

The Business Value of Grafana AI



Shannon Kalvar

Research Director, Enterprise Systems Management, Enterprise Client Platforms, Observability, and AIOps, IDC



Megan Szurley

Business Value Manager,
Business Value Strategy Practice, IDC




Table of Contents



Click any title to navigate directly to that page.

Business Value Highlights	3
Executive Summary	4
Situation Overview	5
Grafana AI Overview	6
Study Firmographics	7
Choice and Use of Grafana AI	8
Business Value and Benefits of Grafana AI	9
IT Benefits	11
Engineering/SRE Benefits	14
Operations and Incident Response Benefits	18
Improved Service Reliability and Customer Trust	20
Business Value Summary	21
Challenges/Opportunities	22
Conclusion	23
Appendix: IDC's Business Value Methodology	24
About the IDC Analysts	25

Business Value Highlights

Click each [→](#) to jump to related content. Use "Return to Highlights" in the page header to get back here.

\$1.67 million average annual benefit per organization



31% higher engineering/SRE productivity, valued at just over \$1 million a year



52% more efficient platform administration and management gain, valued at \$327,000 a year



\$325,000 in annual cost avoidance from tool reduction



63% faster MTTR



69% faster for new engineers to reach independent proficiency



46% less time spent gathering context



49% fewer support requests to senior SREs



Executive Summary

IDC conducted interviews that explored the value and benefits for organizations using Grafana AI, a natural language assistant built into the Grafana platform that lets users query, investigate, and document observability data in plain language. Based on these interviews, IDC found that the organizations achieved observability data in plain language and achieved benefits across IT operations and engineering/site reliability engineering (SRE) productivity worth an annual average of \$1.67 million per organization.

The top value drivers of Grafana AI were:

- **Real-time incident context:**
Grafana AI ran its own correlation across metrics, logs, and traces the moment an incident began. That gave responders earlier visibility into what was happening and who was affected. It also let nonexperts join investigations and query observability data without specialized skills.
- **Reduced reliance on senior staff for support and training:**
Other users found their own answers through Grafana AI's natural language interface instead of routing routine troubleshooting and how-to questions to senior engineers. New engineers reached proficiency faster with its guidance, freeing senior staff for higher-value engineering work.
- **Automated documentation and stakeholder communication:**
Grafana AI generated dashboards, alerts, and post-incident summaries directly from natural language requests and prepared status updates and incident write-ups for internal and customer-facing stakeholders. That cut administrative overhead after every investigation.
- **Lower-cost, easier-to-maintain platform:**
Interview participants consolidated separate tools for metrics, logs, incident response, and front-end monitoring onto one Grafana platform, cutting licensing costs. With configuration as code and Grafana AI handling routine upkeep, IT ran observability with far less effort.
- **Improved service reliability and customer trust:**
Faster, more visible incident handling supported service-level agreement (SLA) commitments and smoother renewal and upsell conversations.

Situation Overview

Digital estates now generate operational signals faster than the people responsible for them can absorb them. Metrics, logs, traces, and events arrive from containers that live for minutes or are completely stateless, services that change several times a day, and dependencies that cross clouds, providers, and teams. Most organizations still run dashboards and procedures built for an estate someone could hold in their head. Few estates fit that description anymore.

The market has responded by building a durable knowledge layer beneath the signals. Topology maps, entity models, knowledge graphs, and context stores record what exists, what depends on what, and what happened the last time something failed. Some organizations consolidate this layer on a single platform, and others federate it by querying data where it already lives. Either arrangement gives machines a structure to reason over when the system state has grown too large and too ephemeral for direct human comprehension.

Observability vendors now position natural language assistants and investigation agents as the primary way to interact with that layer. The claim they make is consistent across the market: Anyone can ask the estate a question in plain language and receive an answer grounded in live telemetry. This claim raises a practical question about who is now doing the asking.

Until recently, the answer was a small specialist cohort. Querying an observability platform meant knowing a separate query language for each signal type, the estate's shape, and the history of its failures. This knowledge was concentrated in a few senior engineers, who spent a large share of their week sharpening and answering others' questions. Natural language access moves correlation, triage, and context gathering to the machine. This sorting work on emergent conditions is the class most safely delegated because it operates on patterns already present in the data.

