dynatrace

Log management and analytics

Simple, fast, and efficient — unified log observability and automated analytics for actionable insights

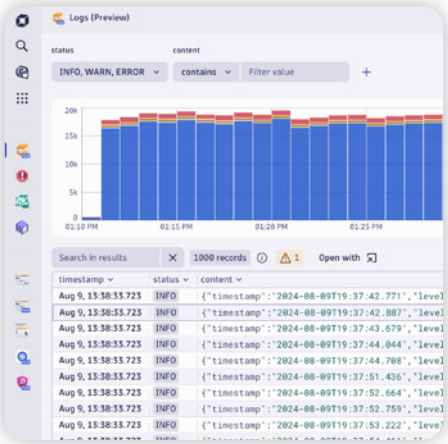

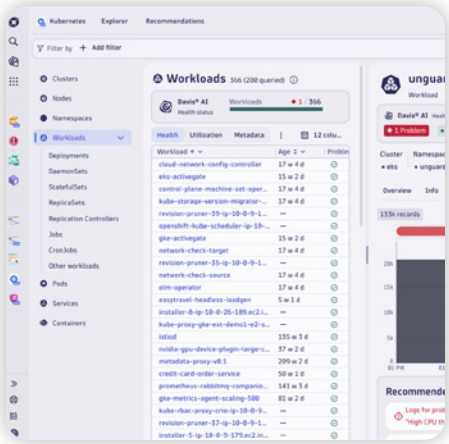

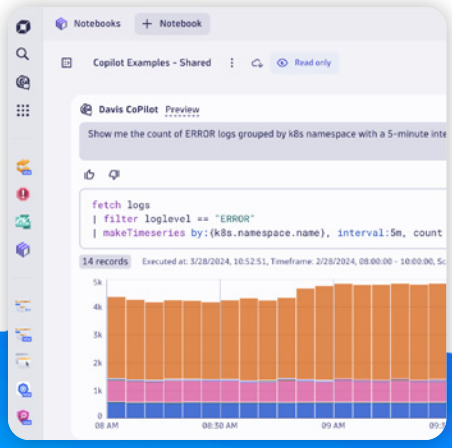

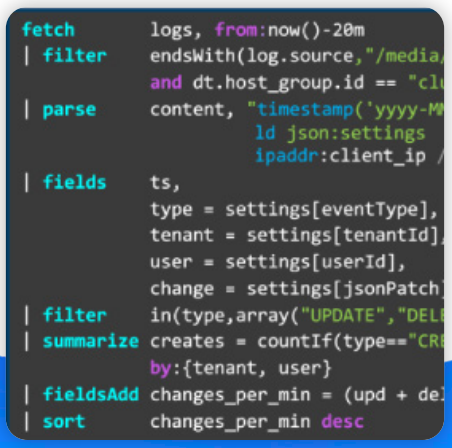

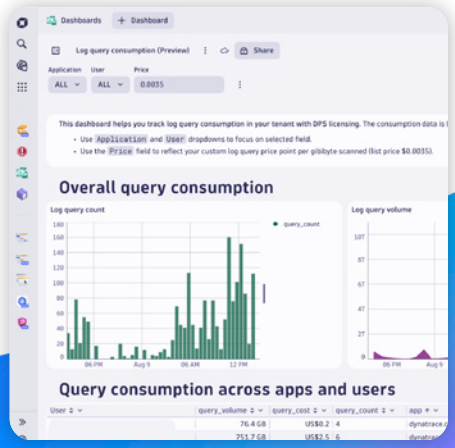

Introduction

Traditional log management is manual and siloed, limiting scalability and innovation. Dynatrace's full-stack approach automatically integrates traces, infrastructure monitoring, and the customer experience, eliminating data silos and helping you get more value out of your data.

Our AI-powered solution offers a unified view of application health and business impact, empowering organizations to not just solve issues faster, but also focus on the things that mater most: cutting costs while delivering better customers experiences.

Traditional log management tools require tradeoffs

Traditional monitoring tools and platforms impose cost, speed and insights tradeoffs.

