



Cloud Management Platform

AI-native Cloud Management Platform — multi-provider, sovereign and white-label

Product presentation document

Version 1.9.4 · July 2026

Purpose of this document

This document describes, from a product and service perspective, the Cloud Management Platform developed by Epic Edge: who it is aimed at, which needs it addresses, how it presents itself to users and what value it brings to a national cloud and public administration context. It does not contain implementation details or technical-level architectural descriptions, which are covered by separate documentation.

1. Summary

Epic Edge's Cloud Management Platform (CMP) is the single point of access through which an organization governs its own cloud resources, regardless of the technology on which those resources are built.

Today most organizations do not have a single cloud: they have a private platform based on open source technologies, a traditional virtualization environment and sometimes resources hosted with international public operators. Each of these environments has its own console, its own language, its own access logic and its own way of measuring consumption. The result is a fragmentation that generates training costs, operational errors, reporting difficulties and, above all, a loss of overall control.

Epic Edge's CMP resolves this fragmentation by presenting all environments through a single, coherent interface, with a single access point, a single logic for displaying consumption and two integrated artificial intelligence components: EdgeForge AI, which carries out the operations requested by the user in natural language, and CloudSentinel AI, which monitors the state of the infrastructure, anticipates its failures and optimizes its energy consumption. The platform is entirely based on open source technologies, can be installed at the customer's premises and involves no dependency on non-European services.

In one sentence

A single platform, customizable with the brand of the operator that adopts it, which makes cloud environments that differ from one another simple and uniform to govern, providing the user with a natural-language assistant for day-to-day operations and the operator with an intelligent monitoring and optimization system.

2. The context and the needs it addresses

The project stems from Epic Edge's direct experience in designing and managing cloud infrastructures for enterprise customers and service operators. Four needs emerge consistently from this experience.

- **Simplified operations.** The native consoles of cloud platforms are tools designed for specialists. Most users need to perform a few recurring operations — creating a virtual machine, adding disk space, opening a network port — without having to know the conceptual model of the underlying platform.
- **Unified view of consumption.** Those who deliver the service need to know, at any moment, how much capacity is available, how much is committed and by whom. Those who use the service need to know how much of their plan they have consumed. Today these two pieces of information live in separate systems.
- **Technological independence.** Tying oneself to a single technology or a single supplier is a strategic risk, all the more so for entities that deliver services of public interest. The ability to add, replace or combine different environments without rewriting the management platform is a requirement, not an option.
- **Data sovereignty.** Data, credentials, metrics and automations must remain within the organization's perimeter and under national and European jurisdiction. This also applies — and particularly stringently — to the artificial intelligence components.

Epic Edge's CMP was designed starting from these four needs, and not as an adaptation of a pre-existing product.

3. What the Epic Edge CMP is

The platform presents itself as a web portal which the user accesses with their corporate credentials. From that moment on, everything they see is determined by who they are and by what they have actually purchased or been assigned: the platform does not expose generic menus to be filtered, but builds the interface around the user.

Feature	Description
Multi-provider	It manages cloud environments of different kinds — open source platforms, traditional virtualization, public operators — presenting them all in the same way to the user.
White-label	The platform is designed so that the operator adopting it can offer it to its own customers with its own brand and visual identity; full white-label customization is in the release phase. Epic Edge remains the technology supplier and, in that mode, does not appear to the end user.
Data-driven	Adding a new environment or a new catalog product does not require a new release of the portal: the platform draws what it is told.
AI-assisted	Two integrated components: EdgeForge AI for carrying out operations in natural language and CloudSentinel AI for intelligent monitoring and energy optimization, both operating within the organization's perimeter.
Open source	The entire technology chain is based on open source components, with no binding proprietary licenses and no dependencies on non-European suppliers.
Installable on site	The platform can be installed within the customer's infrastructure, keeping data and credentials under their exclusive control.

3.1 How it works: the architecture in brief

Between the interface and the managed environments operates a multi-provider middleware: it is the component that makes it possible to present different technologies in the same way. It translates and normalizes the APIs of each backend, exposes a uniform catalog of capabilities to the platform and delegates execution to an automation engine based on Ansible and AWX. The AI assistants never communicate directly with the infrastructure: they produce a plan, which the middleware executes in a traced manner.

It is this abstraction layer that makes it possible to add or replace a provider without rewriting the assistants or the interface.

The model proposes. The automation engine executes.

The platform is also data-driven: the portal does not contain screens developed ad hoc for each provider or product. Each environment describes its own capabilities through the middleware, and the CMP dynamically builds the interface, the catalog and the available operations. Introducing a new service therefore means extending the data catalog and the middleware adapters, not rewriting the user interface.

