middleware

Cloud Native Observability

The New Normal for Enterprises

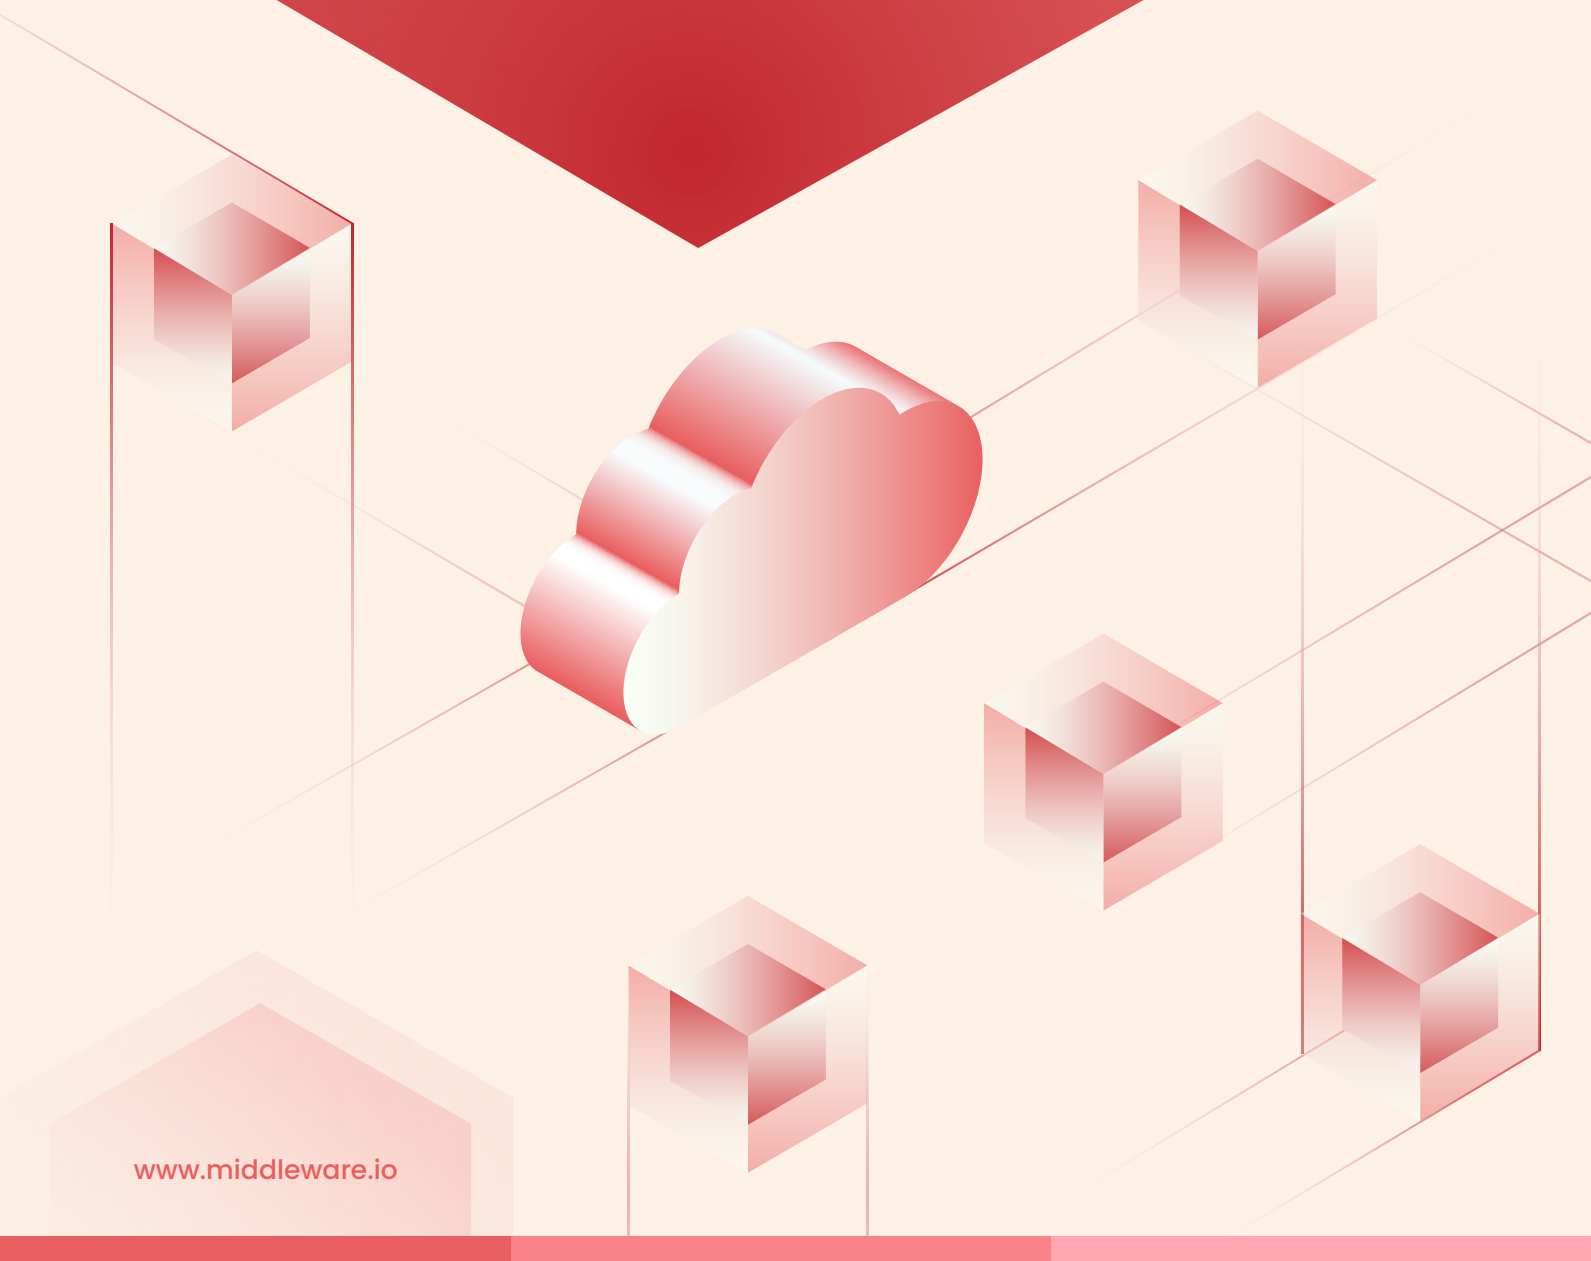

Table of contents

According to research from Gartner, by 2025, **70% of new applications developed by organizations will be cloud-native.** And with good reason: these systems are more agile, elastic, and scalable than outdated predecessors.

Unfortunately, however, many organizations are not investing in the necessary observability and monitoring tools to keep up with these changes. This results in the following problems:

- Higher costs and resource inefficiencies
- Clunky and frustrating user experiences
- Lack of responsiveness, reducing time-to-market and increasing MTTR
- Inability to adapt to more distributed business models (e.g. Web3)

As more companies adopt microservices, containers, and other key cloud-native components, monitoring solutions built for monolith systems will not be able to keep up. Data will continue to remain siloed, in separate warehouses—resulting in blind spots and increasing application errors.

This whitepaper will walk through a solution to this problem—cloud-native monitoring—and how it can and will transform modern application architecture.

> *"As more companies adopt microservices, containers, and other key cloud-native components, monitoring solutions built for monolith systems will not be able to keep up—cloud-native monitoring is the solution to this problem."*

1 What are the latest trends in cloud- native?

In 2022, global cloud revenue is estimated to total $474 billion, up from $408 billion in 2021. Over the coming years, Gartner analysts estimate cloud revenue will surpass non-cloud revenue for relevant enterprise IT markets.

Additionally, some reports indicate that as many as 75% of companies are focusing on new cloud native applications. By 2025, 70% of new applications will be developed using low- or no-code technologies, up from less than 25% in 2020.

More than

75%

of the organizations using cloud services, indicate they have a cloud-first strategy.

As more platforms and systems become cloud native, IT teams that don't respond accordingly will put themselves at a sharp competitive disadvantage. There are a variety of appropriate responses, but one critical one is the presence of elastic, scalable, and on-demand observability & monitoring.