The same access moves orientation outward. Support staff, managers, and junior engineers can establish what is happening, who is affected, and where a fault originates without waiting for a specialist. The agents increasingly act as well, building dashboards and alerts and running scheduled tasks with the permissions of whoever deployed them. This redistribution of who can act, among people and between people and machines, is the story buyers understand but have difficulty measuring.

What buyers can measure is what happens when they compress the chaos: shorter time to resolution, fewer interruptions for specialists, and faster time to action for new staff. Compression reduces turbulence within the existing operating pattern, producing

gains measurable in the current operating model. Changing the productivity slope, the capacity to absorb continued growth in complexity without a proportional rise in human strain, depends on redesigning the arrangements around the tools. Those arrangements set who describes the desired state, how agent action is bounded, how incidents become lasting improvement, and how shared understanding is kept.

Enterprises that bank compression without that redesign tend to spend the freed capacity on broader coverage, and the human load returns at a larger scale. Any buyer evaluating an AI-enabled observability platform therefore faces two questions: How much compression does the product deliver now, and how well does it support the redesign that follows?

Grafana AI Overview

Grafana AI is what Grafana Labs named the AI and machine learning capabilities across Grafana Cloud. They operate at three levels: generative and agentic, predictive, and prescriptive.

Generative and agentic: Grafana Assistant, generally available since October 2025, is a large language model (LLM)-based agent that writes and explains PromQL, LogQL, TraceQL, and SQL queries and builds dashboards and alerts from natural language requests. It works inside Grafana, in Slack, and through an model context protocol (MCP) server that external agents can call. Assistant Investigations, generally available since July 2026, runs specialist agents in parallel across metrics, logs, traces, and profiles and returns a report of hypotheses with the queries behind them. Automations run scheduled agent tasks, and Watchers, in public preview, monitor scoped signals continuously.

Predictive: Statistical models forecast metric values, detect outliers among peer series, and drive ML-based alerting. The knowledge graph models entities and their relationships and flags saturation, errors, and anomalies across them.

Prescriptive: Adaptive Telemetry analyzes how metrics, logs, traces, and profiles are actually used and recommends what to aggregate, drop, or sample. Infrastructure memory, team-authored skills and rules, and search across prior investigations give Grafana Assistant operational history to draw on. The assistant queries more than 30 data sources in place, including third-party observability platforms and databases.

In every case, Grafana Assistant acts with the permissions of the user who invokes it. The deepest capabilities, among them Investigations, infrastructure memory, and Adaptive Telemetry, operate on telemetry held in Grafana Cloud, and self-managed deployments receive a subset. The four organizations IDC interviewed for this study have used these capabilities long enough to measure what has changed.

Study Firmographics

To quantify the business value of Grafana AI, IDC conducted in-depth interviews with four organizations actively using the AI solution. The study explored each organization’s deployment in detail, covering the selection process through day-to-day observability work and incident response. IDC asked participants quantitative and qualitative questions about their environments before and after adoption, the outcomes they observed, and the capabilities they relied on most.

Table 1 (below) presents the firmographics of the interviewed organizations. Participants averaged 979 employees and 195 IT employees, with average annual revenue of \$340 million. Organizations were headquartered in the Netherlands and the United States and represented the healthcare, software, and technology industries.

Table 1
Firmographics of Interviewed Organizations

Firmographics	Average	Median	Minimum	Maximum
Number of employees	979	950	615	1,400
Number of IT employees	195	35	8	700
Annual revenue	\$340M	\$309M	\$90M	\$649M
Countries	Netherlands, United States (3)			
Industries	Healthcare, Software, Technology (2)			

n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

Choice and Use of Grafana AI

Interviewed organizations adopted Grafana AI to close a gap that had been building for years: Observability data kept growing, but the specialized knowledge needed to query it remained concentrated in a handful of people. One participant pointed to the sheer number of product-specific query languages they had to remember just to analyze metrics, logs, and traces. Another described Grafana AI as unmatched when it first came to market since nothing else could make Grafana's own metrics and logs as consumable. Some participants were already Grafana customers and adopted the assistant as soon as it reached preview. They watched it mature from mixed early results into a tool reliable enough to trust with dashboards, alerts, and incident investigation. Cost factored in too. Grafana AI remained free during public preview, and once it started costing money, participants weighed that cost against the value it had already demonstrated. They decided it was worth paying for.

