



# The Converged Cloud-Native Platform

Pure open-source. AI-native. Zero lock-in.  
Infrastructure, data, AI, and communication — deployed in  
hours, not years.

PURE OPEN SOURCE

AI-NATIVE


CNCF MATURE

COMMUNITY = ENTERPRISE

openova.io · info@bankdhofar.com · 2026


# The \$6.15 Trillion Question

## GLOBAL IT SPEND — 2026 (GARTNER)



Enterprises spend **\$6.15 trillion** on IT in 2026, yet **75¢ of every dollar** feeds fragmentation — patching, licensing, vendor management, and technical debt. The market is punishing legacy vendors at historic scale, and those vendors pass the pain to customers too slow to escape.

## MARKET DESTRUCTION — Q1 2026

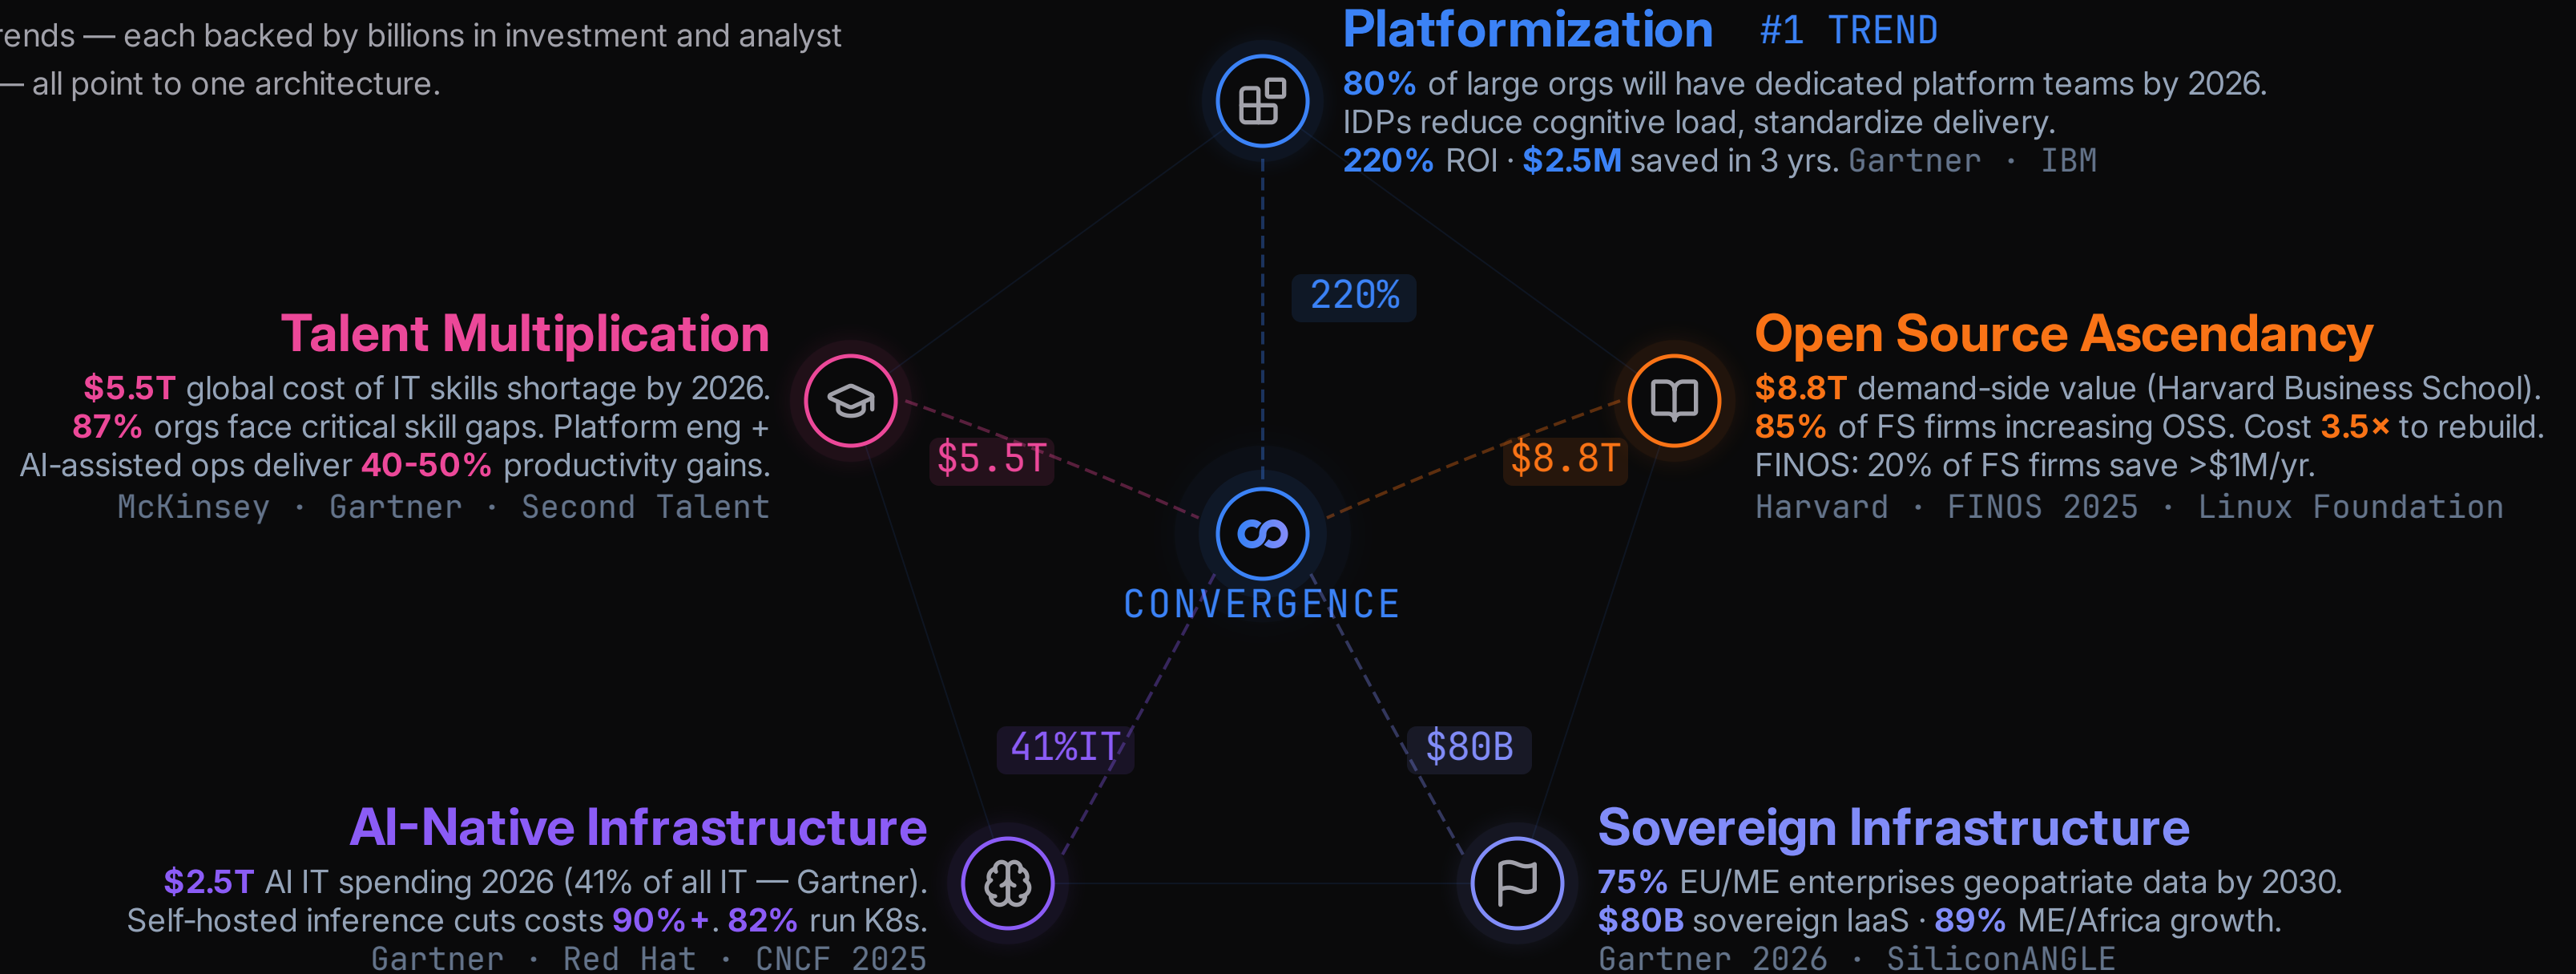


The technical response is **platformization** — converging fragmented tools into a single open-source stack. Open alternatives now outperform proprietary at a fraction of the cost. The shift from maintenance to innovation starts here.

- 75%** **The Maintenance Tax**  
75¢ of every IT dollar goes to maintenance, not innovation.
- 72%** **Infrastructure Blocks AI**  
CIOs say fragmented infra is the #1 barrier to AI ROI.
- 94%** **Modernization Fails**  
Core modernization projects exceed timelines and budgets.
- 95%** **GenAI Pilots Fail ROI**  
Fragmentation makes convergence impossible for AI pilots.
- 53%** **Software Licenses Wasted**  
Over half of enterprise software licenses sit completely unused.
- 27%** **Cloud Spend Wasted**  
\$450B+ global cloud spend — over a quarter is pure waste.
- 85%** **Platformization Savings**  
Converged platforms reduce operational overhead by up to 85%.
- 15x** **Vendor Lock-in Tax**  
Legacy virtualization vendor post-acquisition license hikes.

# Five Megatrends. One Convergence.

Five megatrends — each backed by billions in investment and analyst consensus — all point to one architecture.



# Building Blocks — Concept View

Ten building blocks, one converged ecosystem. Each solves a distinct infrastructure challenge — independently adoptable, collectively transformative.

