dynatrace

Seven criteria for unified observability and security

Summary

The world's leading companies and public-sector organizations are building the digital services we rely on using hybrid-cloud and multicloud applications and infrastructure. While agile, these environments are increasingly complex and challenging to manage.

IT, development, and security teams typically rely on multiple monitoring tools, data repositories, and time-consuming manual processes to tame this complexity. However, this reactive approach is unsuitable for rapidly changing cloud environments and evolving security threats because it's not scalable, can expose organizations to security risks, and steals time from already overstretched teams.

Observability and security in modern cloud environments have become mandatory. They should drive the optimization of those ecosystems to ensure flawless and secure digital services that scale. They should also create *order from the chaos* by enabling teams to have better control of their technology ecosystems.

Organizations should seek a unified observability and security platform rather than a collection of disparate tools. The platform they select needs to leverage the full value of the data in their clouds to deliver precise answers and extensive, intelligent automation of workloads and processes.

Unifying observability and security analytics and delivering intelligent automation will enable organizations to reduce toil and accelerate innovation—enhancing users' and customers' experiences and creating a better working environment for development, security, and IT professionals.

The essential criteria for unified observability and security are the focus of this paper.

What's inside

The rise of multicloud and cloud-native computing

Digital transformation is accelerating. To keep pace, organizations have shifted from on-premises to hybrid and multicloud environments and microservices-based, Kubernetes-orchestrated, cloud-native architectures.

Regardless of the stage where organizations are in their transformation journey, cloud computing is now integral to the services they deliver—and the ones the world depends on—for banking, retail, healthcare, government services, workplace productivity, and home entertainment.

The shift to hybrid, multicloud, and cloud-native architectures yields undeniable benefits, from greater and faster innovation to increased operational efficiencies. However, the dynamic nature and scale of these architectures and the explosion of data they produce make them too complex to manage with static dashboards, alerts, and manual troubleshooting.

The landscape of monitoring and analytics tools has also expanded. Indeed, for each new solution they adopt, IT, development, and security teams often implement another monitoring tool or take a "do it yourself" (DIY) approach to bring insights together. Fragmented or DIY-type tools and techniques result in disparate data analytics, increased manual effort, and misaligned teams.

The hours teams spend managing or stitching together tools and chasing problems are stolen from innovation and reduce operational efficiency. In addition, fragmented tools and DIY approaches make it difficult to meet data governance and privacy requirements, which highly regulated industries, such as financial services, manufacturing, and healthcare, increasingly require.

Multicloud and cloud-native computing on the rise

90% of large organizations say digital transformation has accelerated in the last 12 months.

— Dynatrace CIO research, Jan 2023

Through 2026, cloud **spending will grow at a CAGR of 15.8%** and is forecast to **exceed $1 trillion** in spending worldwide.

— 2022 Gartner®, The Future of Cloud Computing in 2027

By 2026, more than **90% of enterprises** will extend their capabilities to multicloud environments, which is **an increase from 76%** in 2020.

— 2023 Gartner, A CTO's Guide to Multicloud Computing

Looking forward, increased demand for digital innovation will lead to increased reliance on clouds, exacerbating these challenges for the teams tasked with managing them.

IT, development, and security teams need a new, more mature approach to monitor, manage, and maximize the value of their investments in cloud ecosystems. To tame the complexity and enormous amount of data emanating from their clouds, they need a modern observability and security platform that delivers predictable, trustworthy, and precise insights in real time. Those capabilities are essential to teams' ability to automate processes and workflows to help enhance and secure end users' experiences and reduce manual workloads, freeing more time for innovation.

Not all observability platforms, however, are equal. It's only possible to achieve the meaningful insights teams need to manage modern clouds when platforms meet specific criteria.

