



How modern cloud observability addresses the complexity wall



What's inside

Introduction

Chapter 1

What is observability?

Chapter 2

The role of observability in digital transformation

Chapter 3

How modern observability tames cloud complexity

Chapter 4

Essential attributes of a modern cloud observability platform

Chapter 5

How Dynatrace delivers modern cloud observability

Introduction



Many organizations are hitting an IT complexity wall, which is fueling an urgent need for observability.


As organizations adopt modern cloud technologies to become more competitive, efficient, and innovative, they have also encountered the increased complexity these technologies bring.

Multicloud applications and infrastructure use distributed, containerized architectures, open source software, and dynamic microservices that come and go in seconds. At enterprise scale, this complexity often becomes an opaque wall. In response, IT operations, development, and security teams have come to recognize they need observability of their distributed modern architectures to effectively manage and secure them.

In this e-book, we explore what observability is and its role in managing modern cloud environments. We also examine why observability achieved using automation and artificial intelligence (AI) has become key to taming multicloud complexity and automatically delivering answers to make software perform better—and more securely—at scale.

CHAPTER 1

What is observability?



Observability is the ability to measure a system's current state based on the data it generates and the ability to predict the future behavior of systems with data analysis and other capabilities.

Legacy approaches to observability have centered on monitoring a small number of applications and infrastructure components. These approaches rely on reactive alerting (or merely responding to alerts after they happen). As a result, organizations can fix issues only after they occur and have negatively affected the user experience. As applications have modernized using cloud-native technologies, traditional monitoring approaches can't keep up.

The three pillars of observability

Observability comprises three major elements: logs, metrics, and traces. These three elements provide a picture of activity taking place in an IT environment and information about the relationships among its components.




Logs

A log is a data file that contains a record of activity in an IT environment. Every entity in a data environment generates logs, including applications, services, containers, open source software, operating systems, and infrastructure.



Metrics

Metrics are the myriad measurable data points that IT nodes generate that enable teams to determine elements such as application performance, real user activity, and cloud platform statistics.



Traces

Traces provide the relationships of transactions and their dependencies among various services, apps, and infrastructure.

In a distributed microservices-based environment that spans multiple clouds and open source technologies, however, the three pillars of observability are often not enough. Observability must also include context and user data.

Context

To provide actionable intelligence from the deluge of data generated by logs, metrics, and traces, analysts also need access to the context of an event or an alert. Is the event just part of a routine correlation of conditions, or is a critical function being blocked by a malfunctioning code snippet? Understanding context enables analysts to pinpoint root causes and filter out alerting noise.

User data

Analysts also need to understand how the transit of data affects the end-user experience. By tracking real user sessions, analysts can gain real-time insight into user experiences.

A modern observability strategy

To overcome the cloud complexity wall and satisfy customers, enterprises must rethink their observability strategy.

Reactively responding to issues by sifting through a sea of data and guessing at root causes based on correlations is not effective at cloud scale.

Teams need a modern approach to observability that enables them to move from reactive firefighting to proactive problem solving that fosters increased automation. With a platform approach to observability, one that uses artificial intelligence (AI) and automation to discover and monitor all microservices, teams can pinpoint answers proactively before users are affected. By simplifying observability, teams free up time they used to waste on manual tasks so they can accelerate innovation and improve customer experience.

A modern observability strategy involves three key capabilities: sensing, thinking, and acting.



Sensing

IT teams need the ability to gather data automatically, in real time, and in context to the other applications, infrastructure, and services with which it interacts.



Thinking

Teams then need a way to automatically analyze and map that data and the potentially hundreds of components and millions of dependencies it touches.



Acting

With this contextualized data, teams need a way to automatically generate real-time answers so they can automate responses and have more time to innovate.

Figure 1: Following this formula, teams can adopt an approach to observability that can scale with them on their digital transformation journey.

CHAPTER 2

The role of observability in digital transformation



To understand why observability has become so crucial for modern IT environments, it helps to have a clearer picture of the digital transformation that's driving the need for it.

The forecasted global spending on digital transformation will reach \$2.8 trillion in 2025, according to the International Data Corporation (IDC) Worldwide Digital Transformation Spending Guide. This is double the amount allocated in 2020.

Let's explore some of the components of multicloud architecture that advance digital transformation but also contribute to cloud complexity.

Serverless platforms

Dynamic multiclouds enable enterprises to adopt serverless functions and Kubernetes infrastructure across public clouds such as Amazon Web Services, Google Cloud Platform, and Microsoft Azure.

Microservices architecture

These platforms facilitate microservices architecture, an approach to application design that breaks applications into small, independent services that work together to deliver flexible functionality. Microservices make development more agile and enable teams to scale performance more effectively and reliably.