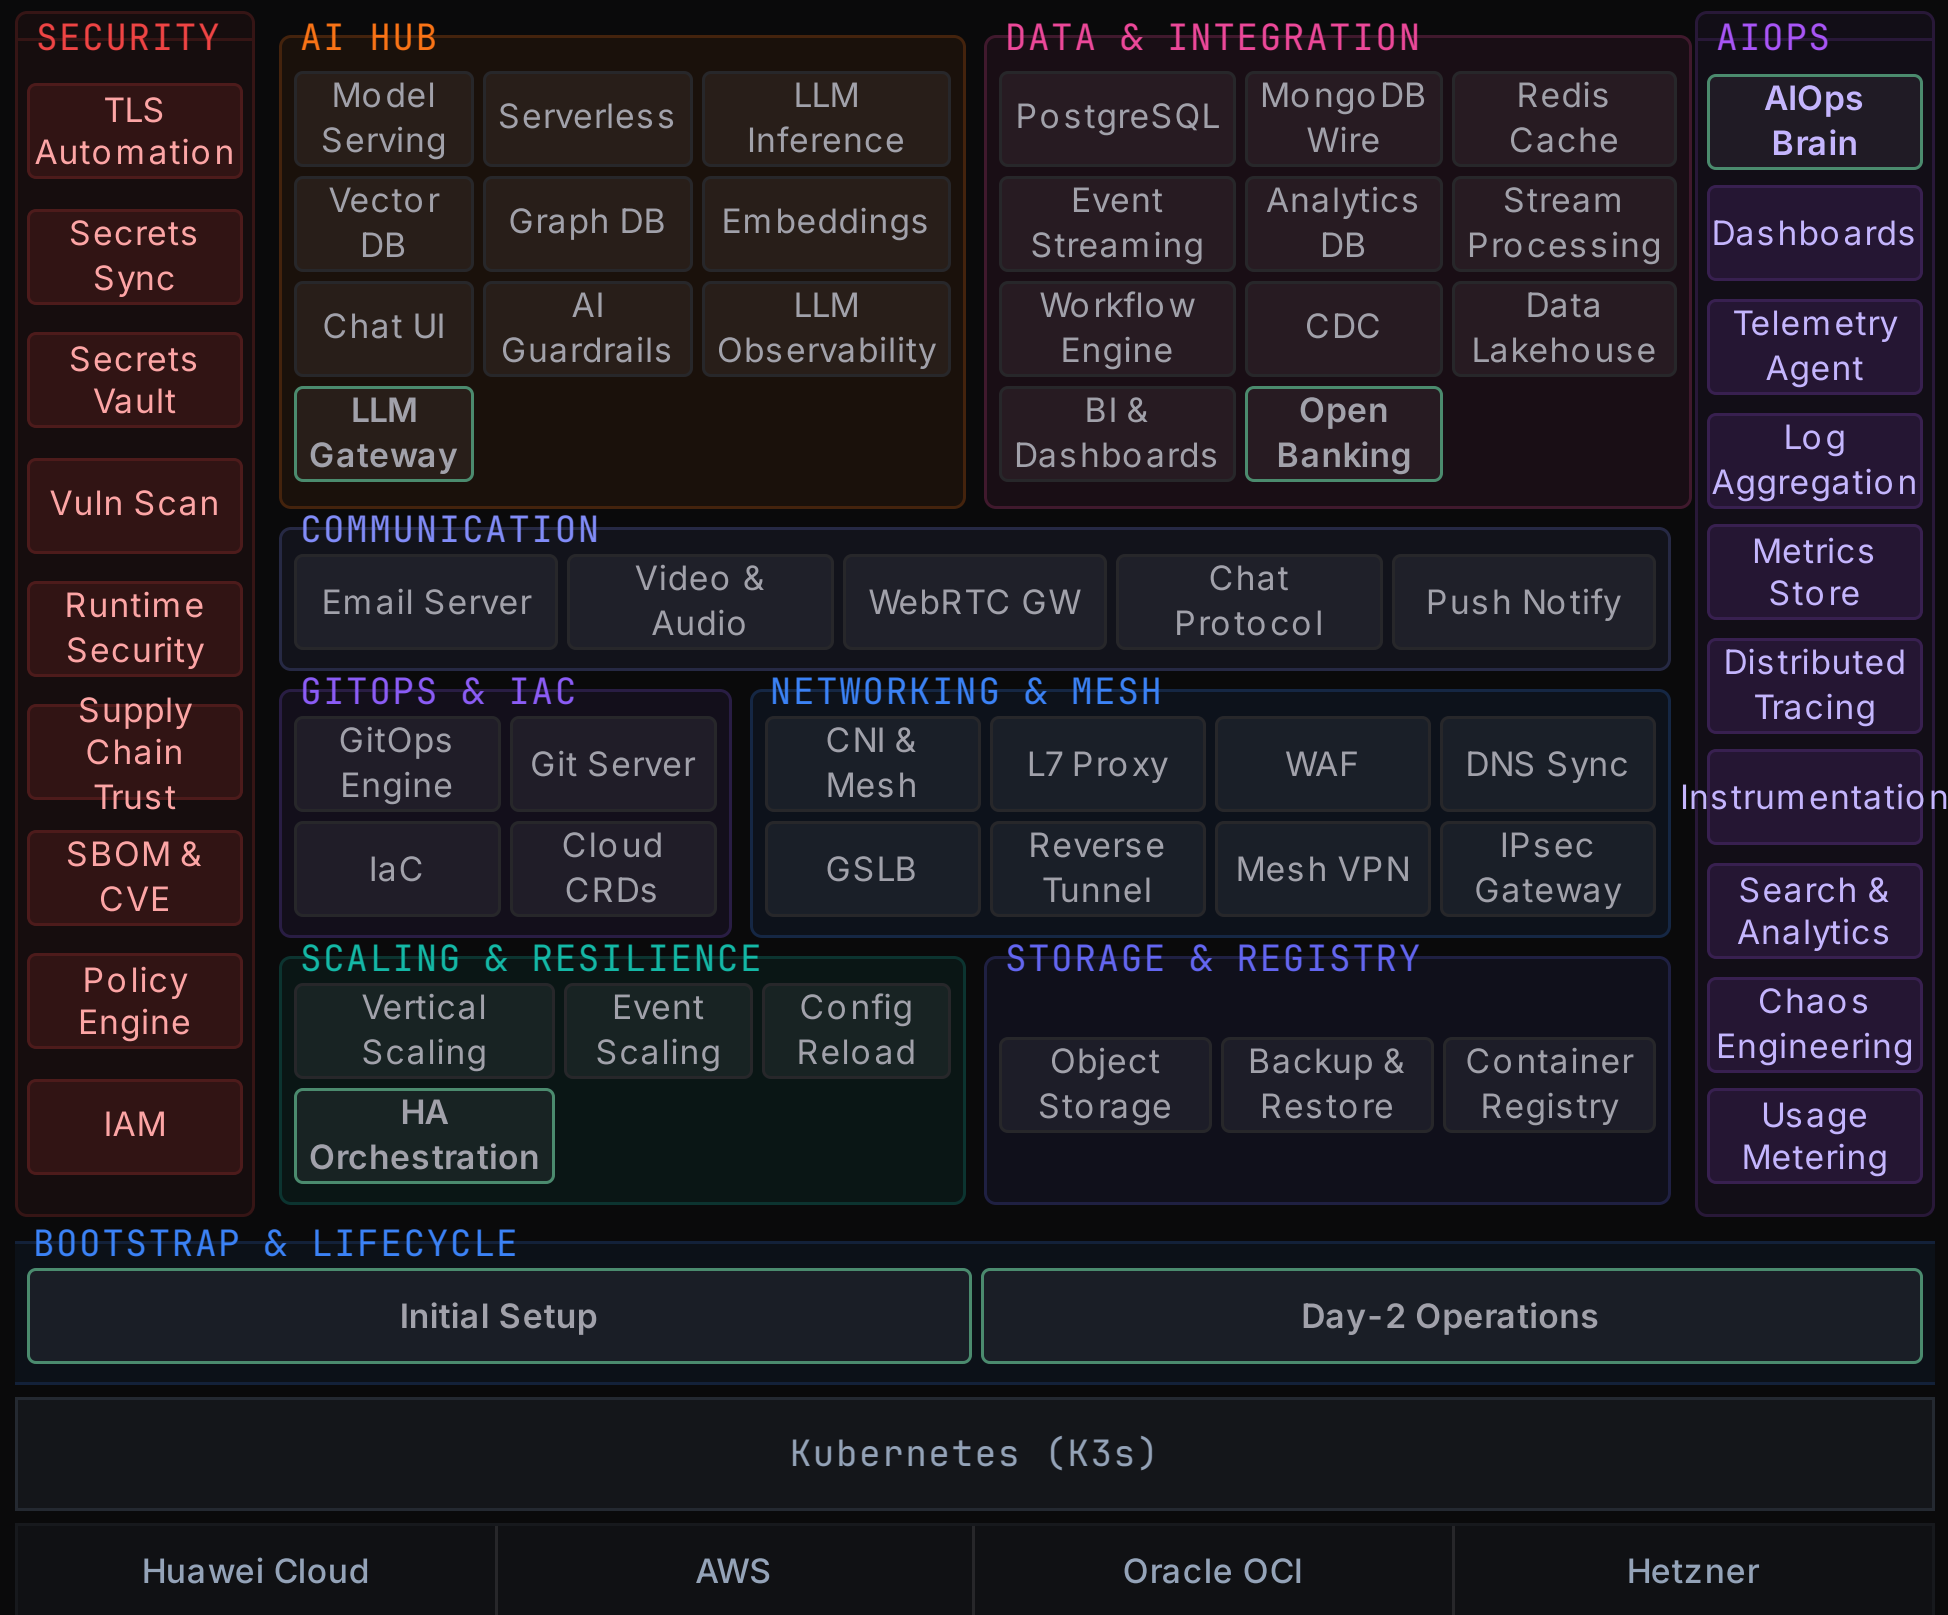
**Ten domains, one control plane.** Each domain is a self-contained building block with its own operator, CRDs, and lifecycle — sharing a single identity layer, observability pipeline, and policy engine.

**Adopt one, adopt ten — your pace.** Every block is independently deployable. Start with the domain that hurts most; existing tools plug in alongside — nothing gets ripped out on day one.

**Container-native, VM-aware.** Runs as Kubernetes operators. For legacy workloads, Cloud CRDs manage VMs through the same GitOps workflow.

**Declarative from edge to cloud.** Every resource is declared in Git and reconciled by Flux. Drift detection is automatic. Rollbacks are a single commit.

**Multi-cloud by default.** Runs identically on any cloud, bare-metal, or air-gapped on-prem — same manifests, same operators. Sovereign deployment is a configuration choice.



# Building Blocks — Tech View

Every component is real upstream open source — CNCF-graduated, battle-tested, community-backed. No proprietary forks. No rebranded wrappers.

