.

Data strengthens the observability-business connection

Organizations understand now more than ever that software-related decisions have sweeping ramifications on customer experience, brand perception, and much more. They're using observability data to answer the question, "How can we use the data from our applications to make business decisions?" rather than, "What's broken?"

To inform those decisions, observability practices are now making it a top priority to capture business metrics. Nearly three-quarters (74%) say that monitoring critical business processes is *moderately* to *very* important. Respondents reporting high ROI from their observability solutions are particularly emphatic on this point, selecting critical business processes more frequently than any other option — suggesting it is a top observability capability.

Did the business experience a spike in revenue because of a new feature release, or because the marketing team launched a killer campaign? Once upon a time, practitioners could only uncover the answer with a lot of time, a keen analytical eye, and saint-level patience. Now, observability is the catalyst that turns application telemetry into business action. Practitioners can use observability data to tell you *why* revenue dipped, *where* customer friction is happening, and *how* product performance is influencing business growth.

Product teams should be working in lockstep with engineering to inform roadmap decisions and which features they should prioritize based on insights from telemetry. Democratizing that business data is the best way to make that magic happen. If teams are waiting for a business analyst to pull data from three separate dashboards, you've already missed the moment.

— Greg Leffler, Director of Developer Evangelism, Splunk

The importance of observability capabilities to the business

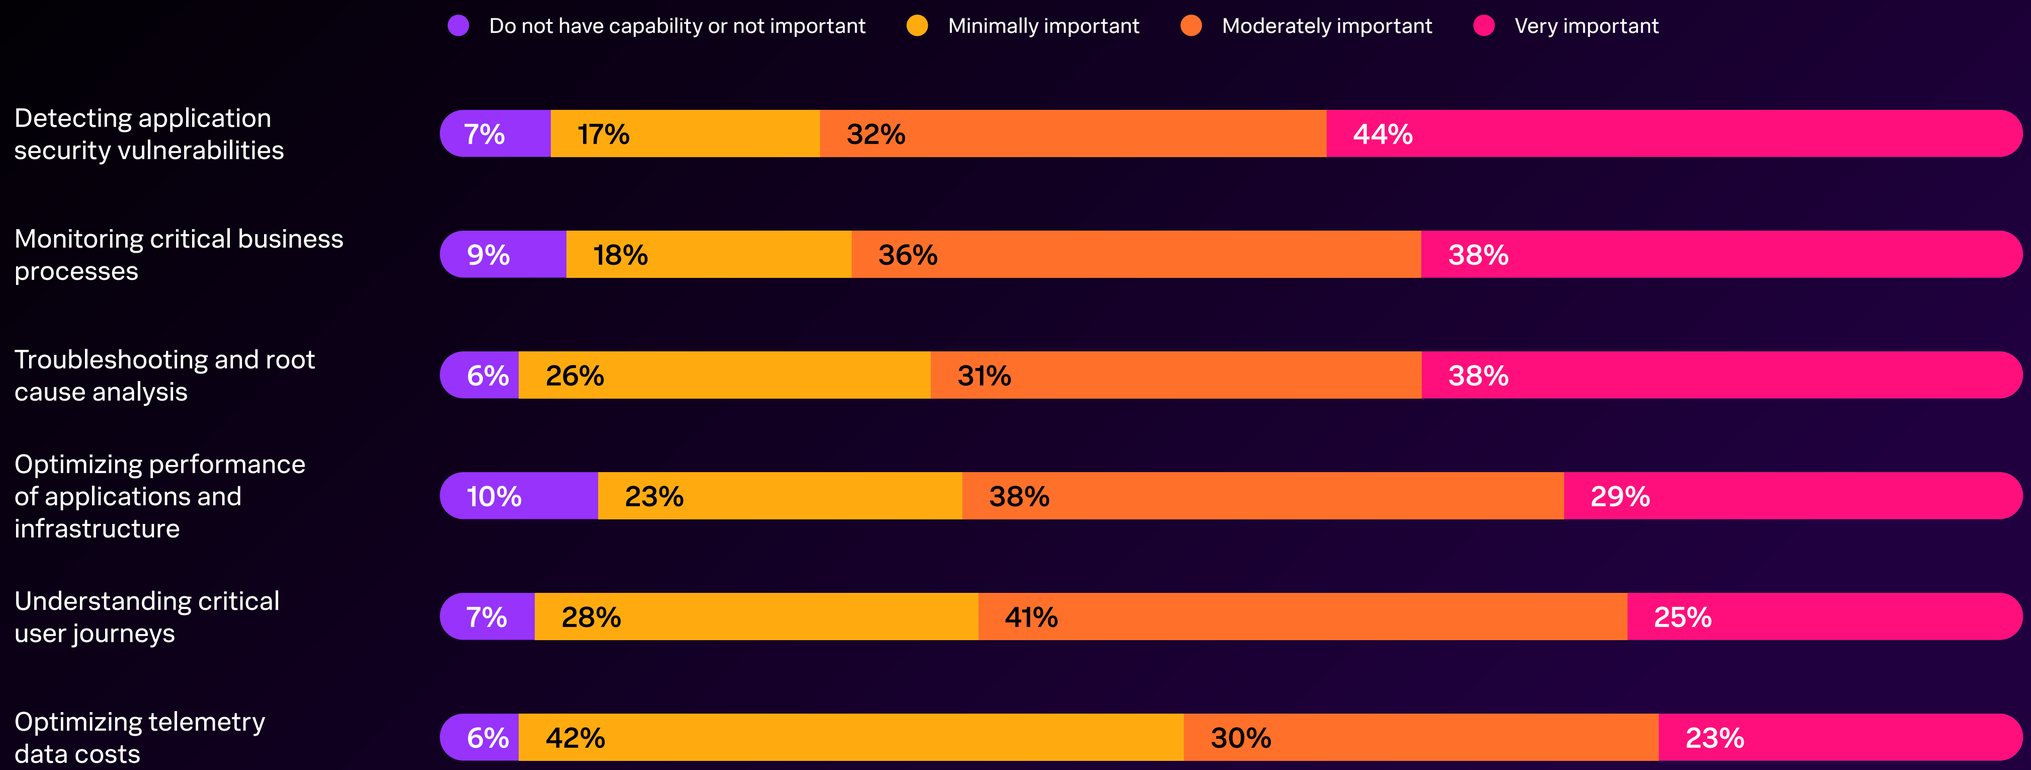

Percentages may not add up to 100% due to rounding.

An observability practice should not only have deep insight into business metrics, but also help drive outcomes like faster performance, better user experiences, and greater revenue. When they have reliable metrics, engineering and ITOps teams can turn their attention to tasks that directly impact the business, like understanding critical user flows, building executive real-time dashboards that inform both technology and business strategies, and correlating application performance issues and opportunities to revenue. An e-commerce company, for example, can use observability to visualize the full journey from website visit to order fulfillment, and prevent revenue loss by ensuring key business systems stay online. So it tracks that 65% of respondents rate their observability solution's ability to understand critical user journeys as *moderately* to *very* important to the overall business.

Organizations are delivering on those demands; 65% of respondents say their observability practice positively influences revenue. And 64% say their observability practice positively impacts product roadmaps. Product teams can rely on real user monitoring (RUM) data to understand how long it takes a page to render completely, or how quickly users can interact with it, and correlate this data with app performance metrics to draw conclusions — for example, adding functionality to a website could slow performance and increase cart abandonment rates. Or, a software release with more aggressive fraud protection may cause revenue to dip.

The impact of observability on the business

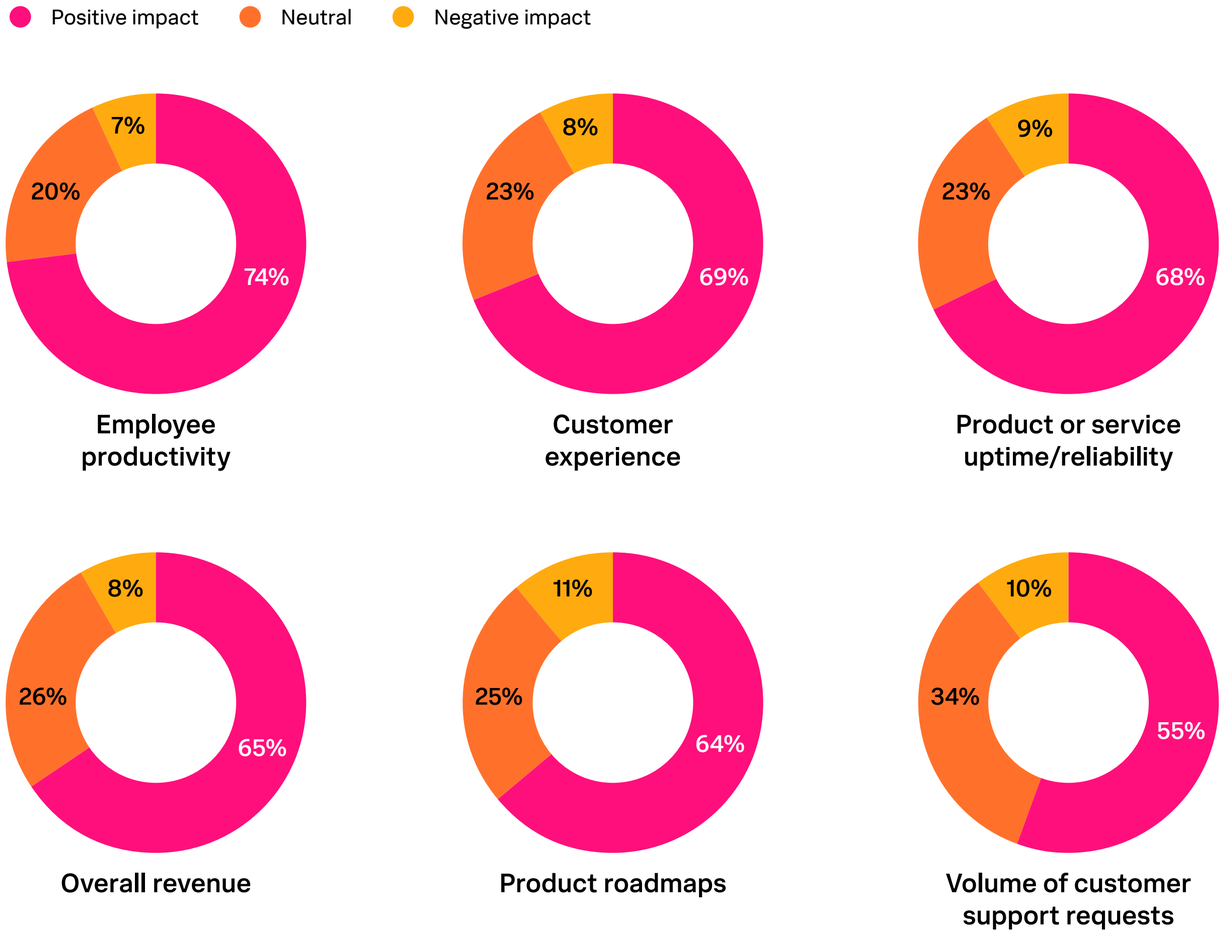

Percentages may not add up to 100% due to rounding.

