



Observability

FORECAST _ 2024

Contents

3 Executive summary

5 Key findings

6 State of observability

8 Current deployment

8 › Observability capabilities deployed

12 › Number of monitoring tools

13 › Open-source usage

15 › Unified or siloed telemetry data

17 › Data integration

19 › Best practices employed

21 › Annual observability spend

23 Strategy and organization

24 › Single platform or multiple point solutions preference

26 › Observability vendor criteria

27 › Trends driving observability

29 › Purpose of observability

31 › Challenges preventing full-stack observability

32 Outages, downtime, and cost

33 › Outage causes

34 › Outage frequency

36 › Mean time to detection (MTTD)

38 › Mean time to resolution (MTTR)

40 › Total downtime

41 › Outage cost

43 › Detection of interruptions

44 › Time spent addressing disruptions

46 › MTTx change

50 › Influencers of lower MTTx by capability

51 › Downtime reduction

53 Observability benefits

54 › Primary observability benefits

55 › Total value of observability

56 › Return on investment for observability

57 Future of observability

58 Observability deployment plans

60 Data integration plans

62 Value maximization plans for observability

64 Summary

65 About this report


66 What's new

67 Definitions

71 Methodology

72 Demographics

73 Firmographics



Executive summary

The *2024 Observability Forecast* provides insights into the **evolution of observability**, identifying key areas of growth and stagnation, and uncovering how external forces are shaping adoption and investment strategies. With input from **1,700 technology professionals across 16 countries**, it stands as the largest and most comprehensive study in the observability industry.

With digital experiences and business growth at the forefront for businesses, the findings highlight the tangible business value of observability. IT professionals are seeking ways to reduce unplanned downtime, improve uptime, and boost reliability, all while managing key performance indicators (KPIs) through smarter investments in automation and preventative measures. The report shows that **organizations prioritizing observability have a significant advantage** when it comes to operational efficiency and overall business performance.



This year's data reveals that observability delivers a 4x median return on investment (ROI). With 79% less downtime and 48% lower outage costs for those using full-stack observability, the case for investing further in observability has never been stronger.

Furthermore, business observability—the ability to correlate telemetry data with business outcomes in real time—has emerged as a top priority. Organizations that adopt business observability see 40% less annual downtime, 24% lower hourly outage costs, and spend 25% less time managing disruptions compared to those without it.

In addition, the rising adoption of artificial intelligence (AI) technologies—such as AI monitoring, machine learning (ML) model tracking, and AI for IT operations (AIOps)—reflects the growing importance of observability in supporting innovation. Organizations that deploy AI-driven observability report higher overall business value and ROI.

In summary, the report confirms that observability is not just a technical practice, but a strategic imperative that drives measurable business outcomes. By investing in observability, organizations can ensure more reliable digital experiences, achieve operational efficiencies, and set the stage for future growth.

Key findings

The median annual downtime from a high-business-impact outage is 77 hours, or \$146 million.

The median total annual downtime across high-business-impact levels was 77 hours (approximately 3 days). That adds up to a median annual outage cost of \$146 million. On average, respondents estimated the median percentage of engineering team time spent addressing disruptions is 30%, which is 12 hours based on a 40-hour work week.

41% of respondents said they plan to consolidate tools in the next year.

There was a 2-to-1 preference for a single, consolidated platform compared to multiple point solutions. In fact, the number of respondents using a single tool increased by 37% year-over-year (YoY). And the average number of tools decreased by 11% YoY. While 45% were still using 5+ tools, 41% said they plan to consolidate tools in the next year.

Organizations deploying business observability experience 40% less annual downtime.

The ability to correlate business outcomes with telemetry data and report them in real time (business observability) was one of the most important observability vendor criteria—the third choice overall. In fact, 40% had deployed business observability. On average, those who had deployed business observability experienced 40% less annual downtime, spent 24% less on hourly outage costs, and spent 25% less time addressing disruptions compared to those who hadn't.

Organizations with full-stack observability experience 79% less downtime per year, saving \$42 million each year.

On average, those with full-stack observability experience 79% less downtime per year than those without (70 hours compared to 338 hours) and spend 48% less on outage costs per hour (\$1.1 million compared to \$2.1 million). There's also a strong association between less downtime and costs and several other factors.

Median observability ROI is 4x: doubling YoY.

The median annual observability spend across all respondents was \$1.95 million. However, the median annual value received from observability was \$8.15 million, and the median ROI was 4x. That means the median ROI doubled YoY from 2x to 4x. In addition, those who had deployed at least five observability capabilities estimated a higher annual value received and ROI from their observability investment than those with four or fewer deployed.

Organizations are embracing observability to capitalize on AI technologies.

The adoption of AI technologies was the top strategy or trend driving the need for observability (41%). About two in five (42%) had deployed AI monitoring, 29% machine learning (ML) model monitoring, and 24% AIOps capabilities. Notably, those who deploy these capabilities estimated a higher annual total value received from observability than those who hadn't deployed them.

State of observability

This section explores deployment trends, organizational strategies, outage impact, downtime costs, and the key benefits driving adoption.


The current state of observability reflects a growing emphasis on **optimizing technology investments** to drive **better business outcomes**. Organizations are shifting from fragmented monitoring practices and towards consolidated observability platforms.




While most organizations have not yet achieved full-stack observability, there was a significant year-over-year (YoY) increase in the deployment of observability capabilities. More organizations are nearing or achieving full-stack observability, which is key to unlocking its full potential.

Although outages remain a frequent and costly issue, observing more of the tech stack, achieving full-stack observability, and implementing observability best practices help organizations improve service-level metrics and get the most business value out of their investments.


Highlights:




62% said high-business-impact outages cost at least \$1 million per hour of downtime




45% were using 5+ tools for observability



41% said the adoption of AI technologies is driving the need for observability



34% said they receive \$10 million or more in annual value from their observability investment




25% had achieved full-stack observability

Current deployment

This section covers the observability capabilities deployed, how many tools were used for those capabilities, open-source usage for those capabilities, whether telemetry data is unified or siloed, what kinds of data are being integrated with telemetry data, best practices employed, the annual observability spend at the time of the survey, and how often respondents use observability.

Highlights:



Observability capabilities deployed

Capabilities, not to be confused with characteristics or tools, are specific components of observability. Survey respondents share which of 19 different observability capabilities they deployed. Below are findings by capability, by number of capabilities, and by how many have achieved full-stack observability.

- By capability
- By number of capabilities
- By how many have achieved full-stack observability

By capability

Survey respondents indicated their organizations deploy observability capabilities by as much as 58% (security monitoring) and as little as 24% (artificial intelligence for IT operations [AIOps] capabilities).

- At least half had deployed each of the core observability capabilities, including security monitoring (58%), network monitoring (57%), database monitoring (55%), alerts (55%), dashboards (54%), infrastructure monitoring (54%), log management (51%), and application performance monitoring (APM; 50%).
- More than a third had deployed key digital experience monitoring (DEM) capabilities—including browser monitoring (44%), error tracking (43%), and mobile monitoring (35%)—as well as AI monitoring (42%) and business observability (40%).
- Less than a third had deployed each of the more advanced capabilities, including AIOps capabilities (24%), synthetic monitoring (26%), distributed tracing (29%), Kubernetes (K8s) monitoring (29%), machine learning (ML) model monitoring (29%), and serverless monitoring (30%).

Organization size insight
 Large organizations were the most likely to deploy all capabilities except for AI monitoring, business observability, and serverless monitoring.

Regional insight
 Respondents surveyed in Asia Pacific were more likely to deploy AI monitoring, AIOps, and synthetic monitoring, but the least likely to deploy all other capabilities. Those surveyed in Europe were the most likely to deploy most capabilities.

Industry insight
 IT respondents were generally more likely than average to deploy most capabilities. Media/entertainment respondents were the most likely to deploy AI-related capabilities and DEM capabilities.




Figure 01. Current capability deployment

■ 2024 respondents

By number of capabilities

Survey respondents said their organizations deploy eight capabilities on average. Three-quarters (75%) had deployed at least five capabilities, including 37% who had deployed 10 or more and 10% who had deployed 15 or more.

37%
had deployed **at least 10**
observability capabilities

Organization size insight
Large organizations were the most likely to deploy 10 or more capabilities (40%), followed by midsize (35%) and small (27%).

Regional insight
Respondents surveyed in Europe were the most likely to deploy 10 or more capabilities (46%), followed by the Americas (42%) and Asia Pacific (29%).

Industry insight
IT respondents were the most likely to deploy 10 or more capabilities (53%), followed by healthcare/pharma (48%) and services/consulting (43%).




Figure 02. Number of capabilities currently deployed


Those who had deployed at least five capabilities were more likely than average to experience less annual downtime, spend less on outages per year, and spend less time addressing disruptions than those who had deployed four or fewer:

- **5+ capabilities:** 45% lower median annual downtime, and 24% less engineering time spent addressing disruptions
- **10+ capabilities:** 74% lower median annual downtime, 32% lower median hourly outage costs, and 41% less engineering time spent addressing disruptions
- **15+ capabilities:** 80% lower median annual downtime, 47% lower median hourly outage costs, and 39% less engineering time spent addressing disruptions

Deploying more observability capabilities is associated with better business outcomes.

Full-stack observability prevalence

Based on our definition of full-stack observability, a quarter (25%) of survey respondents' organizations had achieved it.



Notably, organizations that had achieved full-stack observability experienced 79% less median annual downtime, 48% lower median hourly outage costs, and 44% less time spent addressing disruptions than those that hadn't achieved it. They also had a 27% lower annual observability spend and were 51% more likely to learn about interruptions with observability. And they were more likely to employ all observability best practices and experience most benefits and business outcomes.

Organization size insight

Large organizations were the most likely to have achieved full-stack observability (27% compared to 23% for midsize and 20% for small).

Regional insight

Those surveyed in Europe were the most likely to have achieved full-stack observability (32% compared to 29% for Asia Pacific and 28% for the Americas).

Industry insight

IT respondents were the most likely to have achieved full-stack observability (35%), followed by healthcare/pharma (34%) and services/consulting (31%). Education respondents were the least likely to have achieved full-stack observability (11%), followed by telco (14%), energy/utilities (15%), and government (15%).


Figure 03. Proportion of 2024 respondents who had achieved or not achieved full-stack observability

Those that have achieved full-stack observability generally employ more best practices and experience better business outcomes.

Number of monitoring tools

When asked about the number of tools, not to be confused with capabilities or characteristics, they use to monitor the health of their systems, survey respondents overwhelmingly reported using more than one.

- Most (88%) were using multiple tools, including 45% who were using five or more tools (compared to 52% in 2023 and 73% in 2022) and 3% who were using 10 or more tools.
- The average (mean) number of tools used was 4.5, which is 11% fewer than in 2023 (5.1) and 24% fewer than in 2022 (5.9). Similarly, the median number of tools went from six in 2022, to five in 2023, to four in 2024. The most common answer (mode) for 2024 was three tools (18%), followed by five tools (15%).
- Only 6% used just one tool. However, the proportion of respondents using a single tool increased by 37% year-over-year (YoY).



Compared to those using multiple tools for observability, those using a single tool experienced the following benefits:

- 65% lower median annual observability spend (\$700,000 compared to \$2 million)
- 18% less median annual downtime (249 hours per year compared to 305 hours per year)
- 45% less on median hourly outage costs (\$1.1 million per hour compared to \$2.0 million per hour)
- 50% less engineering time spent addressing disruptions (about seven hours compared to 13 hours based on a 40-hour work week)

Organization size insight

Small organizations were much more likely to use a single tool (17%) than midsize (7%) or large (4%) organizations, while large and midsize organizations were much more likely to use 5+ tools (47% for both compared to just 26% for small).

Regional insight

Respondents surveyed in Europe were the most likely to use a single tool (8% compared to 6% for those in both the Americas and Asia Pacific), while those surveyed in Asia Pacific were the most likely to use 5+ tools (55% compared to 43% for those in Europe and 35% for those in the Americas).

Industry insight

Healthcare/pharma respondents were the most likely to use a single tool (13%), followed by education (10%) and energy/utilities (9%). Media/entertainment respondents were the most likely to use 5+ tools (60%), followed by financial services/insurance (57%) and telco (55%).

45% were using 5+ tools for observability

Figure 04. Number of monitoring tools used for observability in 2022, 2023, and 2024

- 2024 respondents
- 2023 respondents
- 2022 respondents

There's a clear multi-year trend toward using fewer tools. We expect this trend to continue as 41% said they plan to consolidate tools in the next year.

