

2024

State of Observability

Charting the course to success

splunk>
a CISCO company

Your engineering team finally ramps a brand-new feature for your company's software application. After months of planning, coding, and testing, you're all hoping for a successful release with no bugs or major issues.

The new update ships out, and you all wait. But ... nothing bad happens.

No frantic messages in a temporary Slack channel. No late-night runs to the corner store for energy drinks. No calls from angry customers. Just a team of individuals working in peaceful harmony. (Cue the sound of crickets chirping.)

Instead, a sense of calm washes over the office. Your SREs get to work on a pet automation project. Developers pop the next story from the backlog and start writing new features. Platform engineers put the finishing touches on some documentation. The coffee drinkers refill their mugs.

This solace is no accident; it's the power of observability at work. Observability unlocks these delightfully boring experiences by keeping the lights on, acting as a foundation for every digital success. And when teams don't need to worry about their systems failing catastrophically, they can focus on innovating, building, and improving resilience.

Contents

- 4 The view is worth the climb
- 6 What it means to build a leading observability practice
- 13 Creating a flexible future with telemetry data
- 17 Platform engineering ushers in a new DevOps future
- 21 AI in observability starts to take shape
- 24 Earn your spot on the observability leaderboard
- 27 Become an observability leader with Splunk

It may sound like a dream, but for an elite group of IT operations and engineering professionals we surveyed, these successes are very much a reality.

Within a leading observability practice, innovation and resilience aren't mutually exclusive. Issues happen (they always will), but leading organizations find and fix those issues faster, even within minutes. Developers spend their time building and experimenting. They ship their products faster and with more confidence than a wing and prayer. They see the value of OpenTelemetry to maintain flexibility and ownership over their data, and AI to surface insights.

In last year's State of Observability, we declared that "observability has arrived." This year's data reveals that observability is the baseline. A leading observability practice sets organizations apart from their competition, enabling them to find and fix unknown unknowns and deliver incredible digital experiences to their end users.

Let's find out how they make it happen.

How leaders innovate

2.8x

faster at finding issues than beginning organizations



72%

more code shipped on demand than beginners



78%

adopt OpenTelemetry





The view is worth the climb

“There’s no future in which digital systems emit less telemetry data.”

Tom Casey, SVP, Products and Technology, Splunk

Observability isn't the new kid on the block anymore. Nearly half (47%) of survey respondents say they've used observability tools for two years or longer, up from 36% a year ago. And we won't wax poetic about its benefits. You already know: Less time troubleshooting in war rooms. Higher alert fidelity. Faster time to detection. Improved development velocity. Less downtime. The list goes on, and later in this report, we'll show you how leading organizations are reaping those benefits more than their peers.

But even leaders contend with complexity. It's inevitable.

In a field entirely centered around data, guess what the biggest challenges are? They're all data-related. Specifically, respondents' main concern is that they spend too much time correlating it across fragmented sources (30%). Data-related problems appear to be intensifying; the percentage of respondents reporting each challenge increased across the board in 2024.


Unpacking observability

When an organization has observability, it has visibility into everything — an arduous but admirable pursuit, considering the maze that is the modern technology stack. This includes network (owned and unowned), infrastructure (on-premises, cloud, and hybrid), homegrown and third-party applications, and digital customer experience.

To achieve observability, organizations need to overcome a few blockers:


- **Too much noise.** Observability isn't about knowing the success and failure of every component. In fact, information overload can give practitioners headaches, while creating ambiguity and even complacency that harkens back to the boy who cried wolf. How are you supposed to know when it's real? Over half (57%) of respondents say alert fatigue is either somewhat or highly problematic. These problems can range from burnout — 66% say critical staff members have left for this reason — to missed issues that turn into downtime.
- **Dispersed environments.** Where do an organization's assets live? Here, there, and everywhere, and according to this survey, only 24% can correlate all or almost all their data across their application environment. No surprise, considering the state of business application sprawl across public cloud, on-premises, and hybrid infrastructures. Seeing it all is an ongoing ordeal. Historically, gaining insights from network infrastructure not owned or operated by the organization has been difficult and it continues to be the case, with only 26% saying that visibility over these assets is excellent.
- **Too many tools.** Organizations deploy an average of 23 different observability tools, spanning both modern and legacy tech. That's astronomical, especially if organizations expect one person to toggle between all of them. And 23 tools is just the average; 8% of respondents admit to using over 50 tools. That's a lot of chair swiveling.

Our respondents recognize that observability is the path forward. Respondents say their observability solutions deliver value that is 2.42x their costs, and 86% say they'll increase spending for observability solutions over the next year. That sort of value makes observability a non-negotiable.



say alert fatigue is problematic

Business applications live here, there, and everywhere



What it means to build a leading observability practice

“Building a leading observability practice means being obsessed with delivering incredible digital experiences to your customers, and embedding that mindset into every decision you make.”

Patrick Lin, SVP and GM, Observability, Splunk

Sherlock Holmes once said, “It is my business to know what other people don’t know.”

A leading observability practice operates much like a detective, overseeing and understanding their organization’s entire digital footprint, from the big picture down to the nitty-gritty details. Observing every component of the environment is the only way to investigate mysteries and bumps in the night that will inevitably emerge as a business operates.

Solving those mysteries is the key not only to observability success, but also to a cascade of other benefits: employee happiness, developer productivity, customer satisfaction, and higher profits. A better business.

What does it mean to build a leading observability practice? It’s not just about deploying the right tech. Best-in-class teams have a desire, knowledge, and ability to excel in their observability endeavors.

This means creating a culture that cares about observability for the sake of creating excellent digital experiences, not to simply avoid the bad ones. It is a culture of not accepting things simply as they are, but pushing for what they could be. These teams use this desire to continually educate themselves about strategies to achieve observability and then put this knowledge into action through tools, training, and processes.

In other words, observability isn’t something you have. It’s something you *do*.



In organizations with immature observability practices, figuring out where an issue is coming from involves putting everyone in a war room. Putting hundreds of people on a conference bridge is a really bad way to solve problems, and leading teams are finding ways to rise above that.

Annette Sheppard, Director of Product Marketing for AIOps, Splunk

Path to a leading observability practice

Leading observability practices aren't born — they're built. We believe that process unfolds in four stages of increasing sophistication: foundational visibility, guided insights, proactive response, and unified workflows. To understand more about how leaders take their practices to the top, we asked respondents a series of questions aligned to those four stages.

Foundational Visibility

Can you see across any environment and all your stacks?

Leaders have excellent visibility across their entire environment, from on-premises infrastructure to private and public cloud, to containers and applications.

Guided Insights

Can you detect issues with context?

For leaders, alerts aren't another thing to wrangle and decipher; they're actionable insights. And they lean on AI and ML to help recommend solutions, determine root causes, correlate events, and even recommend solutions.

Proactive Response

Can you get ahead of issues?