Chapter 2

Relieving observability pressure points

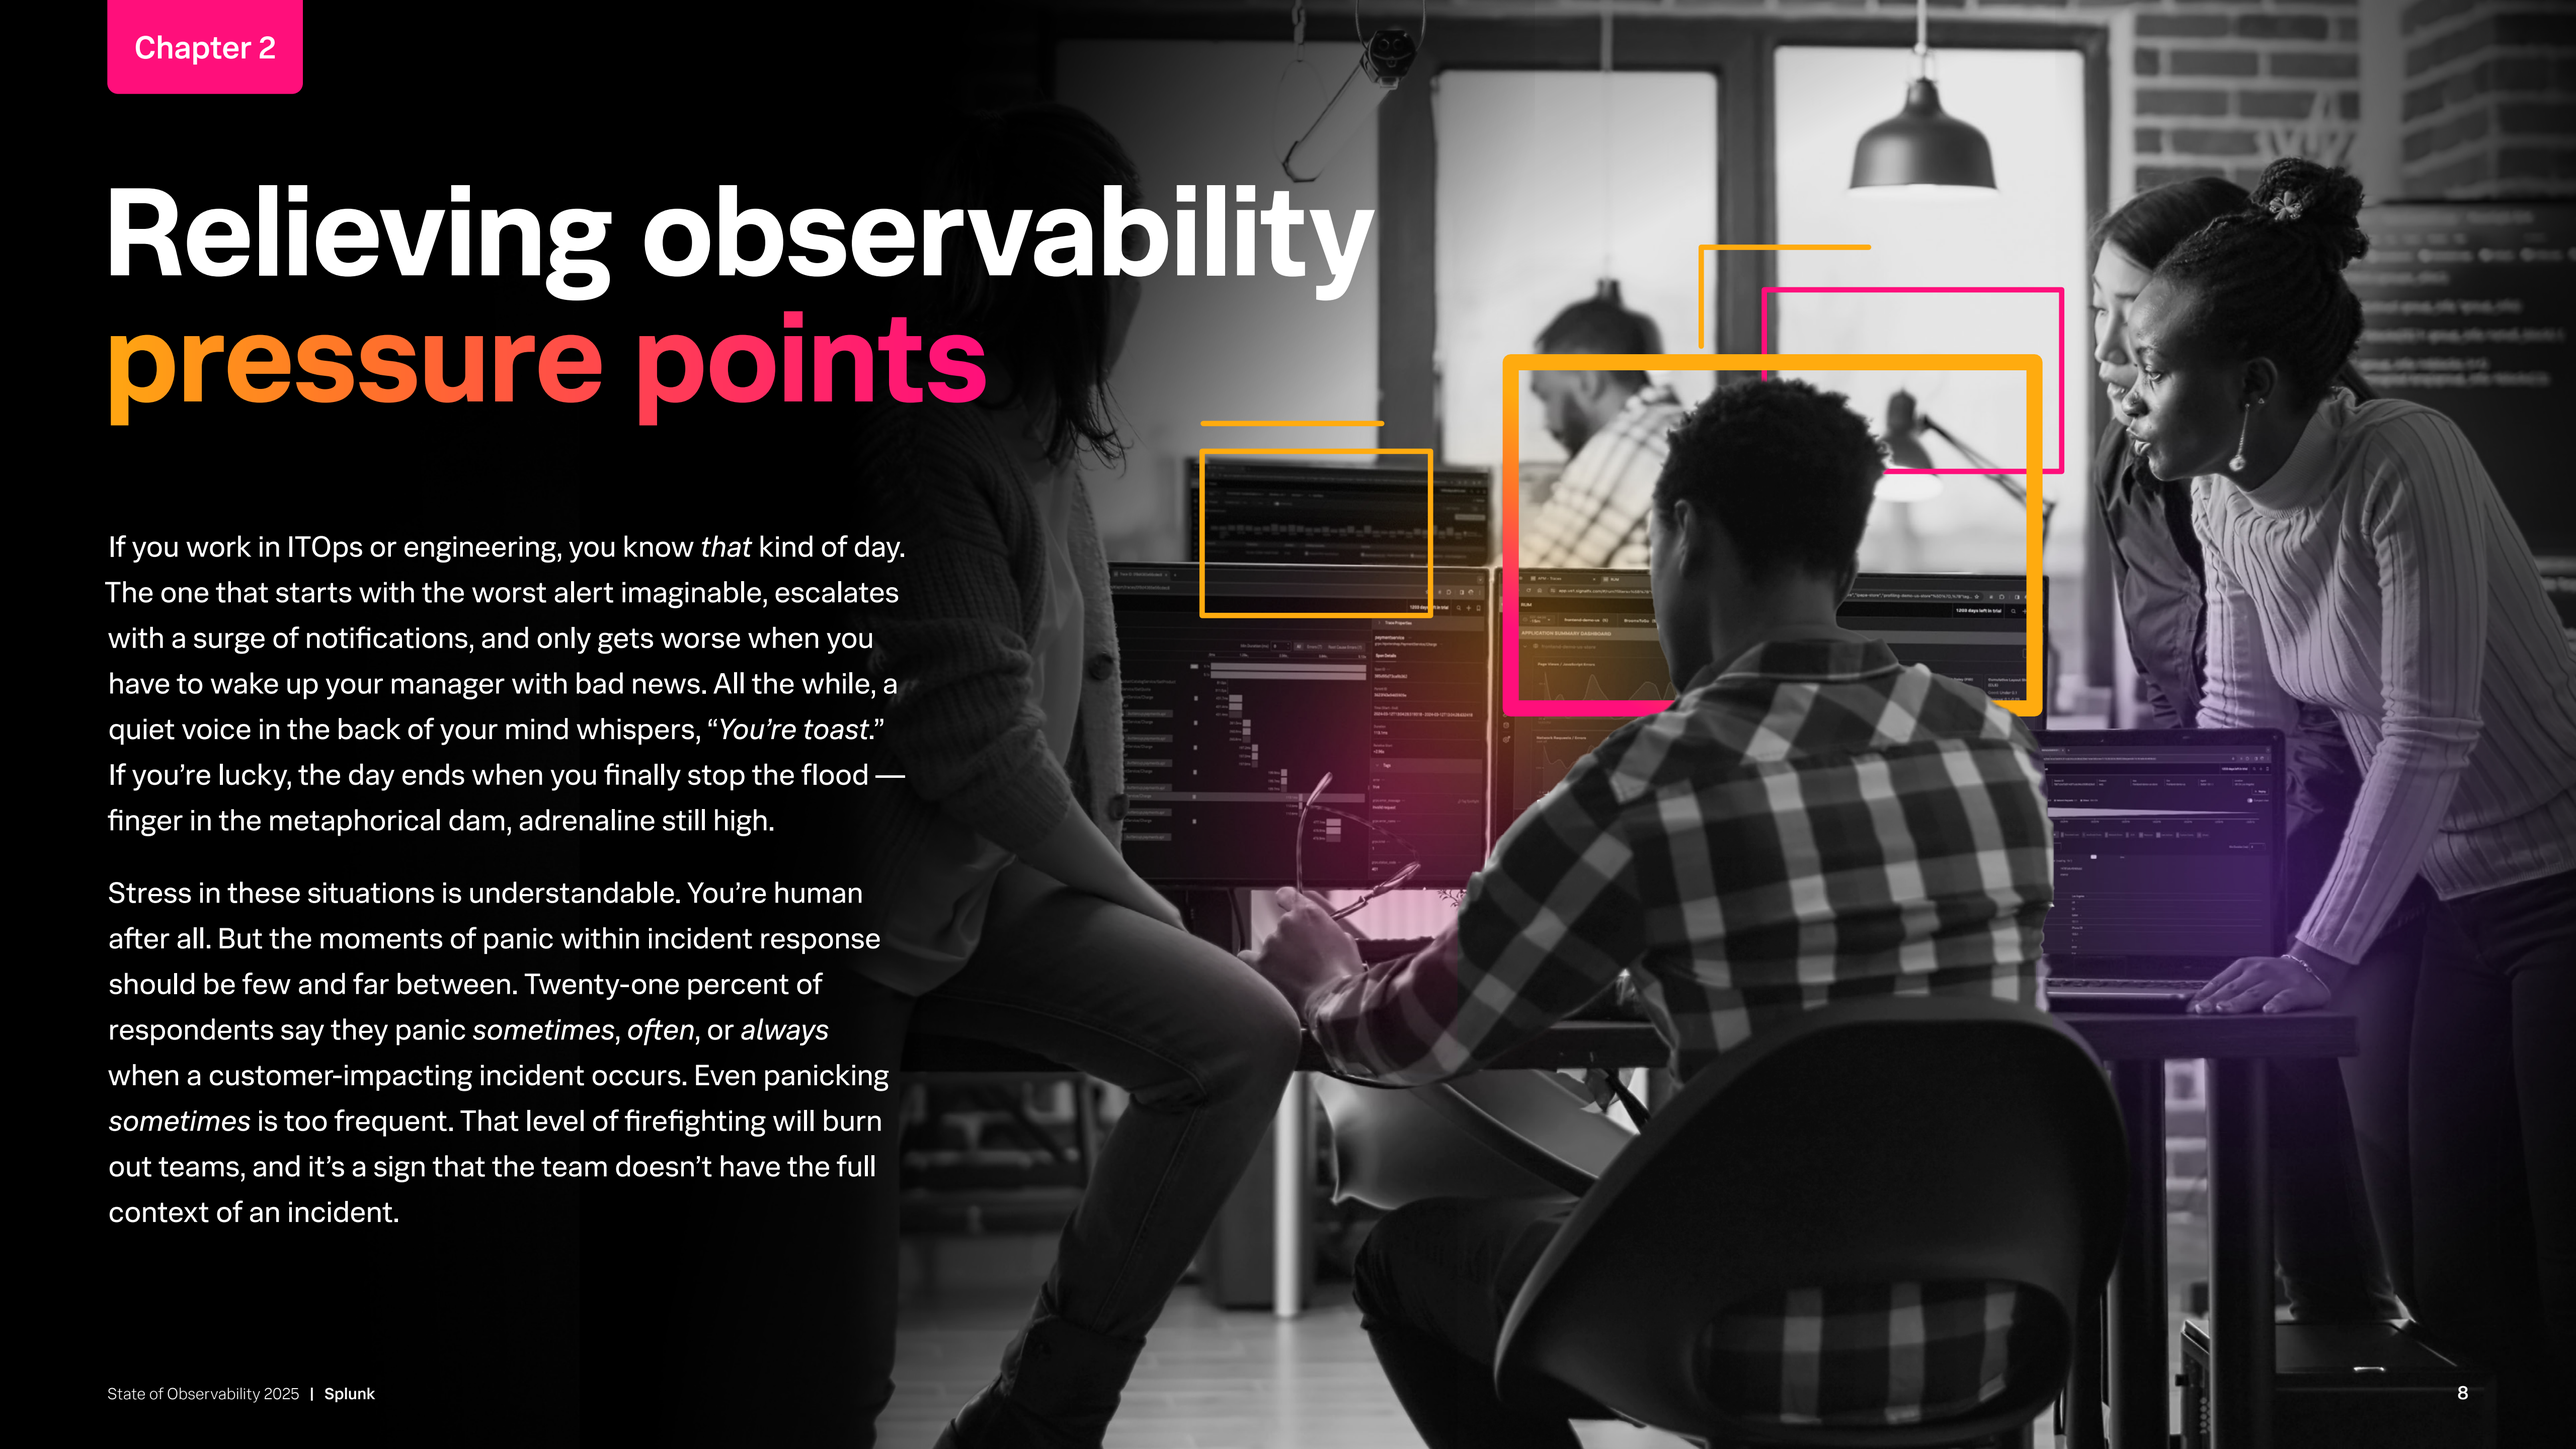

If you work in ITOps or engineering, you know *that* kind of day. The one that starts with the worst alert imaginable, escalates with a surge of notifications, and only gets worse when you have to wake up your manager with bad news. All the while, a quiet voice in the back of your mind whispers, "*You're toast.*" If you're lucky, the day ends when you finally stop the flood — finger in the metaphorical dam, adrenaline still high.

Stress in these situations is understandable. You're human after all. But the moments of panic within incident response should be few and far between. Twenty-one percent of respondents say they panic *sometimes*, *often*, or *always* when a customer-impacting incident occurs. Even panicking *sometimes* is too frequent. That level of firefighting will burn out teams, and it's a sign that the team doesn't have the full context of an incident.

Solid plans keep panic at bay

Pressure is a given when it comes to working in ITOps or engineering. An observability practice is measured by how well it can respond to incidents — and prevent future ones from occurring — so it's important for teams to remain calm and capable when alerts are flying, and everything feels like it's on fire.

Runbooks, response plans, and post-incident reviews are all methodical, strategic approaches to alleviate panic. Over half (54%) of respondents *often* or *always* develop a detailed response plan, and 71% say they *often* or *always* perform a detailed post-incident review.

There's also strength in numbers, and sometimes it's comforting to turn to a colleague in the trenches with you. But there's a stark difference between strategic collaboration and all-hands-on-deck reactivity. War rooms can easily spiral into a "too many cooks in the kitchen" scenario, and burn through precious company resources as valuable team members get pulled into investigative black holes.

Yet 20% say they *often* or *always* start a war room that includes members of many teams until the issue is resolved — a sign that reactive firefighting is still common and that it's hard to know where to look when cascading problems occur.

"War rooms proliferate when an organization's tooling isn't effective enough to help teams isolate the problem domain," says Patrick Lin, SVP and GM, observability at Splunk. "Observability software has evolved to the point where teams should be able to restore services without getting 50 people on a call together."

A smarter approach is to isolate the incident to a specific team, and rely on that team to resolve it, which only 22% say they do *often* or *always*. This is a sign of a mature observability practice with advanced collaboration practices (we'll get to that later).

73%

experienced outages due to ignored or suppressed alerts

Post-incident reviews bring closure, and that can be incredibly cathartic. Knowing that an incident — no matter how anxiety-inducing it was to begin with — won't happen again does wonders for everyone's mental health.

— Caitlin Halla, Developer Evangelist, Splunk

False alerts and poorly managed tools drain morale and ROI

Incidents are undoubtedly stressful, but they are not the biggest detractors to teams' mental health compared to other factors. Only 25% of respondents say incident response negatively impacts morale, whereas 59% say tool sprawl is a source of angst.

Following closely behind is a problem that's plagued ITOps and engineering teams for years: the volume of false alerts. The two problems are intractably intertwined. The more tools an organization accrues, the higher likelihood those tools will generate false alarms — especially if teams are stretched so thin that they can't dedicate the time to fine-tune alerting rules, correlate signals across systems, or validate what truly matters.

Tool sprawl is a real challenge, but what truly undermines ROI is the poor quality of detections across those tools. When alerts are noisy, redundant, or lack context, even the most advanced toolsets can't deliver meaningful value.

