EDGE DELTA

Understand & Control Your Observability Data

Observability Pipelines

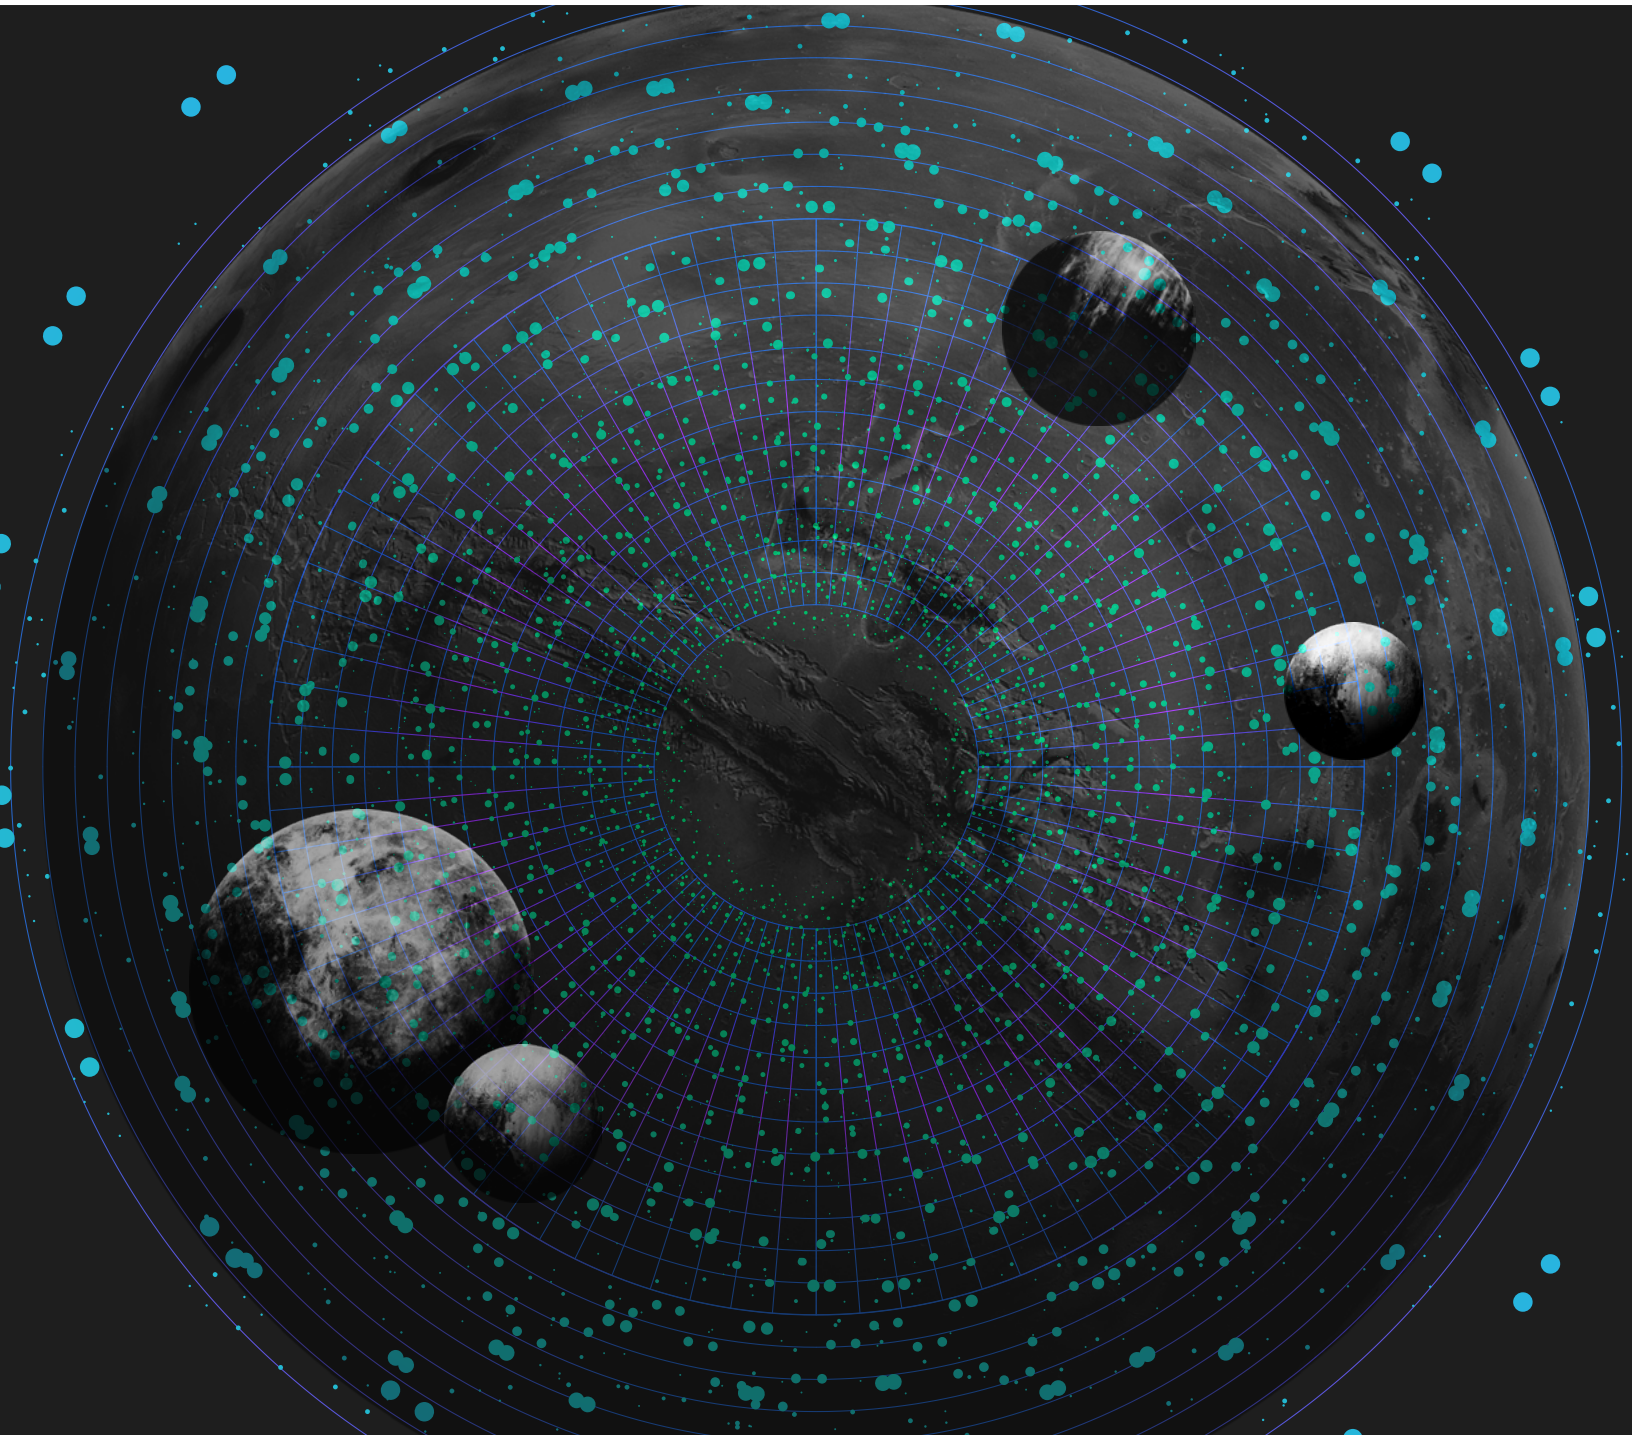

THE CHALLENGE:

Log Volumes Have Outgrown Traditional Architectures

Typically, logging architectures are built using a “centralize-then-analyze” approach, meaning that your raw data must be centralized before users can analyze and search the data. Since the rise of cloud computing and microservices, log volumes have exploded and this approach can no longer keep up.

This approach can create:

Unpredictable Costs

Indexing complete raw datasets can be cost-prohibitive. One simple mistake, like leaving debug logs on, can cause a massive overage.