Downtime happens even to the most resilient organizations. But leaders experience downtime less. They also resolve issues faster so downtime isn't as detrimental — and can even prevent similar issues from happening altogether in the future.

Unified Workflows

Is observability embedded into your organization?

Leaders rely on disciplines like platform engineering to standardize the way teams build, deliver, and operate software, unlocking greater developer productivity and operational efficiencies.

The value of visibility: Where leaders win

Our data reveals that leaders make headway over their peers in multiple areas.

Drive developer innovation and speed

Customer expectations are at an all-time high. Companies that launch innovative products quickly not only meet these demands but also outpace their competition. This doesn't happen without strong engineering teams.

At leading organizations, developer productivity and output drive profitability. Over three-quarters (76%) of leaders push the majority of code on demand, while only 30% of beginners say the same. When it comes to getting new functionality into the hands of users, 60% of leading organizations say they're typically on the leading edge — 8.6x the rate of beginning organizations.

Leading teams hustle, but they push code thoughtfully and with few mistakes. Engineering teams at leading organizations achieve a 22% higher change success rate (a key DORA metric) for production application code. What's more, the majority of leaders say these changes are successful 90% of the time or more.

Developers want to build and experiment — and not just for their personal benefit. Post-it notes, microwaves, penicillin, and Gmail were all born out of good ol' tinkering around. At leading organizations, that's precisely what developers do. Compared to their peers at beginning organizations, developers in leading organizations spend about 38% more time on innovation versus routine tasks like maintenance, alert handling, and configuration.

Find and fix issues faster

Speed not only gives leading organizations an edge in development, but also strengthens resilience that helps them minimize downtime. Sixty-eight percent of respondents at leading organizations say their teams are aware of application problems within minutes or seconds of an outage or slowed performance, 2.8x the rate of beginning organizations. And leaders are only improving — 57% say they are finding the root cause significantly faster than a year ago.

Confidence and speed go hand in hand. Leading organizations estimate that upwards of 80% of alerts are legitimate. Their teams know with certainty that an alert isn't a false alarm. Contrast this scenario to what goes on at beginning organizations: They report just over half (54%) of alerts are tied to a real issue. That's essentially a coin flip. Heads it's real, tails it's not. That leaves a lot of room for conjecture, suspicion, and useless fire drills. After a while, engineers may start to just ignore alerts altogether.

Prompt detection leads to prompt resolution. Leading organizations are 2.3x more likely than beginning organizations to measure their mean time to resolve (MTTR) in minutes or just a few hours, while beginning organizations are 2.4x more likely to measure their MTTR in days, weeks, or even months — too much time to be in the dark as the incident reverberates through an organization. The difference between hours and days is a particularly wide chasm, especially considering that downtime costs organizations \$540,000 per *hour*, according to [The Hidden Cost of Downtime](#).

Leaders push the majority of their code on demand **2.6x more** than beginners.

Leaders are aware of application problems **2.8x faster** than beginners.

Embrace and harness AI

Leaders are fully welcoming AI for observability in all its forms: ML, AIOps, and generative AI. They're adopting these technologies at significantly higher rates than their peers, with 64% extensively using AIOps in their toolsets — over 10x the rate of beginning organizations. Sixty-five percent of leaders lean on AIOps to pinpoint and remediate the root cause of incidents with greater intelligence and automation.

Leaders aren't just adopting AI more; they're also seeing greater benefits from it. Forty percent agree the ROI of AIOps tools far exceeded expectations, compared to only 6% of beginners who agree.

Generative AI is an emerging realm for observability, and leaders are blazing the trail. One-third are using generative AI within observability tools, compared to just 6% of beginners. They're also more likely to appreciate how generative AI-enabled chatbots can strengthen an observability practice, specifically with data analysis (85%) and recommendations to resolve issues (81%).

Prioritize telemetry pipeline management

Data is the foundation of a top-tier observability practice. A sound telemetry pipeline management strategy is a competitive differentiator, especially as telemetry data mounts, compliance rules shift, and use cases multiply.

Leading observability practices are more likely to recognize that data management techniques like tiering (57%) and aggregation (55%) are “critical” to control costs. Perhaps these organizations release more products, maintain more applications, and store more data, causing leaders the pain of ballooning costs faster than their peers. After all, nothing delivers a reality check quite like an unexpected, exorbitant service bill.

Leaders adopt AIOps at **over 10x** the rate of beginners.

Leaders are **2x more likely** than beginners to say that data tiering is critical to control costs.

Share data and resources

SecOps, ITOps, and engineering teams need the same context to root out issues. An alert about a network traffic spike, for example, may indicate a flash sale, a new feature, or a DDoS attack in progress. Without a shared view of data, different teams may spin their wheels and fruitlessly point fingers in the absence of full context.

Breaking down silos and sharing dashboards can answer common questions — “Whose problem is this? Is this a security incident or an application performance problem?” — and get to the root of the problem faster. Nearly three-quarters of leaders (73%) improved MTTR (compared to only 39% of beginners who say the same) thanks to bringing security and observability tools and workflows together.

“Ideally, someone from security should be in every war room,” says Cory Minton, field CTO at Splunk. “Whether organizations want C-level visibility with a unified platform or just a basic understanding between security and observability teams, not being able to share a common language is a problem.”

Organizations working towards leading observability practices will converge security and observability data and tools as part of the process, whether intentionally or unintentionally. That “sharing is caring” mentality bodes well in two key areas. Seventy percent of leaders share data and tools during troubleshooting, and 80% empower software engineers to routinely detect security vulnerabilities throughout the software development lifecycle (SDLC).

Get more value from observability

Less downtime, faster resolution time, more innovation — that’s all music to the ears of any organization. Leaders’ successes go straight to the bottom line, with an annual return that’s 2.67x their spend.

Leaders report greater benefits from observability solutions than their peers. Nearly all leaders agree their observability solution improves problem detection time (95%) — compared to only 67% of beginners — helping to bring systems back online and minimize downtime. And according to 92%, observability compresses application development time (versus 64% of beginners), which means new products and experiences are in market faster.

73% of leaders improved MTTR when they converged observability and security tools and workflows.

Leaders’ annual return on observability solutions is 2.67x their spend.


Being an observability leader pays off

Leaders experience greater success in the following areas.


- Beginners
- Emerging
- Evolving
- Leaders



Percentage of developer time spent innovating



Aware of a problem in seconds



Use AIOps tools extensively



Improve MTTR when converging observability and security workflows



Creating a **flexible future** with telemetry data

“Engineers haven’t had control of how their observability data is collected, transformed, and used. OpenTelemetry addresses that need for flexibility in an open-source, vendor-neutral way.”

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

Flexibility is the future of observability. Thorny data residency requirements, new data sources, and an expanding array of tools means that your data use cases will only grow. Unlocking those use cases and uncovering the true value in telemetry data lies in how effectively you can control and use it. How you capture data, where you store it, and how you transform it should all be up to you. Let's dive into the tools that will help you control your data destiny.