— Stephanie Elsesser, Director of Observability Strategists, Splunk

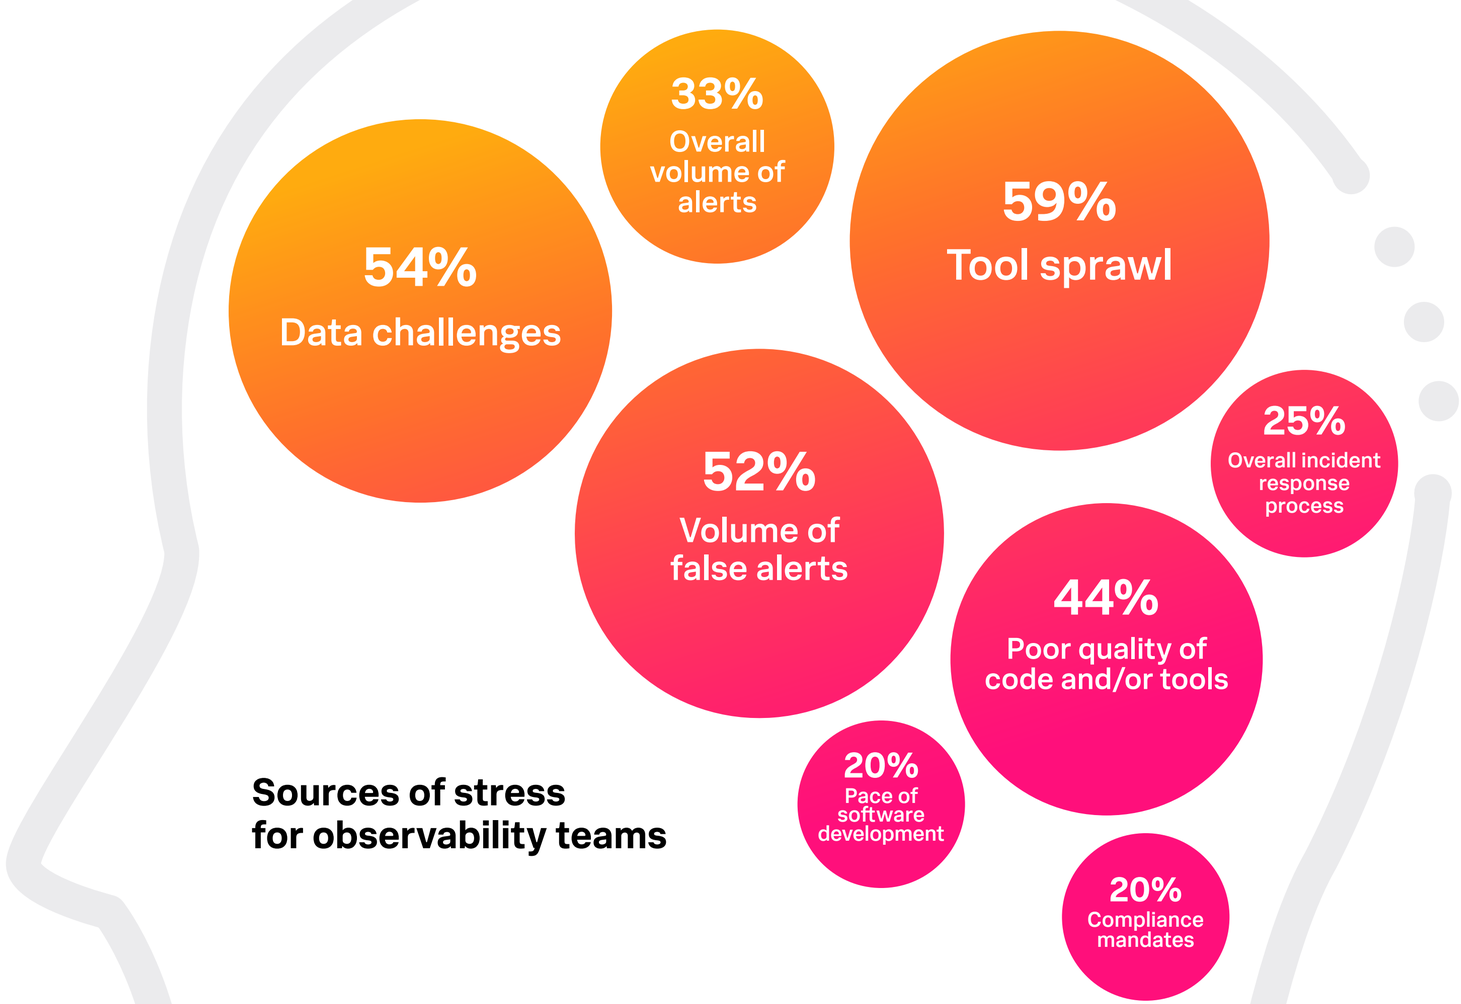

Sources of stress for observability teams

Respondents could select all that apply

False alerts don't simply stress teams out. They have wide-ranging consequences, including a direct impact on the bottom line. Over half (54%) of respondents say the quality of their alert detections is a top driver of their observability ROI, and 47% say alerts *significantly* influence security decisions within their organization.

In an attempt to escape the noise of false alerts, some teams are resorting to risky methods. Thirteen percent say they *often* or *always* ignore or suppress alerts. Even more alarmingly (pun not intended), 73% have experienced outages due to ignored or suppressed alerts.

"A high-functioning observability practice should not suppress any alerts, period," says Greg Leffler, director of developer evangelism at Splunk. "Ideally, alerts should only indicate an immediate problem that has business impact."

Whether teams are chasing down false alarms, triaging a barrage of alerts, investigating an unclear incident, or reconfiguring thresholds on the fly, it's clear that alerts are taking up too many productivity cycles, not to mention causing too many mental health days. Forty-three percent of respondents admit they spend "more time than they should" responding to alerts.

If you walk by any alert board, you'll probably see a slew of unread, unactioned alerts lingering at the bottom. That's reality, but not the ideal state. Alerts are meant to spark immediate attention and contain enough context to take action quickly, not sit idle like background noise.

"Alerts will always be a central part of observability," says Mike Simon, staff developer evangelist at Splunk. "But to make them actionable at scale, you need to focus on improving signal quality, not just reducing noise. Correlate what's meaningful, understand business impact, and surface what truly needs attention so teams can dive deeper when it matters. That's how you recover wasted time and refocus engineers on what matters: building better software."