"As many as 75% of companies are developing cloud native technologies."

2 How is cloud native observability different from legacy approaches?

With each subsequent advancement in cloud computing, there have been numerous advantages to businesses, particularly when it comes to expanding observability functions. ***"Cloud native monitoring is the new normal,"*** **says Middleware CEO Laduram Vishnoi.** "To understand how we got to this point, we need to understand the evolution of observability over the past several decades."

Cloud observability

As companies embraced cloud applications as core functions, observability transformed. Resources could be run only when needed, resulting in greater control with lower costs. Additionally, the cloud opened the door to automation, which became more prevalent for cloud operations and management, including provisioning, configuration, scaling, and self-healing.

Pre-cloud observability

Pre-cloud software development followed a waterfall approach. Software took years to create, with infrequent updates and new releases. Development and QA, were manual, limited in scale, and often siloed. As a result, pre-cloud observability was slow, tedious, and limited in the insights available to developers and engineers.

Cloud-native observability

Now we are in the cloud native era. Application monoliths are being replaced with new software architectures, including microservices and serverless functions. As infrastructures become more distributed and abstract, developers and engineers must approach observability in a new way.

3 What is cloud native observability?

To effectively monitor cloud native applications, observability has undergone a new evolution. This new approach is built for distributed services, microservices, containers, serverless computing, and other key aspects of cloud native infrastructure.

One of the main reasons for a new approach is that cloud native applications can exist in public, private, and hybrid clouds, and do not have the same level of centralization that pre-cloud or cloud applications do. This presents a number of benefits:

- Designed for distributed computing resources
- Elastic, responding dynamically to workload changes
- Easy to deploy and manage on-demand
- Resilient and redundant to protect against failure

Unexpected consequences:

