EDGE DELTA

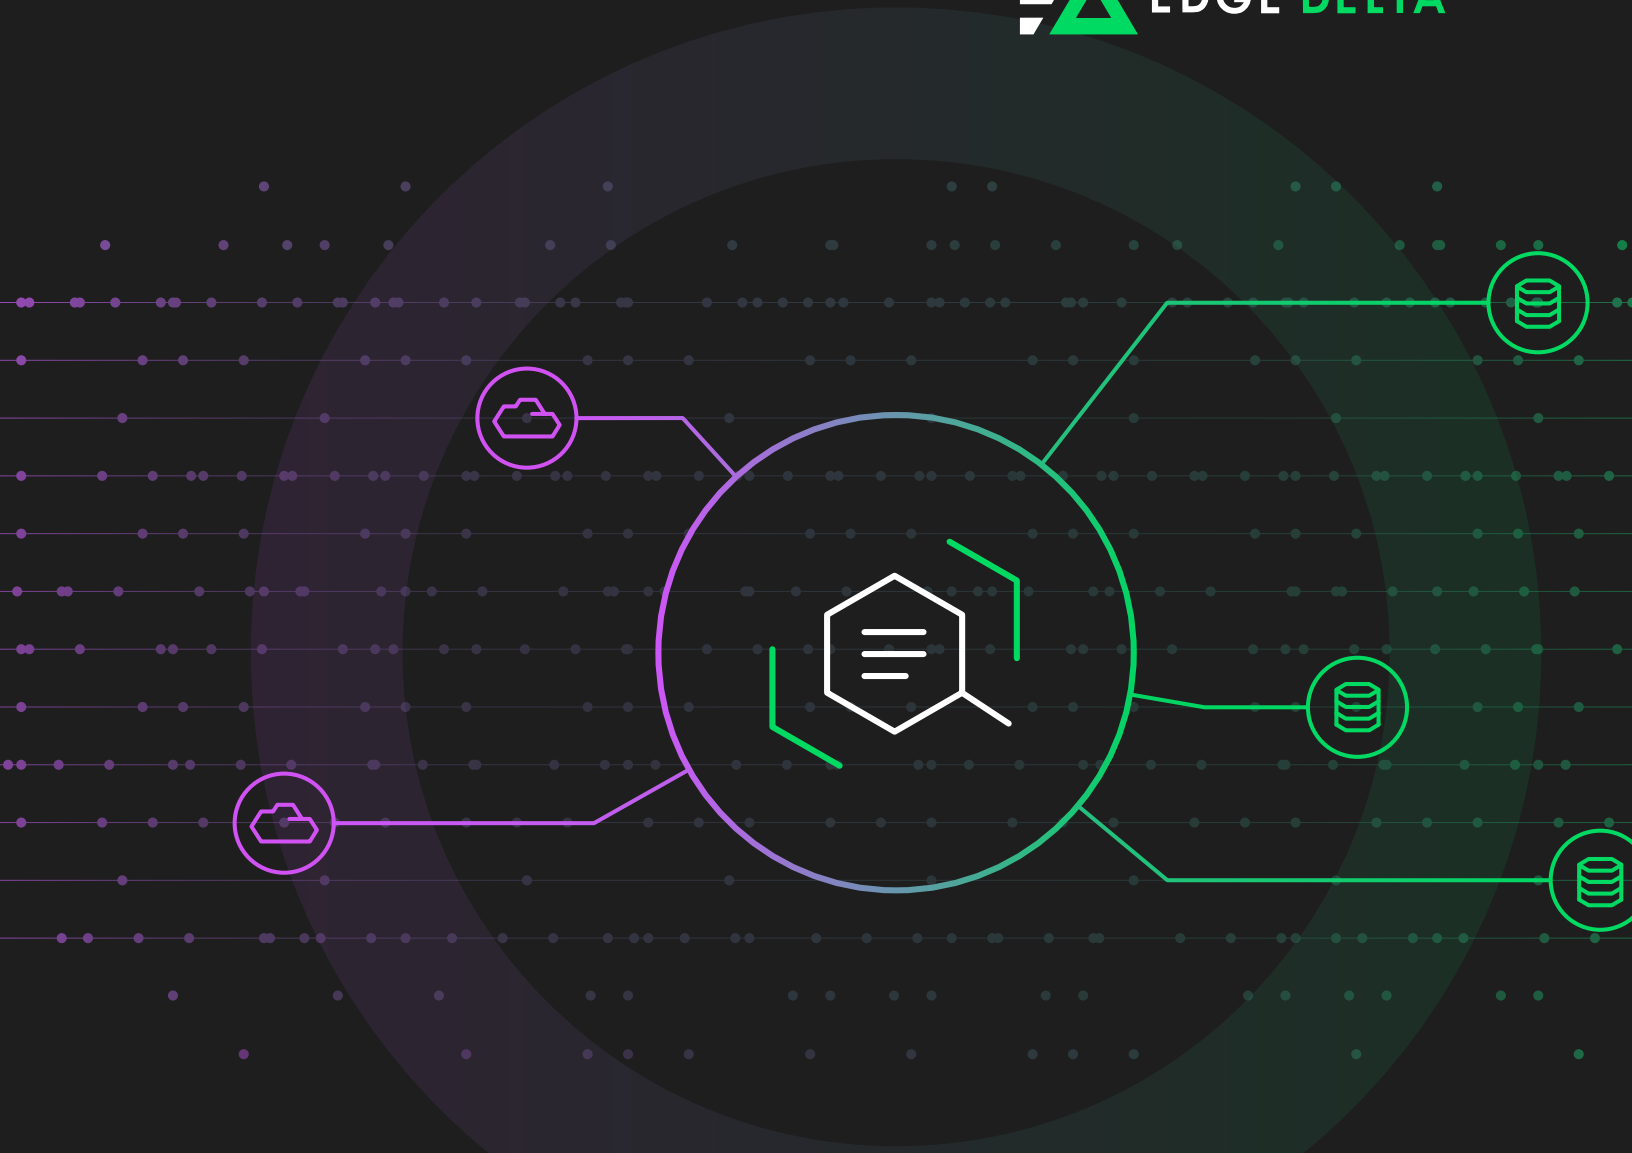

WHITE PAPER

Introducing Hybrid Log Search

Query data in your own cloud data store with the performance you need for real-time use cases.

We've reimagined logging architecture to combine the ease and scalability of SaaS with the data protection and cost of self-hosting.

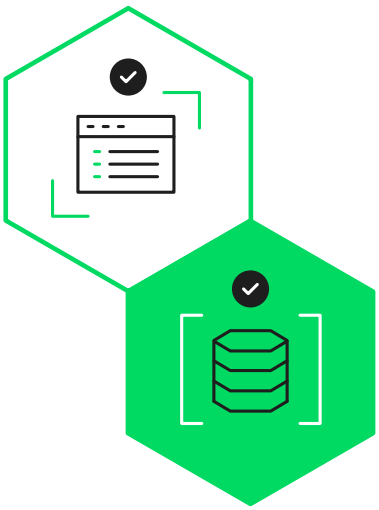

Edge Delta Hybrid Log Search empowers you to query logs directly from the cloud data store your organization owns and controls. With Hybrid Log Search, logs never need to leave your environment, and you get the performance you need for real-time monitoring and troubleshooting use cases. This product is purpose-built for teams that generate large volumes of log events with data residency and compliance requirements.

Deploying Hybrid Log Search is as simple as running a couple of commands. From there, you can query data that resides in a customer-owned cloud data store from the Edge Delta web app you're already familiar with. The only information that leaves your environment is health check data, which Edge Delta uses to help provide visibility into your Hybrid Log Search resources.

Edge Delta Hybrid Log Search is currently in private beta.

The Challenges of Self-Hosted Log Search

Historically, engineering teams have opted for self-hosted log search products for two main reasons:

- **Data Residency and Compliance Requirements:** Organizations in highly regulated industries that handle sensitive data often cannot forward logs to an external or third-party SaaS observability tool.
- **Cost-Efficiency:** Once data volumes exceed a certain threshold, the cost of centralizing observability in third-party SaaS is no longer feasible without going over budget.