Factors having the greatest impact on observability ROI

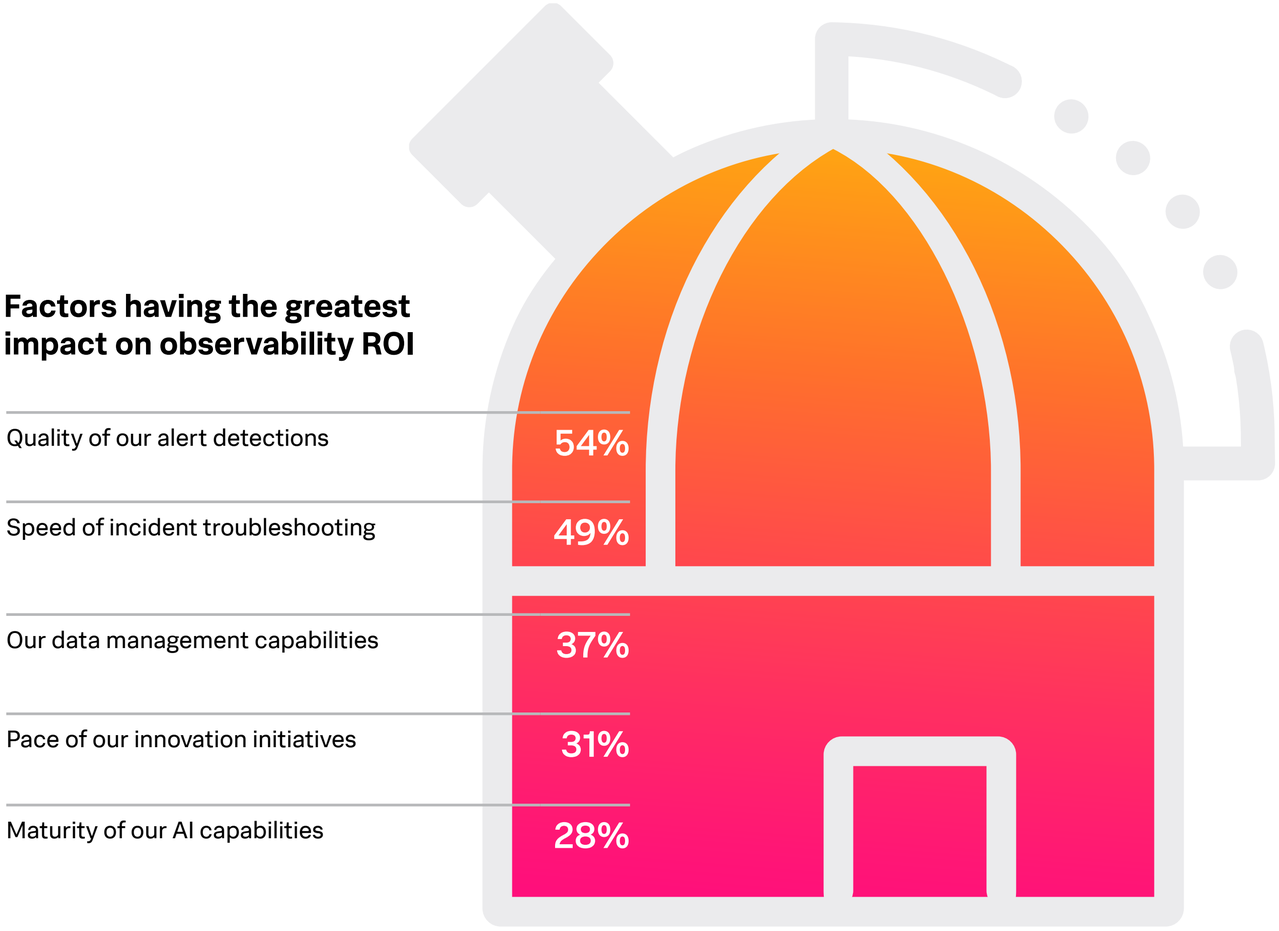

Respondents could select all that apply

Chapter 3

Collaboration with security expands observability's influence

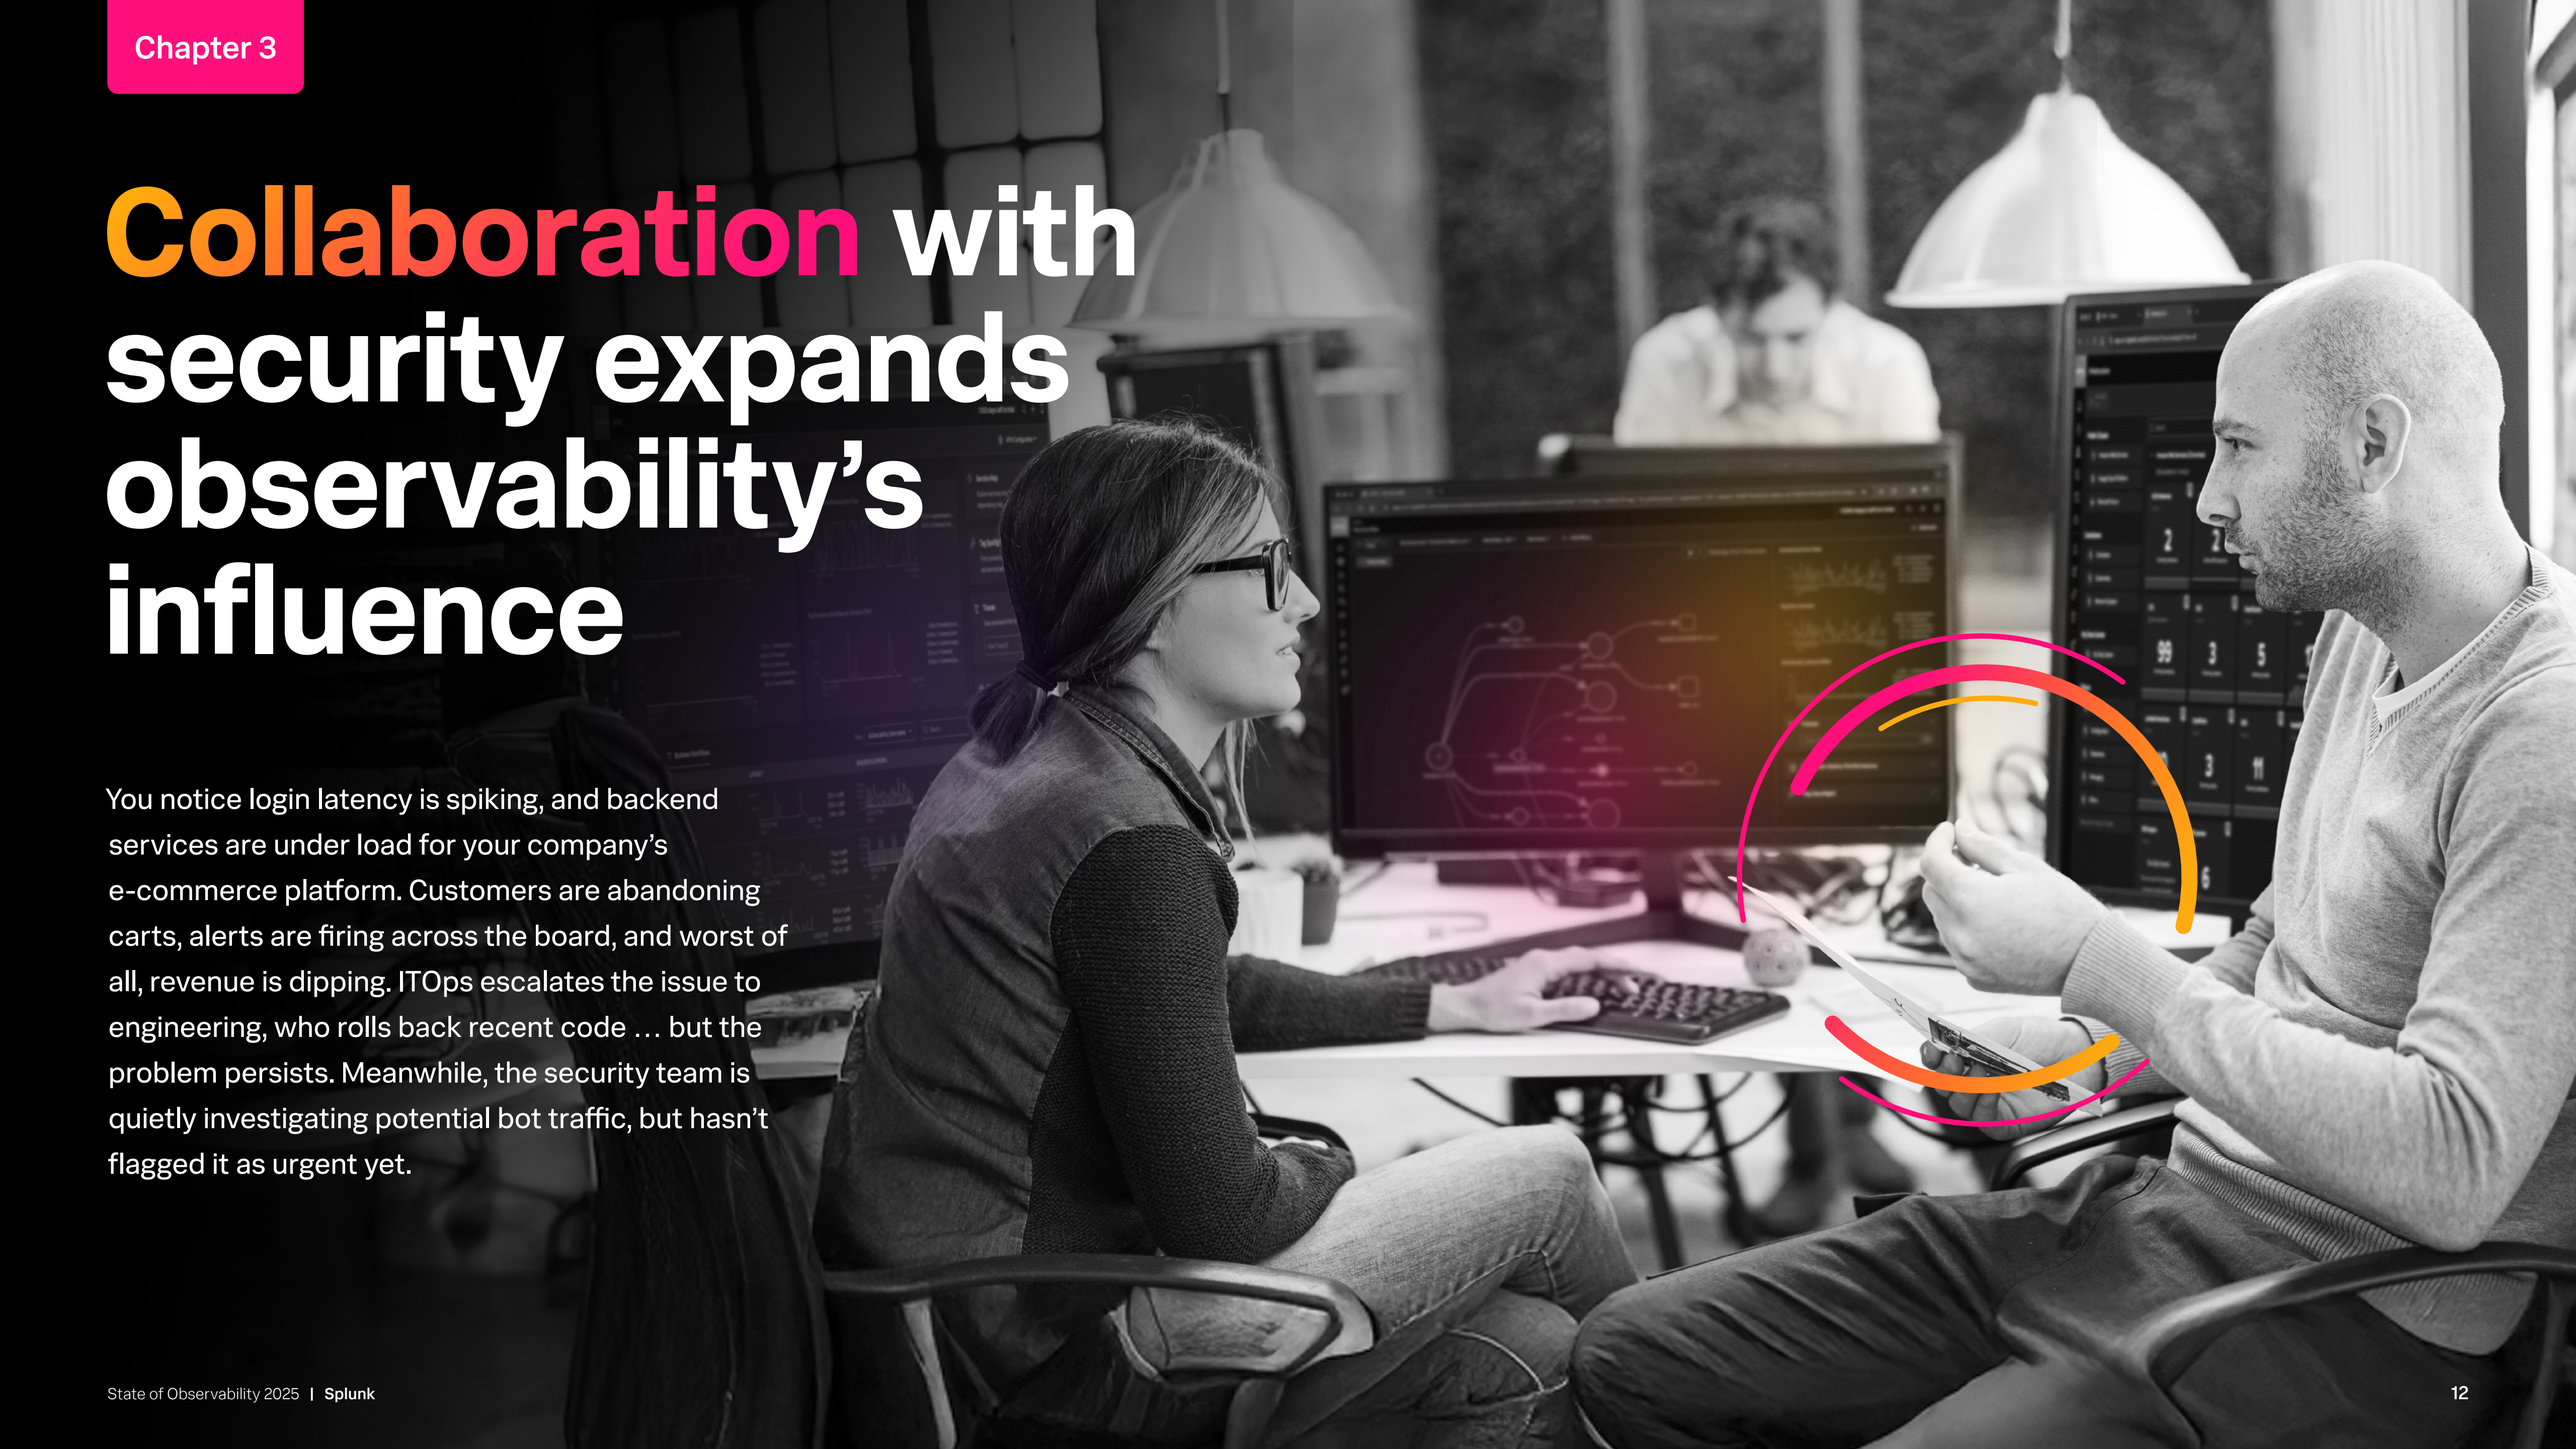

You notice login latency is spiking, and backend services are under load for your company's e-commerce platform. Customers are abandoning carts, alerts are firing across the board, and worst of all, revenue is dipping. ITOps escalates the issue to engineering, who rolls back recent code … but the problem persists. Meanwhile, the security team is quietly investigating potential bot traffic, but hasn't flagged it as urgent yet.

Each team could go down separate rabbit holes, burning time and wasting effort. Or they could use shared data, dashboards, navigators, and context within the observability platform to troubleshoot in parallel. By working together, they could quickly uncover the root cause — a credential-stuffing attack that is overwhelming backend resources — and solve the problem quickly to mitigate customer impact.

The payoff is real for observability practices that collaborate with security teams. Nearly two-thirds (64%) of all respondents encounter fewer application and infrastructure performance issues, 54% improve data quality, and 54% say they waste less time chasing down issues, which translates into faster MTTD and MTTR.

That benefit extends to the overall business, too. Sixty-four percent say collaboration with security teams has led to fewer incidents that affect customers. Organizations are realizing that the value of observability data doesn't end with ITOps and engineering teams. Over three-quarters (76%) say the ability for their observability solution to detect application security vulnerabilities and threats is *moderately* to *very* important to their organization's overall business.

Collaborating with security teams pays off

Respondents reporting collaboration benefits

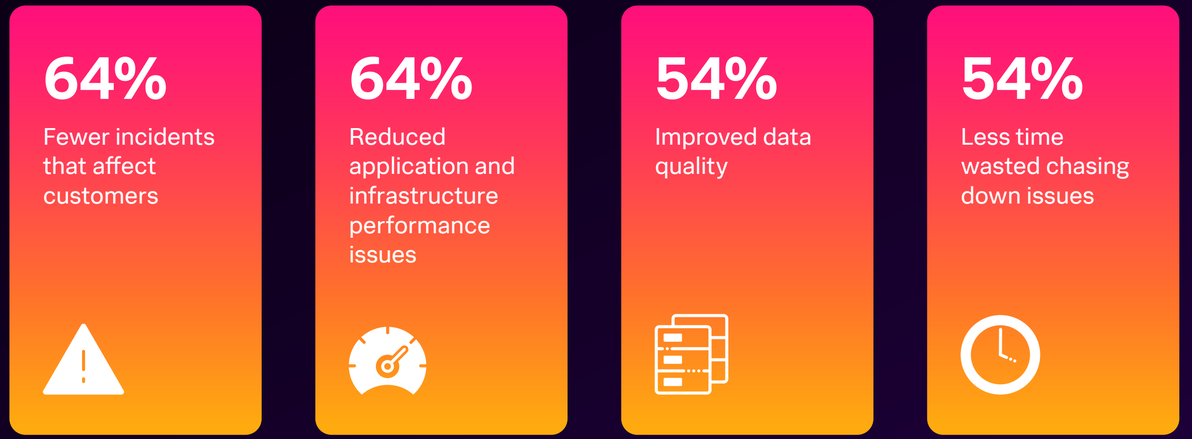

Respondents could select all that apply

Less advanced security and observability teams use their own tools, have distinct priorities, and often communicate only when absolutely necessary, usually during an incident. But now those walls — of data, tooling, and communication — are breaking down. Observability practices that haven't already taken steps to enable that collaboration will be left behind, especially with the growing use of AI.

— Patrick Lin, SVP and GM, Observability, Splunk

Partnership powers faster resolution

Collaboration, synergy, synchronicity, good old-fashioned teamwork … Whatever you call it, working with security teams is a deliberate, intentional process. It takes more effort than simply tossing observability data over the wall and hoping it lands.

Seventy-four percent of respondents say their observability and security teams share and reuse data — a fundamental first step toward collaboration. Meanwhile, 68% report that both teams use the same set of tools.

These practices should be table stakes. Working together in real time surfaces context you just can't get from dashboards alone. Let's say engineering rotated the API key of a backend service, but they didn't update an upstream service to use the new key. As they roll out the new version, user requests start to fail, leading to retries and increased latency. It often takes merging latency spike data with security logs to spot this — a level of correlation not typically visible in most observability dashboards.

Passing data back and forth is fine, but real teamwork happens when observability and security teams are on the virtual frontlines together from the start, rather than waiting for issues to slowly filter through siloed workflows.

"When observability and security software are siloed, there is no deep linkage between tools, and collaboration becomes a painful manual exercise," says Mark Maslach, vice president of global technical sales, Splunk observability at Cisco. "That's often a sign of organizational issues."

More advanced forms of collaboration reflect true organizational maturity, as they require teams to actively break down silos and work closely together. For example, 62% say their team troubleshoots and solves issues in tandem with their security team, and 63% say they can rapidly distinguish between application and infrastructure performance issues that have security root causes.

"Security, ITOps, and engineering teams will likely always remain distinct to some degree because their skill sets and incentives are simply too different," says Craig Robin, field CTO at Splunk. "However, we see mature observability practices having incidents triaged and sent to the appropriate specialized team as quickly as possible, making sure they have the right data at their fingertips to resolve issues efficiently. That's the right way to deal with business-impacting issues."

Top ways observability teams and security teams work together

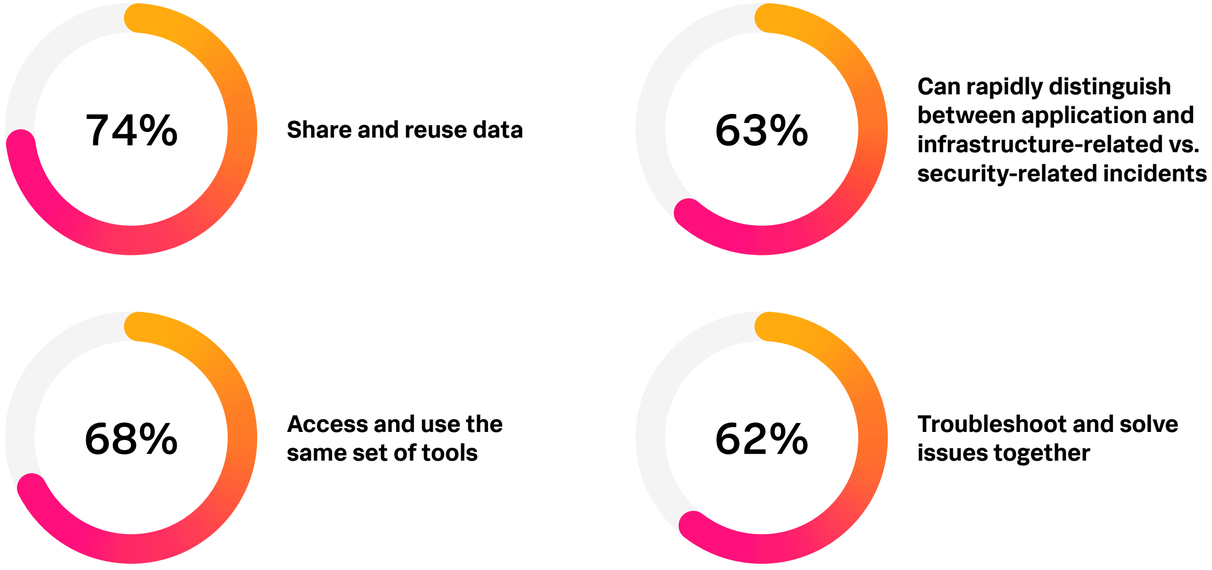

Skills gaps and silos stifle collaboration

Observability and security teams partnering together during incident response may be a sign of a mature observability practice, but these practices aren't widespread yet. The biggest barrier to improving collaboration with security is resistance to change, cited by 59% of respondents.

Security and observability organizations take fundamentally different approaches to incidents. While security teams show their value by closing tickets ("We found 2,000 possible attacks and mitigated them"), ITOps and engineering teams aim to keep incident counts low. They may even butt heads over whether something qualifies as an incident in the first place. Resistance can also manifest as finger-pointing, blame-shifting, and the age-old game of 'Who gets the ticket?'

Knowledge gaps are another top reason that security and observability teams aren't collaborating effectively, with 41% of ITOps and engineering teams reporting a lack of technical expertise and relevant skill sets as a challenge.

"SREs and NOC engineers have very little insight into what security problems are because they were never trained on it," says Leffler. "Meanwhile, security teams aren't as worried about application performance, as long as no one is hacking it."

A third (34%) of respondents point to their software or lack of technological maturity as an obstacle to their collaboration efforts. Many organizations still operate with siloed security and observability platforms, which makes it hard to correlate signals across teams and systems in real time. For example, when a security incident and event management (SIEM) platform triggers a DDoS alert for a customer-facing application, some observability software may only show the performance issues that stem from that attack, like increased latency, error rates, or resource utilization — leading ITOps and engineering teams to chalk up the incident as performance problems.

"The most advanced organizations have realized that observability data is security data, and they're adopting a unified approach that allows technology to do the heavy lifting in identifying security impacts," says Robin.

