groundcover

THE OBSERVABILITY IMPERATIVE REPORT:

Technical and Strategic Forces Reshaping Infrastructure in 2026

Research Report

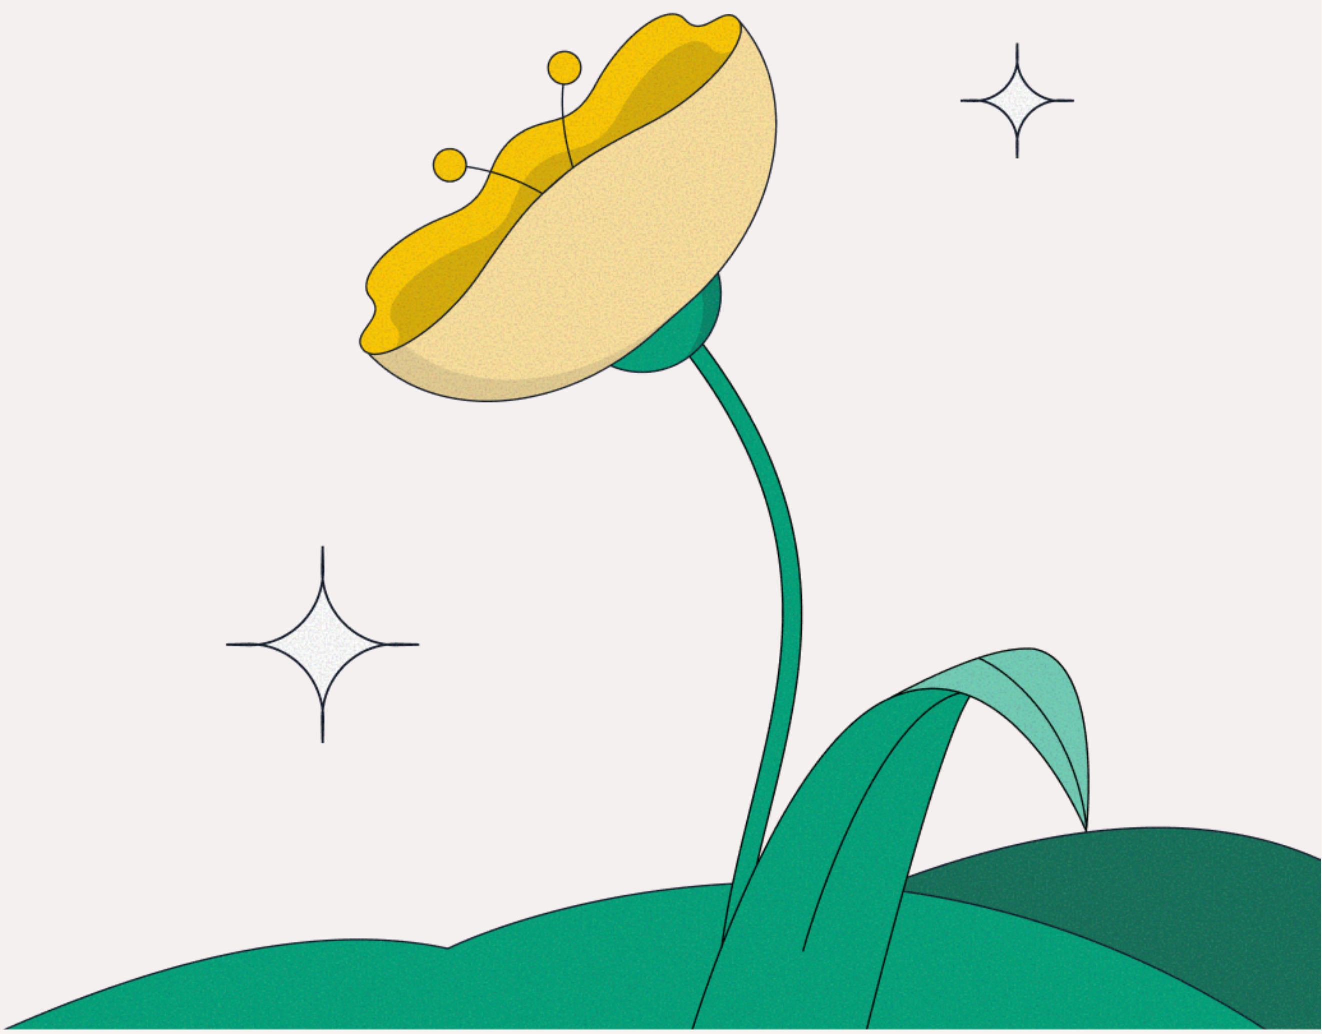

Some thoughts from

Shahar Azulay, CEO groundcover

A quick note before you dig in.

The 500 engineers and observability leaders in this survey are describing the same shift we're watching from our side of the product.

Adoption of groundcover's MCP server has moved faster than anything we've ever shipped. Teams aren't using it to chat with their telemetry. They're wiring it into their own investigation agents. AI isn't sitting next to observability anymore. It's consuming it as raw input.

The data in this report tells you how widespread that shift already is. 89% of teams use observability data for forward-looking decisions. 87% have AI integrated into their workflows.

"You can sample traditional telemetry and mostly get away with it. You cannot do that to AI observability data."

And then the finding that should stop you: only 34% trust that integration to act.

Everyone wants to call that a model problem. It isn't. It's a data problem.

AI agents are pattern-matchers. They reason over the telemetry you captured. Sample the interesting parts away to keep the bill down, and the agent has nothing to reason over.

We've always believed lots of data is the point. Traditional telemetry tolerates sampling. Dashboards smooth over the gaps. AI observability data does not. The signals are too sparse. The failures are too non-deterministic. The missing span is always the one that mattered.

The rest of this report is the clearest picture I've seen of where teams are caught between cost and visibility — and what the ones getting it right are doing differently.

Worth your time.

Executive Summary

Observability (o11y) has crossed a threshold. It is no longer a tool used primarily to investigate failures after they happen. It is a continuous input to AI pipelines, capacity planning, automated routing, and product roadmapping.

