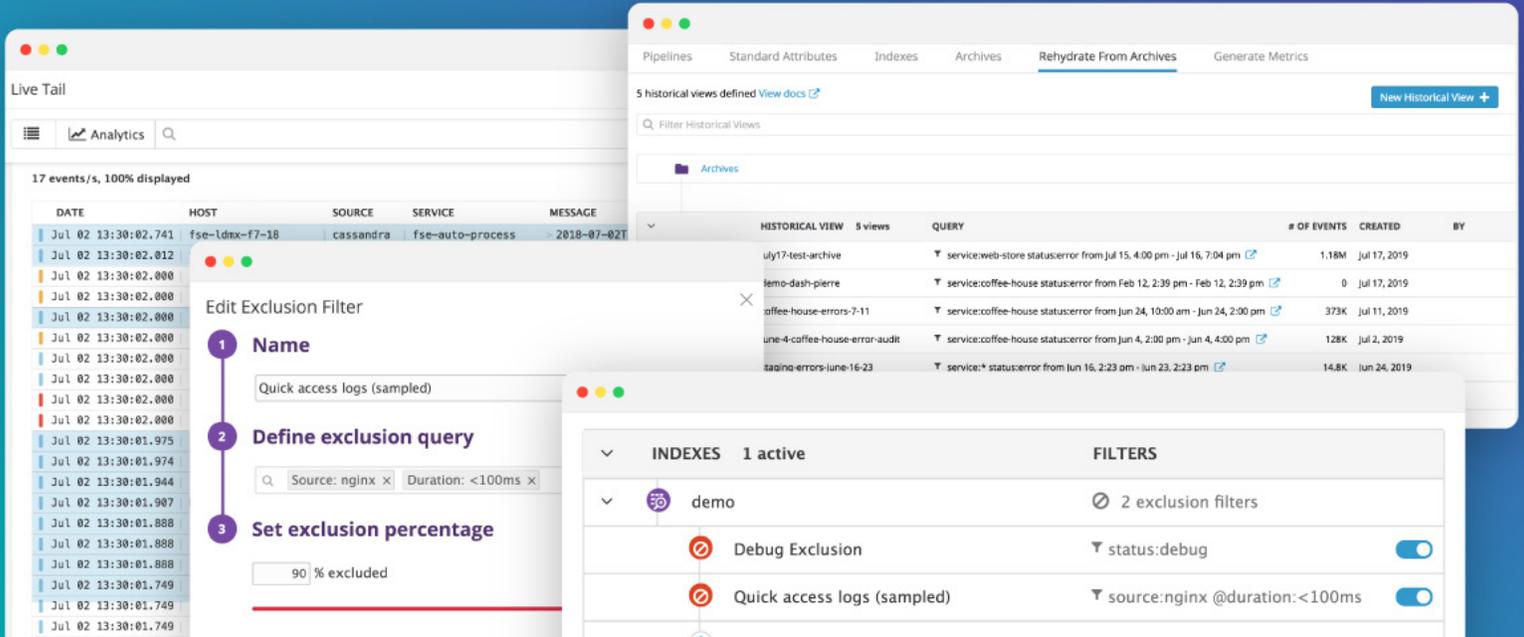


Logging without Limits™



Manage large volumes of logs cost-effectively with no trade-offs on coverage or visibility



Introduction

Organizations today need to manage petabytes of logs per day from myriad systems, applications, and cloud services. This **large volume** makes retaining all of the logs for analysis cost prohibitive. Business seasonalities and incidents such as unplanned outages, production fixes, and new deployments lead to **large variations in log volume and cost**. Some logs that are usually perceived to be low value (e.g., debug logs), can suddenly become critical for troubleshooting or analysis during a production incident. **An effective log management solution needs to address these large variations in volume and value and solve the associated pain points.**

Problems with traditional log management tools

Traditional log management vendors offer a “one-size-fits-all” solution that forces you to either:

- **Decrease the volume of logs upfront to avoid incurring high costs:** However, not having all your logs within the log management solution leads to visibility gaps that make you “blind” at the worst possible moments (e.g., production incidents, unplanned outages, etc.)
- **Accept the high cost of indexing all logs and the lack of control over variations in log volume and value:** Since you commit to a fixed pricing plan in gigabytes per day, you also get penalized for enriching your logs with contextual data that increases their size

Neither of these options are satisfactory, and the outcome is a log management solution that provides partial visibility and therefore leads to limited user adoption.

Ingest all logs, and index only when you need them

Collecting log data should be easy and affordable, and not having the data when you need it for analysis and troubleshooting can be very expensive. Our approach is based on a simple idea: **you should have access to all of your log data when you need it**. By **decoupling** the **ingestion cost** and the **indexing cost**, we solve one of the major problems in the industry related to the prohibitive cost of managing extremely large volumes of logs.

Thanks to Logging without Limits, you can:

INGEST, PARSE, LIVE TAIL & ARCHIVE ALL LOGS

- **Send all your logs** to Datadog at \$0.10/GB, thereby maintaining the full coverage of logs cost-effectively, without filtering any of your logs upfront
- **Parse all your logs at the time of ingestion** and enrich with contextual data into a structured format; this ensures fast performance since logs don't have to be processed at query time, and standardizes all of your log formats and attributes into a single taxonomy
- **Live tail** the full stream of ingested logs for real-time visibility during an outage or deployment
- **Maintain a complete archive**, with the full history of your operations for any audit or historical analysis

