EPIC EDGE

SERVICE CATALOGUE

Epic Edge Services

Design, build and operation of open-source cloud infrastructure

OpenStack · Ceph · Kubernetes · Proxmox VE — enterprise-grade, sovereign, with reversibility by design

Version 1.0 · August 2026

The service model

Epic Edge designs, builds and operates cloud infrastructure based on open-source technologies for telecommunications operators, cloud and service providers, public administration, industry and enterprise organisations. Our work covers the entire life cycle of a platform: assessment of what is already in place, architectural design, implementation, workload migration and continuous operation.

The catalogue is organised across three layers of abstraction, which can be adopted individually or as a complete stack. No layer presupposes the others: Epic Edge can be engaged for design alone, for the operation of an existing platform alone, or for the entire journey.

Layer	Scope	Services
Infrastructure	Compute, storage, network, virtualisation	OpenStack Cloud IaaS · Proxmox VE · Ceph Enterprise Storage · Enterprise Private Cloud · Sovereign Cloud · Edge Computing
Platform	Containers, orchestration, multi-environment governance	Managed Kubernetes · Multi-Cloud Hub · Provider Solution
CloudOps	Operation, evolution, specialist support	Managed Services and 24/7 On-call · Ceph Support · OpenStack Lifecycle & Support · Migration & Implementation · VMware Migration

The principle that holds the catalogue together

The goal is not to replace the technologies a client already owns, but to make them governable within a single operating model. The infrastructure core of our services is entirely open source: no proprietary licence on the platform layer, standard formats and interfaces, reversibility designed in from the architecture onwards. Any Epic Edge software modules are additional and optional components, distinct from the core: the platform remains operable without them.

Catalogue at a glance

Service	Layer	In brief	Status	Page
OpenStack Cloud IaaS	Infrastructure	Multi-tenant IaaS cloud with compute, GPU and distributed storage	Delivered	5
Proxmox VE Hyper-converged	Infrastructure	Hyper-converged KVM/LXC cluster	Delivered	6
Ceph Enterprise Storage	Infrastructure	Distributed block, object and file storage	Delivered	7
Managed Kubernetes	Platform	Private high-availability clusters, operation delegated	Delivered	8
Enterprise Private Cloud	Infrastructure	Single-tenant private cloud on dedicated hardware	Delivered	9
Sovereign Cloud	Infrastructure	Architecture built around residency and auditability requirements	Available by project	10
Provider Solution	Platform	White-label platform for MSPs, CSPs and telcos	Early release / pilot	11
Multi-Cloud Hub — CMP	Platform	Unified governance of heterogeneous environments	Early release / pilot	12
Edge Computing	Infrastructure	Distributed compute nodes in the field	Available by project	13
Migration & Implementation	CloudOps	Assessment, design and workload transfer	Delivered	14
VMware Migration	CloudOps	Exit from proprietary virtualisation	Delivered	15
Ceph Support	CloudOps	Support for production distributed storage clusters	Delivered	16
OpenStack Lifecycle & Support	CloudOps	Health checks, upgrades, expansion, tuning	Delivered	17
Managed Services and 24/7 On-call	CloudOps	Continuous operation with defined coverage levels	Delivered	18

Adoption paths

Services can be adopted individually, but the most effective path starts from the need to be solved. The table below shows, for the six most common contexts, where it makes sense to start and which services join along the way.

Need	Entry point	Services that follow
Build a cloud offering to resell — ISPs, MSPs, telecommunications operators	Provider Solution	OpenStack Cloud IaaS, Ceph Enterprise Storage, Managed Services and 24/7 On-call
Modernise the data centre and exit proprietary virtualisation	VMware Migration	Enterprise Private Cloud on OpenStack or Proxmox VE, Migration & Implementation, Ceph Enterprise Storage
Operate and evolve a platform that already exists	OpenStack Lifecycle & Support	Ceph Support, Managed Services and 24/7 On-call, Managed Kubernetes
Govern heterogeneous environments from a single point	Multi-Cloud Hub	Managed Kubernetes, Managed Services and 24/7 On-call
Meet residency, jurisdiction and auditability requirements	Sovereign Cloud	Ceph Enterprise Storage, Enterprise Private Cloud, Managed Services and 24/7 On-call
Bring processing and storage close to the data source	Edge Computing	Managed Kubernetes, Multi-Cloud Hub

How it starts

In most cases the first step is an assessment of the existing environment with a contained and verifiable scope: it establishes what is actually worth doing and in what order, before committing budget to a complete programme.