Atomik Research surveyed 500 observability leaders and engineers to understand how AI is reshaping observability. 89% of respondents already use observability data for forward-looking decisions, and 47% say their telemetry is continuously consumed by systems outside the observability platform itself. Observability is no longer just tooling in the background. It has become foundational infrastructure that modern systems actively depend on to plan, operate, scale, and recover when things go wrong.

87% have integrated AI into observability. only 34% trust it. **Adoption is outpacing trust by 53 points.**

However, the o11y infrastructure is not yet equal to the role it has been given. AI and automation are integrated into observability workflows at 87% of respondent organizations, but only 34% describe that integration as fully operational and trusted. The 53-point gap between adoption and trust is the most consequential finding in the survey. It is not, in our reading, primarily a **model-quality problem**. It is a **data-fidelity problem**. 43% of respondents name sampling and incomplete telemetry as their top barrier to advancing observability. More than expertise, more than tooling fragmentation. AI features in observability tools reason over the data the platform captures. When the platform has sampled most of the relevant spans away, the AI is reasoning over a redacted document, and engineers can sense it.

43% of respondents name **sampling** and **incomplete telemetry** as their **top barrier** to advancing observability.

Cost pressure is making the data fidelity problem worse, not better. AI workloads now drive 26-50% of observability spend at half of respondent organizations, and 53% experienced budget overages of 10% or more in the last fiscal year. The default response is to sample more aggressively and retain less data, but that approach cuts deepest where visibility matters most: AI agents and their data-thirsty workloads. The choice between cost and visibility, presented as a tradeoff, is in practice a false one for these workloads: the visibility teams give up to save cost is the visibility they need to debug the failures that drive the next cost spike.

Teams are responding.

79%

These respondents have implemented adaptive sampling, changed retention policies, renegotiated contracts, or begun exploring open source alternatives.