Cloud-native impact on observability data volumes

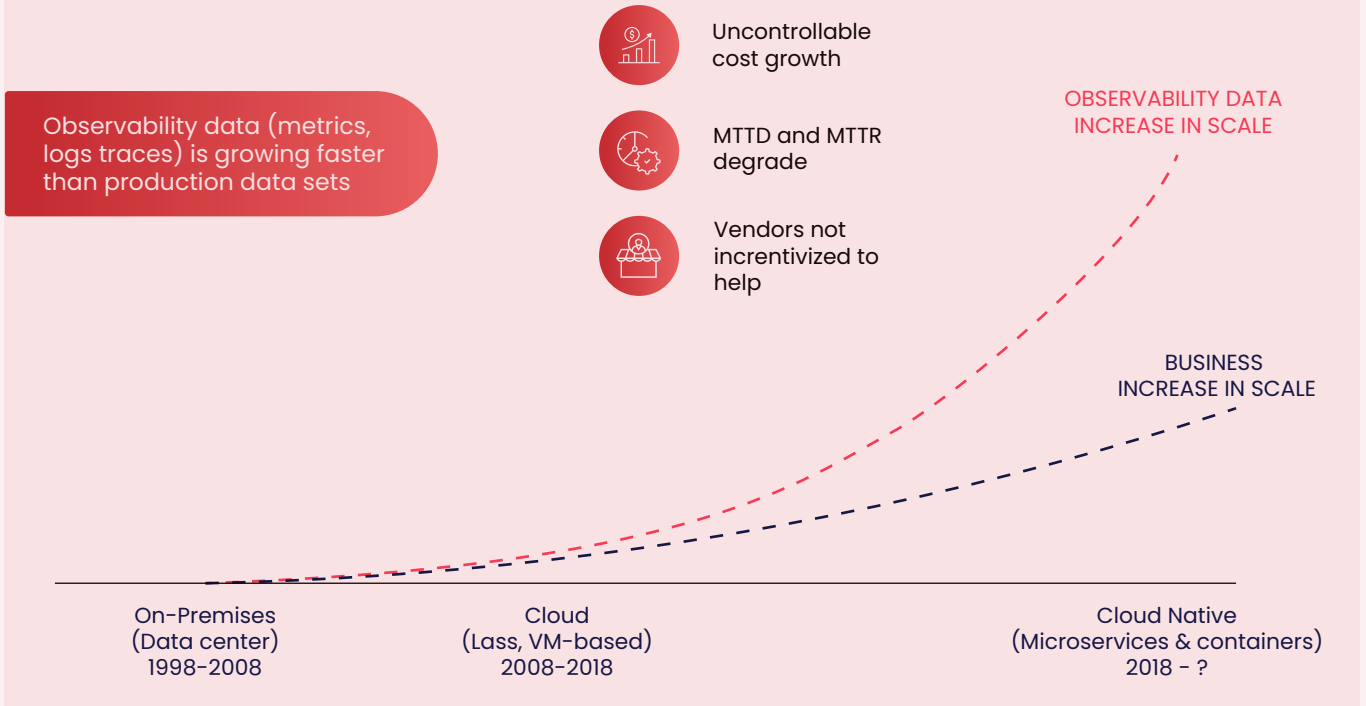

Cloud native observability is built to adapt to the unique challenges that developers and engineers face when monitoring cloud native applications:

- Developer fatigue and unbalanced call rotations
- Communication and collaboration hiccups and growing pains
- Lack of contextual, intelligent insight into performance, usage, and process changes
- Limited default observability capabilities predetermined by cloud suppliers
- Striking the balance between collecting machine data and drawing meaningful insights from it
- Growing concerns about security and cost in distributed systems

According to Vishnoi, "Putting metrics, traces, and logs in a unified dashboard provides unprecedented observability power. Developers and DevOps don't need to go back and forth to see—and fix—problems. They can go directly to the dashboard, select the timeline, and pull up the relevant metrics, traces and logs within seconds."

As users expect continuous innovation and unparalleled responsiveness, business systems must become more strategic and increasingly flexible. Cloud native observability enables organizations to achieve transformational digital and business outcomes without expending too many resources. Here are some specific functions that make that possible.

> *"Cloud native observability enables organizations to achieve transformational digital and business outcomes without expending too many resources."*

Metrics

Metrics are the numerical values that represent and describe the overall behavior of a service or component measured over time. Examples include timestamps, names, and values. Because they are structured by default, they're easy to query and optimize for storage.

For observability purposes, metrics enable engineers to define what is normal and what is not. However, they do have limits. When triggered, they can indicate when maximum or minimum thresholds are reached, not why the issue occurred or what the user experiences on the front end. Those insights require additional pillars of observability.

Traces

If metrics tell you that an issue is occurring, traces can help you investigate the precise service causing the issue, enabling developers and engineers to identify and fix the root cause quickly. Through traces, engineers can analyze request flow and understand the entire lifecycle of a request in a distributed application, as each operation is encoded with critical data related to the microservices performing that operation. For cloud-native applications, traces are critical because of the increased distribution of many systems.

Logs

Logs are immutable, timestamped records of discrete events that happened over a set time frame within an application. Developers can use logs to uncover emergent and unpredictable behaviors within each component in a microservices architecture.

There are three types of logs:

Plain text

Structured

Binary

Every component of a cloud native application emits one of these log types. This can lead to a lot of noise. Thus, it is critical for engineers to centralize and analyze these logs to uncover actionable insights from them.

Serverless monitoring

Serverless is an architecture style that enables developers to write code on a platform that functions as a service (FaaS), using event-driven architectures and various backend-as-a-service (BaaS) models.

U.S. serverless architecture market size, by service, 2015-2025 (USD Million)