Open-source usage

When we asked survey takers whether they were using an open-source solution in addition to a proprietary solution for each of the 19 observability capabilities listed above, we found that:

- More than half (51%) of respondents were using an open-source solution for one or more observability capabilities. But only about 1% were using *only* open-source.
- Of the three open-source solutions included in the survey, 38% were using Grafana, 23% were using Prometheus, and 19% were using OpenTelemetry for one or more observability capabilities.
- More than a quarter were using an open-source solution for AI monitoring (31%), synthetic monitoring (28%), distributed tracing (28%), K8s monitoring (27%), APM (27%), and AIOps capabilities (26%).

Organization size insight

Large organizations were much more likely to use an open-source solution for one or more observability capabilities (55%) compared to midsize (46%) and small (39%) organizations.

Regional insight

Respondents surveyed in Asia Pacific were much more likely to use an open-source solution for one or more observability capabilities (61%) compared to those surveyed in the Americas (44%) and Europe (42%).

Industry insight

Government respondents were the most likely to use an open-source solution for one or more observability capabilities (65%), followed by telco (65%) and financial services/insurance (61%). Services/consulting respondents were the least likely (37%), followed by energy/utilities (43%) and healthcare/pharma (45%).




Figure 05. Open-source solution usage by observability capability

■ 2024 respondents

51%

were using an open-source solution for one or more observability capabilities, but only **about 1% were using only open-source solutions.**

Respondents were least likely to select open-source support (portability) as the most important observability vendor criteria (16%), and only 21% said that adoption of open-source technologies is a strategy or trend driving the need for observability. However, 24% said they are most likely to use open source in the next year to maximize value from their observability investment.


Many organizations start by using open-source solutions for observability to avoid licensing costs and benefit from its wide adoption and support by a vibrant community. However, most use open source in tandem with a proprietary observability solution.

Which open-source solution they use the most is associated with how long that solution has been in the market and its primary function.

For example, Grafana—best known for its dashboarding solution—is the most used open-source solution and has been around the longest (since 2014).

Prometheus—the ubiquitous time series database and metrics monitoring tool—is the second most used and has been around since 2016.

And OpenTelemetry—a set of APIs, SDKs, and tools that collects and exports telemetry data to a proprietary observability solution or visualization tool—is the newest (since 2019) and least used, but it's experiencing massive growth and adoption as it becomes the standard protocol for open-source telemetry.



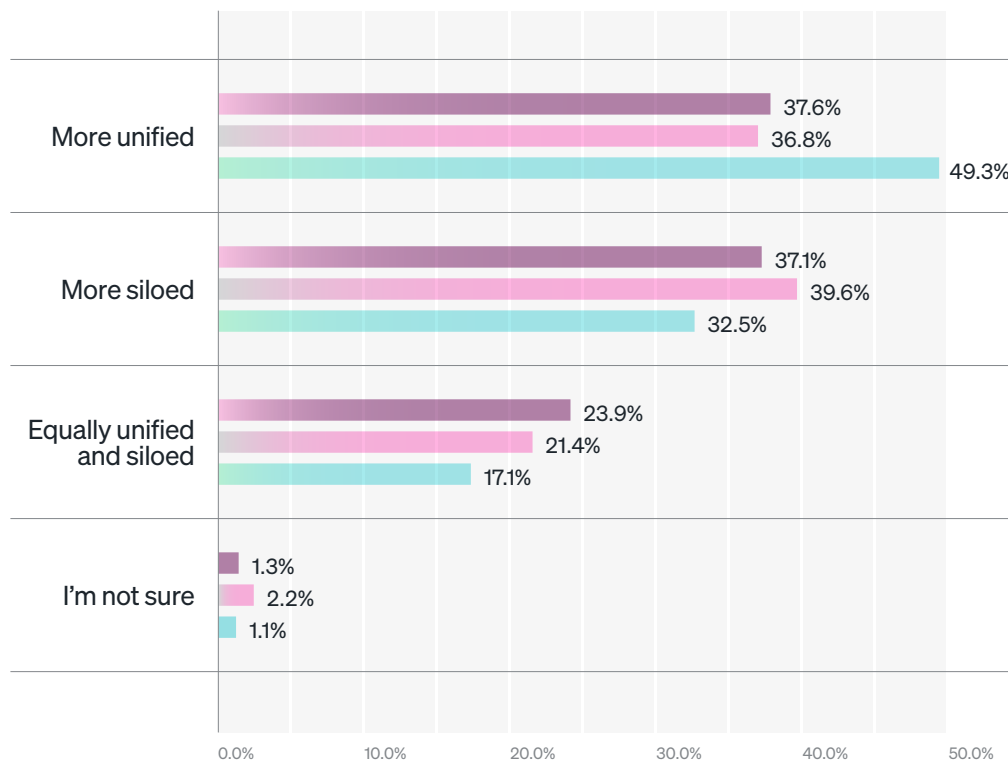
“I’m managing more open-source technologies because the industry is moving to an eBPF or OpenTelemetry type of approach. But to do that, I have to use a collector. I’m using New Relic as the collector for OpenTelemetry, and New Relic provides AIOps capabilities too.”

Senior Director of IT Infrastructure,
Large Fintech Enterprise, United States

Unified or siloed telemetry data

When we asked survey respondents about how unified or siloed their organizations' telemetry data (metrics, events, logs, and traces, or MELT) is, we found:

- Collectively, 38% had more unified telemetry data (increased by 2% from 2023), compared to 37% with more siloed telemetry data (decreased by 6% from 2023)—a roughly even split.
- Only 12% said they had mostly unified telemetry data (they unify telemetry data in one place), and 11% said they had mostly siloed telemetry data (they silo telemetry data in discrete data stores).
- About a quarter (24%) said their telemetry data is roughly equally unified and siloed (increased by 12% from 2023).



Those with five or more tools were 13% more likely to say they have siloed telemetry data to some extent (64%) compared to those with one to four tools (57%).

Compared to respondents with more siloed telemetry data, those with more unified telemetry data on average.:

- Experienced 78% less annual downtime (107 hours per year compared to 488 hours per year)
- Spent 11% less engineering time addressing disruptions (28% compared to 32%)
- Had a 4% higher median ROI (302% compared to 290%)

Organization size insight

Large organizations were the most likely to have more siloed telemetry data (38%), followed by midsize (37%) and small (33%) organizations.

Regional insight

Respondents surveyed in Asia Pacific reported more siloed telemetry data (42%) than those in Europe (39%) or the Americas (30%).


Industry insight

The industries with the highest rates of siloed telemetry data were government (53%), media/entertainment (47%), and education (45%). Those with the highest rates of unified telemetry data were telco (53%), retail/consumer (52%), and services/consulting (41%).

Figure 06. Unified versus siloed telemetry data in 2022, 2023, and 2024

- 2024 respondents
- 2023 respondents
- 2022 respondents

61%
said their telemetry data is siloed
to some extent



Telemetry data is roughly equally unified and siloed this year. The data shows an association between using more tools for observability and having more siloed data. It also shows that more unified data leads to more desirable business outcomes, including less downtime and a higher ROI.

Figure 07. Telemetry data unification in 2024

Data integration

To practice true business observability, organizations must integrate their business-related data with their telemetry data (MELT). In reviewing the types of business-related data they said they currently integrate:

- Most (87%) had integrated at least one business-related data type with their telemetry data, including 77% who'd integrated at least two and 35% who'd integrated at least five. Just 4% had integrated all 10.
- Operations data (43%) and customer data (41%) were the most likely to be integrated.
- Product research and human resources data (both 32%) were the least likely to be integrated.

Regional insight

Respondents surveyed in Europe were the most likely to integrate 5+ types of business-related data with their telemetry data (39% compared to 34% of those in both the Americas and Asia Pacific).

Industry insight

IT respondents were the most likely to have 5+ types of business-related data integrated with their telemetry data (47%), followed by media/entertainment (41%) and healthcare/pharma (38%). Education respondents were the least likely (19%), followed by energy/utilities (25%) and government (27%).




Figure 08. Types of business-related data currently integrated with telemetry data

Compared to those who had less than five business-related data types currently integrated with their telemetry data, those who had integrated five or more:

- Spent 32% less on hourly outage costs (\$1.5 million compared to \$2.2 million)
- Experienced 63% less annual downtime (139 hours compared to 370 hours)
- Spent 27% less engineering time addressing disruptions (11 hours compared to 15 hours based on a 40-hour work week)

35%
had integrated **5+ business-related data types** with their telemetry data

IT professionals realize the importance of using telemetry data to better understand real-world impacts on business outcomes (business observability).




Figure 09. Number of business-related data types integrated with telemetry data

Best practices employed

We once again asked survey takers which of nine different observability best practices listed in the chart below they employ. We found that:

- Most (83%) had employed at least two best practices, but only 16% had employed five or more.
- Respondents were most likely to say their software deployment uses CI/CD practices (40%) and their infrastructure is provisioned and orchestrated using automation tooling (39%)—but less likely than previous years.
- Compared to last year, 24% more said their telemetry data includes rich metadata and business context to quantify the business impact of events and incidents, 18% more said users broadly have access to telemetry data and visualizations, 13% more said their telemetry is unified in a single pane for consumption across teams, 8% more said their telemetry is captured across the full tech stack, and 2% more said they can query data on the fly.



Organization size insight

Small organizations were more likely to employ 5+ best practices (24%) than large (17%) and midsize (11%) organizations.



Regional insight

Respondents surveyed in the Americas were the most likely to employ 5+ best practices (21%) than those in Asia Pacific (13%) and Europe (12%).



Industry insight

Services/consulting respondents were the most likely to say they employ 5+ best practices (23%), followed by financial services/insurance (21%) and healthcare/pharma (20%).





Figure 10. Best practices employed

2024 respondents 2023 respondents 2022 respondents

On average, compared to those employing one to four, those employing five or more observability best practices:

- Experienced 19% less annual downtime (239 hours compared to 294 hours)
- Spent 35% less on hourly outage costs (\$1.3 million compared to \$2.0 million)
- Spent 38% less engineering time addressing disruptions (21% compared to 34%)
- Were 36% more likely to say MTTD improved to some extent since adopting an observability solution (72% compared to 53%)
- Were 38% more likely to say MTTR improved to some extent since adopting an observability solution (77% compared to 56%)
- Spent 20% less on observability per year (\$1.6 million compared to \$2.00 million)



Organizations should consider employing more best practices as it's strongly associated with better business outcomes, including less downtime and lower costs.

Figure 11. Count of observability best practices employed in 2024

Annual observability spend

The median annual spend on observability per year was **\$1.95 million**. More than two-thirds (67%) were spending at least \$1 million per year on observability, and only 2% were spending \$5 million or more. Just 13% were spending less than \$500,000 per year.




Figure 12. Annual observability spend by annual revenue

- <\$1M
- <\$1M-\$1.99M
- <\$2M-\$2.99M
- <\$3M-\$3.99M
- <\$4M-\$4.99M
- ≥\$5M
- I'm not sure

67%
were spending **at least \$1 million**
per year on observability

Seven factors were associated with a lower median annual observability spend, including:

- ✓ **Using a single tool for observability:** Those using a single tool for observability spent 67% less on observability than those using two or more tools (\$700,000 compared to \$2.00 million).
- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less they spent on observability. For example, those who had deployed five or more observability capabilities spent 13% less on observability per year than those with four or fewer (\$1.90 million compared to \$2.18 million). Those who had deployed 10 or more observability capabilities spent 30% less on observability per year than those with nine or fewer (\$1.50 million compared to \$2.15 million).
- ✓ **Achieving full-stack observability:** Those with full-stack observability spent 27% less on observability per year than those without full-stack observability (\$1.50 million compared to \$2.05 million).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability spent 23% less on observability per year than those who didn't (\$1.70 million compared to \$2.2 million).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices spent 20% less on observability per year than those employing four or fewer (\$1.60 million compared to \$2.00 million).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data spent 10% less on observability per year than those who integrated one to four types (\$1.85 million compared to \$2.05 million).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data spent 5% less on observability per year than those who had more siloed telemetry data (\$1.90 million compared to \$2.00 million).

Organization size insight

As expected, those from large organizations reported higher median spend (\$2.20 million) than those from midsize (\$1.85 million) and small (\$650,000) organizations.

Regional insight

Those surveyed in Asia Pacific reported a higher median spend (\$2.50 million) than those in Europe (\$1.75 million) or the Americas (\$1.30 million).