Study participants described their reasoning:

Simplified observability (Technology):

"We decided to use Grafana AI because it resolves challenges around remembering observability product-specific use cases and analyzing the different signals, logs, metrics, and traces."

Unmatched capability (Technology):

"When Grafana AI came into the market, there was nothing else like it that could read the metrics and logs in Grafana and make them more consumable. We were already using Grafana tools, and Grafana AI let us read those metrics better."

Increased reliance (Healthcare):

"We've been a Grafana customer since 2022 and started using Grafana AI when it launched in preview in late 2025, initially with mixed results. By early 2026, it had become more reliable and faster, and we expanded its use to dashboards, alerts, and incident investigation. Since February and March, we haven't done anything by hand in Grafana; all of it now goes through Grafana AI."

Proven value (Software):

"We got access to Grafana AI through our Grafana Cloud license and let teams use it once we saw it could be useful during the public preview. It stayed free while it was in public preview, but that changed once it started costing money, and we had to decide if it was worth it. At that point, we decided it was worth paying for."

Table 2 (below) shows the Grafana AI environment that the interviewed organizations operate. Study participants used Grafana AI to monitor an average of 2,607 applications and services across 60 connected data sources. They managed an average of 649 dashboards and 1,073 configured alerts, and Grafana AI handled an average of 63 terabytes of observability data per month.

Table 2
Grafana AI Environment

Environment	Average	Median
Number of applications and services monitored	2,607	646
Number of connected data sources	60	45
Number of dashboards managed	649	549
Number of alerts configured	1,073	1,250
Volume of observability data managed (TB/month)	63	28

n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

Business Value and Benefits of Grafana AI

IDC's analysis of the four organizations it interviewed about their use of Grafana AI identified measurable benefits across two value dimensions: IT benefits and engineering/SRE benefits. IT teams gained a leaner, easier-to-maintain platform. Engineering and SRE teams gained back time that used to disappear into manual querying, dashboard building, and routine troubleshooting. Participants described speed as a significant benefit of Grafana AI deployment. They gained agility by no longer needing to memorize a specific query language for every signal, and improved visibility helped them meet the SLA numbers they had promised customers. Investigating error log alerts got faster too since Grafana AI could explain who was impacted and where in the code a failure originated, then propose a resolution. Participants also pointed to a benefit that did not show up on a spreadsheet. Grafana AI provided access to non-experts, including managers who were unfamiliar with the platform. This gave them direct self-serve access to observability data they previously had to request from someone else.

Study participants described the most significant overall benefits of Grafana AI as:

Faster querying (Technology):

"Prior to Grafana AI, we had to remember a very specific query language for each signal. Now Grafana AI manages that for us, and it makes querying much faster."

Improved SLA performance (Technology):

"On the business side, Grafana AI let us expand our coverage to all our product lines and increase our overall efficiency managing that infrastructure. We were able to hit the SLA numbers we've promised customers because we now have better visibility."

Faster root cause analysis (Healthcare):

"One big benefit of Grafana AI is investigating error log alerts. We can ask it for more context right in Slack, and it tells us who is impacted, what's happening, and where in the code the failure is occurring. It also helps with root cause analysis and can propose a resolution."


Broader platform access (Software):

"Grafana AI speeds up triage and troubleshooting and improves dashboard creation. It has also enabled people who wouldn't normally be familiar with the platform, like managers, to ask questions and get the details they need."

Figure 1 (next page) quantifies those benefits at \$1.67 million (factoring in deployment time) per organization per year, comprising two distinct categories:

- Engineering/SRE benefits (\$1,027,264 per organization): A more productive engineering and SRE team, freed from routine querying, dashboard work, and troubleshooting, drove this category, the larger of the two.
- IT benefits (\$650,410 per organization): A leaner platform administration team and retired legacy monitoring tools drove this category, as Grafana AI absorbed the routine upkeep IT previously handled manually.