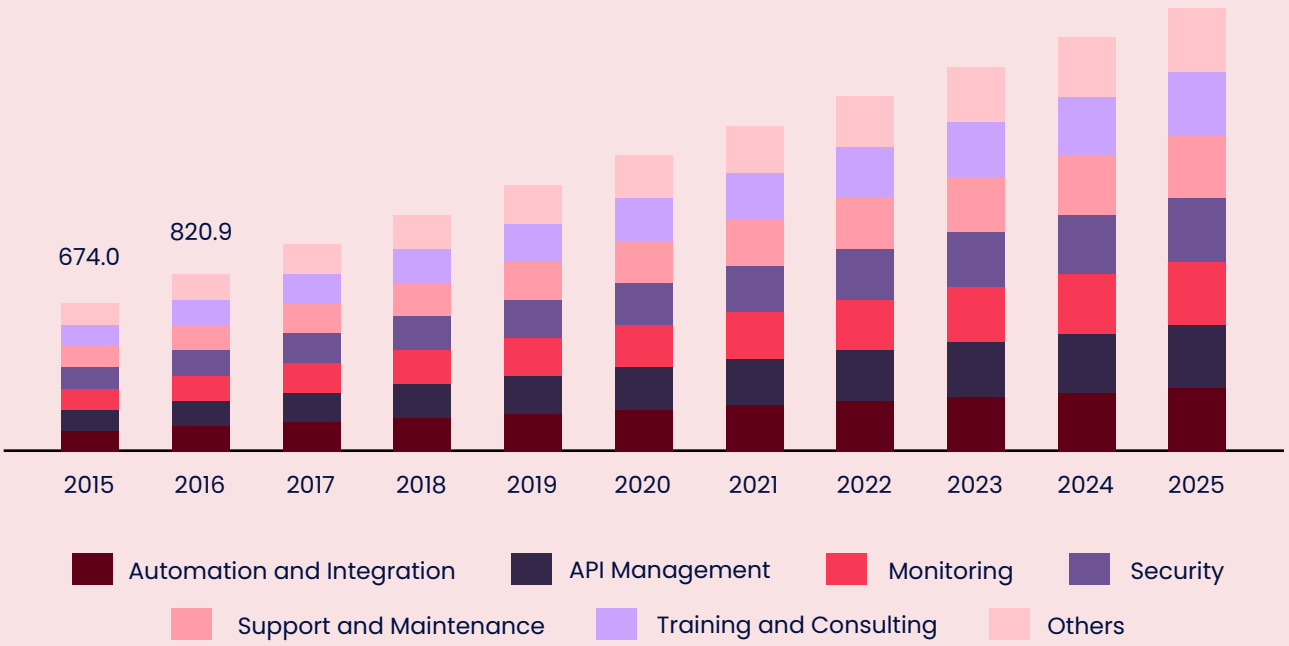

The primary benefit of serverless is that applications are broken up into small pieces of code, then dynamically scheduled and run on demand when triggered by an event. This architecture eliminates the need to worry about provisioning, patching, scaling, security, high availability, etc., since code is called and executed only when needed.

Serverless functions are also highly resource-efficient, since you only pay for resources used during the execution duration.

"Serverless architecture eliminates the need to worry about provisioning, patching, scaling, security, high availability, etc., since code is called and executed only when needed."

Container monitoring

Containers are a fundamental component of a cloud native architecture. They can help accelerate the development and deployment processes, make workloads portable and even mobile, and improve your overall agility when deployed. For distributed cloud native environments understanding which containers are responsible for which issues is key.

4 How are enterprises overcoming cloud native monitoring challenges?

The rise of cloud native applications has spurred many companies to reexamine their monitoring and observability infrastructure to keep up with emerging needs.

According to data from Dynatrace, **74% of CIOs believe** increased complexity could make it extremely difficult to manage performance, and 76% of CIOs say that they don't have complete visibility into application performance in cloud-native architectures.

What's more, IBM published research in 2021 stating that **58% of leaders believe** the complexity of managing hybrid/multi-cloud environments is their leading challenge.

What practical, technical, or cultural challenges have you experienced or do you foresee using these projects?