Industry insight

Media/entertainment respondents reported the highest median annual observability spend (\$2.60 million), followed by financial services/insurance (\$2.50 million) and telco (\$2.35 million). Education respondents reported the lowest spend (\$1.00 million), followed by healthcare/pharma (\$1.20 million) and services/consulting (\$1.40 million).




Those who observe more of the tech stack, employ more observability best practices, use a single tool for observability, and learn about interruptions with observability actually spend less on observability.

Strategy and organization

This section looks at the preference for a single, consolidated platform or multiple point solutions, the most important criteria when choosing an observability vendor or solution, the strategies and trends driving the need for observability, whether observability is more of a key enabler to achieving core business goals or for incident response/insurance, and the challenges preventing full-stack observability.

Highlights:



53%

preferred a single, consolidated observability platform




41%

said the adoption of AI technologies is driving the need for observability




32%

said breadth of features is the most important observability vendor criterion



Single platform or multiple point solutions preference

When it comes to the number of tools in use for observability, more than half (53%) preferred a single, consolidated platform to some extent, which is about the same as last year. More than a quarter (27%) preferred multiple point solutions to some extent, which is 6% less than last year, and one in five (20%) had no preference, which is 24% more than last year.



Organization size insight

Those from small organizations were more likely to prefer a single, consolidated platform to some degree (56%) than those from midsize (53%) and large (52%) organizations.

Regional insight

Respondents surveyed in Europe were more likely to prefer a single platform to some degree (58%) than those in the Americas (52%) and Asia Pacific (51%).

Industry insight

Education respondents were the most likely to prefer a single, consolidated observability platform (64%), followed by IT and government (both 60%). Healthcare/pharma respondents were the most likely to prefer multiple point solutions (36%), followed by services/consulting (34%) and telco (33%).

Figure 13. Observability preferences in 2022, 2023, and 2024

- 2024 respondents
- 2023 respondents
- 2022 respondents

53%


preferred a single, consolidated observability platform

For the second year in a row, there's an almost 2-to-1 preference for a single platform over multiple point solutions. This makes sense since there was an association between using a single tool for observability and spending less on observability, experiencing less downtime, spending less on outage costs, and spending less time addressing disruptions.

Even though more respondents said that they prefer a single, consolidated platform, 88% were using two or more monitoring tools, and just 6% were using a single tool for observability.

When asked what challenges prevent them from achieving full-stack observability, more than a third (34%) said they have too many monitoring tools and siloed data.

While tool sprawl persists, there's a clear multi-year trend toward using fewer tools. In fact, the number of respondents using a single tool increased by 37% year-over-year (YoY). And the average number of tools decreased by 11% YoY. In addition, 41% said they plan to consolidate tools in the next year.




Observability vendor criteria

Nearly a third of respondents said breadth of features (32%) and affordability (31%) were the most important criteria when choosing an observability vendor or solution. In addition, more than a quarter said business observability (27%), deployment model (27%), and learning and help (26%).

32% said **breadth of features** is the most important observability vendor criterion

While breadth of features, affordability, and business observability were the top three observability vendor criteria overall, the top choices varied greatly by role, organization size, region, and industry.

Figure 14. Most important observability vendor criteria



Organization size insight

Those from midsize and large organizations were more likely to cite cost as the most important criterion (both 32%) than those from small organizations (24%).

Regional insight

Cost (affordability) was the top choice for respondents surveyed in Europe (39%) and the Americas (38%), but was notably a lower priority (seventh choice) among those surveyed in Asia Pacific (22%). Business observability was the top choice for those surveyed in Asia Pacific (32%) compared to fifth choice for those in the Americas (25%) and seventh choice for those in Europe (21%).

Industry insight

Cost (affordability) was the top choice for education (43%), services/consulting (43%), IT (38%), energy/utilities (32%), and retail/consumer (31%) respondents. The top two criteria for telco were business observability and best-in-class offering and reputation (both 39%). Business observability was also the top choice for financial services/insurance, along with breadth of features (both 34%). Breadth of features was also the top choice for media/entertainment, along with scalability (both 34%). Ease of implementation and deployment model were the top criteria for government (both 28%).

Trends driving observability

The data shows that the adoption of artificial intelligence (AI) technologies and an increased focus on security, governance, risk, and compliance were the most commonly cited drivers for observability (both 41%). About a third of survey respondents cited integration of business apps into workflows (35%), cost management (33%), and the development of cloud-native application architectures (31%).

41% said the adoption of **AI technologies** is driving the need for observability



Organization size insight


Respondents from small organizations were less likely to cite migration to a multi-cloud environment (17% compared to 25% for midsize and 31% for large), cost management (27% compared to 33% for large and 34% for midsize), adoption of AI technologies (34% compared to 41% for large and 44% for midsize), and an increased focus on security, governance, risk, and compliance (31% compared to 39% for midsize and 44% for large).

Regional insight

Respondents surveyed in Asia Pacific were less likely to cite cost management (25% compared to 35% for Europe and 42% for the Americas), the adoption of AI technologies (36% compared to 41% for Europe and 48% for the Americas), and an increased focus on security, governance, risk, and compliance (34% compared to 43% for Europe and 48% for the Americas).

Industry insight

The adoption of AI technologies or an increased focus on security, governance, risk, and compliance were the top drivers for respondents from all industries except one—the integration of business applications into workflows was the top driver for energy/utilities respondents (43%).




With the exception of the adoption of AI technologies and cost management, all other strategies and trends driving observability decreased YoY. This data reflects the growing interest in AI.

Figure 15. Technology strategies and trends driving the need for observability in 2022, 2023, and 2024

- 2024 respondents
- 2023 respondents
- 2022 respondents

Purpose of observability

Half (50%) of survey takers thought observability is more of a key enabler for achieving core business goals—a 25% increase YoY. In addition, nearly a third (30%) indicated that observability enables business goals and incident response equally in their organizations—a 5% decrease YoY, with one in five (20%) saying observability is more for incident response or insurance—a 25% decrease YoY.



Collectively, 79% said observability is a key enabler for achieving their organization’s core business goals to some extent (compared to 71% in 2023 and 78% in 2022). While 49% said it’s for incident response or insurance to some extent (compared to 57% in 2023 and 49% in 2022).

Organization size insight

Those from small organizations were the most likely to say observability is more for core business goals (59%), followed by those from midsize (51%) and large (47%) organizations.

Regional insight

Those surveyed in Asia Pacific were the most likely to say observability is more for core business goals (57%), followed by those in Europe (47%) and the Americas (42%).

Industry insight

Government respondents were the most likely to say observability is for core business goals (65%), followed by media/entertainment (62%) and IT (59%) respondents. Education respondents were the most likely to say it’s more for incident response or insurance (32%), followed by healthcare/pharma and energy/utilities respondents (both 27%).

Figure 16. Unified versus siloed telemetry data in 2022, 2023, and 2024

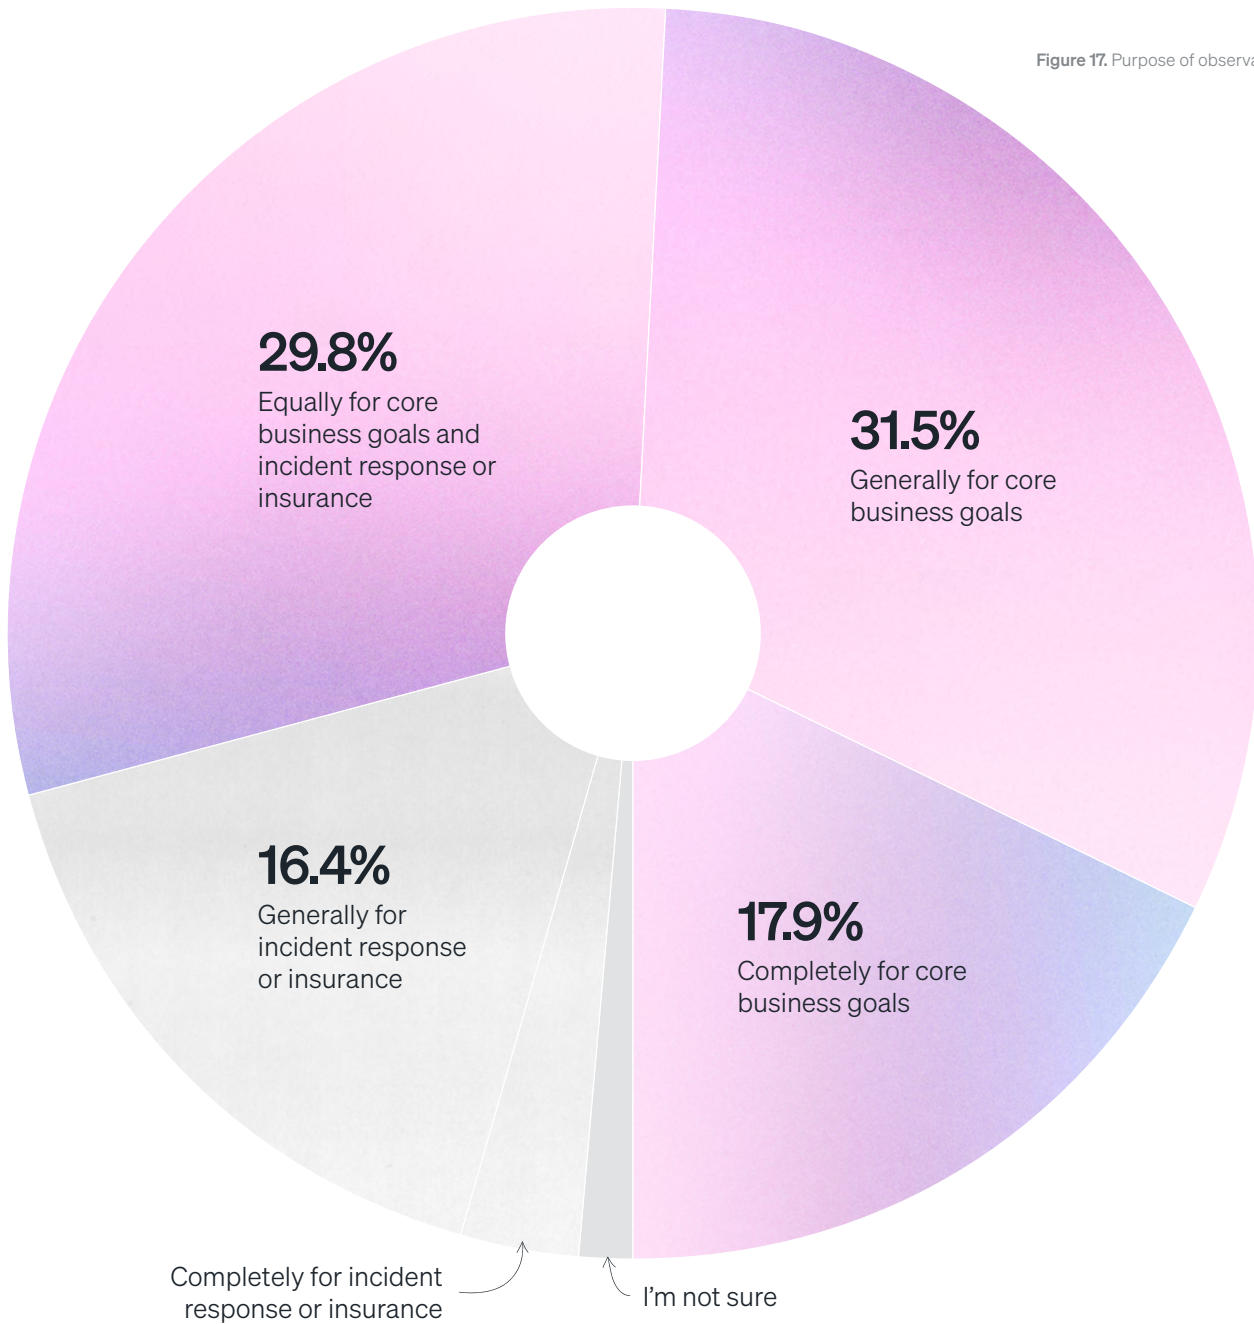
- 2024 respondents
- 2023 respondents
- 2022 respondents



79% indicated observability is a **key enabler for achieving core business goals** to some extent

These results indicate a clear shift among respondents viewing observability as a key enabler for achieving core business goals rather than just for incident response or insurance.