Site reliability engineering (SRE) teams can gain deeper insights through higher volumes of logs, metrics, and traces that are retained over long periods of time. This detailed data is associated with improved analytics and AIOps.

But increased data retention also creates higher storage cost, so monitoring platforms often sample a portion of data, and customers often retain data for shorter periods of time than they would prefer.

Some monitoring solutions enable customers to archive older logs to cold storage, a less expensive form of storage. But now teams cannot easily query or quickly access the data, meaning it will take longer time to find an answer.

The process of requesting access to log data from cold storage and moving it to hot fast storage is often referred to as rehydration. Depending on the storage type, rehydration retrieval times can range from hours to days. Once the data is retrieved, it must be indexed so it can be queried — an expensive process that can also take hours to days.

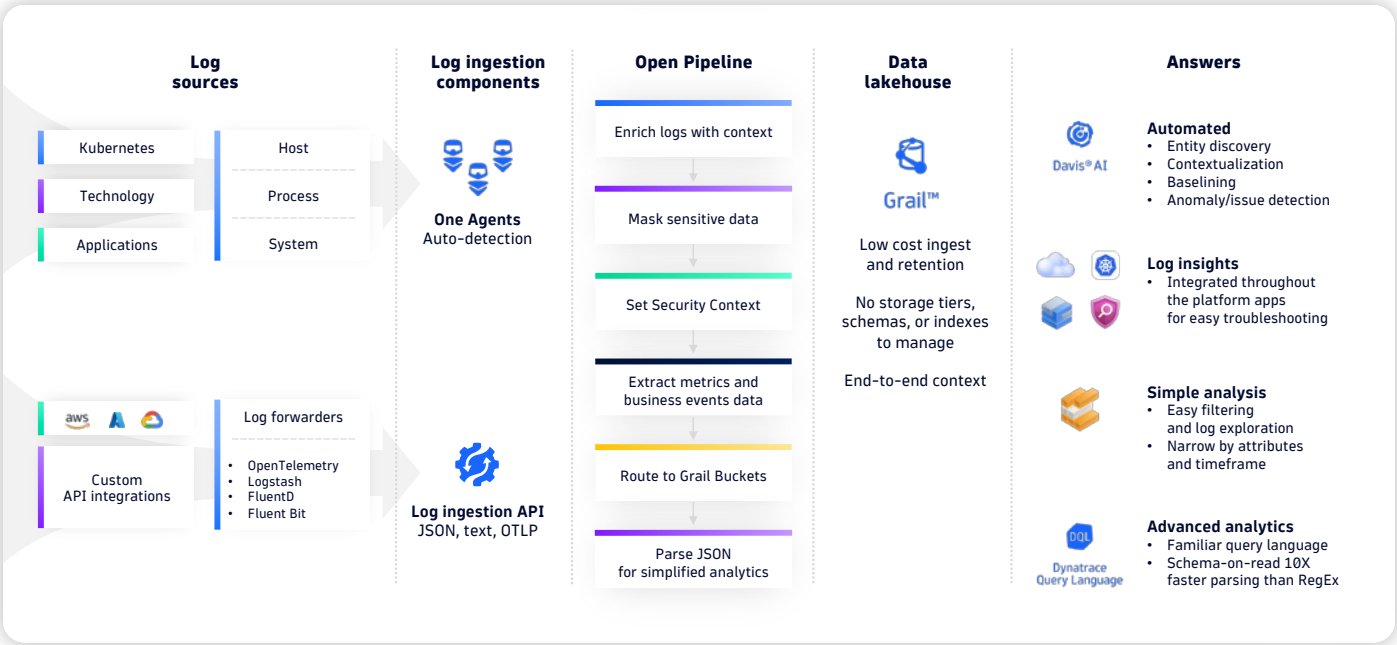

Introducing Logs management and analytics from Dynatrace

Dynatrace breaks the traditional triangle of cost, speed, and insight tradeoffs by providing maximum data context, data fidelity, and speed of insights at the lowest possible cost and effort.