Open source software

Open source software, ranging from operating systems to JavaScript libraries, offers developers flexibility and access to configurable modules of code they can use to speed development and increase functionality.

Containers

A key enabler of these new environments is the move to containerized application models. Introducing containerized architecture typically requires a significant transformation, but offers new benefits in scalability, reliability, agility, and performance.

All these technologies enable teams to be more agile and adaptable. But they also increase the number of elements teams need to track by orders of magnitude.

Observability to cut through modern complexity

As microservice environments become increasingly dynamic and scale to hundreds of thousands of hosts, the real challenge becomes making sense of data in real time, within the context of the entire technology stack. Practitioners need to quickly understand the impact on users and prevent deleterious issues from proliferating. This can be a daunting task that quickly surpasses the capacity of even the most skilled and experienced human operators.


At the same time, the increasingly rapid pace of organizational change has become a key driver in the need for observability. The same flexibility and diversity that enable cloud-native technologies to deliver technical and business benefits for fast-moving organizations also create operational complexity challenges and overwhelming choices.

In the report *“Bridging the Cloud Transformation Gap”*, which evaluates the findings of Aptum’s *“Global Cloud Impact Study”*,

62% of respondents cite complexity and a profusion of choice as a hindrance when planning a digital transformation.

The consensus today about digital transformations is that complexity will be the major enterprise challenge for the foreseeable future. Enterprises that address complexity issues before or during the transformation will fare better than those that address it after the fact.

Practitioners need to quickly understand the impact on users and prevent deleterious issues from proliferating. This can be a daunting task that quickly surpasses the capacity of human operators.



of respondents cite complexity and abundance of choice as a hindrance when planning a digital transformation using multicloud technologies

Source: Bridging the Cloud Transformation Gap

CHAPTER 3

How modern observability tames cloud complexity



The cloud complexity wall (also known as the cloud observability wall) is the proliferation of data that accompanies a distributed cloud-native architecture. There is no doubt that moving to a containerized, microservices-based approach to software development provides benefits for DevOps initiatives and increases reliability and scale.

But the sheer scope of telemetry generated by these multicloud environments produces a volume, variety, and velocity of data that is beyond what humans can manage unaided.

The three Vs of data

Volume

With hundreds of services, users are generating more data than ever before. IT environments generated 79 zettabytes of data worldwide in 2021.

Variety

Telemetry data varies widely both within clouds and between clouds. Teams need a way to embrace all data types, no matter where they come from, in a way that identifies causal relationships.


Velocity

With containers spinning up and down in milliseconds, modern systems generate data at an increasingly rapid pace. Teams need a way to capture and automatically analyze all that information so they can have visibility into issues.

Data silos

Another key challenge of the complexity wall is that most organizations operate in silos. These silos prompt teams to use different tools in dynamic multicloud applications. The Global CIO Report echoes the impact of silos: "49% of CIOs say IT and business teams work in silos." As a result, when problems arise, teams end up in lengthy war rooms and finger-pointing exercises since there is no central source of truth.

The typical impact of hitting the complexity wall is that teams become consumed with troubleshooting application issues resulting in negative customer experiences. But this also slows innovation because development teams are consumed with triaging and fixing issues versus innovating new features and capabilities.



of CIOs say IT and business teams work in silos

Source: Global CIO Report

How complexity undermines business impact

Ultimately, hitting the complexity wall hampers business results. The business value of digital transformation is to create a foundation that supports speed-to-solution and agility. If teams are consumed with tactical issues generated by complexity, then the business will quickly see a negative impact.

As Figure 1 illustrates, the amount and types of cloud complexity will worsen for many organizations compared to their digital transformation goals. Without the ability to mediate this complexity, true digital transformation will become impossible.




Figure 2: Reasons why cloud complexity is becoming a huge issue for enterprises.

When they hit the complexity wall, teams often progress through different stages on their journey through digital transformation.

Stage 1: Continue using legacy tools

Teams often start out using the same monitoring tools they've always used. The advantage is minimal disruption to existing practices, but modern multicloud environments quickly outpace the capacity of legacy tools.

Stage 2: Do it yourself

Many teams progress to adopting a host of specialized tools, open source technologies, and net-new development, then patching together the results. These technologies are more tailored for modern environments, but they often result in data silos and blind spots.

Stage 3: Adopt a platform approach

The final stage is to take a platform approach to observability that is purpose-built to deal with cloud-native complexity. A platform approach that uses automation and AI can ingest data from just about any source and integrate it into a single view that provides automatic answers to issues (see Figure 3).




Figure 3: As complexity becomes more of a challenge, organizations often progress through stages to the ideal cloud-native observability solution.

The first step to cloud-native observability is to consider the capabilities that define a best-in-class observability platform. Most important is to evaluate a technology that deals most effectively with the complexity wall. This ideal platform has various attributes (see Figure 4).