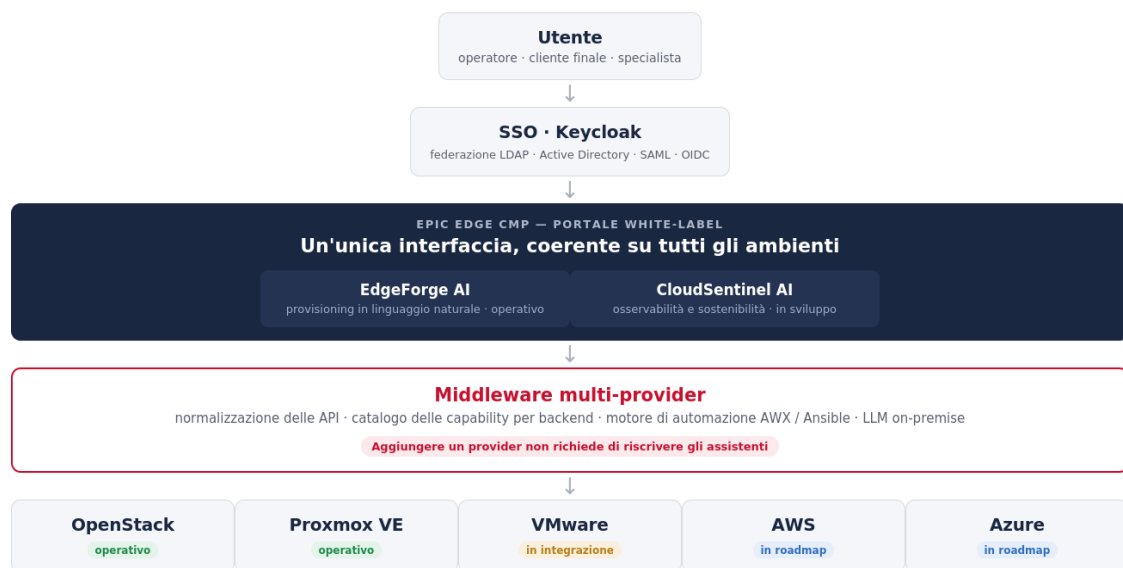


Figure 0 — Logical architecture: federated access, single portal with the AI assistants, multi-provider middleware, heterogeneous backends.

The CMP does not replicate the functionality of the underlying platforms: it offers a uniform operating model for them, leaving the native consoles available for all specialist functions.

How it differs from a traditional CMP. Established multi-cloud management platforms exist, but almost all of them delegate conversational automation and analysis to external services. Here the difference does not lie in the number of functions: it lies in the fact that the artificial intelligence operates entirely on-premise — the model does not leave the organization's perimeter — on an entirely open source and reversible chain. It is this combination that makes AI automation compatible with sovereignty requirements that a commercial CMP, by the way it is built, cannot satisfy.


4. Access to the platform

Entry takes place through a centralized authentication system that can be integrated with the identity systems already in use within the organization. The user uses a single credential for all environments, without having to know or manage the technical accounts of the individual underlying platforms. Authentication is based on a standard identity

provider (Keycloak): the platform can therefore federate with existing corporate directories — LDAP, Active Directory — and with third-party providers via SAML or OpenID Connect, reusing the identities and access policies already in use within the organization.

Once authentication is complete the platform already knows the user's profile, their role and the list of services that pertain to them. It is this information that determines the entire subsequent experience.

Figure
1 —
Single
sign-
on



access: one credential for all environments, federable with corporate identities.

