

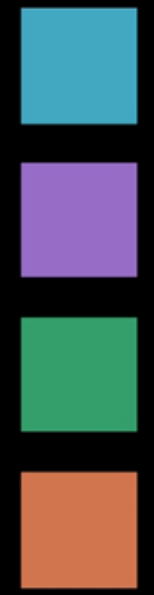

The State of Observability

2022 Report, Observe & CITE Research

Table of Contents

What Is the State of Observability?

Observability has been gaining in popularity and it can be a boon in addressing the challenges that modern businesses are facing. However, achieving observability it is not without its own challenges.

In 2021 Observe started the State of Observability report to better understand the struggles and benefits IT and software development professionals encounter as they work toward more Observable systems.

Now we're on the second edition of the State of Observability and we've partnered with CITE Research to talk to 250 professionals to understand their use of observability, how they're measuring its value and the what challenges they've encountered along the way.

Become your company's observability hero!

Summary of Findings

Adoption of emerging technologies such as microservices and serverless continues to grow, and these technologies – and the complexity they bring – are driving factors in observability. For example, Kubernetes users are more likely to consider observability a high priority.

Many organizations (90%) claim to practice observability, however confusion around the term observability and what it means persists. Many are still in earlier stages of maturity as only 35% cite widespread usage of observability. Additionally, few orgs (38%) say they currently have a way measure the impacts of observability.

Data is crucial to achieving observability but ingesting large volumes can present its own challenges. 78% of organizations are collecting more than 100GB of data a day. However, nearly half (46%) are discarding data they would otherwise keep because of cost concerns.

Tool sprawl itself might not be top of mind as a challenge, but it might be causing complexity and data silos. This can eat away at your budget without providing better MTTR.

Reducing MTTR is more important than ever, with 70% of organizations citing 6+ incidents occurring per month and 30% saying incidents can take a few days to investigate.

Part 1. Adoption Status

Site Reliability Engineering Sees Gradual Adoption Increases

DevOps continues to be commonplace as most organizations now practice it. Though rarer, Site Reliability Engineering (SRE) adoption is up 8% since 2021, at 61%.

SREs are often at the forefront of observability usage. Although SREs are still in short supply at many organizations we can expect their numbers to continue to grow as observability becomes more commonplace.

Does your organization currently practice each of the following?

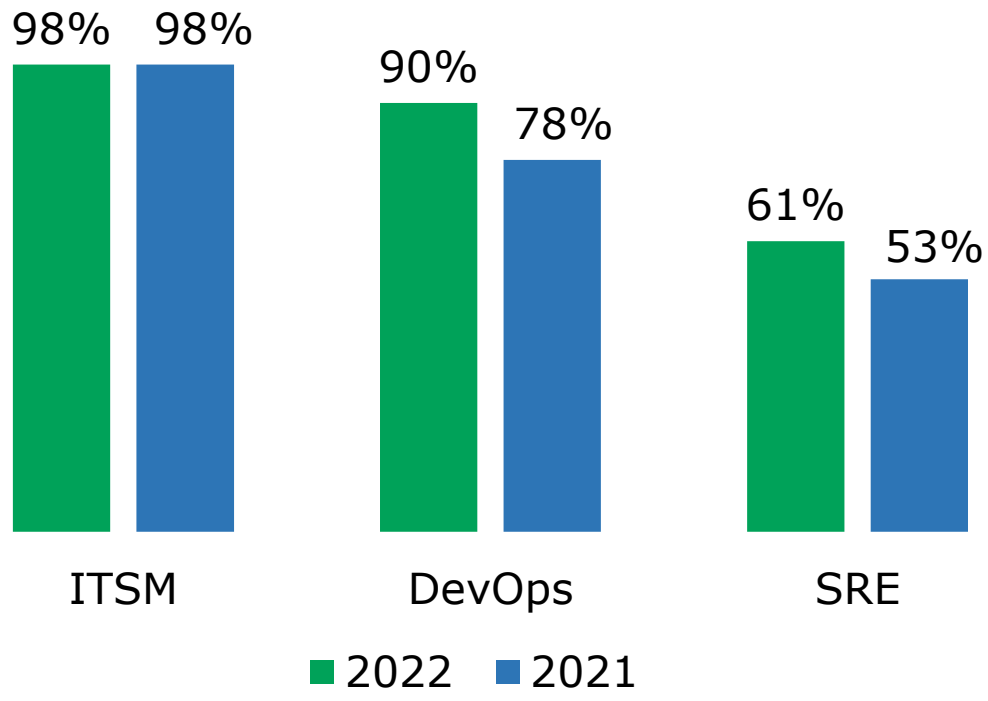

Adoption of Emerging Technologies Is on the Rise

Emerging technologies such as microservices and serverless have become mainstream with around ¾ of organizations citing adoption.

The high adoption of these technologies have made them essential use cases for observability.

The modest gains YoY may be an indicator that organizations lagging in adoption are those taking a more conservative approach to new tech adoption.

Is your organization running any of the following in production?

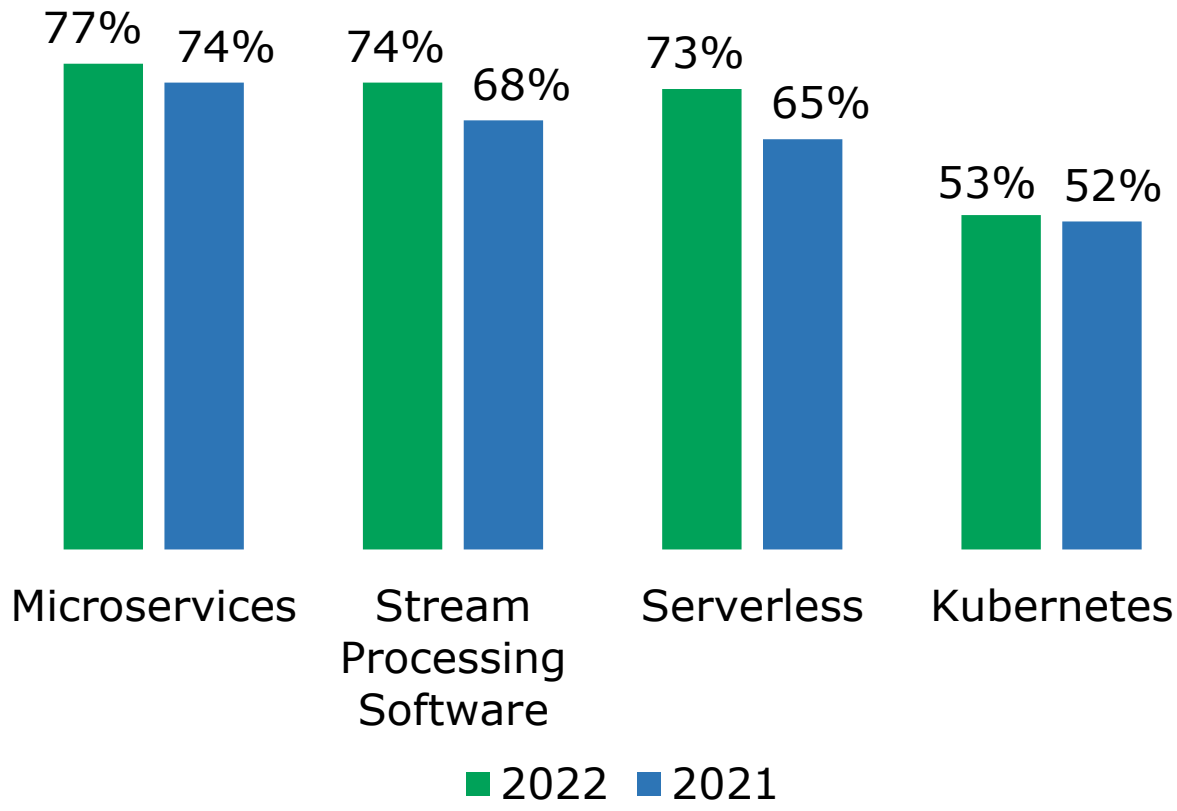

Part 2. Observability Usage

94% of Orgs Say They Are Familiar With Observability

Nearly all organizations claim familiarity with the term Observability, but fewer (67%) feel they are very familiar with it.

Since 90% say they are currently practicing it, we can infer that the gap compared to familiarity indicates some organizations are still grappling with what constitutes Observability and how it differs from past approaches to monitoring.

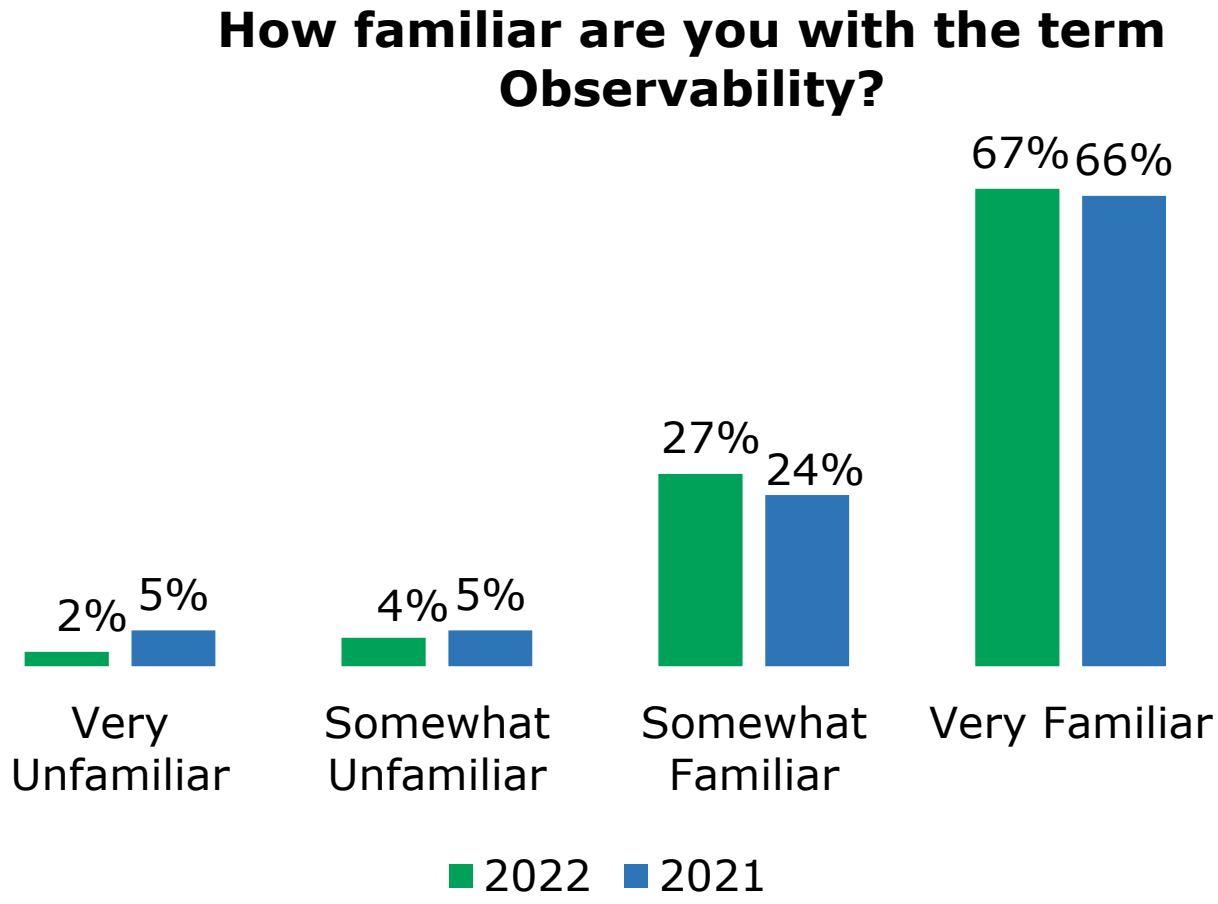

Adoption Appears High, but Maturity Is in the Earlier Stages for Many

Do any teams in your organization practice Observability?

