ONEedge5G

Intelligence and Automation for the Operation of Distributed Edge Systems on 5G Advanced Infrastructures

Edge and 5G in Spain and Europe State of Play and Future Trends Version 1.0 – August 2025

Abstract
This white paper examines how edge computing and 5G/5G-Advanced jointly enable low-latency, secure, and energy-aware digital services across Europe, with a focused deep dive on Spain. Combining desk research with an Expert Committee program of interviews (April–August 2025), the study maps the cloud–edge continuum (on-prem, metro/public edge, hyperscale interconnects), private-5G/5G-SA capabilities (including slicing and exposure APIs), and operational enablers (zero-touch automation, observability, security, and energy telemetry). Findings indicate that edge and 5G are mutually reinforcing, that adoption is ROI-driven and vertical-led (manufacturing, rail/transport, health, logistics, venues), and that public edge supply risks under-utilization without demand-anchored planning. The Spanish context highlights the catalytic role of Red.es in digitalization and AI, the maturation of private 5G pilots, and early, scoped Open RAN activity, alongside persistent needs in interoperability, portability, skills, and energy efficiency. Technically, the paper references the ONEedge5G software outputs as an illustrative platform for distributed edge orchestration and zero-touch operations. Contributions include a capability map for 5G edge, stakeholder-specific recommendations, an indicative 36-month roadmap, KPI frameworks for scale-up and oversight, key risk mitigations, and research/standardization priorities. Overall, the analysis argues for coordinated, open, and sovereign approaches that pair operational excellence with clear demand signals to move from pilots to production at a European scale.

Keywords: Edge computing; 5G/5G-Advanced; Open RAN; private 5G; OpenNebula; EdgeOps/MLOps; zero-touch automation; Red.es; Spain; Europe.

OpenNebula

Content

Glossary

5G	Fifth-Generation Mobile Network
AGV	Automated Guided Vehicle
AI	Artificial Intelligence
B5G	Beyond 5G
Ceph	Ceph Distributed Storage
CI/CD	Continuous Integration / Continuous Delivery
CCTV	Closed-Circuit Television
CN	Core Network
CNF	Cloud-native Network Function
DPDK	Data Plane Development Kit
DRS	Distributed Resource Scheduler
EdgeOps	Edge Operations
ENIA	Spain's National AI Strategy
FL	Federated Learning (in reference to AI)
IOMMU	Input-Output Memory Management Unit
IPCEI	Important Project of Common European Interest
KPI	Key Performance Indicator
KVM	Kernel-based Virtual Machine
LVM	Logical Volume Manager
MEC	Multi-access Edge Computing
ML	Machine Learning
MLOps	Machine Learning Operations
MoU	Memorandum of Understanding
NSA	Non-Standalone (EN-DC)
NUMA	Non-Uniform Memory Access
OEE	Overall Equipment Effectiveness
O-RAN	Open Radio Access Network
OneDRS	OpenNebula Distributed Resource Scheduler
OSS	Operations Support System
PCI	Peripheral Component Interconnect
PIS	Passenger Information System
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RIC	RAN Intelligent Controller
ROI	Return on Investment
RRF	Recovery and Resilience Facility
SA	Standalone (5G Core)
SBOM	Software Bill of Materials
SLA	Service-Level Agreement
SLO	Service-Level Objective
SME	Small and Medium-sized Enterprise
SRE	Site Reliability Engineering
SR-IOV	Single Root I/O Virtualization
TCO	Total Cost of Ownership
UEFI	Unified Extensible Firmware Interface
UP	User Plane
UPF	User Plane Function
vGPU	Virtual Graphics Processing Unit
VET	Vocational Education and Training (Formación Profesional)
VIM	Virtual Infrastructure Manager
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
WASM	WebAssembly
ZTP	Zero-Touch Provisioning

PART A. Introduction

1. Objectives of this whitepaper

The convergence of 5G and Edge computing is meant to revolutionize the industry, increasing productivity and enhancing business competitiveness. 5G, with its high speed and ultra-low latency, provides the backbone of connectivity. However, when 5G synergizes with Edge computing, bringing computation and storage closer to the data source, it makes possible a new generation of applications with the potential to transform the economy and bring prosperity and well-being to society.

Smart factories, connected vehicles, immersive AR/XR experiences, or personalized healthcare are just a few examples of these applications. Such innovative use cases are fueling the demand for the combination of 5G and Edge, with AI amplifying this demand, driving the need for edge-based processing to handle the massive amount of data generated by AI-driven applications. In this sense, the precise location and distribution of these 5G and Edge resources play a critical role in meeting this growing demand.

On the supply side, Europe is actively developing its Edge computing infrastructure, and Spain stands out with a significant contribution. The ambitious target of reaching 10.000 sustainable nodes by 2030 requires a joint effort from industry and governments across various sectors. Not only the “usual suspects” like telcos or cloud providers will be involved, but also other sectors like manufacturing, transportation, or energy will collaborate in the provision of infrastructure to spread the Edge network. Key technologies, including advanced processors, optimized algorithms, and AI-powered infrastructure management, are crucial for building and operating these networks efficiently. Open-source technologies will play a vital role in fostering innovation and interoperability.