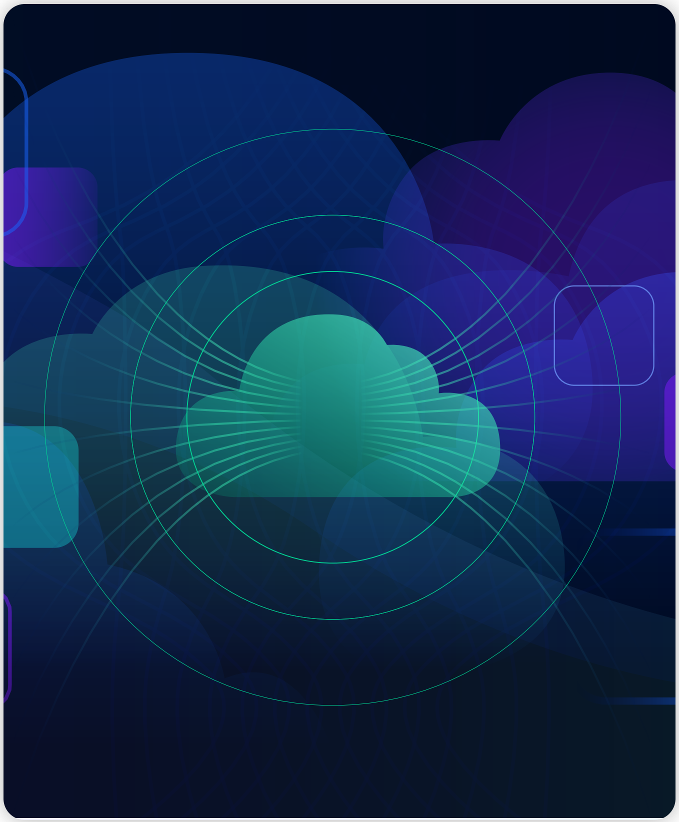

Cloud computing challenges

71% of CIOs say the explosion of data produced by cloud-native technology stacks is beyond human ability to manage.

— Dynatrace CIO Research, Sept 2022

59% of CIOs say their teams will become overwhelmed by cloud complexity without a more automated approach to ITOps.

— Dynatrace CIO Research, Sept 2022

On average, **teams use 10 monitoring tools** across their stacks but have observability into just 9% of their environment.

— Dynatrace CIO Research, Sept 2022

81% of IT leaders say that using Kubernetes has made their infrastructure more dynamic and challenging to manage.

— Dynatrace CIO research, Jan 2023

Seven criteria for unified observability and security

The solution that organizations select for their hybrid, multicloud, and cloud-native observability and security analytics and automation should deliver on each of these criteria:

1. One unified analytics and automation platform for observability, security, and business data
2. Ability to capture and process all data from all sources while retaining topological and dependency mapping context
3. Ability to deliver cost-effective and scalable data analytics
4. AI at the platform's core, combining multiple techniques—predictive, causal, and generative AI
5. Ability to deliver trustworthy automation of business, development, security, and operations workflows
6. Ability to detect and mitigate security vulnerabilities in runtime environments, block attacks in real time, and conduct data-driven security analytics
7. Extensibility to use observability, security, and business data to power custom digital business use cases

What is observability?

Observability is a term that describes the ability to collect and analyze telemetry data from endpoints and services in hybrid and multicloud environments. This telemetry data, coupled with specialized analytics capabilities, enables an observability solution to measure a system's health, performance, and behavior.

Observability solutions enable teams to gain insight into their on-premises, hybrid, and multicloud systems and analyze how adjustments to or within these technologies—even the most minor changes in microservices, code, or newly discovered security vulnerabilities—affect end-user experiences and business performance.

Observability begins with three core pillars:

- **Logs:** A record of what's happening within the software.
- **Metrics:** Counts, measures, and calculations regarding application performance and resource utilization.
- **Traces:** The path a transaction takes throughout a system's applications, services, and infrastructure from one node to another.

In addition to these classic pillars, observability data includes events and other signals generated by the system's components and services.

CRITERION 1

One unified analytics and automation platform for observability, security, and business data

Many organizations are converging their observability and security practices to help innovate quickly and without introducing unacceptable risks. They are adopting processes to shift security left into development and right into runtime or active production environments.

Organizations can further accelerate secure innovation and improve collaboration across their development, security, and IT operations teams by reducing their use of fragmented monitoring tools and point solutions. Instead, they should arm these teams with a unified platform that provides data-backed insights and drives trustworthy automation from the observability and security data emanating from their hybrid and multicloud ecosystems.

Organizations should look for a unified platform that automatically and continuously ingests, stores, and processes data in real time and at scale from all major cloud platforms, including Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and IBM Red Hat. The platform must also integrate with the other solutions across each organization's technology ecosystem to enable extensive automation of custom, or organization-specific development, security, and operations tasks, without excessive manual intervention.