→ **Figure 1**
Average Annual Benefits per Organization



n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

IT Benefits

For the organizations IDC interviewed, Grafana AI impacted IT administration workflows and overall costs. Participants had often accumulated separate tools for metrics, logs, incident response, and front-end monitoring over time, each requiring its own upkeep. Grafana AI's open source, code-based approach took over the manual configuration work a proprietary system used to demand. It required less platform administration and management and, as a result, often enabled the redistribution of skilled staff into higher-value work. Others shut down scattered open source deployments entirely. That consolidated operations and engineering onto a single pane of glass and ended the runaway infrastructure costs those environments had quietly built up. OpenTelemetry also played a role, freeing one participant from a proprietary data-collection model that had trapped them with a single vendor. For a few organizations, cost was the deciding factor from the start: Grafana let them cover their entire product portfolio at a price their previous solution couldn't match.

Participants described that shift:

Reduced staffing needs (Healthcare):

"We needed almost double the staff to support our old platform because it required so much manual configuration work, so we've been able to reduce the team we need by 75%."

Consolidated observability (Software):

"The idea with Grafana Cloud and Grafana AI was to get everything onto a single pane of glass, so we shut down our open source environments and brought everything into the cloud. Previously, our operations team used a different tool, while engineering ran their own local open source Grafana deployments, and different teams saw different metrics, which caused a lot of confusion. Moving to the cloud also ended the runaway infrastructure costs of those open source environments, so it's been a big savings."

Reduced vendor lock-in (Technology):

"We moved from a proprietary data collection model that was essentially trapping us into staying with that company. Grafana's OpenTelemetry-based approach to collecting metrics reduced that lock-in, and we now stay because we think it's the best product for our price point, not because it's painful to switch. That's been a 15% reduction in the total cost of our observability solution."

Cost-driven migration (Technology):

"We moved to Grafana Cloud and Grafana AI mainly because of cost. We couldn't expand to all of our product lines at our old solution's sticker price, whereas with Grafana we could cover the entire landscape."

Interviewed organizations centralized fragmented tooling and legacy point solutions onto Grafana Cloud. Grafana AI's ability to unify analysis across signals supported that consolidation, allowing organizations to avoid the licensing and infrastructure costs of running multiple overlapping tools.

IDC calculated that this tool reduction was worth \$325,000 in annual cost avoidance per organization, stemming from five specific changes:

- Metrics platforms were retired once Grafana AI fully covered their capabilities.
- Log management software was replaced with natural language log search built into Grafana.
- On-call coordination was consolidated and moved into Grafana’s built-in incident response, retiring a third-party tool.
- Open source stacks were retired and consolidated, which cut license and cloud infrastructure costs.
- Renewal price increases were avoided by scaling monitoring affordably across the full portfolio.

Beyond retiring tools, Grafana AI took over much of the hands-on work of building and maintaining dashboards, alerts, and queries. By generating this work directly from natural language requests and running analysis in parallel with staff, Grafana AI reduced the platform administration effort needed to keep the environment running. As **Table 3 (below)** shows, IT teams were 52% more efficient with Grafana AI, freeing up time equivalent to 3.3 FTEs that could be redirected to higher-value work. This capacity was valued at \$327,222 per year.

Table 3
IT Team: Platform Administration and Management Efficiency Gains

Efficiency Gains	Before Grafana AI	With Grafana AI	Difference	Benefit
FTE count	6.3	3.1	3.2	52%
Value of staff time per year	\$633,333	\$306,111	\$327,222	52%

n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

Engineering/SRE Benefits

Engineering and SRE teams felt Grafana AI's impact most directly. These teams had previously carried the full weight of writing observability queries by hand and building and maintaining dashboards. They also walked newer engineers through workflows one question at a time. Grafana AI's natural language interface absorbed much of that day-to-day work, letting the same staff support more applications and take on higher-value engineering instead. Writing up a root cause analysis or a customer-facing status update went from a manual chore to something Grafana AI drafted directly. Teams benefited from the assistant living natively in Slack, so they could ask it questions and generate post-incident summaries without ever leaving the channel where an incident was already being discussed. Writing observability queries got faster too, with one participant describing it as a second brain that could investigate one thing while they looked elsewhere. Participants noted that Grafana AI's usability enabled new engineers to reach full productivity within a week instead of the extended ramp-up that used to require constant senior guidance.

