



Mastering AI observability with Dynatrace and AWS

An IT leader's guide

What's inside

INTRODUCTION

Observability into the new age of intelligence

SECTION ONE

Combatting the challenges of AI observability

SECTION TWO

Dynatrace and AWS: A unified solution for AI-powered visibility

REAL-WORLD USE CASES

AI observability in action

CONCLUSION

Confidently scale your AI and agentic innovations



INTRODUCTION

Observability into the new age of intelligence

AI adoption is accelerating across every industry, transforming applications and automating operations at an unprecedented scale.

From integrating generative AI into customer-facing products to deploying intelligent automation for internal operations, organizations are leveraging AI to innovate faster and solve for new and existing use cases at scale.

However, this rapid evolution introduces a new layer of complexity to the cloud-native landscape. With that in mind, how can teams monitor, secure, and optimize these new AI-integrated systems?

Traditional monitoring approaches are no longer sufficient, relying on siloed logs, metrics, and traces. AI workloads behave differently than traditional applications, with unpredictable performance patterns, opaque “black box” models, and unique security vulnerabilities.

These systems create an explosion of telemetry data that conventional tools cannot process effectively, leading to blind spots, delayed insights, and increased operational risk. According to the Dynatrace 2025 State of Observability report, AI capabilities are now the No. 1 criterion for selecting an observability platform – ahead of cloud compatibility and data collection.

This guide offers a set of best practices for implementing end-to-end observability for agentic AI, generative AI, and large language models (LLMs). Discover the critical challenges that IT leaders, operations teams, site reliability engineers (SREs), and developers face when managing AI-powered applications – and how the strategic partnership between Dynatrace and Amazon Web Services (AWS) can help you overcome them.

Combatting the challenges of AI observability

While the promise of AI is transformative, integrating it into applications and infrastructure introduces a new class of operational challenges.

AI-powered systems are fundamentally different from traditional software, rendering legacy monitoring tools and siloed approaches inadequate. Understanding these unique complexities is the first step toward building a robust observability strategy.

AI workloads are dynamic and non-deterministic. Unlike conventional applications with predictable execution paths, AI models can behave in unexpected ways based on user inputs, data drift, and continuous retraining. This creates significant hurdles for teams tasked with maintaining performance, reliability, and security.


Key challenges



Unpredictable workloads



Opaque models



Complex interactions



Security and compliance risks



Telemetry data

Key challenges in observing AI workloads

Organizations face several distinct obstacles when trying to gain visibility into their AI-driven services:

1. UNPREDICTABLE WORKLOADS AND PERFORMANCE

AI applications, particularly those using generative AI, often exhibit volatile resource consumption. Workload spikes can be sudden and massive, putting immense pressure on underlying infrastructure.

Traditional latency patterns and performance baselines no longer apply, making it difficult to proactively identify performance degradation or manage costs effectively.

2. OPAQUE “BLACK BOX” MODELS

Many AI models operate as “black boxes,” where the internal logic and decision-making processes are not transparent. When a model produces an incorrect, biased, or slow response, it is incredibly difficult to trace the root cause. This lack of explainability leaves operations and SRE teams without the context needed to debug issues quickly.

3. NON-LINEAR, COMPLEX INTERACTIONS

Modern AI applications often rely on agents that orchestrate multiple tools, services, and other AI models to complete a task. These interactions are non-linear and can vary significantly with each user prompt.

Tracing a single request across this dynamic web of dependencies is nearly impossible with traditional tools that were built for more straightforward, predetermined architectures.

4. NEW SECURITY AND COMPLIANCE RISKS

Generative and agent AI introduce a new attack surface. Risks like prompt injection, model misuse, and the inadvertent exposure of sensitive data in model responses require constant monitoring.

Traditional security tools are not equipped to analyze AI-specific behaviors, leaving organizations vulnerable to new threats and compliance violations.

5. THE EXPLOSION OF TELEMETRY DATA

AI-powered, cloud-native architectures generate an unprecedented volume and variety of telemetry data. Simply collecting more logs, metrics, and traces is not the answer. Without the ability to contextualize this data and unify it, teams are left drowning in information but starved of actionable answers.

Dynatrace and AWS: A unified solution for AI-powered visibility

To confidently scale AI-powered applications, organizations need a platform that provides deep visibility, contextual understanding, and intelligent automation.