Infrastructure platforms

The services in this family concern the construction and management of the layer that delivers compute, storage and network. The choice of virtualisation platform is not one of technological preference but of scale and operating model: OpenStack where real multi-tenancy, a service catalogue and APIs are required, Proxmox VE where compactness and operational simplicity matter. Distributed storage is treated as a standalone service, because Ceph is adopted both as the backend of one of these platforms and as an independent storage system.

INFRASTRUCTURE

OpenStack Cloud IaaS

Multi-tenant IaaS platform on an entirely open-source stack

We design and build production OpenStack clouds: virtualised and bare-metal compute, accelerated instances, Ceph distributed storage, software-defined networking and value-added network services. The platform is multi-tenant by construction, with federatable identity, per-project quotas and traceability of administrative operations.

Deployment is containerised and automation-driven, so that installation is deterministic and repeatable and each service can be upgraded individually. It is the same approach with which we have brought multi-region platforms into service at European telecommunications operators and cloud providers.

What it includes

- **Elastic compute —** virtual instances, bare-metal nodes and GPU instances; live migration and automatic re-placement on hypervisor failure depend on the configuration and on the high-availability components being enabled, and are defined at design stage
- **Distributed storage —** Ceph block, object and file storage, with protection profiles differentiated per pool (replication or erasure coding)
- **Software-defined networking —** isolated networks, routing distributed across the hypervisors, native IPv6, load balancing, firewall and VPN delivered as services
- **Advanced network performance —** SR-IOV and DPDK for latency-sensitive workloads and for network function virtualisation
- **Identity and governance —** centralised authentication, federation with corporate directories, per-project roles, audit trail and encryption of data at rest and in transit
- **Automation and integration —** dashboard, CLI and APIs; Terraform templates and Ansible playbooks for repeatable and verifiable provisioning
- **Metering —** metric collection and per-project consumption measurement, as a basis for capacity planning and internal chargeback

Platform	OpenStack in containerised, automation-driven deployment
Storage	Ceph — block, object and distributed file system
Network	distributed SDN backend, VxLAN, SR-IOV/DPDK, LBaaS and FWaaS
Automation	Terraform, Ansible, AWX; configuration versioned in Git
Observability	Prometheus and Grafana, with metric collection for reporting

Who it is for

Telecommunications operators and ISPs, cloud and service providers, public administration, enterprises running a private cloud with a high volume of provisioning requests.

INFRASTRUCTURE

Proxmox VE Hyper-converged

Hyper-converged cluster combining compute, storage and network in a single solution

A hyper-converged architecture places virtualisation and distributed storage on the same nodes, reducing the number of devices to purchase and maintain. It is the configuration that delivers a complete and resilient platform within a contained footprint, while retaining high availability, live migration and native backup.

We state the cost of that choice at design stage: in a hyper-converged configuration compute and storage compete for processor, memory, network and I/O, and sizing must reserve resources for storage. It is a manageable constraint, but one to be considered up front rather than discovered in operation.

What it includes

- **Virtualisation —** KVM virtual machines and LXC containers with near bare-metal performance
- **High availability —** cluster with integrated quorum and automatic re-placement of virtual machines on node failure; adding and removing nodes can be planned without stopping the cluster, subject to relocating workloads and verifying quorum
- **Flexible storage —** Ceph, ZFS, iSCSI, NFS; snapshots, live cloning and thin provisioning
- **Data protection —** Proxmox Backup Server for incremental, deduplicated and encrypted backups, with retention policies and periodic verification; copies can be synchronised to a second datastore or a separate site. This is a backup function, distinct from replication of the virtual machines' storage
- **Network and security —** VLANs, bonding, software-defined networking and an integrated firewall with granular rules
- **Automation —** web console, REST API and Terraform provider for integration with existing toolchains

> **Who it is for**
>
> Organisations modernising the data centre, MSPs, established SMEs, and anyone seeking a complete platform on a contained number of nodes.