This whitepaper aims to provide value to the reader by:

- Identifying the key use cases driving **demand** for combining 5G with the Cloud-Edge Continuum.
- Understanding the impact of AI inference, training, and fine-tuning on this **demand is crucial.**
- From a **supply** perspective, this analysis provides a comprehensive overview of the current state of 5G and Edge computing in Europe, with a particular focus on Spain.
- In **market (**or platform**)** terms, understanding the value and benefits of automated optimal resource placement, as well as the role of advanced 5G deployments, and AI in optimizing Edge infrastructure.
- Exploring the strategic role of open-source technologies in the Edge **ecosystem**.
- And finally, offering actionable recommendations for stakeholders across the ecosystem.

By understanding the complex interplay of demand and supply, and by adopting a collaborative approach, Europe and Spain can unlock the full potential of 5G and Edge computing, driving innovation, economic growth, and a brighter digital future.

1.1. OpenNebula - as a Key Technological Enabler for 5G Communications

OpenNebula is a **simple yet powerful open-source solution for building and managing enterprise clouds and Edge environments**. It combines virtualization and container technologies with multi-tenancy, automatic provision, and elasticity to offer on-demand applications and services.

OpenNebula offers a single, feature-rich, and flexible platform that provides unified management of IT infrastructure and applications, thereby avoiding vendor lock-in and reducing **complexity, resource consumption, and operational costs**. OpenNebula manages:

- **Any Application**: Combine containerized applications from Kubernetes with Virtual Machine workloads in a shared environment to offer the best of both worlds: mature virtualization technology and orchestration of application containers.
- **Any Infrastructure**: Open cloud architecture to orchestrate compute, storage, and networking driven by software.

- **Any Cloud**: Unlock the power of a true hybrid, edge, and multi-cloud platform by combining your private cloud with infrastructure resources from third-party virtual and bare-metal cloud providers such as AWS, OVHcloud, IONOS, and Equinix Metal, and manage all cloud operations under a single control panel and interoperable layer.
- **Any Time**: Automatically add and remove new clusters to meet peaks in demand, implement fault tolerance strategies, or meet latency requirements.

OpenNebula provides the necessary tools for running containerized applications from Kubernetes, while ensuring that your DevSecOps practices meet enterprise requirements. It enables organizations to easily adopt Hybrid and Edge Computing, allowing them to scale their Enterprise Cloud on demand with infrastructure resources from third-party Public Cloud and bare-metal providers, such as AWS and Equinix Metal. This disaggregated cloud approach enables a seamless transition from centralized private clouds to distributed, edge-like cloud environments. Companies can expand their private cloud by leveraging resources at cloud and edge data center locations, thereby meeting peaks in demand or the latency and bandwidth needs of their workloads. This approach involves a single management layer, allowing organizations to continue using existing OpenNebula images and templates, maintain complete control over their infrastructure, and avoid vendor lock-in.

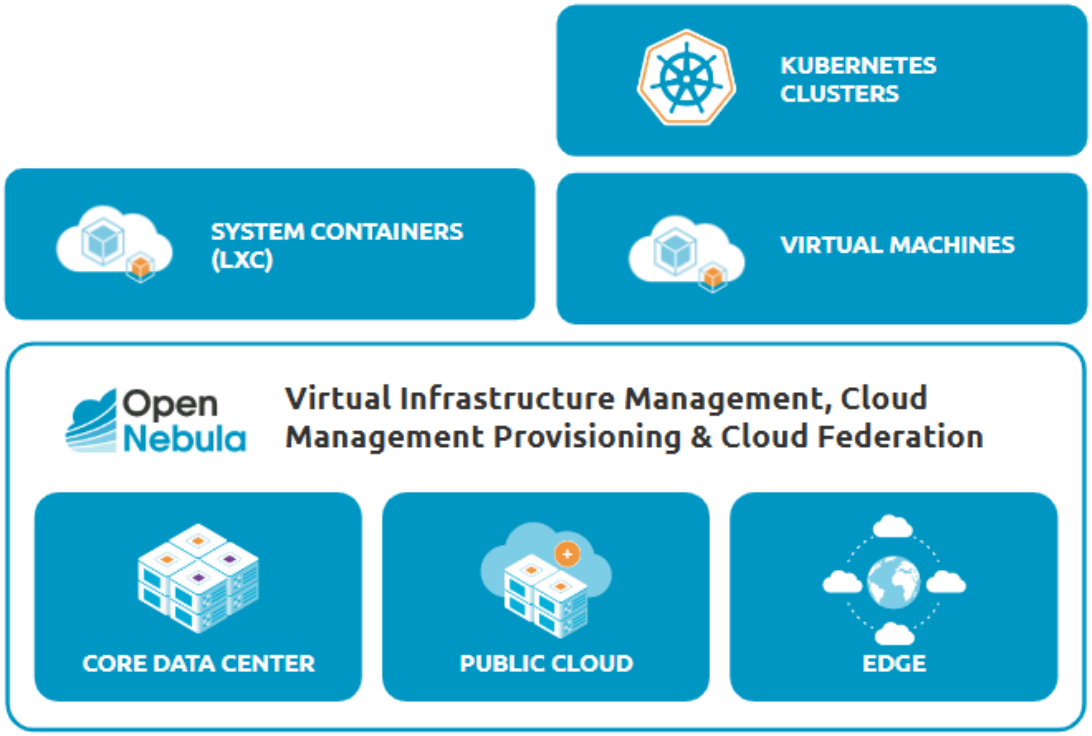

Figure 1. OpenNebula's flexibility.