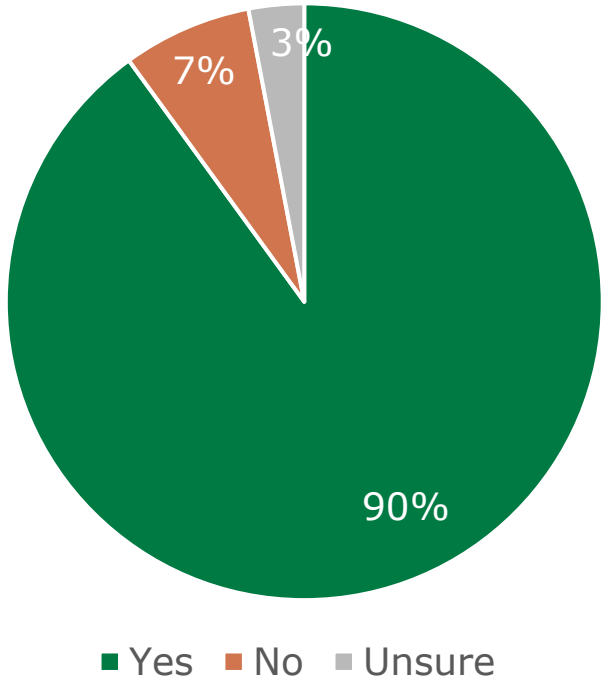

How would you rate your organization's maturity with Observability?

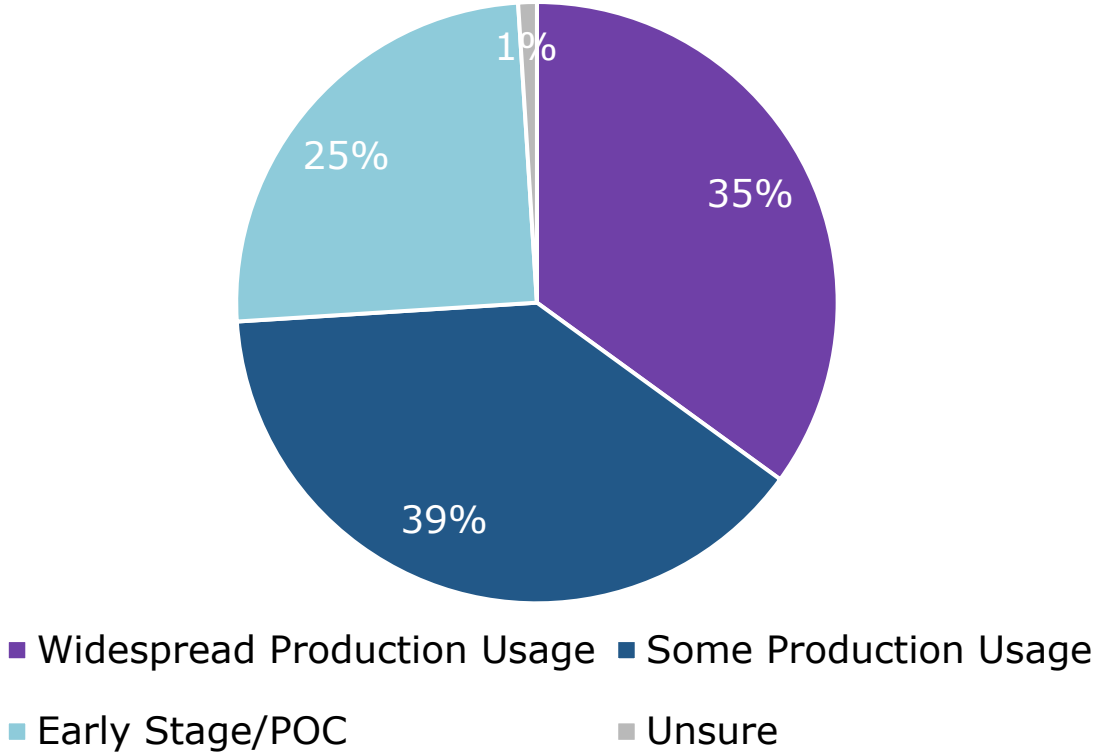

Organizational Buy In Is Less of a Challenge, but Complexity Remains the Primary Hurdle

Which of the following challenges has your organization faced when it comes to Observability?

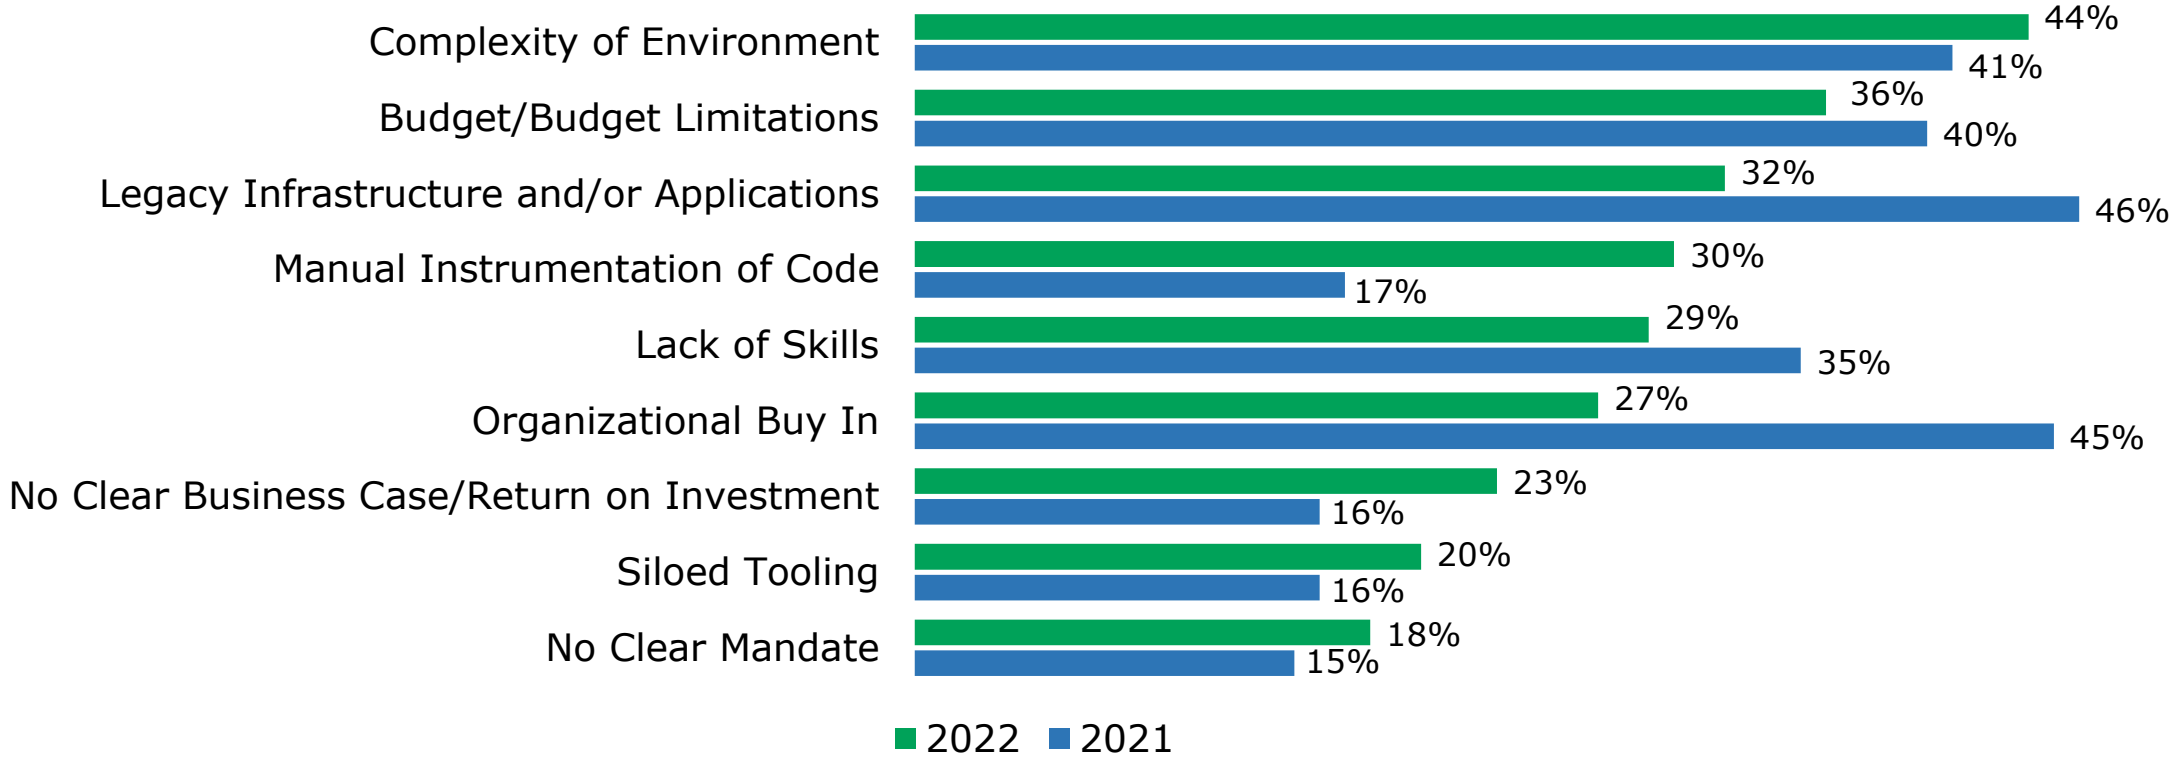

Observability Is a High Priority for Half of Organizations

Observability is a high priority for most organizations (53%). The advent of Microservices and Serverless have made complexity an even greater challenge driving the immediate need of observability.