Limited visibility

To help control costs, teams might sample loglines or drop datasets altogether, creating blindspots that only get larger as datasets grow.

Poor Performance

Architectures aren’t built to support today’s data volumes. It’s not feasible to collect, compress and ship terabytes of data in real time.

On top of these challenges, logging architectures have grown increasingly complex. They can include several data collection agents, different routing and transformation components, and in some cases multiple backend analytics platforms. This complexity can be a headache for operations teams to manage and can be difficult to scale as data volumes continue to grow.

At Edge Delta, we’ve spent a lot of time talking to customers about these challenges, and we’ve built a new way to solve them. We provide a unified architecture to collect, transform and route data to your observability platform in a manner that maximizes your control over what gets indexed.

This whitepaper will walk you through our approach. But first, let’s define our product as a whole.

What is Edge Delta?

Edge Delta is a platform that helps you better manage your observability data. It can be used for log management or to create observability pipelines. In this whitepaper, we'll focus on the latter. Our platform has two components:

A software agent deployed within your environment (or as close to the data source as possible).

A SaaS backend where you can manage and monitor your agents, and visualize log analytics.

Our unique approach processes and analyzes log data as it's created at the source – before anything is indexed downstream. From the agent, you can stream optimized datasets to your preferred observability platform. This approach can help you:

Simplify Your Logging Architecture: Eliminate complexity at the collection and routing levels of your observability stack.

Control Your Data Footprint: Fit new datasets into your platform and index what's most important to you.

Store and Search All Logs: Gain access to all your datasets by routing full-fidelity data to object storage – no more dropping or sampling loglines.

