

Seamless Migration to Elastic Cloud: Empowering Scalable and Cost-Efficient Search and Analytics for a Leading IT Services Provider

1,000+ Indices and 7TB of Data Migrated

Downtime limits ensuring business continuity.

Improved Performance, Enhanced Data Access, and Future-proof Scalability.

CLIENT

The client is a leading IT services and integrator they provide licenses and knowledge around the core infrastructure like Cisco Networking, VMware, Veeam – and managed services

PROJECT BRIEF

The clients existing Elasticsearch setup managed approximately 50-60 GB of data ingestion daily, totaling 7TB. With 55 agents and 11 agent policies, the migration required careful planning to minimize downtime and maintain data integrity. The project was executed in phases with specific milestones to meet operational goals, including snapshot restoration, ILM policy updates, and configuration migration

RESULTS

- Security and agent configurations were replicated accurately, maintaining operational continuity.
- The new cloud deployment offered scalable storage and optimized costs through tiered ILM policies.
- The new cloud environment positions clients to scale efficiently as data grows.

PROBLEM STATEMENT

1. Data Volume and Complexity:

Managing the migration of over 1,000 indices and handling 7TB of historical data, including system indices and live ingestion, required a robust strategy.

2. Configuration Migration:

Critical settings such as security configurations, agent policies and index lifecycle management (ILM) required precise replication in the cloud environment.

3. Infrastructure Alignment:

The cloud deployment needed to support current and future data storage requirements, particularly in optimizing frozen tier storage for cost efficiency.

SOLUTION DELIVERY

1. Test and Validation:

- A test cloud deployment was created to validate the architecture and configuration before full-scale migration.
- Snapshots of system indices were restored in the test environment to ensure compatibility.

2. Phased Snapshot Restoration:

- **Phase 1:** Migrated system indices (e.g., .kibana, .security) to establish foundational settings.
- **Phase 2:** Migrated historical and live data.

3. Custom Scripting for Configuration:

- Security, fleet, and space configurations were migrated using custom scripts to ensure seamless integration.

4. ILM Optimization:

- Updated ILM policies to align with the cloud's architecture, optimizing hot and frozen tier retention periods.

5. Downtime Management:

- Migration activities were carefully planned to fit within the client's acceptable downtime window of 4-6 hours.

6. Scalable Cloud Deployment:

- The cloud deployment was sized to accommodate future data growth, with an emphasis on frozen tier storage for historical data.