Organizations using Kubernetes in production are even more likely to consider observability a high priority (64%).

How much of a priority is Observability for your organization?

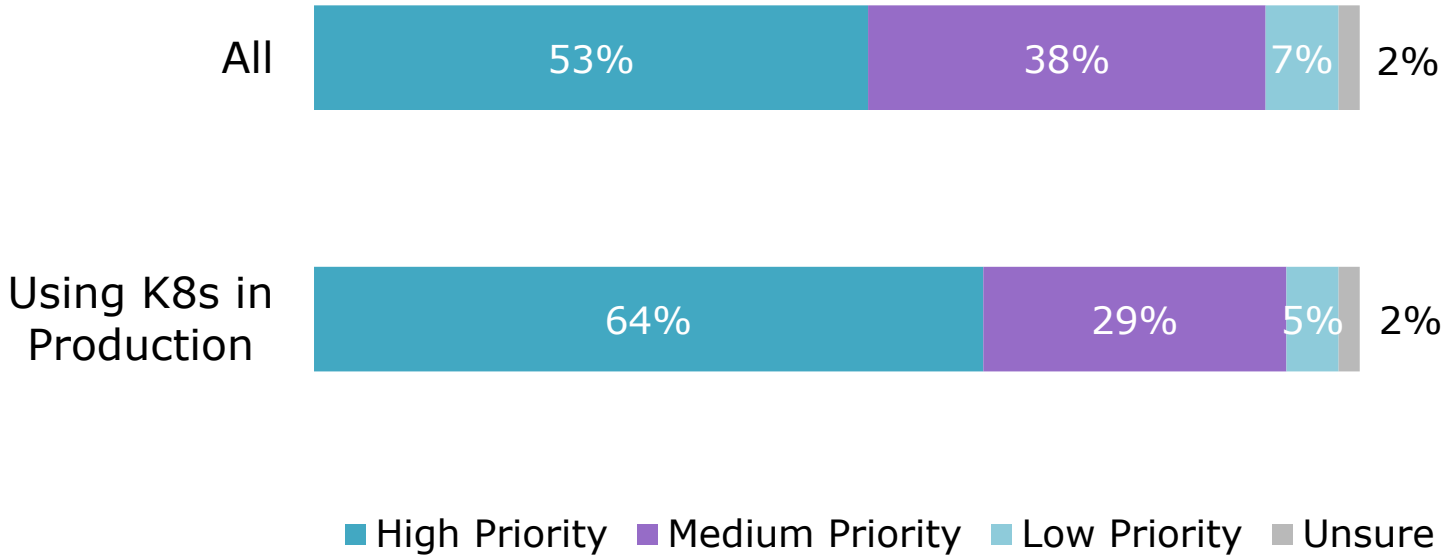

Part 3. Measuring Impact

Better Service Reliability Is the Top Perceived Benefit for Observability

What are the top three benefits you perceive Observability brings to your organization? Please rank the top three benefits, with 1 being the most beneficial.

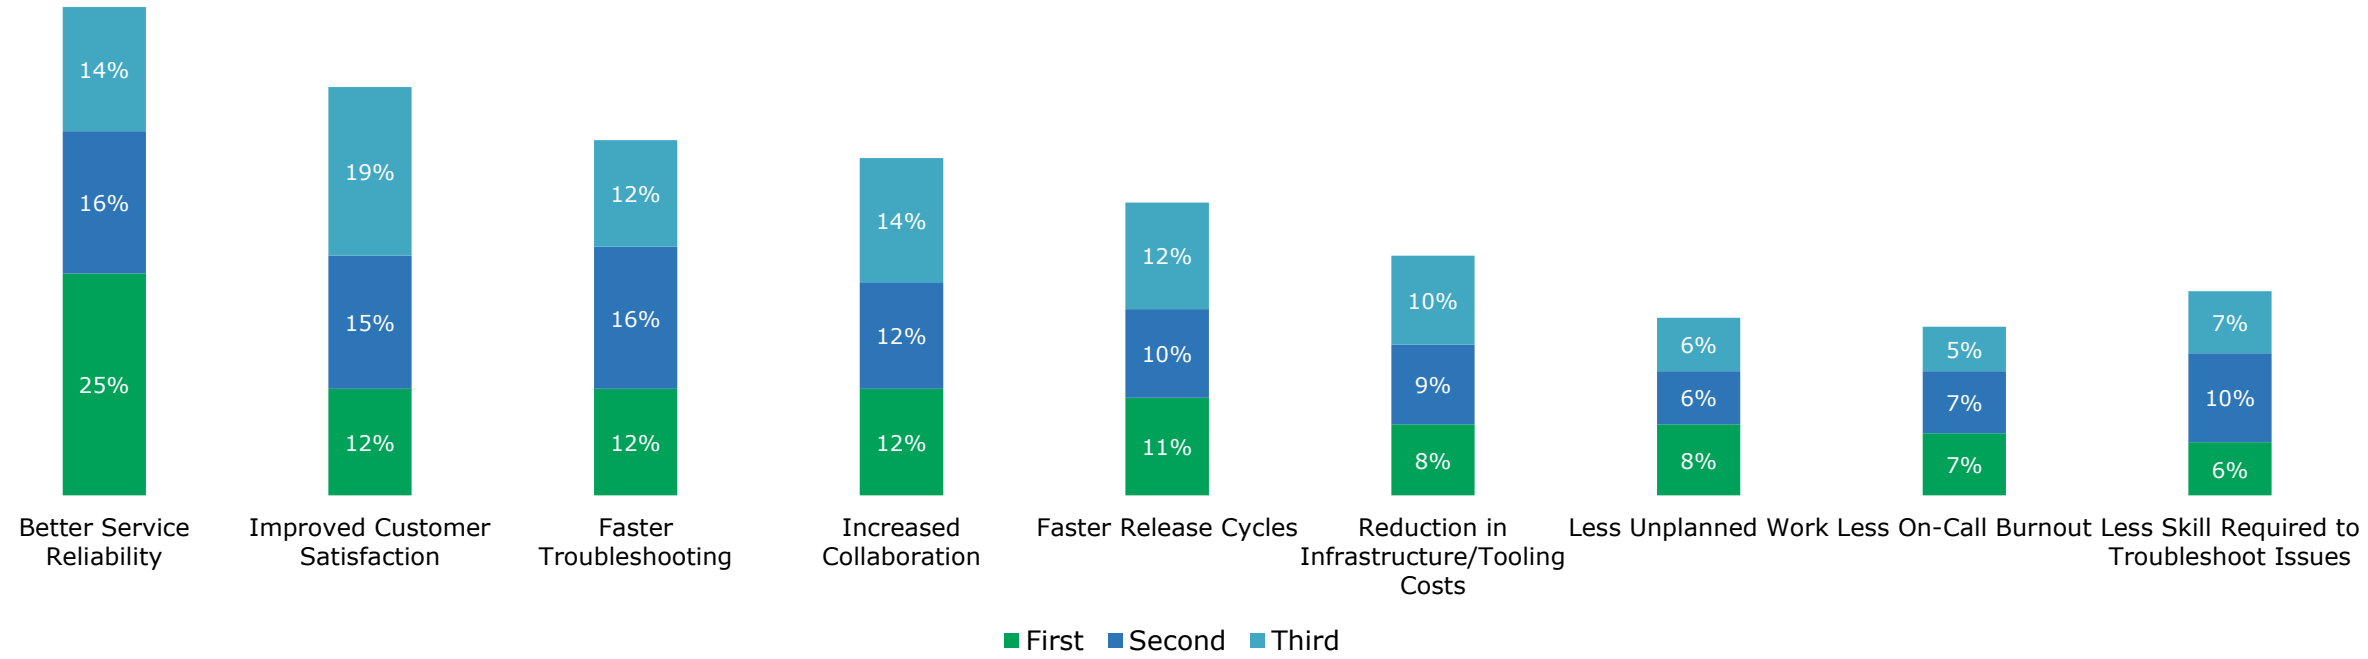

Orgs That Prioritize Observability Are More Likely To Measure Its Impact

While many organizations cite familiarity with Observability, many are not equipped to measure its impact.

Being able to measure results is crucial to ensuring that the perceived benefits of observability are lining up with the actual realized improvements.

Organizations that are prioritizing observability are more likely to have measuring its impact in their plans and may have more success with their o11y strategy as a result.

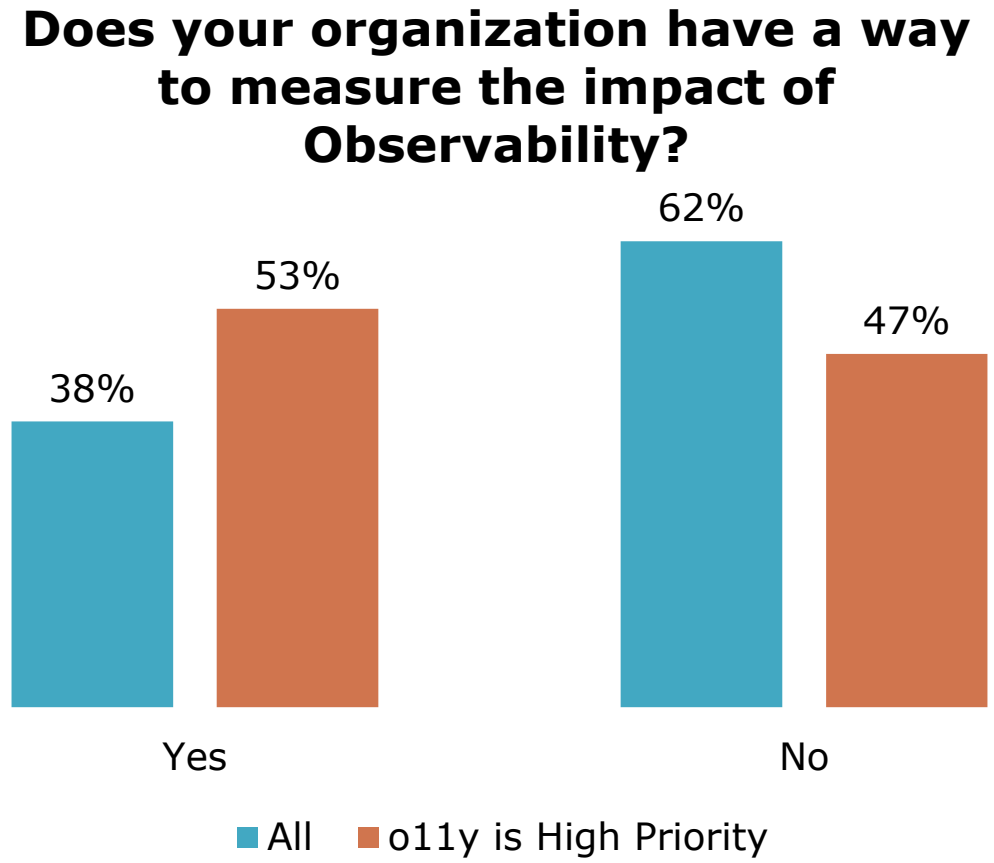

Customer Satisfaction and Service Availability Are Measurable Benefits

Since adopting Observability, has your organization had any measurable improvements in any of the following?