The toughest obstacles to collaboration

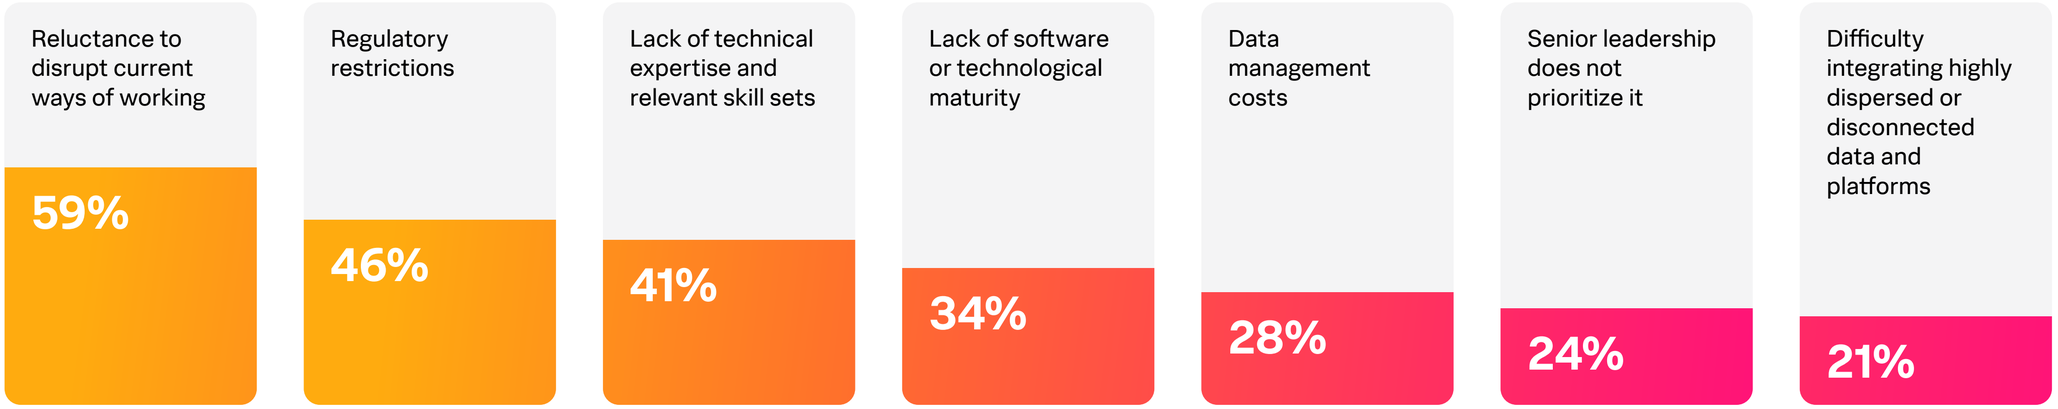

Respondents could select all that apply

Chapter 4

Observability in the AI era

AI has piqued the curiosity of even the most hardened skeptics as observability practitioners realize that it delivers incredible value when properly implemented (emphasis on *properly*).

The bulk of ITOps and engineering teams have embraced AI, with 76% of respondents regularly using AI in their everyday workflows. However, adoption rates vary across the different flavors of AI, with 54% *often* or *always* using AIOps — an ITOps staple that's been around for nearly a decade — and 39% *often* or *always* using generative AI.

Only 18% *often* or *always* use emerging AI technologies like fast-growing agentic AI, which can learn, reason, adapt, and act autonomously, enabling agents to complete entire workflows like coding and debugging software. However, agentic AI adoption will likely increase sharply over the next few years.

Observability practices lean on AI

How often respondents use different types of AI

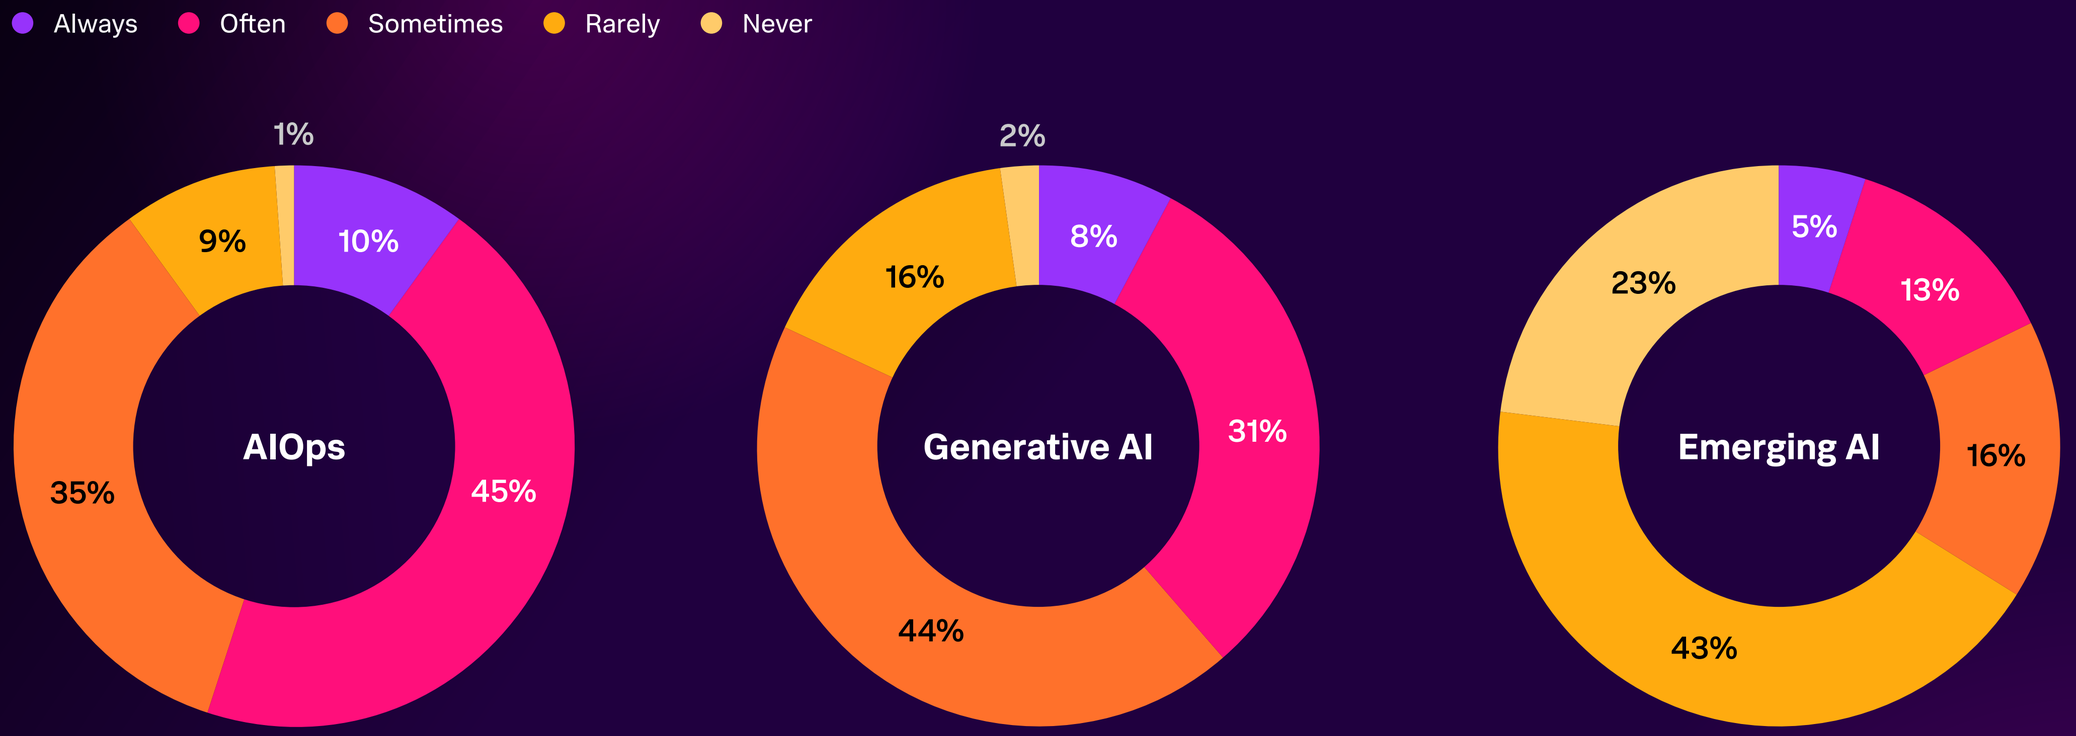

AI enables innovation, speeds troubleshooting

ITOps and engineering teams recognize AI's ability to boost productivity; 78% say AI has enabled them to spend more time on innovation than maintenance, helping them deliver better business outcomes. With those time savings, teams can focus on a range of high-impact initiatives, from implementing microservices and serverless technologies to developing new digital products.

This should be music to the ears of teams who struggle to balance priorities, as 42% admit they spend more time than they should on app maintenance like editing code and enabling feature flags. Nearly half (45%) of respondents say they spend less time than they should building new software, with 12% of this group saying they spend *significantly* less time than they should on it — the highest rate of any task we asked about.

A generative AI assistant can help by answering questions about applications and infrastructure — which is particularly powerful for less experienced team members who might otherwise spin their wheels trying to complete maintenance tasks. For instance, a junior engineer could ask a generative AI assistant to analyze a trace ID during a service interruption to get remediation recommendations and even a full incident report.

"Give your junior analysts tools that can analyze logs, metrics, and traces with precision, and let generative AI do the heavy lifting on context and pattern recognition," says Cory Minton, field CTO, Splunk. "That way, your elite engineers — the ninjas — can focus on the work that really matters, like building automation and engineering systems that scale."

Respondents expect AI to add value in areas that are most critical to the business. Most frequently, they cited detecting application security vulnerabilities and threats as an impactful observability capability, with 58% claiming AI will have a positive impact in this area.

Similarly, 69% say the observability capability of troubleshooting and root cause analysis is *moderately* to *very* important to the business — and respondents expect AI to help the most here, with 60% saying it will have a positive impact. AIOps in particular can accelerate root cause analysis by discovering granular trends contributing to service issues and identifying underlying code-level issues.

The rapid growth of generative AI has paved the way for agentic AI to take on even more complex, autonomous roles in observability. We're moving toward a future where AI agents can manage entire incident workflows end-to-end.

— Julie Gibbs, Vice President of Splunk AI and Integrations Product Marketing, Splunk

78%

spend more time on innovation than maintenance with AI's help

Data quality impacts AI readiness

Reaping the benefits of AI requires far more than plugging it in and letting it do its thing. Successfully adopting AI, or even just achieving AI readiness, involves so much more — like incorporating it into the team's daily operations, understanding its output, measuring its value, using it sustainably, and ultimately seeing the payoff (or not).

When it comes to AI's success, data quality and quantity are equally critical. Low data quality is the main barrier to AI readiness, with 48% of respondents citing it as one of the biggest challenges.

"Oftentimes, nobody is explicitly responsible for maintaining data quality," says Leffler. "Instead, a dev or SRE team will say, 'Let's just collect the golden signals — latency, errors, saturation, and traffic — with these four metrics,' and the data quality just needs to be good enough for troubleshooting purposes."

So who's in charge, then? Those who care deeply about observability can form a hands-on center of excellence to define and enforce data quality standards across the organization. This includes collaborating with key stakeholders like the compliance team to ensure the data serves everyone's needs.

Top barriers to AI readiness

1. Lack of data quality
2. Cost of AI infrastructure
3. Lack of expertise or understanding across teams
4. Reluctance to disrupt current ways of working
5. Low data visibility

Observability teams adjust to new AI dynamics

AI is a proverbial double-edged sword for observability practitioners; it helps get more done, but it also means spending more time monitoring the workloads it produces. Nearly half (47%) of respondents say monitoring AI workloads has made their jobs more challenging.

Yet the ability to understand and capture LLM data is crucial, especially since AI's reach extends far across the business.

AI workloads are incredibly dynamic, and often change as models are retrained or updated. Additionally, subtle changes in the data — also known as data drift — can degrade the model's performance without triggering traditional alerts.

AI also isn't a typical workload; it involves specialized infrastructure that often sits outside typical application stacks. Teams need to capture certain intricacies associated with an AI workload. For example, are the GPUs maxed out? How fast are tokens being generated and used? What is the model's response time? Did the model's behavior suddenly drift after a retraining? And most importantly, how much will all of this cost?

These are questions that an ITOps or engineering team may not be able to easily answer. Lack of expertise or understanding is a major obstacle, with 40% citing it as a major challenge to achieving AI readiness.

"It's important that a single team has all the context needed to monitor the performance of the entire application, including the AI," says Annette Sheppard, director of product marketing for observability at Splunk. "This means observability teams need to upskill their existing practitioners, and train them on the nuances they need to pay attention to."

47%

say monitoring AI workloads has made the job more challenging

If your AI systems aren't observable, they're already a liability. When models go sideways, it can happen quickly and sometimes quietly. AI needs observability more than most of your digital systems because it evolves in ways you don't expect.

— Cory Minton, Field CTO, Splunk

Chapter 5

OpenTelemetry evolves from a standard to a strategy

Over the past few years, OpenTelemetry has cemented itself as the industry standard for collecting observability data in a consistent, easy-to-understand format. Effectively, every observability vendor (more than 40 of them) supports OpenTelemetry, and many other applications are released with built-in support.

The technical benefits of OpenTelemetry are well established. *State of Observability 2024: Charting the Course to Success* revealed that OpenTelemetry enables organizations to access a broader tech ecosystem, meet data residency requirements, and more easily adopt modern cloud frameworks. But this year, its advantages are extending well beyond the observability practice. The vast majority who use OpenTelemetry at least *sometimes* say it positively affects revenue growth (72%), operating margins (71%), and brand perception (71%).

OpenTelemetry power users derive deeper insights

How, exactly, does OpenTelemetry's reach extend so far? OpenTelemetry captures distributed traces, metrics, logs, and profiles with little effort, and enriches these with standardized metadata, making it simple to unify data across different environments, languages, and platforms. OpenTelemetry also makes it easy to capture additional custom data that represents what matters to your business, or to modify the data being sent through it. Having all of this rich telemetry data enables teams to solve unique problems that would otherwise fly under the radar — or even worse, surface only when customers complain about them.

Let's say an organization runs a site with a very high traffic volume, but a proportion of visitors using a specific browser are experiencing login errors. The organization would have no insight into the problem without certain metadata.