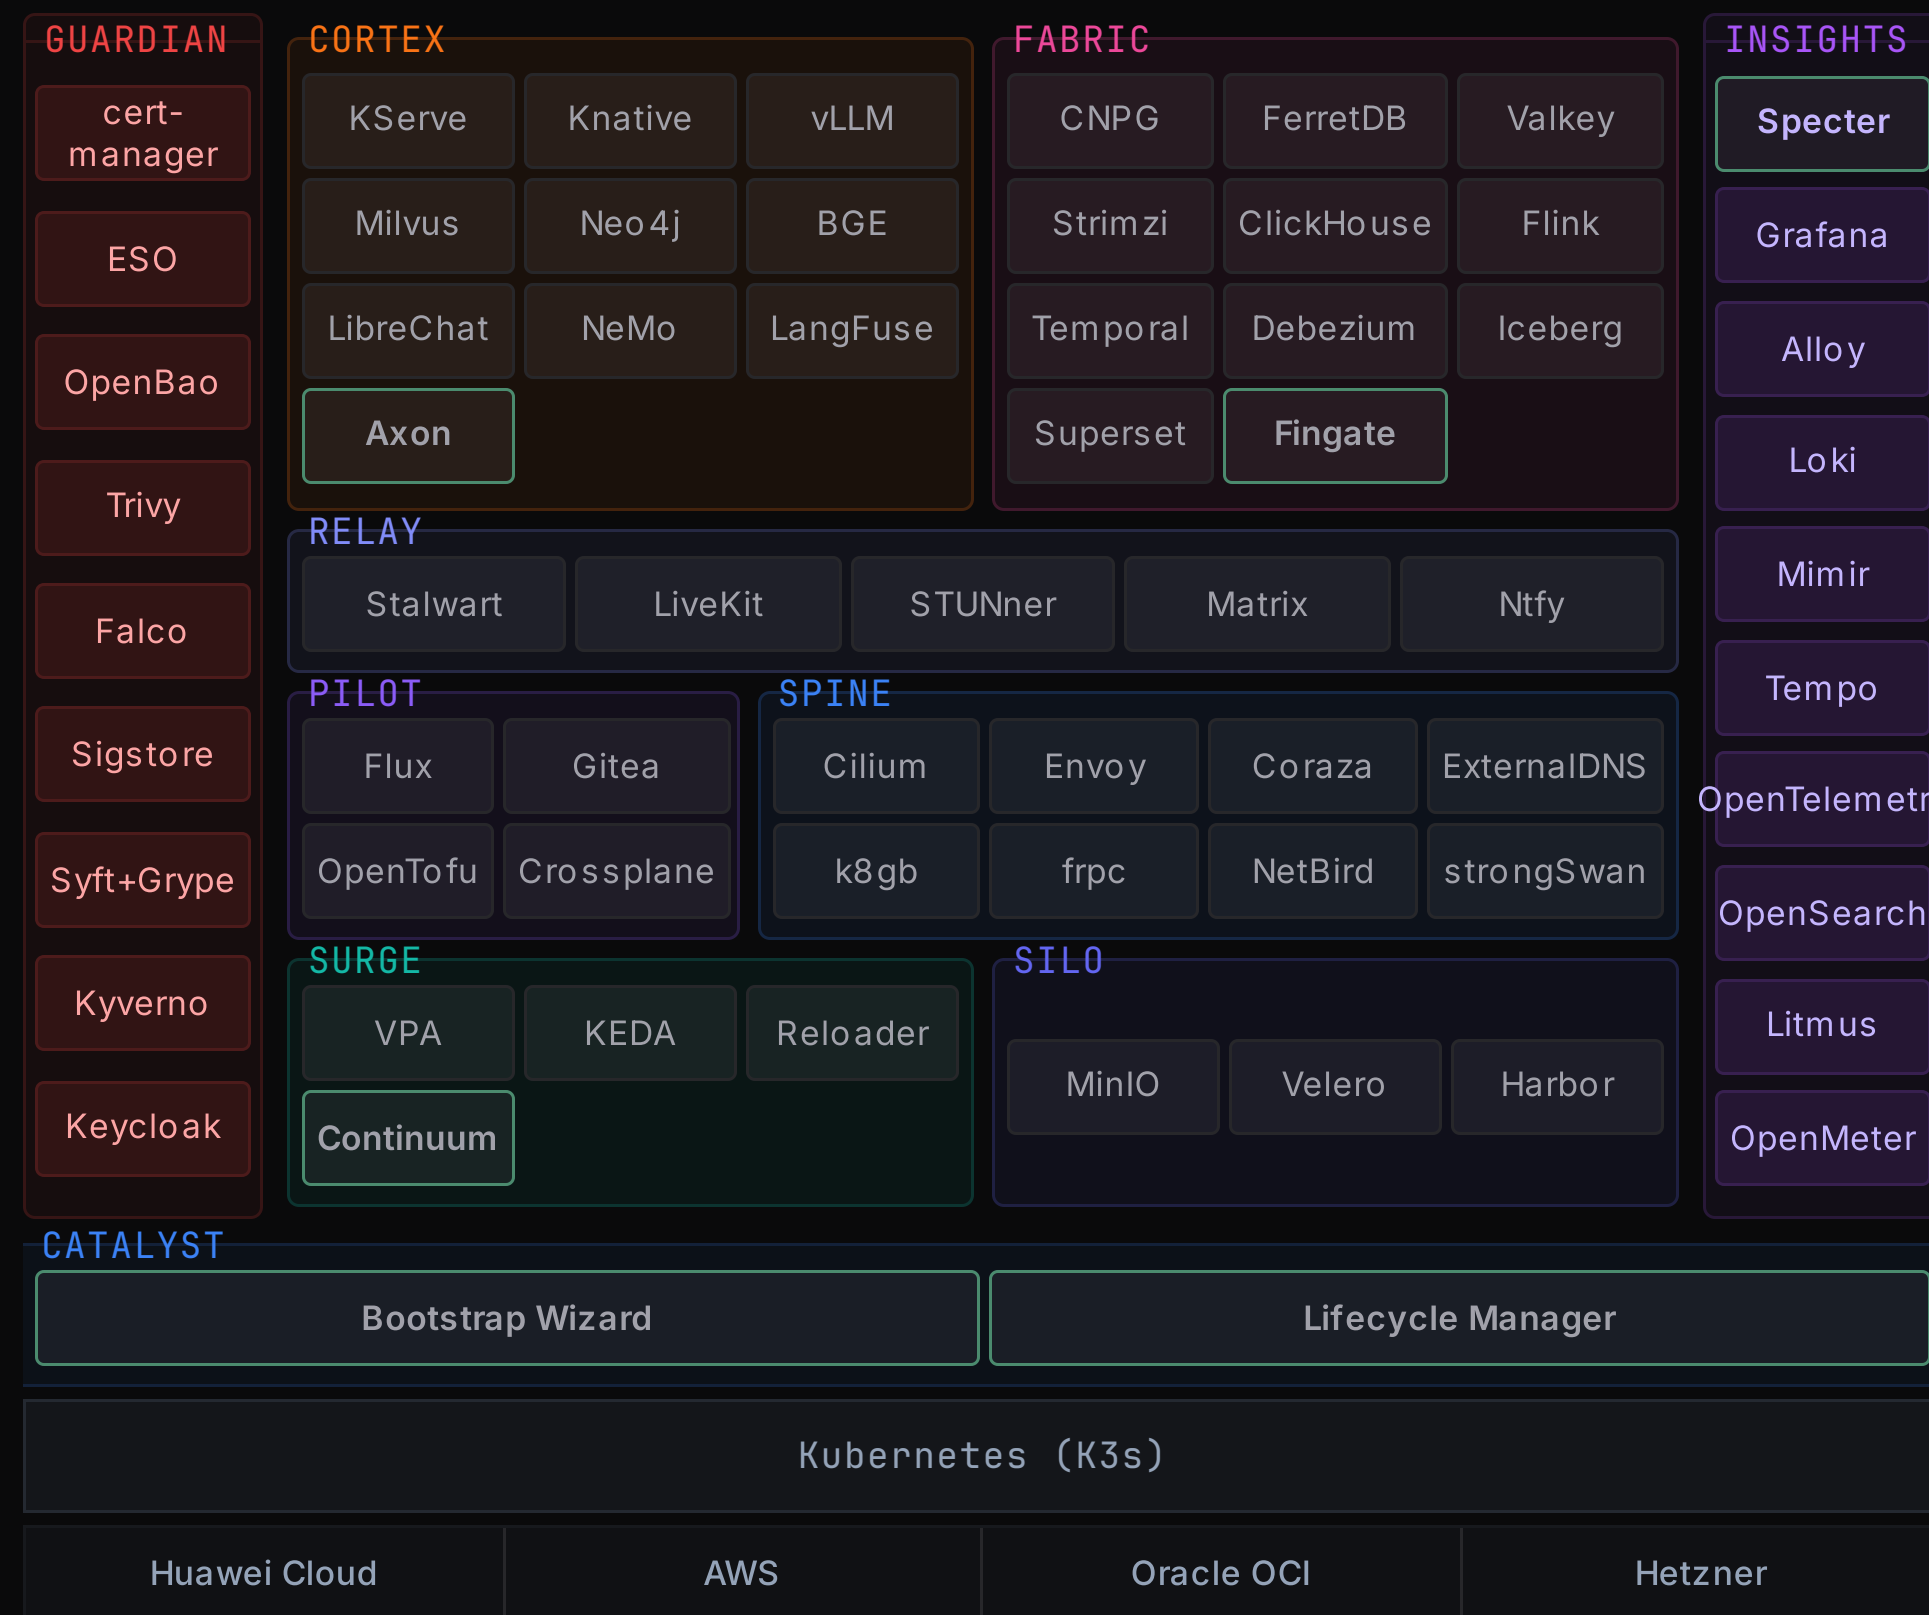
**Real upstream, not rebranded forks.** Every component — Cilium, Falco, Flux, Grafana, Strimzi, KServe, Keycloak — is the genuine upstream project. No proprietary patches, no vendor-specific APIs.