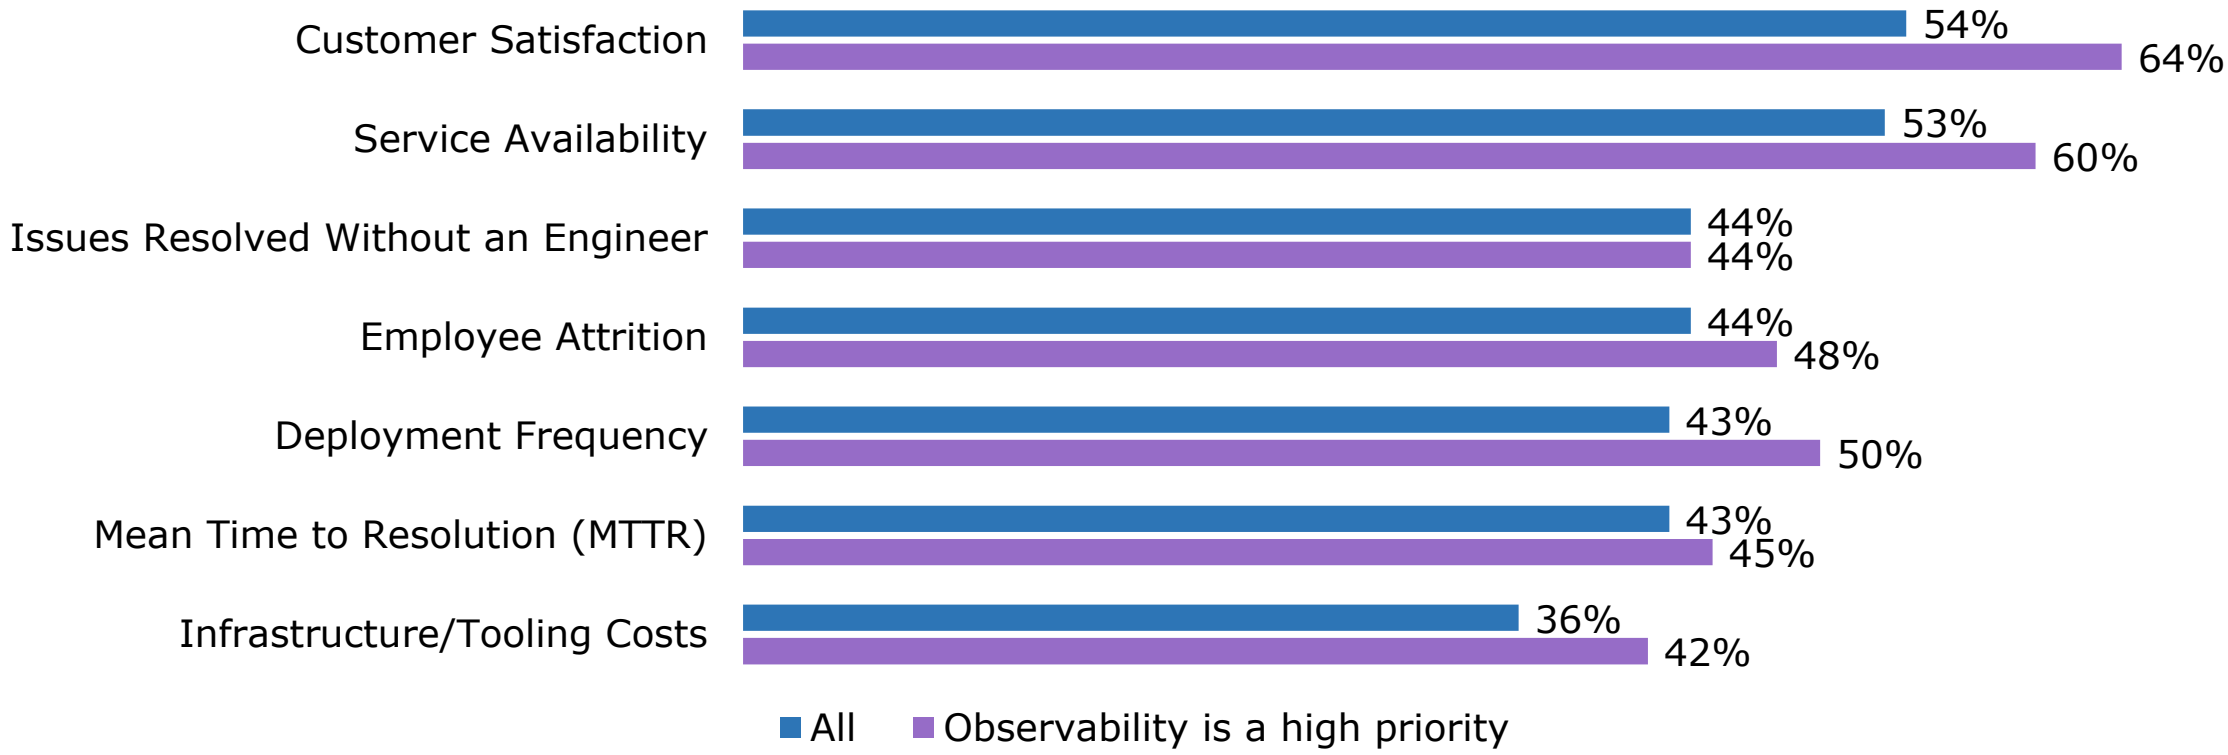

Part 4. The Data Challenge

78% of Orgs Ingest 100GB+ of Observability Data Per Day

Many organizations are already ingesting large amounts of data daily as 78% are collecting more than a 100GB.

Observability needs data to be effective but knowing which data to collect, and how much (if any) can be discarded is challenging.

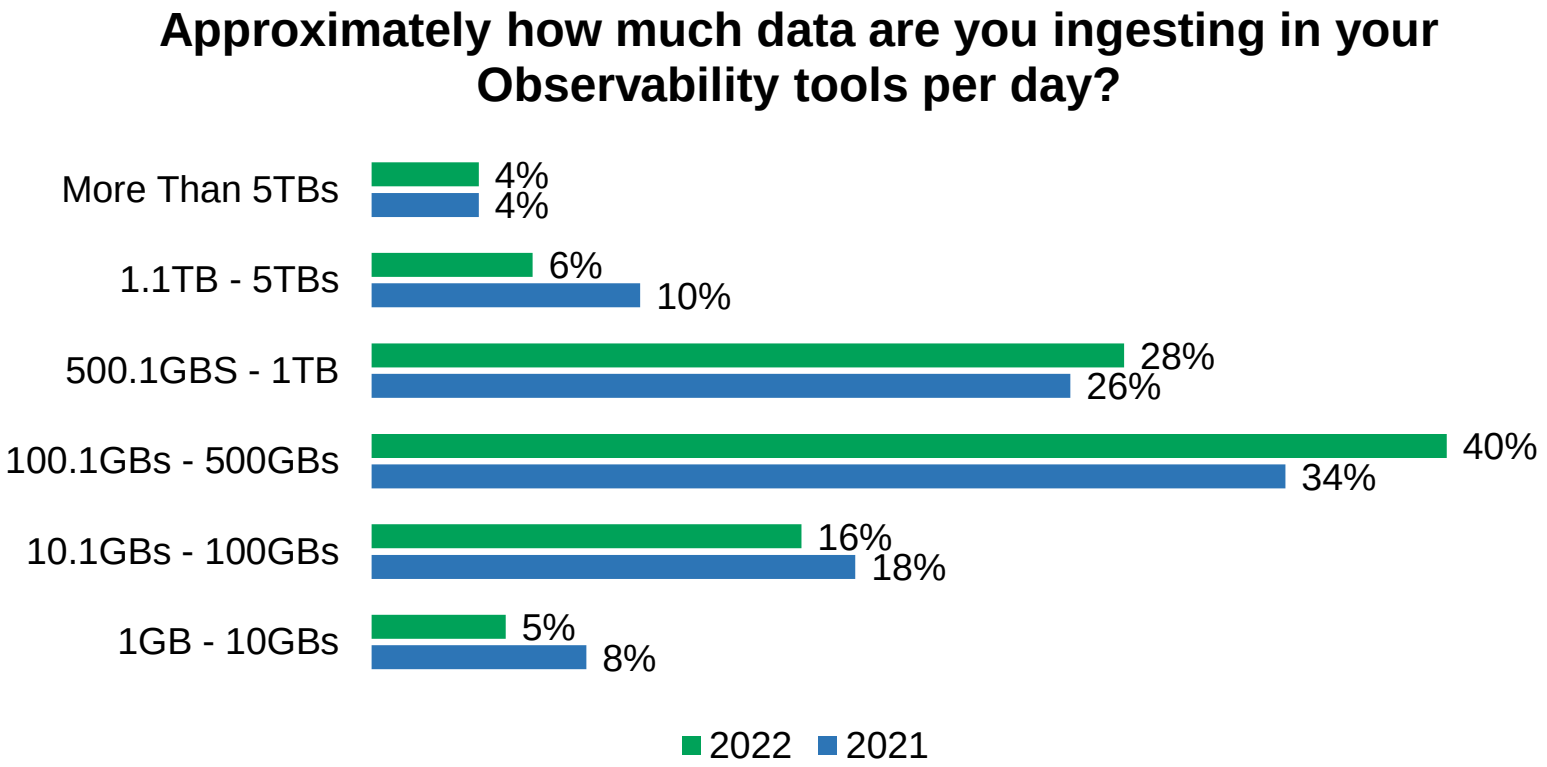

Data Is Coming From a Growing Number of Sources

Has your organization adopted any of the following data forwarders?