OpenTelemetry benefits extend beyond the observability practice

Respondents who reported positive impact on business results

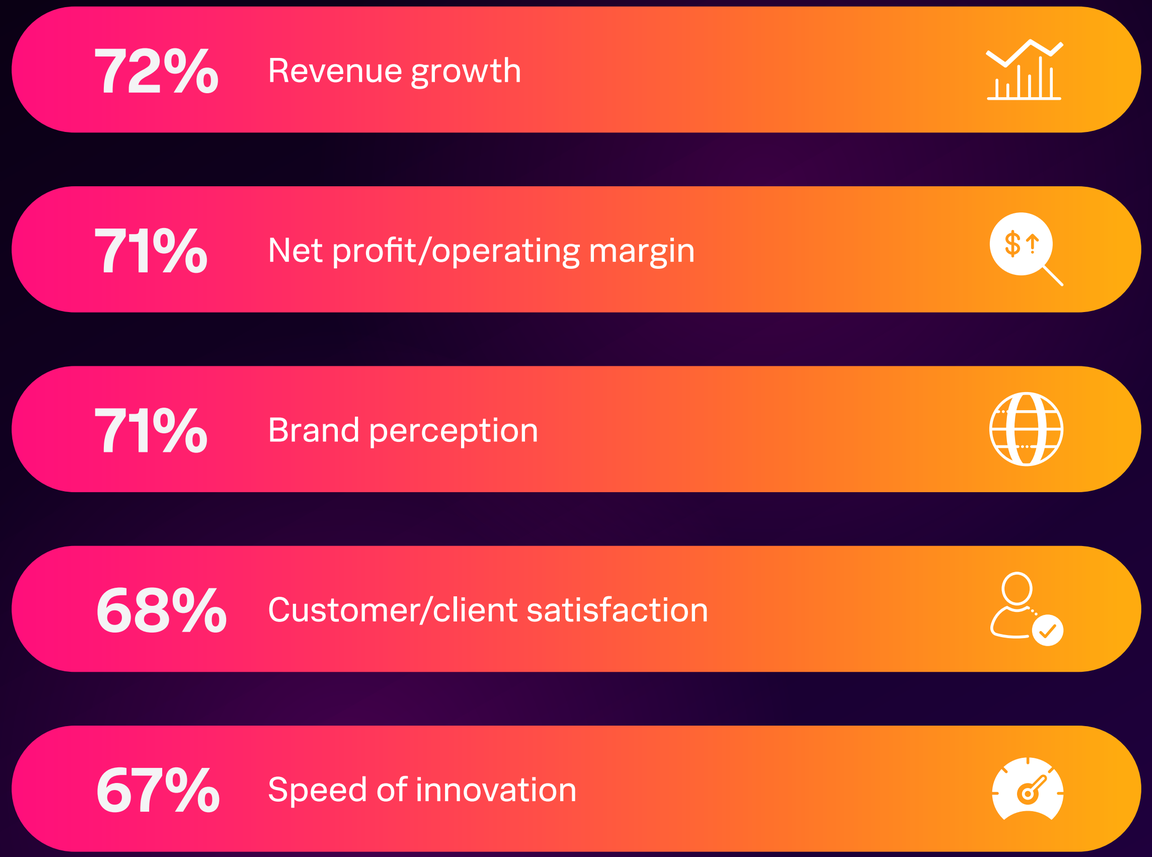

"When teams adopt OpenTelemetry, it usually means they've hit a turning point. They're not just collecting signals — they're investing in understanding how their systems really work," says Morgan McLean, senior director of product management at Splunk. "That mindset shift is what maturity looks like in modern engineering."

The deeper teams go with OpenTelemetry, the more benefits they reap, with frequent users — respondents who say they *often* or *always* use OpenTelemetry — managing incidents more calmly and systematically. In fact, 47% say they *never* panic during customer-related incidents, compared to just 32% of those who *rarely* or *never* use OpenTelemetry.

These power users are three times more likely to say their observability practice *significantly* impacts employee productivity compared to OpenTelemetry laggards, and are twice as likely to say their observability practice *significantly* impacts customer experience.

OpenTelemetry adopters tend to be more forward-thinking about other technologies, too — likely because they've built a culture that fosters intellectual curiosity and supports modern tooling. OpenTelemetry power users are far more likely to use generative AI, ChatOps, observability-as-code, and automatic remediation. For example, the majority (57%) of frequent OpenTelemetry users *often* or *always* use observability-as-code, compared to only 10% of OpenTelemetry laggards.

When teams standardize with OpenTelemetry, they're collecting richer data and laying the groundwork for better generative AI outcomes. A unified telemetry pipeline — combined with business-related tags such as customer ID and marketing campaign group — means richer, more consistent data to inform AI models, and that translates to more context-driven insights, better recommendations, and fewer blind spots.

OpenTelemetry power users boost business results even more

3x

impact on employee productivity

2x

impact on customer experience

OpenTelemetry is the ultimate foundation for any observability solution. It is the most capable, extensible, and future-proofed standard for telemetry.

— Morgan McLean, Senior Director of Product Management, Splunk and Co-Founder of OpenTelemetry

Chapter 6

Leading observability practices boost revenue and ROI

Observability isn't just about keeping systems running anymore. It plays a meaningful role in shaping the business. ITOps and engineering practitioners now influence high-stakes metrics like revenue and customer experience, and the teams making the biggest impact are acting as business catalysts that accelerate performance across the entire organization.

Our research singled out a distinct group of respondents rising above the rest that consistently achieves better outcomes than its peers. These state-of-the-art observability leaders are expanding their influence to the entire organization. They're nearly twice as likely as their peers to say that their observability practice *significantly* improves overall revenue, employee productivity, and product roadmaps. They also generate an annual 125% ROI from their observability practice — 53% higher than their peers.

What they have in common is a top-tier technology foundation; these respondents *often* or *always* use forward-leaning technologies — namely, OpenTelemetry, code profiling, and observability-as-code.

Leaders widen their circle of influence

Observability practices have a *significant* positive impact on the business

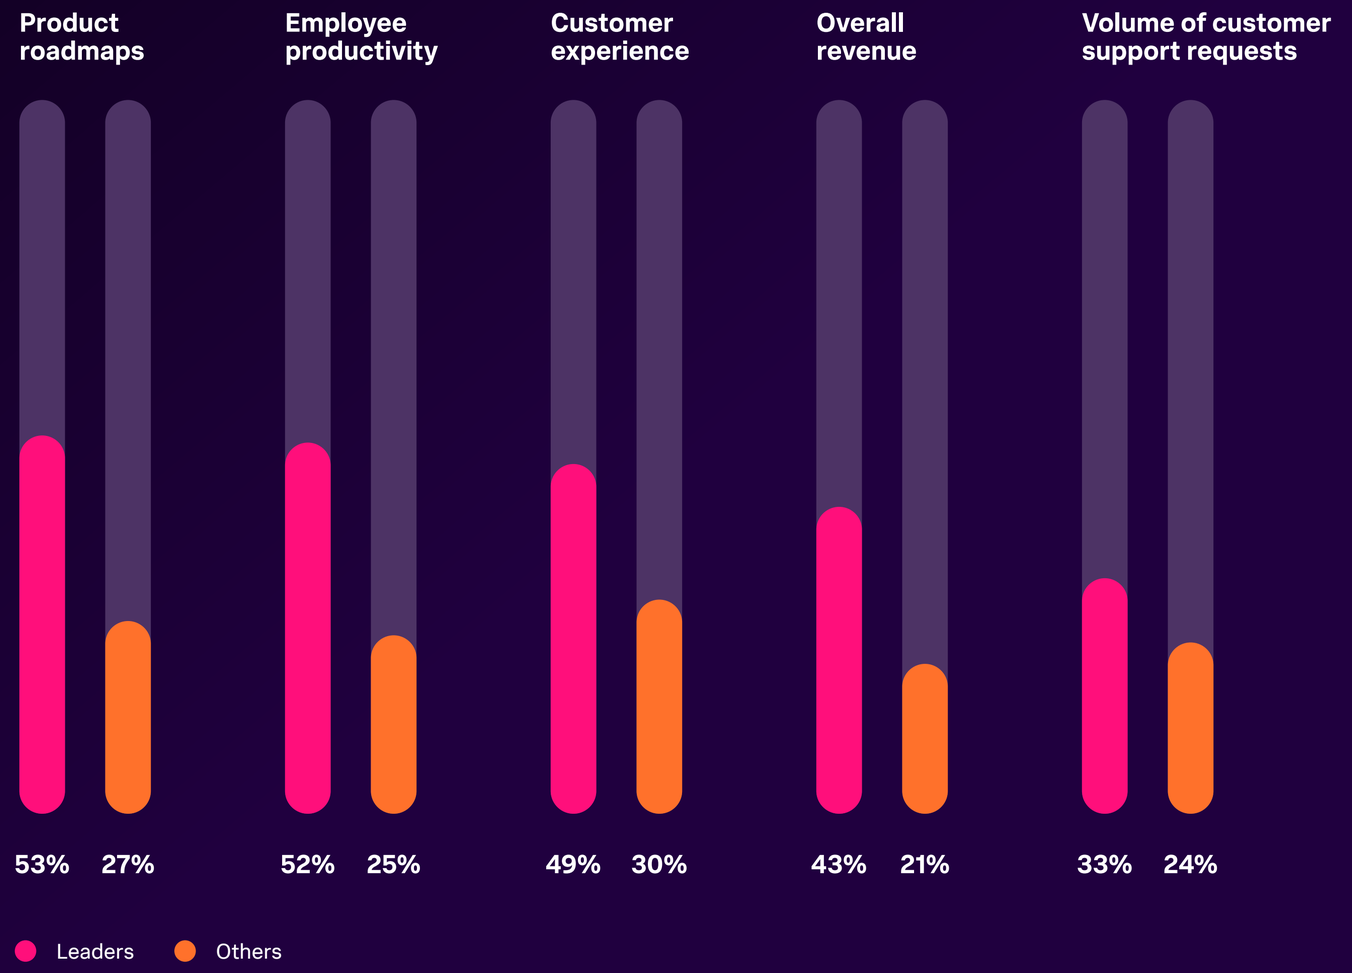

Code profiling and observability-as-code unlock better outcomes

We talked extensively about OpenTelemetry in chapter five — and with nearly three-quarters of respondents (72%) reporting it positively affects revenue growth, its value is clear. Now let's dig deeper into code profiling and observability-as-code.

Code profiling unlocks another level of granularity during troubleshooting by enabling teams to identify the problematic source code file (and more granularly, the call and its associated line of source code), so teams know which engineer to contact and how to fix the issue. A whopping 78% of leaders say code profiling helps find root causes faster to a *significant* or *transformative* extent.

Every minute of delay equals customer churn and revenue loss, so specificity is everything. "Without the ability to drill down into code performance issues, it's like a firefighter knowing there's a fire somewhere on the block but not knowing which house," says Leffler. "Code profiling gives you that clarity — it's like pinpointing the exact house, the exact floor, and even the room where the fire is burning."

Observability-as-code is a DevOps approach that treats observability configurations like code, which means teams can track changes, collaborate, and roll back these configurations using version control systems. It also enables teams to create dashboards, alerts, and other observability components using the same language and methods they use to create the applications themselves. It means that software engineering teams treat observability as a core part of the development process, not an afterthought. All of this translates to consistency, standardization, and scalability.

"Observability-as-code is one of the clearest signs of a mature observability practice," says Lin. "It shows that observability is baked into the development process, with telemetry collection and interpretation being treated with the same discipline as the rest of your code — making observability versioned, automated, and consistent."

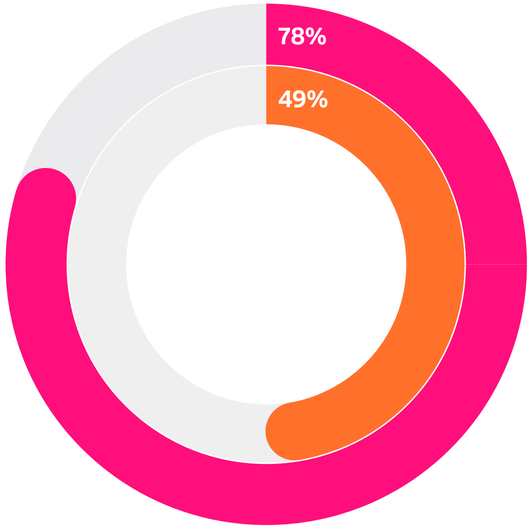

Leaders find a shortcut to root causes with code profiling

Increased speed is *significant* or *transformative*

Leaders Others

Without the ability to drill down into code performance issues, it's like a firefighter knowing there's a fire somewhere on the block but not knowing which house. Code profiling gives you that clarity — it's like pinpointing the exact house, the exact floor, and even the room where the fire is burning.

— Greg Leffler, Director of Developer Evangelism, Splunk

Cultivating a culture that drives the business

The tech — code profiling, observability-as-code, and OpenTelemetry — certainly helps, but let's be clear: its adoption is a *symptom* of maturity, not the cause. People are always behind the decision to invest in forward-looking tech. And that decision highlights progressive qualities, like an interest in innovation, a commitment to excellent digital experiences, a concerted effort to tune into the broader observability landscape, and the tenacity to learn and stay up-to-date on skills.

"To me, it shows that the organization has people who care deeply about the craft of observability, and that someone is setting the right culture," says Robin. "These are the teams willing to do the hard work of making observability part of the organization's cultural DNA — adopting new technology, researching and learning, nudging internal teams, and justifying the investment in both time and money."

Besides investing in tooling, let's take a closer look at how business catalysts actually drive better results.

Leading observability practices outpace their peers

Align closely with security

Leaders tend to collaborate more effectively with their organizations' security teams on things that matter most. They share and reuse data more than their peers (59% compared to 45%), but that's just the starting point for collaboration. Nearly half (44%) *strongly* agree that their ITOps and engineering teams troubleshoot and solve issues with their security teams, compared to 29% of other respondents.

This collaboration is likely easier due to the tooling they've put in place. OpenTelemetry, for example, gives observability and security teams a common language to work together, using shared signals and context. Only 16% of leaders say immature software is a roadblock for collaboration, compared to 35% of other respondents. A mere 7% of leaders report difficulty integrating highly dispersed or disconnected IT, engineering, and security platforms, which is three times less than their peers.

As a result of this tighter collaboration, leaders' telemetry data may have greater value across teams. Each type of observability data we asked about — metrics, events, traces, and logs — plays a more substantial role for security teams than their peers' data. In particular, leaders are 2.6 times more likely than others to say that traces *significantly* influence security decisions.

This is tangible evidence that leaders are breaking down silos. If security teams are actively using traces, it means data is not only being shared — it's being understood and embraced by teams beyond engineering.