Figure 17. Purpose of observability



Challenges preventing full-stack observability

When examining what’s preventing organizations from achieving full-stack observability, more than a third of respondents said a complex tech stack and too many monitoring tools and siloed data (both 34%). About a quarter cited a lack of budget (27%), having adequate IT performance (25%), a lack of strategy (25%), too expensive (25%), and resistance to change (24%), with just 4% claimed to have already achieved full-stack observability.

🏢 Organization size insight

Respondents from large organizations were the most likely to struggle with too many monitoring tools and siloed data as well as a complex tech stack. Those from midsize organizations were the most likely to say it’s too expensive. And the top challenge for those from small organizations was a lack of budget.

🌍 Regional insight

Respondents surveyed in Europe were the most likely to say it’s too expensive, but least likely to cite too many monitoring tools and siloed data. Those surveyed in the Americas and Asia Pacific were more likely to struggle with a lack of strategy and resistance to change. Those surveyed in Asia Pacific were also more likely to say they don’t have the skills and that their IT performance is adequate.


🏭 Industry insight

The top two challenges for those from most industries was a complex tech stack and too many monitoring tools and siloed data—with a few exceptions. Education respondents struggled most with a lack of budget (51%) and cost (35%). Lack of budget was also the second choice for government (32%), telco (32%), and services/consulting (28%) respondents. And the second choice for media/entertainment respondents was a lack of strategy (34%).

Figure 18. Challenges preventing organizations from achieving full-stack observability

These results suggest a number of different hurdles and pain points when it comes to achieving full-stack observability, with tool sprawl, data silos, and complex tech stacks topping the list.

34% said too many monitoring tools and siloed data are barriers to achieving full-stack observability



Note: There were several changes to answer options from 2023 to 2024 for this question, so YoY, apples-to-apples comparisons are difficult to make.

Outages, downtime, and cost

Developers and engineers often use observability to solve three key business and technical challenges: reducing downtime, reducing latency, and improving efficiency.

Outage frequency, mean time to detection (MTTD), and mean time to resolution (MTTR) are common metrics used in security and IT incident management.

This section covers outage causes, frequency, and costs; and MTTD and MTTR trends.

Highlights:



62%

said high-business-impact outages cost at least \$1 million per hour of downtime



59%

said their MTTR improved to some extent since adopting observability



38%


experienced high-business-impact outages at least once a week

Promo Service (3s)
Service - APM
ECOMMERCE
Backend
Golden signals
This entity is:
hosted by 6 hosts
Show dependencies

Outage causes

More than a third (35%) of respondents said network failure was the most common cause of unplanned outages at their organization in the last two years. More than a quarter said third-party or cloud provider services failure (29%), someone making a change to the environment (28%), and deploying software changes (27%) were the most common causes.

35% said **network failure** was the most common cause of unplanned outages in the last two years



Organization size insight

Those from large organizations were more likely to say someone making a change to the environment was a top cause (31% compared to 24% for both small and midsize). Those from small organizations were more likely to cite capacity constraints (26% compared to 19% for large and 16% for midsize).

Regional insight

Those surveyed in Europe and the Americas were more likely to say someone making a change to the environment was a common cause (32% and 31% respectively compared to 23% for those in Asia Pacific). Those surveyed in Asia Pacific were more likely to contend with capacity constraints (21% compared to 18% for those in the Americas and 15% for those in Europe) and unexpected traffic surges (22% compared to 18% for those in the Americas and 16% for those in Europe).

Industry insight

Network failure wasn't the top choice for all industries. It was tied for first place with security failure for government respondents (34%), and with power failure for media/entertainment respondents (32%). Power failure was also the top choice for energy/utilities respondents (35%).

Figure 19. Most common causes of unplanned outages in the last two years

More than a quarter attributed the most common causes of unplanned outages at their organization in the last two years to human error (someone making a change to the environment or deploying software changes). But most attributed these outages to forces beyond their control.

Outage frequency

When we asked survey takers how often they experience low-, medium-, and high-business-impact outages, the median annual outage frequency across all business impact levels was **232 outages**. Low-business impact outages occurred the most frequently—more than half (57%) experienced them at least once a week, and 15% dealt with them daily. While high-business-impact outages happened the least frequently, 38% still experienced them at least once per week, and 12% said they occur at least once per day.

Organization size insight
 Those from small organizations experienced substantially more outages per year (410) compared to those from large (234) and midsize (183) organizations.

Regional insight
 Those surveyed in the Americas experienced the fewest outages per year (94) compared to those surveyed in Europe (207) and Asia Pacific (272).

Industry insight
 Government organizations experienced the most outages per year (419), followed by media/entertainment organizations (413). Services/consulting organizations experienced the fewest outages per year (55), followed by retail/consumer organizations (118).




Figure 20. Outage frequency across all business impact levels

38%

experienced high-business-impact outages at least once a week

Seven factors were associated with less frequent outages, including:

- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data experienced 77% fewer annual outages than those who had more siloed telemetry data (96 outages compared to 409 outages).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability experienced 71% fewer outages per year than those who hadn't (74 outages compared to 252 outages).
- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the fewer outages they experienced per year. For example, those who had deployed five or more observability capabilities experienced 47% fewer annual outages than those who had deployed four or fewer (196 outages compared to 370 outages). Those who had deployed 10 or more experienced 62% fewer annual outages than those who had deployed nine or fewer (96 outages compared to 252 outages). And those who had deployed 15 or more experienced 69% fewer annual outages than those who had deployed 14 or fewer (74 outages compared to 234 outages).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability experienced 69% fewer annual outages than those who used more manual detection methods (114 outages compared to 366 outages).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data experienced 47% fewer annual outages than those who integrated one to four types (134 outages compared to 252 outages).
- ✓ **Using a single tool for observability:** Those using a single tool for observability experienced 9% fewer annual outages than those using multiple tools (214 outages compared to 234 outages).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices experienced 8% fewer annual outages compared to those who had employed four or fewer (214 outages compared to 232 outages).



Although outages happen fairly frequently, full-stack observability and other factors have an enormous positive impact on outage frequency.

Mean time to detection (MTTD)

The mean time to detect an outage is a common service-level metric used in security and IT incident management. The data shows that the median number of hours spent on MTTD per year across all business impact levels was **134 hours**—which is approximately six days. The median MTTD for high-business-impact outages was 37 minutes, and more than a quarter (29%) of respondents said MTTD was an hour or more for high-business-impact outages.

Organization size insight

On average, midsize organizations spent less time detecting outages per year (101 hours) than large (138 hours) and small (163 hours) organizations.

Regional insight

On average, those surveyed in Asia Pacific spent the most time detecting outages per year (219 hours), followed by those surveyed in Europe (110 hours) and the Americas (42 hours).

Industry insight

The industries that spent the least amount of time detecting outages per year included services/consulting (23 hours), retail/consumer (61 hours), and education (64 hours). The industries that spent the most amount of time detecting outages per year included media/entertainment (331 hours), government (269 hours), and financial services/insurance (227 hours).

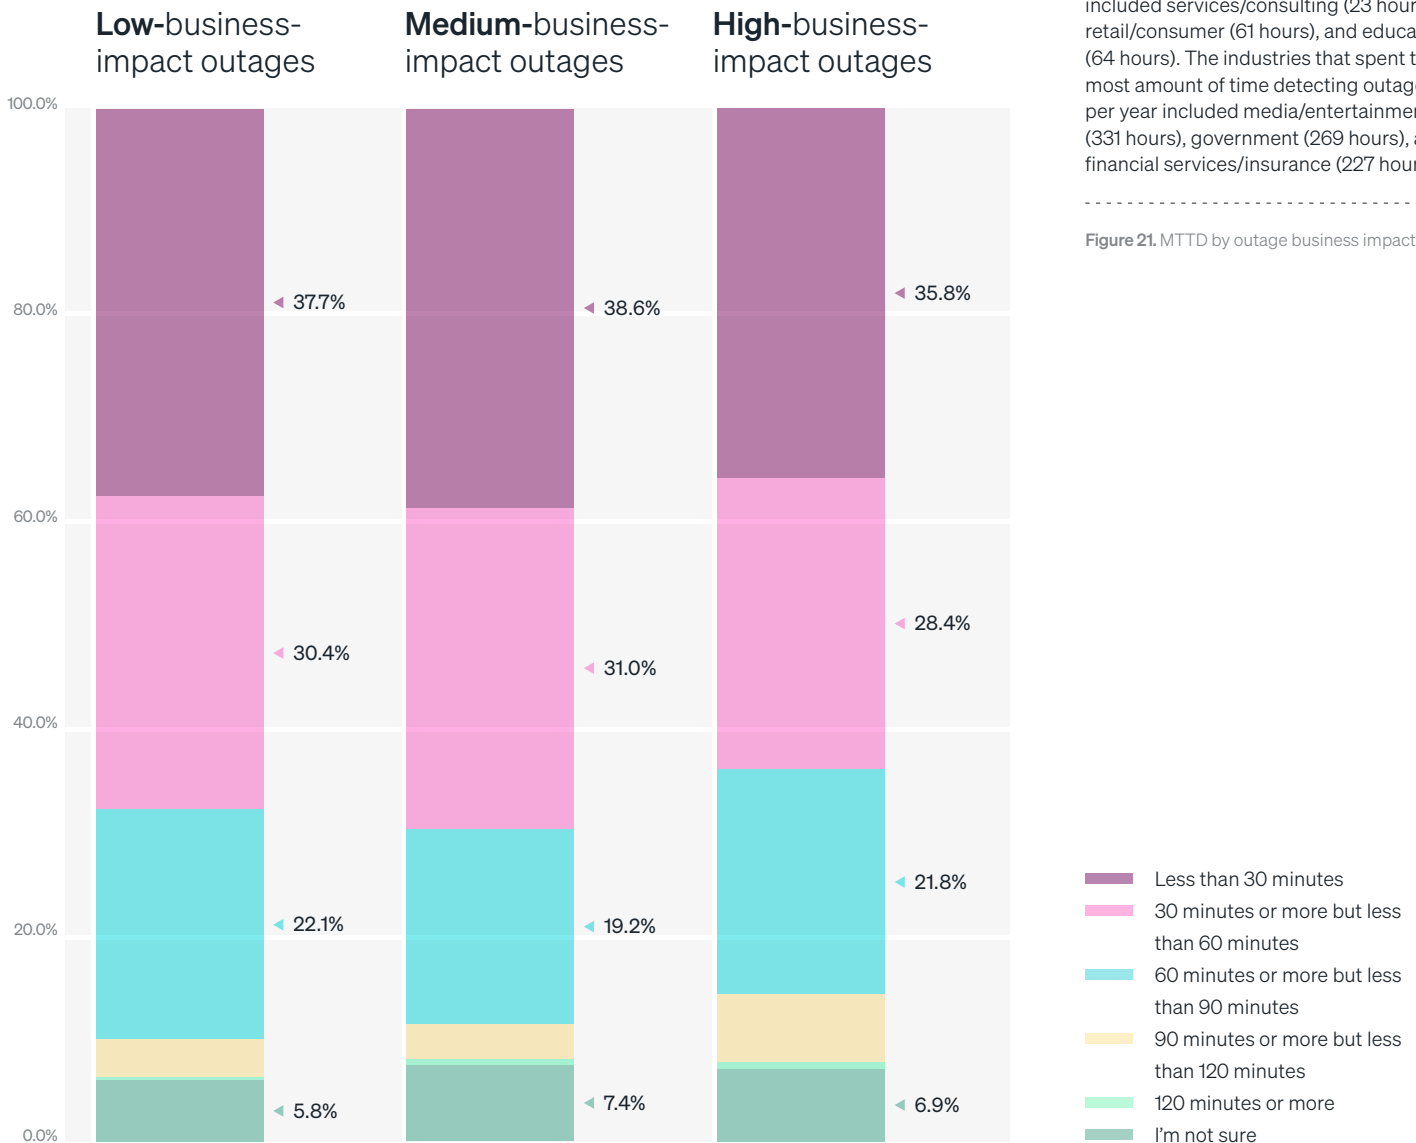


Figure 21. MTTD by outage business impact level

29%

took at least an hour to detect high-business-impact outages

Seven factors were associated with faster MTTD, including:

- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability spent 85% fewer hours detecting outages per year than those who hadn't (23 hours compared to 155 hours).
- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less time they spent detecting outages per year. For example, those who had deployed five or more observability capabilities spent 52% less time detecting outages than those who had deployed four or fewer (95 hours compared to 195 hours). Those who had deployed 10 or more spent 77% less time detecting outages per year than those who had deployed nine or fewer (39 hours compared to 170 hours). And those who had deployed 15 or more spent 84% less time detecting outages per year than those who had deployed 14 or fewer (22 hours compared to 138 hours).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data spent 79% less time detecting outages per year than those who had more siloed telemetry data (28 hours compared to 225 hours).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability spent 78% less time detecting outages per year than those who used more manual detection methods (48 hours compared to 216 hours).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data spent 65% less time detecting outages per year than those who integrated one to four types (57 hours compared to 162 hours).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices spent 35% less time detecting outages per year compared to those who had employed four or fewer (90 hours compared to 138 hours).
- ✓ **Using a single tool for observability:** Those using a single tool for observability spent 15% less time detecting outages per year than those using multiple tools (117 hours compared to 138 hours).