OpenTelemetry emerges as the new standard

OpenTelemetry (OTel) — the industry-standard, open-source approach to collecting observability data — might not be the magic wand that cures all data woes, but it is capable of some serious data wizardry. Supported by the Cloud Native Computing Foundation (CNCF), it provides APIs and SDKs in over a dozen languages to help developers and engineering teams take complete control of their data, allowing them to skip vendor lock-in and proprietary agents.

We'll say it loud and proud: OpenTelemetry is the future of observability. OTel is quickly becoming recognized as *the* way to collect telemetry data, with over half (58%) of respondents saying their primary observability solution relies on OpenTelemetry. Since OpenTelemetry support is widely considered a sign of innovation, it's no surprise that leaders are more advanced in its adoption, with 78% embracing the platform.

Leaders are also more likely to recognize the benefits OpenTelemetry affords. Better control and ownership of data is up there at 65%, but one benefit stood above the rest: access to a broader ecosystem of technologies, which, according to 72%, is their top reason for using OpenTelemetry. OpenTelemetry offers unparalleled flexibility and customization. It provides libraries for nearly every common programming language, enabling developers to instrument their applications regardless of the tech stack they use. And since it can integrate with an extensive list of frameworks and libraries, adding observability to existing applications is a simpler process.

Leaders reap the rewards of OpenTelemetry

● Leaders ● Beginners

Access a broader technology ecosystem



Meet data residency requirements



Adopt modern cloud frameworks more easily



Lower observability costs



Get better control and ownership of data



Avoid vendor lock-in



Unlocking the potential of OpenTelemetry

Technology as powerful as OpenTelemetry usually isn't simple, but its quirks are worth conquering. Early adopters understand the end value is worth it. This may be why leaders experience the challenges of OpenTelemetry more acutely than their peers. Forty-four percent of leaders describe implementing OpenTelemetry as "very challenging" for their organization, compared to 21% of beginners who feel the same way. The learning curve, or lack of staff familiar with OpenTelemetry, is the biggest pain point for leaders at 55%. But even if it's hard now, it won't be forever as OpenTelemetry training becomes more available.

Other stumbling blocks include a lack of support for particular infrastructure or frameworks, and over half (54%) of leaders agree. But this is likely to resolve itself as more vendors recognize the demand for OpenTelemetry support. With only five years since its inception, there's a lot of opportunity for growth. Plenty of new products and frameworks now ship natively with OpenTelemetry support. And the addition of [support for profiling](#) as a new signal type, giving code-level insights into resource utilization in a vendor-neutral way, shows that OpenTelemetry is not done innovating yet.

Those who find it easier to implement OpenTelemetry claim its built-in support is why: for programming languages, for cloud-native technologies, and from the community of practitioners. This abundance of help can shrink the learning curve: Practitioners are unlikely to be alone in their confusion because someone out there will have an answer.



An open-source project is always evolving. Right now is the best time to dive into OpenTelemetry.

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

Taming telemetry pipelines

Data tells a story. Logs, metrics, and traces are all instrumental in investigating incidents because they tell the who, what, when, where, and why. Since the advent of cloud, microservices, and now AI, observability teams are grappling with a groundswell of telemetry data that makes it difficult to uncover what that story may be.

Navigating this data labyrinth demands not just sophisticated tools, but a strategy centered around data ownership. As data volumes and costs continue to climb, data management techniques are emerging as a route to more granular control over telemetry data. Leaders recognize this more than their peers:


- **Transformation and redaction.** Transforming data — for example, changing a 900-character log line and associated metadata into a 50-character timestamp and metric — can reduce volume, make it more cost-effective to store, and simplify analysis. Ninety-one percent say transformation is important to controlling observability data costs, with leaders more often specifying that it's "critical."
- **Data tiering.** Not all data is created equal, and understanding the best location depending on how often you'll access it — like cold storage for audit data, for example — helps organizations avoid paying for expensive storage when it's unnecessary. Ninety percent agree with this approach, describing data tiering as an "important" way to control costs.
- **Aggregation.** Aggregation — collapsing multiple data points into fewer using statistical methods — combines data from diverse sources and transforms it into actionable information, revealing patterns and insights. It also reduces storage needs, which is likely why 92% agree that it's "important" or "critical" to cost control.

These data management capabilities are packaged as dedicated point solutions or within a broader observability suite, giving organizations the flexibility to decide the best way to consume them.

Leaders have a clear preference here: Less is more. Best-in-class observability teams suffer tool sprawl more intensely than their peers and tend to prefer a single suite for all three capabilities: data tiering (63%), data transformation (62%), and aggregation (61%).

"Telemetry pipeline management is a set of capabilities, not a product," says Tom Casey, SVP of products and technology at Splunk. "When pipeline capabilities are baked into a product that's already integrated within your toolset, organizations can streamline operations and get deeper insights without the overhead of managing separate solutions."

Adding another tool to the stack is like playing Jenga; it requires finesse to avoid toppling the tower. It involves training staff on a new system, creating a succession plan, and possibly investing in certification programs. Introducing a dedicated point solution can also slow incident response as teams swivel to yet another dashboard.



Platform engineering ushers in a new DevOps future

“When you have a platform engineering division, you attract top talent eager to be on the vanguard. The challenge is convincing your CTO that it’s worth it to create a whole new team.”

Annette Sheppard, Director of Product Marketing for AIOps, Splunk

The platform engineering revolution is here. Nearly three-quarters (73%) of respondents practice platform engineering, either extensively or for select projects, and another 20% will implement it over the next year.

This comes at a welcome time for overextended ITOps and engineering teams. While some may get a thrill from the pressure of troubleshooting, others consider it a genuine reason to find a new job. Sixty-six percent of respondents say critical staff left due to burnout in the past year, and 70% have been short-staffed.

Demystifying platform engineering

At its core, platform engineering is an approach to building toolchains, workflows, and self-service platforms for software engineers so that they can spend more time building cool and reliable software, and less time managing their tools.

Despite what you may have heard, platform engineering is not:

- Another term for a site reliability engineer.
- The death of DevOps.
- A silver bullet.

Here is what a typical day in the life of a platform engineer might really look like:

- Defining how the organization deploys software, down to the VM
- Determining how the organization writes software by setting coding standards and running test machinery
- Standardizing how apps are instrumented

Sound familiar?

Platform engineering is a broad discipline, but it's reminiscent of when SREs were often confused for DevOps engineers when they first came on the scene. As the discipline continues to drive value, the roles and responsibilities will naturally solidify — but organizations shouldn't wait to bring clarity to their own teams.



Figure out where there's the most friction in your current development process — whether it is in the deployment pipeline, a Tower of Babel's worth of different programming languages, or a lack of standardization around observability — and focus your platform team's efforts there.


Greg Leffler, Director of Developer Evangelism, Splunk

What drives the platform engineering revolution

Organizations with platform engineering teams most often say their top achievement is increasing IT operations efficiency, improving tasks like scaling, monitoring, and troubleshooting (55%). In fact, improving efficiency appears to be where platform engineering shines, with 40% also claiming it makes developers more efficient and productive.