5. Two usage profiles, a single platform

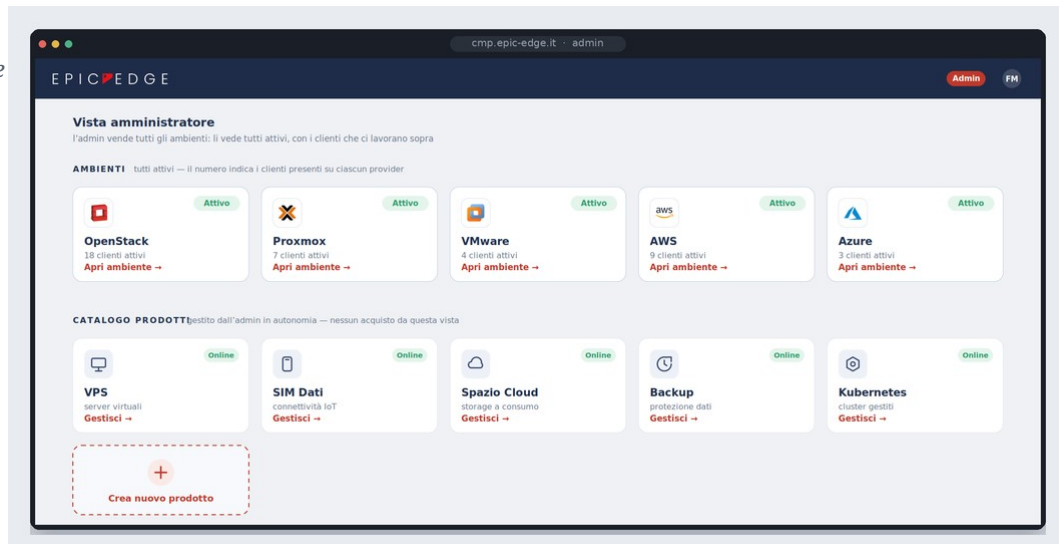
The platform distinguishes two profiles, which access the same interface but with different purposes and visibility.

Profile	Role and scope of visibility
Administrator	This is the operator that delivers the service. It governs the catalog of products offered and accesses all environments in monitoring mode: it sees the overall capacity of each environment and the customers operating on it, with their respective consumption.
End customer	This is the user of the service. They see the entire available offering, operate on the environments that have been activated for them and, for those not yet active, are directed to the purchase path defined by the operator.

5.1 The administrator's view

On the home screen the administrator finds two distinct bands: the technological environments offered to its customers, each showing how many customers operate on it, and the catalog of commercial products it publishes and manages independently.


Figure 2 —



Administrator view: the environments offered and the product catalog, governed from a single point.

On entering an environment, the administrator does not see its own resources but the overall picture of the environment: the available and committed capacity for memory, computing power and storage space, the number of running instances and the list of active customers with the consumption of each. It is the view of those who have to plan capacity and report on the service.

Figure 3 —



Administrator view of an environment: available capacity and consumption per individual customer, for planning and reporting.

The platform also exposes, within the same portal, the observability dashboards of the environment, built on the standard tools already present — Prometheus and Grafana. Metrics are collected by leveraging the observability infrastructure already present in the environment, without introducing proprietary agents into the virtual machines and

without duplicating collection. Alongside the aggregated capacity and consumption data, the administrator has detailed metrics for each individual resource — CPU, memory, network and disk of the virtual machines — and a capacity-planning-oriented reading: occupancy per environment, trend over time and remaining margin, as a basis for planning expansions. The platform is also designed to offer an aggregated cross-provider view that brings together in a single place the metrics of technologically different environments, without opening a separate monitoring console for each; this functionality is currently being implemented in the unified portal.