35%

These respondents are using cost as a triggering event to consolidate tools or rethink vendor relationships.

The next wave of observability contracts will be shaped less by feature checklists than by whether vendors can credibly deliver high-fidelity, broadly-accessible telemetry at a cost that does not force teams to choose between visibility and budget. The organizations that find that combination first will close the trust gap.

SECTION 1

Introduction: The Contradiction at the Center of Observability in 2026

Observability in 2026 is caught in a contradiction of its own making. Teams have elevated telemetry data from an operational backstop into a strategic input where 89% of respondents now use observability data to inform forward-looking decisions, not just to triage incidents. At the same time, the cost pressure of running observability at AI-era scale is forcing teams to sample more aggressively, retain less, and instrument selectively. The visibility required for strategic decision-making is being undermined by the cost controls teams are forced to implement.

> At the same time, the **cost pressure** of running observability at AI-era scale is forcing teams to **sample more aggressively, retain less, and instrument selectively.**

This contradiction is the structural condition this report describes. Every finding in the survey either causes it, follows from it, or describes how teams are attempting to resolve it. The trust gap between AI integration and AI operationalization, the persistence of incomplete telemetry as the top barrier to advancing observability, the rise of AI workloads as a non-deterministic cost driver, the consolidation pressure on multi-vendor stacks; none of these are independent findings. They are connected pieces of the same problem.

The organizations that resolve the contradiction first will operate with materially better signal than organizations that don't. The organizations that don't adjust will keep paying more for less visibility into systems that are themselves becoming less deterministic. This report explores that contradiction in detail, highlighting the engineering and operations teams most affected by the tradeoffs it creates, and examining where organizations that successfully resolve it are likely to begin.

SECTION 2

Adoption Has Outpaced Trust

AI is integrated into observability workflows. It is not yet trusted with autonomous action. The reason is structural, not a function of model quality.

2.1 The 53-point gap between adoption and operationalization

87% of respondents report AI or automation are integrated into their observability workflows. Only 34% describe that integration as fully operational and trusted. The remaining 53% sit in an assistive role: AI helps draft a query, summarizes a trace, or proposes a correlation, but a human still ratifies the action before anything happens.

The gap between integrated and trusted is the finding that matters. Teams have not been blocked on adoption. They have been blocked on the cost of verifying AI suggestions against the underlying data. When an AI tool proposes a root cause, an engineer still has to confirm it against the trace, and that confirmation often takes longer than reading the trace would have taken in the first place. The productivity gain implied by AI integration is partially absorbed by the verification overhead it creates.

87%

Integrated AI into observability workflows

34%

Describe it ass fully operational and trusted

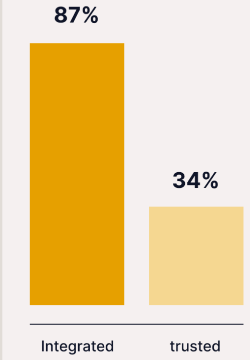

53-points trust gap

2.2 Confidence in the underlying signal is the precondition

The instinct is to read the trust gap as a model-quality problem. That the AI features are not yet good enough to act autonomously, and as models improve, the gap will close on its own. The survey data suggests something different. Only 22% of respondents say they are "very confident" that their observability tools capture enough high-fidelity signals like logs, traces, reasoning steps, model outputs to detect real AI issues. Only 50% say they are confident at all. For teams responsible for mission-critical systems, that is alarming on its face, but the deeper implication is what it does to AI features built on top of that signal.

AI in observability is a pattern-matcher. It reasons over the data the platform has captured. If the platform's confidence in its own signal is low, the confidence ceiling for AI features built on that signal is also low, regardless of how capable the underlying models are.

SECTION 3

Strategic Imperatives —Why Are Organizations Starting to Shift Their Focus

Growing organizations are already responding to the agentic age and deploying their observability with a markedly different strategy. That shift in approach will likely yield significant benefits for them as the scale of development escalates this year, and for the foreseeable future.