The development of OpenNebula follows a bottom-up approach driven by the real needs of sysadmins, DevOps, and corporate users. OpenNebula is an **open source product** with a healthy and active community, commercially supported by OpenNebula Systems through its OpenNebula Subscription program. New versions are released regularly and delivered as a single package with a smooth migration path. OpenNebula defines its short-term roadmap and plans the features for the next release guided by the demands of its Sponsors, Customers, Users, and Partners. A detailed list of planned features for the upcoming release of OpenNebula is available at the GitHub OpenNebula/One issues page. For more information on the benefits of running an OpenNebula cloud, please visit the Key Features page.

The reader is invited to explore OpenNebula's diverse Case Studies and Success Stories to learn more from users about how they are leveraging OpenNebula.

1.2. The ONEedge5G Project: Strengthening OpenNebula Key Features for 5G

The **ONEedge5G** project focuses on AI-driven, **zero-touch** operations for highly distributed edge systems on **5G/5G-Advanced** infrastructures. Its goal is to harden OpenNebula for low-latency, mobility-aware, and energy-optimized deployments at scale, while keeping the platform open and sovereign by design. These objectives directly feed into OpenNebula's next-generation evolution. They are reflected in the ONEedge5G software outputs, which emphasizes intelligent automation, edge readiness, hybrid/multi-provider integration, and GPU-accelerated workloads for AI at the edge.

From a platform standpoint, **ONEedge5G** introduces a modernized control plane and runtime capabilities that are especially relevant for 5G and MEC use cases. Highlights include the **a Distributed Resource Scheduler (OneDRS)** for predictive, policy-based workload placement and automated migrations; **resource usage forecasting** to anticipate hotspots; **cluster-level quotas** for granular multi-tenancy across geo-distributed sites; and a consolidated **FireEdge-based Sunstone** UI as the standard web console for large fleet operations. These features enable tighter SLO/SLA control for latency-sensitive slices and services while simplifying day-2 operations at the edge.

On the **architecture** side, ONEedge5G's **Edge Cloud Reference Architecture**—built around deployable **Edge Clusters**—allows operators to run both **VM** and **Kubernetes-based container** workloads across on-premises and public cloud resources with unified management and portability. This model supports telco scenarios that need to place network functions and data processing close to users, while preserving hybrid flexibility and avoiding lock-in, key for European sovereignty and cost control.

For **performance-critical network functions** (e.g., UPF, user-plane analytics) and **AI inference** at the edge, it strengthens **platform awareness and hardware acceleration**. Native support for **PCI passthrough** (including SR-IOV/IOMMU) and **NVIDIA vGPU** enables high-throughput packet processing (DPDK/SR-IOV) and GPU-backed inference, while KVM/networking and storage enhancements (UEFI, NUMA-aware PCIe topology, NetApp/Ceph/LVM improvements, and Veeam integration) improve determinism and resilience in distributed environments. These capabilities align with ONEedge5G's aim to operationalize low-latency, energy-efficient edge services.

Importantly, the documentation recognizes that the development of **next-generation features** is supported by **Spain's Ministry for Digital Transformation and Civil Service** and the EU NextGenerationEU RRF through**ONEedge5G**. This confirms the close link between the project's research results and the platform's plans. Practically, this means the zero-touch, AI-assisted operations, multi-site orchestration, and energy-aware controls being explored by **ONEedge5G** are expected to be included in open-source projects, like OpenNebula, speeding up adoption in 5G edge deployments. docs.opennebula.io

2. Methodology

This whitepaper complements secondary sources with semi-structured interviews (April–May 2025) covering edge–5G convergence, use cases, drivers/barriers, operations/automation, standards, and EU targets. Interviews provide ground-truth on deployments, costs, operational pain points, and regulatory friction.

This whitepaper has been built upon the fundamental insights provided by an expert committee, whose members have been selected to ensure that all relevant stakeholders are represented and their different perspectives are captured. These stakeholders can be clustered considering six relevant groups:

1. Telecommunications operators;
2. Use cases;
3. Research and Technology Organizations (RTO);
4. Academia;
5. Infrastructure Providers; and

6. Vendors.

The committee counted with individual representatives of these organizations; the experts often share personal views rather than representing their companies' views and vision:

- Neutroon Technologies S.L.
- IKERLAN, S. Coop.
- Universidad Carlos III de Madrid
- Trebide, Ikusi, S.L.U.
- Fundació i2CAT
- Hospital Clínic de Barcelona

Use cases may be a transversal group, meaning that an RTO, for instance, could be a use case at the same time.

They received a survey carefully prepared with a set of questions addressing both demand and supply. Such questions were designed to collect relevant information to assemble the white paper.

The part about demand was made to obtain the following information:

1. Identify the most common use cases driving demand for Edge Computing and 5G, as well as the verticals poised to dominate related applications.
2. Identify the main features that are driving Edge deployments, as well as likely technology enablers.