What is shift left and shift right?

Shift left is the practice of moving software quality testing, security, and performance evaluation to the early stages of development, even before developers write any code. This practice helps teams anticipate changes that could affect their software's performance so they can proactively resolve those issues and ensure that software meets customer requirements.

Shift right is the practice of performing quality testing, security, and performance evaluation in production under real-world conditions. This practice ensures that applications running in production are resilient and reliable and can withstand real-user load without exposure to newly discovered vulnerabilities. The goal is to detect and remediate issues that would be difficult to anticipate in development environments.

> Only **50% of CISOs** are fully confident that the software delivered by development teams has been completely tested for vulnerabilities before going live in production environments.
>
> — Dynatrace CISO Research, April 2023

CRITERION 2

Capture and process all data from all sources while retaining topological and dependency mapping context

The data originating from cloud environments are a rich source for insights that can drive automation to optimize user experiences and increase operational efficiency. However, organizations often struggle to operationalize their data because it is trapped in silos across the storage repositories and monitoring tools different teams use. This siloed data also lacks context reflecting the relationships and dependencies across hybrid and multicloud ecosystems. Without this context, it's difficult to distinguish between the symptoms and cause of an issue, resulting in time wasted chasing false positives, duplicate alerts, or low-priority issues.

Getting the precise answers and context needed to drive real-time analytics and automation requires a platform that can automatically capture and process all data from cloud environments, regardless of source or format. This includes all logs, metrics, and traces, as well as data from the latest open source standards, including OpenTelemetry, and business events, such as conversions, cart abandonment, and sentiment analysis. Given the dynamic nature of modern cloud technology stacks with microservices and containers that come and go in seconds, the platform needs to be informed by a precise visualization, or map, that discovers and tracks the topology of each organization's technology ecosystem as it changes in real time.

Optimizing cloud-native environments requires observability in context

"Sprawling and dynamic cloud-native and multicloud environments are an ecosystem of various technologies and services, and the composition changes by the second. This paradigm makes it critical for organizations to acquire a platform with advanced AI, analytics, and automation capabilities. The platform must be able to ingest all observability, security, and business data, put it in an accurate context in real time, and facilitate access to data-backed insights when needed."

— Stephen Elliot, group vice president, I&O, Cloud Operations, and DevOps, IDC

CRITERION 3

Cost-effective and scalable data analytics

The cost constraints of conventional storage models often force organizations to be selective about which observability, security, and business data they keep, discard, or move to “cold storage,” where they can no longer access or query it easily. To extract value from this data using traditional solutions, ITOps, DevOps, SRE, and security teams rely on time-consuming manual processes. They try to structure their datasets to reflect the questions they expect to ask in the future, and they attempt to rehydrate or resurrect data that’s locked in cold storage. These procedures are time-consuming and don’t adapt quickly to rapidly changing application and cloud environments or evolving security threats.

Teams should not have to decide which data to keep or what to do with it before storing it. They should be empowered with a platform that enables them to ask any question about any data anytime, without rehydration.

Teams can achieve cost-effective and scalable data analytics only with a unified platform that has a data storage technology built specifically for analyzing observability, security, and business events data at a massive scale. A causational data lakehouse with a massively parallel processing (MPP) analytics engine delivers this by harnessing and unifying data from even the most complex multicloud architectures.

What is a data lakehouse?

Data warehouses have existed for many years and were created to store structured data, often categorized as quantitative data organized into a specific format to access, manage, update, and search easily using software.

Data lakes are a more recent concept created to store structured and unstructured data, with the latter referring to data such as text, social media posts, and logs that do not have a predefined model. Data lakes have advantages, including storing large amounts of data in raw or original formats. They also have some issues. For example, people might dump all their data into the data lake, creating problems ranging from cost management to governance and privacy.

A **data lakehouse** is a data management architecture that combines the contextual querying capabilities of a data warehouse with a data lake’s flexibility. People can query and analyze data within the data lakehouse without structuring it or worrying about which data to maintain as “hot” and accessible for analytics or “cold” for long-term archives.