3.1 From Reactive to Proactive

According to the survey respondents, 89% of individuals are already using observability platforms for their forward-looking decisions. That trend is not surprising to the well-initiated in the engineering space. Logs, metrics, and traces are filled with extraordinary context that reflect not just how individual applications are behaving, but also how their interconnected systems are reacting to each other.

The 89% breaks down by frequency of o11y data use. 42% of respondents are frequently using observability data to inform decision-making (beyond incident response), and 47% are continuously leveraging it - embedding it into AI, product or operational decision workflows. Regardless of how often respondents consult the new data signal, it has clearly become a tool for proactive roadmapping within their organizations.

The shift to a proactive approach signals a dramatic change in perspective from how teams and their leadership have historically viewed observability data. What was once seen as telemetry data to populate dashboards and automatic alarm bells to trigger when the ecosystem metaphorically started to smoke, has become a guiding signal for what technical teams need to build next. This is a structural change in what observability data is for. The historical role where observability as the diagnostic layer humans consult when something is wrong has not disappeared, but it now sits alongside a second role: observability as a data layer that production systems consume to make automated decisions. The two roles have meaningfully different requirements.

3.2 The New Metrics Era

Surveyed individuals say their organizations are matching a new approach to observability with a new set of metrics. 67% of respondents are using AI observability tooling to methodically and frequently track AI system performance. Metrics highlighted include agent handoff rate, collaboration latency, and workflow completion reliability. That's a digression from historical metrics, like latency at each hop, request rate/throughput, and log error volume.

It is significant that coordination-specific performance metrics are so important to these teams and individuals. It points to an era of systems, their thinking and coordination, and the frailty of those connections when not measured and optimized. Throughput still matters, but on its own it's no longer enough to measure the health and performance of the complex systems organizations are building today.

Engineering and platform teams now focus on more than traditional infrastructure metrics. For example, respondents' top priority fror this year was reasoning accuracy (48%). However, respondents' other priorities look familiar: tool/API reliability (44%), and cost efficiency (38%). It's clear teams are in a transition phase where they are blending their always on, business critical focus with new initiatives.

TOP PRIORITIES FOR ENGINEERING AND PLATFORM TEAMS

Reasoning accuracy leads, but reliability and cost remain top of mind

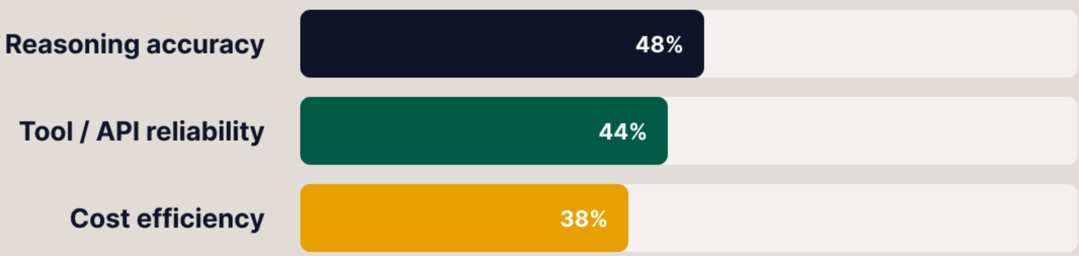

PERCENTAGE OF RESPONDENTS SELECTING EACH AS A PRIORITY. MULTI-SELECT; TOTALS EXCEED 100%.

But this is key.

When telemetry was primarily a diagnostic input, a five-minute outage in the observability backend was inconvenient but not consequential to the systems it observed. When telemetry is a production input such as feeding an autoscaler, an incident router, or an AI agent that takes action on its content, then a five-minute outage in the observability backend is a production outage in every downstream system that depends on its data. The observability platform's availability implicitly becomes the availability floor of every system that consumes its telemetry.