The information provided by the stakeholders is complemented by that obtained from sources like the European Commission's Competitiveness Compass, that helps in understanding the importance of Edge to make the most of the possibilities brought by 5G, how AI can be used to optimize the performance of 5G and Edge and identifying the use cases driving such a demand with special focus on those sectors targeted by the European Commission as key for recovery and growth. Some examples of such sectors include digital, clean, and sustainable technologies; Manufacturing (Industry) 4.0; and education and skills (with a focus on capacity building), to name but a few. These use cases present a series of requirements that are translated into the capabilities of cloud-edge continuum resources, enabling the implementation of these use cases.

The questions related to the supply side aim at capturing the following information around the following areas:

1. Status of the Edge in Spain;
2. Digital targets for 2030;
3. Practical examples of Edge use cases and their development.

The information provided by the stakeholders is complemented by that obtained from sources like the Edge Observatory Report or GSA Reports.

The expert committee members were invited to a plenary session of approximately 30 minutes, during which all necessary context information was provided, and the survey was presented. They were given three weeks to digest the information and warm up, and subsequently, they were invited to interviews of around 30 minutes, whose objective was to fill out the survey. The respondents have provided a rich variety of perspectives on the different topics, with insightful and complementary visions that help in constructing the overall picture.

The demand and supply sides come together on platforms designed to provide solutions for the various challenges they pose.

With all these elements, we obtained an accurate analysis of the state of play in Europe and subsequently zoomed in and broke it down for the specific case of Spain. Additionally, we moved forward with identifying future trends on both the demand and supply sides.

3. Scope and structure of the document

3.1 Scope

This whitepaper examines how **edge computing** and **5G/5G-Advanced** can be combined to deliver low-latency, secure, and energy-aware digital services across Europe, with a **deep dive into Spain**. The analysis covers the **cloud–edge continuum** (on-premises edge, metro/public edge, and hyperscale interconnects), **network capabilities** (private 5G, 5G SA core features, slicing, and exposure APIs), and **operational enablers** (automation, observability, security, and energy telemetry). The scope explicitly includes **open and interoperable architectures** (e.g., containerized VNFs/CNFs, Open RAN domains where relevant), and **AI at the edge** (MLOps/EdgeOps) as cross-cutting capabilities. The primary **time horizon** is 2025–2028 (with forward-looking notes to 2030), focusing on **industrial automation, transport/rail, health, logistics, and venues**, and on public-sector use cases where data sovereignty is material.

3.2 Intended Audience and Use

The document is aimed at:

- **Policy-makers and public agencies** (EU, national, regional) that are evaluating instruments for cloud-edge/5G adoption, SME enablement, and sovereign infrastructure.
- **Network operators and neutral hosts** that are planning private 5G offerings, edge platforms, and operational automation.
- **Enterprises and SMEs** prioritizing ROI-anchored use cases and seeking guidance on architecture, governance, and scale-up.
- **Vendors, integrators, and research bodies** that contribute to open ecosystems and standardization.

Readers can use the whitepaper to **prioritize use cases**, **shape roadmaps**, and **align investment** across infrastructure, platforms, and talent. Where relevant, we provide **KPI frameworks** and a **36-month roadmap** to support execution.

3.3 Methodological Boundaries

Findings are derived from (i) **desk research** on European policy, market, and technology sources; (ii) **expert interviews** conducted for ONEedge5G (personal communications, April–August 2025); and (iii) **project-driven technical work** using **OpenNebula v7.0** capabilities for distributed edge orchestration and zero-touch operations. **Personal communications are cited in text but not listed in the references**. Unless otherwise noted, data and policy references reflect the state of knowledge as of **21 August 2025**. The whitepaper is **technology-agnostic** but emphasizes **open, portable** solutions to mitigate lock-in and to respect EU data-sovereignty requirements.

3.4 Structure of the Document

- **PART A — Introduction & Methods.** Sets objectives, context, and this scope; summarizes the methodological approach and how interview insights inform the analysis.
- **PART B — Europe: State of Play.** Maps EU-level drivers (policy, spectrum/licensing patterns, ecosystem, and standardization), supply-side capabilities across the cloud–edge continuum, demand signals by vertical, and emerging gaps (skills, portability, security/energy).
- **PART C — Spain: State of Play and Future Trends.** Provides national context, including the role of **Red.es** in digitalization and AI enablement, relevant trials/pilots (private-5G, MEC, and—where applicable—O-RAN domains), and **interview-derived conclusions** on challenges and opportunities.
- **PART D — Conclusions and Next Steps.** Consolidates strategic conclusions, **stakeholder-specific recommendations**, a **36-month roadmap**, KPIs, risks/mitigations, and priorities for research/standardization, and a brief outlook.

- **Annexes. Glossary/Abbreviations**, from interviews to findings, and auxiliary materials (e.g., interview guide).

3.5 How to Read and Apply This Whitepaper

Readers pressed for time can start with the **Executive Summary** and **PART D** for an action-oriented view (recommendations, roadmap, KPIs). Stakeholders planning deployments should pair **PART B/C** (context and trends) with the **Annexed KPIs and traceability** to ensure that pilots are **demand-anchored**, **portable across edge tiers**, and **operationally sustainable** (ZTP, observability, SBOM/signing, remote attestation, energy telemetry). Policymakers may focus on **program design** (SME enablement, neutral-host/co-investment models, licensing simplification, skills).