Teams maximize the value of event logs when automatically contextualized with other observability data — including application, infrastructure, business and security data — then analyzed in real-time using causation-based AI.

This gives IT teams a single source of truth that reflects full stack modeling of applications, services, processes, hosts and underlying infrastructure — with comprehensive topology and dependency mapping — so they can see the precise paths a user or request took during each step of an emerging issue. No war rooms required.

Root-cause analysis, remediation and optimization is fast, precise, and actionable.

Get answers to any question, in real time, at scale, and on budget

Dynatrace is the only platform to provide observability, security, and business analytics in context with no indexes, rehydration, or sampling.

Faster troubleshooting: Gain real time access to log events data in context of your traces, metrics, and topology to reduce mean time to repair by 90%.

Cost-effective architecture: With no more schemas to manage or indexes to wait on, teams no longer have to pay a premium for data they'll never use.

Actionable analytics: Query and parse logs to identify, metric, and reduce major outages and degradations by 60%.

Situational awareness: Collaborate securely with live data and generate information-rich tables, charts, and dashboards.

Unified observability: Replace distracting war rooms and alert storms with AI-powered root cause analysis to increase DevOps efficiency by over 40%.

Overview of Logs management and analytics features

Logs Application

Simplified log management and analytics

- Quickly analyze logs with a Log distribution dashboard
- Powerful search and customizable filters without requiring a query language
- Find surrounding logs before and after events with a single click

Ingest support for over 600+ logs sources

- Dynatrace One Agent add context to all ingested logs
- Native integration with all AWS services via AWS firehose
- Open Telemetry support
- Native Fluent Bit support for Kubernetes with Dynatrace Operator
- Support for common log sources like Syslogs including networking, Fluent Bit, and FluentD
- API integration for Azure and Google Cloud
- See Hub for a complete list of logs sources supported

Logs Insights included in Infrastructure, Kubernetes, Cloud, Security Investigator and Database platform apps

- Provides Logs distribution dashboard within the app
- Provides the ability to search and easily filter during investigations
- Provides easy ability to click to see surrounding logs before and after log under review
- Automatically associates distributed traces with logs
- Do threat detection and incident analysis using logs within Security Investigator App

Dynatrace Query Language (DQL) and Notebooks including logs

Ability to do deep analytics instantly and collaborate on results

- Analyze any data in Grail (Dynatrace Data Lakehouse) without the need for predefined schemas, indexes, or storage management to move data
- Query across multiple observability signals: traces, logs, business events, metrics, Davis AI insights, and more
- Supports complex multi-step analysis in real-time simplifying queries
- Ability to share log analysis or security forensics investigations with a dynamic, repeatable and collaborative resource like Notebooks

Davis AI provides predictive, causal and generative AI for logs

- Davis AI provides automated anomaly detection across your entire IT and application landscape through causal AI
- Davis AI provides automated problem detection and root cause analysis with 100% accuracy through causal AI
- Davis AI predictive capabilities when matched with Dynatrace workflow automation can help customers automatically respond to forecast trends (like when you need to add host to your cloud environment)
- Generative AI provides the ability for customers to create dashboards, write Dynatrace Query Languages, create notebooks, explain log content and more all with natural language

Logs, traces, events and metrics in a unified view

One Agent and Smartscape topology work together to provide context for logs

- Additional context is added to logs automatically during ingest based on Smartscape topology that maps all relationships of all entities in Grail
- Davis AI problem cards associate traces, events, metrics and logs to a specific problem to decrease time to MTTI and MTTR
- Dynatrace Grail data lakehouse is scalable, fast single source of truth for all observability signals

Out-of-the-box dashboards provided for Logs

Customers can now take advantage of ingest and cost-control dashboards

- Gain complete transparency to Dynatrace Platform Subscription cost of ingest, retention, and queries
- Top query volumes per application and per user make recommendations on how to optimize cost, such as converting logs to metrics for dashboards
- Gain insights into ingest volume totals, top ingest sources, health indicators and bucket/retention usage

Benefits of Logs management and analytics features

Hassle-free management

Store petabytes without schemas, indexing, or rehydration

- Instantly access petabytes of data without needing to reconstitute and reindex from cold archives.
- Just run it — with no rigid schemas, no expensive indexes to manage, and no need to know what you want to query upfront.
- Pick the retention that fits your business and compliance needs, from debug to audit.