**15.6M developers behind every component.** The CNCF ecosystem represents the largest concentration of infrastructure talent in history. Each graduated project has passed rigorous security audits and governance reviews.

**Auditability is a feature, not a cost.** Regulators can inspect every line of code. SBOMs from real source trees, vulnerability scans against public CVE databases, supply-chain integrity via Sigstore.

**Community edition equals enterprise edition.** Same codebase, same features. 24/7 support and SLAs — no "enterprise-only" tier. Switch providers without switching a single binary.

**Velocity through shared innovation.** When Cilium adds WireGuard encryption or KServe adds a new model runtime, it flows into OpenOva automatically. No proprietary platform can match this R&D output.



# CATALYST: Bootstrap & Lifecycle

From initial platform setup to Day-2 operations — guided, automated, and continuously managed. Runs on K3s across any cloud.

Platform engineering is the #1 strategic technology trend (Gartner 2024), yet 94% of core modernization projects exceed timelines and budgets (IBM IBV). The bottleneck isn't technology — it's the operational complexity of assembling, configuring, and maintaining dozens of cloud-native components across environments.

CATALYST eliminates this complexity. A guided bootstrap wizard provisions a production-ready platform in hours, not months. Day-2 lifecycle management — automated upgrades, health monitoring, drift repair — keeps every component current and healthy without manual intervention. Deploy on any cloud or bare metal, from edge to enterprise.




## Bootstrap Wizard

Guided platform provisioning with interactive configuration and automated deployment



## Lifecycle Manager

Continuous Day-2 operations with automated upgrades, health checks, and drift repair



## K3s

Lightweight Kubernetes distribution optimized for edge and production workloads



## Cloud Providers

Multi-cloud support: Huawei Cloud, AWS, Oracle OCI, Hetzner, and bare metal

# PILOT: GitOps & IaC

Git-driven delivery and infrastructure as code. Every change is auditable, reversible, and automated.

The shift to GitOps is accelerating — 76% of organizations now use or evaluate GitOps practices (CNCF 2025). Yet many rely on proprietary Git hosting and IaC tools with restrictive licensing changes. License compliance risk is real: the BSL/SSPL trend forces enterprises to re-evaluate their entire delivery toolchain.

PILOT delivers a fully integrated, license-safe GitOps stack. Every infrastructure change is Git-versioned, auditable, and reversible.

Crossplane Cloud CRDs extend GitOps beyond Kubernetes — managing cloud resources, VMs, and databases through the same declarative workflow. Zero vendor lock-in on your delivery pipeline.



## Flux

Continuous GitOps delivery engine that reconciles cluster state with Git as the single source of truth



## Gitea

Lightweight internal Git hosting with built-in CI/CD pipelines and issue tracking



## OpenTofu

MPL 2.0 licensed infrastructure-as-code for declarative cloud provisioning and management



## Crossplane

Kubernetes-native cloud resource management using custom resource definitions and reconciliation

# SPINE: Networking & Mesh

eBPF-powered service mesh, L7 proxy, WAF, DNS automation, global load balancing, and secure connectivity.

Enterprise networking is fragmenting across sidecar proxies, proprietary load balancers, managed WAFs, and per-seat VPN licenses. Each adds latency, cost, and operational burden. Traditional service meshes inject sidecar containers into every pod — adding 15-30% memory overhead and significant complexity.

SPINE leverages eBPF to eliminate sidecars entirely. Cilium provides CNI, service mesh, mTLS, and L7 observability in a single kernel-level data plane. Add WAF protection, automated DNS, global load balancing, and secure connectivity across hybrid and edge environments — all managed declaratively through Kubernetes.



## Cilium

eBPF-powered CNI with built-in service mesh, mTLS, and L7 visibility



## Envoy

High-performance L7 proxy for traffic management, load balancing, and observability



## Coraza

Open-source web application firewall implementing OWASP Core Rule Set protection



## ExternalDNS

Automated DNS record synchronization between Kubernetes resources and DNS providers



## k8gb

Global server load balancing with authoritative DNS for multi-cluster traffic management



## frpc

Reverse tunnel proxy for exposing services behind NAT in hybrid and edge deployments



## NetBird

Zero-config WireGuard mesh VPN for secure peer-to-peer connectivity across clusters



## strongSwan

IPsec VPN gateway for site-to-site encryption and legacy network integration

# GUARDIAN: Security & Identity

End-to-end security: TLS, secrets, scanning, runtime protection, supply chain integrity, policy enforcement, and identity management.

The largest breaches of the decade came from proprietary, closed-source software — a supply-chain attack on a network monitoring platform (\$100B+ impact), an endpoint agent update that crashed 8.5M devices, a file-transfer zero-day (\$10B+ damages). Meanwhile, open-source code is audited by millions. Backdoors can't hide in code the world can read.

GUARDIAN wraps your platform in defense-in-depth: automated TLS, zero-trust secrets, eBPF runtime detection, supply chain signing, SBOM tracking, policy-as-code, and enterprise IAM. Every layer is open source, auditable, and integrated — from build pipeline to production runtime to identity.



## cert-manager

Automates TLS certificate lifecycle with Let's Encrypt and private CA integration



## ESO

Syncs secrets from external vaults into Kubernetes with automatic rotation support



## OpenBao

MPL 2.0 secrets vault with per-cluster isolation and dynamic credential generation



## Trivy

Comprehensive vulnerability scanner for container images, filesystems, and IaC templates



## Falco

eBPF-based runtime threat detection for containers and Kubernetes workloads



## Sigstore

Keyless container signing and verification for supply chain integrity assurance



## Syft+Grype

Automated SBOM generation paired with vulnerability matching and CVE tracking



## Kyverno

Policy engine for validating, mutating, and generating Kubernetes resources at admission



## Keycloak

Enterprise identity and access management with OIDC, SAML, and FAPI support

# INSIGHTS: AIOps & Observability

Full observability stack with AI-native operations. Specter's 6 autonomous agents replace legacy SOAR/SIEM — DevSecOps, DevOps, SRE, FinOps, Compliance, and AI Ops.

Enterprises spend \$300-500K/yr on proprietary observability — per-host pricing that scales linearly with infrastructure. As cloud-native environments grow to thousands of containers, cost becomes prohibitive. Meanwhile, alert fatigue and siloed tools mean MTTR stays high despite massive spending.