3.6 Limitations and Updates

The market and regulatory environment for edge/5G evolves rapidly. Performance claims depend on **workload profiles**, **site topology**, and **operational maturity**. While interview insights enrich the analysis, they reflect the **context and constraints** of the participating organizations at the time of discussion. The ONEedge5G team recommends periodic updates aligned with **major policy releases**, **standards milestones**, and **platform versions** (e.g., OpenNebula and related edge/5G components).

PART B. State of Play in Europe

Interviewed experts consistently emphasize that edge computing and 5G/beyond-5G are mutually reinforcing: localized compute is pivotal to achieving ultra-low-latency and high-reliability promises, while also reducing backhaul and strengthening privacy and data sovereignty (Interviewee 1, personal communication, April 22, 2025; Interviewee 2, personal communication, April 23, 2025; Interviewee 6, personal communication, May 21, 2025).

Observed architectural patterns include on-premises edge for regulated sites (e.g., hospitals, factories), multi-access edge computing (MEC) integrated with 5G for offloading heavier AI inference, and cloud–edge continuum strategies for dynamic workload placement (Interviewee 2, personal communication, April 23, 2025; Interviewee 4, personal communication, May 7, 2025).

Market-validated or piloted use cases include:

- Industrial automation & logistics (AGVs, robotics coordination, process control): wired + private 5G with on-prem edge for deterministic behavior.
- Construction sites: AR overlays, asset tracking, safety haptics (vests/jackets), and real-time analytics via private 5G + edge.
- Stadiums: ultra-low-latency video adaptation at the edge to improve accessibility for visually impaired fans (social-impact use case).
- Healthcare: surgical AR/VR and remote collaboration under strict on-prem privacy constraints; streaming/orchestration kept local.
- Railways: on-board edge gateways (CCTV, PIS, telemetry), AI predictive maintenance, and 5G backhaul; critical control segregated from non-critical services.

Primary motivations: privacy/data sovereignty and resilience (local fallback), deterministic latency/QoS (often with slicing), and operational efficiency from local pre-processing; energy/climate constraints increasingly shape design choices (Interviewee 2, personal communication, April 23, 2025; Interviewee 3, personal communication, April 30, 2025).

Key constraints: lack of a horizontal ‘killer app’ for mass public edge footprints; total cost of ownership (CapEx + OpEx); operational complexity at scale (zero-touch, updates, observability, security); fragmented spectrum/licensing; and portability gaps across the cloud–edge continuum (Interviewee 1, personal communication, April 22, 2025; Interviewee 2, personal communication, April 23, 2025).

Geography: comparatively higher maturity in Germany, the Netherlands, and Nordics; the UK is cited as policy-forward (e.g., neutral-host), while Southern Europe advances but faces bureaucratic and investment hurdles (Interviewee 2, personal communication, April 23, 2025; Interviewee 3, personal communication, April 30, 2025).

On EU targets (e.g., "10,000 climate-neutral, highly secure edge nodes" and "≤100-km availability" by 2030), experts stress that supply-led rollouts risk stranded capacity without demand-anchored planning and cross-stakeholder coordination; otherwise operators may repurpose capacity for internal workloads (Interviewee 1, personal communication, April 22, 2025; Interviewee 2, personal communication, April 23, 2025).

Although it is a well-known fact that Europe has been losing technological competitiveness in a scenario where non-EU technologies have progressively increased their share in the European market, only the publication of the so-called Draghi report managed to generate an intense wave of reactions and awareness-raising about this issue. The dominance of giants, primarily from the US and China, poses risks to Europe, as it creates dependency by relying on them to operate strategic infrastructures, granting them access to sensitive digital assets, and, above all, allowing enormous volumes of data to cross European borders and be stored overseas. This has a substantial potential impact on the security, freedom, and autonomy of European citizens.

The European Commission has acknowledged the need to regain competitiveness and strengthen the position of European technology providers. Europe possesses a solid, stable, and predictable legal environment, carefully built over the years in complete alignment with the human-centric values of the European Union. However, considering the weak position the Draghi report thoroughly describes, it becomes challenging to enforce this legal framework and apply the rule of law to the non-EU providers.

The Competitiveness Compass highlights the importance of investing in modern and state-of-the-art digital infrastructure. The Digital Networks Act proposes solutions to enhance market incentives for building these future networks, where Edge and, not only 5G but also B5G, will play a pivotal role.

In comparison to cloud nodes in centralized data centers, edge nodes are gaining momentum as they are physically closer to their intended users, are energy-efficient, reduce connectivity costs, and can enhance security protocols. The EU is making a remarkable effort to adopt, with significant investments, reflected in a notable acceleration in the deployment of Edge nodes. With an ambitious target of 10,000 Edge nodes deployed by 2030, the deployment rate has skyrocketed from 498 in 2022 to 1836 in 2024. This represents not only an exponential increase but also a significant step towards achieving 20% of the final target, with seven years remaining. Additionally, this effort is not focused solely on the dominant economies; on the contrary, the geographical figures indicate a distributed effort across the EU, benefiting from a single market.