“Developers shouldn’t need to think about every ancillary detail: ‘How can I make this app observable? What framework will I use to talk to the database? Will this be FedRAMP compliant? Can we profile it while it’s running?’” says Leffler. “With a platform engineering team, that’s all decided and handled for you, enabling you to focus on delivering the right output.”

Standardization fuels efficiency. A larger organization may have multiple development teams working on different projects, each using its own set of tools and processes. When platform engineers standardize a common source code repository or CI/CD pipeline, for example, it can go a long way in improving collaboration and reducing deployment time. Security and compliance is where platform teams drive the most value here, likely because they provide rigor and certainty around how software is built and deployed, which is instrumental to achieving high-demand certifications like HIPAA and FedRAMP.



say their platform engineers’ efforts to standardize operations are successful

Top 5 benefits of platform engineering

55% Increase IT operations efficiency



42% Improve app reliability/performance



40% Improve developer productivity



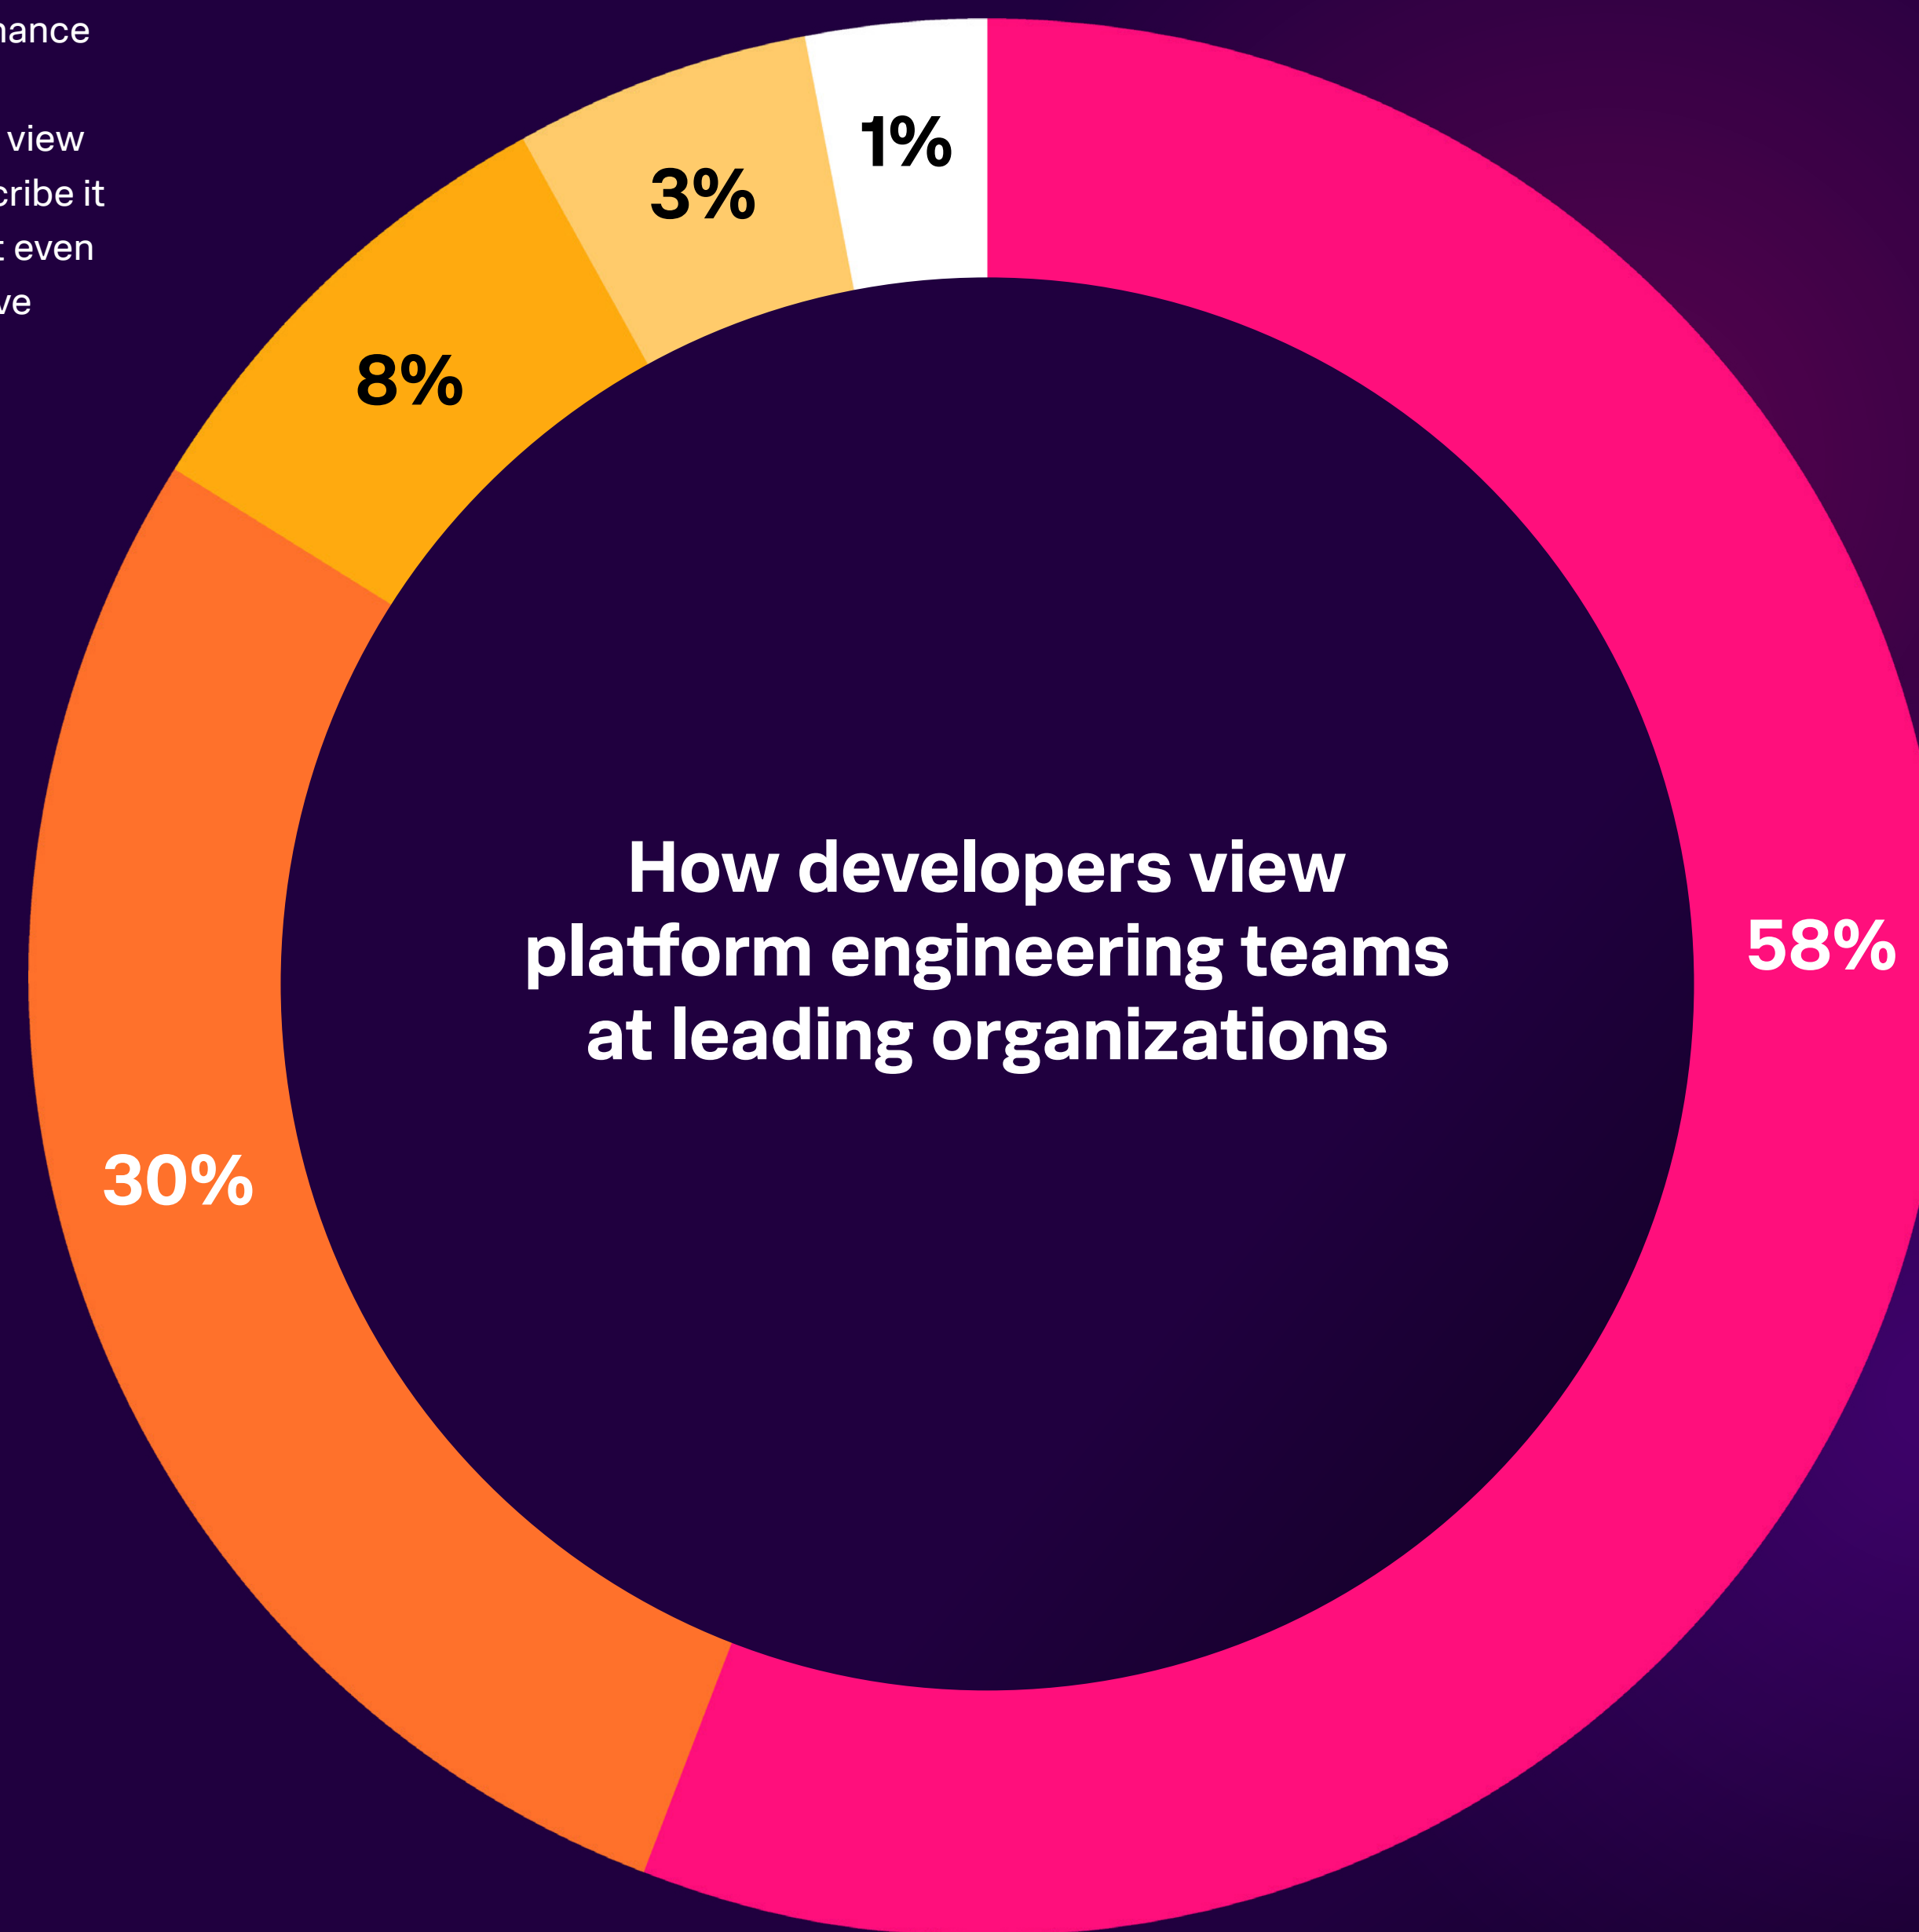
38% Optimize infrastructure for emerging workloads



31% Reduce tech debt



Since the ultimate goal of platform engineering is to enhance productivity, improve code quality, and generally make developers' lives easier, it makes sense that developers view platform teams in a positive light. Nearly half (47%) describe it as a "valuable service bureau." Some respondents went even further. Fifty-eight percent of leaders call it a competitive differentiator, compared to only 18% of beginners.



- Competitive differentiator
- Valuable service bureau
- Innovation hindrance
- Cost center
- Don't know

AI in observability starts to take shape

“Generative AI assistants are the key to democratizing observability domain knowledge. The ability to ask questions in natural language unlocks a completely new layer of insights and intelligence.”

Hao Yang, VP of AI, Splunk

What if engineers could instantly access information about their entire stacks, instead of squinting at dozens of charts, graphs, and dashboards as they send rapid-fire messages in a virtual war room? For decades, developers and engineers have been searching for this nirvana with the help of machine learning (ML), traditional AI, AIOps, and now generative AI. Have they discovered it yet?