It takes about the same amount of time to detect outages for all business impact levels, but full-stack observability and other factors can reduce MTTD time drastically.

Mean time to resolution (MTTR)

There are similar patterns with MTTR, another common service-level metric used in security and IT incident management. The median number of hours spent on MTTR per year across all business impact levels was **141 hours—which is about six days**. The median MTTR for high-business-impact outages was 51 minutes, and more than a third (39%) of respondents said MTTR was an hour or more for high-business-impact outages.

Organization size insight

Midsized organizations had the lowest median MTTR (118 hours) compared to large (155 hours) and small (167 hours) organizations.

Regional insight

Respondents surveyed in Asia Pacific had the highest median MTTR (245 hours), followed by those surveyed in Europe (125 hours) and then Americas (53 hours).

Industry insight

The industries with the lowest median MTTR included services/consulting (48 hours), retail/consumer (75 hours), and education (97 hours). The industries with the highest median number MTTR included government (302 hours), media/entertainment (284 hours), and financial services/insurance (277 hours).




Figure 22. MTTR by outage business impact level

39%

took at least an hour to resolve high-business-impact outages

Seven factors were associated with faster MTTR, including:

- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability spent 76% fewer hours resolving outages per year than those who hadn't (41 compared to 168).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data spent 76% less time resolving outages per year than those who had more siloed telemetry data (62 hours compared to 258 hours).
- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less time they spent resolving outages per year. For example, those who had deployed five or more observability capabilities spent 41% less time resolving outages than those who had deployed four or fewer (113 hours compared to 191 hours). Those who had deployed 10 or more spent 70% less time detecting outages per year than those who had deployed nine or fewer (53 hours compared to 179 hours). And those who had deployed 15 or more spent 75% less time detecting outages per year than those who had deployed 14 or fewer (38 hours compared to 150 hours).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability spent 74% less time resolving outages per year than those who used more manual detection methods (63 hours compared to 240 hours).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data spent 57% less time resolving outages per year than those who integrated one to four types (77 hours compared to 178 hours).
- ✓ **Using a single tool for observability:** Those using a single tool for observability spent 20% less time resolving outages per year than those using multiple tools (124 hours compared to 155 hours).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices spent 10% less time resolving outages per year compared to those who had employed four or fewer (130 hours compared to 145 hours).



High-business-impact outages take the longest to resolve. However, full-stack observability and other factors can lead to resolving outages of all business impact levels much faster.

Total downtime

Given the relative frequency of outages and time to detect and resolve them as noted above, this adds up to considerable downtime for organizations. The data show that the median annual downtime across all business impact levels was **77 hours—which is about 3 days**.

Several factors were associated with less annual downtime, including:

- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less downtime they experienced per year. For example, those who had deployed five or more observability capabilities spent 45% less time resolving outages than those who had deployed four or fewer (223 hours compared to 409 hours). Those who had deployed 10 or more spent 74% less time detecting outages per year than those who had deployed nine or fewer (95 hours compared to 371 hours). And those who had deployed 15 or more spent 80% less time detecting outages per year than those who had deployed 14 or fewer (60 hours compared to 299 hours).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability experienced 79% downtime per year than those who hadn't (70 hours compared to 338 hours).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data experienced 78% less downtime per year than those who had more siloed telemetry data (107 hours compared to 488 hours).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability experienced 73% less downtime per year than those who used more manual detection methods (118 hours compared to 445 hours).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data experienced 63% less downtime per year than those who integrated one to four types (139 hours compared to 370 hours).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices experienced 19% less downtime per year compared to those who had employed four or fewer (239 hours compared to 294 hours).
- ✓ **Using a single tool for observability:** Those using a single tool for observability experienced 18% less downtime per year than those using multiple tools (249 hours compared to 305 hours).



Organization size insight

Small organizations had the highest median annual downtime (372 hours, which is about 16 days), followed by large (300 hours, which is about 13 days) and then midsize (230 hours, which is about 10 days) organizations.



Regional insight

Respondents surveyed in Asia Pacific had the highest median annual downtime (467 hours, which is about 19 days), followed by those surveyed in Europe (227 hours, which is about nine days) and the Americas (97 hours, which is about 4 days).



Industry insight

The industries with the highest median annual downtime included media/entertainment (608 hours, which is about 25 days), government (564 hours, which is about 24 days), and financial services/insurance (528 hours, which is about 22 days). The industries with the lowest median annual downtime included services/consulting (80 hours, which is about three days), education (158 hours, which is about a week), and retail/consumer (164 hours, which is about a week).


There's a strong association between less downtime and several factors, including achieving full-stack observability and deploying more capabilities in general.

“Downtime is expensive. You can spend more time from a human capital perspective to go in and solve a specific problem if you don't have the right tool in place.”

Senior Director of IT Infrastructure, Large Fintech Enterprise, United States

Outage cost

For low-business-impact outages, the median outage cost per hour of downtime was \$1.3 million. For medium-business-impact outages, the median outage cost per hour of downtime was \$1.6 million, and for high-business-impact outages, the median outage cost per hour of downtime was **\$1.9 million**.



62% said high-business-impact outages cost at least **\$1 million per hour of downtime**

Six factors were associated with a lower median outage cost for high-business-impact outages, including:

- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less they spent on outage costs per hour. For example, those who had deployed five or more observability capabilities spent 5% less on outages per hour than those who had deployed four or fewer (\$1.9 million compared to \$2.0 million). Those who had deployed 10 or more spent 41% less on outages per hour than those who had deployed nine or fewer (\$1.3 million compared to \$2.2 million). And those who had deployed 15 or more spent 50% less on outages per hour than those who had deployed 14 or fewer (\$1.0 million compared to \$2.0 million).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability spent 48% less on outages per hour than those who hadn't (\$1.1 million compared to \$2.1 million).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability spent 19% less on outages per hour than those who used more manual detection methods (\$1.7 million compared to \$2.1 million).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data spent 32% less on hourly outage costs than those who integrated one to four types (\$1.5 million compared to \$2.2 million).
- ✓ **Using a single tool for observability:** Those using a single tool for observability spent 45% less on outages per hour than those using multiple tools (\$1.1 million compared to \$2.0 million).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices spent 35% less on outages per hour compared to those who had employed four or fewer (\$1.3 million compared to \$2.0 million).

Organization size insight

Large organizations had higher median hourly outage costs for high-business-impact outages (\$2.1 million) than midsize (\$2.0 million) or small (\$1.3 million).

Regional insight

Respondents surveyed in Asia Pacific had the highest median hourly outage costs for high-business-impact outages (\$2.3 million) compared to those in Europe (\$1.7 million) and the Americas (\$1.4 million). Industry insight.

Industry insight

The industries with the highest hourly outage costs for high-business-impact outages included government (\$2.3 million), media/entertainment (\$2.2 million), telco (\$2.2 million), and financial services/insurance (\$2.2 million). The industries with the lowest median annual outage costs included services/consulting (\$1.3 million), education (\$1.3 million), and healthcare/pharma (\$1.3 million).




“On average, I can say that a minute of downtime can cost \$10,000 or more. Every single minute of downtime can be lost revenue for the company. If you’re down for one hour, it can cost millions.”

Senior Director of IT Infrastructure, Large Fintech Enterprise, United States

Detection of interruptions

While respondents were still more likely to say they learn about interruptions with observability (54%) than without observability (45%), this is 26% less than last year. And 12% more said they learn about interruptions with one observability platform compared to last year (17% compared to 15% in 2023).



Regional insight

Those surveyed in the Americas were the most likely to learn about interruptions with observability (63% compared to 55% for those in Europe and 46% for those in Asia Pacific). Conversely, those surveyed in Asia Pacific were the most likely to learn about them without observability (54% compared to 44% for those in Europe and 36% for those in the Americas).

Industry insight

Services/consulting respondents were the most likely to learn about interruptions with observability (74%), followed by healthcare/pharma (60%) and IT (58%). Media/entertainment respondents were the most likely to say they learn about them without observability (57%), followed by energy/utilities (56%) and telco (53%).

Figure 24. How respondents learned about software and system interruptions in 2022, 2023, and 2024

2024 respondents
 2023 respondents
 2022 respondents


45% still learn about interruptions through less efficient methods

Compared to those who learned about interruptions without observability, those who learned about them with observability:

- Experienced 73% less annual downtime (118 hours compared to 445 hours).
- Spent 19% less on hourly outage costs (\$1.7 million compared to \$2.1 million).
- Spent 28% less engineering time addressing disruptions (10 hours per week compared to 16 hours per week based on a 40-hour work week).

Time spent addressing disruptions

The median percentage of engineering team time spent addressing disruptions was 30%, which works out to 12 hours per week based on a 40-hour work week. Nearly half (45%) of respondents said their engineering team spends less than 30% of their time addressing disruptions, or less than 12 hours per week based on a 40-hour work week.



Organization size insight
Midsize and large organizations spend more time addressing disruptions (32% and 31% respectively) than small organizations (25%).

Regional insight
Respondents surveyed in Asia Pacific estimated the most time spent addressing disruptions (41%), followed by those in Europe (30%) and then the Americas (20%).

Industry insight
The industries with the highest time spent addressing disruptions included media/entertainment (49%), government (43%), and financial services/insurance (40%). The industries with the lowest time spent addressing disruptions included education (20%), services/consulting (20%), and healthcare/pharma (24%).

Figure 25. The percentage of engineering team time spent addressing disruptions was correlated with annual downtime (correlation value of 0.516).


- 0% to 9%
- 10% to 19%
- 20% to 29%
- 30% to 39%
- 40% to 49%
- 50% to 59%
- 60% to 69%
- 70% to 79%
- 80% to 89%
- 90% to 99%
- I'm not sure

Seven factors were associated with a lower percentage of engineering team time spent addressing disruptions, including:

- ✓ **Using a single tool for observability:** Those using a single tool for observability spent 50% less engineering time addressing disruptions than those using multiple tools (17% compared to 33%, or seven hours compared to 13 hours based on a 40-hour work week).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability spent 44% less engineering time addressing disruptions than those who hadn't (20% compared to 36%, or eight hours compared to 14 hours based on a 40-hour work week).
- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the less they tended to spend on engineering time per year. For example, those who had deployed five or more observability capabilities spent 24% less engineering time addressing disruptions than those who had deployed four or fewer (29% compared to 38%, or 12 hours compared to 15 hours based on a 40-hour work week). Those who had deployed 10 or more spent 41% less engineering time addressing disruptions than those who had deployed nine or fewer (22% compared to 38%, or nine hours compared to 15 hours based on a 40-hour work week). And those who had deployed 15 or more spent 39% less engineering time addressing disruptions than those who had deployed 14 or fewer (20% compared to 33%, or eight hours compared to 13 hours based on a 40-hour work week).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices spent 38% less engineering time addressing disruptions compared to those who had employed four or fewer (21% compared to 34%, or eight hours compared to 14 hours based on a 40-hour work week).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability spent 38% less on outages per year than those who used more manual detection methods (25% compared to 40%, or 10 hours compared to 16 hours based on a 40-hour work week).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data spent 27% less engineering time addressing disruptions than those who integrated one to four types (27% compared to 37%, or 11 hours compared to 15 hours based on a 40-hour work week).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data spent 11% less engineering time addressing disruptions than those who had more siloed telemetry data (28% compared to 32%, or 11 hours compared to 13 hours based on a 40-hour work week).

29%

said their engineering team spends at least half of their time addressing disruptions




Engineers spend a considerable amount of time addressing disruptions. But full-stack observability and other factors can help free up engineers' time so they can focus on higher-value work.