Ensuring equitable deployment is a significant merit, given the EU's large and heterogeneous territory, which poses challenges at both Union and National levels. Some examples include the constraints posed by orography, climate, and the availability of infrastructure such as power and connectivity, as well as the maturity of existing networks. These latter two factors are closely tied to the country's wealth and its capacity for significant investments in digital infrastructure. In this sense, the European Commission considers targeted economic interventions in underdeveloped countries, disadvantaged areas, or in general in regions where the deployment is progressing more slowly than planned. In parallel, it is necessary to evolve both hardware and software to stay up to date, create scalable solutions, and integrate Artificial Intelligence and Machine Learning to enhance Edge Nodes' capabilities, thereby meeting the requirements posed by rollouts in different countries. Again, it is relevant to stress, as far as hardware is concerned, the dependence on components manufactured in countries like the US, Japan, South Korea, or China, which shows, once more, the current weak strategic European position. In addition, despite massive efforts to harmonize the legal and regulatory framework, a heterogeneous landscape persists across different countries in areas such as data protection, spectrum allocation, and environmental protection. This landscape

negatively impacts cross-border deployment activities and EU-level integration, increasing costs and complexity.

5G and Edge Computing are tightly coupled. 5G networks have improved connectivity and the availability of bandwidth to support more sophisticated services that enable the connection of more IoT devices and a significant increase in the volume of data generated, which, in turn, needs enhanced processing capabilities at the edge. This sets the scene for the definition and development of a myriad of use cases that will utilize these capacities. In this sense, the timely and successful deployment of the targeted number of Edge nodes across the EU is crucial, as they will impact specific sectoral agendas. One representative example is the energy vertical and the transition to a sustainable model, as edge computing will be integrated into energy management systems to enhance the efficiency and responsiveness of renewable energy grids. This is just an example; it is evident that this effort to boost the deployment of edge nodes is motivated by the existence of a flourishing European Edge Computing market, which, according to the Edge Observatory, is expected to grow from approximately €11.122 billion in 2022 to €34.363 billion by 2028, representing a CAGR of 20.6% from 2022 to 2028. Not only the energy sectors, but also a wide range of verticals, will be seduced by the undeniable advantages of Edge Computing and will ultimately integrate and adopt it in their daily operations, benefiting from the described infrastructure deployment.

Interviews reveal a two-sided gap: demand is currently concentrated in verticals with clear ROI (industry, rail, health, venues), while supply (public edge footprints) lacks a unifying 'killer app.' To bridge the gap, experts point to:

- Anchoring public targets to concrete demand models and measurable KPIs to avoid stranded capacity.
- Prioritizing zero-touch operations, observability, and security baselining to make distributed edge operable at scale.
- Standardizing workload portability/policies across the cloud–edge continuum and harmonizing private-5G spectrum procedures.
- Energy-aware designs and telemetry to meet climate-neutrality constraints while keeping TCO under control.
- Co-investment and neutral-host models to extend access beyond top customers and enable SME participation.

PART C. State of Play and Future Trends in Spain

Spain's progress on 5G and edge computing builds on a broad national digital agenda that aligns funding, spectrum, and public–private programs. Under Spain Digital 2026 and the Recovery and Resilience Plan, 5G coverage and standalone (SA) upgrades have advanced alongside experimentation with edge-enabled use cases across regions and verticals. This policy frame is the reference for regional pilots, SME enablement, and the integration of AI across public and private services.

On the ground, nationwide roll-outs by major operators have been complemented by flagship pilots that combine radio upgrades with localized compute and orchestration. These initiatives have accelerated learning on operational models (on-prem versus metro edge), partner ecosystems, and the scalability requirements for real-world deployments in manufacturing, mobility, public safety, healthcare, and venues.

A distinctive feature of Spain's approach is the emphasis on SME digitalization. Through programs such as **AceleraPyme** and the **Kit Digital**, thousands of smaller firms have accessed subsidized solutions and, more recently, AI-enabled capabilities. This demand-side vector increases the addressable market for edge-assisted applications —from analytics and automation to cybersecurity—while strengthening adoption readiness for 5G/edge services beyond flagship pilots.

Red.es has acted as a key execution arm for this agenda. It has financed and coordinated emblematic 5G pilots (e.g., urban mobility and public-safety trials). It manages SME programs (Digital Kit, Kit Consulting) that channel EU funds toward concrete digitization outcomes. Red.es has also supported data-space and AI

initiatives in coordination with the State Secretariat for Digitalization and AI, helping align connectivity with data governance, skills, and trustworthy AI.

Concerning 5G Open RAN (O-RAN), Spain has seen targeted activity rather than broad-scale deployment. Orange Spain has successfully demonstrated a 5G SA data connection using Open RAN in its Pikeo experimental environment in Madrid with partners such as HPE, Casa Systems, Mavenir, and Dell Technologies. Telefónica has pursued pre-commercial Open RAN pilots in Spain and signed a Cloud RAN MoU to evolve toward more open, disaggregated architectures. At the same time, Vodafone Spain and academic partners have explored 5G SA networks with O-RAN in testbeds. Commercial momentum is emerging—e.g., vendor agreements referencing Open RAN approaches—but nationwide roll-outs in Spain remain measured as operators balance maturity, performance at scale, and vendor diversity.