Figure 4: When considering the right platform in support of observability, there are a few key attributes you need to look for.

CHAPTER 4

Essential attributes of a modern cloud observability platform



When evaluating a modern observability platform, IT buyers should look for the following key capabilities:

End-to-end view

“Shift-left” visibility into applications in development and “shift-right” visibility to identify issues with live applications in production. Having an end-to-end view of the entire software delivery lifecycle also gauges how applications affect user experience.

Open and extensible platform

Because organizations operate a range of open source and cloud-native technologies, an observability platform should be extensible and be able to ingest telemetry from all sources. A solution should use AI and automation to deliver precise answers about the entire ecosystem.

Continuous automation

The dynamic nature of modern cloud environments puts new pressure on observability. As a result, a modern observability platform is always up to date and continuously discovers and maps your environment in real time.

Customized experience

Because organizations have teams with different priorities, an observability solution should provide a customized experience so each team has the data they need. A common source of truth enables cross-team collaboration.

Deterministic AI

Because telemetry data volumes have exploded, an observability platform should use deterministic, cause-based AI to deliver answers, not just more data on dashboards. Observability should detect issues in real time and offer immediate data-driven answers.

Cloud scale

With the availability of hyperscale platforms and serverless services, a modern observability solution should enable organizations to break through the cloud observability wall and ensure that all applications and services are observable.

With a modern approach to observability, IT pros can overcome cloud complexity and accelerate innovation that delivers the most compelling user experiences. Instead of being mired in manual tasks and time-consuming fire-fighting, IT teams can automatically ensure applications are available, reliable, and fast across every channel, including mobile, web, Internet of Things, and application programming interfaces. As a result, IT teams can be free to turn their attention to strategic, revenue-generating activities rather than tactical problem solving.

A modern observability platform also empowers development teams by identifying issues in testing and development environments and addressing these issues automatically.

Because Dynatrace gathers data in real time in multicloud environments, businesses get complete visibility of their entire application ecosystem. The combination of AI and automation ensures complete coverage at scale in even the most rapidly changing environments.

With code-level root-cause analysis, cloud practitioners can get to the heart of issues before they undermine system performance or digital experience. Consider the December 2021 Amazon Web Services outage, which affected a variety of streaming services in the Northeast U.S. region. Cloud outages can have ripple effects and require precise, swift root-cause analysis to restore customers' services.

With the ability to contextualize interdependencies throughout the cloud stack and causal AI to provide actionable answers to problems in real time, IT practitioners can address issues before they affect system performance rather than after the fact.

As a result, multicloud observability can drive real business value, enabling cloud practitioners to focus on software innovation and business growth.

Unparalleled value across stakeholders



DEVELOPMENT

Innovate faster with
higher quality

4x

greater innovation
throughput - AGII

42K

more hours per year for
innovation - TIAA

0

major errors in post-production
code - Rack Room Shoes



OPERATIONS

Operate more
efficiently

60%

greater operational
efficiency - SAP

99%

fewer support tickets
- Kroger

80%

faster issue resolution -
Porsche Informatik



BUSINESS

Drive better business
outcomes consistently

20%

higher cart value -
Mitchells & Butlers

65%

lower customer
churn - Vitality

32%

increase in
conversions - BT

Dynatrace makes essential government services always on

This article features the following takeaways:

- The Minnesota Department of Employment and Economic Development (DEED) successfully navigated an 850% year-over-year increase in unemployment benefits applications in 2020—and without disruption in providing essential unemployment benefits to Minnesota workers.
- DEED developed new DevOps (or development or operations teams) best practices that enabled it to collaboratively resolve issues with its benefits application.
- By automatically identifying anomalies, pinpointing the root cause of issues, and alerting when problems with true business impact occur, DEED could reduce manual effort and focus on driving innovation, without risking the performance of existing services.
- Modern observability thus enables organizations with multicloud environments to sense, think, then act automatically on the root causes of issues in their environments.

Government services provide an essential safety net for citizens. Government agencies—and their IT infrastructure—can get strained in times of crisis.

As COVID-19 made its way across the globe, societies and economies were disrupted, with nearly 60 million U.S. workers filing for unemployment benefits in 2020.

The Minnesota Department of Employment and Economic Development (DEED), which manages unemployment benefits for the state's citizens, saw exponential increases in filings for unemployment.

In 2019, for example, DEED processed some 80,000 applications. By August 2020, it had processed 709,000 applications—an 850% year-over-year increase with a total of 26 million logins.

"We never thought we would hit numbers like this," said Kailey Smith, application architect for the Minnesota IT Services department.



Unemployment [benefits are] an essential service. It's our responsibility to keep the app running and performing well."

Kailey Smith, application architect,
Minnesota IT Services

These unprecedented loads in activity taxed DEED's Java-based web unemployment app, which Smith says has a "large and complex code base."