INFRASTRUCTURE

Ceph Enterprise Storage

Design and implementation of distributed block, object and file storage

Ceph brings together in a single platform three access modes — block volumes, object storage compatible with the market-standard APIs and a distributed file system — with replication or erasure-coding policies that can be differentiated per pool. That is why it is adopted both as the backend of a cloud platform and as a storage system in its own right.

We design the cluster starting from the workload profile rather than from the hardware catalogue: required latency, the read-to-write ratio and cost per terabyte determine the number of nodes, the separation between performance tiers and the data protection factor. The choices and their consequences are stated up front, together with the resulting usable capacity figures.

What it includes

- **Sizing —** definition of the topology, the number of nodes and disks and the data distribution hierarchy, starting from the workload profile
- **Block storage —** volumes for virtual machines and databases, integrable with OpenStack, Proxmox VE and Kubernetes
- **Object storage —** gateway compatible with the market-standard APIs, multi-tenancy, versioning, lifecycle policies and object immutability
- **Distributed file system —** POSIX shares for workloads requiring concurrent access, with snapshots and quotas
- **Data protection —** replication or erasure coding differentiated per pool, with the resulting usable capacity stated at design stage
- **Configurations —** dedicated cluster on storage nodes, or hyper-converged on the same compute nodes, with the constraints of either choice made explicit
- **Multi-site —** asynchronous replication of block volumes between clusters via RBD mirroring; native multi-site configurations for object storage via RGW; for the distributed file system, replication between sites requires dedicated strategies, defined case by case
- **Acceptance testing —** measurement of real performance and verification of behaviour under failure conditions before entering service

> **Who it is for**
>
> Cloud and service providers, telecommunications operators, multi-datacentre enterprises, computing centres and anyone who has to build large-capacity distributed storage without per-terabyte licensing.

PLATFORM

Managed Kubernetes

Private high-availability clusters, with operation delegated to our team

We build and operate private, isolated Kubernetes clusters, on the client's own infrastructure, on existing OpenStack or Proxmox environments, or on infrastructure designed and operated by Epic Edge. The service covers the redundant control plane, the worker nodes, integration with existing storage, observability and the upgrade cycle — that is, everything that makes a cluster usable in production and not merely installed.

The scope of the service is stated explicitly before activation: we cover the platform and its components — cluster, network, storage, ingress, monitoring — not application development nor the workloads running on top. It is the distinction that keeps a managed service sustainable over time, for the party delivering it and for the party receiving it.

What it includes

- **Isolated clusters —** high-availability control plane per client, workload separation through namespaces and access policies
- **Observability —** centralised metrics, logs and proactive alerting
- **Network and security —** granular network policies and automatic certificate management
- **Persistent storage —** integration with the distributed block and file storage already in place
- **Operation —** patching, major version upgrades, application backup and periodic restore verification

Design options

The following capabilities are not part of the base service: they are designed, implemented and quoted case by case, according to the workloads and the client's operating model.

- **Node autoscaling —** expansion and reduction of capacity based on system or application metrics
- **Advanced release strategies —** progressive rollouts with canary and blue-green deployments and rollback to the previous version
- **GitOps model and pipeline integration —** versioned manifests and continuous delivery integrated with the existing toolchain
- **Service mesh —** traffic management between services, mutual authentication and communication policies
- **Distributed tracing —** correlation of calls between services in support of performance analysis

> **Who it is for**
>
> Service providers, development teams without in-house Kubernetes skills, organisations running microservice or AI/ML workloads in production.

Delivery models

The same technologies can be delivered under different models depending on who uses them and which constraints apply. The services in this family describe those configurations: a dedicated private cloud, a platform with stringent sovereignty requirements, a solution resellable under the operator's brand, the governance of multiple heterogeneous environments, and the distribution of compute into the field.