INSIGHTS unifies logs, metrics, traces, and SIEM into a single stack with zero per-host fees. Specter's 6 autonomous AI agents — DevSecOps, SRE, FinOps, Compliance, DevOps, and AI Ops — replace the traditional SOC/NOC model. They detect, diagnose, and remediate issues before your team even wakes up. Chaos engineering and usage metering round out a complete operational platform.



## Specter

Six autonomous AI agents for DevSecOps, SRE, FinOps, and continuous compliance



## Grafana

Unified dashboards for metrics, logs, traces, and alerts in a single pane of glass



## Alloy

Next-generation OpenTelemetry collector and telemetry pipeline agent by Grafana Labs



## Loki

Horizontally scalable log aggregation system optimized for Kubernetes metadata queries



## Mimir

Long-term metrics storage with global query capabilities and multi-tenancy support



## Tempo

Distributed tracing backend with native OpenTelemetry support and deep Grafana integration



## OpenTelemetry

Vendor-neutral instrumentation framework for automatic trace and metric collection



## OpenSearch

Search and analytics engine powering hot-tier SIEM and security event correlation



## Litmus

Chaos engineering platform for proactive resilience testing of Kubernetes workloads



## OpenMeter

Real-time usage metering and billing engine for consumption-based pricing models

# SURGE: Scaling & Resilience

Intelligent autoscaling, configuration management, and high-availability orchestration.

Over-provisioning wastes 27% of cloud spend (\$44.5B globally), while under-provisioning causes outages. Traditional scaling is reactive — you pay for peak capacity 24/7 or risk downtime. Manual capacity planning fails in dynamic, event-driven architectures.

SURGE combines vertical right-sizing (VPA), event-driven horizontal scaling (KEDA), and continuous availability orchestration into an intelligent autoscaling engine. Workloads scale from 10 to 10,000 pods based on actual demand — queue depth, HTTP traffic, custom metrics. Zero-downtime deployments with automatic failover and configuration-change detection.



## VPA

Automatically right-sizes pod CPU and memory requests based on actual utilization patterns



## KEDA

Event-driven horizontal autoscaling triggered by queue depth, HTTP traffic, or custom metrics



## ReLoader

Watches ConfigMaps and Secrets for changes and triggers rolling restarts automatically



## Continuum

Continuous availability orchestration ensuring zero-downtime deployments and failovers

# SILO: Storage & Registry

S3-compatible object storage, backup & disaster recovery, and container artifact registry.

Data sovereignty demands self-hosted storage. Cloud egress fees (\$0.08-0.12/GB) and proprietary registry lock-in make multi-cloud and hybrid deployments expensive and inflexible. Backup strategies that depend on a single cloud provider fail the sovereignty test.

SILO provides S3-compatible object storage, Kubernetes-native backup with cross-region replication, and a full container/Helm registry — all self-hosted. Run it on-prem, in your sovereign cloud, or alongside public clouds. Your data stays where regulations and business logic dictate. No egress surprises.



## MinIO

S3-compatible high-performance object storage for cloud-native data and backup targets



## Velero

Kubernetes-native backup, restore, and disaster recovery with cross-region replication



## Harbor

Enterprise container and Helm chart registry with vulnerability scanning and RBAC

# CORTEX: AI Hub

Self-hosted AI infrastructure. Full control over models, prompts, and costs. No data leaves your environment.

\$2.5T will be spent on AI in 2026, with 90% going to inference (Gartner). Yet 95% of GenAI pilots fail ROI (MIT) — not because of the models, but because enterprises lack the infrastructure to run them. API-based AI services leak data, impose per-token fees, and create hard vendor lock-in on the most strategic technology of the decade.

CORTEX gives you full control. Self-hosted multi-model inference cuts token costs by 90%+. RAG pipelines with vector and graph databases. AI safety guardrails. LLM observability. A ChatGPT-like interface for business users. All running on your infrastructure, with your data, under your governance. Switch models freely — no API lock-in.



## KServe

Kubernetes-native model serving with autoscaling, canary rollouts, and multi-framework support



## Knative

Serverless platform for deploying event-driven and scale-to-zero inference workloads



## vLLM

High-throughput inference engine with PagedAttention for optimal GPU utilization



## Milvus

Distributed vector database for similarity search in RAG and recommendation pipelines



## Neo4j

Graph database for knowledge graphs, entity relationships, and AI-powered reasoning



## BGE

State-of-the-art embedding and reranking models for retrieval-augmented generation



## LibreChat

ChatGPT-like interface supporting multiple LLM backends for enterprise AI chat



## NeMo

AI safety guardrails framework for input/output filtering and responsible AI deployment



## LangFuse

LLM observability platform with prompt tracing, cost analytics, and evaluation tools



## Axon

Managed LLM gateway providing subscription-based inference access and rate limiting

# FABRIC: Data & Integration

Event-driven data platform: PostgreSQL, streaming, workflows, CDC, data lakehouse, and BI analytics.

Data infrastructure is the most expensive line item in enterprise IT. Proprietary database licensing alone costs \$100-300K/yr per organization, while managed event streaming adds another \$80-120K. Per-core pricing punishes scale. License audits punish growth. And each vendor creates its own data silo.

FABRIC replaces fragmented data infrastructure with a unified, open-source data platform. Cloud-native PostgreSQL with HA and automated backups. Event streaming for real-time architectures. Workflow orchestration for distributed transactions. A complete data lakehouse for analytics. Self-service BI. All integrated, all open source, all under your control.



## CNPG

Cloud-native PostgreSQL operator with HA, automated backups, and declarative management



## FerretDB

MongoDB wire protocol compatibility layer running on PostgreSQL for document workloads



## Valkey

Redis-compatible in-memory cache and data store with full cluster mode support



## Strimzi

Apache Kafka operator for production event streaming with automated topic management



## ClickHouse

Columnar OLAP database for real-time analytics on billions of rows with SQL interface



## Flink

Unified stream and batch processing engine for complex event processing pipelines



## Temporal

Durable workflow orchestration for sagas, long-running processes, and distributed transactions



## Debezium

Change data capture platform streaming database changes to Kafka in real time



## Iceberg

Open table format enabling data lakehouse architecture with time-travel queries



## Superset

Self-service BI platform with interactive dashboards, SQL Lab, and data exploration



## Fingate

Open banking platform with FAPI authorization, OBIE APIs, and usage metering

# FINGATE: Open Banking

End-to-end open banking solution compliant with FAPI and OBIE standards. Built entirely on open-source components from the OpenOva stack.

Open banking mandates (PSD2, OBIE, Saudi Arabia SAMA) require FAPI-certified APIs — yet proprietary platforms charge per-API-call fees that scale exponentially with transaction volume. Banks face a choice: pay ever-growing vendor fees or build from scratch. Neither option is sustainable for the 89% growth in ME/Africa digital banking.

FINGATE delivers a complete open banking solution built entirely on OpenOva's open-source stack. FAPI-certified authorization, OBIE-compliant APIs, real-time metering, payment orchestration with saga patterns, and financial-grade API protection. Zero per-API-call fees. Full regulatory compliance. Built for the next decade of digital finance.



## Keycloak (FAPI)

Financial-grade API authorization server with Strong Customer Authentication for PSD2



## OBIE APIs

Open Banking Implementation Entity standards for accounts, payments, and consent management



## Temporal

Payment orchestration engine for PSD2 consent flows and settlement workflows



## OpenMeter

API consumption metering with real-time billing and partner settlement reporting



## Envoy + Coraza

API gateway with WAF protection, rate limiting, and financial-grade traffic management



## CNPG

PostgreSQL for transactional ledger data with point-in-time recovery and replication



## Strimzi + Debezium

Kafka event streaming with CDC for real-time balance synchronization and audit trails

# Engagement Model

Two service tracks — platform and consultancy — take you from discovery to autonomous operations.