Participants described Grafana AI's engineering benefits below:

Faster documentation (Technology):

"We now communicate more regularly, both internally and out to customers, using Grafana AI to build summaries of technical issues instead of doing that manually. These are relatively infrequent activities, but updating a support ticket or writing a root-cause analysis is now massively more efficient than before."

Slack-native workflow (Healthcare):

"We really like the Slack integration, since alerts and incident work happen in Slack and teams can interact with Grafana AI right there to get answers faster. The dashboard side panel and new workbench screens also make it easy to work in natural language, and we can generate post-incident investigation summaries very quickly."

Faster query writing (Software):

"Writing observability queries takes half the time with Grafana AI because it lets me have it investigate one thing while I'm looking elsewhere. I use it almost like a second brain."

Faster onboarding (Technology):

"With onboarding engineers, we've seen them up and running within a week because Grafana AI is so much easier to figure out and understand."

Building on these testimonials, IDC studied Grafana AI's impact on engineering and SRE productivity. Participants found that the solution absorbed much of the day-to-day investigation, query, and dashboard work that had consumed prior time. This allowed the same skilled staff to scale to support more applications and take on higher-value work. As **Table 4 (below)** shows, Grafana AI helped engineering and SRE teams work at a 31% higher productivity level, which amounted to the equivalent level of having 10.3 additional FTEs. That gain was worth \$1,030,125 per organization per year.

Table 4
Engineering/SRE Team Productivity Gains


Productivity Gains	Before Grafana AI	With Grafana AI	Difference	Benefit
Equivalent productivity level, FTEs	33.5	43.8	10.3	31%
Value of staff time per year	\$3,350,000	\$4,380,000	\$1,030,125	31%

n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

Grafana AI enabled these productivity gains by supporting routine operational tasks. The tool generated dashboards, triaged alerts, and wrote queries directly from natural language requests, which cut the most time from those activities that had previously required the most manual effort. Preparing status updates and stakeholder reports improved the most, at 73% less time, followed by building and maintaining dashboards at 64% and writing queries at 61%. **Figure 2 (next page)** details the full breakdown across six observability tasks.


New engineers felt that same effect from their first week on the job. Rather than requiring extended instruction from senior staff, new hires could ask Grafana AI directly and learn by doing. This significantly shortened the time it took them to operate independently. **Figure 3 (next page)** shows that new engineers achieved independent proficiency with workflows in 3.1 days on average, down from 10.0 days before Grafana AI, a significant 69% reduction.

Figure 2
Engineering/SRE Operational Key Performance Indicators



n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

Figure 3
Time to Independently Operate Grafana Workflows



n = 4; Source: IDC Business Value In-Depth Interviews, September 2026

That faster ramp-up fed directly into a separate benefit: fewer interruptions for the senior staff who used to field those questions. Grafana AI let nonexperts resolve routine observability questions on their own and absorbed the repetitive requests that had previously gone to senior SREs. A software organization noted, *"I get a lot fewer questions now, about 70% fewer random, low-level how-to questions, and the questions I do get are more in-depth."*

Across all the organizations IDC interviewed, this cut the volume of internal support requests directed at senior SREs and observability experts by 49% for three related reasons:

- **Plain-language access:**
Engineers, support staff, and even nontechnical managers could ask observability questions directly, without routing through a senior SRE.
- **Self-serve answers:**
Dashboards, queries, and configuration guidance came straight from Grafana AI's natural language responses, which were available immediately regardless of who asked or when.
- **Routine questions absorbed:**
Grafana AI took on the repetitive how-do-I and basic troubleshooting requests once escalated to senior staff, leaving only substantive technical questions for them.

Interviewed organizations noted that this time saving had a significant impact and often meant:

- **Senior time reclaimed:**
Freed from constant interruptions, senior SREs and observability experts redirected that reclaimed time to higher-value engineering and platform work.
- **No more waiting on a specific person's availability:**
Grafana AI responded immediately regardless of who asked or when.
- **Scale without headcount:**
Freeing senior staff from constant interruptions allowed organizations to support a larger, self-sufficient user base.