The staggering volume threatened to crash the application, which could have had catastrophic consequences for the state's unemployed workers.

"Unemployment [benefits are] an essential service," Smith says. "Regardless of challenges, it's our responsibility to keep the app running and performing well."

Cloud observability is key to application performance

Application performance is critical for organizations like DEED. But the unemployment app's cloud-based architecture—which is distributed, complex, and sensitive to peaks in traffic —places new burdens on teams like Smith's to ensure application uptime. That's because organizations now rely so heavily on public clouds such as AWS, Azure, and Google, as well as hybrid clouds.

According to recent Dynatrace research, 99% of organizations use cloud computing.

"Organizations today are under the gun to keep their IT infrastructure and applications up and running, and to minimize the amount of downtime they suffer," wrote Alex Woode in Datanami.

"This has always been the goal, but it's more difficult to achieve today due to modern architectures (microservices, containerization, hybrid-cloud deployments, growth of the edge) as well as the manner in which they are developed."



Dynatrace was our go-to source of information."

Kailey Smith, application architect,
Minnesota IT Services

Single source of truth alleviates crisis

DEED users developed customized tiles to monitor the overall health of the application and surrounding environment, including application servers, web servers, FileNet servers as well as batch services and databases to the number of live users. Users can also access a dashboard tile that indicates all open problems regardless of where they occur in the system.

The health tiles indicate the status of servers, applications, and databases with a green or red circle. "This gave us an immediate visual indication of problems across our application," Smith says.

"Dynatrace tells you how many users are affected, what services were affected, what metric anomalies were detected. And in many cases, the Dynatrace AI is able to identify the root cause of the problem," Smith says.

As a result, DEED used the data from Dynatrace as a single source of truth to identify root causes of application problems and resolve them. DEED's implementation of observability helped it to sense, think, then act on problems within its environment to keep its systems up and functioning.

"Dynatrace was our go-to source of information," Smith said.



Dynatrace drove cross-team collaboration."

Kailey Smith, application architect,
Minnesota IT Services

Cloud observability brings coss-team collaboration

According to recent Dynatrace research, 70% of DevOps teams say that application complexity is a barrier to DevOps success.

In the process of gaining visibility into application issues, DEED forged cross-team collaboration among infrastructure, development, middleware, and database teams. The various dashboards enabled teams to work together to identify issues that might affect application performance.

“Dynatrace drove cross-team collaboration,” Smith recalls. “We had to ... quickly come to a consensus on the changes needed. We had to ensure we were all on the same page and looking at the same data. The Dynatrace dashboard was the solution.”

With a single source for data, teams could collaborate to identify the true source of issues when their key metrics diverged.

“When team members observed different metrics, it led to conversations. Teams were all looking at the same data and working together to identify solutions,” Smith says.

“The level of cross-team collaboration was invaluable to support the application,” Smith says.


A single source of truth identified system anomalies and enabled workers to collaborate even with stay-at-home orders.

While traditionally, teams would have met in a conference room to hash out issues, they instead met virtually and workshopped the data, which was only possible with a unified platform.


“War rooms were just not an option,” Smith recalls. “We had daily online meetings to coordinate our efforts ... to handle the high volume and troubleshoot problems ... and to share information.”

Ultimately, Dynatrace enabled DEED to provide essential services without skipping a beat.

“In the most turbulent times, the state of Minnesota could rely on Dynatrace,” Smith says. “It gave us a single view ... with precise answers to help us ensure our system was always available and performing well.”



of respondents say a unified platform that integrates toolchains is critical to scale DevOps efforts



of DevOps teams say that application complexity is a barrier to DevOps success

Source: “2021 Global DevOps Report,” Dynatrace Inc.

The Dynatrace difference

Go beyond alerts and correlations with Dynatrace's deterministic AIOps platform. By embracing and building on open standards such as OpenTelemetry, Dynatrace provides immediate, actionable identification of root-cause issues. With real-time causation-based analysis, Dynatrace does much more than simply tell your team that something might be wrong. It frees engineers to focus on what really matters — productive, value-creating work, not time-consuming reactive work. In turn, Dynatrace software intelligence helps teams build a culture of collaboration and shared tooling across the IT organization.

Learn more by requesting a [demo of Dynatrace in action](#).



About Dynatrace

Dynatrace provides software intelligence to simplify cloud complexity and accelerate digital transformation. With automatic and intelligent observability at scale, our all-in-one platform delivers precise answers about the performance and security of applications, the underlying infrastructure, and the experience of all users to enable organizations to innovate faster, collaborate more efficiently, and deliver more value with dramatically less effort. That's why many of the world's largest enterprises trust Dynatrace® to modernize and automate cloud operations, release better software faster, and deliver unrivalled digital experiences.