COMILLAS
UNIVERSIDAD PONTIFICIA
ICAI

MASTER'S DEGREE IN INDUSTRIAL ENGINEERING

FINAL MASTER THESIS

Virtual Data Concentrator: Testing, validation, and documentation of containerization/virtualization architecture for telemanagement Data Concentrator application in distributed environments

Author: Emilio Juan Valdivielso Suárez

Director: Aurelio Sánchez Paniagua

Co-Director: Javier Matanza Domingo

Madrid

August 2025

This dissertation was typeset with LaTeX and compiled in TeXstudio using the TeXLive 2017 distribution. The font families used are Bitstream Charter, Utopia, Bookman and Computer Modern. Unless otherwise noted, all figures were created by the author using Microsoft PowerPoint®, Grafana®, Inkscape® and Python®.

AI tools were used to enhance the clarity and quality of the writing throughout this document.

Declaro, bajo mi responsabilidad, que el Proyecto presentado con el título

Virtual Data Concentrator: Testing, validation, and documentation of containerization/virtualization architecture for tele management Data Concentrator application in distributed environments

en la ETS de Ingeniería - ICAI de la Universidad Pontificia Comillas en el

curso académico 2025 es de mi autoría, original e inédito y

no ha sido presentado con anterioridad a otros efectos.

El Proyecto no es plagio de otro, ni total ni parcialmente y la información que ha sido

tomada de otros documentos está debidamente referenciada.

Fdo.: Emilio Juan Valdivielso Suárez Fecha: …15…/ …08…/ …2025…

Autorizada la entrega del proyecto

LOS DIRECTORES DEL PROYECTO

Fdo.: Javier Matanza Domingo Fecha: …15…/ …08…/ …2025…

Firmado por MATANZA DOMINGO JAVIER - ***4397** el día 14/08/2025 con un certificado

Fdo.: Aurelio Sánchez Paniagua Fecha: …15…/ …08…/ …2025…

Virtual Data Concentrator: Testing, validation and documentation of containerization/virtualization architecture for tele management Data Concentrator application in distributed environments

Emilio Juan Valdivielso Suárez, *Author*, Aurelio Sánchez Paniagua, *Director* and Javier Matanza Domingo, *Director*

***Abstract*—This project aims to technically and functionally validate the virtualization of the Data Concentrator (VDC) in distributed power grid environments, as an alternative to traditional embedded hardware. A comparative analysis was conducted between Minsait's previously tested platform and Barbara's edge computing solution, currently under evaluation at Iberdrola i-DE, as well as against the conventional embedded concentrator. The study assessed their ability to support critical substation applications.**

A complete deployment of a Barbara-based edge node was carried out on industrial-grade hardware. Functional equivalence with the embedded solution was demonstrated through tests involving configuration, interconnectivity, data acquisition, and scheduled task execution. Additionally, a voltage regulation use case using On-Load Tap Changers (OLTC) was developed and implemented directly on the virtualized node, showcasing the platform's ability to support new grid automation functionalities.

Scalability tests confirmed the VDC's capacity to manage over 7,000 smart meters without compromising operational integrity. An economic analysis further highlighted the feasibility of the virtualized approach under specific regulatory frameworks.

Beyond technical validation, this work contributes to the E4S Alliance by proposing interface definitions that enhance interoperability among edge computing components. The results lay the groundwork for modular, scalable, and interoperable architectures in future smart grid deployments.

***Keywords*—Distributed Processing Platforms (DPP), Edge Computing, Edge Node, Intelligent Secondary Substations, Interoperability, Smart Grids, Virtualized Data Concentrator.**

I. INTRODUCTION

Distribution networks are evolving into "smart grids", driven by digitalization and DER integration to achieve more efficient, sustainable, and secure power delivery. In this context, modernizing secondary substations (SSs), which connect medium-voltage networks to low-voltage consumers, is crucial. Traditionally, SS infrastructure was static and hardware-bound, now it faces challenges such as DER integration, predictive maintenance, and managing real-time operations with growing smart meter data.

Currently, SS functions rely on proprietary devices (e.g., data concentrators and remote terminal units) for aggregation, control, and monitoring. These vendor-specific systems are highly customized, hindering multi-sourcing and slowing adaptation to new requirements. For instance, i-DE's network manages ~120,000 secondary substations and ~11.6 million meters yet requires up to four years for a full upgrade cycle, impeding timely innovation. As a result, utilities are shifting to intelligent electronic devices (IEDs) on general-purpose hardware that consolidate DC, RTU, and low-voltage supervisory functions via virtualization.