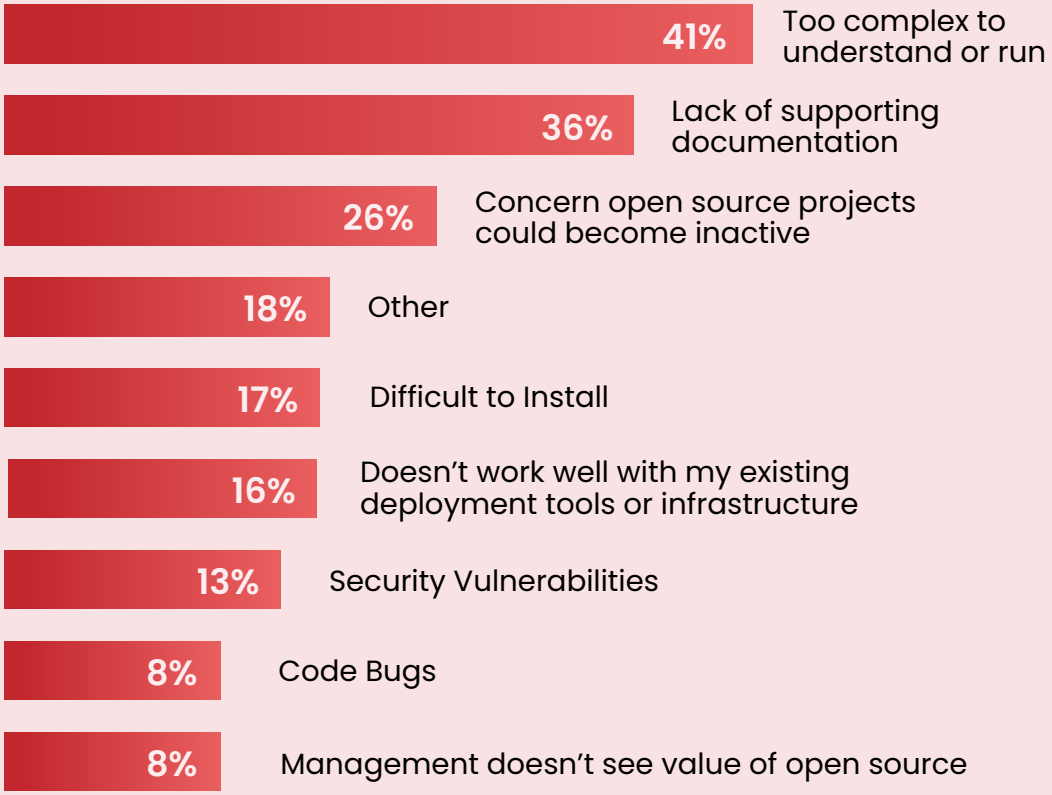

A 2018 study by DORA (DevOps Research and Assessment team) showed that teams that met the following criteria were more likely to be successful in cloud native implementation:

- Adopted essential cloud characteristics (on-demand self-service, broad network access, resource pooling, rapid elasticity, measured service)
- Leveraged infrastructure as code and modern deployment and pipeline practices
- Took advantage of platform-as-a-service (PaaS)
- Adopted cloud native design practices

One tool at many enterprises' disposal is open source software. Open source capabilities will continue to play a significant role, even in those cases where it's not front and center, but integrates into other essential tools.

Aside from open source, other solutions include leveraging existing vendors, new vendors, and even creating DIY tooling. Many organizations are leveraging at least one of these four tactics. Some are also turning to larger vendors, who have more resources and capabilities to address these needs.

Among these options, 83% of organizations prefer to buy monitoring and incident response tools from a single vendor when possible. This is likely because of the simplicity of leveraging an already established relationship, as well as the ease with which a single vendor can provide holistic observability across all of an organization's assets.

> *"83% of organizations prefer to buy monitoring and incident response tools from a single vendor when possible."*

5 Benefits of implementing full-stack cloud native monitoring tools

At Middleware, we provide a full-stack cloud native monitoring solution, bringing all your metrics, logs, and traces into a single timeline, and empower your developers and DevOps to debug and fix issues faster.

According to data from New Relic, **66% of respondents** who indicated that they had already prioritized or achieved full-stack observability **experienced high-impact outages less frequently than those who had not.**

The same study demonstrated that 68% of respondents who said they had already prioritized or achieved full-stack observability took fewer than 30 minutes to detect outages, and 61% took less than 30 minutes to resolve them.