## Platform Services

Subscription-based, self-service + managed

### Community Edition

Full platform, self-managed. Community = Enterprise — same code, same features.

### Enterprise Subscription (ESA)

24/7 enterprise support, SLAs, Expert Network access (80hr/quarter), continuous updates, quarterly business reviews.

### Fully Managed SOC/NOC

Specter AI agents + human oversight. 24/7 autonomous operations — DevSecOps, SRE, FinOps, Compliance.

### Pay-As-You-Go (Axon)

Managed LLM inference — consumption-based pricing, no GPU procurement, instant access.

## Consultancy Services

Engagement-based, expert-led

### Discovery & Strategy

Full landscape audit — vendor contracts, license obligations, integration dependencies. TCO/ROI modeling. Board-ready strategy with milestones and risk quantification.

### Statement of Work (SoW)

Fixed-scope delivery: platform deployment, workload migration (Exodus methodology), compliance alignment, team enablement, and knowledge transfer.

### Time & Materials (T&M)

Embedded experts that scale with your transformation. Architecture reviews, performance optimization, security hardening — ramping up or down as phases progress.

### CIO & DT Advisory

C-suite strategic partnership: technology roadmap aligned to business KPIs, change management, transition planning, and quarterly board reporting.

# The Guild — Expert Network

A distributed team of senior engineers, architects, and consultants across **6 cities** and **3 continents**. Follow-the-sun coverage ensures 24/7 responsiveness.

## TEAM STRUCTURE

### Platform Engineering

Kubernetes, infrastructure, observability, security

### AI & Data Engineering

ML pipelines, LLM inference, data platforms

### Consulting & Advisory

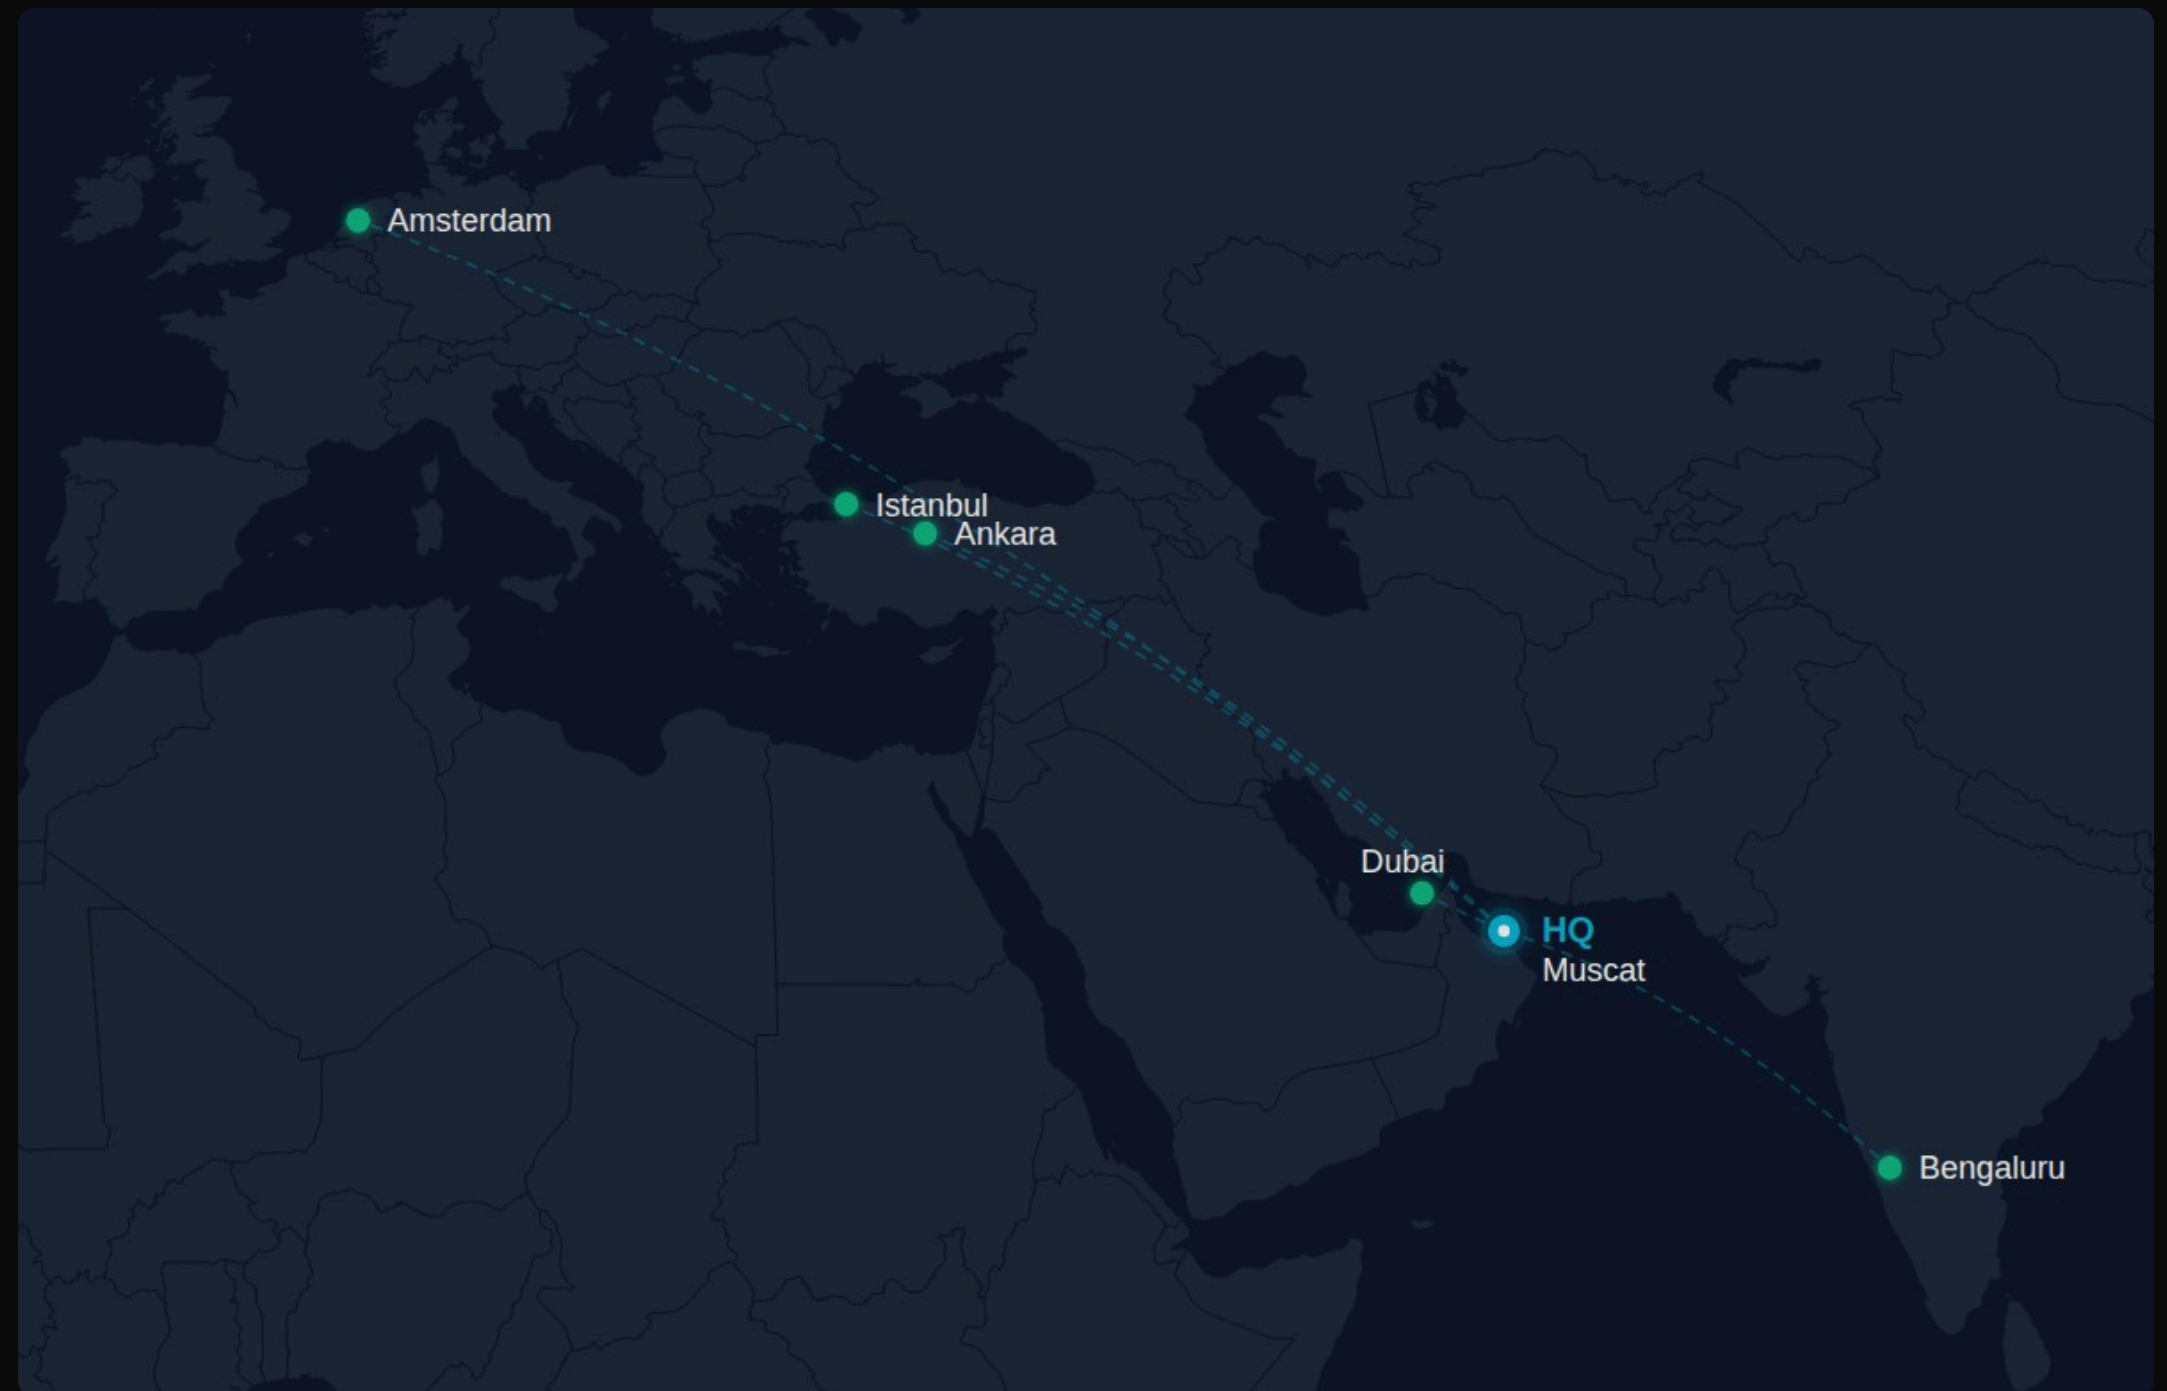
Strategy, migration planning, change management


