CASE STUDY

SQUARESHIFT

Secure and scalable observability solution for a government department, with capacity for ingestion of 1.5 TB of data every day.

1.5 TB/day	270 TB	20
Data Ingestion	Data Retention	Nodes

CLIENT

The client is one of the largest digital identity program run by a government.

This program issues, maintains and facilitates verification of digital identity for millions of residents.

TECHNOLOGY STACK

elasticsearch

Kibana

elastic apm

fluentd

kafka

DATA CENTER

PROJECT CONTEXT

The client decided to move away from exiting legacy monitoring solution. Given the sensitive nature of the data, the data migration was carried out by in-house team.

The in-house team required expert consultation & engineering support for deployment of on-premise Elasticsearch for observability and SIEM (Security Information and Event Management).

PROJECT OBJECTIVES

- Deploy Elasticsearch observability and security for infra monitoring, APM, SIEM, ML based anomaly detection.
- Capacity planning for 1.5 TB data ingestion everyday
- Capacity planning and storage optimisation for 270 TB data retention across hot, warm, cold and frozen tears.
- High availability configuration that spans across two data centers located in two different cities in India.

SOLUTION DELIVERY

- SquareShift delivered the entire project with onsite team, given the sensitive nature of data and data security requirements of the GoI.
- Deployed Elasticsearch to meet the requirements and developed bespoke dashboards and reports in Kibana.
- Deployed Kafka Mirror Maker setup to ensure both clusters are fed with information from both data centres.
- Performed high-volume benchmark tests to test the capacity of the cluster.