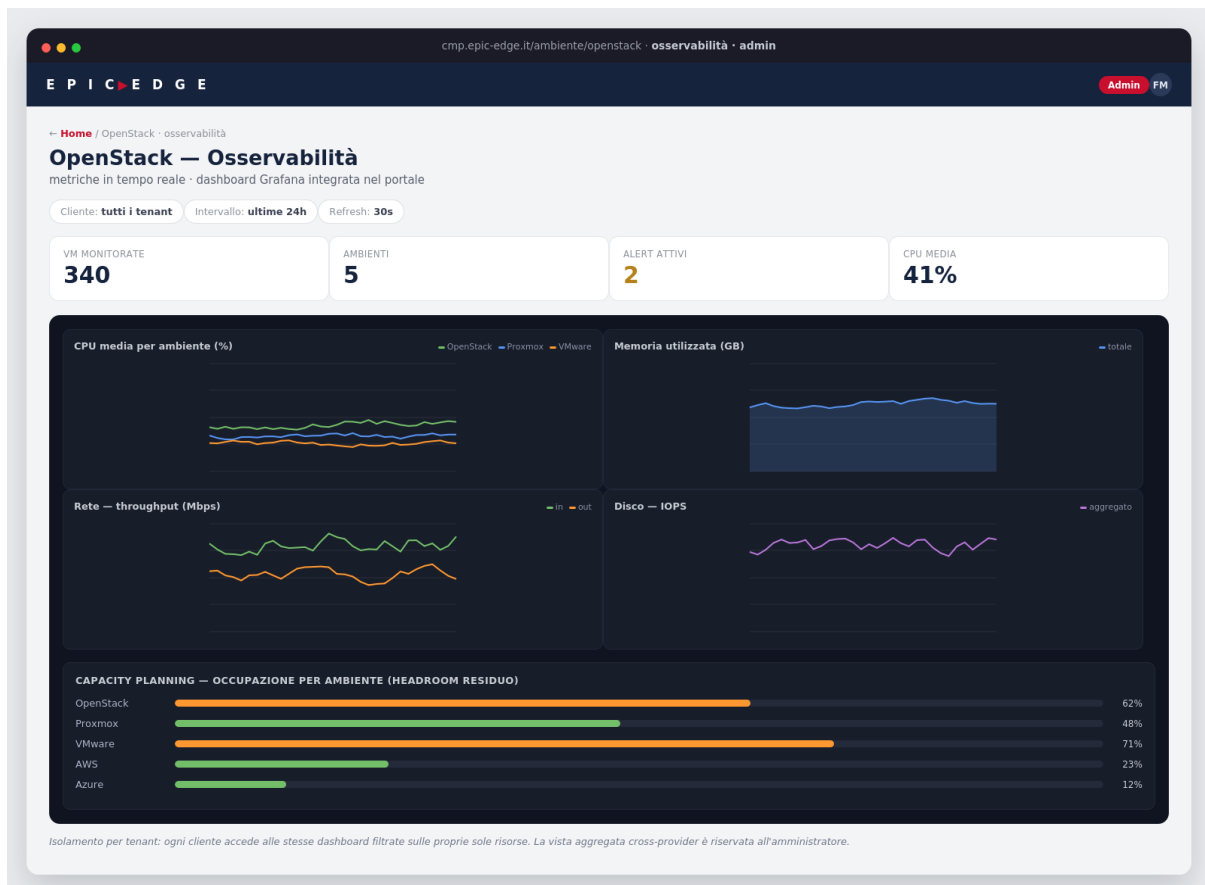
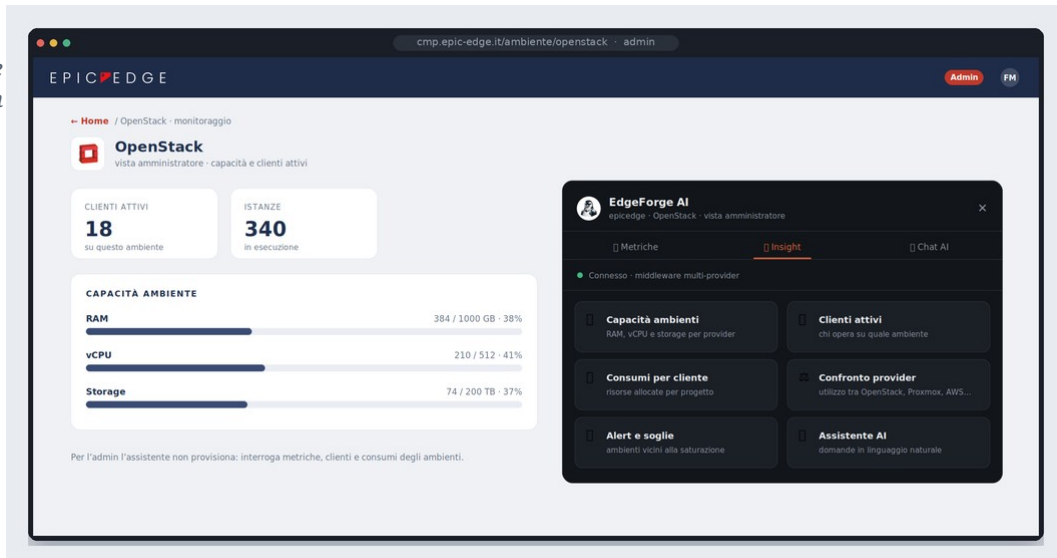


Figure 3-bis — Integrated observability: Grafana dashboard with the detailed resource metrics and the capacity planning reading per environment.

The EdgeForge AI assistant is also available to the administrator; in this profile it does not perform provisioning operations but answers monitoring questions: which environments are approaching saturation, how the different providers compare with one another, which resources are allocated to which customer. It is in this same area that CloudSentinel AI, described in chapter 7, will be integrated.

Figure 4 — In the




administrator view the AI assistant answers monitoring questions: it does not perform provisioning.

5.2 The end customer's view

The end customer sees the operator's entire offering. The environments already activated are immediately accessible; those not yet active remain visible and, if selected, lead to the purchase page set up by the operator. The same logic applies to the product catalog. In this way the interface is not only an operational tool, but also a commercial channel.


Figure 5 —



Customer view: access to the active environments and to the catalog of purchasable services, in the same place.

Within an active environment the customer finds only what belongs to them: their own machines, their own networks, their own volumes and the usage of the subscribed plan against the agreed limits. The structure of the page is identical for all environments: the content changes, not the way of reading it.

