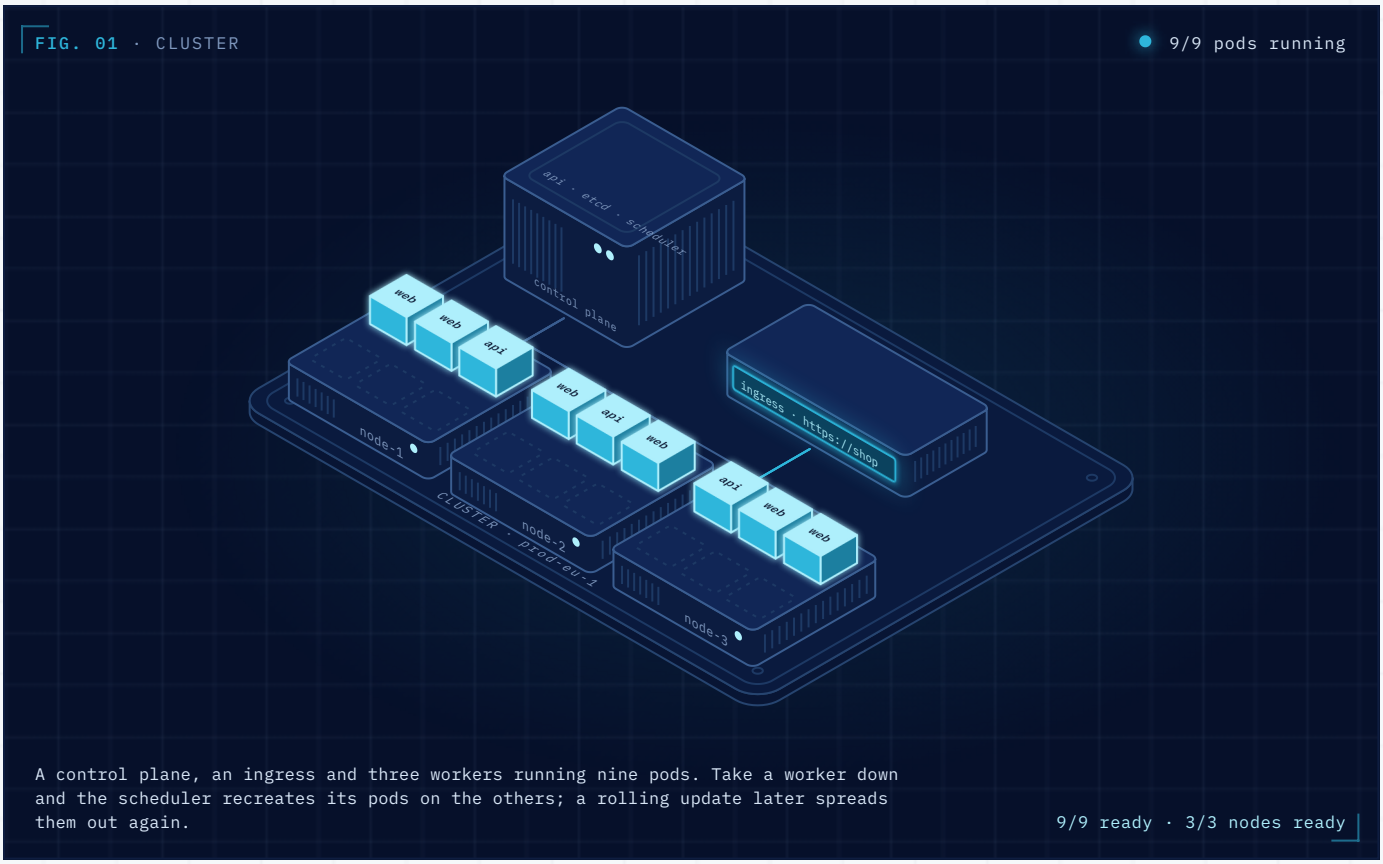
TRAINING & CONSULTING · PLATFORM

Kubernetes your team can *run*, not just install.

Consulting and training for platform teams running Kubernetes on-premises and in the cloud. We design clusters, harden them, make deployments boring and train the people on call — including running stateful systems such as Elasticsearch, Kafka, Redis and Qdrant on Kubernetes.