Self-hosted solutions can be complex and time-consuming to deploy, maintain, and scale. They often require specialized resources to maintain their availability during peak periods of demand to avoid platform outages. Moreover, teams that generate large data volumes often sacrifice long-term visibility by reducing retention to work within resource constraints.

How Others Have Approached Hybrid

Other observability vendors have offered "federated log search" and "search in place" products to their customers. With these products, the customer can store log data in their own cloud data store and query data from a vendor-hosted SaaS backend.

These approaches have one major shortcoming: returning queries requires the user to send data from their environment to the vendor-hosted SaaS backend. This results in poor performance – shipping data from the customer environment to the vendor-hosted backend incurs latency and slows queries.

Existing federated search products help you run one-off queries at a very low cost – they are ideal for non-critical datasets. However, these aren't designed for use with frequently queried, high-volume datasets that are often essential for troubleshooting.

The Benefits of Edge Delta Hybrid Log Search

Edge Delta Hybrid Log Search combines the benefits of self-hosted and SaaS log search products. By doing so, Edge Delta provides the optimal experience for teams that work in highly regulated industries and create large-scale datasets.

- **Protect Sensitive Data:** Meet data security requirements, control where your data resides, and ensure sensitive data stays within your "four walls."
- **Reduce Observability TCO:** Leverage Edge Delta's optimized storage and eliminate egress charges to drive new levels of efficiency compared to legacy providers.
- **Minimize Ongoing Maintenance:** Receive support and consulting to reduce the burden of managing a self-hosted logging platform.
- **Adopt Open Standards:** Data is written to your storage target in the Open Telemetry (OTel) schema – not a proprietary format.