Figure 6 —

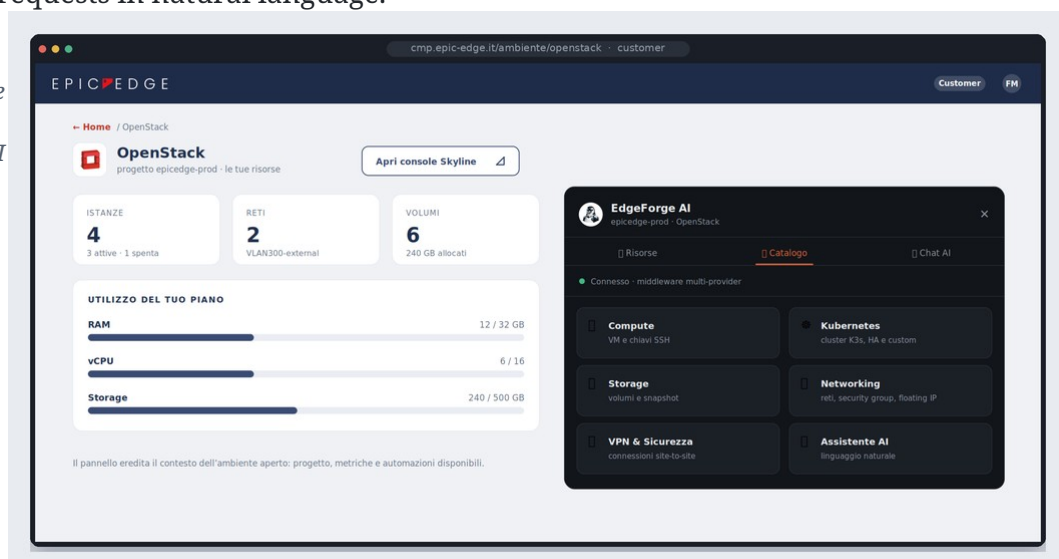


Customer view of an environment: only their own resources and plan usage, with an identical reading for every provider.

In the same way, the customer accesses the observability dashboards of their own resources, with the detailed metrics of their own instances and services — CPU, memory, network and disk. An aggregated view of overall consumption across the different activated environments is also planned, being implemented together with the unified portal. The dashboards are isolated per tenant: each customer sees exclusively their own resources, never those of other customers on the same environment.

From this same page the customer can open the EdgeForge AI assistant, which automatically inherits the context of the environment they are in and proposes the available automations grouped by functional area, as well as the possibility of making free requests in natural language.

Figure 7 —
The AI



assistant inherits the context of the active environment and proposes only the automations actually available.

6. EdgeForge AI: the operational assistant

EdgeForge AI is the component that most distinguishes the platform. It is not an informational chatbot placed alongside the portal, but an operational assistant that actually carries out the user's requests on the infrastructure, through verified and traced automations.

The user can describe what they want in natural language — creating a test environment, adding space to an existing machine, installing an application on an already active server — and the assistant translates the request into a sequence of operations, executes it and reports its progress in real time. Alternatively, for recurring operations, the assistant offers guided paths that accompany the user step by step. In every case, EdgeForge AI always proposes the operational plan to the user before execution, allowing verification and approval: no resource is created without explicit approval.