INDEX ONLY YOUR MOST VALUABLE LOGS

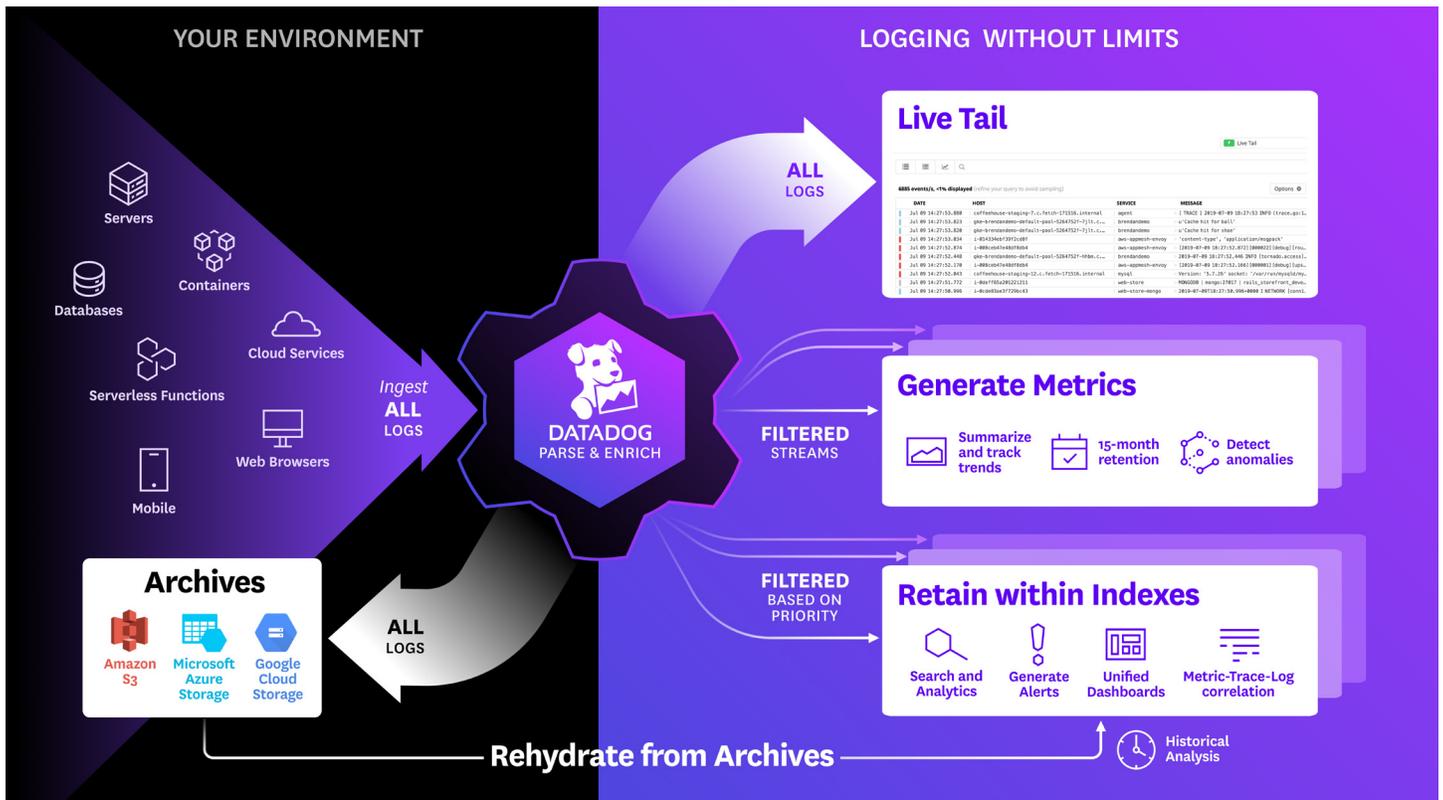
- Benefit from full, fine-grained control over indexing costs:
 - **Index only the logs that matter**, define team-based quotas upfront, and let end users prioritize which logs get retained within the Datadog platform
 - Vary the retention period per log index, based on log value or business priority
 - Create an index from any log pipeline, at any point of time, with the flick of a switch (using exclusion filters)

REHYDRATE ARCHIVED LOGS

- **Log Rehydration™** enables you to rehydrate archived logs into indexes when required for audits or analysis
 - Incentivizes end users to retain fewer logs within the platform or decrease the retention period per log index which further reduces costs
 - Empowers end users to quickly access old logs (from a few weeks or a few months ago) that become extremely valuable for a security, business, or technical audit

GENERATE METRICS FROM ALL YOUR LOGS

- **Generate metrics from all ingested logs** to track and retain anomalies and trends
 - Instead of indexing hundreds of millions of log events per day, retained for 7 or 15 days, you can retain one custom metric, **for 15 months with no rollups**
 - If you notice a sudden spike in the custom metric or anything abnormal, the associated logs can be seamlessly rehydrated from archives



Correlate logs with metrics and traces automatically

Datadog’s Log Management solution is built on the same SaaS platform as Infrastructure and Application Performance Monitoring (APM), thereby unifying metrics, traces and logs, the three pillars of observability. The entire platform is designed to be easy to use and utilizes a common tagging structure. When you notice a sudden spike in any metric on a dashboard, you can **find the related logs with a single click**.

Datadog automatically brings together all the logs for a given request and **links them seamlessly to tracing data from that same request**. When your users are encountering errors or high latency in your application, drilling down to view the logs from a problematic request can reveal exactly what went wrong. By pulling together all the logs pertaining to a given request, you can see in rich detail how it was handled from beginning to end so you can quickly diagnose the issue. With auto-instrumentation for Java, Python, Ruby, Go, Node.js, .NET, PHP, and many associated frameworks, you can start correlating logs and request traces without touching your application code.

In summary, **Logging without Limits eliminates the upfront trade-offs that force you to sacrifice log coverage to save costs**. You gain full flexibility without sacrificing visibility.





DATADOG