Hybrid Search Components

You can deploy Edge Delta Hybrid Search using a Terraform package and a Helm chart. When you do so, the following resources will be deployed in your environment:

- An isolated instance of the **Edge Delta Log Search service** instance, which consists of...
    - An **on-premise Edge Delta API** to communicate between the Log Search service and the other local resources
    - A **hybrid token** which is specified during installation as a secret

- Column-oriented **data store** that reads queries from the API and returns queries from an Amazon S3 bucket
- Two **Amazon S3** buckets – one for storing your log files and one for database operations
- **Prometheus** to track the health of the other local Edge Delta resources

In addition to the resources above, you will deploy the Edge Delta agent on the resources you'd like to collect data from.

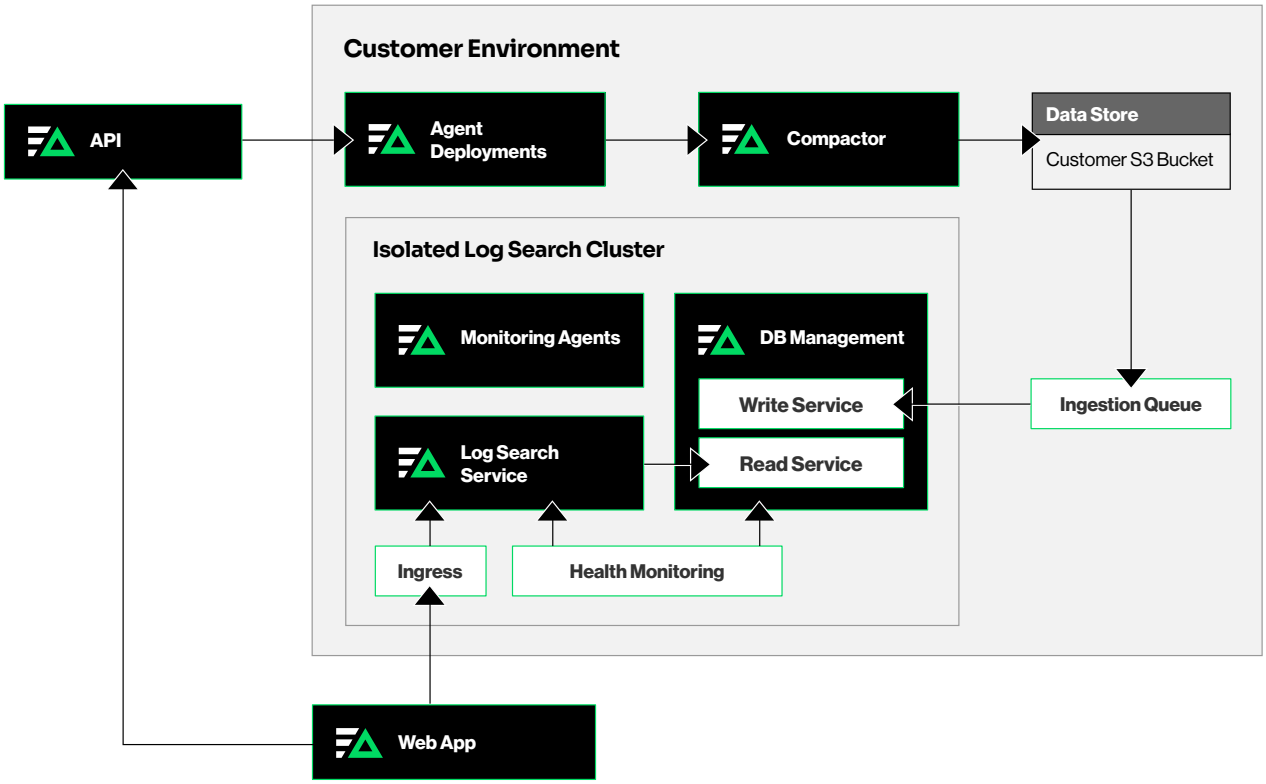

Figure 1: A reference architecture depicting the components and data flow of an Edge Delta Hybrid Log Search deployment.

How Edge Delta Keeps Your Data Private

Edge Delta runs as a single-page application in your browser, and verifies if you are a Hybrid Search customer via Edge Delta's central API on login.

This process, apart from confirming your deployment mode, only transmits health metrics of Edge Delta resources running in your environment and query statistics. Once the deployment mode has been confirmed, all subsequent requests from the user's browser are sent to the on-premise API only.

Your logs will remain in the cloud object store of your choosing (for now, Amazon S3 only). To get started, you'll need to provision two Amazon S3 buckets, one of which Hybrid Log Search uses to read and write logs and the other is used to manage data store operations.

Notification and queue services running in your environment facilitate log ingestion from the S3 buckets into the Hybrid Log Search data store. Query responses are prompted when the Edge Delta API queries this system.

Both on-premise and central Edge Delta APIs use short-term JSON Web Tokens (JWTs) with a one-hour time-to-live (TTL), aligning with Microsoft and Google's best practices.

Getting Started

Edge Delta Hybrid Log Search is currently in private beta. Participate in the beta program by **signing up here**.