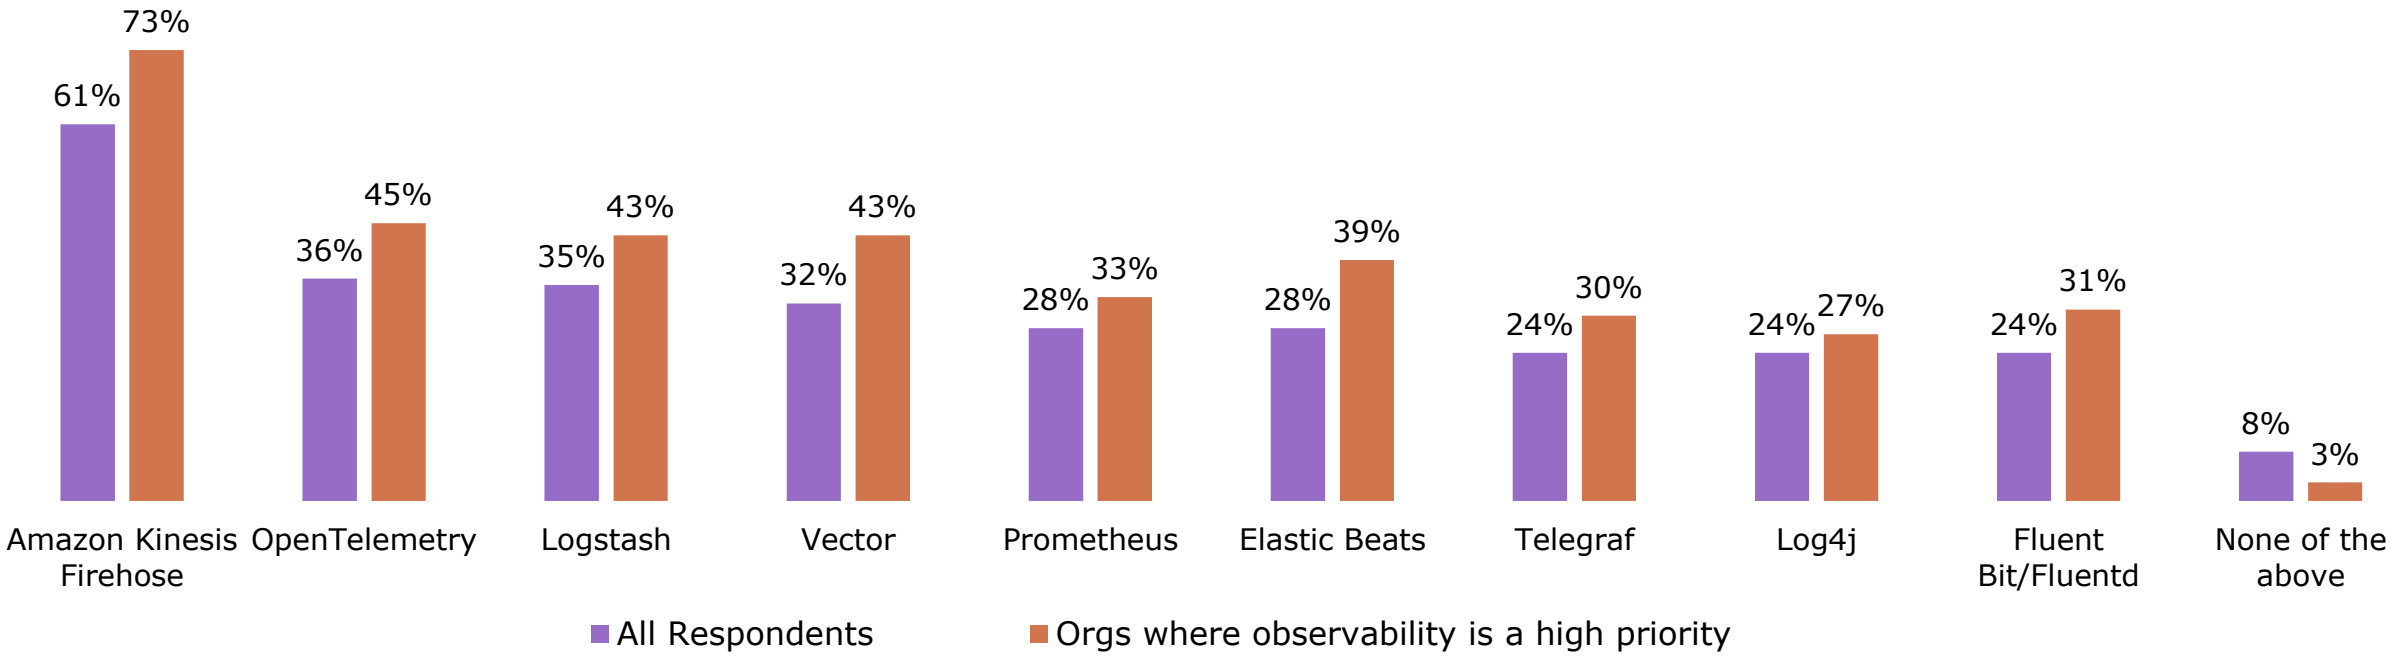

Few Users Emit the Majority of Logs With Structure

63% of organizations are emitting half or fewer of their logs with structure. Data will not always be structured in the ideal format when it comes time to query it.

It’s imperative that users have observability tooling that can help them derive usable insights from all manner of data, including whatever they are generating now whether its structured or not.

What percentage of your logs are emitted with a structure? Please give your best estimate.

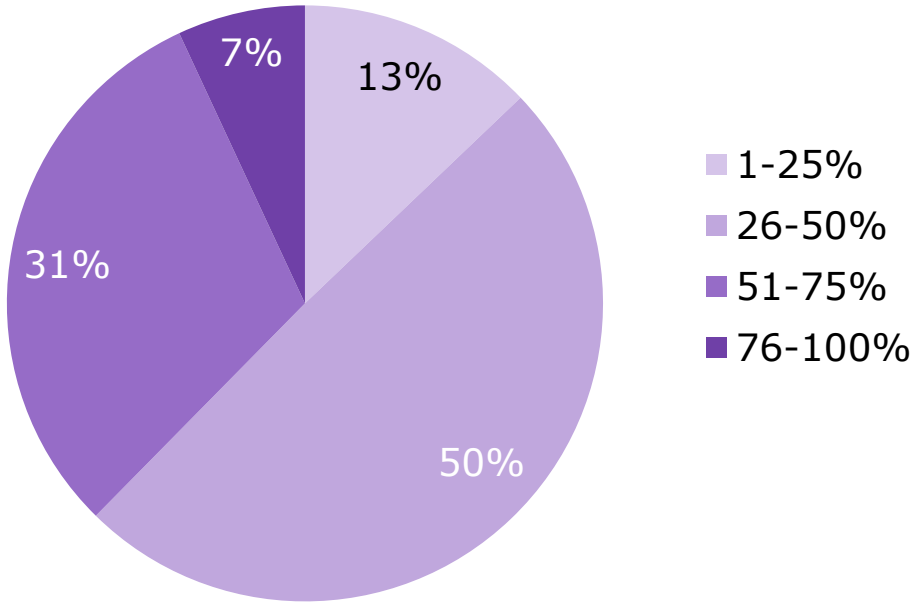

Nearly Half of Users Are Spending Valuable Time, 6-10 Hours, Maintaining Tags

On average, how much time do you spend adding or maintaining tags to correlate data per month?

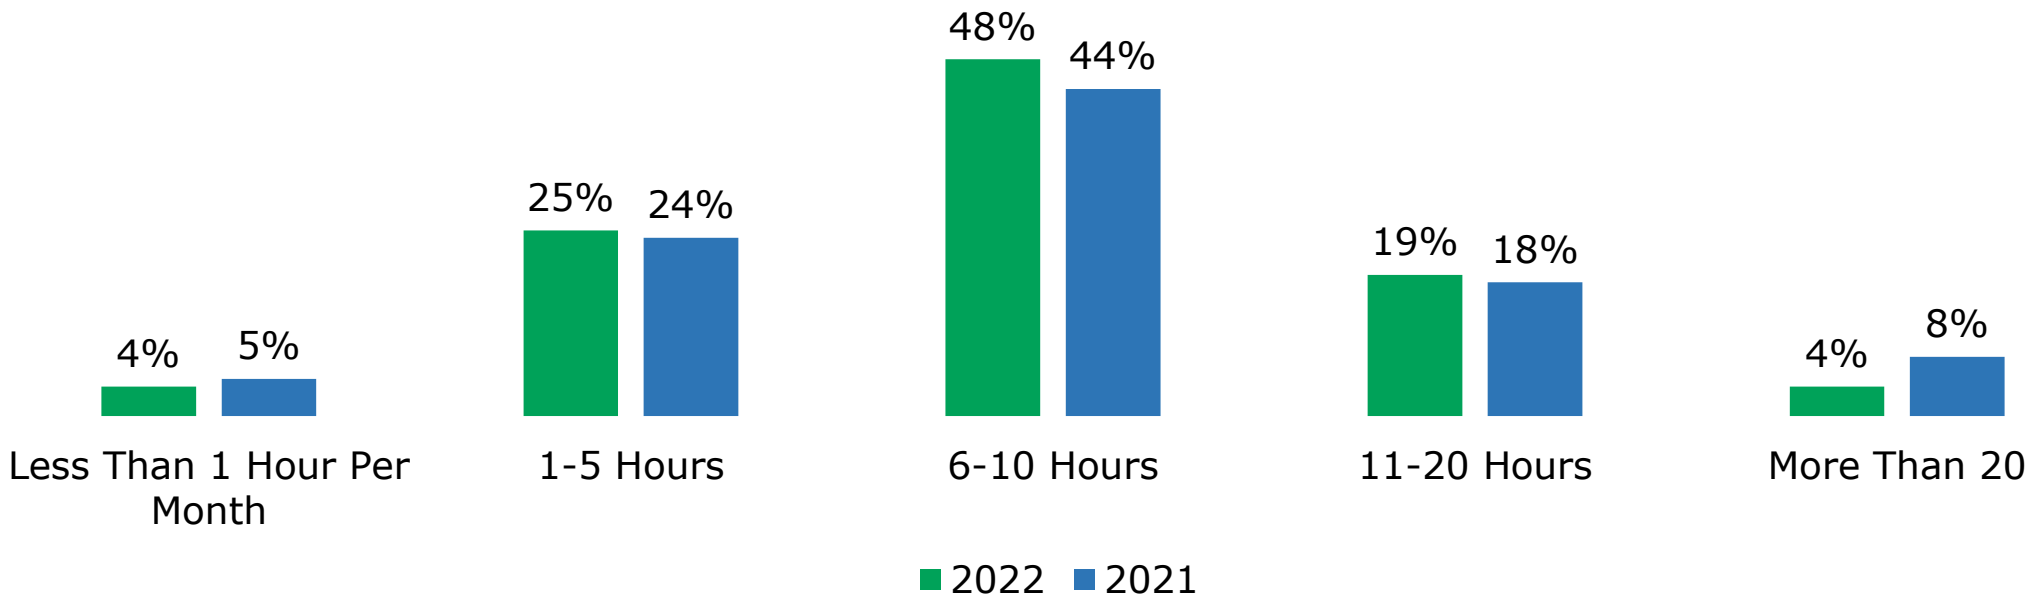

Longer-Term Data Retention Is Very Important To Most Organizations

94% of organizations consider longer-term telemetry data retention to be important, with 70% calling it very important.

