Techstrong Group
a Futurum company

Observability in Modern Application Environments: Evaluating Unified, AI-Powered Approaches

Executive Summary

Organizations that once measured success by quarterly revenue now find their fate determined by milliseconds—the time it takes for an artificial intelligence (AI) model to respond, a cloud service to scale, or a security breach to propagate across their infrastructure.

The stakes have never been higher. A single AI hallucination can trigger customer churn. An undetected model drift can cascade into millions of dollars in costs or lost revenue. A fragmented monitoring approach can leave engineering teams firefighting instead of innovating. Traditional monitoring tools, built for predictable, monolithic applications, can crumble under the weight of today's dynamic, AI-driven architectures.

Yet, complexity doesn't have to be the enemy. Organizations that master unified, AI-powered observability can help transform these daunting challenges into competitive advantages. This can enable them to turn data chaos into strategic clarity, unpredictable AI costs into optimized investments, and reactive troubleshooting into proactive innovation.

The Mounting Challenges of Modern Application Environments

The Observability Challenge: Growth, Innovation, and Complexity

When Simple Becomes Impossible

Consider a typical enterprise today: hundreds of microservices spread across multiple clouds, dozens of AI models making autonomous decisions, and terabytes of data flowing through systems that change by the minute. Netflix modernized its monolithic application stack to a microservices architecture with over 700 components.[1] Retail giants like Walmart use advanced personalization systems that ingest, process, and respond to real-time data feeds on inventory, pricing, digital engagement, and supply chain disruptions. The company's systems constantly adapt its online and in-store offerings based on fluctuating data from millions of sources. But these aren't outliers—they're the new normal.

Organizations are drowning in 149 zettabytes of data created globally in 2024, with projections hitting 394 zettabytes by 2028.[2] The number of apps just keeps growing, with large organizations now averaging over

1,000 apps in use.[3] This explosion of app development and use comes at an increasing secondary cost: Engineering teams are spending more than half of their time firefighting issues rather than innovating,[4] and critical systems are failing while dashboards show everything is green.

Meanwhile, modern cloud workloads behave more like living ecosystems than predictable machines. Auto-scaling triggers cascade effects. AI agents influence each other in ways their creators didn't anticipate. A single configuration change can ripple through dozens of interconnected services. Traditional monitoring tools, designed for static environments, struggle to keep pace with this dynamic reality.

Most importantly, the app is now the organization. The cost of this complexity isn't just operational—it's existential. Every company now has something in common with companies like Airbnb or DoorDash, where the digital experience defines the brand. Application performance problems directly affect image, revenue, stock price, and market position.

The Hidden Costs of AI Innovation

AI promises transformation but may also deliver operational and financial surprises. Generative AI compute costs often exceed traditional cloud services by significant margins, with organizations sometimes experiencing unexpected budget impacts. For instance, one company reported that their bill surprisingly grew into thousands of dollars every month as their application scaled.[5]

In the AI era—especially with agentic AI—it's hard to escape the sense that infrastructure monitoring becomes critical for survival. GPU temperatures, memory saturation, and token consumption patterns directly affect both service availability and cost optimization. Organizations need to monitor everything from inference latency to energy consumption, especially as sustainability regulations tighten and energy costs soar.

The Reliability-Quality Dilemma

One interesting, and in fact rather frightening, aspect of all this is that when AI systems fail, they fail *creatively*. Unlike traditional software that breaks in a predictable way—like taking forever to respond or ceasing to function entirely—AI models hallucinate, drift, or produce confidently wrong answers that don't look like failure from a traditional standpoint, but nevertheless can destroy user trust at bewildering speed.

Large language models (LLMs) introduce non-deterministic behavior. This makes traditional testing inadequate, as the same prompts generate different responses and make it impossible to predict system behavior through conventional means. Model drift, meanwhile, occurs essentially silently—quality degrades gradually until suddenly the system produces unacceptable results.

Fragmented Data and Siloed Tooling

All of this development, expansion, and innovation results in enterprise IT metrics as a virtual Tower of Babel. Teams use one tool for infrastructure,

another for applications, and a third for logs, and then develop and maintain multiple custom dashboards for operational business metrics. Each element stores data in its own format and provides its own incomplete view.

Fragmentation like this can create dangerous blind spots. A recent analysis found that 78% of critical outages were caused by cross-domain network failures, 15% by internet data center issues, and the remaining by server issues.[6] These failures span multiple monitoring domains—exactly the scenarios where siloed tools struggle to provide a comprehensive picture. Teams waste hours correlating data from different sources while customers experience degraded service.

The business impact is real: Engineering teams spend roughly 30% of their time on managing outages, typically using manual actions that could be automated with proper observability.[7] The impact of this situation goes beyond mere inefficiency to create opportunity costs. Every hour spent hunting through fragmented logs is an hour not spent building features that drive growth or achievement of goals.

How is the Next Stage of AI Affecting Observability?

