



An observability platform buyers' guide: Five key criteria

Buying an observability platform requires careful consideration about platform capabilities. In this observability buyers' guide, we review five key buying criteria for successful observability in today's cloud environments.



What's inside

| | |
|---|----|
| Introduction | 3 |
| Choosing a unified observability platform | 4 |
| Gathering high-quality trusted data | 5 |
| Developing multi-tiered AI and automation | 6 |
| Establishing integrated application security | 7 |
| Developing analytics for IT and business decision making | 9 |
| Evaluating a modern observability platform for an organization's needs | 10 |

Introduction

Every enterprise IT organization strives to meet the goals of performance, reliability, and security. But as modern IT environments become more complex with containerized microservices, multicloud architectures, and open source software, organizations need a way to integrate, manage, and analyze it all automatically.

Observability tools have emerged to answer this call, deriving answers in various ways from [logs, metrics, and traces](#).

Buying an observability platform requires careful consideration about platform capabilities to select the right technology for today's evolving digital ecosystems. As application architectures become more complex, organizations are drowning in data that needs to be processed. Because the digital environments that support applications continue to advance, the observability technologies that provide answers must also evolve.

Why traditional observability needs to evolve

Traditionally, IT teams collected and analyzed observability data—metrics, logs, and traces—in silos from different tools on static dashboards. With cloud migration resulting in an [explosion of data](#) from cloud platforms, container environments, and microservices, this approach is no longer effective. It's impossible for a human to sift through petabytes of data to consistently make fast and accurate data-driven decisions.

Organizations now require an integrated and contextual approach to observability

Data at the scale of modern cloud environments is where artificial intelligence becomes crucial. AI-powered capabilities are essential for processing and identifying relevant data, enabling organizations to have a complete and accurate picture of their systems. Without AI, organizations become overwhelmed working with fragmented observability data instead of obtaining the comprehensive insights they need for decision making.

With architecture changes come observability challenges

Selecting the right observability solution for an organization's needs can be challenging. There are various obstacles to achieving end-to-end observability, such as tool sprawl, siloed information, lack of actionable insight from data, and complex technology setup. Choosing the right observability solution is vital; a mismatch can increase complexity, lose time, and boost costs for organizations.

Considerations for choosing an observability solution

Modern observability goes beyond traditional logs, metrics, and traces. It takes a holistic approach that includes security, user experience, and business impact. Buying an observability platform requires thoughtful consideration about a technology platform's capabilities. In what follows, we review five key considerations for successful observability in today's cloud environments:

1. A unified approach to collecting, processing, and analyzing data
2. An approach that enables high-quality, trusted data
3. An AI-enabled system
4. An integrated observability and application security approach
5. Analytics that enable IT and business decision making

Let's dive into the world of observability and uncover the essential factors to consider when choosing an observability solution.

BUYING CONSIDERATION 1

Choosing a unified observability platform

Observability aims to understand digital ecosystems and provide answers so users can make informed decisions. The quality of answers, however, is only as good as the data inputs. Unified data sets enable informed decision making, particularly in areas like application architecture, user experience (UX), and key business decisions where context provides richer insights.

Automated discovery

Automated topology discovery and mapping serve as initial steps in comprehending an application. These capabilities are particularly valuable in systems with ephemeral components because mapping relationships can be complex when compute instances are spinning up and down.

Full-stack and end-to-end visibility

Full-stack visibility encompasses the vertical layers of system topology and application architecture. End-to-end visibility includes the horizontal relationships among all entities, from individual requests to code-level details. Since problems arising in one application layer may originate from issues in another layer, it is critical to understand the interdependencies among all the entities.

Hybrid, multicloud, and polyglot architectures

Many teams have hybrid on-premises and multicloud architectures that support multilanguage polyglot microservices. As a result, an observability solution must support a range of technologies—from container platforms to large cloud service providers—to enable a complete understanding of applications across layers and their dependencies.

Open and extensible

Organizations are increasingly adopting open source tools to solve specific problems, accelerate innovation, and reduce costs. An observability platform should integrate seamlessly with a range of open source solutions used throughout the software development lifecycle (SDLC). It should also provide a potent application programming interface for observability of all custom implementations.

