

SentinelForge AI

AI-OPS ECOSYSTEM

The AI ecosystem for infrastructure operations — from provisioning to optimization

FORGE · PROTECT · OPTIMIZE — on-premise · sovereign · open-source



Figure 1 — The SentinelForge AI ecosystem: three modules (Forge, Protect, Optimize) on the same substrate — on-premise LLM engine, traced execution, autonomy policies, audit, multi-provider middleware.

The problem

Managing a cloud infrastructure requires three distinct skills at three different moments: building the environments, keeping them healthy, optimizing their consumption. Today each of these phases lives in separate tools, with knowledge concentrated in a few people and processes that do not talk to one another.

Added to this is a constraint that in the public and regulated sector is often non-negotiable: the AI automation solutions available on the market require data, telemetry and credentials to be sent out to external services. Intelligent automation and data sovereignty, today, appear to be alternatives.

The solution

SentinelForge AI is Epic Edge's AI ecosystem for the entire life cycle of the infrastructure. Not three disconnected applications, but modules that share the same substrate — on-premise language engine, traced deterministic execution, autonomy policies, audit — and cover the three phases: build (Forge), monitor (Protect), optimize (Optimize).

One module is adopted at a time, on the infrastructure already in place; the value grows when the modules work together on the same foundation. At every phase the artificial intelligence operates entirely within the organization's perimeter: no data leaves the network.

The model proposes. The automation engine executes.

The three modules

EdgeForge AI FORGE AVAILABLE

Natural-language infrastructure provisioning, on OpenStack and Proxmox VE. The user describes what they need; the assistant designs the architecture, sizes the resources, estimates the cost and — after approval — carries out the deployment in an automated and traced manner.

CephSentinel AI PROTECT FIRST PILOT

Continuous monitoring and operational automation for Ceph fleets, including multi-cluster and on remote clusters. It correlates, deduplicates and prioritizes signals, and proposes — executing them on confirmation — the day-to-day operations, with three levels of autonomy configurable per cluster.

CloudSentinel AI OPTIMIZE UNDER DEVELOPMENT · 2027

Intelligent monitoring and energy optimization on the physical infrastructure: anomaly detection, automatic root cause analysis and workload consolidation to reduce consumption. Under development; collection of operational data already under way.

The common substrate

The modules are not separate products with different technologies: they all rest on the same foundation. It is this that makes the ecosystem more than the sum of its parts — and that allows new domains to be added without rebuilding from scratch each time.

- **On-premise language engine** — open-weight models run within the customer's perimeter (Ollama): no data, no telemetry, no credential leaves the network.
- **The model never executes** — it generates a plan; a separate deterministic execution engine (AWX + Ansible) carries it out. It is what makes audit, dry-run and revocation before execution possible.
- **Autonomy policies** — the same three levels across the whole ecosystem: notification, execution subject to approval, automatic execution limited to low-risk actions.
- **Native traceability** — every playbook is versioned on Git and linked to the signal that generated it: what was done, where and why always remains open to inspection.
- **Multi-provider abstraction** — an intermediate layer normalizes the underlying technologies: adding a domain (OpenStack, Ceph, Kubernetes, Proxmox...) means extending the platform, not rewriting it.

Key features

- **Sovereign by design** — AI automation does not require exposing data and infrastructure to non-European services.
- **Modular** — you adopt the module you need, when you need it; starting with one does not preclude the others.
- **Coherent** — the same approval, autonomy and audit logic in every module: you learn it once, it applies to the whole ecosystem.
- **Open-source at the base** — built on the same stacks that power European sovereign and telco clouds: no proprietary lock-in.
- **Defensible in an audit** — every action, manual or assisted, leaves a complete trail: a requirement, not an extra, in regulated contexts.

Benefits for the customer

One ecosystem, not three tools

Building, monitoring and optimizing live in the same platform, with the same experience — not in separate silos with different skills.

Sovereignty preserved

Artificial intelligence enters operations without giving up control of the data: everything stays within the perimeter.

Distributed expertise

The operational knowledge is in the system, not in a few people: faster onboarding and less dependency on rare profiles.

Low-risk adoption

You start with one module on a real environment, with everything supervised; autonomy widens as trust grows.