AI and ML elevate efficiency


At this point, using AI and ML within observability tools is table stakes. Nearly all respondents (97%) currently use AI/ML-powered systems to enhance observability operations — a significant jump from last year, when 66% reported adoption.

Respondents — and especially leading organizations — report high success using traditional AI and ML to improve staff efficiency, specifically through correlating events and prioritizing alerts (56%) and recommending solutions (53%). Leading organizations in particular lean on the latter use case much more heavily than their peers, at 91%.

Alert fatigue is real — ask any SRE. But it's also where AI and ML can offer some relief. Fifty-seven percent of respondents agree their volume of alerts is problematic, and AI and ML's correlation capabilities are a potential remedy. Many report using AI- and ML-powered systems to actually remediate alerts; 34% say that half or more of their observability-related alerts per month are resolved in this manner. Once again, leaders are on the precipice, with 85% claiming that AI and ML remediate half or more of their alerts. In stark contrast, only 16% of beginning organizations say the same.

Root cause analysis is another area ripe for AI and ML because grouping data together is a puzzle, and not necessarily the fun kind. One-quarter of respondents agree data correlation and analysis is the hardest part of root cause analysis, followed by data collection and aggregation (22%). AI and ML can help with both, so it's no surprise that over half (55%) of respondents use AI- and ML-enabled tools to complete investigations and determine root cause.

Leaders lean on AI/ML to resolve alerts



● 50% or more ● Less than 50%, more than 0% ● 0% ● Don't know

Generative AI curiosity continues

Generative AI is shiny and new, and observability teams are curious about how it will transform their work. Most top-of-mind for observability teams are chatbots or AI assistants. These AI assistants comb through mountains of data to find the most relevant information, and give useful analysis or recommendations. Users can ask questions in plain speak to dig deeper, get context, and ultimately determine root cause. This reduces the learning curve of understanding query languages, enabling junior-level employees or even business users to gather insights when they need them.

A full 84% of respondents have explored generative AI features within observability platforms, specifically as chatbots or embedded AI assistants. They envision using these capabilities for data analysis (66%) and recommendations to resolve issues (60%). Despite the initial interest, only 13% have actually adopted these capabilities.

A lot of experimentation without actual adoption — but the reasons for that are valid. A big one: Generative AI-enabled observability solutions in the current market are brand-spanking new. Vendors are making plans but only a few are shipping a generally available product. Observability teams are still figuring out how to take advantage of these capabilities and incorporate them into their daily workflows. Another route is to build your own homegrown models to incorporate into an observability platform, but it requires a significant amount of time, money, and effort.

Many organizations are still working on their internal policies and procedures, so observability teams may be unclear on whether generative AI usage is even company sanctioned. In our report [State of Security 2024: The Race to Harness AI](#), 34% of security professionals say their organization doesn't have a completed generative AI policy.

Many explore generative AI, but few reach the summit

13%

Have actually adopted these capabilities

84%

Are interested, evaluating, or piloting generative AI within observability tools

3%

Not interested/
don't know



Earn your spot on the observability leaderboard

“Becoming an observability leader isn’t just about keeping the lights on, it’s about spearheading innovation. Those who can turn data into actionable insights will not just lead the industry — they will redefine it.”

Patrick Lin, SVP and GM, Observability, Splunk

Leading organizations have better visibility across their environment and, as a result, experience less downtime. When they do have downtime, they bounce back swiftly. They also have a lasting impact on the business by enabling their developers to innovate, not drown in drudgery, and launch products faster. That's just a sliver of what leading organizations can achieve. How do they rise above the rest and elevate an observability practice to be best-in-class?

1. Recognize OpenTelemetry champions within your engineering team.

OpenTelemetry's widespread adoption among leaders (78%) signals that it's a standard worth investing in with a range of benefits: better control over data, prevention of vendor lock-in, and access to a broader ecosystem of tech. However, respondents say implementation is challenging, mainly because they lack the staff familiar with OpenTelemetry. Mobilizing your existing workforce solves this, among several things: It signals to your developers that your company cares about keeping up-to-date with modern technology, employee learning and growth, and it sets you up for success by transforming your engineers into OpenTelemetry experts. It also imbues a sense of innovation to attract other talent in the future.

To inspire your development team to learn about OpenTelemetry, lean on your observability vendors — or better yet, build an internal training program. One way to get started is to find champions within passionate and motivated team members. Set them loose on the OpenTelemetry's project website and GitHub, and encourage them to join communities like the CNCF Slack group.

2. Position platform engineering as everyone's dream team.

Platform engineering is the key to embedding observability practices across teams, as well as enabling software engineers and SREs to do what they do best (building and shipping code) rather than wrangling toolchains. Overlooking its value would be a big miss; three-quarters (73%) of respondents employ platform engineering within their organization. But creating a team is just one step. It's equally important to set your platform team up for success by getting buy-in from executives, management, and practitioners.

Fifty-eight percent of leading organizations say that developers view platform engineering as a competitive differentiator. Gaining access to this differentiator first involves coming to an understanding that a platform engineering team would add value. For engineers and developers, this value is better efficiency and productivity (48% of leading organizations agree). For leadership, the value lies in business outcomes such as faster time to market, reduced technical debt, and more reliable code.

3. Share observability and security data for powerful outcomes.

Nearly three-quarters (73%) of leading organizations improved their MTTR by converging observability and security workflows and tools. When both teams have better context and data, they can more easily determine why an incident happened.

However, it's important to recognize that these teams have different goals, and those goals are sometimes in conflict. Approaching this collaboration incrementally can help ease the tension. First find common data sources and workflows that both security and observability teams rely on. Eventually, you can consider consolidating tools used by these teams, but it's best to start with a few high-priority data sources first and finetune the workflows surrounding them. Giving the network team access to all of this data will ensure everyone has even more context when troubleshooting.

4. Control telemetry pipeline costs before they control your business.

When it comes to your data, being obsessed with control is a good thing. Having control over how much data you emit, how you send it, and where you send it can help with wrangling costs — and squeeze more value out of your existing data.

Take a cue from leading organizations that embrace data management tactics like tiering (57%) as critical, cost-controlling strategies. Data tiering recognizes that not all metrics, logs, and traces are created equal. Move lower-priority data into less expensive cold storage to reduce costs. When you need it, rely on capabilities like federated analytics to access the data where it lives.

5. Lean on AI to alleviate alert fatigue and discover unknown unknowns.

Fifty-seven percent of respondents say that the alert fatigue associated with their observability solution is problematic. Thankfully, AI- and ML-enabled tools excel at detecting anomalies and correlating events, two use cases that lend themselves well to the alert fatigue problem. In fact, 85% percent of leading organizations use AI and ML to remediate half or more of their alerts. AI-enabled tools can reduce the noise in a few ways, either by aggregating alerts into meaningful groups that form a single incident or using smarter analytics to generate fewer alerts overall.