MTTx change

We also wanted to know how respondents thought their organization's MTTx (MTTD and MTTR) for outages had changed since adopting an observability solution.

MTTD change

For MTTD, data show more than half (56%) of respondents indicated some degree of improvement in MTTD since adopting an observability solution, including 29% who said it improved by 25% or more. About one in five (19%) said it remained the same.



Regional insight

Respondents surveyed in the Americas were much more likely to say their MTTD had improved to some extent since adopting observability (69% compared to 48% for those in both Asia Pacific and Europe).

Industry insight

Services/consulting respondents were the most likely to say their MTTD has improved to some extent since adopting observability (65%), followed by retail/consumer (63%), healthcare/pharma (63%), and media/entertainment (60%) respondents.

Figure 26. MTTD change since adopting observability

- Worsened to some extent
- Remained the same
- Improved to some extent
- I'm not sure

Most respondents detected outages faster after adopting observability. And several factors lead to an even faster MTTD, including achieving full-stack observability and employing more best practices.

56% said their their **MTTD improved** to some extent since adopting observability


Six factors were associated with improved MTTD, including:

- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the more likely they were to say their MTTD improved to some extent. For example, those who had deployed five or more observability capabilities were 62% more likely to say it improved than those who had deployed four or fewer (61% compared to 38%). Those who had deployed 10 or more were 44% more likely to say it improved than those who had deployed nine or fewer (69% compared to 48%). And those who had deployed 15 or more were 34% more likely to say it improved than those who had deployed 14 or fewer (72% compared to 54%).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability were 37% more likely to say their MTTD improved to some extent than those who hadn't (70% compared to 51%).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices were 37% more likely to say their MTTD improved to some extent than those who had employed four or fewer (72% compared to 53%).
- ✓ **Learning about interruptions with observability:** Those who learned about interruptions with observability were 35% more likely to say their MTTD improved to some extent than those who used more manual detection methods (63% compared to 47%).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data were 33% more likely to say their MTTD improved to some extent than those who integrated one to four types (66% compared to 50%).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data were 15% more likely to say their MTTD improved to some extent than those who had more siloed telemetry data (63% compared to 55%).



MTTR change

For MTTR, data shows that the majority (59%) of respondents indicated some degree of improvement in MTTR since adopting an observability solution, with less than a quarter (22%) saying it remained the same.



Organization size insight

Small (64%) and large organizations were more likely to experience improved MTTR (64% and 61% respectively) than midsize organizations (55%).

Regional insight

Respondents surveyed in the Americas were much more likely to say their MTTR improved to some extent since adopting observability (67%) compared to those surveyed in Asia Pacific (59%) and Europe (47%).

Industry insight

Media/entertainment respondents were the most likely to say their MTTR has improved to some extent since adopting observability (73%), followed by education (71%), healthcare/pharma (66%), services/consulting (65%), and financial services/insurance (62%).

Figure 27. MTTR change since adopting observability

- Worsened to some extent
- Remained the same
- Improved to some extent
- I'm not sure

Most respondents resolved outages faster after adopting observability. And several factors lead to an even faster MTTR, including achieving using a single tool for observability and employing more best practices.

59% said their their **MTTR improved** to some extent since adopting observability

Seven factors were associated with improved MTTR, including:

- ✓ **Deploying more observability capabilities:** The more capabilities they deployed, the more likely they were to say their MTTR improved to some extent. For example, those who had deployed five or more observability capabilities were 30% more likely to say it improved than those who had deployed four or fewer (63% compared to 48%). Those who had deployed 10 or more were 27% more likely to say it improved than those who had deployed nine or fewer (69% compared to 54%). And those who had deployed 15 or more were 23% more likely to say it improved than those who had deployed 14 or fewer (71% compared to 58%).
- ✓ **Employing more observability best practices:** Those who had employed five or more observability best practices were 36% more likely to say their MTTR improved to some extent than those who had employed four or fewer (77% compared to 56%).
- ✓ **Achieving full-stack observability:** Those who had achieved full-stack observability were 23% more likely to say their MTTR improved to some extent than those who hadn't (69% compared to 56%).
- ✓ **Integrating more types of business-related data with telemetry data:** Those who had integrated five or more types of business-related data with their telemetry data were 20% more likely to say their MTTR improved to some extent than those who integrated one to four types (67% compared to 57%).
- ✓ **Having more unified telemetry data:** Those who had more unified telemetry data were 14% more likely to say their MTTR improved to some extent than those who had more siloed telemetry data (65% compared to 57%).
- ✓ **Learning about interruptions with observability:** Those who learn about interruptions with observability were 13% more likely to say their MTTR improved to some extent than those who used more manual detection methods (63% compared to 55%).
- ✓ **Using a single tool for observability:** Those using a single tool for observability were 11% more likely to say their MTTR improved to some extent than those using multiple tools (65% compared to 59%).



Influencers of lower MTTx by capability

The data show there is a positive association between a lower than average MTTD and MTTR and 11 observability capabilities:

- **Business observability and error tracking are statistically significant within 5% significance levels.**
- Alerts and dashboards have had a positive association for three years in a row (2022–2024).
- Error tracking and log management have had a positive association for two years in a row (2023 and 2024).
- APM, database monitoring, and security monitoring have had a positive association twice (2022 and 2024).
- AI monitoring (new this year), browser monitoring, business observability (new this year), and network monitoring had a positive association for the first time this year.




Figure 28. Observability capabilities associated with a lower than average MTTD and MTTR

Organizations looking to reduce MTTx may improve their odds by prioritizing the deployment of strategic observability capabilities, especially business observability and error tracking.

Downtime reduction

At least a third of respondents said conducting root cause analysis (RCA) and post-incident reviews (37%), monitoring DORA (DevOps Research and Assessment) metrics (34%), monitoring the golden signals (33%), and tracking, reporting, and incentivizing MTTx (33%) have helped their organization reduce downtime.

About a quarter said implementing service-level management (28%), providing organization-wide access to observability data (26%), using dashboards to report detailed performance and health KPIs (22%), and configuring automated alerts for critical incidents (22%) have helped their organization reduce downtime.

37%

said conducting root cause analysis and post-incident reviews helped reduce downtime




Figure 29. Observability practices that have helped organizations reduce downtime

Developers and engineers are proactively taking steps to improve uptime and reliability.



Organization size insight

Respondents from large organizations were notably more likely to say conducting RCA and post-incident reviews, monitoring DORA metrics, and monitoring the golden signals than those from midsize and small organizations.

Regional insight

Respondents surveyed in Asia Pacific were much more likely to say monitoring DORA metrics helped reduce downtime. Respondents surveyed in the Americas were much more likely to say monitoring the golden signals and using a centralized log management system.

Industry insight

Nearly half said monitoring DORA metrics helped reduce downtime for the following industries: media/entertainment (48%), telco (47%), government (44%), and financial services/insurance (42%). At least a third said monitoring the golden signals helped reduce downtime for the following industries: education (39%), financial services/insurance (39%), retail/consumer (37%), media/entertainment (35%), telco (35%), energy/utilities (35%), and industrial/materials/manufacturing (33%).

Observability benefits

This section covers the primary benefits of observability, the advantages of achieving full-stack observability, how much total value organizations receive from their observability investment per year, and the median return on investment (ROI) for observability.

Highlights:



58%

Said they receive \$5 million+ in total value per year from their observability investment



34%

said they receive \$10 million or more in annual value from their observability investment




28%

said they received a higher annual value from their observability investment when they deployed AIOps capability



Primary observability benefits

Respondents saw clear benefits as a result of their current observability solution—nearly half (46%) cited improved system uptime and reliability. More than a third said increased operational efficiency (42%), reduced security risks (39%), and an improved real-user (customer) experience (36%). Nearly a third (32%) said they experience improved developer productivity, cost optimization (30%), business and/or revenue growth (29%), and the ability to handle traffic surges (28%).



Organization size insight

Respondents from large organizations were generally more likely to cite each observability benefit, followed by those from midsize and then small organizations.

Regional insight

Respondents surveyed in the Americas were generally the most likely to say they experienced benefits. Those surveyed in Asia Pacific were the most likely to cite the ability to handle traffic surges, business and/or revenue growth, and accelerated rate of innovation or competitive advantage. Those surveyed in Europe were the most likely to cite cost optimization and regulation compliance.

Industry insight

Improved system uptime and reliability was the top choice for most industries. However, increased operational efficiency was the top choice for media/entertainment (46%), IT (45%), and energy/utilities (39%). Reduced security risks was the top choice for education (54%).

Figure 30. Benefits enabled by observability solution


Organizations that had achieved full-stack observability experienced more benefits than those that hadn't:

- 51% more likely to improve system uptime and reliability (62% compared to 41%)
- 44% more likely to increase operational efficiency (55% compared to 38%)
- 30% more likely to optimize costs (36% compared to 28%)
- 26% more likely to reduce security risks (47% compared to 37%)
- 15% more likely to improve the real-user (customer) experience (40% compared to 35%)

Total value of observability

The median annual value an organization received from its observability investment was **\$8.15 million**. More than half (58%) said they receive \$5 million or more in annual value, and over a third (34%) said they receive \$10 million or more in value from their observability investment.

Those who had deployed five or more observability capabilities received a higher annual value (\$8.20 million) from their observability investment than those who had deployed one to four (\$7.93 million).



In addition, a higher total annual value received was associated with deploying the following observability capabilities:

- 28% higher for those deploying artificial intelligence for IT operations (AIOps) capabilities (\$9.85 million compared to \$7.70 million for those who hadn't deployed AIOps capabilities)
- 17% higher for those deploying synthetic monitoring (\$9.15 million compared to \$7.83 million for those who hadn't deployed synthetic monitoring)
- 14% higher for those deploying AI monitoring (\$8.75 million compared to \$7.65 million for those who hadn't deployed AI monitoring)
- 14% higher for those deploying Kubernetes (K8s) monitoring (\$9.00 million compared to \$7.93 million for those who hadn't deployed K8s monitoring)
- 5% higher for those deploying machine learning (ML) model monitoring (\$8.43 million compared to \$8.05 million for those who hadn't deployed ML model monitoring)
- 3% higher for those deploying mobile monitoring (\$8.35 million compared to \$8.13 for those who hadn't deployed mobile monitoring)

Organization size insight

Respondents from large organizations reported the highest annual value received (\$9.15 million), followed by those from midsize (\$8.15 million) and small (\$2.65 million) organizations.

Regional insight

Respondents surveyed in Asia Pacific reported a much higher median annual value received (\$10.08 million) compared to those surveyed in Europe (\$7.05 million) and the Americas (\$5.40 million).

Industry insight

The industries that saw the highest median annual value from observability were financial services/insurance (\$10.15 million), government (\$10.08 million), and media/entertainment (\$10.00 million).

Figure 31. Benefits enabled by observability solution



Organizations are experiencing substantial value from their observability investments. This total value includes all benefits, such as downtime avoidance, tool optimization, employee productivity, and so on.

Return on investment for observability

The median ROI calculation was based on annual observability spend and annual value received estimates.

The median ROI for observability across all respondents was **4x (295%)**. In other words, for every \$1 spent, respondents believe they receive \$4 of value.

“How much I want to spend versus the benefits that I’ll receive is always a hot topic. I acknowledge that I have a really complex environment to maintain, but I also need to factor in the ROI and how much I’m really willing to compromise, considering that we have a B2C type of environment. That means we’re customer-facing, the retail side. Uptime resiliency is something that is really important for me. I cannot compromise on downtime or taking more time for MTTR.”

Senior Director of IT Infrastructure, Large Fintech Enterprise, United States

Regional insight

Respondents surveyed in the Americas had a slightly lower median ROI (3.8x) than those surveyed in Asia Pacific or Europe (both 4x).

Industry insight

Government respondents had the highest median ROI (4.1x), followed by telco, retail/consumer, IT, and financial services/insurance (all 4x). Education respondents had the lowest median ROI (3.7x), followed by energy/utilities and healthcare/pharma (both 3.8x).

Future of observability

This section explores observability deployment and data integration plans for next year and the next two to three years, and what steps organizations are most likely to take in the next year to **get the most value out of their observability spend.**

Highlights:



83%

expected to deploy 6+ new observability capabilities by next year



59%

planned to integrate 5+ business-related data types with their telemetry data in the next 1-3 years



41%

planned to consolidate observability tools in the next year



Observability deployment plans

In addition to asking survey-takers about their current observability deployment, we inquired about their deployment plans for next year and the next two to three years.