Most respondent organizations have not yet adjusted their operational expectations to match. Observability backends are still typically operated with the SLAs appropriate to their historical role, not the role they have inherited. Where AI agents consume observability data to make autonomous decisions, this gap is particularly consequential: the trust problem is compounded with the availability problem here, because gaps and delays in the telemetry feed translate directly into gaps and delays in the actions the AI is supposed to take.

3.3 Cost as a Strategic Constraint

According to respondents, observability costs are no longer just a budget concern. They’re actively limiting how teams monitor, troubleshoot, and scale their systems. Cost predictability plays a predominant role in that. 42% of survey respondents experienced moderate overages on their observability spend. Another 13% of respondents experienced significant overages (anything over 30% greater than their forecasted observability costs) in the last 12 months. Together, that’s a majority of respondents (53%) who were billed at least 10% over forecasted costs.

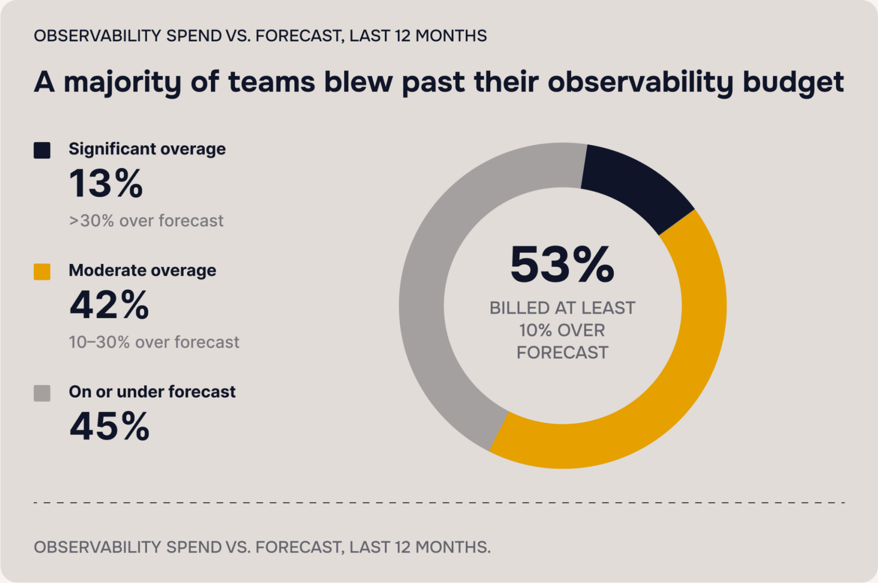

61% of respondents in this group say their organizations spend at least $1 million annually on observability, and 22% spend more than $5 million. At that scale, overage charges become a major operational and financial burden.

Evaluated in the context of these costs, it's not surprising that 27% of respondents are already implementing adaptive sampling and telemetry controls. Another 34% are optimizing ingestion and retention policies. Only 4% of respondents answered that no, costs are not a priority. All of this data suggests costs are a significant source of friction for platform, dev, and engineering teams.

ANNUAL OBSERVABILITY SPEND

61%
spend at least $1M annually on observability

22%
spend more than $5M annually

AMONG ORGANIZATIONS RUNNING OBSERVABILITY AT SCALE.

It becomes clear that there are two competing priorities. 43% of respondents cited sampling limitations or incomplete telemetry reducing visibility into AI workflows, as the biggest obstacle preventing their team from using observability data for proactive or AI-driven decision-making. At the same time, incomplete data is a severe obstacle to delivering their strategic commitments; only 4% of respondents are free of the burden of cost and can leverage the data they need.

Winning organizations are responding in the AI-first era by using observability data from within their ecosystems to try and plan, and yet they are battling the realities of their costs which form a painful strategic constraint.

3.4 The Platform Consolidation Imperative

Telemetry data gives invaluable organizational context to the platforms that leverage it, but it's clear from respondents that each tool also creates its own walls around the data it accesses. As a result, observability practitioners are limited by the value they can deliver when leaning on individual tools. When polled, 34% of respondents suggested that data silos or fragmented tooling prevent their teams from using observability data for their proactive decision-making.

