EPICEDGE | PRODUCT SHEET

ClusterSentinel AI LIVE

The Kubernetes cluster lifecycle, governed by the platform

A module of the SentinelForge AI ecosystem · on-premise · sovereign · open-source · AI-native

The problem

Running Kubernetes in production means creating clusters, adding and removing them, scaling them, protecting their state, restoring and upgrading them. Every operation goes through Rancher or the command line and requires skills that few people in the organization truly possess.

The cost is not only time: cluster credentials multiply and circulate, etcd copies stay on the very node they are meant to protect, and an irreversible operation — a restore, an upgrade, a deletion — carried out by mistake cannot be undone. Governing the fleet stays in the hands of a few specialists.

The solution

ClusterSentinel AI is the SentinelForge AI ecosystem module dedicated to the lifecycle of Kubernetes clusters. It lives within the CMP as the panel's Kubernetes catalog: from there the user creates, governs, scales, protects, restores and upgrades their own clusters — without opening Rancher or using kubectl, and without the cluster credentials passing through the platform.

Like EdgeForge AI, ClusterSentinel AI never communicates directly with the infrastructure: it translates every request into a plan executed by the automation engine in a traced manner, and irreversible operations are protected by an explicit confirmation. The same operations can also be requested in natural language, through the assistant shared with EdgeForge AI.

The model proposes. The automation engine executes.

How it works

1 · CREATE OR IMPORT

The user creates a K3s or custom cluster, or brings an already-active cluster under management without moving it.

2 · THE PLATFORM DECLARES

The operations available on a cluster are declared by the platform: whatever is not supported simply does not appear.

3 · VERIFY AND CONFIRM

Irreversible operations — restore, upgrade, deletion — require retyping the cluster name, not a simple confirmation.

4 · TRACED EXECUTION

The plan passes through the automation engine (AWX + Ansible): every action is recorded and inspectable.

“Create a 3-node K3s cluster and schedule a nightly etcd copy, with an S3 bucket as destination.”

ClusterSentinel AI →

K3s cluster · 3 nodes · scheduled etcd snapshot · external S3 destination (the secret stays in the customer’s Rancher)

Operational plan proposed before execution | no irreversible action without confirmation

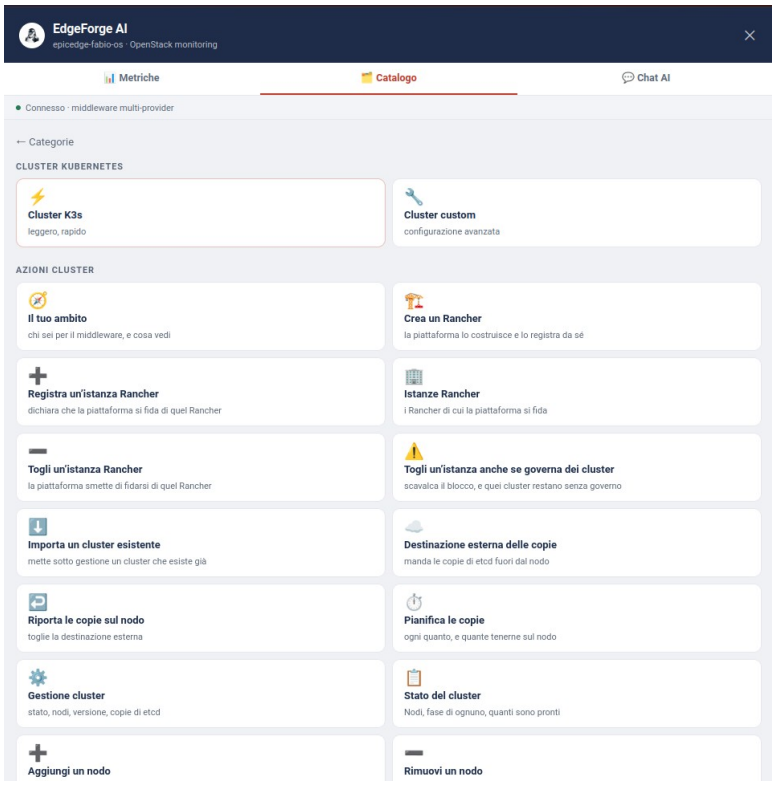

Figure 1 — ClusterSentinel AI’s Kubernetes catalog: creating, scaling, protecting state, restoring and bringing clusters under management, from the panel.