Stack and technologies

Infrastructure domains	OpenStack · Ceph · Kubernetes · Proxmox VE
LLM engine	Open-weight models run on-premise via Ollama
Execution and automation	AWX + Ansible, with the playbooks versioned on GitLab
Observability	Prometheus and Grafana — leveraging the tools already present
Identity and access	Keycloak — SSO and federation with corporate directories

Who it is for

- **Telco & ISP** — operators with production open-source infrastructure and high volumes of repetitive operations.
- **Cloud & Service provider** — those who deliver cloud services and seek to reduce the operational cost across the entire life cycle.
- **Public administration** — AI automation compatible with data sovereignty and traceability requirements.
- **Finance & Banking** — contexts where every infrastructure intervention must leave an audit trail.

Status and adoption

EdgeForge AI (Forge)	Available — provisioning operational on OpenStack and Proxmox VE
CephSentinel AI (Protect)	First pilot available — multi-cluster observation and supervised actions
CloudSentinel AI (Optimize)	Under development — 2027 roadmap; collection of operational data under way
Adoption model	One module at a time, on the same substrate: starting with one does not preclude the others

Perché è diverso da un'AI gestita in cloud

La differenza non è nel numero di funzioni: è in dove vivono i dati, il modello e il controllo.

	Epic Edge SentinelForge AI	AI-as-a-Service gestita in cloud
AI eseguita on-premise, nel perimetro del cliente	● Sì	● No
Nessun dato o telemetria inviati all'esterno	● Sì	● No
Stack open-source, senza lock-in proprietario	● Sì	● No tecnologia proprietaria
Audit completo e tracciabile di ogni azione	● Sì Git + audit trail	● Parziale log applicativi
Sovranità e giurisdizione del dato	● Sì	● Dipende dal provider e dalla region
White-label per l'operatore nazionale	● Sì in rilascio	● No
Avvio immediato senza infrastruttura propria	● No richiede il proprio ambiente	● Sì

La sovranità ha un prezzo, ed è dichiarato: serve un'infrastruttura propria. In cambio, dati, modello e controllo non lasciano mai il perimetro.

Figure 2 — Epic Edge vs. cloud-managed AI: the difference lies in where data, model and control live. Sovereignty requires infrastructure of your own; in exchange nothing leaves the perimeter.

Roadmap di prodotto

Dove sta andando la piattaforma — con lo stato reale di ciascun elemento, non promesse.

Disponibile in produzione e in pilot	Roadmap in sviluppo - 2027	Visione strategica direzione della piattaforma
● EdgeForge AI provisioning in linguaggio naturale — live ● CephSentinel AI operazioni sulle flotte Ceph — primo pilot ● CMP — portale unificato multi-provider — in implementazione ● OpenStack · Proxmox VE provisioning operativo ● Osservabilità Grafana metriche per ambiente — disponibile	● CloudSentinel AI sorveglianza predittiva ed efficienza energetica ● VMware provisioning — in integrazione ● Capacity planning cross-provider vista aggregata nel portale ● White-label completo per l'operatore nazionale ● Agente CephSentinel per cluster dietro firewall	● AWS · Azure gestiti governo degli estate hyperscaler ● Osservabilità estesa oltre le metriche: log e tracce ● Operazioni autonome il cloud che si gestisce da sé, sotto controllo
● disponibile ● pilot / in implementazione ● pianificato 2027 ● visione		

Figure 3 — The product roadmap: available today, planned for 2027, the direction of the platform — every item with its real status.

Architectural principles

Building an AI automation platform for infrastructure means holding together things that usually live separately: integrating technologically different providers, maintaining independent adapters that evolve without touching the rest, guaranteeing a complete audit of

every action, preserving data sovereignty and — above all — separating the engine that decides from the engine that executes.

This is why SentinelForge AI was designed as a modular ecosystem: not a collection of functions, but an architecture designed by people who have lived with those constraints in production.

Why Epic Edge

25+ years of operations

Real production infrastructures, not concepts on slides. The expertise comes from the field.

Already in production

EdgeForge AI is live and CephSentinel AI is in pilot: the ecosystem is not a roadmap, half of it is already running.

Sovereign and open-source

On-premise AI, open stack, no dependency on non-European services.

Which phase would you like to start from?

Forge, Protect or Optimize: we identify the module with the fastest return on your environment and start there, with a low-risk pilot.

andrea.martra@epic-edge.it · www.epic-edge.it/sentinelforgeai

EdgeForge AI LIVE	CephSentinel AI PILOT	CMP IN IMPLEMENTAZIONE	CloudSentinel AI ROADMAP
----------------------	--------------------------	---------------------------	-----------------------------

Epic Edge S.r.l. · Via Michele Buniva, 26 — 10064 Pinerolo (TO), Italy · info@epic-edge.it · www.epic-edge.it

OPEN-SOURCE • SOVEREIGN • ENTERPRISE-GRADE • AI-DRIVEN