Historical data can be increasingly useful, especially as more organizations are beholden to regulatory compliance and security-related concerns.

Cost of observability and monitoring tools, as well as infrastructure supporting them, has continued to be a major inhibitor of longer retention.

Tools that rely on indexing can encourage users to keep less data to avoid costs, resulting in less data on hand when it is needed.

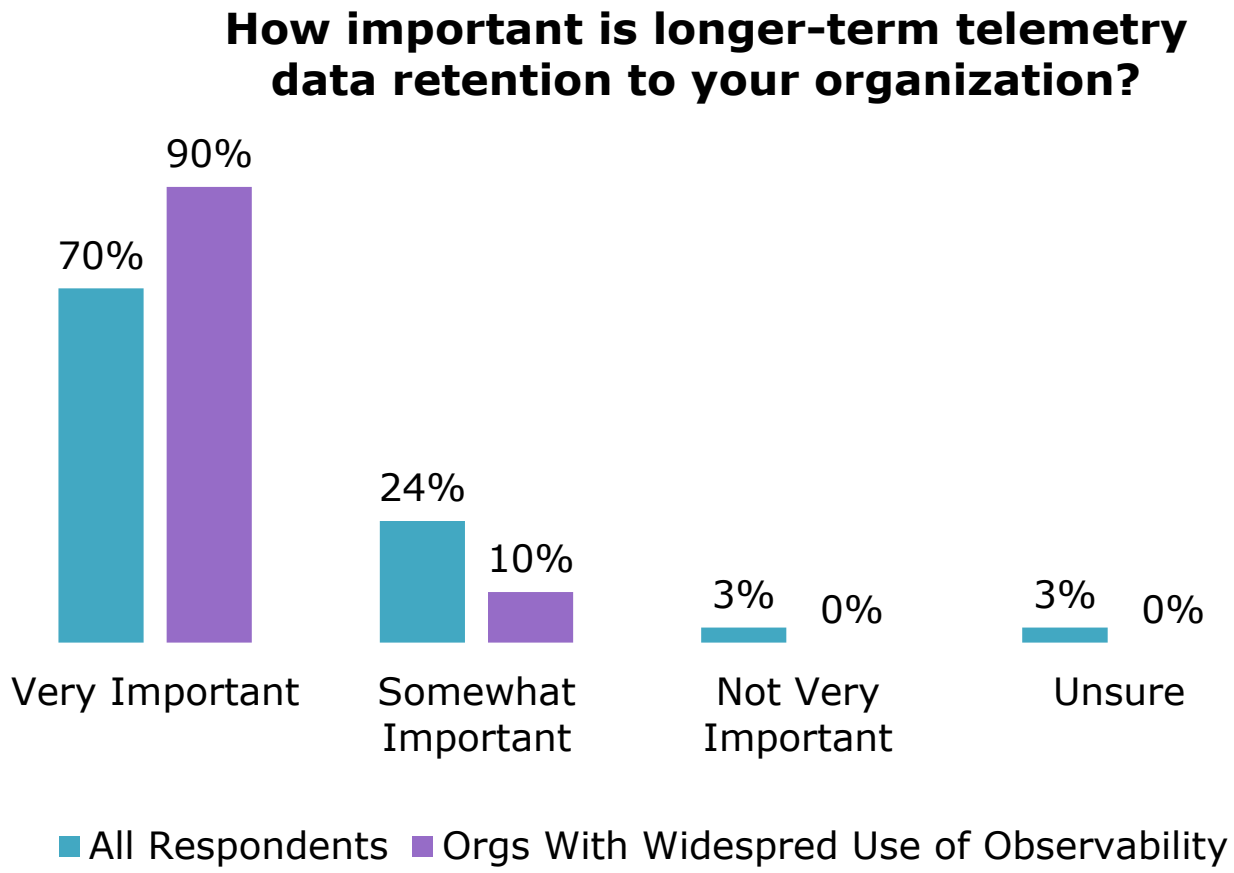

Almost Half of Orgs Throw Away Data Because of Cost

Not all telemetry has significant historical value, but if data is discarded because of cost and not because the data has limited value then there is a negative impact to the insights that can be gained.

For tools that charge based on ingest volume, users often end up sampling their data heavily when the cost to index is high. Users may have unintentionally "sampled out" the very thing they'll need later on.

Ideally, cost should not outweigh the desire to attain beneficial business outcomes.

In the past year has your organization discarded telemetry data you would have otherwise kept because of cost-related concerns?

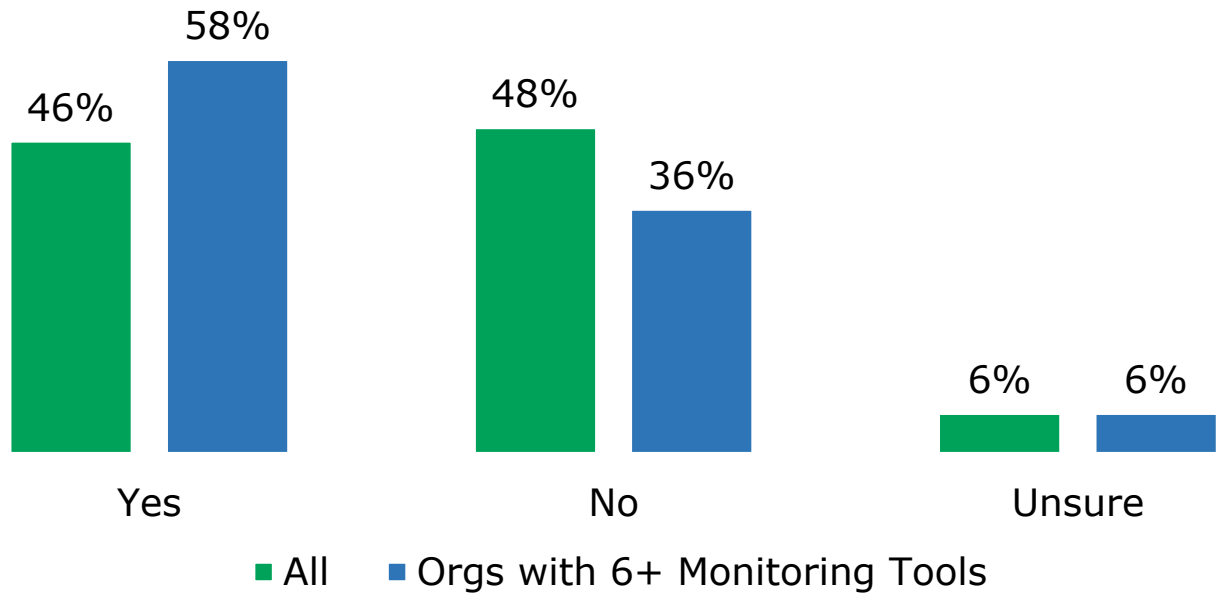

Part 5. Monitoring Tool Sprawl

Most Companies Plan To Adopt New Tools – Even Those With Many In Use

Most organizations (81%) indicate they plan to adopt a new observability tool in the next year. This has decreased slightly YOY, potentially the result of more organizations now looking to combat tool sprawl.

However, 89% of organizations with 11+ monitoring tools in their environment indicate they plan to adopt new tooling in the next year.

Despite the abundance of tooling that many organizations are experiencing, the search for new options indicates their existing or emerging needs are not being met.

Is your organization considering adopting a new observability tool within the next 12 months?

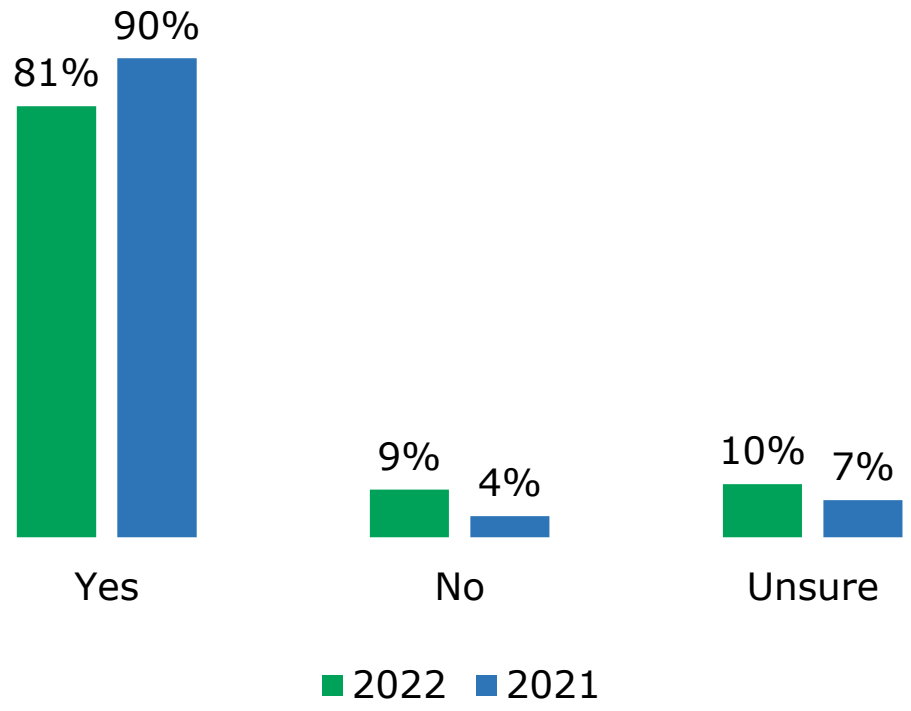

Tool Sprawl Continues To Be Pervasive, Worse at Orgs Earlier in Their Observability Maturity

Its common, but not always advantageous, to have many monitoring tools in use.

57% of orgs have 6 or more tools in their environment.

Orgs earlier in their observability journey are more likely to have greater tool sprawl.

How many different monitoring tools are in place across your organization?