With one of the main AIOps use cases addressed, you can focus on other use cases such as accelerating root cause analysis, proactively spotting “unknown unknowns” and predicting potential problems before they turn into customer incidents.

6. See your network — even when you don’t own it.

Without visibility across your entire infrastructure, you don’t have end-to-end observability — and trends like containers, public cloud, and distributed networks make it harder to achieve. Leading organizations reported excellent visibility across public cloud environments (77%), on-premises infrastructure (74%), and applications (71%), but network infrastructure they didn’t own and operate was a gap, with only 45% saying they had excellent coverage.

Filling this gap includes prioritizing tools that give you insight into not only your owned network, but also the applications and services your internal users consume. While you may not have any control over this infrastructure itself, having that added visibility — whether it’s knowing an issue lies within a third-party API, ISP network, or border gateway protocol (BGP) — can help you spot problems like performance degradation and adjust accordingly. This may mean shifting certain workloads to other environments or rerouting traffic through use of software-defined networking.


Become an observability leader with Splunk



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 Hidden Costs of Downtime

Discover how resilient organizations avoid the \$400 billion dollar problem facing global businesses.

[Get the report](#)

Industry highlights

We identified key insights across four select industries worldwide.

Manufacturing

When it comes to investments in efficiency, manufacturers lead. Forty percent mostly or completely automate event grouping and alert correlation (compared to 32% across all industries). Manufacturers are more likely to use AI and ML technologies to address alerts; more than half report use cases such as anomaly detection (55%).

Improving efficiency has measurable payoffs. Eighty percent said they find issues within their applications in hours or less. The competitive advantage is even clearer in manufacturers' confidence in meeting application reliability and performance objectives: Half say they were "extremely confident," ahead of the 41% average.

When asked about the top observability challenges, data volume came out as the most frequent roadblock, experienced by a third of manufacturers. The good news is that manufacturers are at least aware of some solutions to their data-related conundrums. As an industry, they are more likely to agree that data aggregation (44%) and data tiering (43%) are critical for keeping costs manageable.

Communications and media

Organizations in the communications and media industry tend to have more mature observability practices: About a quarter are evolving (24% compared to the 17% average across all industries), whereas 10% are leaders.

Impressively, most communications and media organizations are recent adopters, with 59% adopting observability within the past two years. But within these past two years, they've racked up an average of 24.6 tools in their suites (more than the 23.1 industry average), as well as \$3.9 million in business value from their observability investments on average each year. That's half a million more than the industry average, \$3.4 million.

One secret to the success of the industry is its widespread use of platform engineering. Seventy-five percent have a dedicated team, and of these, 35% report they employ platform engineering extensively (compared to the 27% average). As such, it's no surprise that 76% of comms and media organizations say they're on the leading edge or in the early majority when they launch new products and features, well ahead of the 66% industry average.

Public sector

In the observability adoption journey, organizations in the public sector lag: More than half are beginners (54%, versus the 45% average). This squares with how they are much less likely to report having excellent visibility into their environments:

- Private cloud: 38% versus 47% across all industries
- Public cloud: 35% versus 47%

Since having visibility is foundational to observability success, it makes sense that public sector organizations are less likely to reap the outcomes of observability in areas such as application development (23% versus 34%) and problem detection times (19% versus 34%).

About a third feel underfunded when it comes to observability tools and technologies (29%, as compared to the 15% average). The public sector should keep developing and maturing their practices, especially since their current investments have already generated measurable payoffs. Spending \$1.2 million on average each year for observability tools and solutions has yielded \$2.4 million in business value.

Financial services

Relative to other industries, financial services organizations have an edge when it comes to visibility. They consistently reported having excellent visibility across infrastructure types at above-average rates:

- Network infrastructure they own and operate (53% versus 49% across all industries)
- Public cloud infrastructure (50% versus 47%)

Having established these foundational capabilities, a majority of financial services organizations have proceeded to make other investments, such as AIOps. Fifty-seven percent use AIOps tools for use cases like improving visibility through consolidating data from multiple monitoring systems. Taking this leap has paid off, as 67% of financial services organizations report the ROI exceeds expectations.

As the industry makes headway along the observability journey, it's encountered compatibility challenges with OpenTelemetry adoption (reported by 57%, compared to 46% across all industries). However, these compatibility gaps tend to resolve themselves as the project matures and as more vendors adapt their offerings to work well with OpenTelemetry. And despite the growing pains of OpenTelemetry, those within the industry are already reaping the benefits. Close to half (45%) report experiencing better control and ownership over their data.

Country highlights

Snapshots from across ten countries across the globe.

Australia and New Zealand

Over a quarter of Australia and New Zealand (ANZ) organizations are leaders — 26%, to be exact, compared to 11% across all countries. Organizations in ANZ spend more on observability tools than any other country surveyed (\$2.3 million per year). It's no wonder that they're reaping a hefty \$5.6 million annually in business value. That's more than two million ahead of other nations, which average \$3.4 million.

ANZ organizations tend to have excellent visibility across their environment, whether in private cloud environments (61%, compared to 47% globally) or network infrastructure they don't own or operate (44%, compared to 26%). Not only can ANZ teams spot a greater number of issues, but they also have 67% alert fidelity — the highest of all countries.

Sixty-five percent of ANZ organizations create hybrid observability and security roles (versus 53% across all countries). Converging these workflows and data enables teams to tap into more resources and resolve incidents faster. It makes sense that an overwhelming 86% of ANZ organizations have an MTTR in hours or minutes.

Developers in Australia and New Zealand spend 56% of their time on innovation and experimentation, which supersedes all other countries. Devoting more bandwidth has its payoffs, as ANZ app development teams launch an average of 13.3 new digital products and services each year, more than most other nations.

France

Organizations are still early in their observability journeys, with half in the beginner stage. Without a robust observability practice in place, an organization is less likely to be equipped with the means to detect, respond, and bounce back when incidents strike. This aligns as 36% of organizations in France have yet to implement a formal approach to digital resilience.

When we asked respondents about the total business value they're getting from their observability solutions, those in France reported the lowest average: \$2.6 million annually, compared to a \$3.4 million average.

To make headway in their observability practices, organizations in France should start by addressing their data volume and correlation challenges. Thirty-one percent report the amount of data collected exceeds human capacity to digest, and that correlating data from multiple sources is too time-consuming.

Another opportunity lies in OpenTelemetry, which provides better data control, prevents vendor lock-in, and enables access to a broader technological ecosystem. Over half (56%) do not currently use this industry standard in their observability solution, but there is certainly consensus on its value; 92% of organizations in France agree that it's important or critical for their observability vendor to contribute code to OpenTelemetry.

Germany

German respondents tend to be more nimble in their observability approach, with fewer tools in use and a tendency to prefer platform solutions over point solutions. They use an average of 17.8 different observability tools, well below the 23.1 average. Meanwhile, more than half prefer a single platform for observability capabilities like data transformation (57%) and data tiering (62%).