Looking out to 2025, most (91%) expected to deploy at least one new capability in the next year, more than half (59%) expected to deploy one to five, and more than a third (37%) expected to deploy seven or more capabilities.

When we look at the summary for one year out, capability deployment expectations are at least 80% for capabilities like security monitoring and network monitoring. At least a third expected to deploy artificial intelligence for IT operations (AIOps) capabilities (39%), AI monitoring (36%), machine learning (ML) model monitoring (34%), distributed tracing (33%), and serverless monitoring (33%).

Organization size insight

Small organizations (64%) are more likely to have 5+ observability capabilities deployed next year than midsize (60%) or large organizations (58%).

Regional insight

A much higher proportion of organizations in Asia Pacific (74%) plan to deploy 5+ observability capabilities by next year than those in Europe (45%) or Americas (49%).

Industry insight

More than three-quarters (77%) of respondents in the media/entertainment industry plan to deploy 5+ observability capabilities by next year, more than any other industry.




Figure 32. Number of observability capabilities planned for next year

- 0 new observability capabilities
- 1-5 new observability capabilities
- 6-10 new observability capabilities
- 11-19 new observability capabilities

By mid-2027, 75% or more expected to deploy each of the 19 different observability capabilities. Very few of our survey respondents did not expect to deploy these observability capabilities (up to 13%).




Figure 33. Observability capabilities deployment summary for 2024 through 2027

- We currently deploy
- We do not currently have deployed, but have plans to add in the next year
- We do not currently have deployed, but have plans to add in the next 2-3 years
- We do not currently have deployed, and have no plans to add
- I'm not sure

Data integration plans

To practice true business observability, organizations must integrate their business-related data with their telemetry data (MELT). We already reviewed what types of data they currently integrate. Now let’s review the types of business-related data they said they plan to integrate in the next one to three years:

- About half planned to integrate each data type in the next one to three years.
- Most (89%) planned to integrate at least one data type, including 59% who planned to integrate five or more. Only 11% had no plans to integrate business-related data with their telemetry data.
- These results mean that by 2027, 89% will have integrated five or more business-related data types, including 57% who will have integrated all 10.

Regional insight

Respondents surveyed in Asia Pacific were the most likely to say they plan to integrate 5+ types of business-related with their telemetry data in the next 1–3 years (72% compared to 50% for the Americas and 46% for Europe).

Industry insight

Media/entertainment respondents were the most likely to say they plan to integrate 5+ types of business-related with their telemetry data in the next 1–3 years (75%), followed by energy/utilities (72%) and telco (66%). IT respondents were the least likely (45%), followed by healthcare/ pharma (49%) and education (52%).




Figure 34. Types of business-related data currently integrated or planned to be integrated with telemetry data through 2027

- We currently integrate with our telemetry data
- We do not currently integrate with our telemetry data, but have plans to integrate in the next year
- We do not currently integrate with our telemetry data, but have plans to integrate in the next 2–3 years
- We do not currently integrate with our telemetry data, and have no plans to integrate
- I'm not sure

59% planned to **integrate 5+** business-related data types with their telemetry data **in the next 1–3 years**




Figure 35. Telemetry data integration plans for the next 1–3 years


- 0 business-related data types
- 1-4 business-related data types
- 5+ business-related data types

Organizations plan to diversify the data types that they integrate in their observability practice in the coming years.

Value maximization plans for observability

We were interested to know what steps organizations are most likely to take in the next year to get the most value out of their observability spend. The survey results showed that:

- Nearly half (47%) planned to train staff on how to best use the observability tools they have.
- Roughly two in five (41%) planned to consolidate tools.
- More than a third planned to optimize their engineering team size (37%).
- Nearly a third planned to reduce spending across the board (31%).
- The rest planned to switch to a more affordable vendor (28%), use open source (24%), or observe less of their tech stack (23%).
- Just 1% said they won't take steps to get the most value out of their observability spend.



Regional insight

Respondents surveyed in Asia Pacific were the most likely to say they plan to integrate 5+ types of business-related with their telemetry data in the next 1–3 years (72% compared to 50% for the Americas and 46% for Europe).

Industry insight

Media/entertainment respondents were the most likely to say they plan to integrate 5+ types of business-related with their telemetry data in the next 1–3 years (75%), followed by energy/utilities (72%) and telco (66%). IT respondents were the least likely (45%), followed by healthcare/pharma (49%) and education (52%).

Figure 36. Types of business-related data currently integrated or planned to be integrated with telemetry data through 2027

■ 2024 respondents
■ 2023 respondents

“The majority of the time, my desire is to have a single tool. This way people would be reporting things in a uniform way across the board.”

Senior Manager of Engineering, Large Retail/Consumer Enterprise, India

41%
planned to **consolidate**
observability tools in the next year




Figure 37. Telemetry data integration plans for the next 1-3 years

- 0 business-related data types
- 1-4 business-related data types
- 5+ business-related data types

Organizations plan to diversify the data types that they integrate in their observability practice in the coming years.

Summary

The **benefits of observability are becoming clear**, with nearly half (46%) of survey respondents citing improved system uptime and reliability as a top advantage. Other key benefits include increased operational efficiency (42%), reduced security risks (39%), and enhanced customer experience (36%). Organizations also reported gains in developer productivity, cost optimization, business growth, and the ability to handle traffic surges. Notably, companies that achieved full-stack observability saw even greater benefits, including a 51% higher likelihood of improving system uptime and a 44% higher likelihood of increasing operational efficiency.

In terms of financial impact, the median annual value organizations received from their observability investments was \$8.15 million, with more than half of respondents reporting annual value of \$5 million or more. Those that deployed five or more observability capabilities realized even higher returns, averaging \$8.20 million annually. Specific observability capabilities, such as artificial intelligence for IT operations (AIOps) and synthetic monitoring, further amplified the total value received, with organizations deploying these technologies **seeing up to 28% higher annual value** than those that had not.

Companies with full-stack observability, unified telemetry data, and a single observability platform experienced notably fewer outages, with some reporting up to **77% fewer annual outages** than those with more fragmented or siloed observability practices.



About this report

New Relic partnered with Enterprise Technology Research (ETR) to conduct a survey and analysis for this fourth annual **Observability Forecast** report.



What's new

We conducted the survey in Brazil and updated the regional distribution. We repeated many of last year's questions to compare trends year-over-year (YoY) and added seven new ones to gain additional insights.

For industries, we broke out media/entertainment, separated IT from telco, and removed unspecified.

We added two additional observability capabilities—artificial intelligence (AI) monitoring and business observability—bringing the total number of observability capabilities covered in this year's report to 19. We also provided survey takers with brief definitions for each observability capability to improve clarity, which was deemed especially necessary for emerging or contested areas. Providing these definitions inherently constrained survey-takers' interpretations of each of these capabilities.

We included sections in the survey for the average revenue cost per hour of downtime for each outage business impact level (low, medium, and high).

We updated the answer format for all currency-, time-, and quantity-related questions to sliders instead of ranges. The sliders enabled survey takers to select a more precise amount, which enabled us to tabulate the median amount across all respondents for these questions. For example, we've tabulated the median annual observability spend, median annual downtime, median annual outage cost, median total annual value from their observability investment, median engineering time spent addressing disruptions, and median return on investment (ROI).

In addition to mean time to resolution (MTTR), we asked survey takers about how their organization's mean time to detection (MTTD) for outages changed since adopting an observability solution.



Definitions

We've defined common terms and concepts used throughout this report.

Observability

To avoid bias, we did not define observability for survey takers.

Observability enables organizations to measure how a system performs and identify issues and errors based on its external outputs. These external outputs are called telemetry data and include metrics, events, logs, and traces (MELT). Observability is the practice of instrumenting systems to surface insights for various roles so they can take immediate action that impacts customer experience and service. It also involves collecting, analyzing, altering, and correlating that data for improved uptime and performance.

Achieving observability brings a connected, real-time view of all data from different sources—ideally in one place—where teams can collaborate to troubleshoot and resolve issues faster, prevent issues from occurring, ensure operational efficiency, and produce high-quality software that promotes an optimal customer/user experience and competitive advantage.

Software engineering, development, site reliability engineering, operations, and other teams—plus managers and executives—use observability to understand the behavior of complex digital systems and turn data into tailored insights. Observability helps them pinpoint issues more quickly, understand root causes for faster, simpler incident response, and proactively align data with business outcomes.

A subset of observability, organizations use monitoring to identify problems in the environment based on prior experience that's expressed as a set of conditions (known unknowns). Monitoring enables organizations to react to these conditions and is sufficient to solve problems when the number and complexity of possible problems are limited.

Organizations use observability to determine why something unexpected happened (in addition to the what, when, and how), particularly in complex environments where the possible scope of problems and interactions between systems and services is significant. The key difference is that observability does not rely on prior experience to define the conditions used to solve all problems (unknown unknowns). Organizations also use observability proactively to optimize and improve environments.

Many tools are purpose-built for observability and include capabilities such as:

Analysis and incident management

- AIOps (artificial intelligence for IT operations) capabilities
- Alerts
- Error tracking

Apps and services

- Application performance monitoring (APM)
- Distributed tracing
- Serverless monitoring

Artificial intelligence (AI)

- AI monitoring
- Machine learning (ML) model monitoring

Business impact and visibility

- Business observability
- Dashboards

Digital experience monitoring (DEM)

- Browser monitoring
- Mobile monitoring
- Synthetic monitoring

Infrastructure

- Database monitoring
- Infrastructure monitoring
- Kubernetes (K8s) monitoring
- Network performance monitoring

Log management

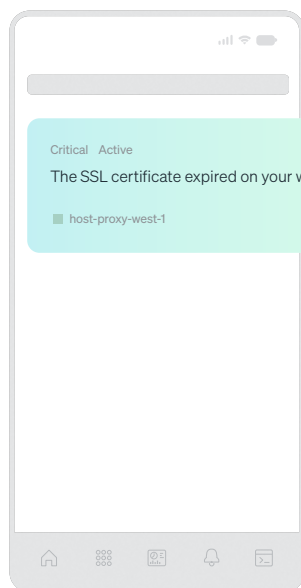
Security monitoring

Monitoring tools alone can lead to data silos and data sampling. In contrast, an observability platform can instrument an entire technology stack and correlate the telemetry data drawn from it in a single location for one unified, actionable view.

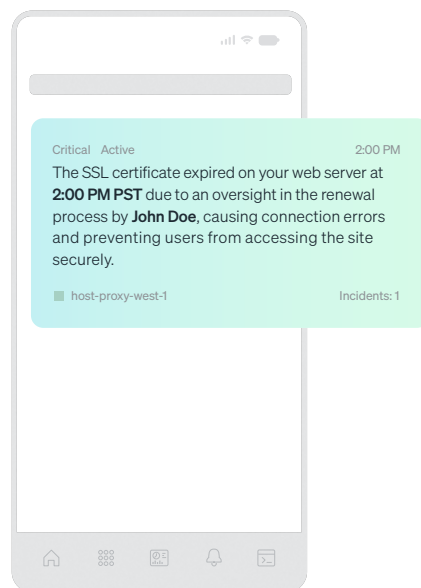
Monitoring	Observability
Reactive	Proactive
Situational	Predictive
Speculative	Data-driven
What + when	What + when + why + how
Expected problems (known unknowns)	Unexpected problems (unknown unknowns)
Data silos	Data in one place
Data sampling	Instrument everything

Business observability takes that a step further by offering comprehensive visibility and quantifying business impact.

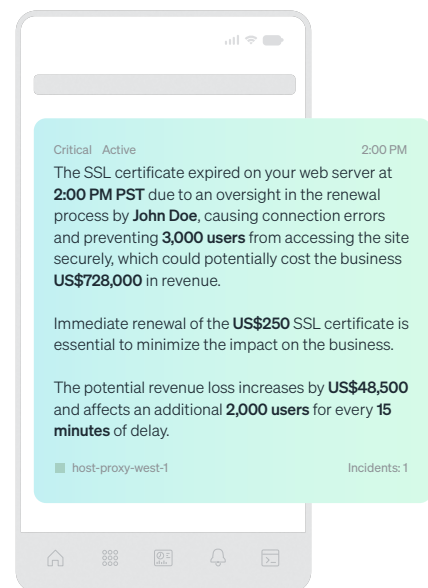
★ ★ ★
Monitoring: Good
 IT teams



★ ★ ★
Observability: Better
 Operations teams



★ ★ ★
Business observability: Best
 Business managers



Full-stack observability

The ability to see everything in the tech stack that could affect the customer experience is called full-stack observability or end-to-end observability. It's based on a complete view of all telemetry data.