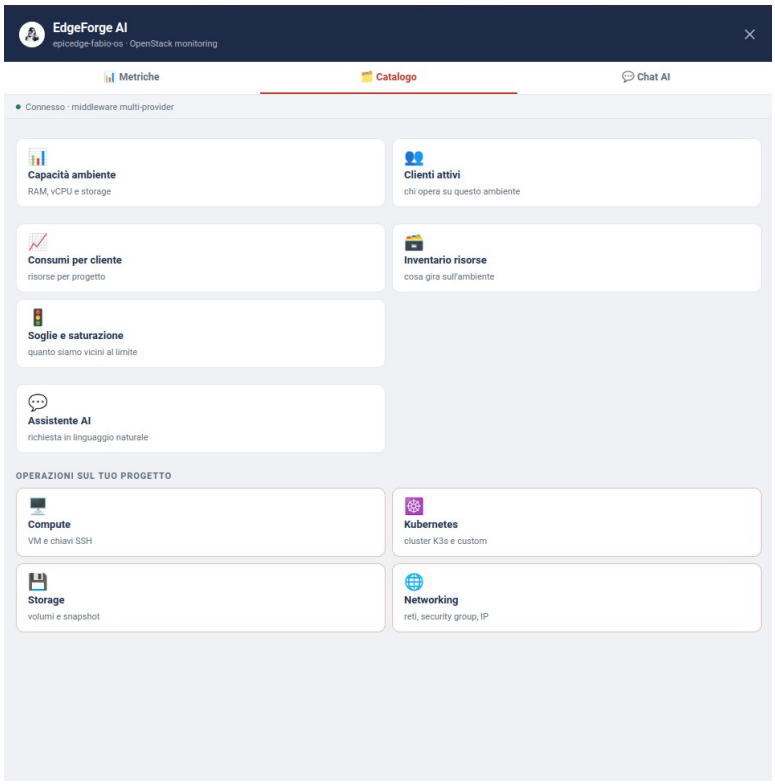

Figure 2 — ClusterSentinel AI lives in the unified operational panel as the Kubernetes catalog, alongside compute, storage and network — same experience, same assistant.

Key features

- **Full cluster lifecycle** — creating, scaling, protecting state, restoring and upgrading, all from the platform panel.
- **Capabilities declared by the platform** — the operations available on a cluster are declared by the platform, not by the interface: a cluster whose machines' location is unknown cannot be scaled, and the command does not appear.
- **Storage credentials never in transit** — to send etcd copies to an S3 bucket, the secret with the keys stays in the customer's Rancher; only the name of the secret is communicated to the platform.
- **Irreversible operations protected** — restore, upgrade and deletion require retyping the cluster name, not a simple confirmation.
- **State protection (etcd)** — on-demand or scheduled copies, with configurable retention on the node and an external destination on S3 storage.
- **Execution decoupled from the model** — the model generates the plan, a separate execution engine (AWX + Ansible) runs it. This is what makes audit, dry-run and revocation before execution possible.
- **No Rancher or kubectl for the user** — the user operates the clusters from the panel; the cluster credentials never pass through the platform.
- **Operable in natural language too** — the same operations can be requested through the assistant shared with EdgeForge AI.

Customer benefits

Kubernetes without dedicated specialists

Creating and governing clusters no longer requires mastering Rancher and kubectl: the operations are in the panel.

State truly protected

etcd copies are scheduled and taken off the node: recoverable even if the master is lost.

Fewer irreversible errors

Restore, upgrade and deletion are protected by an explicit confirmation by cluster name.

Centralized fleet governance

All clusters — new and imported — under a single view and a single operational logic.

Data sovereignty preserved

Cluster credentials do not pass through the platform, and the AI processing stays within the customer's perimeter.

Full traceability

Every operation passes through the automation engine and remains recorded and inspectable for audits.

Stack and technologies

Orchestration	Kubernetes (K3s and custom distributions), governed through Rancher
LLM engine	Open-weight models run on-premise via Ollama — no dependency on external APIs
Cluster state	etcd — on-demand or scheduled copies, retention on the node and external destination on S3
Execution engine	AWX / Ansible Automation Controller
Versioning	Git — every generated playbook is saved, versioned and inspectable
Interface	Kubernetes catalog native in the CMP panel (Epic Edge Platform)

Who it's for

- **Telco & ISP** — operators with multiple production Kubernetes clusters to govern in a uniform way.
- **Cloud & Service Provider** — those who deliver Kubernetes-as-a-Service and want to reduce the operational cost per cluster.
- **Public Administration** — cluster governance compatible with sovereignty requirements: credentials and data stay within the perimeter.
- **Finance & Banking** — contexts where every operation on the clusters must leave a complete audit trail.
- **Enterprise IT** — platform teams managing cluster fleets with staff undersized relative to demand.