CRITERION 4

Platform with AI at its core that combines multiple techniques — predictive, causal, and generative AI

Artificial intelligence (AI) should help teams dramatically reduce manual work. But how can teams trust their AI is drawing the correct conclusions, making the right decisions, and implementing the right automation?

It's important to remember that not all AI is created equal, and some types are better suited to certain tasks than others.

Traditional approaches to AI for observability and security use cases often require additional tools, which teams add or bolt onto their analytics and automation platform. These tools often use machine learning (ML) models that correlate data to produce probabilistic insights and alerts, which they display on dashboards. This approach can be difficult to scale because it relies heavily on human intervention. In addition to the challenge of managing "tool sprawl," teams adhering to this approach must also train their correlation-based AI models. This training is the manual process of feeding the model or algorithm with data, testing the outcome, and adjusting to help ensure the model delivers the desired outcome. Moreover, correlation-based AI doesn't provide continuous, real-time insight into an issue's precise root cause. Instead, its insights reflect patterns from historical events, which may or may not reflect future outcomes.

Organizations should select a platform with AI at its core that combines multiple AI techniques, with each one excelling in specific capabilities:

1. **Forecasting —** using **predictive AI** models that recommend future actions based on data from the past, including sales and customer experience trends, seasonality, cloud application health, and other metrics.
2. **Determining —** using **causal AI** to deliver fact-based, deterministic, and precise answers and intelligent automation based on analyzing dependencies across large sets of observability, security, and business data while retaining an accurate context that reflects each data point's source.
3. **Creating —** using **generative AI** that is automatically fueled by causal and predictive AI insights and humans' natural language prompts to deliver precise recommendations for solving specific tasks in the context of each organization's unique environment and situation.

Combining AI techniques is more effective

Gartner released its 2023 Hype Cycle™ for Artificial Intelligence report, which states, "Combining AI techniques is much more effective than relying only on heuristics or a fully data-driven approach."

The report describes **causal AI** as an emerging technology that "identifies and utilizes cause-and-effect relationships to go beyond correlation-based predictive models and toward AI systems that can prescribe actions more effectively and act more autonomously. It includes different techniques, such as causal graphs and simulation, that help uncover causal relationships to improve decision making."

The report also states, "AI's ultimate value comes from helping people take better actions. Machine learning (ML) makes predictions based on statistical relationships (correlations), regardless of whether these are causal. This approach is fine for prediction, but predicting an outcome is not the same as understanding what causes it and how to improve it. Causal AI is crucial when we need to be more prescriptive to determine the best actions to influence specific outcomes. Causal AI techniques help make AI more autonomous, explainable, robust, and efficient."

> **"Generative AI** is already proving useful in broadening the accessibility of operations insights to new personas and speeding workflows for users of observability solutions. However, when combined with other forms of AI, generative AI has the potential for additional notable impact. For instance, leveraging other forms of AI to feed generative AI with more than just user inputs can deliver more value for customers and help maximize the value of generative AI for business, development, security, and operations use cases."
>
> — **Nancy Gohring,** Research Director for Enterprise System Management, Observability and AIOps, IDC

CRITERION 5

Trustworthy automation of business, development, security, and operations workflows

Delivering precise and trustworthy answers from unified observability, security, and business data is only part of the journey. Ensuring flawless and secure digital experiences amid modern cloud ecosystems' complexity and scale requires extensive and intelligent automation of error-prone manual tasks. This automation should include continuous discovery and instrumentation of applications and infrastructure, proactive vulnerability and anomaly detection, and optimization across the software lifecycle.

Organizations need a platform that uses causal AI to create and extend trustworthy automation across their ecosystems and processes, from software development to cloud operations and application security. The platform should also integrate with their business, development, security, and operations toolsets.

A platform with these capabilities empowers teams to automate workflows with confidence. For example, teams can automate workflows that react to changing user experiences or seasonality. They can also automate workflows to turn off features for security or quality reasons or to enhance software orchestration to reflect myriad external factors, from weather forecasts to energy consumption and supply chain delays. As a result, organizations can free their teams from many manual tasks and enable them to focus on driving innovation, growth, and customer impact.

AI-driven automation tames cloud complexity

"Cloud-native and multicloud technology adoption is accelerating, driving platform fragmentation and an explosion in observability and security data variety and volume. AI/ML and automation are increasingly useful in taming this complexity, enabling organizations to find the right answers and automate fixes."