Without Edge Delta

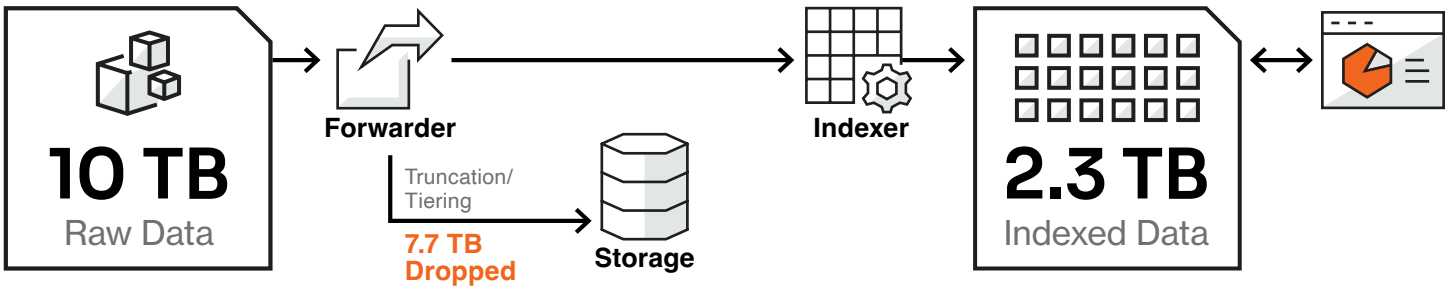

With Edge Delta

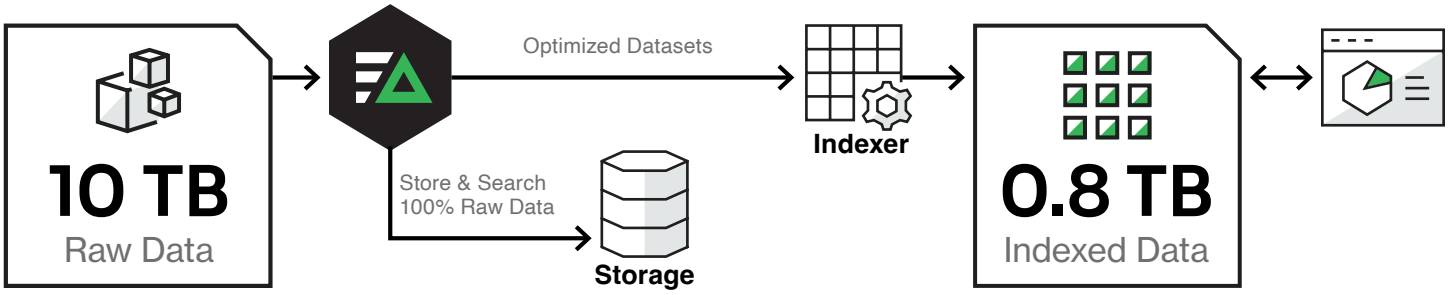

DATA COLLECTION

Maximize Vendor Flexibility, Performance and Management

To collect data, teams previously had to choose between vendor-specific or open-source offerings, both of which have their shortcomings:

Vendor-specific agents deliver superior performance and management capabilities but lead to vendor lock-in.

Open-source agents help improve vendor flexibility, but can't keep up with today's log volumes and offer no capabilities to manage or monitor agent fleets.

Frequently, teams also have multiple different agents deployed in their environment to support routing to different streaming destinations.

You can replace and/or consolidate your existing log collection agents with Edge Delta. By doing so, you can support each of your team's unique needs with one offering. Moreover, you can avoid the shortcomings of vendor-specific and open-source agents. Specifically, Edge Delta's agent is:

Vendor-agnostic: Route data from any source to any destination – and support multiple streaming destinations – with a single agent.

Highly performant: Collect and process data at virtually any volume to support real-time use cases.

Easy to manage: Manage configurations, control where data is routed and monitor agent health – all from our SaaS backend.

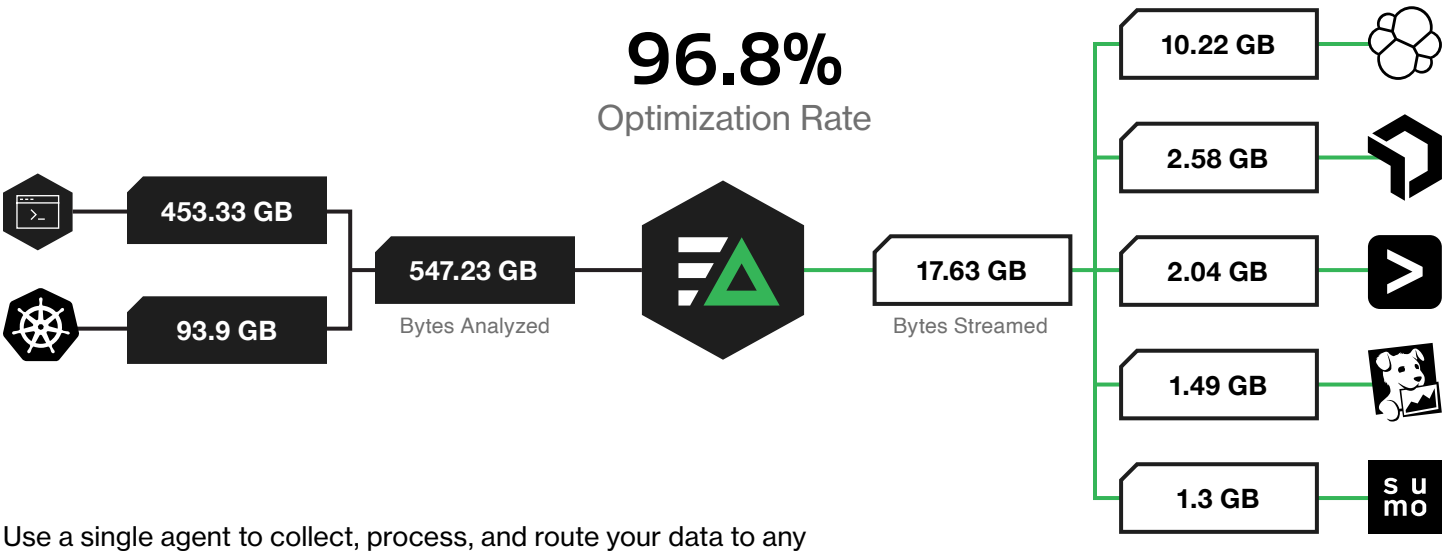

Use a single agent to collect, process, and route your data to any (or multiple) streaming destination(s).