Dynatrace on AWS delivers exactly that – a unified observability solution on a global infrastructure designed for the demands of modern cloud and AI environments.

By combining the AI-powered automation and advanced observability of the Dynatrace platform with the flexibility and scale of AWS offerings like Amazon Bedrock and Amazon Bedrock AgentCore, enterprises can finally close the visibility gaps created by AI workloads. This collaboration empowers teams to monitor, optimize, and secure their AI initiatives from day one.

Key capabilities to consider



Unified
telemetry



AI-powered
anomaly detection



End-to-end
distributed tracing



Guardrail and quality
monitoring



Intelligent
cost analytics



Autonomous
remediation

Key capabilities leaders should consider

Dynatrace on AWS is built on a foundation of openness, intelligence, and automation. It provides the critical capabilities needed to manage the entire lifecycle of AI-powered applications, including the following:

1. UNIFIED, REAL-TIME TELEMETRY

The Dynatrace Platform ingests high-fidelity metrics, logs, and traces from across the AWS environment, including critical AI-specific signals from services like Amazon Bedrock.

By leveraging open standards like OpenTelemetry, it captures everything from user prompts and model responses to token consumption and infrastructure performance. All this data is unified within the Dynatrace Grail® data lakehouse, providing a single source of truth for all teams.

2. AI-POWERED ANOMALY DETECTION AND ROOT-CAUSE ANALYSIS

Simply collecting data isn't enough. Dynatrace Davis® AI continuously analyzes telemetry in real time to understand dependencies and baseline normal performance.

It automatically detects anomalies in application behavior, AI model latency, and even cost, moving beyond simple error and latency metrics. When a problem is detected, Davis AI pinpoints the precise root cause, eliminating war rooms and enabling teams to resolve issues in minutes, not hours.

3. END-TO-END DISTRIBUTED TRACING AND DEPENDENCY MAPPING

In the world of AI agents and non-linear workflows, understanding data flow is paramount. Dynatrace provides end-to-end distributed tracing that maps every interaction, from the initial user input to model reasoning, tool calls, and vector database queries.

This creates a real-time topological model of an entire system, allowing teams to instantly pinpoint bottlenecks, errors, or costly loops that would otherwise be invisible.

4. COMPREHENSIVE GUARDRAIL AND QUALITY MONITORING

Building trustworthy AI requires robust governance. The Dynatrace Platform offers out-of-the-box monitoring for AI guardrails, enabling you to track metrics related to toxicity, personally identifiable information (PII), and denied topics.

This facilitates efforts to ensure that AI models are behaving safely and meet compliance requirements. By correlating these guardrail outcomes with user experience and performance data, you gain a holistic view of an AI application's quality and reliability.

5. INTELLIGENT COST AND PERFORMANCE ANALYTICS

AI workloads can lead to unpredictable costs. Dynatrace provides detailed analytics and forecasting for token consumption, prompt caching and API usage.

By leveraging the predictive power of Davis AI, organizations can forecast spend, detect cost anomalies, and correlate resource usage with business outcomes. This enables teams to optimize their AI deployments for both performance and ROI.

6. AUTONOMOUS REMEDIATION AND OPTIMIZATION

Agentic AI transforms observability into action by automatically detecting anomalies and executing corrective measures without human intervention. It continuously optimizes resource allocation to maintain performance and cost efficiency while enforcing compliance guardrails in real time. This proactive approach supports resilient, secure, and cost-effective operations at scale.

By integrating Dynatrace's deep observability capabilities directly with AWS's foundational AI services, organizations can move past the challenges of traditional monitoring. This partnership provides the answers and automation needed to innovate securely, manage costs effectively, and deliver exceptional experiences with AI-powered applications at scale.

Achieve operational excellence with faster resolution

In dynamic AI environments, speed is everything. The longer an issue goes undetected, the greater the impact on customers, revenue, and brand reputation. Dynatrace on AWS is designed to deliver answers, not just data, enabling teams to resolve issues at an unprecedented speed.

- The Dynatrace **Davis® AI** engine automates root-cause analysis, eliminating the need for manual investigation and lengthy war room calls. One financial services customer reported reducing incident resolution times from hours to just minutes, freeing up valuable engineering resources to focus on innovation instead of firefighting.
- With **real-time anomaly detection and dependency mapping**, teams can identify performance degradations and potential failures before they affect users. This proactive stance minimizes downtime and supports a consistently high-quality service.
- By embedding **agentic AI** with the AWSDevOps Agent the platform goes beyond insights to enable autonomous troubleshooting and remediation. Developers and platform engineers can resolve issues proactively without waiting for manual intervention, accelerating delivery and reducing operational friction. This agentic approach breaks down silos and fosters a culture of collaboration and innovation.