User experience

An observability platform should also address issues from the user perspective. With deep observability of real-user activity, teams can resolve performance issues and innovate based on user experience findings. Beyond performance, observability data is linked to business-level metrics such as key performance indicators (KPIs), conversion goals, and revenue contribution, allowing companies to make better business decisions and optimize user experiences.

Integrated security

With the escalating sophistication of security threats, organizations are increasingly integrating application security throughout the SDLC. An integrated approach allows organizations to quickly resolve security breaches by understanding the problem's location, how to resolve the issue, and assessing the business impact.

BUYING CONSIDERATION 2

Gathering high-quality trusted data

Making informed decisions and building a mature automation practice relies on having access to the right type of data and answers. So, how can an observability platform help teams collect high-quality data?

Real-time data in context

Because of the dynamic nature of microservices, multicloud environments—and highly distributed systems—having access to real-time data has become critical. Rapid environment changes require quick access to answers and immediate visibility into the full context as an incident occurs. Furthermore, an observability solution must properly map real-time data to other relevant events and entities on a session-by-session basis for root-cause analysis.

On-demand querying

Conventional storage methods force organizations to choose which data to keep and which to have less immediate access to (where “hot storage” is data that is available in real time and “cold storage” is less immediately available). This choice significantly reduces data fidelity and availability for immediate querying. Real-time querying enables fast and actionable answers. Waiting for data sets to rehydrate from cold storage and get indexed can render results almost useless in critical and fast-acting scenarios. In addition, the ability to query on demand and perform exploratory analytics can lead to proactive and predictive monitoring, detecting issues before they escalate and become problems.

High cardinality and high dimensionality

Observability solutions should also enable teams to handle high-cardinality and high-dimensionality data sets. High cardinality refers to data fields with many possible unique values, such as the ability to query a trace ID or a user's email address. High dimensionality involves a data set with a very high number of variables, such as gene sequencing. Both types of data can cause the cost and complexity to quickly escalate. As a result, most observability platforms have restrictions on querying and how users can manipulate collected data. However, empowering users to manipulate and query data freely enables them to obtain specific and targeted answers based on their needs.

Indexless data storage and management

A conventional database arranges data using some form of indexing based on a predefined schema. Indexing can improve database performance, but also requires schema management and additional storage for each index. A modern observability solution should give teams access to data in real time without having to index it first, which dramatically increases database speed and efficiency. Such storage should likewise require no schema and lose no data. Indexless, schemaless, and lossless data storage enables teams from across the organization to query data with its full details in place without suffering the expense, inefficiency, and data loss of a conventional database.

BUYING CONSIDERATION 3

Developing multi-tiered AI and automation

Because of the increasing volume, velocity, and variety of data in today's microservices-based, multicloud environments, teams need AI that delivers certainty.

An observability solution that uses a multi-tiered approach to AI for IT operations (AIOps) can deliver less noise and more precise answers. By combining causal AI for root-cause certainty, predictive AI to determine the net effects, and generative AI to accelerate reliable responses, an observability platform can help teams further automate the SDLC and speed DevOps maturity.

Causal AI based on fault-tree analysis

Fault-tree analysis determines system-level failures based on component-level failures, unlike correlation-based machine learning, which calculates probabilities based on statistics. An observability platform that also draws on supporting data—such as relationships, dependencies, and other contextual information—can form the foundation for reliable recommendations from generative AI technologies.

Predictive analytics

AI can use data not only to understand problems that have already occurred but also to predict outcomes that are likely to occur based on past and current conditions. Predictive analysis identifies issues before they become costly problems. Analysis using predictive AI learns from patterns in your environment and delivers insight about potential risks. Some examples include forecasting disk utilization, predicting cloud application health, and prescriptive recommendations for potential customer experience issues.

LLM-based generative AI

With the rise in popularity of generative AI tools like ChatGPT and Bard across the globe, it should come as no surprise that the next step in automation is using large language models (LLMs).

However, generative AI can't deliver accurate and reliable results if the LLMs have to rely on imprecise correlations that require manual verification. An observability platform that uses causal AI and real-time entity-relationship data provides full context that enables generative AI to produce precise answers unique to the environment.