1. State of Play

The state of play compiled from the Expert Committee interviews is resumed here:

- Adoption of private 5G + edge is growing via pilots and focused deployments in industry, logistics, and mobility; traction is concentrated in large players, while SMEs face cost and operational barriers (Interviewee 3, personal communication, April 30, 2025).
- The predominant model is on-prem edge nodes connected to the operator backbone; commercial "metro/public edge" offers are still perceived as expensive and not truly pay-per-use for SMEs (Interviewee 3, personal communication, April 30, 2025).
- Railway stands out as an advanced vertical: on-board edge gateways with 5G backhaul for CCTV/PIS/telemetry and predictive maintenance, with strict segregation between safety-critical control and non-critical services; containerized AI models are distributed for on-board inference (Interviewee 4, personal communication, May 7, 2025).
- Spanish use cases mirror the European pattern: industrial automation, campus/lab networks, AGVs/logistics, and selected venue deployments; the key drivers are privacy, reliability/resilience, and deterministic latency—often coupled with slicing (Interviewee 2, personal communication, April 23, 2025).
- Operational frictions persist, notably the need for zero-touch provisioning/updates, observability, and security hardening for distributed fleets. Furthermore, local teams often lack mature tools and processes (Interviewee 1, personal communication, April 22, 2025).
- Regulatory compliance and data sovereignty shape on-prem architectures and component selection (minimal Linux, containers, vetted libraries); open source is a lever, but enterprise-grade support and certifications are required in regulated verticals (Interviewee 4, personal communication, May 7, 2025).
- Energy and climate-neutrality are emerging as design constraints for distributed nodes; energy telemetry and efficient AI pipelines are increasingly important (Interviewee 2, personal communication, April 23, 2025).

2. Future Trends

The future trends compiled from the Expert Committee interviews are summarized here:

- Growth will remain vertical-led and ROI-driven, with emphasis on rail, health, manufacturing, logistics, and venues; a horizontal 'killer app' is not yet in sight (Interviewee 1, personal communication, April 22, 2025; Interviewee 2, personal communication, April 23, 2025).
- Demand for local AI inference and MLOps/DevOps practices across distributed fleets will increase, driven by the need for efficient models, pruning/quantization, and hardware-specific compilation. Consider federated learning where privacy prevails (Interviewee 2, personal communication, April 23, 2025; Interviewee 3, personal communication, April 30, 2025).

- Automation is non-negotiable: zero-touch lifecycle, policy-driven orchestration, end-to-end observability, and near-real-time control loops (including Open RAN) are required to scale nodes across many sites (Interviewee 6, personal communication, May 21, 2025).
- Regulatory and public programs (e.g., IPCEI/CIS) and harmonized private-spectrum licensing could accelerate scale; without them, adoption will remain concentrated in large actors with on-prem edge that selectively interconnect to metro/public edge (Interviewee 3, personal communication, April 30, 2025; Interviewee 4, personal communication, May 7, 2025).
- Standardization for workload portability/interoperability across the cloud–edge continuum and neutral-host/co-investment models are levers to widen SME access (Interviewee 2, personal communication, April 23, 2025).
- Energy-aware orchestration and renewable-powered micro-datacenters are promising avenues to meet climate-neutrality objectives while keeping TCO under control (Interviewee 2, personal communication, April 23, 2025).

PART D. Conclusions

D.1 Strategic conclusions

- Edge and 5G are mutually reinforcing: localized compute is essential to deliver ultra-low-latency, reliability, and privacy/data-sovereignty benefits identified across interviews (Interviewees 1–6, personal communications, April–May 2025).
- Adoption is vertical-led and ROI-driven; industrial automation, rail, health, logistics, and venues concentrate near-term demand. A horizontal consumer “killer app” has not yet emerged.
- Public edge footprints will underperform if supply-driven; demand-anchored planning and coordination are needed to avoid stranded capacity and to align with EU objectives.
- Operational excellence (zero-touch lifecycle, observability, security by design, energy telemetry) is as decisive as footprint size for sustainable scale-up.

D.2 Recommendations by stakeholder

Policy & ecosystem (EU, national, regional)

- Tie node-deployment targets to concrete demand models/KPIs and prioritize co-investment/neutral-host approaches to extend reach beyond top customers.
- Harmonize private 5G licensing and simplify procedures for pilots scaling to production; promote interoperable, open stacks across the cloud–edge continuum.
- Back vertical programs (rail, health, manufacturing, logistics, venues) with procurement templates that include energy and security baselines.
- Expand SME enablement (e.g., vouchers, advisory) to include edge-AI, security posture, and observability toolchains.
- Network operators & neutral hosts
- Productize zero-touch operations (ZTP, declarative configs/GitOps, automated patching) and expose observability/SLA telemetry to customers.
- Stage Open RAN adoption via pre-commercial domains and non-critical layers while hardening performance, energy efficiency, and security at scale.
- Offer tiered edge options (on-prem, metro/public) with clear portability policies and price structures aligned to workload profiles.
- Adopt energy-aware placement and measure kWh per workload/inference; publish reference values for typical use cases.