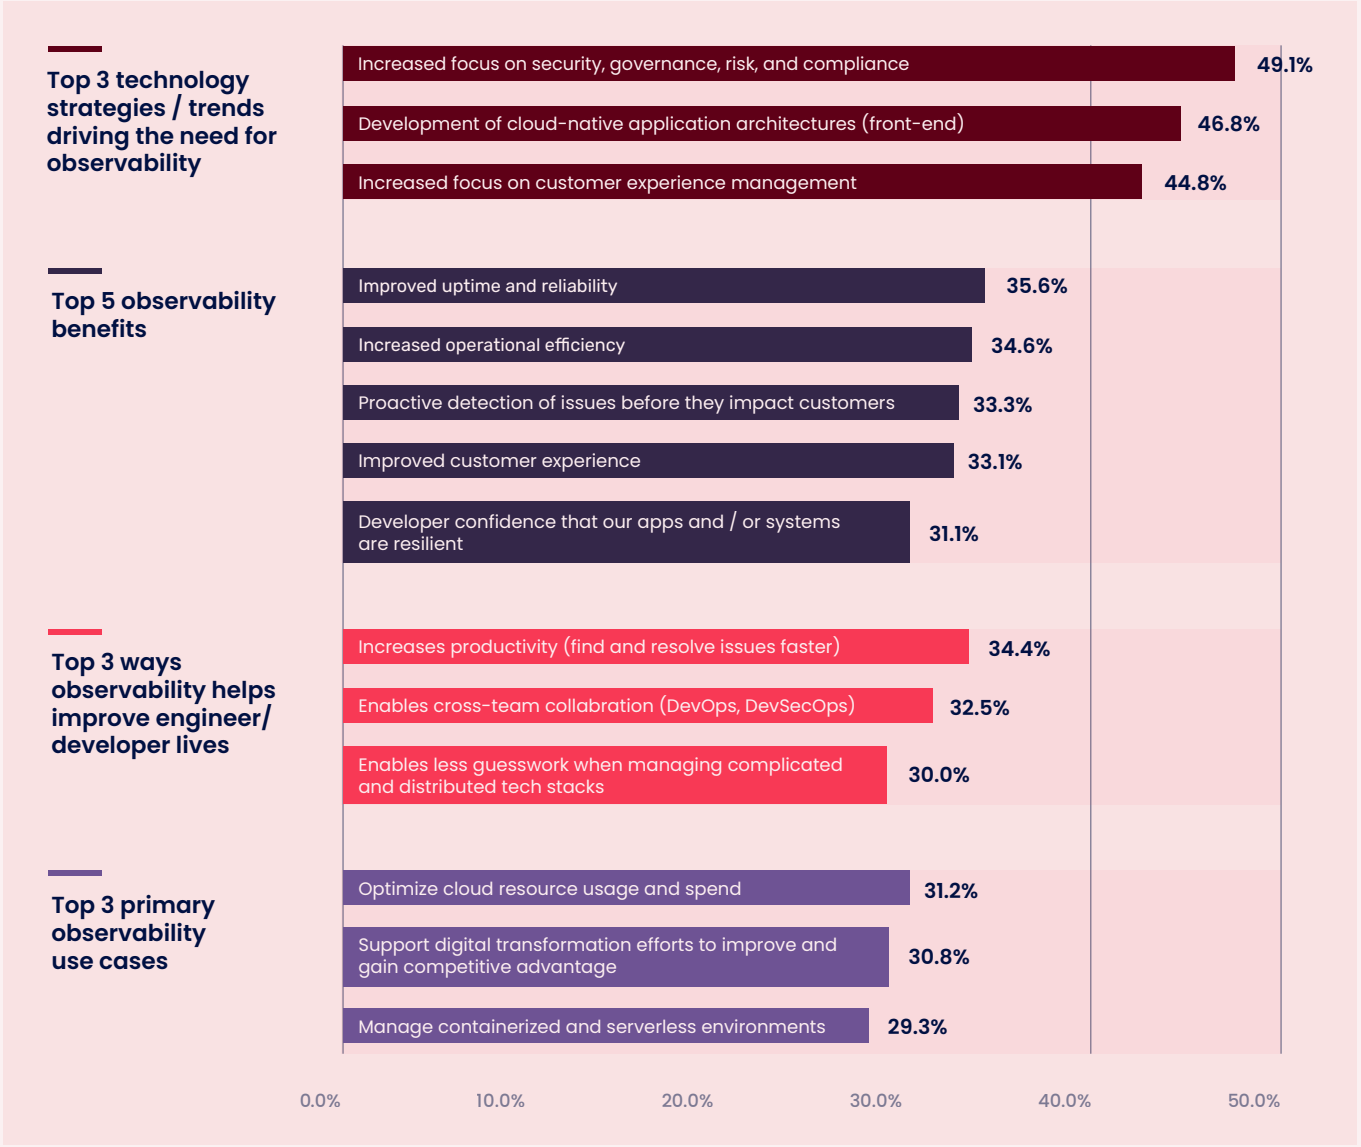

"68% of respondents who said they had already prioritized or achieved full-stack observability took fewer than 30 minutes to detect outages"

There are a number of reasons why full-stack cloud native observability enables organizations to detect, identity, and resolve issues quickly. Here are a few.