Having fewer tools could mean less context switching for teams, resulting in more efficient investigations. Organizations in Germany report high confidence in their ability to accelerate MTTD over time (85%, versus 78% across all countries). They also suffer downtime far less frequently than all other nations surveyed: just 5.9 instances each year on average, compared to 8.5.

When it comes to maturity, the most represented contingent in Germany is emerging (32%). Adopting AIOps can yield benefits for areas like data correlation and could be a way for German organizations to further mature their observability practices. Respondents in Germany use AIOps less than other countries to correlate data from different source types (38%, versus 46% across all countries). This may address a felt need, as 37% say only a limited amount of data can be correlated (versus 28% across all nations).

India

In India, organizations are twice as likely to be leaders (20% versus 11% across all countries). This could be due to a focus on innovation: 96% invest in solutions to reduce the amount of time developers spend dealing with incidents and alerts, so they have more bandwidth to spin up new products and features.

Respondents in India are also ahead of the game when it comes to adopting new technologies. Seventeen percent already use generative AI within their observability toolsets, with another 14% currently piloting the technology. Sixty-seven percent say their primary observability solution uses OpenTelemetry, ahead of the 58% adoption rate globally. And the benefits are already materializing, as 61% report that it's given them better control and ownership over their data.

By prioritizing innovation and embracing new technologies, organizations in India maximize the uptime and quality of their applications. Over half (52%) are “extremely confident” in their ability to meet application reliability and performance objectives, compared to 41% across all countries.

Italy

Two-thirds of respondents in Italy are beginners when it comes to observability, and a mere 5% are leaders. Since their practices are less developed, they're also less likely to reap the benefits that leaders do, such as improved MTTR and innovation. Organizations in Italy are less likely to report that observability has accelerated their problem resolution time significantly (22% versus 34% in the aggregate). Meanwhile, 11% of them are in the late majority when it comes to launching new digital products, which is almost double the average rate globally (6%).

Organizations in Italy are behind the curve, and the delta is only growing. Only 59% in Italy accelerated root cause identification across the past twelve months, compared to 71% worldwide. Meanwhile, when asked how optimistic they were about their ability to speed MTTR over time, 68% in Italy are confident (versus 80% across all countries).

Just 15% of organizations in Italy say their business allocates all the funding they need, compared to 28% globally. Highlighting the ROI of observability investments so far can help Italian organizations prove their value and improve these funding deficits as they continue their journey to maturity.

Japan

Two-thirds of Japan's organizations fall into the beginner category (66%). However, it has a strong 12% that have achieved leader status.

One possible barrier to observability success, especially for those in the beginner stage, is low confidence that an alert indicates a real problem. Organizations in Japan report only 49% of their alerts are accurate, compared to 61% across all nations. Alert fidelity issues may lead to end users, not an observability solution, reporting incidents. In Japan, 51% of incidents are flagged by the customer, the highest rate reported out of all countries.

Although organizations in Japan struggle with staffing — 55% say their observability teams have been left short-staffed multiple times — they are taking action. Nearly two-thirds (63%) will increase spending in observability staffing in the upcoming year, and resources are also being poured into AI/ML to increase efficiency. Already, 60% use AI/ML to detect anomalies (also ahead of the 51% worldwide average). These developments are impressive; last year organizations in Japan were behind in this regard, with just 15% using AI/ML within their observability toolsets.

Singapore

Eighteen percent of Singaporean respondents are leaders in observability (versus 11% across all countries). As a result, they're better able to innovate and evolve. Application development teams in Singapore launched an average of 13.5 new digital products and services this past year, more than any other country surveyed.

To sustain the pace of innovation, Singaporean organizations may need to address bandwidth issues among their observability teams. Eighty-two percent reported critical members have left because of burnout. Low alert fidelity could explain why teams are overtaxed; 31% of respondents in Singapore experience a "highly problematic" number of false positives, compared to 13% globally.

Further steps to improve staff bandwidth include reducing the number of observability tools, as Singaporean organizations use an average of 28 tools. They may also benefit from automating event grouping and alert correlation, since 20% say that process is completely manual right now (versus 7% globally).

U.K.

More than a third of organizations in the U.K. are in the emerging stage of observability (34%), a promising sign that they've advanced beyond the beginner stage. At the same time, organizations in the U.K. still have plenty of room for improvement, as only 8% are leaders.

Given that they're less mature in their observability practices, respondents in the U.K. are also slower when it comes to innovation. Releasing an average of 10.8 new products and services annually, they lag behind most other nations surveyed when it comes to the pace of innovation.

On the whole, U.K. organizations should keep investing in observability, as it's producing measurable dividends. Spending \$1.3 million annually on observability tools has yielded \$3.5 million in business value (above average, relative to other nations).

An area of potential investment U.K. organizations can explore is identifying root cause. Sixty-five percent have seen the time to identify root cause accelerate this past year, compared to 71% worldwide. To detect incidents faster, organizations in the U.K. could benefit from AIOps adoption, which brings automation and greater intelligence to investigations. Just 46% are adopters, compared to 52% across all countries.

U.S.

When it comes to observability maturity, the U.S. is in line with the global average.

One area of strength for U.S. organizations is its investments in staff efficiency. An overwhelming 93% have explored generative AI to enhance their observability team's productivity, and of these, 84% report it already makes a positive "significant impact" (compared to 74% across all countries).

Observability teams in the U.S. are also more likely to use AIOps within their toolsets (60%, compared to 52% globally), most commonly for use cases like determining root cause and incident remediation. And 73% say its ROI has exceeded expectations.

One sign of an effective observability practice is the ability to detect issues before the customer does. And that's the case for many U.S. organizations. Fifty-eight percent of incidents are generated by their observability solutions, not the end customer — the highest rate of any other nation surveyed. It's likely that prevailing AIOps adoption could have played a part, since it enhances root cause detection and incident remediation with greater intelligence.

Methodology

Researchers surveyed 1,850 ITOps staff, managers, and executives, as well as developers, engineers, architects, and SREs in May and June of 2024. Respondents were in Australia, France, Germany, Italy, India, Japan, New Zealand, Singapore, the United Kingdom, and the United States. They also represented 16 industries: Aerospace and defense, business services, consumer packaged goods, education, financial services, government, healthcare, life sciences,

manufacturing, technology, media, oil/gas, retail/wholesale, telecom, transportation/logistics, and utilities.

Respondents were placed into one of four stages of observability maturity, based on a new, four-level observability maturity framework assessing 20 self-reported data points.

The breakdown of organizations that represented each stage is as follows:

- Stage 1: "Beginning" organizations (45%)
- Stage 2: "Emerging" organizations (27%)
- Stage 3: "Evolving" organizations (17%)
- Stage 4: "Leading" organizations (11%)

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, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk LLC. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2024 Splunk LLC. All rights reserved.