The signal correlation problem is real and worsens with scale. When the trace that captured a failed agent request lives in one tool, the metric that captured the cardinality spike lives in another, and the log line that captured the provider error lives in a third, the engineer triaging the incident is reconstructing the request across three platforms. Each tool holds only a fragment of the story.

Putting it another way, the complexity of the observability ecosystem impacts how effectively the respondents can guide the engineering roadmap and long-term technical planning. That challenge is compounded by the fact that 33% say their current tools are not aligned with business or AI KPIs, signaling that many organizations are rethinking the role each observability tool should play in their stack.

Furthermore, 35% of their organizations are using costs as a compelling event to consolidate tools (17%), or renegotiate vendor contracts or explore open source (18%). If the value delivered by the traditional observability vendors isn't commensurate with the costs and ability to meet strategic imperatives, the next wave of contracts will shift to vendors delivering a complete and integrated AI-native platform.

SECTION 4

Snapshots by Organization Type

Some notes on the methodology for this particular report section:

- Averages are rounded down to the closest whole number.
- This report uses the self-reported revenue indicator to categorize organization type into early-stage, mid-market, and enterprise.

METHODOLOGY — ORGANIZATION TYPES

Revenue-based segmentation used throughout this report

TIER 1

Early-stage

$10M – $99.99M annual revenue

- $10M – $49.99M subgroup
- $50M – $99.99M subgroup

TIER 2

Mid-market

$100M – $2B annual revenue

- $100M – $500M
- $500M – $1B
- $1B – $1.5B
- $1.5B – $2B

TIER 3

Enterprise

>$2B annual revenue

- Self-reported revenues greater than $2B annually

Organization type is based on self-reported revenue indicators. Averages are rounded down to the closest whole number.

4.1 Early-Stage and Scaling Organizations

Less than 10% of the respondents are from firms with fewer than 100 employees.
Only 34% of respondent businesses are less than $100M in revenue.

Current state:

Across the two early-stage groups, an average of more than 54% of firms in this category spend at least $1M annually on observability. Only 15% would claim that their end-to-end observability is excellent (the highest distinction). Additionally, in the average between them, only 22% would say they are very confident that they enjoy comprehensive observability (again, the highest rating). An average of 14% of these companies have highly variable/untracked or insufficient visibility data to detect and resolve AI application incidents.

Challenges:

These firms are pressed to quickly show growth and to demonstrate their ability to iterate under market pressures. The speed with which they need to deliver will be challenged by their observability stack's overall time to deployment. That speed is balanced against noted organizational concerns.

Organizational concerns preventing broader AI-driven automation in observability include lack of transparency or explainability (31% average), false positives (43% average), compliance or governance (26% average), integration complexity (40% average), and limited visibility into AI reasoning paths (36% average).

Strategic move:

Early stage and scaling organizations need to quickly establish their observability ROI measurement, and commit to it before beginning their aggressive scaling. They also need to decide whether they have mission critical production workloads that require observability, or whether to consider alternative observability approaches.

4.2 Mid-Market Teams

45% of respondents were from firms with between 100 and 500 employees. 47% of respondents are between $100M and $2B in revenue.

Current state:

On an average across the four groups in the mid-market category, 68% of firms spend at least $1M annually on observability. On average, only 21% would claim that their end-to-end visibility observability is excellent, with full visibility across AI systems. 26% would say they are very confident that they enjoy comprehensive observability (the highest rating) when asked how confident they are that their observability tools capture sufficient high-fidelity data to detect real AI issues. An average of 14% of these companies report having highly variable/untracked or insufficient visibility data to detect and resolve AI application incidents.

68% of firms spend at least **$1M** annually on observability.

Challenges:

These firms are forced to deliver continued growth and innovation. Observability tooling can't hamper their strategic considerations towards growth, and requires full-stack visibility across both nascent and legacy applications or systems.