Operations and Incident Response Benefits

Grafana AI's impact on incident response came from a simple change. Instead of an engineer manually forming queries and searching across disconnected data sources, Grafana AI began correlating signals the moment an incident started. This gave responders a head start on almost every investigation, whether they were senior staff or not. One participant reached findings about 20 times faster once nobody had to stop and figure out how to phrase a log query. For another, that head start arrived before anyone even sat down. Feedback on what was happening, where, and who was affected often showed up in Slack before an engineer was hands on keyboard. That made Grafana AI a constant presence across nearly every investigation. It ran in parallel like a second brain and opened the platform to nonspecialist staff, including managers, who could query it directly.

Participants described how this changed day-to-day operations:

Faster investigations (Technology):

"Significantly more people can use Grafana AI to engage in investigations, and the quality of investigations is better. We also get to findings about 20 times faster because we're not stuck figuring out how to form log queries."

Always-on investigation (Healthcare):

"We use Grafana AI in almost every investigation and incident and often get feedback on what's happening, where, and who's affected right in Slack before we're even hands on keyboard. As people investigate, they use it actively to sort through metrics, logs, and profiles and look for patterns across data sources. It is pervasive."

Broader query access (Software):

"Our primary use case is investigation, triage, and troubleshooting, which is about 99.9% of our usage. Natural language querying lets us run investigations in parallel and acts like a second brain, and it also lets non-specialist staff, like managers, query the platform directly."

That head start compressed the entire incident life cycle. Grafana AI correlated signals across metrics, logs, and traces as soon as an incident began, often surfacing likely causes before an engineer was fully engaged. By running its own investigation in parallel with an engineer's manual troubleshooting, it cut the hands-on work needed to reach a resolution. Grafana AI's access to expert knowledge on demand prevented long-term disruption and ultimately shortened degraded-service windows.

As a result, interviewed organizations reduced mean time to resolution (MTTR) by 63% for four related reasons:

- **Real-time signal correlation:**
Grafana AI began correlating metrics, logs, and traces the moment an incident started, often surfacing what was happening and who was affected before an engineer was fully engaged.
- **Causes surfaced in Slack:**
Likely root cause hypotheses appeared directly in Slack, giving responders a starting point instead of manually searching across data sources.
- **Parallel investigation:**
Grafana AI ran its own investigation alongside an engineer's manual troubleshooting, catching what they missed or advancing the analysis independently.
- **Fewer manual queries needed:**
Responders could ask direct questions in natural language and get an immediate answer.

On-call engineers felt a related benefit during the shift itself. On-call staff no longer had to piece together what was happening, who was affected, and where an issue originated from by hand across dashboards and logs. Instead, they could ask Grafana AI directly for context on an active alert. This mattered because it let teams act faster and allowed a broader range of staff to address issues directly.

IDC determined that Grafana AI cut the share of an on-call shift spent on context gathering and manual correlation by 46% for four related reasons:

- **Direct queries to Grafana AI:**
On-call engineers could ask Grafana AI directly for context on an active alert instead of piecing it together manually across dashboards and logs.

- **Carried-forward pattern knowledge:**

On-call staff no longer needed one person's memory of past incidents to start investigating, since Grafana AI carried that pattern knowledge forward. A technology participant supported this sentiment: *"Grafana AI has lowered the experience level needed to respond to incidents so we can use more junior people. Previously you needed historic knowledge just to connect the dots, but now Grafana AI does that trend analysis for you. It recalls when we had the same issue 30 days or 6 months ago."*

- **On-demand senior-level detail:**

Senior-level context was available during the shift itself without escalating to first-line support or a specific senior colleague.

- **Broader on-call staffing:**

Once historical knowledge no longer lived only with a few senior people, organizations could staff on-call rotations more broadly.

Improved Service Reliability and Customer Trust

The benefit of faster, more visible incident handling with Grafana AI did not stop at the engineering team. Study participants connected it directly to their ability to meet service-level commitments and preserve customer trust. They found that it protected against relationship and reputational risk, which is much harder to recover from than a single missed metric.