Davis® AI



Anomaly detection



Agentic AI

Optimize costs and maximize ROI


AI workloads can introduce unpredictable and substantial costs, from API calls to infrastructure consumption. Gaining control over this spending is critical for ensuring a positive return on investment.

- That starts with the need to eliminate costly blind spots. Dynatrace provides complete visibility into **AI-related spending**, including token consumption and API usage. By tracking these metrics in real time, organizations can understand exactly where their budget is going.
- Dynatrace provides out-of-box **cost dashboards** for AWS Bedrock to accelerate AI FinOps best practices
- Dynatrace AI Model version and A/B testing – compare models and versions to **validate improvements and spot bottlenecks** across token usage and cost.
- Additionally, the predictive capabilities of Davis AI allow leaders to **forecast future costs** based on usage patterns and detect anomalies in spending. This empowers teams to attach automated remediation playbooks to cost anomalies, preventing budget overruns before they occur.
- And by correlating cost data with performance and business outcomes, teams can make informed decisions about **resource allocation** and better understand how infrastructure spend contributes to value creation.


Enhance security and support compliance efforts

AI introduces new security and compliance risks that traditional tools cannot address. The Dynatrace Platform extends observability into the security domain, helping teams protect applications and data.


- Out-of-the-box monitoring for **AI guardrails** enables teams to track and address issues such as prompt injection, model misuse, and sensitive data exposure. While we focus on monitoring, AWS Bedrock's ability to enforce guardrails complements this approach. Together, the partnership delivers total visibility and the capability to enforce these guardrails, promoting trustworthy AI and compliance.
- Further, by providing a **unified platform for observability and security**, Dynatrace helps teams identify and remediate vulnerabilities faster. The platform's real-time topology model provides the context needed to understand the full impact of any potential security threat.
- It's critical to maintain a **comprehensive audit trail** by documenting all inputs and outputs, including full data lineage from prompts to responses. Dynatrace enables long-term storage of prompts and provides dashboards to visualize AI system behavior.



AI guardrails



Unified platform



Audit trails

Deliver superior user experiences

Ultimately, the goal of any application is to deliver value to its users. AI observability ensures that AI-powered services are not only functional, but also provide a high-quality, relevant, and satisfying experience.

- Go beyond basic operational metrics to measure what truly matters: the **quality and relevance** of AI-generated responses. By establishing baselines for “good” behavior, teams can quickly identify when AI is failing to meet user expectations.
- End-to-end distributed tracing connects every user interaction to its underlying technical dependencies. This allows teams to **instantly trace a poor user experience** back to its root cause, whether it’s a slow model, a failing API call, or an infrastructure bottleneck.

REAL-WORLD USE CASES

AI observability in action

While the principles of AI observability are compelling, its value is most clearly demonstrated through tangible outcomes.

Organizations across various industries are already leveraging Dynatrace on AWS to transform their operations, enhance efficiency, and drive business growth.

The following use cases illustrate how moving from fragmented monitoring to intelligent, automated observability empowers teams to not only solve problems but to innovate with confidence.




Use case 1: Accelerating incident resolution in financial services

A leading financial services firm faced a common but costly challenge. During operational incidents, they were forced to assemble large-scale “war rooms,” often involving over 50 teams on hours-long calls to diagnose and resolve a single issue. This approach was expensive, both in terms of personnel hours and the opportunity cost of pulling engineers away from strategic projects. Furthermore, prolonged downtime directly affected their end users and business performance.

Solution: The organization implemented an AI-powered root-cause analysis agent built on AWS and monitored by Dynatrace. When an alarm is triggered, Davis® AI automatically begins an investigation. It analyzes telemetry data, identifies dependencies, and pinpoints the likely root cause of the problem.

Outcome: By the time human operators log in to address the issue, Davis® AI has already proposed the most probable cause and recommended next steps. Incidents that previously took hours to resolve are now handled in minutes. This has translated into huge savings, faster application recovery, and has enabled their teams to focus more on building value and less on troubleshooting.