INFRASTRUCTURE

Enterprise Private Cloud

Single-tenant private cloud on dedicated hardware, on the client's premises or in a data centre

Infrastructure entirely reserved to a single organisation, hosted in a data centre or directly at the client's site. Dedicated compute, storage and network resources mean greater performance predictability and direct control over patching and firmware policies and over maintenance windows — a necessary condition where strict compliance requirements or mission-critical workloads with guaranteed performance profiles exist.

The platform integrates with what the organisation already owns: identity systems, observability tooling, security appliances, connectivity to existing data centres and public clouds.

What it includes

- **Full isolation —** dedicated cluster, no sharing of resources with other clients
- **Platform selected on context —** OpenStack with Ceph for multi-service private clouds, with catalogue and APIs; Proxmox VE with Ceph or ZFS for compact, hyper-converged virtualisation. In both cases high availability and live migration, with the features enabled according to the agreed configuration
- **Hybrid connectivity —** layer 2 and layer 3 links, IPSec VPN, dynamic routing and peering towards public clouds
- **Integrated security —** network segmentation, firewalling, centralised log collection and integration with the organisation's SIEM
- **Business continuity —** encrypted backup, replication between sites and recovery plans documented and periodically tested
- **Cost governance —** consumption measurement per organisational unit and resource lifecycle management

> **Who it is for**
>
> Large organisations, industrial groups, finance and healthcare, bodies with compliance requirements that rule out shared infrastructure.

INFRASTRUCTURE

Sovereign Cloud

Architecture designed around agreed requirements for residency, jurisdiction and auditability

AVAILABLE BY PROJECT

An architecture designed according to agreed requirements for data residency, applicable jurisdiction, encryption key management, auditability and reversibility. Sovereignty does not follow from the software stack: it follows from placement decisions, from the key custody model and from the contractual obligations of the parties involved, all of which are defined at design stage together with the client.

Reversibility is the point on which this configuration stands out: standard formats and interfaces, an open-source infrastructure core, and documented procedures for exporting data and configurations through the available interfaces. Exit from the platform is designed as a property of the architecture, not entrusted to a contractual clause.

What it includes

- **Residency and jurisdiction —** placement of data and metadata defined at design stage, with non-transfer constraints formalised contractually
- **Key management —** encryption at rest and in transit, with an agreed key custody model and, where required, keys held exclusively by the organisation
- **Compliance —** control framework and documentation prepared in support of the controller's obligations, aligned with information security standards
- **Traceability —** recording of administrative operations with log immutability designed and documented according to the agreed requirements, and reporting prepared for reviews and audits
- **Surveillance —** centralised collection and correlation of security events, with documented response procedures
- **Reversibility —** export of data, snapshots and configurations through standard interfaces, with scope and procedures defined and documented at design stage

Who it is for

Public administration, healthcare, defence, finance, and organisations subject to data localisation constraints.

PLATFORM

Provider Solution

White-label cloud platform for MSPs, CSPs and telecommunications operators

EARLY RELEASE / PILOT

Designed for those who do not consume cloud but resell it. The platform is built so that the operator adopting it can offer it to its own clients under its own brand and visual identity: Epic Edge remains the technology supplier and does not appear to the end user.

The portal, the catalogue and the operator-branded customisation rest on the Epic Edge Cloud Management Platform, currently in early release: the capabilities described are available for pilot projects and are being consolidated together with the first group of operators. The design is that of a commercial offering rather than an internal infrastructure: tenant isolation, a product catalogue governed by the operator, per-client consumption measurement and integration with the management and billing systems already in use.

What it includes

- **Operator-branded portal —** customisable visual identity, domain and language, with self-service journeys for the end client
- **Tenant isolation —** private networks, resource quotas, roles and access control lists per project
- **Consumption measurement —** reporting per client and per resource, exportable to existing administrative systems
- **Governed catalogue —** the operator publishes and manages the products offered independently
- **Integration —** APIs and Terraform provider for connection to operational and commercial management systems
- **Modular growth —** progressive expansion planned so as to avoid general service unavailability, subject to the architecture and the residual capacity, following a replication unit designed and validated once
- **Enablement —** training of technical staff, operating procedures and third-level specialist support

> **Who it is for**
>
> MSPs, cloud and service providers, telecommunications operators and ISPs that want a cloud offering of their own without building it from scratch.