Participants described that connection in their own words:

Record SLA performance (Healthcare):

"SLA plays a key role in customer retention, satisfaction, and upselling opportunities. Since onboarding Grafana AI alongside our platform stability improvements, SLAs have held steady at record highs, and customer retention and satisfaction have followed."

Stronger customer trust (Technology):

"The business sees improvement in signal to noise and in debugging back to customers, so we're building better customer trust, which drives retention because we're getting better information out of our observability."

Lower reputational cost (Software):

"For any product, the faster you can figure out what's causing errors or why it's breaking, the faster you can resolve it and the less cost or reputation cost you incur. So even if you're not seeing direct money savings, you're still saving money."

IDC identified that Grafana AI supported service reliability and customer trust in four ways:

- **Earlier detection:**
Grafana AI helped detect and address issues before they escalated far enough to reach or noticeably affect customers, protecting service level performance.
- **Better customer-facing investigations:**
It improved signal to noise in customer-facing investigations, giving teams more confidence and better information when resolving customer-reported problems.
- **Proactive, portfolio-wide analysis:**
The solution supported proactive analysis across an organization's entire customer base at a scale that would have been impractical to do manually, catching patterns before they became renewal risks.
- **Stronger SLA foundation:**
It gave organizations a stronger foundation for meeting SLA commitments consistently, which respondents connected directly to customer retention and upsell conversations.

Business Value Summary

IDC's interviews with four organizations using Grafana AI found that these benefits were worth an average of \$1.67 million per organization each year. That value came from three forces working together. Retiring redundant monitoring tools and running a leaner platform produced direct cost savings. A smaller IT team could then administer the Grafana environment without adding headcount. And engineering and SRE teams resolved incidents and completed routine work faster, which freed them to spend less time teaching others and more time on higher-value work.

For the organizations IDC interviewed, that combination meant observability kept pace with a growing environment without a growing team. Engineering and SRE capacity freed up by Grafana AI went toward higher-value work instead of routine troubleshooting. And faster, more visible incident handling fed directly into steadier SLA performance and stronger customer trust.

Across all four organizations, the consistent theme was a platform that made observability data usable by more people; it wasn't just more powerful for the few who already knew how to query it. A healthcare participant summarized the overall effect:

To summarize, Grafana AI has increased teams' abilities to self-serve on incident investigations, from low to high complexity, without cross-team dependencies. It speeds up dashboard creation, alerting, and reporting across the whole incident life cycle, from detection through investigation to writing up summaries afterward. The whole life cycle is faster and accessible to more people with Grafana AI.

Challenges/Opportunities

The capabilities that produce compression also shape whether it becomes real change. IDC assesses that against four principles, taken in dependency order: Later steps need the owners and intent, earlier ones supply, and each yields value alone.

- **Baselines describe the past:**

Descriptive authority lets people describe the desired future state and leaves the system to resolve the variables and behaviors that reach it. Past patterns suit sorting, and direction needs a described future. Watchers calibrate normal from current signal, and Adaptive Telemetry sets what the knowledge layer keeps from past query habits, so in an estate that changes daily both anchor to what it has been. SLOs trigger Investigations and can carry the description.

- **Controls are procedural:**

Rules, skills, and automations are instructions for known cases. The token cap, the only quantified limit on autonomous runs, bounds spend. A policy envelope also bounds scope, reversibility, blast radius, and named authority. Existing roles can carry autonomous, supervised, and reserved decision classes, and pull-request routing for write actions would supply the decision record audit logs omit.

- **Investigations stop at restoration:**

Recalling prior incidents shortens the next. Reducing the failure class takes a policy change someone owns. Routing findings into proposed skill, alert, and service-level objectives (SLOs) revisions, with human-led review beside the AI summary, would leave policy stronger after each cycle.

- **Broader orientation weakens on the capacity it depends on:**

Managers and junior engineers orient to incidents directly, without the sharpening and mentoring given by senior engineers. Correlation is safe to delegate, and reliance on it still erodes the skill needed to check it; participants who described freed capacity spent it on broader coverage, the signature of work intensification. Past investigations and the knowledge graph can supply practice scenarios, and judgment capacity belongs beside MTTR on the scorecard. Each remedy uses capabilities Grafana already ships, which leaves a question of sequence.