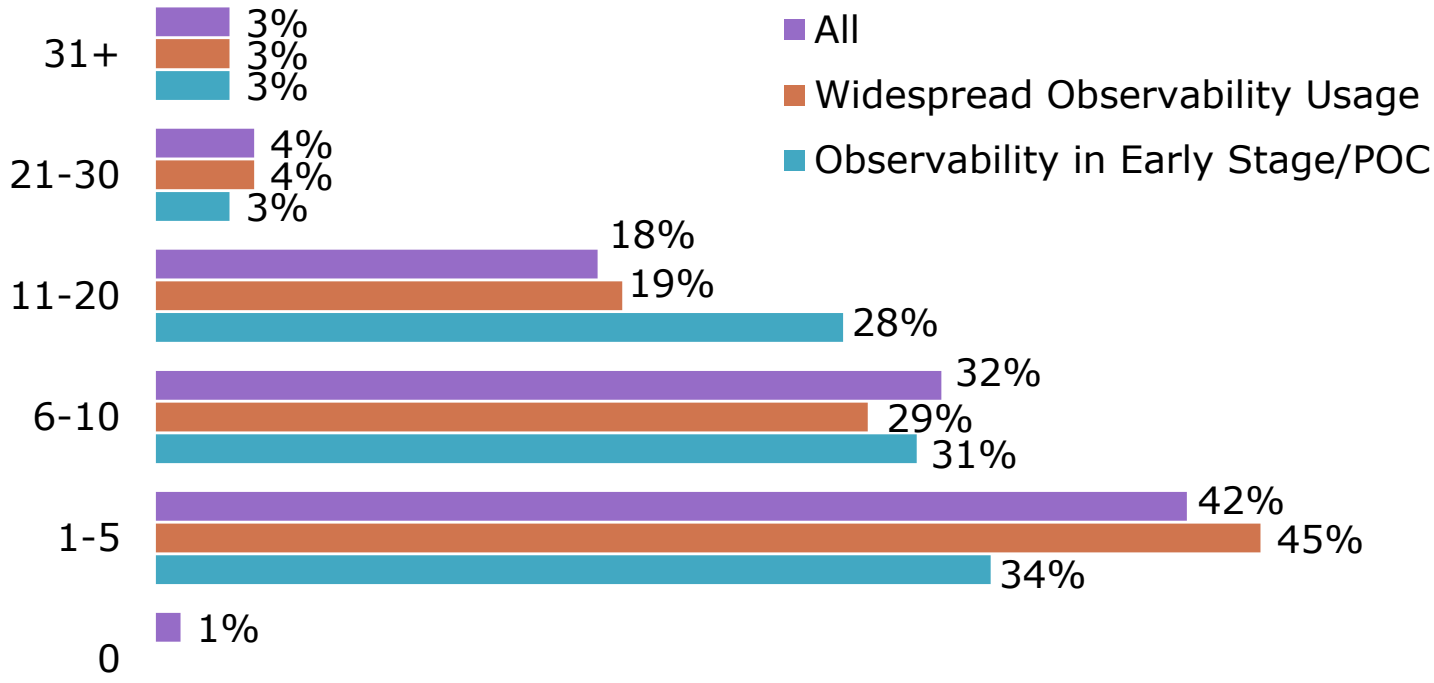

Using Many Tools in Investigations Is Common

Concerns around tool sprawl tend to be anchored around cost and wasted budget. However, the trend can result in silos that impact investigations.

70% or organizations are using 3+ tools on average to investigate an incident. The greater the number of tools, often times the more prolonged and complex an investigation can become.

Ideally users will be able to respond to alerts and dig into relevant data quickly without having to hop between different tools and dashboards.

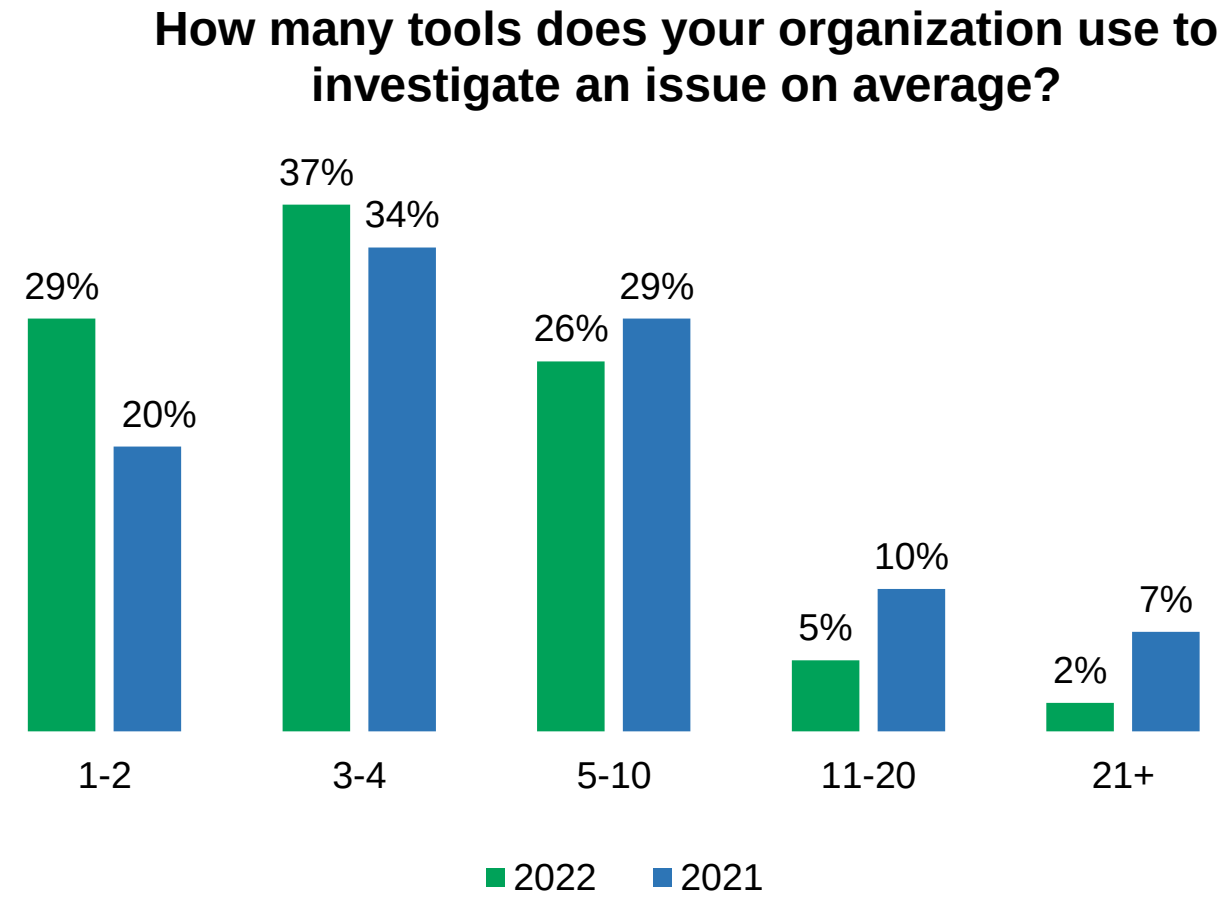

Complexity Is Now the Leading Challenge for Monitoring Tools

Tool complexity as a challenge with existing tools has increased significantly (24%) YoY.

As tools are applied to more complex use cases, legacy tooling itself can become more complicated, impacting user experience.

What are your biggest challenges with existing monitoring/troubleshooting tools? Please select the top 3 aspects that are most challenging.

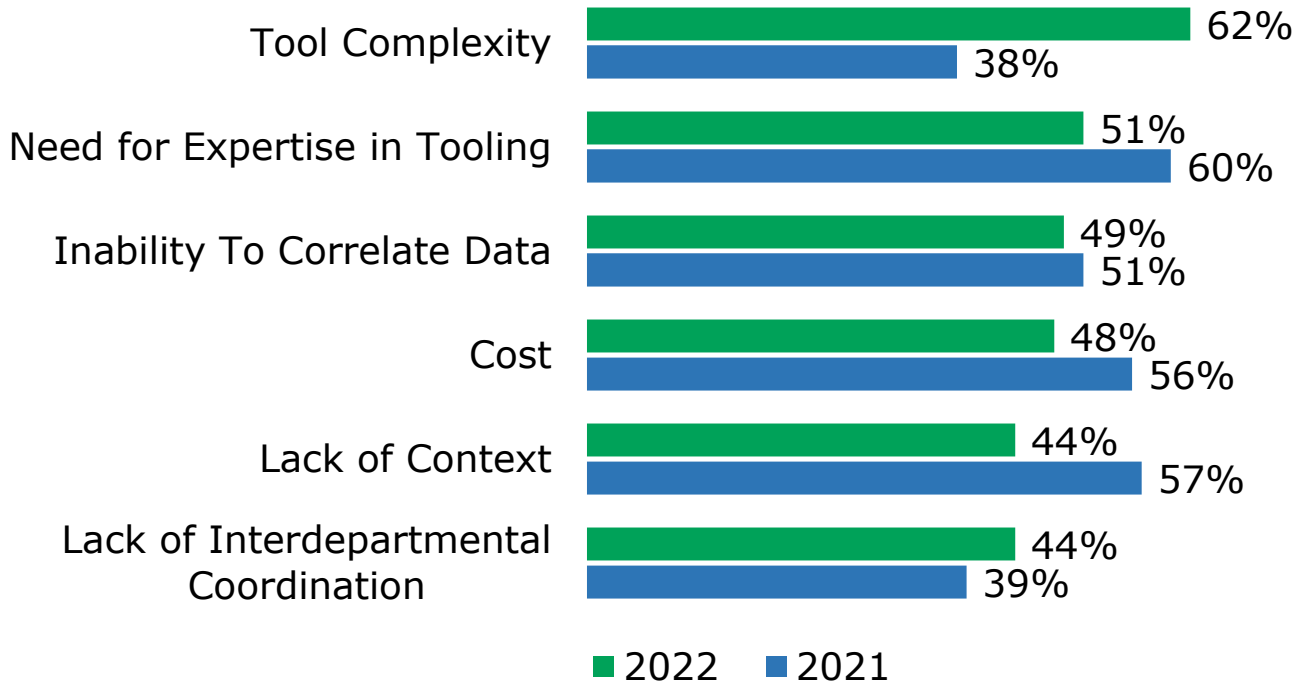

Part 6. Improving MTTR

What is MTTR?

Typically, MTTR is an abbreviation for Mean Time To Resolution/Remediation. It is the average time it takes for an incident to be detected and ultimately fixed. All organizations want to strive for shorter MTTR to minimize downtime and reduce the impact that IT incidents can have to the business. Achieving lower MTTR can be complex and include many variables including teams and tools needed to troubleshoot.

Most Organizations Experience 6 or More Incidents per Month

Incidents are relatively commonplace and most organizations (70%) experience 6+ a month.

As more organizations strive for increased release velocity we can expect that the quantity of incidents will remain high, underscoring the importance of tooling that can help reduce the time to troubleshoot.

