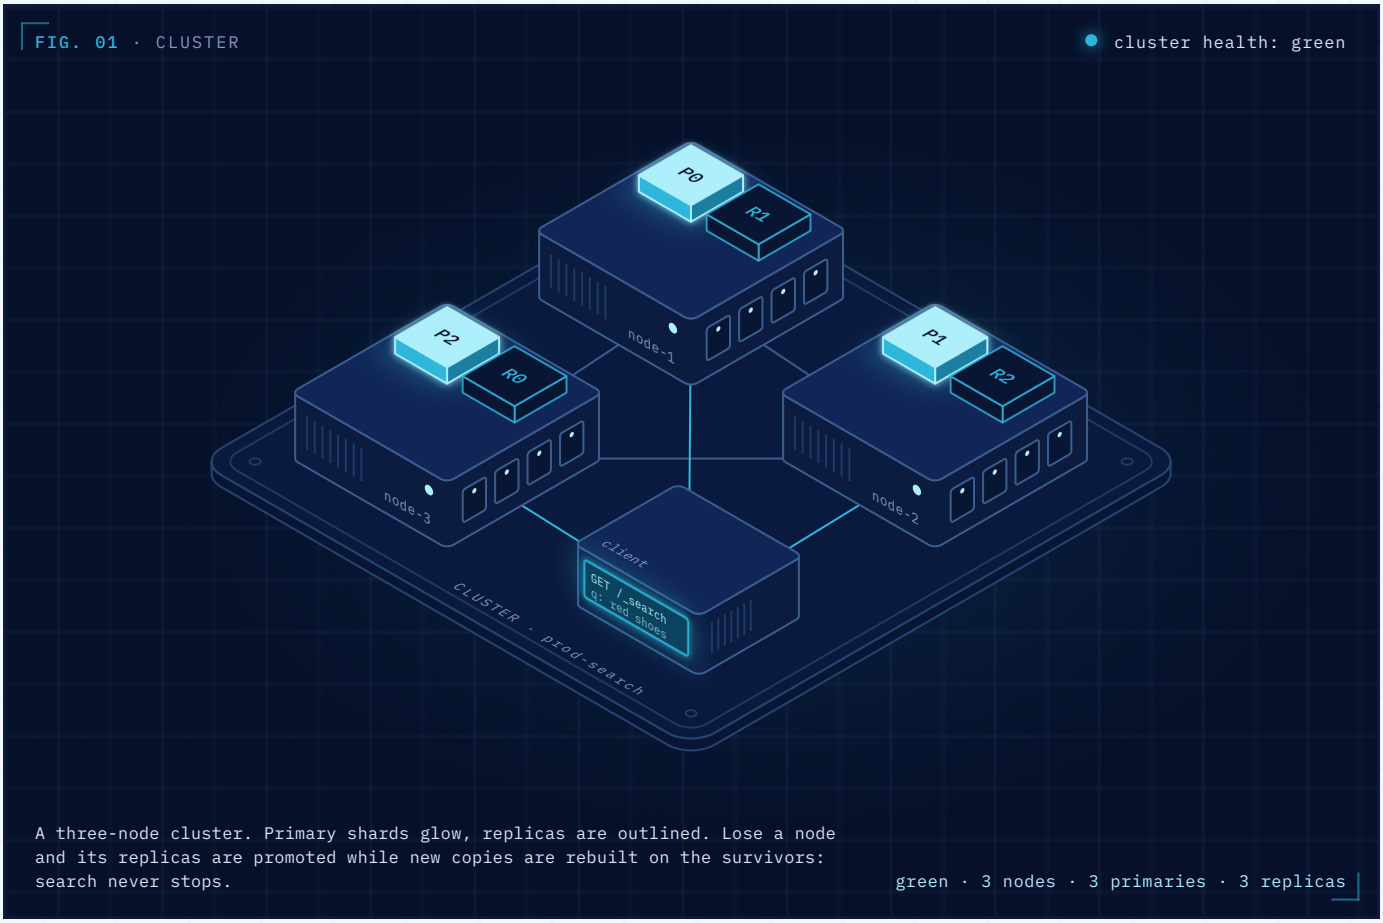
TRAINING & CONSULTING · SEARCH

Elasticsearch you can *trust* in production.

Consulting and training for teams that run Elasticsearch and OpenSearch. Our engineers have designed, tuned and rescued search clusters since 2014 — for banks, public institutions, defence and telecom teams and large e-commerce catalogues — and trained the people who keep them running.



2014
Running Elasticsearch in production since the 1.x releases

3
Programmes, from the first index to relevance engineering

≤15
People per class, with a cloud lab for each of them

EN · TR
Delivered in English or Turkish, remote or on-site

CONSULTING

Fix what hurts now. Health checks, relevance, performance and migrations — done with your team, on your data, measured before and after.

TRAINING

Leave a team that can run it. Programmes built from field work, with labs on realistic data and your own use cases.

SUPPORT

Stay close after go-live. Retained hours for reviews, upgrades and incidents, with response times agreed up front.

