



Digital commerce platform migrates terabytes of sales analytics data from self-hosted Graylog to Elastic Cloud 8.3 in 8 weeks.

2.4 terabyte
Data Migrated

9000+
Indices Migrated

2.2 Billion
documents

CLIENT

The client is an international digital commerce agency.

The client enables over 45 global fashion brands to setup and operate large scale online stores without a hitch!

AT-A-GLANCE

Challenges

- Legacy analytics system.
- No changes or impact to digital storefront.
- Lacked expertise, migration was time sensitive and to be completed before upcoming holiday sales season.

Solution

- Architected solution using AWS Load Balancer and AWS SQS for data ingestion.
- Configured AWS SQS as a buffer to store events during migration.
- Solution included migration from ES 2.x → 5.x → 7.x and finally to ES 8.3

PROJECT CONTEXT

The client had implemented self-hosted Graylog built on Elasticsearch 2.x for Real User Monitoring (RUM), Alerts and Sales Analytics.

The legacy & out-of-support solution was not scalable anymore, and the client experienced availability issues, especially during the holiday season.

PROJECT OBJECTIVES

- Migrate from Graylog to Elasticsearch 8.3 on Elastic cloud
- Migration should be carried out without any change to digital storefronts, without data loss and with zero downtime.
- 2.4 Terabyte of Historical data with over 9000 indices & 2.2 billion documents to be migrated.

SOLUTION DELIVERY

- Two weeks of solution discovery & POC to identify the best feasible solution for migration that meets project objectives.
- Removed Graylog from the existing solution and re-architected the solution to ingest current GELF events as JSON payloads in Elasticsearch 8.3.
- Two-step migration process to handle Historical & live data migration.
- Post-migration support and training was provided to enable the Live-Ops team to monitor and manage the Elastic cloud deployment.

TECHNOLOGY STACK