Standards like IEC 61850 have improved interoperability by standardizing substation communication and data models, but they do not fully resolve the above issues. Therefore, the FUTURED consortium proposed a "Smart Automated Secondary Substation with Computing and Telecommunications Infrastructure" (SASCTI) architecture to guide this transition in [1]. Concluding that, virtualization and edge computing, offer a promising solution. On one hand, edge computing moves data processing to the SS, minimizing latency and enabling real-time local control. On the other hand, virtualization decouples software from dedicated hardware, allowing multiple substation functions to run on one generic platform. This consolidation reduces capital expenditures and improves flexibility for updates. For example, virtualizing the data concentrator (VDC) allows it to adapt to evolving communication protocols and new functions without retrofitting thousands of substations.

Accordingly, edge orchestration platforms such as Barbara IoT [2] and Onesait Phygital Edge by Minsait [3] have emerged to deploy and manage containerized applications across distributed nodes to enable secure and scalable deployment of SS applications. Industry alliances such as E4S are also working on common edge architectures to ensure interoperability.

In this context, i-DE is undertaking a project to validate virtualization of core SS functions, particularly the data concentrator, on an industrial edge platform (Barbara IoT). This approach is expected to enhance flexibility, interoperability, and scalability while reducing costs. However, realizing these benefits requires overcoming challenges in legacy integration, cross-platform interoperability, and communication security. This project addresses these issues by implementing and testing a virtualized SS architecture in an operational environment.

II. STATE OF THE ART

High penetration of DERs (e.g., electric vehicles and solar PV) is disrupting the traditionally one-way operation of distribution grids, prompting the smart grid paradigm to enhance monitoring, control, and flexibility by moving intelligence to the network edge, that is, to the secondary substations [4]. Although recent smart meter rollouts (mandated by EU policy) drove the deployment of telecommunication links and data concentrators in many SSs, their internal systems remain largely vendor-specific, hindering interoperability and making upgrades difficult. Therefore, many SSs now have remote monitoring and control via SCADA, yet still rely on rigid proprietary architectures that are costly to update and

Manuscript submitted August 13, 2025.

E.Valdivielso is with the Escuela Técnica Superior de Ingeniería (ICAI), Comillas Pontifical University, Madrid, Spain (e-mail: emiliovaldivielso-suarez@alu.comillas.edu).

lack local intelligence for rapid fault response. To address this, virtualization decouples software from hardware. This allows functions like the data concentrator and RTU to run as software on generic edge devices, greatly increasing flexibility.

Virtualization in SSs can be implemented either with virtual machines (VMs) or using containers. Virtual machines function under a hypervisor providing strong isolation for critical operations, but imposing higher resource overhead. While containers, are lightweight and easily scalable, which suits many SS applications as studied in [5]. Besides, field demonstrations such as the one carried out at [6] have validated these technologies, where hundreds of virtual protection IEDs were deployed on one server with no performance degradation, meeting IEC 61850 real-time requirements. They also showed that processing high-speed substation data at the edge can off-load the central SCADA. Similarly, an i-DE pilot [7] demonstrated that local edge computing in SSs can handle real-time load balancing and improve fault response using containerized applications.

Furthermore, a comprehensive Request for Information (RFI) analysis conducted by Iberdrola in [8] defines the essential architectural components that a DPP should have to fulfill all the technical requirements that any electrical utility needs. According to the document, the DPP should be divided in a Platform Software (PS), and Operation Administration and Management System (OAMS) Agent for the PS, an Application Runtime Environment (ARE), an ARE agent, and an OAMS for the ARE software. This RFI analysis gathered and classified some of the best commercial and open source DPPs and identified key players such as Barbara IoT, Minsait, Canonical and Linux-based platforms.

In this context, the E4S alliance is defining requirements that standardized hardware-agnostic edge computing architecture for SSs must fulfill to achieve full interoperability between different vendors [9]. Some of these requirements include the definition of interfaces between different softwares or the use of already built interfaces like the ones presented on initiatives such as OpenAPI 3.0 [10] and the W3C Web of Things [11].

III. Project Definition

The main objective of this project is to technically validate a **Virtual Data Concentrator (VDC)** implemented on an edge platform, and to propose a modular architecture that facilitates its integration into distributed and multi-vendor environments aligned with the E4S initiative.