Multi-Agent AI and LLM Intricacies

The frontier of AI involves multiple models collaborating with autonomous agents to solve complex problems. Think about an agentic customer support system: one model analyzes sentiment. Another accesses knowledge bases. A third generates responses. And a fourth ensures compliance. And all of this is happening in real time with interdependencies that change based on context.

Agentic AI has created observability challenges that didn't exist a mere six months ago. Traditional application performance management tools can't trace prompt flows through multiple LLMs. They can't correlate token consumption with business outcomes. They can't detect when agent interactions create cascading performance issues.

Risk Mitigation, Compliance, and Trust in AI

New AI regulations are developing rapidly alongside AI adoption. The EU, Colorado, Texas, and Arkansas have all passed AI bills, and many other U.S. states are in the process of developing AI regulations. These laws, as well as emerging federal guidelines, establish compliance requirements that many organizations are still working to address. In many industries, non-compliance carries significant legal and operational risks that can affect business continuity and market access.

Optimizing Observability

Unified Observability as a Strategic Opportunity

Modern enterprises face an increasingly complex operational landscape, where effective observability is less about adding more tools and more about improving the clarity and usefulness of insights. A unified, AI-enabled platform can help to transform massive volumes of operational data into actionable understanding; realize concise, high-value insights that enable automated actions; accelerate innovation; and generally improve all kinds of decision-making.

The Value of Holistic Visibility

Holistic visibility is the outcome of unifying discovery, monitoring, and analysis across the full technology stack—from legacy monolithic systems to modern container-based applications. Automatically mapping every component and dependency as it emerges builds a complete, real-time view of the digital ecosystem, eliminating blind spots and reducing manual effort.

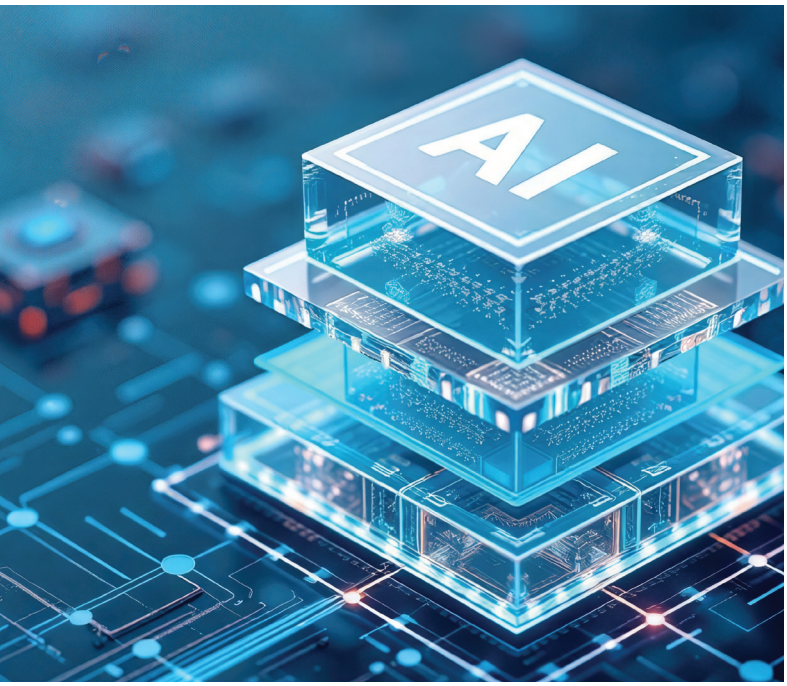

Seamless integration with cloud services and on-premises applications, combined with lightweight agents that detect processes and extend monitoring capabilities, ensures consistent data collection across diverse environments. This establishes a single source of truth, linking technical performance directly to business outcomes. With this level of visibility, teams can identify root causes faster, correlate user experience with backend performance, and make informed decisions that optimize both operations and strategy—resulting in a more resilient, efficient, and agile business.

Intelligent Resource and Cost Optimization

Managing resources and controlling costs now must comprehend AI in two senses. One is the addition of real-time visibility into critical generative AI metrics, such as token consumption, request duration, and model performance. The other is AI-driven analytics, especially on unified dashboards, to more effectively detect and respond to usage pattern changes and to forecast cost increases before they affect budgets.

Predictive cost forecasting, combined with historical usage analysis, helps identify expense drivers and recommends optimizations that preserve service quality while reducing spending. Granular infrastructure monitoring—covering metrics such as temperature, energy usage, utilization rates, and GPU performance—supports both cost efficiency and sustainability goals, enabling organizations to balance performance, operating costs, and environmental impact. The result is a more predictable, efficient, and environmentally responsible operation that is aligned with financial goals, sustainability objectives, and AI strategies.

Proactive reliability and quality management

Reliability and quality improve when operators identify and address potential issues before they affect users. Continuous monitoring, powered by advanced anomaly detection, enables the early identification of deviations from normal behavior and provides proactive alerts about current and future events.