Automation

With precise answers, teams can automate more confidently. Automation is the key to success in today's ever-changing environments. From continuous dependency mapping to system configurations, automation speeds DevOps processes and responses to customer-facing issues.

Automating can go beyond simple maintenance to automated problem remediation. Auto-remediation using AI is becoming more common, solving real problems as they arise. Automating responses, however, requires that AI provides deterministic answers to ensure precision and that teams can explain the steps involved. By automating processes and responses, teams can free up time and resources they can spend on innovation instead of tedious manual tasks.


BUYING CONSIDERATION 4

Establishing integrated application security

Because application architectures have changed, so have the ways organizations must secure applications. Cloud-based and ephemeral computing require that organizations re-evaluate the best application security strategies at each stage of the SDLC. Here are some of the best strategies to improve overall application security posture before, during, and after an attack.

A unified observability platform should provide a layered approach to security, where organizations can take preemptive action before an attack, take action during an attack, and take post-mortem action to address vulnerabilities.

A platform should enable teams to integrate security into each stage of the software development lifecycle. This means “[shifting left](#)” as code is developed and “shifting right” to identify vulnerabilities that may have made it into a live application. During a security incident, teams should be able to have access to real-time data, to block ongoing attacks, and to provide rapid incident response. Finally, after an attack, teams should be able to apply application forensics to understand what has happened.



Before a security incident

The best time to take preventive security action is prior to an attack. Because systems are now so distributed and can potentially expose more risk, it is more important than ever to integrate application security into each stage of the SDLC. This proactive approach requires shifting left—mitigating known vulnerabilities while code is being developed—and shifting right—monitoring for new vulnerabilities while the application is live in runtime.

Bad actors can exploit previously unknown gaps in security. This is why integrated, observability-based security is so important, gathering full-stack, end-to-end data from every application, tool, and computing environment. An observability platform should also provide context to other sources of security intelligence, such as scans of cloud infrastructure and production environments.

During a security incident


Once a security incident begins to unfold, teams need real-time data, the ability to block any ongoing exploitation, and fast incident response.

- **Real-time data.** An observability platform that examines real-time data helps teams immediately identify attack patterns in the moment, not after the fact. With real-time topology mapping and intelligence about system states, locations, and dependencies, teams can instantly understand all affected entities and the systems they depend on.
- **Automatic blocking.** An observability platform should be able to detect SQL injection attacks, command injection attacks, and JDNI attacks such as Log4Shell or the H2 and MOVEit vulnerabilities. An observability solution shouldn't just rely on vulnerability databases, but rather identify and block such attacks automatically even if they are exploiting unknown weaknesses.
- **Fast incident response.** An observability-based approach to application security provides a unified data set that pinpoints in real time where the exploitation is happening.

An observability platform should use AI to automatically prioritize the most critical apps affected so teams can immediately kick off remediation workflows that prevent attackers from inflicting further damage.

After a security incident

In the event of a zero-day attack of a previously unknown vulnerability, teams need to be able to analyze historical data to ensure they're not affected and to prevent future attacks. Application forensics is a post-mortem analysis after an attack has happened, helping teams understand where an incident occurred, which entities were affected, and its effect on the business. However, many organizations can't conclusively determine their exploitation status because the logs don't contain enough information. An observability platform should provide immediate access to all historical observability and security data to investigate a potential compromise. As a result, teams can put controls in place to prevent future incidents from exploiting applications in the same way.





BUYING CONSIDERATION 5

Developing analytics for IT and business decision making

Ultimately, the end goal of observability ties back to decision making. Often, data sources from infrastructure, logs, applications, and security data may indicate that all systems appear to be working as expected. But user experience data tells a different story. This, in turn, affects an organization's ability to make effective decisions.

The richest answers come from unifying all this data in context. A problem that occurs in an application's infrastructure has the power to influence end-user experience and ultimately business outcomes.

One example of how observability helps understand application end users is with a simple e-commerce application. A typical user journey might look like this:

1. Users navigate to a webpage.
2. Users choose desired items.
3. Users place an item in the cart.
4. Users check out and purchase items in the cart.

But when the team makes a deployment change, it observes a steep decline in the number of checkouts.