Use case 2: Boosting developer productivity and efficiency

An enterprise was building an internal AI agent designed to help its development teams write code more efficiently. The goal was to create a system where developers could ask the agent to find and provide instructions for using pre-built internal libraries, such as authentication modules. The problem was twofold. Developers were often unaware these libraries existed, leading them to rebuild functionality. And the company had no way to measure if the new AI agent was being used effectively or providing value.

Solution: The customer used Dynatrace to gain full observability into the entire workflow of their new GenAI-based agent. This allowed them to trace requests from a developer's prompt to the agent's interaction with various data sources, like Confluence and source code repositories, and back to the final response.

Outcome: The company gained the ability to measure the effectiveness and quality of the agent's responses. More importantly, they could analyze user behavior to determine whether the tool was being used for its intended purpose—solving complex coding challenges—rather than simple queries that could be answered with a web search. This insight allowed them to optimize the agent's performance, control costs associated with its use, and ultimately prove its ROI by increasing developer efficiency.



Use case 3: Autonomous troubleshooting for cloud-native applications

Detecting a performance slowdown is only the first step. Understanding the “why” behind it is what truly matters. One organization integrated AWS DevOps Agent, an autonomous agentic AI DevOps service, with the Dynatrace Platform to create a closed-loop system for event resolution.

Solution: When Dynatrace Davis® AI detects a performance degradation—for instance, a slowdown in a critical CRM application—it automatically triggers an investigation by AWS DevOps Agent. While Dynatrace pinpoints the affected component, such as a database, AWS DevOps Agent autonomously investigates the underlying configuration changes, deployment history, and infrastructure modifications.

Outcome: In one scenario, Dynatrace identified a database as the source of slowness. AWS DevOps Agent then discovered the root cause was a recently dropped database index and provided the exact SQL statement required to fix it. This combination of causal and generative AI delivered an actionable mitigation plan directly within the Dynatrace console, enabling a rapid and precise resolution without extensive manual effort.

CONCLUSION

Confidently scale your AI and agentic innovations

The agentic AI era is here, and it is fundamentally reshaping how businesses operate, innovate, and compete.

Integrating AI into applications and infrastructure unlocks immense potential but also introduces a new class of operational complexities.

Unpredictable workloads, opaque models, and novel security risks mean that traditional, siloed monitoring approaches are no longer viable. Success in this new landscape requires a strategic shift to a modern, unified, and AI-powered observability platform.

Mastering agentic, AI, and LLM observability is not just about collecting more data; it's about getting precise, contextual answers that enable teams to act decisively. It requires the ability to see across an entire technology stack in real time, from user interactions and application code to AI models and the underlying AWS infrastructure.

By embracing this new paradigm, organizations can move beyond reactive firefighting and begin to proactively optimize performance, control costs, and secure their most critical AI-driven services.

The strategic partnership between Dynatrace and AWS provides a clear path forward. By combining the deep, automatic, and intelligent

observability of the Dynatrace platform with the power and scale of AWS cloud and AI services, teams gain the unified solution needed to navigate this complexity with confidence.

This collaboration empowers organizations to:

- **Eliminate blind spots** with end-to-end distributed tracing for complex AI workflows.
- **Resolve issues faster** with AI-powered root-cause analysis that delivers answers in minutes.
- **Control costs and maximize ROI** with intelligent analytics for AI token and API consumption.
- **Build trustworthy AI** by monitoring for quality, safety, and compliance with AI guardrails.

The journey to AI-powered transformation can seem daunting, but you don't have to go it alone. The tools and strategies outlined in this guide provide a blueprint for success, enabling teams to innovate faster while ensuring operational excellence.

Take the next step

Ready to see how Dynatrace and AWS can support your AI observability needs and accelerate your transformation journey?



EXPERIENCE IT YOURSELF

Sign up for a free Dynatrace trial and explore these capabilities firsthand in a sandbox environment.

[Start a free trial →](#)



LEARN MORE

Reach out to your AWS or Dynatrace representative to schedule a personalized demo and discuss how Dynatrace on AWS can address your specific challenges.

ABOUT DYNATRACE

Dynatrace is advancing observability for today's digital businesses, helping to transform the complexity of modern digital ecosystems into powerful business assets. By leveraging AI-powered insights, Dynatrace enables organizations to analyze, automate, and innovate faster to drive their business forward. Learn more at www.dynatrace.com.