Organizational concerns preventing broader AI-driven automation in observability include lack of transparency or explainability (40% average), concerns with false positives or incorrect recommendations (32% average), compliance or governance concerns (36% average), integration complexity (40%), and limited visibility into AI reasoning paths (45% average).

Strategic move:

On average, 19% of mid-market firms will consider consolidating tooling, but they are the group most likely to implement adaptive sampling or telemetry tiering: on average 30% are considering sampling this year. This group will prioritize high value telemetry data until vendors catch up to their concerns about governance and cost controls.

4.3 Enterprise organizations

47% of respondents were with firms between 500 and 1,000 head count. No respondents were from firms with more than 1,000 employees. 19% of respondents fit into this category of revenues greater than $2B.

Current state:

On average across the six subgroups in the enterprise category, 58% of firms spend at least $1M annually on observability tooling, and 29% consider their end-to-end visibility into AI/LLM applications to be excellent. 22% would say they are very confident they enjoy comprehensive observability.

An average of 6% of these companies have highly variable/untracked or insufficient visibility data to detect and resolve AI application incidents. This is the most confident group: 65% of respondents in this category say the average time it takes to detect and then resolve AI application incidents is at most a day. 28% detected related incidents within minutes, and resolved within the hour, while 37% detected within hours, and resolved within the day.

ENTERPRISE — AI INCIDENT RESPONSE

The most confident group resolves AI incidents fast

65% detect and resolve AI application incidents within a day

28% Detected within minutes
Resolved within the hour

37% Detected within hours
Resolved within the day

Enterprise category — average across six subgroups. Self-reported.

Challenge:

Enterprise companies at this scale face enormous operational and cost pressures as their environments grow. They also have to wrangle organizational complexities across teams and departments.

Their organizational concern profile for trends preventing AI-driven automation in observability include lack of transparency or explainability (40% average), concerns with false positives or incorrect recommendations (41% average), compliance or governance concerns (39% average), integration complexity (29%), and limited visibility into AI reasoning paths (30% average).

Strategic move:

27% of firms in this category are already planning to adopt AI-driven recommendations or automation as their organizational use of observability data evolves and AI adoption increases.

With such confidence in their existing detection and remediation capabilities, this group will be the first to close the loop between observation data and the full software development lifecycle.

WHAT'S NEXT — ENTERPRISE

27%

of firms are already planning to adopt **AI-driven recommendations or automation** as observability data use evolves.

Key Takeaways and Recommendations

Eliminate blind spots

Collect all the Telemetry data that you can with and reduce sampling where fidelity is critical.

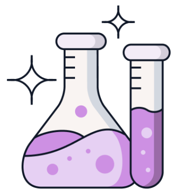

Measure what matters

Audit how much telemetry actually drives decisions.

Simplify the stack

Consolidate tools across your teams to improve signal correlation.

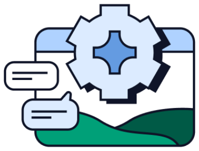

Design for AI systems

Treat AI observability as a first-class requirement.

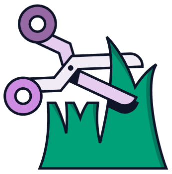

Embrace OTel

Standardize instrumentation to reduce fragmentation and cost.

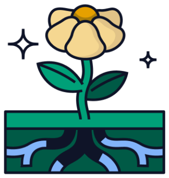

Rethink your o11y arch

Align costs, control, and visibility with modern infrastructure needs.

CONCLUSION

Observability Becomes Infrastructure

The observability market in 2026 is caught in a contradiction of its own making. Teams are forced to cut costs, and therefore visibility. Necessary sampling limits and tighter retention policies directly erode the telemetry fidelity they need to do the thing they say matters most: use observability data to drive strategy, not just triage incidents.

The teams that recognize this first and restructure their observability investment accordingly will grow at their desired scale and pace. The vendors that give them a credible path to observe it all will win the contracts when this wave of renewals hits.