PLATFORM

Multi-Cloud Hub

Epic Edge Cloud Management Platform

Unified governance of technologically diverse environments from a single point

EARLY RELEASE / PILOT

Most organisations do not have a single cloud: they have an open-source platform, a traditional virtualisation environment and, often, resources with public operators. Each with its own console, its own vocabulary and its own consumption accounting. The result is a fragmentation that generates training costs, operational errors and loss of overall visibility.

The service brings these environments back to a uniform operating model: the same user experience, the same access policies, the same reading of consumption. The native consoles remain available for all specialist functions — the platform does not replace them, it sits alongside them.

What it includes

- **Unified view —** resources, networks and storage from different environments in a single coherent representation
- **Federated access —** a single credential across the supported environments and identity systems, integrated with corporate directories
- **Uniform policies —** roles, quotas, tags and lifecycle rules defined once and translated onto each platform within the limits of the capabilities it exposes: the model is common, its application depends on what each environment allows
- **Consumption analysis —** breakdown by project, tenant or application, with a data basis for capacity planning
- **Self-service provisioning —** catalogue of validated configurations with integrated approval workflows
- **Integration —** APIs and notifications towards existing ticketing systems, pipelines and security tooling

> **Who it is for**
>
> Organisations running multiple platforms, groups with distributed data centres, operators managing infrastructure on behalf of third parties.

INFRASTRUCTURE

Edge Computing

Compute and storage capacity distributed in the field, centrally governed

AVAILABLE BY PROJECT

When processing has to happen close to the data source — for latency, for transport cost or because the data must not leave the site — the answer is to place compute and storage nodes locally, while retaining the same automation and the same security policies as the central site.

The service is offered as an architectural model to be adapted to context: the capabilities described below are design objectives, sized and implemented for the specific case, not pre-packaged components already running at third parties.

What it includes

- **Compact nodes —** self-contained micro-infrastructures of compute and storage, sized on the site and the workload
- **Remote activation —** a design objective: minimising on-site intervention, with the node receiving its configuration from the central platform
- **Local processing —** execution of analytical workloads at the site, sending only the necessary data to the centre
- **Encrypted replication —** incremental synchronisation towards the central site or a second peripheral node, designed around the available connectivity profile
- **Remote management —** monitoring, patching and enforcement of security policies from the same tooling used for the other environments

> **Who it is for**
>
> Industry and manufacturing, distributed retail, smart cities and public infrastructure, peripheral telecommunications sites.

Professional and continuous services

Building a platform is one part of the problem; moving existing workloads onto it and keeping it in service for years is the other. The services in this family cover the entry path — assessment, design, migration — and the platform's subsequent life, including specialist support on components that require rare skills.

CLOUDOPS

Migration & Implementation

From the existing system to the new platform, with service continuity

The path is substantially the same regardless of the source technology, and its most delicate part is not only technical: survey of the source environment, compatibility assessment and sizing of the destination, preparation, data synchronisation, cutover in agreed batches, and decommissioning of the source only once verification is complete.

The rule that has proved most useful is to measure rather than estimate. The actual duration of a migration depends on the real path of the data and on the systems involved: a trial on a test environment radically changes the number of maintenance windows to be negotiated, and therefore the plan.

What it includes

- **Assessment —** inventory of workloads, application dependencies and real consumption profiles, with cost analysis and evaluation of the options
- **Destination design —** target architecture, sizing and definition of the operating procedures
- **Sample trial —** pilot migration of a representative workload with measurement of actual throughput before planning
- **Automation —** infrastructure described as code, versioned and reproducible, for an environment that can be rebuilt deterministically
- **Cutover in batches —** progressive transfer with functional validation at the end of each batch and a defined rollback plan
- **Knowledge transfer —** written operating procedures, staff training and support during the period following go-live

> **Who it is for**
>
> Organisations leaving a proprietary platform, consolidating heterogeneous environments or renewing infrastructure that has reached end of life.