Adopting a data-driven approach for end-to-end observability helps empower engineers and developers with a complete view of all telemetry data so they don't have to sample data, compromise their visibility into the tech stack, or waste time stitching together siloed data. Instead, they can focus on the higher-priority, business-impacting, and creative coding they love. And it provides executives and managers with a comprehensive view of the business and enables them to understand the business impact of interruptions.

Full-stack observability, as used in this report, is achieved by organizations that deploy specific combinations of observability capabilities, including apps and services, log management, infrastructure (backend), DEM (frontend), and security monitoring.

See how many respondents had achieved full-stack observability and the advantages of achieving full-stack observability.

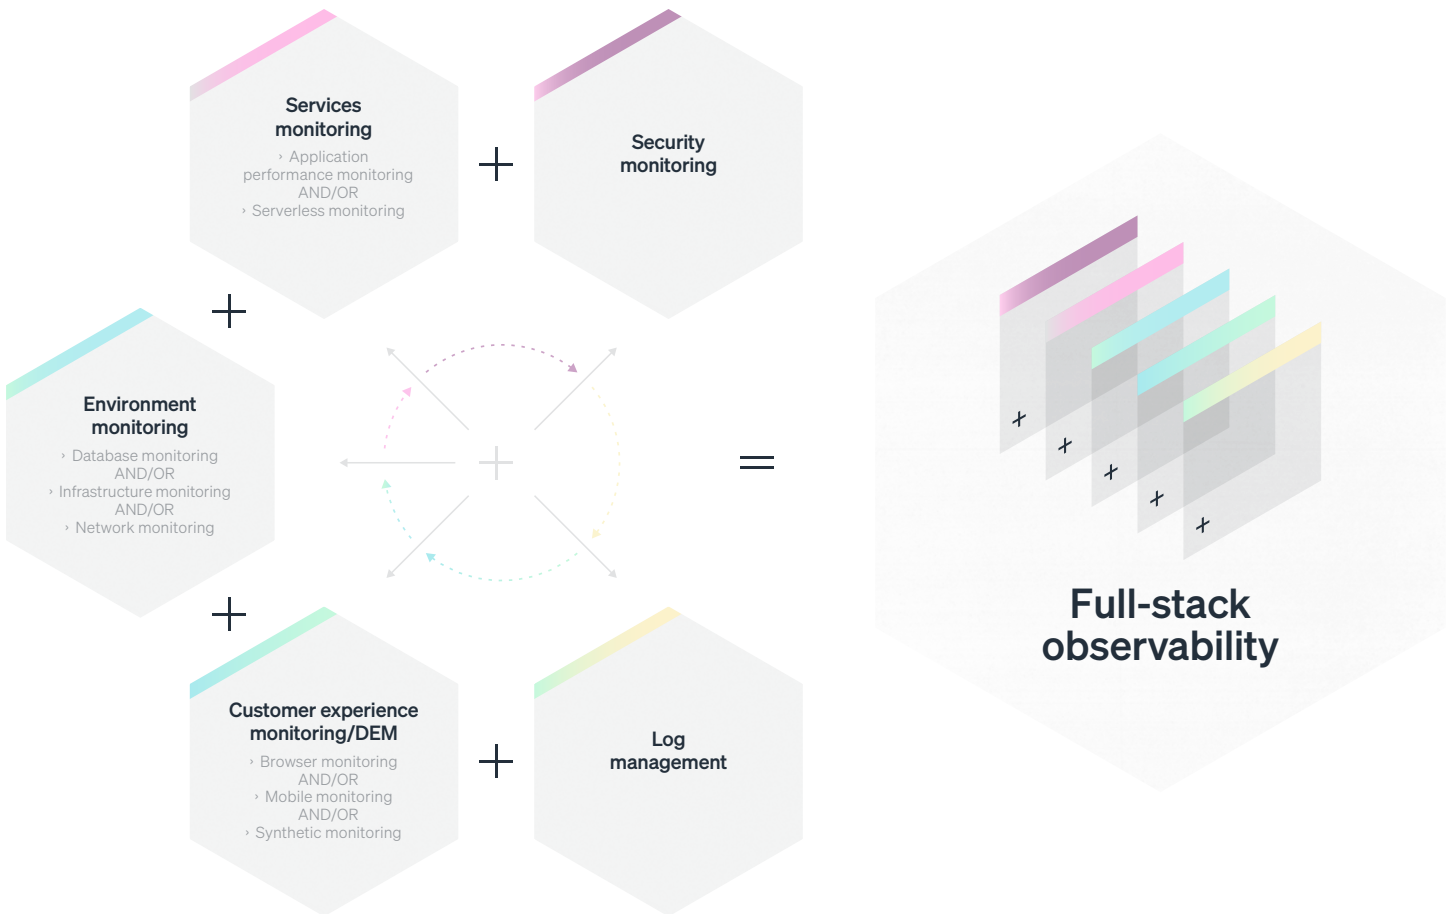


Figure 38. Full-stack observability definition

MTTx

MTTx includes mean time to detection or discovery (MTTD), mean time to identification (MTTI), mean time to acknowledge (MTTA), and mean time to resolution or repair (MTTR).

Organization size

In this report, organization size is determined by employee count unless otherwise noted.

- Small:** 1–100
- Midsized:** 101–1,200
- Large:** 1,201+

Roles

Study participants consisted of practitioners and IT decision makers (ITDMs). Practitioners are typically the day-to-day users of observability tools.

Roles, job titles, descriptions, and common key performance indicators (KPIs) for practitioners and ITDMs:

Roles	Job titles	Descriptions	Common KPIs
Developers	Application developers, software engineers, architects, and their frontline managers	Members of a technical team who design, build, and deploy code, optimizing and automating processes where possible	<ul style="list-style-type: none"> • Cycle time (speed of making changes) • Endpoint security incidents • Error rates • Lead time (speed from idea to deployment) • Mean time between incidents (MTBI) • Speed of software performance • Uptime percentage
Operations professionals	IT operations engineers, network operations engineers, DevOps engineers, DevSecOps engineers, SecOps engineers, site reliability engineers (SREs), infrastructure operations engineers, cloud operations engineers, platform engineers, system administrators, architects, and their frontline managers	<p>Members of a technical team who are responsible for the overall health and stability of infrastructure and applications</p> <p>Detect and resolve incidents using monitoring tools, build and improve code pipeline, and lead optimization and scaling efforts</p>	<ul style="list-style-type: none"> • Availability • Deploy speed and frequency • Error budgets • Error rates • Mean time to detection (MTTD) • Mean time to resolution (MTTR) • Service level agreements (SLAs) • Service level indicators (SLIs) • Service level objectives (SLOs) • Uptime percentage
Non-executive managers	Directors, senior directors, vice presidents (VPs), and senior vice presidents (SVPs) of engineering, operations, DevOps, DevSecOps, SecOps, site reliability, and analytic	<p>Leaders of practitioner teams that build, launch, and maintain customer-facing and internal products and platforms</p> <p>Own the projects that operationalize high-level business initiatives and translate technology strategy into tactical execution</p> <p>Constantly looking to increase velocity and scale service</p>	<ul style="list-style-type: none"> • Customer satisfaction • MTBI • MTTR • On-time project completion • Software development and efficiency • Speed of deployment • Uptime percentage
Executives (C-suite)	<p>More technical focused: Chief information officers (CIOs), chief information security officers (CISOs), chief technology officers (CTOs), chief data officers (CDOs), chief analytics officers (CAOs), and chief architects</p> <p>Less technical focused: Chief executive officers (CEOs), chief operating officers (COOs), chief financial officers (CFOs), chief marketing officers (CMOs), chief revenue officers (CROs), and chief product officers (CPOs)</p>	<p>Managers of overall technology infrastructure and cost who are responsible for business impact, technology strategy, organizational culture, company reputation, and cost management</p> <p>Define the organization's technology vision and roadmap to deliver on business objectives</p> <p>Use digital technology to improve customer experience and profitability, enhancing company reputation as a result</p>	<ul style="list-style-type: none"> • Conversion rates • Cost-effectiveness • Customer satisfaction • Return on investment (ROI) • Speed of deployment • Speed of innovation • Total cost of ownership (TCO) • Uptime percentage

Methodology

All data in this report are derived from a survey, which was in the field from April to May 2024.

ETR qualified survey respondents based on relevant expertise. ETR performed a non-probability sampling type called quota sampling to target sample sizes of respondents based on their country of residence and role type in their organizations (in other words, practitioners and ITDMs). Geographic representation quotas targeted 16 key countries.

To avoid skewing results by industry, subsamples of $n < 10$ are excluded from some analysis in this report.

All quotes were derived from interviews conducted by ETR with IT professionals who use observability.


All dollar amounts in this report are in USD.

Demographics

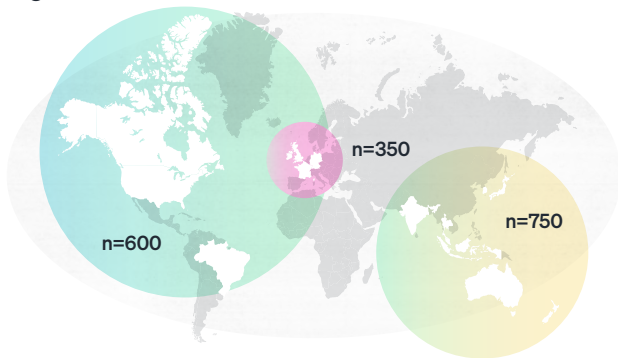
In 2024, ETR polled 1,700 technology professionals—more than any other observability study—in 16 countries across the Americas, Asia Pacific, and Europe. Approximately 35% of respondents were from Brazil (new this year), Canada, and the United States. France, Germany, Ireland, and the United Kingdom represented 21% of respondents. The remaining 44% were from the broader Asia-Pacific region, including Australia, India, Indonesia, Japan, Malaysia, New Zealand, Singapore, South Korea, and Thailand. View the regional highlights.

The survey respondent mix was about the same as in 2021, 2022, and 2023—65% practitioners and 35% ITDMs.

Countries



Regions



Roles




Figure 40. Respondent demographics, regions, countries, and roles

Firmographics


More than half of survey respondents (57%) worked for large organizations, followed by 34% for midsize organizations, and 9% for small organizations.

Less than a quarter (17%) cited \$500,000 to \$9.99 million, 26% cited \$10 million to \$99.99 million, and 57% cited \$100 million or more in annual revenue.

The respondent pool represented a wide range of industries, including IT, financial services/insurance, industrials/materials/manufacturing, retail/consumer, healthcare/pharmaceutical (pharma), energy/utilities, services/consulting, telecommunications (telco), education, government, media/entertainment (new this year), and nonprofit.


Organization size

- Large organizations (1,201+ employees)
- Midsize organizations (101-1,200 employees)
- Small organizations (1-100 employees)



Annual revenue

- \$500K-\$999.99K
- \$1M-\$4.99M
- \$5M-\$9.99M
- \$10M-\$24.99M
- \$25M-\$49.99M
- \$50M-\$99.99M
- \$100M-\$249.99M
- \$250M-\$499.99M
- \$500M-\$999.99M
- ≥\$1B



Industries




Figure 41. Respondent organization size, annual revenue, and industries

About ETR

ETR is a technology market research firm that leverages proprietary data from its targeted ITDM community to deliver actionable insights about spending intentions and industry trends. Since 2010, ETR has worked diligently at achieving one goal: eliminating the need for opinions in enterprise research, which are typically formed from incomplete, biased, and statistically insignificant data.

The ETR community of ITDMs is uniquely positioned to provide best-in-class customer/evaluator perspectives. Its proprietary data and insights from this community empower institutional investors, technology companies, and ITDMs to navigate the complex enterprise technology landscape amid an expanding marketplace.



About New Relic

The New Relic Intelligent Observability Platform helps businesses eliminate interruptions in digital experiences. New Relic is the only platform to unify and pair telemetry data to provide clarity over the entire digital estate. We move problem solving past proactive to predictive by processing the right data at the right time to maximize value and control costs. That's why businesses around the world—including Adidas Runtastic, American Red Cross, Domino's, GoTo Group, Ryanair, Topgolf, and William Hill—run on New Relic to drive innovation, improve reliability, and deliver exceptional customer experiences to fuel growth. Visit www.newrelic.com.



[Learn About New Relic Platform](#)