Conclusion

The four organizations IDC interviewed, with 615 to 1,400 employees in the healthcare, software, and technology industries, report compression at a scale that registers in financial terms. For these four, IDC's Business Value model calculates average annual benefits of \$1.67 million per organization across IT and engineering/SRE teams. IDC's analysis of the interviews shows a 63% reduction in mean time to resolution, 46% fewer on-call hours spent gathering context, and a 49% reduction in support requests to senior SREs and observability experts, and the organizations' new engineers reached independent operation of Grafana workflows in 3.1 days, down from 10.0. With four interviews, these results are directional.

Part of this value cannot be separated from the platform beneath the AI. Participants consolidated monitoring tools and retired point solutions while adopting Grafana AI, and at least one moved to OpenTelemetry-based collection, so four interviews cannot isolate one effect from another. The entanglement is a finding in its own right: The AI reads the estate it is given, and its value grows with the depth of the knowledge layer underneath, whether consolidated or federated.

The larger opportunity lies past compression. The challenges in this paper define a new operating loop: Describe the desired future state, bind agent action with policy, route investigation outcomes into lasting improvement, and maintain the human capacity that evaluates the machine. Grafana AI already ships components each step can use. Organizations that take these steps can build an operation whose capacity grows with its complexity. ●

Appendix:

IDC's Business Value Methodology

IDC's standard Business Value methodology gathers data from organizations currently using Grafana AI as the foundation for its model. IDC bases calculations on several assumptions, which are summarized as follows:

Multiplying time values by burdened salary quantifies efficiency and productivity savings. For this study, IDC uses its standard salary assumptions of an average fully loaded IT salary of \$100,000 per year and a user salary of \$70,000 per year. IDC applies a margin assumption (15%) for other user productivity gains and additional gross revenue attributed to interviewed organizations' use of Grafana AI, resulting in the net productivity and revenue calculations IDC applies to its model. Given the sample size of four organizations, findings should be treated as directional indicators of value rather than statistically representative averages.

Note: All numbers in this document may not be exact due to rounding.

About the IDC Analysts



Shannon Kalvar

Research Director,
Enterprise Systems Management, Enterprise Client Platforms,
Observability, and AIOps, IDC

Shannon Kalvar is research director for enterprise systems management, enterprise client platforms, observability, and AIOps at IDC. His current research addresses the evolution of digital work, AI in both end-user and IT operational contexts, and the transformation of IT management in a rapidly changing world where humans must master the exponential increase in data and complexity while navigating tumultuous economic and geopolitical times. Kalvar is a member of the Future of Work and AI councils, as well as a contributor to IDC's quantum computing research.

[More about Shannon Kalvar →](#)



Megan Szurley

Business Value Manager,
Business Value Strategy Practice, IDC

Megan Szurley is manager for the Business Value Strategy practice, responsible for creating custom business value research that determines the ROI and cost savings for enterprise technology products. Szurley's research focuses on the financial and operational impact of these products for organizations once deployed and in production. Prior to joining the Business Value Strategy practice, Szurley was a consulting manager within IDC's Custom Solutions division, delivering consultative support across every stage of the business life cycle: business planning and budgeting, sales and marketing, and performance measurement. In her position, Szurley partners with IDC analyst teams to support deliverables that focus on thought leadership, business value, custom analytics, buyer behavior, and content marketing. These customized deliverables are often derived from primary research and yield content marketing, market models, and customer insights.

[More about Megan Szurley →](#)

IDC Custom Solutions

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for external use and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

[IDC.com](#)[IDC on LinkedIn](#)[IDC Blog](#)

International Data Corporation (IDC) is the technology intelligence layer of the AI economy. A global leader in research and data for more than 60 years, IDC's expert analysts, proprietary datasets, and rigorous methodologies are trusted by business and IT leaders to guide critical business strategies and IT investments. Today, that intelligence is built into the tools and workflows where work gets done with IDC Quanta, making work sharper, teams faster, and businesses harder to beat.

©2026 IDC. Reproduction is forbidden unless authorized. All rights reserved. [CCPA](#)