CLOUDOPS

VMware Migration

A dedicated path out of proprietary virtualisation

The change in licensing model following Broadcom's acquisition of VMware has made the cost trajectory uncertain for many organisations, prompting them to consider alternatives they would not have considered shortly before. This service addresses that specific transition, with three destination paths depending on scale and workload profile.

An honest assessment must also state where the destination platform does not offer the equivalent of what is being left behind. In the comparisons we have carried out, a few recurring divergences emerge — automated workload balancing, distributed intrusion detection, remote-access VPN — all manageable, all to be declared before the migration rather than after.

What it includes

- **Towards OpenStack —** the path for large private clouds, telecommunications operators and service providers, with multi-tenancy and complete APIs
- **Towards Proxmox VE —** the faster path for smaller organisations and for MSPs, on a hyper-converged architecture
- **Towards Kubernetes —** not a direct migration of virtual machines but an application modernisation path, adoptable selectively for the workloads that benefit from it, typically alongside one of the two paths above
- **Comparative analysis —** documented functional comparison between source and destination platform, with the divergences declared in advance
- **Workload conversion —** conversion of disks and drivers, mapping of networks and security policies, definition of the batches
- **Subsequent operation —** takeover of the destination platform under the managed services, if requested

> **Who it is for**
>
> Organisations with a VMware estate in service that are evaluating, or have already decided on, an exit from the proprietary platform.

CLOUDOPS

Ceph Support

Specialist support for production distributed storage clusters

Ceph is reliable but it is not simple, and the skills needed to run it are scarce on the market. Many organisations discover they do not have them in-house at the worst possible moment: when the cluster is degraded and critical workloads are at risk.

The service intervenes both in emergency and in prevention. Every engagement starts from a survey of the cluster's actual state — version, configuration, topology, data distribution, performance — and closes with a document setting out what was done and why.

What it includes

- **Incident management —** analysis and resolution of critical problems on production clusters, with agreed response times
- **Health check —** in-depth analysis of configuration, capacity, data balance and component status, with a report and corrective actions
- **Performance tuning —** tuning of pools, of the storage engine, of the cluster network and of internal maintenance policies
- **Upgrades —** planning and execution of upgrades, including across several major versions, validated on a test environment before production
- **Expansion —** addition of disks, nodes and hosts with capacity planning and controlled data rebalancing
- **Proactive monitoring —** continuous surveillance with alerting on anomalies, capacity thresholds and replication status: this belongs to a continuous service rather than to on-demand engagements, and is activated with Managed Services and 24/7 On-call

> **Who it is for**
>
> Anyone running Ceph in production: cloud and service providers, telecommunications operators, multi-datacentre enterprises, MSPs.

CLOUDOPS

OpenStack Lifecycle & Support

Specialist support and evolution of OpenStack platforms in production

An OpenStack platform does not end at acceptance: it lives for years, changes version, grows in nodes, acquires services it did not have at the start. This service takes on that part of the life cycle, on platforms built by us or by others.

Major version upgrades are the activity by which this service is really measured. They are prepared with a survey of the actual configuration, tried on a test environment, executed with a preventive backup and a defined rollback plan for each step, and verified against the services in operation before being considered complete.

What it includes

- **Health check —** survey of the configuration, the status of services, capacity and divergences from good practice, with a report and corrective actions
- **Troubleshooting —** analysis and resolution of malfunctions on production environments, with agreed response times
- **Version upgrades —** planning and execution, including across several successive versions, validated on a test environment before production
- **Expansion —** addition of compute and storage nodes, new availability zones or new regions, with verification of the impact on existing services
- **Tuning —** optimisation of overcommit ratios, network parameters, storage integration and scheduler behaviour
- **Onboarding of new services —** activation of components not present in the initial installation — load balancing, firewall and VPN as a service, DNS, orchestration, instance high availability
- **Operating documentation —** update of the as-built specification and of the procedures, as an output of the engagement rather than a separate activity

> **Who it is for**
>
> Organisations running OpenStack in production that do not have in-house skills for major interventions, or that want to complement them with external specialist support.

CLOUDOPS

Managed Services and 24/7 On-call

Continuous operation of the platform, with coverage levels defined per service

The service takes on the daily life of the infrastructure: surveillance, maintenance, upgrades, data protection and capacity planning. It is intended for organisations that want to use a cloud platform without first having to build the department that runs it.