Deployment model

Mode	On-premise, within the CMP and inside the customer's perimeter. No mandatory SaaS component.
Prerequisites	Rancher for the managed clusters; GPU resources dedicated to the on-premise inference engine.
Integration	Kubernetes catalog native in the CMP panel — no parallel portal to adopt.
Activation	Initial assessment, connection of the existing Rancher instances and import of the already-active clusters.
Customization	Autonomy and copy-retention policies configurable per individual cluster and customer.
Support	Delivered by Epic Edge across the whole stack (Kubernetes, OpenStack, Ceph, AI engine), not only on the module.

The Epic Edge ecosystem

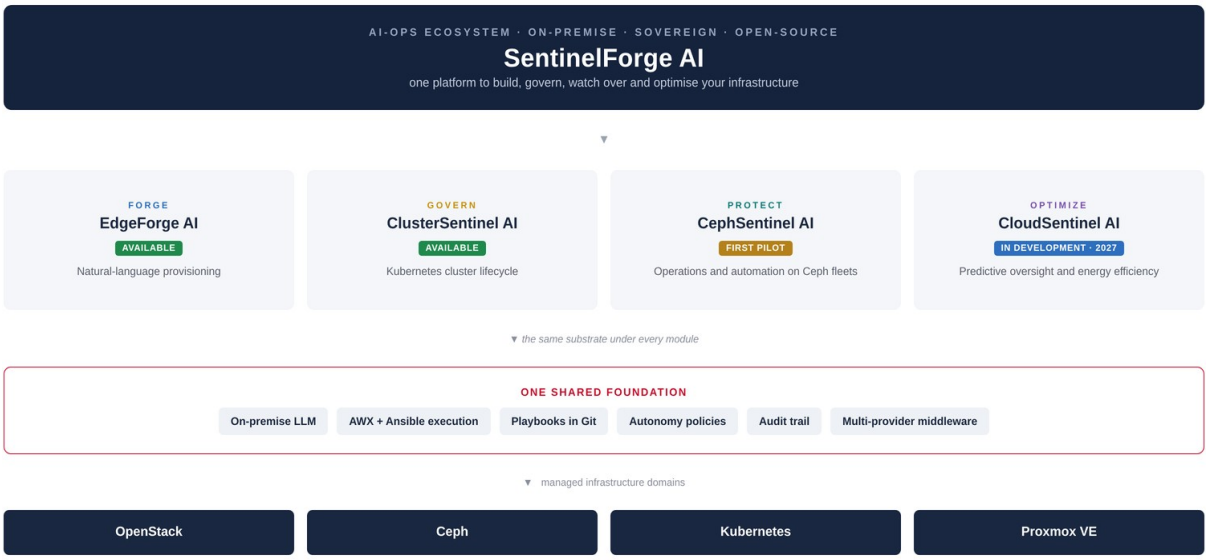

Every Epic Edge product is part of SentinelForge AI: distinct modules on the same substrate — on-premise LLM, traced execution, autonomy policies, audit, multi-provider middleware.

Architectural principles

Building an AI automation platform for infrastructure means holding together things that usually live apart: integrating technologically different providers, maintaining independent adapters that evolve without touching the rest, guaranteeing a full audit of every action, preserving data sovereignty and — above all — separating the engine that decides from the engine that executes.

This is why ClusterSentinel AI is born as a module of the SentinelForge AI ecosystem: not a function added to an existing tool, but an architecture conceived by those who have lived those constraints in production.

Why Epic Edge

25+ years of operations
Real data centers, real SLAs, real enterprise customers. Not a startup imagining the cloud.

Already in production
ClusterSentinel AI is operational today within the CMP — not a prototype, not a demo.

Native to the stack
Built on OpenStack, Ceph and Kubernetes: the stack of European sovereign and telco clouds.

Want to see ClusterSentinel AI on your stack?

We can arrange a demo on your real use case, with a preliminary assessment of your Kubernetes clusters.

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



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