Enterprises & SMEs

- Select ROI-anchored use cases first (e.g., predictive maintenance, quality inspection, computer-vision safety, real-time telemetry) and define target KPIs before deployment.
- Establish EdgeOps/MLOps: artifact signing & SBOMs, model lifecycle (drift monitoring, rollback), and hardware-specific optimization.
- Design for data governance and privacy (on-prem where required) and plan for workload portability across on-prem and metro/public edge.
- Build a talent bridge by upskilling OT/IT teams in observability, security, AI pipeline efficiency, and 5G QoS management.

D.3 36-month roadmap (indicative)

Phase	Key actions

0–12 months	• Identify 2–3 priority verticals per region and launch demand-anchored pilots with measurable KPIs (latency, availability, OEE gains, safety incidents).
	• Deploy a minimal edge platform:
	  ○ ZTP,
	  ○ CI/CD for edge workloads,
	  ○ centralized observability,
	  ○ signed artifacts/SBOMs.
	• Begin energy/CO_2 baselining for nodes and models; enable power telemetry and set initial efficiency targets.
	• Align with Red.es/SME programs and sector regulators to streamline compliance and funding.

12–24 months	• Scale to multi-site fleets; introduce policy-based orchestration and workload portability across on-prem and metro/public edge • Harden security posture (secure boot, remote attestation, key mgmt.) and automate patch SLAs • Expand private 5G coverage/QoS; integrate with slicing where available; connect to sectoral data spaces for cross-organization analytics. • Operationalize cost models (TCO/ROI), including energy per workload; renegotiate SLAs and pricing as utilization grows.
24–36 months	• Integrate public edge marketplaces; enable cross-domain roaming of workloads with policy and compliance enforcement. • Introduce Open RAN in selected domains with mature performance/energy/security metrics; evolve toward cloud-native RAN/edge control loops. • Transition from pilots to production by implementing 3–5 scaled use cases per vertical and institutionalize EdgeOps/MLOps along with governance boards • Publish benchmarks and share best practices across EU networks of competence; formalize talent pipelines.

D.4 KPIs for scale-up and oversight

	KPI
Technical QoS	• 95th-percentile latency, availability, packet loss/jitter, failover time; • per-workload energy (kWh) and carbon metrics.

Business impact	• OEE/throughput gains, • defect rate reduction, • downtime avoided, • time-to-decision, • ROI/TCO per use case.
Adoption:	• Number of sites • Number of nodes • Quantity of production workloads • Percentage of workloads portable across edge tiers • Percentage of automated rollouts.
Security & compliance	• Mean time to patch • Percentage of signed artifacts with SBOM • Coverage of remote attestation • audit pass rates.
AI performance	• accuracy drift, • re-training cadence, • inference cost per decision, • percentage of federated learning where applicable.

D.5 Risks and mitigations

Risk	Mitigation suggestions
TCO overrun and stranded capacity	Demand-anchored planning, phased investments, neutral-host sharing.
Operational complexity at scale	ZTP, standardized observability, automation of updates/rollback, and SRE practices at the edge.
Vendor lock-in/portability gaps	Open interfaces, container/VM/WASM options, contractually defined exit/portability clauses.
Talent gaps	Targeted upskilling, centers of excellence, joint curricula with universities, and VET programs.
Regulatory and security hurdles	Early regulator engagement, privacy-by-design, remote attestation, and SBOM policies.
Energy constraints	Lead to energy-aware schedulers, hardware acceleration when efficient, and renewable-powered micro-datacenters.

D.6 Research and standardization priorities

- Portable policy frameworks for cloud-edge workload placement with compliance and energy constraints.
- RIC/xApps/rApps for closed-loop optimization across RAN and edge compute; verifiable performance at scale.
- Privacy-preserving analytics (FL, secure enclaves) for cross-organization data spaces.
- Benchmark suites and certification for edge security posture (secure boot, attestation, SBOM) and energy metrics.
- Tooling for AI lifecycle at the edge (drift detection, on-device optimization, safe rollback).

D.7 Outlook

Europe's comparative advantages—privacy, data sovereignty, and resilience—align with edge-centric designs. Spain's mix of policy instruments, SME enablement, and targeted trials positions the ecosystem to

move from pilots to scaled operations over the next 24–36 months. As 5G SA matures and Open RAN capabilities harden, the edge will act as the execution fabric for AI-driven automation. Early movers that pair operational excellence with demand-anchored planning will capture outsized value and inform the next wave toward 6G and edge-native intelligence.

Acknowledgements

ONEedge5G (TSI-064200-2023-1) is supported by the Spanish Ministry for Digital Transformation and Civil Service through the UNICO I+D 6G Program, co-funded by the European Union – NextGenerationE through the Recovery and Resilience Facility (RRF)

Funded by the European Union NextGenerationEU

GOBIERNO DE ESPAÑA MINISTERIO PARA LA TRANSFORMACIÓN DIGITAL Y DE LA FUNCIÓN PÚBLICA

Plan de Recuperación, Transformación y Resiliencia

Sign up for updates at OpenNebula.io/getupdated

Rev1.0_20250821