— James Governor, Co-founder, RedMonk

CRITERION 6

Detect and mitigate security vulnerabilities in runtime environments, block attacks in real time, and conduct data-driven security analytics

The complexity of cloud-native applications, the growing use of open source libraries, and higher frequency code releases increase organizations' exposure to security risks. To mitigate these risks, organizations need to institute DevSecOps practices and embrace shift-left and shift-right methodologies. These steps help establish a cross-team culture where security is a shared responsibility.

Teams also need a platform that provides end-to-end visibility into their security posture to detect vulnerabilities and facilitate automated, continuous, and real-time remediation across the software lifecycle. With an automatic, data-driven approach to security, teams can immediately understand whether a detected vulnerability has been exploited and trace the impacts of that exploit to determine which data and applications were affected.

Detecting and mitigating vulnerabilities and blocking attacks in real time requires a platform with the following capabilities:

Runtime vulnerability analysis — Providing teams with a clear understanding of the most critical vulnerabilities to address and eliminating the time they spend chasing false positives.

Runtime application protection — Continuously detecting and blocking critical Open Web Application Security Project (OWASP) threats, including SQL injections and command injections that target critical vulnerabilities.

Security automation — Orchestrating a response to security incidents or detected vulnerabilities across tools and teams through precise insights and automated workflows that act in real time.

Advanced security analytics — Combining observability context with security events to enable teams to swiftly detect and respond to threats by uncovering advanced persitent threats (APTs) and analyzing tactics, techniques, and procedures (TTPs) for actionable insights that support proactive defense.

Support for shift left and shift right — Integrating data from security tools across the entire software lifecycle to provide runtime observability context that enables better prioritization of vulnerabilities and targeted insights for faster remediation.

CRITERION 7

Extensibility to use observability, security, and business data to power custom digital business use cases

Every organization has unique requirements and technology stacks, resulting in many custom analytics and automation use cases. To address these unique needs and use cases, teams need an extensible platform with an easy-to-use, low-code approach for creating custom, compliant, data-driven, and AI-powered apps and automations. To be effective, the platform must provide automatic scalability, runtime application security, safe connections and integrations across hybrid and multicloud ecosystems, and full lifecycle support, including security and quality certifications.

Armed with these capabilities, organizations will be poised to unlock the wealth of insights available in the explosive amounts of observability, security, and business data their modern cloud ecosystems generate. They will also be able to extend precise answers and intelligent automation to numerous business, development, security, and operations use cases, reflecting their unique technology ecosystems, addressing their specific needs, and empowering more people across the organization to make data-backed decisions.

The path to "cloud done right"

While each organization's journey with cloud computing are different, there is no doubt that modern hybrid and multicloud environments and cloud-native architectures are playing an increasingly important role in powering digital transformation. A unified approach to observability and security has become mandatory for organizations seeking to tame cloud complexity and accelerate digital transformation. Identifying a unified observability and security platform that delivers on each of the seven criteria featured in this paper will help ensure the long-term success of these transformation initiatives. Adhering to these criteria will enable organizations to ensure they are on the path to *cloud done right*.

Analyst disclaimers

Gartner, The Future of Cloud Computing in 2027: From Technology to Business Innovation, David Smith, Dennis Smith, 22 October 2022

Gartner, A CTO's Guide to Multicloud Computing, David Smith, 23 March 2023

Gartner, Hype Cycle for Artificial Intelligence, 2023, Afraz Jaffri, 19 July 2023

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the US and internationally, and HYPE CYCLE is a registered trademark of Gartner, Inc. and/or its affiliates and are used herein with permission. All rights reserved.

About Dynatrace

Dynatrace (NYSE: DT) exists to make the world's software work perfectly. Our unified platform combines broad and deep observability and continuous runtime application security with the most advanced AIOps to provide answers and intelligent automation from data at enormous scale. This enables innovators to modernize and automate cloud operations, deliver software faster and more securely, and ensure flawless digital experiences. That's why the world's largest organizations trust Dynatrace® to accelerate digital transformation.

Curious to see how you can simplify your cloud and maximize the impact of your digital teams? Let us show you. Sign up for a free 15-day Dynatrace trial.

dynatrace

dynatrace.com blog **@dynatrace**