Business agility

With correlated data that gives you insight into your entire tech infrastructure, full-stack cloud native monitoring helps you measure and track software delivery and business results, responding to issues in real time.

- Faster launch times for new features and experiences
- Insight-driven innovation and experimentation
- Stronger technology alignment with business goals and outcomes
- Objective decision making based on a single source of truth

With correlated data that gives you insight into your entire tech infrastructure, full-stack cloud native monitoring helps you measure and track software delivery and business results, responding to issues in real time.

Resiliency and reduced downtime

Full-stack cloud native monitoring helps you keep track of your entire array of applications, containers, and functions. When an outage does occur, track interconnecting dependencies to identify the issue and resolve it at the source.

Automated, continuous delivery

By using DevOps automations and features, cloud native applications enable continuous delivery and deployment of software changes. Methods like blue-green and canary deployments also enable application improvement without any major disruptions to the user experience.

Seamless integration of multiple sources

Cloud native monitoring seamlessly integrates data and insights from a wide range of cloud sources, including AWS, GCP, Azure, and more. The agility and elasticity of the cloud native approach enables an organization's observability infrastructure to map directly to each data source, without the need for complex and unwieldy integrations.

"The agility and elasticity of the cloud native approach enables an organization's observability infrastructure to map directly to each data source, without the need for complex and unwieldy integrations."

360-degree data visualizations

Data visualizations are a powerful tool to convert raw numbers into actionable insights. However, many visualizations lack data from all relevant sources, resulting in gaps or misleading information. Cloud native monitoring ensures that all visualizations present a full picture of the organization's tech stack.

Monitor multi-cloud environments as you migrate

For multi-cloud environments, migrations can be a pain. Cloud native monitoring helps developers and engineers track on-prem and multi-cloud monitoring data in one central location, which helps ensure timely project execution.

6 Middleware: One platform. Infinite benefits.

If you're on board with the benefits of cloud native, full-stack visibility, then Middleware is here for you. Ditch your legacy tech and leverage a unified platform to analyze, diagnose, and predict issues across your entire tech stack.

Middleware is an integrated, scalable system that will help you un-silo your data and insights from all your containers, empower you to identify root causes and solve issues in real time, and give you the best value for money with a platform that fits to your specific needs.

Bring all your metrics, logs, and traces into a single timeline, and empower your developers and DevOps to debug and fix the issue faster—reducing downtime and improving the user experience.

- To learn more about what we have to offer, visit https://middleware.io/

References

01 Source: Gartner, 2022: https://www.gartner.com/en/newsroom/press-releases/2021-11-10-gartner-says-cloud-will-be-the-centerpiece-of-new-digital-experiences

02 Source: Gartner, 2022: https://www.gartner.com/en/newsroom/press-releases/2021-11-10-gartner-says-cloud-will-be-the-centerpiece-of-new-digital-experiences

03 Source: VentureBeat, 2022: https://venturebeat.com/data-infrastructure/report-75-of-companies-are-focusing-on-cloud-native-apps/

04 Source: Dynatrace, 2021: https://www.dynatrace.com/global-cio-report/

05 Source: IBM, 2021: https://www.ibm.com/downloads/cas/VKW3RNDP

06 Source: Cision News, 2018: https://www.prnewswire.com/news-releases/devops-research-and-assessment-dora-announces-the-2018-accelerate-state-of-devops-report-300703837.html

07 Source: Chronosphere, 2021: https://go.chronosphere.io/rs/098-HTR-317/images/451_Reprint_CloudNativeMonitoring.pdf

08 Source: New Relic, 2022: https://newrelic.com/blog/nerdlog/2022-observability-report#toc-prioritizing-full-stack-observability-leads-to-fewer-outages-and-faster-mttd-and-mttr-