To achieve this general objective, several specific sub-objectives were defined:

1) **Edge platform analysis:** Critically study the Barbara IoT platform in a laboratory setting, evaluating its architecture, its capacity to orchestrate applications on distributed nodes, and its performance under real operating conditions. A technical comparison was also carried out with another edge computing platform (Onesait by Minsait) to identify strengths, limitations, and suitability for smart grid requirements.
2) **Deployment documentation:** Document the complete process of implementing the VDC in an OT environment, including the installation of the Barbara OS operating system on generic industrial hardware and the configuration of the edge node. This includes the adaptation of the Barbara platform to i-DE's infrastructure and the connection with field devices.
3) **Modular and interoperable architecture:** Define the requirements and architecture that a distributed processing platform must meet to be vendor-agnostic, with an emphasis on interoperability, agent migration, and alignment with emerging frameworks such as the E4S alliance. A modular node architecture is proposed, with standardized interfaces that allow components from different manufacturers to be replaced without altering the system, in line with the E4S vision.
4) **Functional validation in pilot:** Deploy a pilot prototype of the VDC application on an edge node and verify its technical compatibility with the existing measurement infrastructure based on PRIME. This involves establishing and testing bidirectional communication, northbound with the remote management system (RMS) and southbound with the smart meters through the base node, as well as ensuring that the VDC correctly executes the scheduled reading and management tasks as a physical concentrator would. Additionally, develop an advanced use case that leverages edge computing, demonstrating an automation functionality (OLTC voltage regulation) supported by the VDC.
5) **Economic study:** Evaluate the economic feasibility of scaling the edge computing solution within the i-DE network, comparing the scenario of edge nodes + virtual VDC against the traditional deployment of physical concentrators (TGUC), under the regulated remuneration framework in Spain. Estimated capital expenditures (CAPEX) and operational expenditures (OPEX) for both scenarios are analyzed, and financial indicators (NPV, IRR) are calculated considering different regulatory assumptions regarding software remuneration.
6) **Scalability study:** Evaluate the scalability offered by the data concentrator by subjecting it to a stress test to determine the limit of supply points it can manage without saturating the device's resources.

This project is highly relevant in the context of smart grids, as the validation of a virtual concentrator represents the testing of a fundamental component for future distributed automation in low-voltage networks, and the proposal of a modular architecture for edge computing platforms enables the coexistence and incorporation of new advanced services or the rapid adaptation to new regulatory standards.

IV. Study and Analysis

A. Barbara vs Onesait/Minsait Comparison

As part of the project, two industrial edge computing platforms tested by i-DE were evaluated: the solution from Barbara IoT and Minsait's Onesait Phygital Edge. Both offer an environment for deploying and managing applications on edge nodes located in substations, but they present relevant architectural differences:

Minsait uses a central Edge Management System (EMS) installed at Iberdrola i-DE, which acts as a web management console and provides all backend services for remotely managing the edge nodes. This EMS maintains bidirectional control-plane connectivity with each node through secure channels, allowing operators to monitor status, deploy containers, perform firmware/software updates, configure settings, and even access the nodes remotely via console (SSH). Additionally, the Onesait architecture incorporates a central multiprotocol IoT gateway (based on an EMQX broker) that could be integrated with sensors or RTUs, supporting field-to-center protocols such as MQTT-SN, CoAP, LwM2M, etc.

On the other hand, Barbara deploys a node management server known as the Barbara Panel, also within i-DE's corporate network, which provides the central administration interface and is designed for operation in isolated or offline environments. The Barbara Panel is installed on a virtual machine in i-DE's industrial DMZ, serving as the central monitoring point for edge nodes. Both systems separate the control plane from the data plane, but in Barbara all field data is processed on the node, and only aggregated results are sent if necessary, avoiding acting as a gateway for all data to the center.

Both Onesait and Barbara use secure MQTT messaging for communication between the central manager and the nodes, but with different philosophies. In Onesait, the EMS uses an MQTT broker (RabbitMQ) to exchange commands and data with the nodes in a publish/subscribe manner, and also integrates advanced services such as a reverse proxy (Traefik), a Git repository for configurations, a private Docker container registry, OAuth2 authentication, etc., all to support multiple IoT use cases from the central platform. Barbara, for its part, includes an internal MQTT(S) broker in the Panel focused on the control plane, assuming that field data management is handled directly by the applications on each node. Consequently, Barbara does not include a multiprotocol sensor hub in the Panel, but instead delegates the connection to meters or sensors to the nodes themselves.

In summary, Minsait/Onesait presents an edge computing architecture with a centralized control plane rich in services and possibilities for central IoT data integration, ideal for large-scale corporate scenarios, offering a complete and closed service while aiming to align with E4S requirements. Barbara, on the other hand, presents a more open solution intended for the client to customize their solutions using the platform. Both platforms demonstrated compliance with key requirements (container deployment, remote management, secure MQTT/TLS communications), although each option has its pros and cons. The features of the Onesait platform are more aligned with Iberdrola's needs, although it is not as flexible in allowing the entry of different application providers as Barbara is, since Minsait aims to offer a complete service.