Unlock the potential of AI

Can you have a state-of-the-art observability practice without AI? Probably not. Leaders ride the wave of AI innovation while their peers are still paddling. Sixty-four percent *always* or *often* use emerging AI technologies like agentic AI, compared to only 15% of their peers. They also adopt generative AI and AIOps at higher rates.

For leaders, fewer barriers block the path of AI innovation. Data quality isn't as much of an obstacle for AI readiness. About a third (34%) say low data quality is a challenge, compared with nearly half (49%) of others. Knowledge gaps are less severe, too; 25% say lack of expertise prevents AI readiness, compared to 41% of their peers who say the same.

Leaders also believe AI will improve essential operations. Forty-two percent expect AI to have a *significant* positive impact on monitoring critical business processes, compared to 19% of others who say the same.

Handle alerts and incidents with greater precision

Alerts are a source of angst for most organizations, but less so for leaders. Over half (52%) of other respondents say the volume of false alerts has a negative impact on their team morale, whereas only 35% of leaders say the same.

That's likely because leaders tend to have better alert hygiene and incident response processes. Thirty-seven percent say they *never* miss alerts, compared with only 15% of other respondents. They're 2.3 times more likely than their peers to say they *always* develop a detailed response plan when an incident affects customers. They're also more likely to *often* or *always* isolate an incident to a specific team and rely on that team to resolve it (43% vs. 22% of others), rather than unnecessarily involving multiple teams and burning out staff.

How to become a business catalyst

State-of-the-art observability practices are drivers for higher revenue, better customer experiences, and a slew of other meaningful goals. The value of their application data reverberates throughout the business, not just within the observability practice.

With the survey findings in mind, here's some advice on using your observability practice to spark results.

1 Limit war rooms and reactivity

Panicking is rarely the best way to respond to a customer-facing incident, yet 21% of respondents say they do this *sometimes*, *often*, or *always*. And 20% say they *often* or *always* start a war room that involves members of many teams until the issue is resolved, which brings productivity to a grinding halt and strains company resources. Here are some ways leaders avoid this:

- **Isolate the incident to a specific team.** Being able to quickly trace whether an issue is security- or observability-related prevents multiple teams from going down separate rabbit holes. Ideally, ITOps, engineering, and security teams troubleshoot in parallel and share context and insights to pin down the root cause, then hand the incident off to the appropriate team.
- **Make post-incident reviews a habit.** Not only do post-incident reviews help teams learn from prior wins and mistakes, but they boost morale by reassuring teams that history won't repeat itself. Build post-incident reviews into the incident response process so they become the standard, and make sure to treat them as living documents to accommodate policy, tool, or plan changes down the road.

2 Get a handle on alerting

False alerts are one of the top sources of stress for ITOps and engineering teams. And 54% of respondents say their alert quality has the greatest impact on observability ROI — so getting alerting under control will pay dividends.

- **Level up fine-tuning with adaptive thresholding.** Fine-tune thresholds to suit the criticality of the system or service you're monitoring by filtering out noisy false positives and making sure every alert is valid. Adaptive thresholding can take this to the next level by dynamically adjusting baselines based on historical data.
- **Only suppress alerts for *really* good reasons.** Use alert suppression very minimally, if at all — and ideally, have your continuous deployment (CD) system handle it on your behalf. Thoughtfully deliberate the decision to suppress alerts, and base it on a specific reason — for example, for an upcoming deployment or planned maintenance, not when you're experiencing a traffic spike.

3 Set standards for data quality to reap AI benefits

A substantial 78% of respondents say that AI has enabled them to spend more time focusing on innovation than maintenance; yet nearly half (48%) say that poor data quality prevents them from achieving AI readiness.

- **Clear up data quality ownership.** In many organizations, responsibility for application telemetry quality can be murky. Oftentimes, ownership defaults to platform engineering, regardless of the team's bandwidth or expertise. Identify a group of individuals most passionate about the craft of observability, and empower them to develop and drive a set of data quality standards throughout the organization, looping in stakeholders like the compliance team, to meet everyone's needs. Framing this with the advanced capabilities that higher-quality data will unlock can motivate teams to commit to this practice.
- **Inject business context and tags.** One way to enrich the insights that AI (and your engineers) can draw is to use tags with relevant business data. This might include the application that emitted it, version number, environment, or logged-in user. With this added context, teams can uncover patterns tied to business impact — such as whether an issue affects VIP customers — and prioritize alerting and response accordingly.

4 Dip a toe in the water of forward-looking tech

A cohort of respondents we call 'business catalysts' are twice as likely as their peers to say that their observability practice *significantly* improves overall revenue. What did they have in common? A commitment to forward-looking technology; they *often* or *always* used OpenTelemetry, code profiling, and observability-as-code.

- **Start with your biggest bottleneck.** Adopting all three of these technologies simultaneously would be a massive undertaking. Identify your priorities and decide from there. Is your observability practice slow to pinpoint issues? Code profiling might be a good starting point. Or, if your team struggles with collecting data in a consistent format, then OpenTelemetry sounds like the right solution.
- **Host regular knowledge-sharing sessions.** Once your internal champions have had the opportunity to learn, host monthly or quarterly technical forums where they can share updates on observability tools, frameworks, and advancements with the rest of the team. These sessions serve as dedicated spaces for discussing new tooling and addressing any implementation challenges, which helps to facilitate widespread adoption.

Continue your journey to becoming an observability leader

Perspectives by Splunk — by leaders, for leaders

Looking for more insights on observability trends? Learn how leaders tackle today's most pressing challenges including AI, data management, and developer innovation.

Learn more

The New Rules of Data Management: Creating value in the AI era

Discover how to tame data volume and complexity to drive better cybersecurity and observability outcomes with the new rules of data management.

Get the report

Industry highlights

We identified key insights across four select industries worldwide.

Financial services

Financial services organizations report strong links between their observability practices and business impact. Over three quarters (77%) say observability has a positive impact on revenue — well above the 65% average — and 75% say it influences their product roadmaps. This aligns with their priorities for observability capabilities, as 40% say that monitoring critical business processes is *very* important to the overall business.

Overall, the industry is enthusiastic about AI — 46% express excitement about AI's potential (vs. 36% overall) — yet they also recognize the challenges it will bring. More than half (54%) say monitoring AI workloads has made their jobs harder, compared to the 47% average.

On the tooling front, the industry has high OpenTelemetry adoption, with 36% reporting that they use it *often* or *always* (compared to the 26% average), and is more likely to reap the rewards of OpenTelemetry. Three-quarters (75%) who use OpenTelemetry at least *sometimes* say that the technology has had a positive impact on revenue.

However, true collaboration between observability and security teams is limited in the financial services sector. Only 59% say ITOps, engineering, and security teams share the same tools (vs. 68% overall), and just 61% report sharing data across teams (vs. 74%). In a compliance-driven industry, it tracks that 60% point to regulatory restrictions as the top barrier here.

Perhaps these silos contribute to financial services teams' high stress during incidents: 12% say they *often* or *always* panic during customer-impacting events, compared to 9% overall.

Get the financial services action guide.

Manufacturing

Observability practices in manufacturing organizations have a strong influence on the business, particularly on employee productivity; 86% of respondents say their observability practice improves this area, compared to 74% across all industries.

ITOps, engineering, and security teams are far more likely to work together in the manufacturing sector. A whopping 97% say they share and reuse data, and 81% troubleshoot and solve issues with their security teams.

AI plays a significant role in observability practices at manufacturing organizations. Nearly half (48%) express enthusiasm about AI's benefits to their team. Manufacturing respondents tend to face fewer AI readiness challenges, with only 35% citing lack of data quality as an obstacle, compared to 48% overall.

Manufacturers aren't just optimistic about AI, they're using AI in its most advanced forms. A significant 45% report using emerging AI *often* or *always*, compared to just 18% overall. Nearly all respondents (94%) say AI has allowed them to spend more time innovating rather than maintaining systems. Perhaps that innovation time is dedicated to software, as only 39% of manufacturing organizations say they spend less time than they should on building new software, compared to 45% overall.

Manufacturing teams are leaning into advanced observability tooling. They lead in the adoption of many technologies, saying that they *often* or *always* use automated remediation (43%), code profiling (41%), and observability-as-code (39%). These investments, paired with high collaboration and AI maturity, position the manufacturing sector as a forward-thinking observability leader.

Get the manufacturing action guide.

Public sector

Public sector agencies are currently exploring how observability practices can more directly influence their mission outcomes. Compared to other industries, respondents from the public sector are significantly less likely to report their observability practice has a positive influence on budget (30% vs. 65%), product roadmaps (30% vs. 64%), and employee productivity (36% vs. 74%).

For public sector respondents, ROI is most closely tied to operational efficiency, particularly alerting: 69% cite the quality of alert detections as one of the biggest drivers of ROI, well above the 54% average. This aligns with their top source of stress — high volumes of false alerts, cited by 61%.

Collaboration presents a significant area for development within the public sector. Just 46% of public sector teams say they reuse and share observability data, and only 35% report cross-team troubleshooting with security — the lowest of any sector. Contributing to this are major talent and infrastructure gaps: 62% cite a lack of relevant skill sets, and 60% report low technology maturity as obstacles, both significantly above average.

Observability practices within the public sector are in a developing phase, so it makes sense that they're not positioned to take advantage of forward-leaning technologies. Only 35% say they use AIOps *often* or *always* (vs. 54%), and just 10% say they use generative AI *often* or *always*, compared to 39% overall. Only 8% *often* or *always* use observability-as-code (vs. 29%), code profiling (2% vs. 21%), and OpenTelemetry (2% vs. 26%).

Get the public sector action guide.

Communications and media

Communications and media organizations are among the most advanced in their observability practices — and they're seeing outsized business benefits as a result. A striking 88% report a positive impact on overall revenue from their observability practice, compared to 65% across all industries, and 81% say it positively influences their product roadmap (vs. 64%).

Speed is paramount for this sector. Communications and media respondents are most likely to cite incident troubleshooting speed as one of the biggest drivers of observability ROI (68% vs. 49% overall). They also place a high importance on AI, with 51% citing maturity of their AI capabilities as having a major impact on their observability ROI.

Communications and media teams are leaders in AI adoption: 79% *often* or *always* use AIOps, and 68% *often* or *always* use generative AI — both well above average. However, data remains a hurdle, with 56% citing data quality as a barrier to AI readiness, and 69% pointing to data challenges, such as data accessibility, quality, and fragmentation, as their top source of stress.

Despite this, teams in this sector are able to stay focused. Only 27% say they spend too much time responding to alerts (vs. 43% overall), and 73% *rarely* or *never* miss alerts — outperforming the 60% average. Yet 69% also report spending less time than they'd like building new software, which suggests competing priorities persist.

Communications and media organizations are leading adopters of OpenTelemetry, with 67% using it *often* or *always* — more than double the industry average. And it's paying off: 86% say OpenTelemetry contributes to revenue growth, and 83% cite a positive impact on customer satisfaction.

Get the communication and media action guide.

Country highlights

Snapshots from nine countries across the globe.

Australia

Australia stands out as a leader in both AI adoption and advanced observability practices, with strong signs that these investments are already delivering measurable business benefits. Forty-five percent of Australian respondents say they're enthusiastic about the benefits AI can bring to their teams, notably higher than the global average of 36%. And that optimism is backed by action: Australian organizations report higher adoption across all three categories of AI technology, with 21% saying they *often* or *always* use emerging AI (such as agentic AI), compared to just 18% globally.

This forward-thinking approach is paying dividends. A striking 87% of Australian respondents say AI has enabled them to spend more time on innovation rather than maintenance (vs. 78% globally). This may explain why only 37% say they're spending less time than they'd like building new software — well below the 45% average — indicating that engineers in Australia are more likely to have the bandwidth to focus on high-value work.

Australian respondents have higher expectations for AI's impact on observability. Seventy-two percent expect AI to improve troubleshooting and root cause analysis — a full 12% above the global average.

Australian organizations are also using OpenTelemetry at higher rates, with 36% saying they *often* or *always* use it (vs. 26% globally). Importantly, that usage is translating into tangible business results: 79% of those using OpenTelemetry at least *sometimes* say it positively affects revenue growth, compared to 71% across all regions.

France

In France, data challenges loom large — both as a barrier to AI readiness and as a source of day-to-day stress for observability teams. Fifty-eight percent of French respondents cite data issues, such as accessibility and quality, as the top factor negatively impacting their team's stress levels. It's no surprise, then, that 51% also point to poor data quality as their biggest obstacle to adopting AI.

Despite these hurdles, French organizations show signs of operational discipline, particularly in alert management and incident response. Just 8% say they *often* or *always* miss alerts — well below the 13% global average — suggesting strong alert hygiene practices. That likely contributes to their lower incident-related stress: only 4% say they *often* or *always* panic during customer-impacting incidents, compared to 9% globally.

A focus on incident resolution speed further supports this. Fifty-five percent of French respondents say that the speed of incident troubleshooting has the greatest impact on observability ROI — signaling that fast, effective response is a top priority.

To support that speed, French organizations are embracing forward-leaning tooling. Notably, code profiling is more widely adopted in France than in many other countries, with 30% saying they *often* or *always* use it, compared to just 21% globally. And this investment is strategic; 43% of those who at least *sometimes* use code profiling believe it will enhance the effectiveness of their AI capabilities.

Germany

Germany's observability practices are driving strong business results, with 74% of respondents reporting a positive impact on overall revenue — notably higher than the global average of 65%. This performance reflects not only technical maturity, but also a collaborative approach to problem-solving that brings teams together. Sixty-two percent say their observability and security teams troubleshoot and resolve issues together.

Incident response practices, however, reveal a more nuanced picture. On one hand, 74% of German teams *often* or *always* conduct detailed post-incident reviews — slightly above the 71% global average — reflecting a commitment to continuous learning and improvement. On the other hand, 28% say they *often* or *always* launch a war room during customer-impacting incidents, significantly higher than the global average of 20%.

German teams are also leading adopters of OpenTelemetry, with 32% reporting that they use it *often* or *always*. The impact is clear: 79% using OpenTelemetry at least *sometimes* say it positively affects revenue growth, compared to 72% globally.

India

Teams in India are highly collaborative with their security counterparts — 81% say they share and reuse data with security teams, compared to 74% globally. More importantly, 74% say they can accurately trace application and infrastructure performance issues back to security root causes, outpacing the global average of 65%. This suggests not just surface-level collaboration, but meaningful technical alignment across functions.