On-prem

and cloud: kubeadm, RKE2, K3s, EKS, GKE, AKS

4 days

From the first cluster to monitoring and CI/CD

≤15

People per class, with a cluster for each of them

EN · TR

Delivered in English or Turkish, remote or on-site

CONSULTING

Make the platform dependable. Design, hardening, delivery and observability — built as code with your team.

TRAINING

Confidence on call. Core objects, security, scaling and troubleshooting, practised on real clusters.

SUPPORT

Upgrades without fear. Version upgrades, reviews and incident help with response times agreed up front.

02 · CONSULTING

Clusters that are secure, observable and *boring*.

Most Kubernetes trouble is not Kubernetes: it is unclear ownership, hand-made changes and missing limits. We bring structure — code, policies, dashboards — and leave a team that understands why each piece is there.

01

Cluster design & setup

kubeadm, RKE2, K3s or managed; CNI networking, storage classes, ingress and multi-cluster layouts.

YOU GET

A reproducible cluster built from code.

02

Security hardening

RBAC, Pod Security Standards, NetworkPolicies, secrets, image scanning and admission policies.

YOU GET

Findings mapped to CIS benchmarks, and the fixes applied.

03

Delivery & GitOps

Helm and Kustomize, Argo CD or Flux, CI/CD pipelines, progressive rollouts and rollback.

YOU GET

Deployments that are reviewed, repeatable and reversible.

04

Observability

Prometheus, Grafana and Alertmanager; logs with ELK, OpenSearch or Loki; SLOs.

YOU GET

Dashboards and alerts that point to causes.

05

Stateful workloads

Operators for Elasticsearch, Kafka, Spark, Trino, Redis and Qdrant; storage, backups and upgrades.

YOU GET

Data services that survive the loss of a node.

06

Cost & capacity

Requests and limits, HPA and VPA, cluster autoscaling, node pools and spot capacity.

YOU GET

A smaller bill without OOM kills.

SIGNALS IT IS TIME TO CALL

- Pods get OOMKilled or sit in Pending
- Upgrades are postponed because nobody knows what breaks
- Everyone is cluster-admin
- Deploys are kubectl apply from someone’s laptop
- Alerts fire constantly and nobody acts on them
- Databases on Kubernetes have never had a restore test

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 · SaaS · Retail & e-commerce

03 · TRAINING

From the first pod to *production* practice.

Four days for teams new to Kubernetes, and a five-day follow-on for running data platforms on it. Every topic is practised on a cluster per participant.

PROGRAMME 01 · CORE

Kubernetes

4 days · 24 hours

LEVEL Beginner – intermediate

PREREQ. Basic Linux and containers

- **Day 1 · Foundations.** Docker and containers, microservices, Kubernetes architecture (control plane, workers, pods, ReplicaSets, Services); installing with Minikube, kubeadm, K3s, Kubespray, RKE2 or Rancher; the Dashboard
- **Day 2 · Workloads.** Pods and inter-container communication, stateful and stateless apps, Deployments, Services, ConfigMaps and Secrets, updates and rollback
- **Day 3 · Scale & security.** Logging and debugging, scaling and load balancing, RBAC, pod security, NetworkPolicies, monitoring
- **Day 4 · Operations.** Prometheus and Grafana, ELK logging, alerting, CI/CD integration, Helm packages and deployment automation

PROGRAMME 02 · DATA PLATFORMS

Big Data on Kubernetes

5 days · 30 hours

LEVEL Intermediate – advanced

PREREQ. Programme 01 or equivalent

- Batch and stream concepts; HDFS and Hadoop on a modern platform
- Hive for SQL on large data sets
- Kafka for real-time ingestion
- Spark clusters on Kubernetes: installation and configuration
- Kafka and Trino clusters on Kubernetes
- Managing Spark, Kafka and Trino with operators
- Storage, backups and upgrades for stateful services

WHO IT IS FOR

- DevOps and platform engineers
- System administrators moving to containers
- Data teams running Spark, Kafka or search on Kubernetes

WHAT YOUR TEAM LEAVES WITH

- A cluster they can install and operate with confidence
- Applications deployed, scaled, secured and rolled back safely
- Monitoring with Prometheus and Grafana, delivery with Helm and CI/CD

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/kubernetes

Ankara · Türkiye — remote or on-site, worldwide