

Log monitoring powered by AI for operations at petabyte scale

How modern teams transform massive log volumes into contextual intelligence by automating ingestion and analysis to replace months of manual work.

At scale, log volume growth drives the need for unified context

SRE and IT Ops teams today face environments where logs multiply faster than the ability to act on them:

- Explosive log growth driven by cloud-native architectures
- Rising storage and ingestion costs with diminishing returns
- Difficulty manually extracting actionable insight from massive log volumes

At scale, more logs don't equal better operations. They often slow response and increase noise if not unified and contextualized with all your telemetry.

THE OPERATIONAL SHIFT

Precision becomes critical at petabyte scale

When log volumes reach petabyte scale, traditional indexing, sampling, and manual analysis approaches stop working. A key operational shift is already underway:

- Logs continuously analyzed in real-time, not retrospectively searched
- Signal quality matters more than raw volume
- AI surfaces what actually impacts operations



Getting insights into logs has never been easier. Just ask your question like you're asking a buddy. No complex query language, no special syntax."

— ROHAN SHAH, Director of Log Specialist Team, Dynatrace

REFRAMING LOG MANAGEMENT FOR MODERN OPERATIONS

From log storage to operational intelligence

Modern log monitoring is defined by what happens to data before it reaches human analysis: how it's ingested, enriched, and connected to the operational context that gives it meaning.

Real-time analysis	Analyze and enrich data as it arrives, not indexed and forgotten until something breaks.
Automatic enrichment	Every log arrives enriched with the metrics, traces, topology, and ownership context needed to act on it, instead of as a raw event to be correlated later.
Contextual precision	Teams see which patterns matter and which don't, so they respond to real risk instead of drowning in alert fatigue.



What precision looks like at IoT scale

ADT generates billions of IoT events daily across homes nationwide. After consolidating onto Dynatrace, the team used OpenPipeline to extract only the highest-signal business events from the ingestion stream, powering real-time product intelligence without the cost and data sprawl of broad log collection.

1. Consolidated observability onto a single platform, eliminating tool context-switching
2. OpenPipeline filters billions of daily events down to a compact, high-signal dataset for cost-effective processing
3. Real-time product dashboards now query weeks or months of history near-instantaneously
4. Anomaly detection runs directly on live product behavior signals

The core challenge (and the outcome)

As an IoT company, ADT needed to move beyond raw log storage toward operational precision: extracting meaning from massive event streams without drowning in data or cost. The shift to Dynatrace made that possible: not by storing less, but by being smarter about what gets analyzed and how.



By extracting only the most important events through the ingestion pipeline, we're able to do product trending, real-time dashboards, and anomaly detection on things happening within our product.

It's really empowered our ability to see what customers see."

— RYAN GOW, VP Product Engineering, ADT

Human analysis does not scale to cloud complexity

Analysis powered by AI is critical for modern log operations across every team, from SREs to executives:

- Natural language queries across petabyte-scale data – no DQL required
- Detecting meaningful patterns across massive log streams automatically
- Reducing noise and false positives before they reach a human
- Supporting faster triage and root-cause analysis across distributed systems

This directly aligns with SRE realities: telemetry volume has outpaced human capacity. Automation depends on trusted, high-quality signals and logs are foundational when managed correctly.

Precision enables action

When logs are contextualized automatically, analyzed continuously, and integrated into a broader observability platform, they unlock capabilities that weren't previously possible:

- Smarter alerting with far fewer false positives
- Automated operational responses triggered by trusted signals
- Progression toward AI-assisted and autonomous operations

This wouldn't be possible without:

- 1. Indexless data lakehouse:** Logs, metrics, and traces live in a single data model. Query across all of them without moving, copying, or reconciling data first.
- 2. Always-on hot/hot storage:** All your data is hot. Query anything, instantly, regardless of age or volume.
- 3. No compromises at scale:** Designed so teams don't have to choose between cost efficiency, query speed, and scale.

WHY THIS MATTERS

Log strategy is a platform decision

At petabyte scale, SRE and IT Ops leaders must treat log monitoring as more than a tooling choice. It is a core platform capability that determines whether automation and AI are even possible.

Enterprise-scale log monitoring

Infrastructure: Petabyte-scale ingestion with automatic extraction

Governance: Self-service pipelines with central control

Analytics: From natural language to deep forensics

Unified telemetry context	Logs are a foundational part of telemetry, complementing metrics and traces to power automated workflows and AI-driven decisions.
Prerequisite for AI adoption	Operations powered by AI are only as reliable as the data they run on. Structured, governed, actionable logs are what makes automation trustworthy at scale.
Shared system of truth	Logs provide a rich, contextual operational record that, when unified and contextually analyzed, can support teams across the organization, from developers to leadership.

1 PB/day

log ingestion per tenant

10MB

per log event supported

3,000

pipelines per environment

100+

out-of-box processors and technologies

Fragmented or legacy log tools increase risk, cost, and response time as environments scale. The teams that treat logs as a core platform investment alongside metrics and traces are the ones that maintain operational control as complexity grows.

Understand what modern log management really requires

Explore the State of Log Management to learn:

- How leading teams are rethinking logs at petabyte scale.
- Why insights powered by AI matters more than retention.
- What capabilities distinguish modern log platforms from legacy tools.

[Read the report](#)

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.com/blog](#)

 [@dynatrace](#)