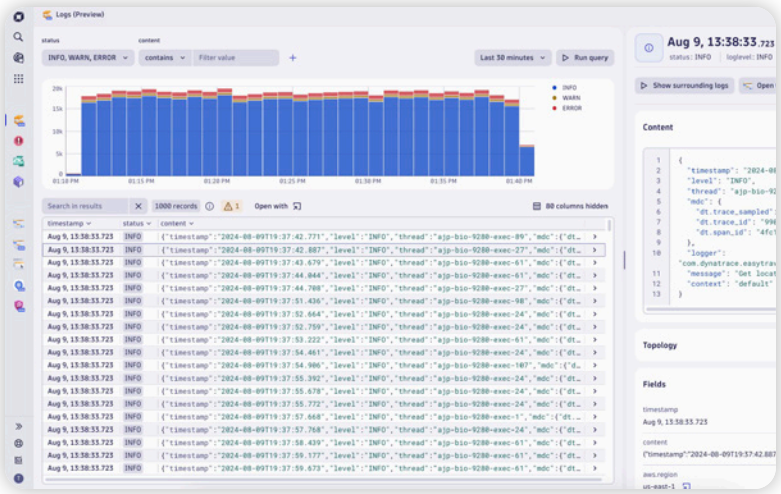

Real-time analytics at scale

Analyze de-siloed data sources at once in full context

- Consolidate data into a single purpose-built data lakehouse to get answers with one ubiquitous query approach.
- Analyze logs in context of traces, user sessions, and topology with intuitive visual diagnostics and powerful queries.
- Pinpoint root cause and eliminate manual correlation with AI-powered analysis to automatically reveal relevant log lines.

Accelerate time to value

Easily turn logs to metrics to dashboards

- Optimize costs with the flexibility to create metrics from log data and pin to dashboards at ingest, independent of retention strategy.
- Turn any query into a metric and dashboard without needing to rehydrate or reindex from an archive.
- Eliminate manual effort and alert storms with auto-baselining, anomaly detection, and root-cause analysis.

Precise answers with embedded expertise

Explore logs in context of auto-discovered entities and topologies with dependency mapping

- Quickly focus on relevant application and cloud components for live troubleshooting and debugging.
- Explore, filter, and search log and event data in context of Kubernetes and cloud platforms.
- Use interactive visual queries to navigate and analyze logs and easily pin to dashboards.

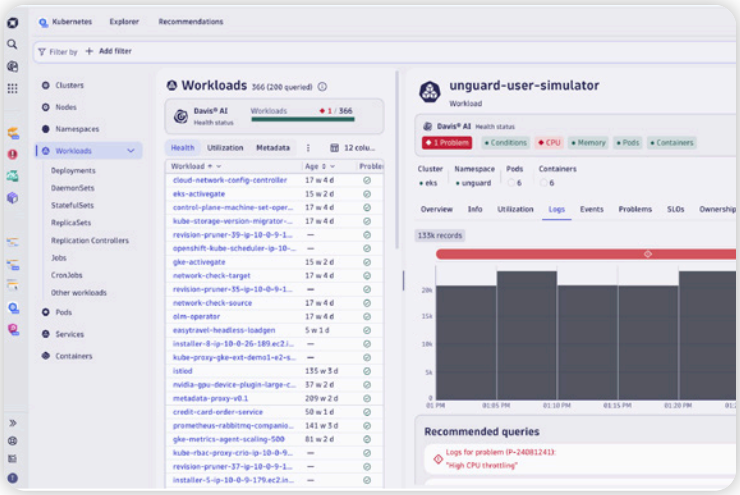


```
fetch       logs, from:now()-20m
| filter    endsWith(log.source,"/media/da
            and dt.host_group.id == "clust
| parse     content, "timestamp('yyyy-MM-d
                      ld json:settings
                      ipaddr:client_ip //I
| fields    ts,
            type = settings[eventType],
            tenant = settings[tenantId],
            user = settings[userId],
            change = settings[jsonPatch]
| filter    in(type,array("UPDATE","DELETE
| summarize creates = countIf(type=="CREAT
            by:{tenant, user}
| fieldsAdd changes_per_min = (upd + del)/
| sort      changes_per_min desc
```