02 · CONSULTING

Hands on your cluster, *measured* on your traffic.

Every engagement starts read-only: mappings, shards, heap, queries and logs come first. Fixes are agreed with your team, applied in small steps and checked against the numbers that matter to you.

01

Cluster health check

Shards, heap, mappings, ILM, snapshots and security reviewed against your real traffic.

YOU GET

A findings report ranked by risk and effort, with the exact settings to change.

02

Search relevance

Analyzers, Turkish stemming, synonyms, boosting and query rewriting — judged on your own queries.

YOU GET

A judged query set and a before/after scorecard (nDCG, zero-result rate).

03

Performance & capacity

Slow queries profiled; shard sizing, hot–warm–cold tiers, JVM and GC tuning, bulk indexing throughput.

YOU GET

A sizing model for the next 12–24 months and tuned settings.

04

Architecture & migration

New designs, major-version upgrades, Elasticsearch ↔ OpenSearch moves, SQL or Solr to search.

YOU GET

A tested runbook with a zero-downtime cut-over and a rollback plan.

05

Vector & hybrid search

Dense vectors and kNN added to your index and fused with BM25 using RRF — measured, not assumed.

YOU GET

A working prototype benchmarked against your current ranking.

06

Observability & support

Monitoring, alerting, Logstash, Beats and ingest pipelines, plus on-call runbooks.

YOU GET

Dashboards and alerts your team understands, and retained support hours.

SIGNALS IT IS TIME TO CALL

- The cluster turns yellow or red under load
- Heap stays above 75% and GC pauses keep growing
- Queries slow down as the indices grow
- Users cannot find products or documents that exist
- Thousands of tiny shards — or a few enormous ones
- A major-version upgrade keeps getting postponed

HOW AN ENGAGEMENT RUNS

01 **Scoping call** — 45 min · free
We hear the problem, agree goals and the numbers success is judged by.

02 **Assessment** — 1–2 weeks
Read-only access, diagnostics, logs and short interviews. Findings in writing.

03 **Fix & tune** — 2–6 weeks
Changes made with your team in small steps, each measured before and after.

04 **Hand-over** — ongoing
Runbooks, a focused workshop for your engineers and optional retained support.

FIELD EXPERIENCE: Banking & finance · Defence · Public sector · Telecom · Retail & e-commerce · Media · SaaS

03 · TRAINING

Programmes built from *field* work.

Short explanations, then labs on a cloud cluster per participant. Content, length and schedule are tailored to your team, and Elasticsearch programmes are delivered by certified engineers.

PROGRAMME 01 · FOUNDATIONS

Elasticsearch Fundamentals

4 days · 24 hours

LEVEL Beginner – intermediate

PREREQ. Basic Linux

- Stack overview: Elasticsearch, Kibana, Logstash, Beats
- Data modelling, mappings and CRUD
- Installation and configuration essentials
- Query DSL: term, match, bool, filters, sorting
- Text analysis: analyzers, tokenizers, filters
- Shards, replicas and high availability
- Index and cluster management, cluster health
- Performance: caching and query profiling
- Kibana dashboards and Dev Tools
- Ingest pipelines and Logstash
- Security (TLS, roles, audit) and monitoring

PROGRAMME 02 · ADVANCED

Elasticsearch Advanced

5 days · 30 hours

LEVEL Advanced

PREREQ. Fundamentals or equivalent

- Multi-node clusters and hot–warm architecture
- Docker and Kubernetes deployments
- Bulk ingestion and large data moves
- Deep Query DSL: fuzzy, wildcard, regex
- Custom analyzers, synonyms, phonetics, n-grams
- Shard rebalancing, cross-cluster replication
- Rollover, ILM and segment merging
- JVM tuning, heap and garbage collection
- Query cache, profiling and reports
- Kibana Lens, Canvas and ML for time series
- RBAC, encrypted snapshots, Prometheus & Grafana

PROGRAMME 03 · MASTERCLASS

Search Relevance Masterclass

2 days · 12 hours

LEVEL Intermediate – advanced

PREREQ. Hands-on search experience

- BM25, IDF and why the inverted index survives
- Precision, recall and desirability
- Judged query sets: nDCG, MRR, session metrics
- Query understanding: segments, entities, scoping
- Expansion and relaxation — and when they hurt
- E-commerce relevance: exact, substitute, complement
- Seasonality, context and personalisation
- Semantic and hybrid retrieval, rerankers
- Running relevance as a weekly practice

DELIVERY
Remote or on-site; daytime, evening or weekend sessions

LABS
A cloud lab for every participant, on realistic data

CLASS SIZE
Up to 15 people, so everyone gets hands-on time

LANGUAGE
English or Turkish, with all materials provided

TAILORING
Content, length and examples fitted to your stack

NEXT STEP

Start with a 45-minute scoping call.

Tell us what is slow, wrong or about to grow. We reply within one working day with a suggested next step — consulting, training or both.

info@scaleyazilim.com

scaleyazilim.com/services/elasticsearch

Ankara · Türkiye — remote or on-site, worldwide