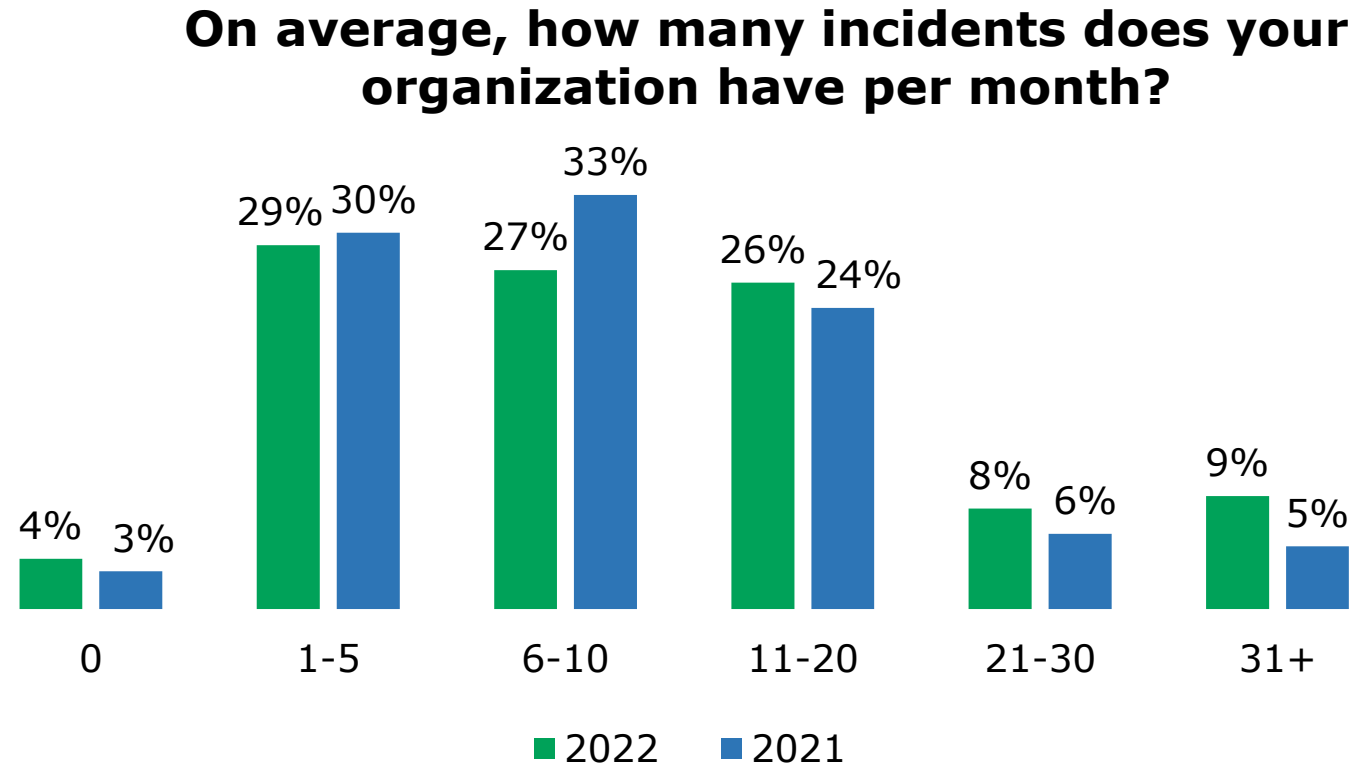

Orgs With More Data Can Experience Longer Investigations

Ideally more data should equate to better troubleshooting, but this isn't always the case.

Organizations ingesting 501GBs-1TB of data per day into their observability tools are more likely to spending a few days to a week on investigations.

Factors include tool sprawl creating data silos and the inability of some tools to effectively correlate data and understand context.

Users need to ensure that their data is an effective asset rather than a burden. Tooling should be able to provide context to make more data as useful as possible.

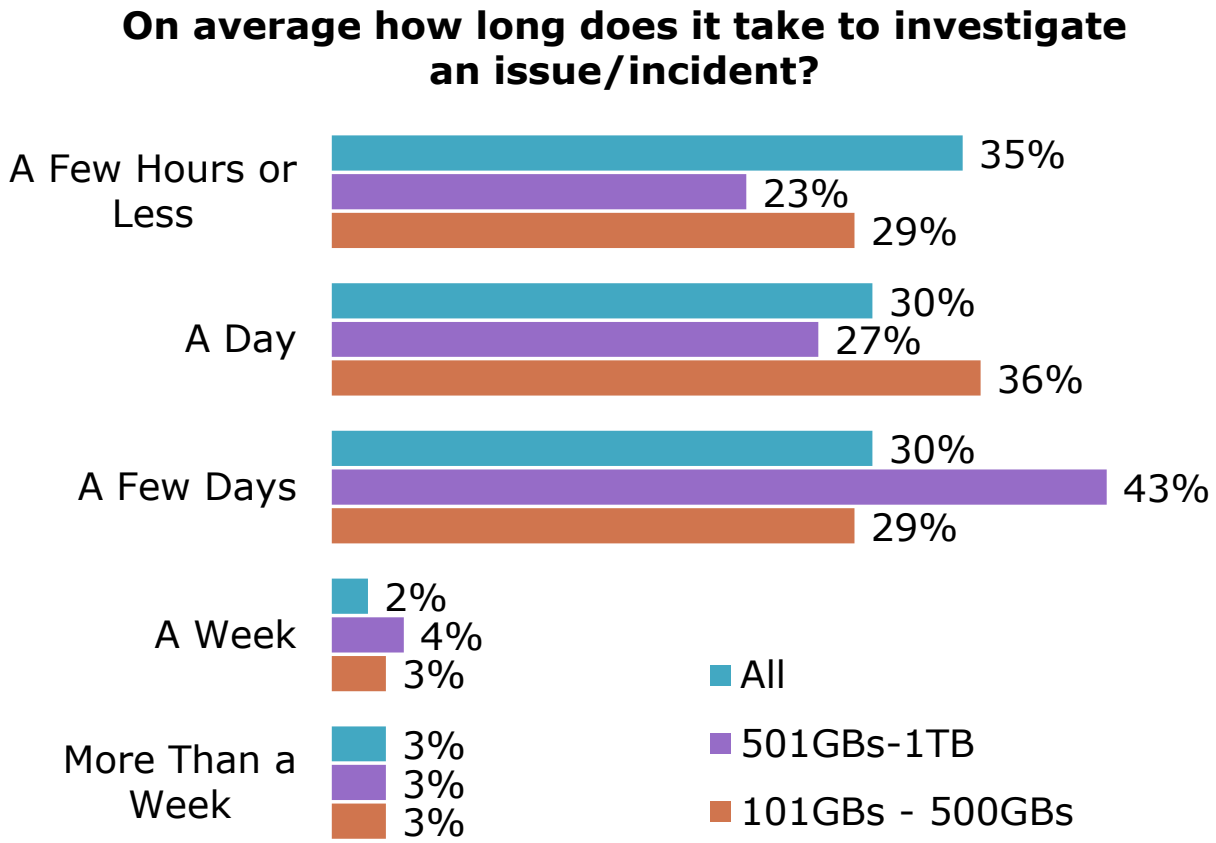

Existing Tools Are Considered Effective For Cloud Native Troubleshooting

More than half of users consider their existing monitoring tools very effective at troubleshooting cloud native apps.

However, similar numbers also report spending a day or more resolving incidents, indicating that what's considered effective may actually be incurring long MTTR.

Emerging technologies such as Kubernetes and the resulting complexity continue to strain existing tools that were not purpose-built for troubleshooting modern environments and applications.

How effective have your existing monitoring tools been for troubleshooting cloud native apps (those built with microservices, Kubernetes or serverless)?

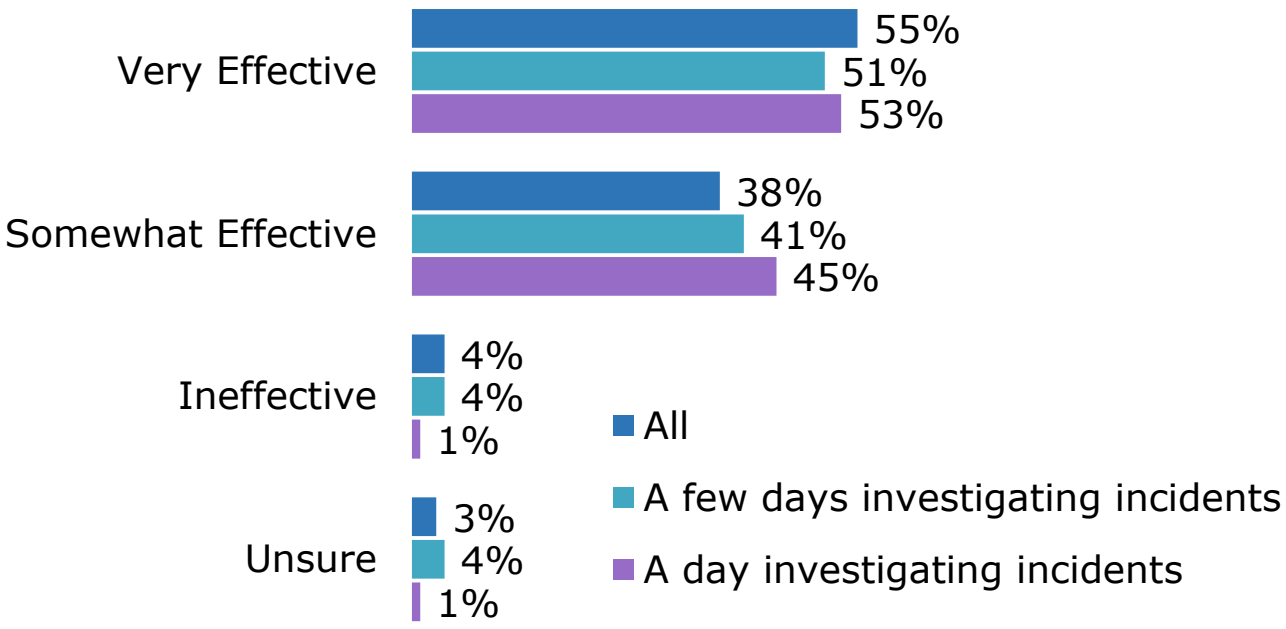

Appendix. Demographics

Demographic Infromation

Current Department	
CEO or C-Suite Executive	7%
IT	85%
Software Developement	8%

Current Role	
Manager, SRE	2%
Director of Observability	3%
Monitoring Engineer	2%
DevOps Engineer	5%
Manager, Operations	16%
CTO	30%
Director of DevOps	36%
Infrastructure Engineer	2%
Architect	4%
Other C-Suite	1%

Size of Engineering Team	
100+ employees	28%
51-100 employees	30%
21-50 employees	23%
11-20 employees	14%
Less than 10 employees	4%

Years in Business	
10 years or more	46%
5 to less than 10 years	33%
3 to less than 5 years	16%
1 to less than 3 years	6%
Less than one year	0%