DATA OPTIMIZATION

Remove Bottlenecks from Your Observability Stack by Moving Compute Upstream

Edge Delta uniquely processes your data as it's created at the source. It allows you to push compute to your data, instead of trying to compress and ship massive raw datasets to an observability platform.

This approach removes bottlenecks from your architectures by allowing you to optimize your datasets before you index them in your observability platform.

With Edge Delta, you can optimize data in two primary ways:

By deduplicating event streams through Patterns

By extracting metrics from your log data

Deduplicating Event Streams With Patterns

Easily detect new behaviors and better control what you index in your observability platform with Edge Delta's Patterns capability. Patterns aggregate every logline into groups of recurring behavior. This view preserves the essence of the event message while representing variant values (e.g., the time stamp) with a wildcard. Along with each aggregate, you can see metrics that provide context into the scope and sentiment of the behavior.

For example, if one log message is generated 100,000 times in five minutes, it's unlikely your team needs to see each logline. They simply need to know what the message says, how frequently it's been occurring, and whether that's abnormal.

Fama reduced Datadog costs by 60+% and gained complete visibility into their log data.

When using Edge Delta to create observability pipelines, you can use Patterns in two ways to better control index volumes.

1. Stream the optimized log messages directly to your observability platform to populate your index and dashboard.

2. Capture and index samples of the raw data making up each Pattern, along with the accompanying metrics, to give your team a snapshot of every behavior.

> “Edge Delta allows our developers to see for the first time what was making the most logs, and what was giving the most errors.”
>
> **Justin Head**
> VP of DevOps
>
> SUPER LEAGUE

By using Patterns, you can index more meaningful data and avoid flooding your platform with repetitive, noisy loglines.