Take full control

Query and parse instantly with Dynatrace Query Language (DQL)

- Leverage DQL, a familiar query language and pipe-based syntax that's simpler and more powerful than SQL and regular expressions.
- Unlock value and new potential for observability, security, and business data with schema-on-read parsing versus ingest alone.
- Enable any query, any time, on large volumes and five or more times faster parsing than regular expressions with purpose-built, high-speed algorithms and massively parallel processing (MPP).

Fits unique environments

Industry-leading cost-value ratio: broadest ingest with full control

- Leverage an open application programming interface (API) with native multicloud support for Kubernetes, Red Hat OpenShift, and Amazon Web Services, Microsoft Azure, and Google Cloud Platform environments.
- Automatically collect logs across the entire application stack, retained in context of all observability signals, using Dynatrace OneAgent.
- Make separate log pipelines obsolete with pre-ingest filtering and data transformation during processing.

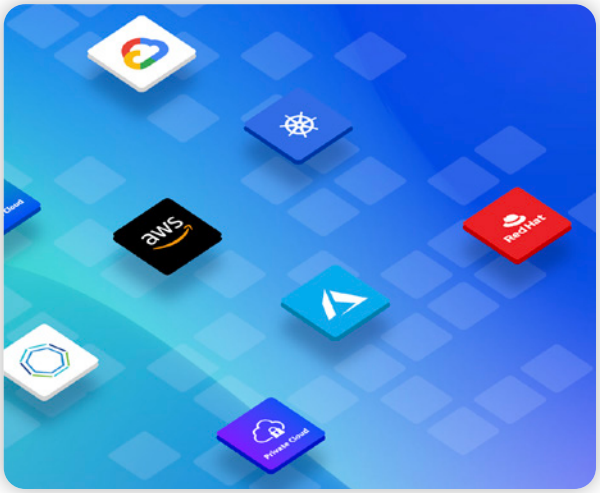

The Dynatrace difference

Dynatrace offers a unified software intelligence platform, with a unique set of capabilities that work together to support any enterprise or agency mission to accelerate value delivery. With Grail, observability, business, and security data sources are efficiently and automatically stored in causational context.

Dynatrace establishes a context for all this data — what talks to what, what depends on what, how the entire infrastructure and application topology works — without human effort. This context is maintained and analyzed across billions of dependencies and entities. This gives Dynatrace unparalleled precision in its ability to cut through the noise and empower teams to focus on what is most critical.

Dynatrace's new approach to log management and analytics brings many business benefits, including the following:

Business revenue returns. IDC reported customers received a 451% return on investment over three years. Large enterprises received 2.6M on average back to the business for every 100 supported applications utilizing Dynatrace observability platform including logs

Create operational efficiencies. According to Forrester,[1] when organizations use Dynatrace to continuously improve application and infrastructure quality, they identify issues 80% faster and reduce meantime-to-recovery (MTTR) by 90%. IDC reports Dynatrace drove improvements that resulted in platform managers being 53% more efficient, IT/DevOps application teams being 28% more efficient and developers to be 18% more effective

Deliver better customer experiences. IDC reported customers using Dynatrace had 37% few Sev1 and Sev2 outages. They spent 56% less time to resolve Sev1 and Sev2 outages and had 72% less revenue lost from Sev1 and Sev2 outages.

1. Forrester Consulting: The Total Economic Impact of Dynatrace
2. State of SRE Report: 2022 Edition
3. IDC 2024 Business Value of 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

08.12.24 BAE13629_WP_USLet_cs