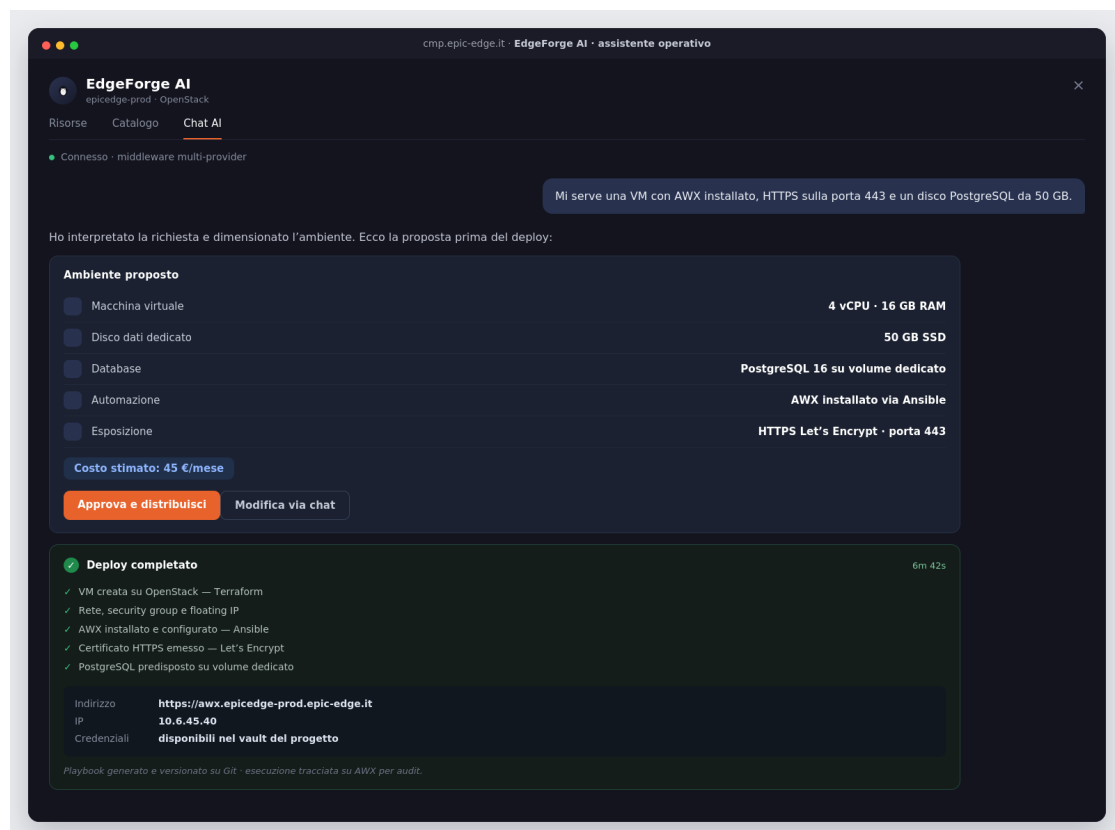


Figure 7-bis — EdgeForge AI in action: from the natural-language request to the proposed environment, through to the automated and traced execution of the approved plan.

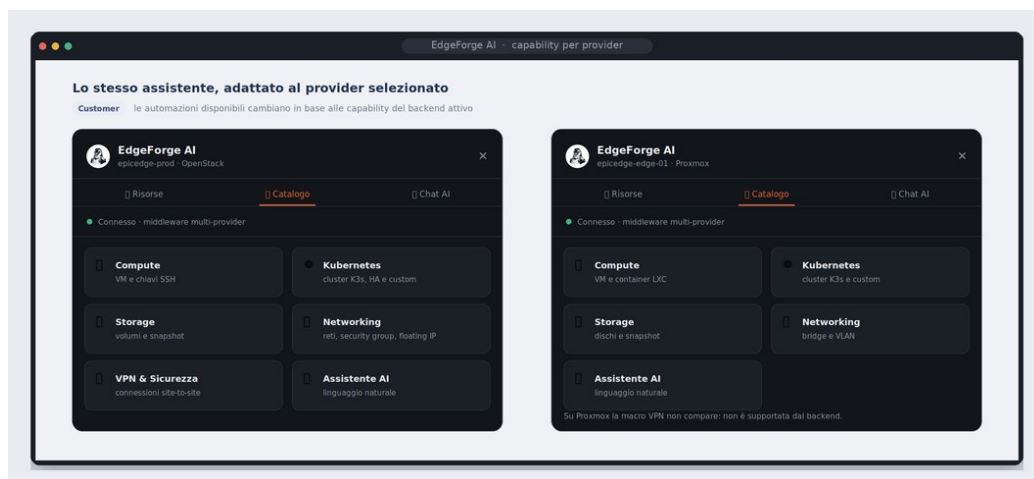
- **Contextual.** It knows which environment the user is in and with what authorizations they operate. There is no need to specify it with every request.
- **Multi-platform.** Automatic provisioning is already operational on several technologies: in addition to open source cloud environments, EdgeForge AI also creates and manages

resources on traditional virtualization platforms, with the same interface and the same language. The user does not need to know which technology they are working on. At present automatic provisioning is fully operational on OpenStack and on Proxmox VE; the extension to VMware is being integrated.

- **Adaptive.** The automations proposed change according to the actual capabilities of the selected environment. If a feature is not supported by the underlying technology, it is simply not offered: the user does not encounter options that could not work.
- **Sovereign.** Processing takes place within the organization's perimeter. No data relating to the infrastructure, to customers or to credentials passes to external artificial intelligence services.
- **Traced.** Every operation performed by the assistant passes through the organization's automation system, which keeps a complete record of it for verification and audit purposes.

The following figure shows the same assistant on two technologically different environments: in both, automatic provisioning is fully functional, but the list of automations adapts to what each platform is actually able to do.