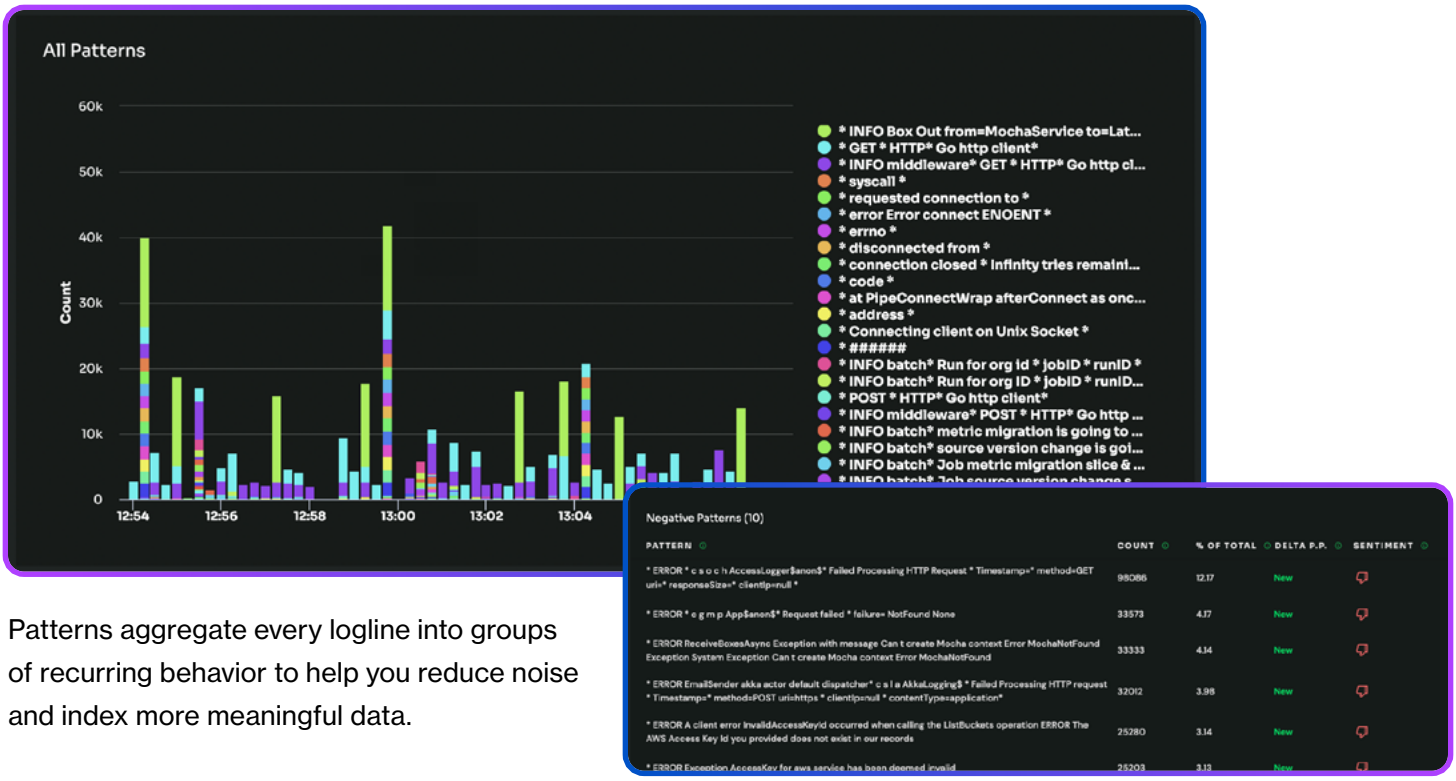

Patterns aggregate every logline into groups of recurring behavior to help you reduce noise and index more meaningful data.

Extracting Metrics From Logs

Populate dashboards in your observability tools with lightweight KPIs instead of complete raw datasets. Edge Delta's logs-to-metrics capability helps you increase visibility in a couple different scenarios:

+ If you primarily use **log management and analytics** tooling for monitoring, you likely only have visibility into the data that you've indexed (and not data that's been sampled out or dropped).
+ If you primarily use a **metrics platform** for monitoring, you might only search log data when an issue occurs. In essence, your log management tool turns into a glorified data lake vs. a real-time analytics platform.

Gain real-time visibility into your log data without adding index capacity or introducing new dashboards to your day-to-day operations. Edge Delta automatically converts logs to metrics giving your teams insightful KPIs that they can track over time. You can extract:

+ **Numeric values** (e.g., latency)
+ **Multi-dimensional values** (e.g., response code, method, and latency for an HTTP endpoint)
+ **Top k values** (e.g., top 10 endpoints with 5xx status codes)

...and more

Additionally, Edge Delta baselines the metrics over time, so it can understand what behavior is "normal" and automatically trigger alerts when anomalies occur.

Converting logs to metrics gives your team better visibility into application behavior.

+ If you use a **log management tool**, you can add analytics from datasets that you couldn't index before, giving your team newfound visibility.
+ If you use a **metrics-based tool**, you can add log analytics, creating a single executive dashboard to check application health.

Access All Your Raw Data With Live Search and Log Forwarding

Optimizing data is helpful when it comes to controlling your observability costs and data footprint. However, when an issue occurs, you'll usually want access to your raw log data.

As Edge Delta collects and processes your data, it automatically routes 100% of your raw logs to the object storage target of your choice. This allows you to retain data for as long as you need. Plus, from the Edge Delta SaaS backend, your team can search on this data at any time.

In addition, to live search, you can ensure that commonly searched datasets are always available in your monitoring platform. To fulfill this requirement, you can use Log Forwarding to pass through data your team is likely to use. This could include all ERROR or WARN level logs, log data tied to deployments or updates, or any other subset of data your team needs.

Lastly, with Anomaly Capture, you can dynamically route data tied to an anomaly to your observability platform. If an alert is fired from Edge Delta, it will automatically ship data before, during and after the anomaly to your monitoring platform for easier troubleshooting.