### Managed Operations

24/7 SOC/NOC, SRE, FinOps, compliance

## EXPERTISE

Banking & financial services, regulated industries, sovereign cloud, enterprise Kubernetes at scale, legacy migration (VMware, Oracle, SAP).





# Let's Build Your Platform

Pure open-source. AI-native. Zero lock-in.  
From assessment to production in weeks, not years.

`hello@openova.io` · `openova.io`

## REFERENCES

# Sources & Citations

**Gartner (Feb 2026)** — Global IT Spending Forecast: \$6.15 trillion for 2026. IT services (\$1.9T), software (\$1.4T), telecom (\$1.4T), devices (\$836B), data centers (\$653B).

**Gartner (2025)** — CIO Survey: 72% of CIOs identify fragmented infrastructure as the #1 barrier to AI ROI.

**Zylo (2025)** — SaaS Management Index: Average large enterprise runs 660 SaaS applications. Organizations above 10,000 employees average \$284M in annual SaaS spend. 53% of software licenses sit completely unused.

**Flexera (2025)** — State of the Cloud Report: 27% of \$450B+ global cloud spend is pure waste. Average organization wastes ~\$17M per year on idle/over-provisioned resources.

**IBM Institute for Business Value** — 94% of core modernization projects exceed their timelines and budgets.

**MIT / Fortune (2025)** — 95% of GenAI pilot programs fail to deliver measurable business value.

**CNBC (Q1 2026)** — Over \$1 trillion in software market capitalization destroyed. Hedge funds made \$24B in profits shorting legacy enterprise vendors in five weeks.

**CNN Business (Jan 2025)** — \$285 billion single-day wipeout in AI sector stocks.

**CIO / Avasant Research** — VMware post-Broadcom acquisition license restructuring: increases up to +1,500% reported across customer base.

**Deloitte** — Enterprise IT budget allocation: ~75% maintenance vs ~25% innovation across Fortune 500.

**Gartner (2025)** — Platform Engineering and Platformization: converged platform approach reduces operational overhead by 60–85%.

**Linux Foundation (2024)** — State of Open Source: \$8.8 trillion economic value contributed by open-source software globally.

**DORA / Google Cloud (2024)** — State of DevOps Report: Platform teams with mature internal developer platforms see 4x faster deployment frequency.

**McKinsey (2025)** — Digital Sovereignty: 67% of EU enterprises prioritize data sovereignty in vendor selection.

**IDC (2025)** — AI Infrastructure: Organizations with unified platforms are 3.2x more likely to achieve production AI deployments.

**Harvard Business School (2024)** — The Value of Open Source Software: \$8.8 trillion demand-side economic value globally. Rebuilding OSS would cost 3.5x current software spend.

**FINOS (2025)** — State of Open Source in Financial Services: 85% of FS firms increasing OSS adoption. 20% report >\$1M annual savings.

**Gartner (Jan 2026)** — AI Spending Forecast: \$2.52 trillion worldwide in 2026, representing 41% of all new technology spend.

**Gartner (Feb 2026)** — Sovereign Cloud IaaS: \$80 billion worldwide in 2026 (+35.6% YoY). ME & Africa 89% growth.

**Gartner (2025)** — Top Strategic Technology Trends: 75% of EU/ME enterprises will geopatiate workloads by 2030.

**CNCF (Jan 2026)** — Annual Cloud Native Survey: 82% of container users run Kubernetes in production (up from 66% in 2023).

**IDC (2024)** — IT Skills Shortage: \$5.5 trillion cost by 2026 in delays, quality issues, and lost revenue. 90% of organizations affected.