Figure 8 —
The same AI



assistant adapts to the provider's capabilities: on Proxmox the VPN macro does not appear because the backend does not support it.

Why this aspect matters

The adoption of artificial intelligence in the public sector often encounters an obstacle of a non-technical nature: the impossibility of letting sensitive data leave the organization's perimeter. EdgeForge AI is designed to operate entirely within that perimeter, making conversational automation compatible with confidentiality and data sovereignty requirements.

The extension to VMware follows the same principle: the assistant presents the native resources of the environment — datastores, distributed switches, templates — and offers only the automations actually supported by the backend, without the user needing to know the underlying technology.




Figure 8-bis — The assistant extended to VMware: the automations mirror the native resources of the backend (datastore, distributed switch, template).

7. CloudSentinel AI: intelligent monitoring and sustainability

If EdgeForge AI deals with building and modifying the infrastructure, CloudSentinel AI deals with what happens afterwards: keeping it healthy, anticipating its problems and reducing its energy consumption. It is the platform's second artificial intelligence component, currently under development, and it will be integrated into the administrator interface alongside the current monitoring functions.

CloudSentinel AI does not replace the monitoring tools already in use within the organization: it brings them together and enriches them with an analytical capability that today requires rare and expensive specialist skills.

7.1 The needs it addresses

Critical issue	Impact on the organization
Unplanned outages	A service outage on critical infrastructure has a very high hourly cost and, in the public sector, a direct impact on citizens. Service continuity commitments allow unavailability margins in the order of a few tens of minutes per year.
Long recovery times	Analyzing the cause of a complex failure today requires hours of work by highly specialized staff, available in limited numbers on the market.

Critical issue	Impact on the organization
Energy waste	The average utilization of servers in data centers is typically far below the installed capacity: much of the energy absorbed produces no service.
Environmental reporting	The regulatory and contractual obligation to document the consumption and emissions of digital infrastructure is growing, without native tools existing to do so on the cloud.

7.2 The planned functions

CloudSentinel AI will operate along four lines, presented to the administrator within the same interface already used for monitoring.

- **Anomaly detection.** It learns the normal behavior of each environment and reports critical deviations before they become service disruptions.
- **Failure prediction.** It analyzes the trend over time of the infrastructure indicators in order to anticipate failures with hours or days of notice, allowing planned interventions rather than emergency reactions.
- **Automatic root cause analysis.** Once an incident has occurred, it correlates the events of the different components and returns a timeline reconstruction, the most likely cause and the suggested actions.
- **Assisted remediation.** It proposes and executes the corrective actions with three levels of autonomy selectable by the organization: simple notification to the operator, execution subject to approval, or automatic execution limited to low-risk actions only.

7.3 The energy sustainability module

Alongside these functions is an energy efficiency module: it consolidates workloads on fewer servers without service interruption, takes operating conditions into account when placing new workloads, and automatically produces consumption and emissions reporting per individual service, in a format designed to be consistent with European sustainability frameworks.

The combined value of the two components


Together, EdgeForge AI and CloudSentinel AI cover the entire life cycle of the infrastructure: describing and building, then monitoring and optimizing. For the organization this means simultaneously reducing the time needed to deliver a service, the time needed to restore it in the event of a failure and the energy used to keep it running —

three quantities that today are governed with separate tools and different skills.

8. Continuity with the technical consoles

The platform does not replace specialist tools, it complements them. For advanced operations the user accesses, without new authentication, the native console of the environment in which they are working, where they find the entire set of functionality of the underlying platform.


Figure 9 — From the CMP, single sign-on access to the



provider's native technical console for advanced operations.

In this context too the EdgeForge AI assistant remains available and retains the same functionality: the specialist user can continue to work with the tools they know and, where convenient, delegate the more repetitive operations to the assistant. The following images document functionality that is already operational in Epic Edge environments.

Figure 10 — The AI



assistant remains available inside the technical console too: functionality already operational in Epic Edge environments.

8.1 Integration with the organization's ecosystem

The platform does not live in isolation: it is designed to fit into the tools already present in the organization, reusing them rather than replacing them. The main integration points are:

Identity. Federation with existing identity providers — LDAP, Active Directory, SAML, OpenID Connect — via the SSO system.

Automation. Execution through Ansible and AWX, with the playbooks versioned on the organization's Git repository.

Monitoring. Connection with the metrics and observability tools already in use, starting with Prometheus.

Application interoperability. REST APIs and webhook notifications for connection with ticketing systems and internal portals.

Some connections are already operational (identity, automation, metrics); others, such as connectors to specific ticketing systems, are planned for subsequent developments.