The fee draws a clear line between two components. The recurring activities necessary to keep the platform running — monitoring, alerting, backup verification, application of security fixes, periodic reporting — are included. Activities requested by the client draw on a stated hours allowance, with a known rate for anything beyond it. It is the separation that prevents the silent drift in which a fixed-fee service becomes unsustainable for the party delivering it or inadequate for the party receiving it.

What it includes

- **Continuous monitoring —** automatic surveillance of the entire stack around the clock, with proactive alerting and event correlation: it is always active, regardless of the human coverage window
- **Maintenance —** patching and upgrades planned within protected windows, with preventive verification and a rollback plan
- **Data protection —** execution and periodic verification of backups and restore procedures
- **Operational security —** vulnerability management, configuration review and incident response following documented procedures
- **Capacity and cost governance —** periodic reporting on consumption, trend analysis and sizing recommendations
- **Single point of contact —** a service manager for service levels, governance and evolution roadmap, with a structured request management process

> **Three distinct levels, not to be confused**
>
> Continuous monitoring is automatic and covers the full twenty-four hours. On-call is the availability of an engineer outside service hours, with a response time agreed per severity class. Active operational staffing is the continuous presence of personnel on shift: it is a further level, sized and quoted separately, and it does not coincide with on-call.

Who it is for

Organisations without a dedicated operations team, operators that want to concentrate their resources on the product, groups outsourcing infrastructure operations.

Engagement models

The same skills are delivered through different contractual vehicles depending on the context. The skill is the constant, the vehicle is the variable: every engagement states from the outset which of these models it belongs to, what its scope is and who is accountable for what.

Model	When it applies	Commercial form
Fixed-scope project	Defined and verifiable scope: assessment, design, implementation, migration	Fixed price against milestones
Specialist consultancy	Technical support on existing environments, side-by-side working, discrete interventions	Days at an agreed rate
Continuous service	Operation of the platform and specialist support over time	Monthly fee with an included hours allowance
Technical subcontracting	Technical scope within a programme run by others	To be agreed with the principal contractor
Delivery under partner brand	Implementation and management proposed to the end client under the partner's identity	To be agreed with the partner
Training and operational handover	Autonomy of the client's team on an existing or newly built platform	Training days, side-by-side working and operating documentation

White-label work and subcontracting

A significant part of our activity has been carried out as the technical scope within programmes run by system integrators, under the brand of technology partners, or as subcontracting within consortia. It is a model we know and are equipped for: declared scope, formal documentation, direct technical dialogue and no visibility towards the end client where the context requires it.

Service levels and scope

Service levels are defined contractually case by case, according to the criticality of the platform, the coverage hours required and the architecture actually implemented. The principle we follow is to declare the limits together with the capabilities: the constraints section is part of the deliverable, not a concession.

Coverage	Defined per service: extended business hours, out-of-hours on-call, continuous staffing
Response times	Agreed per incident severity class and formalised in the operating annex
Continuity objectives	Recovery and data loss objectives defined at design stage and verified through periodic tests
Scope included	Infrastructure, platform and their components: compute, storage, network, identity, observability
Scope excluded	Application development, end clients' workloads and systems not covered by the operating annex
Reporting	Periodic reporting on availability, incidents, consumption and activities carried out

Why Epic Edge

Epic Edge grew out of experience in designing and running OpenStack, Ceph and Kubernetes infrastructure in critical production environments, at telecommunications operators, cloud providers, public administration, industry and supercomputing centres, in Italy and across Europe.

Production, not laboratory	Control of the platform	Complete life cycle
Real clusters, real incidents, real service levels. The skill comes from the field, not from a catalogue exercise.	Open-source infrastructure core, reduced supplier dependency, data and control within the organisation's perimeter.	Design, implementation, migration and operation: the same team from assessment to continuous support.

Let us talk about your environment.

From an assessment of the existing infrastructure to operational takeover: together we identify the service with the fastest return in your context and start there, with a contained and verifiable scope. andrea.martra@epic-edge.it · www.epic-edge.it

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