AI-driven root cause analysis can accelerate resolution by correlating metrics, logs, and traces to a real-time topology model, allowing teams to pinpoint issues in seconds rather than hours. Continuous trace analysis identifies slow requests and errors, helping to reduce model response times and improve service quality. Automated workflows triggered by predefined thresholds or unexpected conditions enable 24/7 remediation, ensuring issues are resolved quickly—often before users notice—while protecting experience, reputation, and operational continuity.

Unifying Data for Deeper Insights and Business Alignment

Consolidating metrics, logs, and traces—while preserving the context that connects technical performance to business results—can be a useful way to create clearer, more complete understandings

of how technology affects organizational outcomes. In turn, consolidation may benefit from a high-performance data lakehouse architecture that can to utilize massive volumes of observability, security, and business data without complex indexing or rigid schema management. Ideally, it can deliver both scalability for deep analytics and speed for real-time responsiveness.

With powerful query languages and collaborative analytics tools, teams can explore data at petabyte scale, ask ad-hoc questions, and receive context-rich answers with visualizations and actionable recommendations. This shared, data-driven understanding can improve coordination between technical and business teams by clearly bridging the gap between operational performance and strategic goals.

Specialized Observability for Agentic AI Applications and Innovation Acceleration

Agentic AI presents unique behaviors, lifecycles, and interaction patterns that demand specialized observability. Seamlessly integrating this capability with broader application monitoring helps ensure full visibility into multi-agent interactions, as well as the performance of integrated agents, without adding complexity.

Code-level insight into generative AI models, vector databases, and orchestration frameworks support comprehensive quality assurance. Monitoring guardrail metrics—such as hallucinations, prompt injections, PII leakage, and toxic content—maintains accuracy, safety, and compliance. For agentic AI, tracking decision rationales, emergent behaviors, inter-agent dependencies, and goal alignment ensures transparency, explainability, and trustworthiness.

Compliance features enable the maintenance of full data lineage from prompt through agent actions to final response, thereby supporting auditability and regulatory requirements. AI-powered developer assistance enables real-time problem resolution in development environments and automatically generates queries, dashboards, and workflows from natural language input. And adding a control layer for managing and optimizing distributed AI agents at scale turns specialized observability data into actionable intelligence—empowering organizations to adopt and scale AI confident in their ability to maintain trust and accelerate innovation.

Moving to Unified, AI-Powered Observability

The landscape of digital transformation continues to evolve rapidly, bringing both significant challenges and opportunities. Cloud-first, AI-driven environments are the modern application design and delivery paradigm, but introduce complexities such as fluctuating AI costs, unpredictable model behaviors, fragmented data sources, and increasing regulatory demands. These factors make traditional monitoring approaches less effective and highlight the opportunity for a more unified, intelligent observability strategy.

Observability that is both unified and AI-driven promises quick technological improvement and can provide a foundation for enduring strategic advantage. By turning data complexity into clear, actionable insights, organizations can shift from reactive troubleshooting to proactive optimization

and automated resolution. This enables faster issue detection, cost-effective AI management, and improved compliance without compromising innovation.

Adopting these advanced observability capabilities may better position an organization to navigate increasing complexity with confidence. It can enhance operational resilience, accelerate innovation, and align technology performance more closely with business goals. As complexity continues to grow, embracing unified observability will be essential for maintaining agility and competitive strength in the years ahead.

About Dynatrace

Dynatrace (NYSE: DT) is the world-leading AI-powered observability platform. We're advancing observability for today's digital businesses, and helping to transform the complexity of modern digital ecosystems into powerful business assets. By leveraging AI-powered insights, Dynatrace enables organizations to analyze every transaction, automate at the speed of AI, and innovate faster and without limits to drive their business forward.

To gain an elevated level of understanding that helps your business meet the challenges of today and prepare for what's to come, sign up for a free 15-day Dynatrace trial at www.dynatrace.com.

1 Pavlo Khropatyy, Intellias, "Microservices Architecture: A Guide to Sustainable Modernization," February 19, 2021.
2 Mark Pittman, Forbes Tech Council, "https://www.forbes.com/councils/forbestechcouncil/2025/06/03/the-massive-implications-of-data-becom...," June 3, 2025.
3 Customer Data Platform Institute, "Average Enterprise Uses 1,000+ Apps: Salesforce Survey."
4 Cisco, "Developers Spending More Time Firefighting Issues Than Delivering Innovation," May 7, 2024.
5 John McBride, OpenSauced, "How We Saved 10s of Thousands of Dollars Deploying Low Cost Open Source AI Technologies At Scale with Kubernetes," May 13, 2024.
6 Wu, H., Ling, B., Yang, B., Wang, X., & Chen, Y. (2021, June). BSODiag: A global diagnosis framework for batch servers outage. In 2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQoS) (pp. 1–10). IEEE, "A Novel Proof-of-Reputation Consensus for Storage Allocation in Edge Blockchain Systems," 2021.
7 Peter Marelas, New Relic, "The Role of Observability within Organizations is Changing," Oct. 22, 2024.