9. The value in a national cloud context

The platform responds directly to some of the principles that guide national and European cloud policies for public administration.

Principle	How the platform responds to it
Data sovereignty	Installation within the customer's infrastructure, Italian and European jurisdiction, no dependency on non-European services, including the artificial intelligence components.
Reversibility	Open source technologies and standard formats. The organization can change supplier or underlying technology without having to redesign its own management platform.
Absence of constraints	No binding proprietary license and no exit cost. The platform is designed to coexist with heterogeneous environments, not to induce concentration on a single technology.
Service continuity	CloudSentinel AI (under development) is conceived to evolve failure management from reactive to predictive, with the aim of supporting compliance with service availability commitments.
Traceability	Every operation performed through the platform, whether manual or assisted, is recorded by the automation system and available for checks and audits.
Operational accessibility	Reduction of the skill barrier required to operate on the cloud, with direct effects on delivery times and on staff training costs.

Principle	How the platform responds to it
Transparency of consumption	Unified view of the available and committed capacity, per individual environment and per individual user, as a basis for planning and reporting.
Environmental sustainability	The energy module of CloudSentinel AI (under development) is designed to reduce the consumption of the infrastructure and to produce the environmental reporting required by European regulatory frameworks.

10. Project status

The project is at an advanced stage of development, with components already in operation and others being built.

- **Already operational.** The EdgeForge AI assistant is functional and in use within Epic Edge environments, integrated into the technical console and connected to the automation system. Automatic provisioning is already active on open source cloud platforms (OpenStack) and on traditional virtualization (Proxmox VE). Figures 9 and 10 document the real situation.
- **Being built.** The unified multi-provider portal, of which figures 1 to 8 represent the user experience design, approved and currently being implemented. The extension of automatic provisioning to further environments, VMware first of all, is being integrated.
- **Under development.** CloudSentinel AI, whose architecture is defined and for which the collection of the operational data needed to train the models has already begun. Integration will take place in the administrator interface, alongside the monitoring functions already present.
- **Planned developments.** Extension of the number of supported technologies, advanced graphic customization per individual operator and expansion of the predictive consumption analysis functionality.

Note on the images

Figures 9 and 10 are images of a real environment in operation. Figures 1 to 8 represent the interface design of the unified portal, currently being built. CloudSentinel AI, described in chapter 7, is under development and is not yet represented in the images. The branding shown in the figures is that of the Epic Edge demonstration instance; in white-label mode it can be replaced with the operator's identity. Figure 8-bis illustrates the extension of the assistant to the VMware environment and is likewise part of the design in progress.

10.1 Status overview

Function	Status
SSO access (Keycloak, LDAP/AD/SAML/OIDC)	● Available
OpenStack — EdgeForge AI provisioning	● Available
Proxmox VE — EdgeForge AI provisioning	● Available
EdgeForge AI — operational assistant	● Available
Unified multi-provider portal	● Being implemented
VMware — provisioning	● Being integrated
AWS — environment management	● On the roadmap
Azure — environment management	● On the roadmap
White-label	● Being released
CloudSentinel AI — observability and sustainability	● Under development
Observability per environment (Grafana dashboard)	● Available
Aggregated cross-provider metrics + capacity planning	● Being implemented

Legend: ● green = available and in use · ● amber = being implemented or released · ● blue = under development.

11. Why this platform is strategic for a national Cloud Service Provider

Beyond the individual functions, the value of the platform for a national operator can be summarized in five concrete points:

A single front-end. OpenStack, Proxmox, VMware and the further platforms integrated over time through the multi-provider middleware, governed with the same user experience, without multiplying consoles and training.

Faster provisioning. Natural-language automation takes the commissioning of a standard environment from days of manual work to an automated deployment of a few minutes.

No dependency on external AI. The model remains within the organization's perimeter: intelligent automation does not require exposing data and infrastructure to non-European services.

White-label for the national operator. The service is offered to the end customer with the operator's brand and identity, not with those of the technology supplier.

A backend that evolves independently. Adding, replacing or decommissioning an underlying technology does not change the end user's experience.

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.

12. Epic Edge

Epic Edge S.r.l. is an Italian company specializing in the design, construction and management of cloud infrastructures based on open source technologies. It operates on critical production environments for enterprise customers and service operators in Italy and Europe, with specific expertise in open source cloud platforms, distributed storage, container orchestration and managed services with guaranteed service levels.

The company's positioning is centered on data sovereignty and technological independence: entirely open source infrastructures, data in Italy and in the European Union, no contractual or technical constraint that hinders reversibility.



Epic Edge S.r.l.

www.epic-edge.it · info@epic-edge.it