Introduction

As organizations embrace Datadog to modernize their observability stack, they often face the challenge of rising costs, particularly in the realm of logging. Unpredictable log volumes, unexpected changes, and vendor lock-in concerns can escalate these expenses.

Fortunately, there are several ways to optimize logging costs and performance with Datadog. This brochure provides an overview of how Apica can help you optimize your Datadog costs for managing log volumes and performance monitoring with our active observability Data Fabric solution.

Challenges

Cost Ballooning

Datadog cost optimization is a recurring issue for enterprises seeking to leverage its benefits while managing expenses effectively. The initial decision to adopt Datadog is a strategic one, allowing teams to redirect their focus from managing self-hosted observability solutions to core organizational activities.

However, the cost of Datadog, especially for logging, can quickly spiral out of control. Log volumes are notoriously unpredictable, subject to fluctuating user and system changes, and unforeseen developer alterations in code. This unpredictability often leads to ballooning Datadog expenses.

Did you know? In 2022, Coinbase, a cryptocurrency exchange platform, spent $65 million on Datadog!

Vendor Lock-Ins

In response to escalating costs, engineering teams often explore alternative solutions. However, the quest for alternatives typically reveals a challenge- Organizations are heavily reliant on Datadog's feature set, and adopting point solutions to address logging expenses can introduce fragmentation and overlaps.

Solutions

Drastically Reduce Costs and Improve Agility

Apica's Data Fabric equips engineering teams with the tools to efficiently stream relevant data to Datadog while minimizing noise. Approximately 95% of data streams tend to be irrelevant in a given context. Apica allows real-time data filtering and dynamic data attribute management to optimize the data volume sent to Datadog. Furthermore, Apica ensures data retention without loss, maintaining a master copy of all data streams in your chosen object store for fast retrieval.

This approach yields dual cost benefits- reducing Datadog's ingest and indexing costs by 70 to 95% and providing an infinitely scalable and active retention layer, eliminating the need for Datadog's two-tiered storage architecture.

Apica

Apica VS Datadog

FEATURE	APICA	DATADOG
Observability Data Storage	★★★ Low cost, infinite scale	★★☆ Unified data platform
Data Level Integrations	★★★ Full pipeline control	★★☆ 600+ integrations
Synthetic Monitoring	★★★ Best of breed ASM	★★☆ Good capabilities
Load Testing & CI/CD Integration	★★★ Best of breed ALT	☆☆☆ None
NPM	★★☆ Solid capabilities	★★☆ Solid NPM capabilities
Native Cloud Monitoring	★★☆ Good cloud functions	★★☆ Good cloud monitoring
Open Telemetry	★★★ Strong OTel functions	★★☆ Built in OTel Collector
Distributed Tracing	★★★ Strong tracing	★★☆ Good capabilities

Application and Infrastructure Optimizers

Apica identifies outliers, patterns, and anomalies, ensuring that only useful data is streamed to Datadog by utilizing AI/ML and rules-based capabilities. Simultaneously, it offers data enrichment for enhanced Datadog analytics. All data streams are parallelly indexed in Apica's operational Data Fabric, leveraging low-cost object storage.

For instance, when collecting logs from a Kubernetes cluster, a few clicks within the Apica rule pack can result in a 70% reduction in Datadog license spending and index size.

Manage Long-Term Retention and Compliance with Ease

Enterprises require a robust system capable of ingesting, storing, and retrieving data at scale and speed. Datadog's tiered storage approach limits access to older data, requiring costly rehydration, and reindexing, and incurring tenfold expenses for indexed data retrieval.

Apica's storageless architecture, built on any object storage, allows organizations to store vast amounts of data without impacting performance and reliability. Data retrieval is instantaneous.

By transitioning long-term retention storage to Apica, organizations can liberate their data and gain better cost control. Additionally, Apica offers a purpose-built automation engine for on-demand data retrieval into Datadog, saving time and money on indexed long-term retention.

Zero Vendor lock-in

With Apica's Operational Data Fabric, concerns about vendor lock-in become a thing of the past. Much like Datadog, Apica enables the on-demand activation of various platforms to meet specific business requirements.

If you need to collect all data but route security data to Splunk and developer data to Datadog, our 1-Click data route management simplifies the process extensively.

Conclusion

Apica's Data Fabric offers a holistic solution for Datadog cost optimization, addressing the challenges of ballooning expenses, vendor lock-in, data noise, and long-term retention management.

With Apica, organizations can regain control over their observability costs, reduce waste, and enhance their Datadog experience.

Optimize, streamline, and take charge of your Datadog investment with Apica today.

**The names of all other brands, product names, or trademarks are the property of their respective owners.
All rights reserved by Apica.**