Still, collaboration doesn't come without its challenges. Over half (53%) of Indian respondents cite regulatory restrictions as the primary barrier to improving collaboration — the highest-ranked obstacle in the country.

AI adoption is another bright spot. Eighty-two percent of Indian respondents say AI has allowed them to spend more time on innovation rather than maintenance, slightly above the 78% global average. Only 36% of respondents in India say they spend more time than they should responding to alerts, compared to 43% globally, which may indicate that AI is beginning to ease some of the operational burden.

Alerts play a critical role in shaping security strategy, with 58% of Indian respondents saying that alerts *significantly* influence security decisions, compared to just 47% globally. Additionally, 62% say that the quality of alert detections has one of the greatest impacts on observability ROI. Yet, alerting isn't without its pain points. Fifty-five percent say the volume of false alerts negatively affects team morale.

Japan

In Japan, observability practices take a cautious but optimistic approach to emerging technologies, especially AI. AI adoption remains slightly below the global average, with 48% of respondents saying they *often* or *always* use AIOps (compared to 54% globally), and only 9% saying they *often* or *always* use emerging AI such as agentic AI (vs. 18% globally).

The top challenge impeding broader adoption appears to be data quality, with 47% of Japanese respondents citing it as the primary barrier to achieving AI readiness. Additionally, 53% say that monitoring AI workloads has made their jobs harder, compared to 47% globally. Despite these hurdles, there is a strong sense of AI's potential. Sixty-two percent believe AI will positively impact the monitoring of critical business processes, slightly above the global average.

Tool sprawl is another major issue facing observability teams in Japan. Sixty-five percent say that the proliferation of disconnected tools negatively impacts team morale — the most commonly encountered morale challenge in the country, and above the 59% global average. This fragmentation may also be contributing to alert fatigue and visibility gaps, with 15% of respondents saying they *often* or *always* miss alerts.

New Zealand

New Zealand's observability practices stand out for their clear, measurable impact on customer experience and business alignment. A striking 82% of respondents report that their observability efforts positively impact customer experience — significantly above the global average of 69%. That success likely stems from a strong organizational focus on the customer journey; nearly half (48%) say understanding critical user journeys is *very* important to their overall business strategy, compared to just 25% globally.

Collaboration between security and observability teams is another strength in New Zealand. Ninety percent of respondents say they can accurately trace application or infrastructure performance issues to security-related root causes, far surpassing the 65% global average. And that collaboration appears to be paying off: 74% say cross-functional teamwork leads to fewer customer-impacting incidents, compared to 64% globally.

New Zealand respondents are also leaning into AI. Forty-four percent are enthusiastic about the benefits AI provides their teams, and 38% say they *often* or *always* use emerging AI technologies like agentic AI — more than double the global average. However, the biggest obstacle to expanding AI readiness isn't data quality, but talent: 50% cite a lack of expertise or understanding across teams as their top barrier, compared to 40% globally.

Despite these strengths, alert fatigue remains a challenge. Fifty-two percent of respondents say they spend more time than they should responding to alerts, suggesting that even advanced teams still face friction in this area.

Singapore

In Singapore, observability practices are evolving rapidly, with a clear focus on speed and efficiency. Sixty-four percent of Singaporean respondents say the speed of incident troubleshooting is a top driver of observability ROI, compared to just 49% globally. This priority reflects a fast-moving IT culture where rapid response is essential.

Singaporean respondents are investing heavily in AI-driven solutions, possibly to support this need for speed. A higher-than-average 61% say they *often* or *always* use AIOps in their observability workflows, and a remarkable 85% report that they regularly use AI as part of their day-to-day work — nearly 10% above the global average. These numbers suggest that AI is not just being explored, but actively embedded into operational practices.

However, the road to operational efficiency isn't without obstacles. Tool sprawl is a significant challenge in Singapore, with 65% of respondents saying it negatively impacts team morale. But the biggest drain on morale isn't the number of tools — it's the volume of false alerts. Half of respondents (50%) say they spend more time than they should responding to alerts, a clear sign that teams are struggling to separate signal from noise even as they adopt sophisticated technologies.

U.K.

In the United Kingdom, observability is a strong enabler of productivity, with 75% of respondents reporting that their observability practices have a positive impact on employee efficiency. This suggests that teams in the U.K. are successfully using observability data to reduce friction, streamline workflows, and empower teams to focus on higher-value work.

The U.K. also reflects a cautiously optimistic stance on AI. Thirty-nine percent of respondents say they're enthusiastic about AI's benefits — slightly above the global average of 36% — and nearly half (48%) describe themselves as optimistic but still seeking more information before fully embracing AI within their teams. Troubleshooting and root cause analysis are seen as key areas where AI could make a difference; 60% expect AI to positively impact these processes.

Still, challenges remain, especially around alerting. Over half (54%) say that the volume of false alerts is negatively affecting their team's stress levels, and when asked to identify the top challenge impacting morale, false alerts topped the list. This noise may be contributing to less-than-ideal alert hygiene, with 15% saying they *often* or *always* ignore or suppress alerts — slightly higher than the global average.

While adoption of OpenTelemetry in the U.K. is in line with the global average (26% *often* or *always* use it today), its impact is particularly pronounced when it comes to brand perception. Over three-quarters (76%) who are using OpenTelemetry at least *sometimes* say OpenTelemetry has improved how their brand is viewed, highlighting the strategic value of modern observability tooling beyond technical outcomes.

U.S.

Observability practices in the United States largely align with global averages across key metrics, but several areas reveal distinct patterns, particularly in how observability data supports security teams and how organizations handle incident response.

Security appears to benefit strongly from observability data in the U.S., with 54% of respondents saying alerts *significantly* influence security decisions — a notable increase over the global average of 47%. Additionally, U.S. respondents are optimistic about the potential of AI to further improve this area; 65% believe AI will have a positive impact on detecting application vulnerabilities and threats, compared to 58% globally. These figures suggest a growing alignment between observability, AI, and cybersecurity functions.

However, alert hygiene presents a challenge in the U.S. Fifteen percent of U.S. respondents say they *often* or *always* miss alerts (vs. 13% globally), and 16% report they *often* or *always* experience outages due to missed alerts — higher than the global average of 11%. Perhaps these alerting challenges play a role in heightened panic during incident response: 12% of U.S. respondents admit they *often* or *always* panic during customer-facing incidents, compared to just 9% globally.

Methodology

Oxford Economics researchers surveyed 1,855 ITOps and engineering professionals from practitioners to VP-level executives (including developers, SREs, systems engineers, infrastructure operations professionals, CTOs, and CIOs) from February through March 2025. Respondents resided in Australia, France, Germany, India, Japan, New Zealand, Singapore, the United Kingdom, and the United States. Respondents represented 16 industries: business services, construction and engineering, consumer packaged goods, education, financial services, government (federal/national, state, and local), healthcare, life sciences, manufacturing, technology, media, oil/gas, retail/wholesale, telecom, transportation/logistics, and utilities.

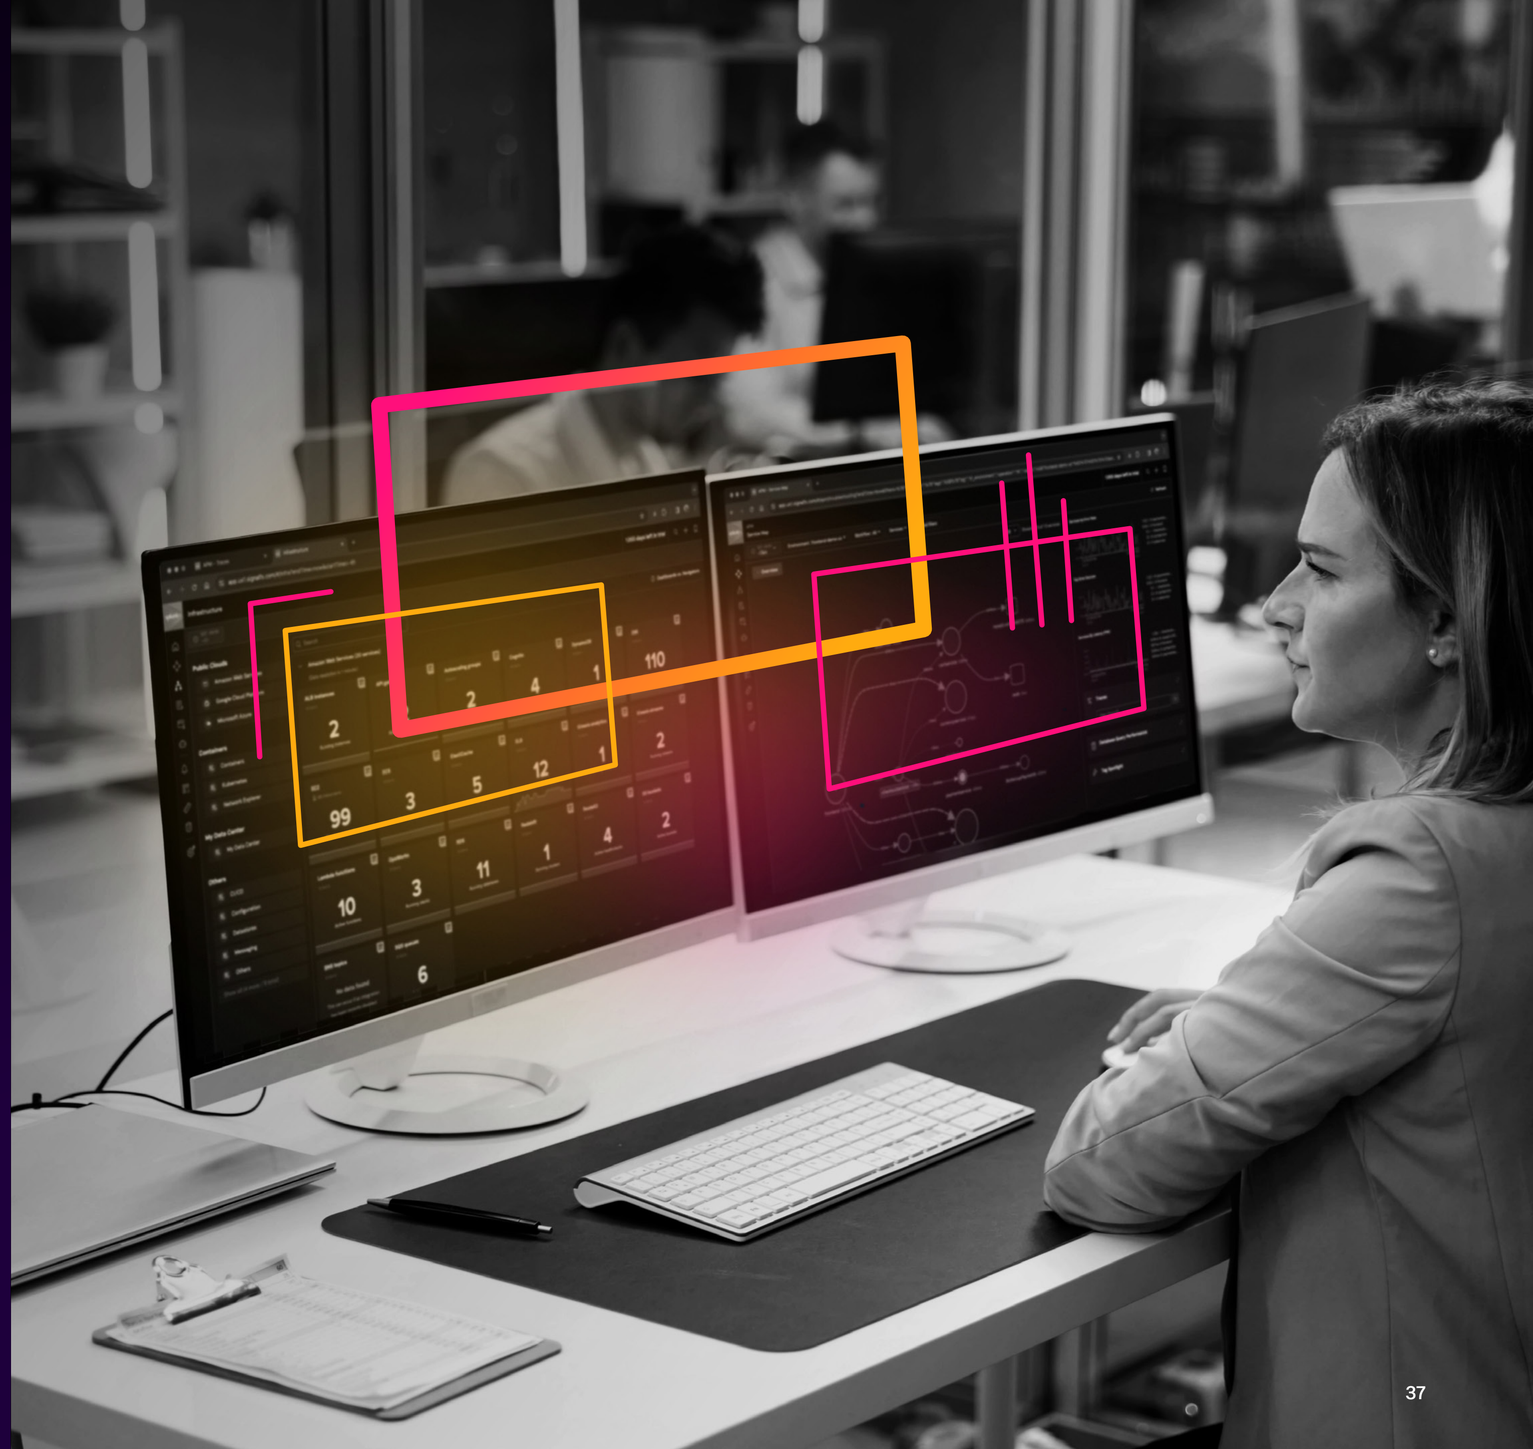

About Splunk

Splunk, a Cisco company, helps make organizations more digitally resilient. Leading organizations use our unified security and observability platform to keep their digital systems secure and reliable. Organizations trust Splunk to prevent infrastructure, application, and security incidents from becoming major issues, recover faster from shocks to digital systems, and adapt quickly to new opportunities.

Keep the conversation going with Splunk.

splunk>
a cisco company

25_CMP_report_state-of-observability-2025_v15