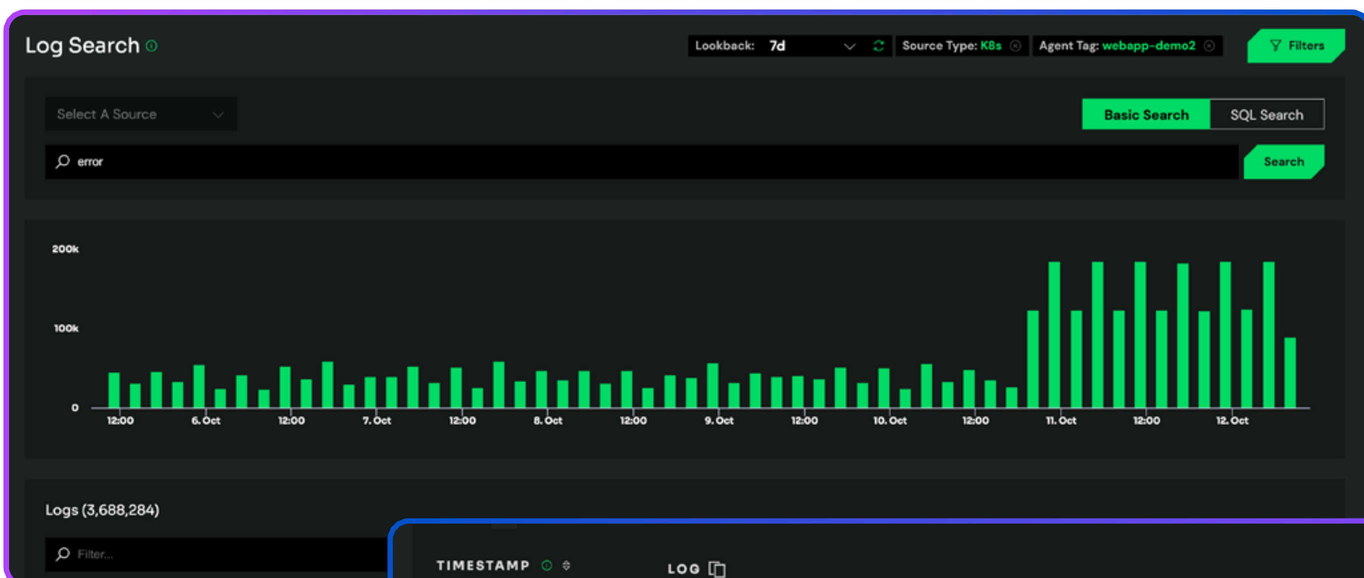

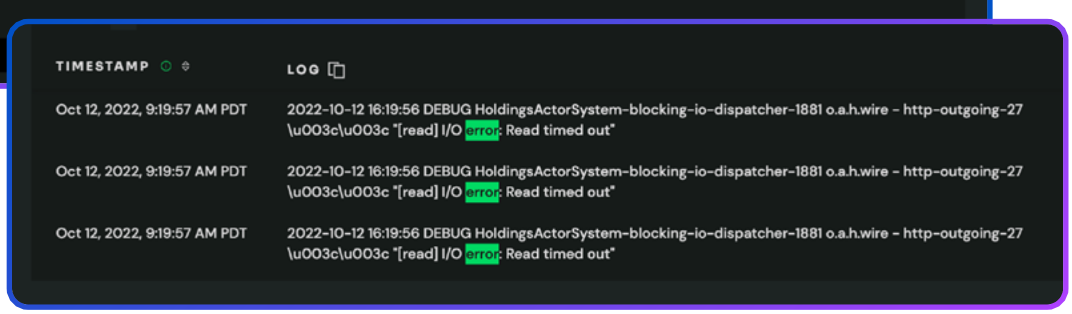

Store 100% of your log data in object storage and search on it directly within Edge Delta.

Shape Your Data Before It's Indexed

In addition to optimizing and routing your data, Edge Delta also provides several capabilities to help you shape your data before it is indexed.

Transforms can be used to format, parse, filter, sample or deduplicate log data.

Drop unnecessary fields from your loglines to reduce noise and decrease index consumption.

Mask sensitive data and Personally Identifiable Information (PII) to improve your compliance posture.

These capabilities give you an additional level of control to help you maximize the value and utility of your log data.

Get Started With Edge Delta Today

With Edge Delta, you can better understand and control your observability data to:

Ensure your team has access to the data they need.

Reduce noise within your datasets.

Get more value out of your observability license.

Want to learn more?
Chat with our experts today to start planning your potential use case.