With a comprehensive observability solution, teams can achieve the following:

- Understand the critical user journey path.
- Identify that there is a problem.
- Identify the root cause of the problem.
- Understand the business impact.

By understanding the user journey, teams can see that users are putting items in their carts, but they are unable to check out successfully. When looking at the checkout page performance, they can see that page loads are taking a significant amount of time, prompting impatient users to abandon the site. A comprehensive observability solution can immediately pinpoint what is causing the slowness and enable teams to quickly remediate it—or autoremediate it. It is even possible to understand the number of users affected and calculate the potential revenue lost from this problem. Understanding the business impact of a problem can help prioritize which problem to address first.

There are endless possibilities for integrating business data and key performance indicators with other observability data to get actionable insights. Combining these data sets aligns both IT and business priorities under one umbrella.

CONCLUSION

Evaluating a modern observability platform for an organization's needs

The features and capabilities outlined in this buyers' guide provide a blueprint for capabilities organizations should consider when choosing an observability solution. Use the chart provided in the next section to evaluate each vendor you consider according to the essential features covered in this ebook. If you need assistance, global system integrators (GSIs) such as Deloitte and Accenture can help you identify and deploy the observability solution that aligns most effectively with your organization's requirements.

Ultimately, organizations should seek a platform with unified observability capabilities. This enables organizations to analyze logs, metrics, traces, and other data sources in one location. A unified platform also enables teams to use both full-stack and end-to-end observability. Moreover, a unified approach enables organizations to apply observability to applications in development and production, preventing teams from introducing code-level vulnerabilities that can make their way to live applications. A truly unified approach can monitor at the deep application code level, but also observe issues in end-user experience, thus providing the most holistic view of issues throughout an organization's environment.

Observability vendor platform scorecard

Use the following evaluation criteria to assess and evaluate observability solutions. Rate each criterion on a scale of 1–5.

Legend

0 = N/A — The platform doesn't provide this capability

1 = Poor — The platform provides some of the capability, but it's incomplete or depends on multiple point solutions. Setup is time-consuming and hard to maintain, and results are not reliable for automation.

2 = Weak — The platform provides some of the capability but requires substantial manual setup and returns spotty or unreliable results. Platform also provides some data that can inform automation efforts, but results require manual validation.

3 = Adequate — The platform provides most of the capability but requires substantial manual work for setup and maintenance. Once set up, results can inform automation efforts with some manual validation.

4 = Good — The platform provides most of the capability with some manual configuration. Saves some time and enables some automation. Provides insights that help the IT organization.

5 = Excellent — The platform provides the full capability with little or no manual configuration. Provides reliable answers that enable automation. Saves time and provides unique insights that can benefit the entire organization.

| Criteria | Vendor A Grade (1-5) | Vendor B Grade (1-5) | Vendor C Grade (1-5) |
|--|-------------------------|-------------------------|-------------------------|
| <p>A unified platform</p> <p>A unified platform provides the following elements:</p> <ul style="list-style-type: none"> • Holistic view of data sources • Full-stack and end-to-end observability • Open and extensible | | | |
| <p>High-quality, trusted data</p> <p>High-quality, trusted data present the following characteristics:</p> <ul style="list-style-type: none"> • Real time, accurate, and mapped • In context, high cardinality, and high dimensionality • Schemaless and indexless • No cold storage or rehydration | | | |
| <p>An AI-enabled system</p> <p>An AI-enabled system delivers the following capabilities:</p> <ul style="list-style-type: none"> • Combination of predictive, causal, and generative AI • Transparent root-cause analysis based on dependency mapping • Automatic processes and problem remediation | | | |
| <p>Integrated application security</p> <p>Integrated application security presents the following traits:</p> <ul style="list-style-type: none"> • Real-time vulnerability management • Vulnerability mitigation during development and in runtime • Automatic issue prioritization • Forensic analysis | | | |
| <p>Analytics for IT and business decision making</p> <p>Analytics for IT and business decision making provide the following attributes:</p> <ul style="list-style-type: none"> • Integrates infrastructure, applications, security business events, user experience data for business insights in real time • Provides key performance indications • Enables shareable data for business users | | | |

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](#).