B. *On-Premise Deployment of the Barbara platform*

Given i-DE's strict cybersecurity policy, the Barbara IoT solution was deployed in an on-premise mode within the company's OT corporate network, instead of using its cloud version. This required adapting the Barbara software, originally offered as SaaS, to a local implementation that meets Iberdrola's OT requirements. The Barbara team, in coordination with i-DE, carried out custom modifications to install the Barbara Panel on a virtual machine within i-DE's environment, ensuring compliance with the IEC 62443 industrial security standards.

The deployment architecture was segmented into three network zones separated by firewalls. Zone 1, or the OT zone, is where the field devices reside (edge nodes, PRIME base node, existing concentrators, etc., in the substation or laboratory); the second zone is an industrial demilitarized zone (DMZ) that hosts the virtual machine with the Barbara platform (Panel and associated services); and the third zone is the corporate IT zone from which authorized users access the Panel's web console via their PCs. This segregation ensures that the control Panel is accessible only internally and that the edge nodes communicate with the Panel exclusively through the DMZ, avoiding direct exposure to the Internet.

Since the Barbara platform originally assumes cloud connectivity for certain functions (container downloads, updates, etc.), several actions were carried out to operate it fully offline, such as preloading all necessary packages, Docker images, OS updates, and certificates locally into the VM, transferring them securely under the supervision of i-DE's cybersecurity team. Key services such as the container registry, the OTA update service, and the certificate authority were replicated in the local instance of the Panel, so that the system could operate without relying on external components. Likewise, communication with the edge nodes was secured at the application layer using mutual TLS. Each node holds a device-specific client certificate and connects to the Panel's management services over encrypted and mutually authenticated TLS sessions, using predefined internal domain names. The Panel server was configured with two network interfaces (one in the OT subnet and another in the DMZ) and with specific DNS and routing settings to properly steer management traffic, allowing a single instance of the Panel to serve both domains while preserving network segmentation.

After these adaptations, the on-premise deployment successfully replicated the functionalities of the cloud solution within i-DE's segregated environment. A system was achieved with secure end-to-end communication and full local operational control of the edge nodes and their applications, without compromising security. This test environment laid the groundwork for the physical installation of the edge nodes in i-DE's laboratory and their integration with the actual equipment of the Secondary Substation.

C. *Proposed Modular Architecture for the E4S*

E4S proposes a layered model for edge platforms in secondary substations, aiming for modularity and interoperability as key attributes. Based on this reference, the existing E4S architecture, shown in Fig. 1, was analysed, and improvements were proposed in order to contribute to the definition of an architecture as interoperable as possible.

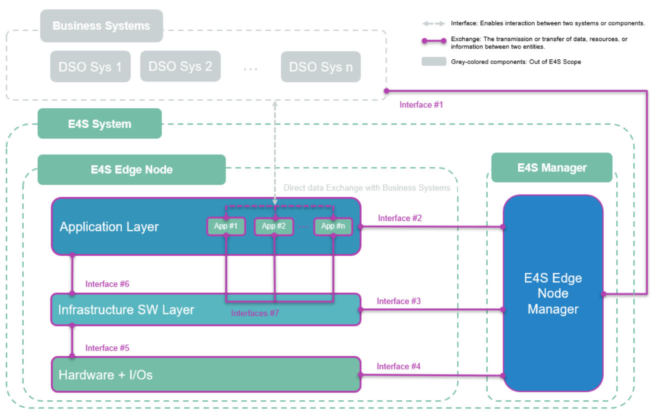

Fig. 1: Schematic representation of the architecture proposed by E4S from [12].

In the conception of E4S, the system is divided into two main subsystems: on the one hand, the Node Manager (central edge management platform), and on the other, the Edge Node (industrial device deployed in the field). The Edge Node, in turn, consists of several layers: the hardware layer, which defines the physical device and its I/O resources; the infrastructure software layer (or virtualization layer), which acts as an orchestrator of the hardware resources for the applications (including the operating system and base software that manage containers); and the application layer, where the DSO's operational applications reside, usually encapsulated in containers. The Node Manager, for its part, supervises all these layers in each node and coordinates their operation at a central level.

The reason for adopting an architecture with well-defined layers is to enable massive deployments of edge nodes while ensuring that each layer fulfills a specific function and that the components at each level can interoperate and be interchangeable without vendor dependency. In this way, the aim is to avoid vendor lock-in and promote standardization and competitiveness in the devices, which in the long term reduces costs. To achieve this, E4S identifies a set of 7 standardized interfaces between the layers and subsystems, numbered from 1 to 7. Each interface specifies the communications or interactions between adjacent components, promoting these connection points as open and common for any implementation.

The project first carried out a detailed analysis of the E4S interfaces defined in the original proposal, evaluating their purpose and importance within the operation of edge computing. Subsequently, adjustments and extensions were proposed to improve the architecture based on the experience gained from the implementation of the VDC, such as clearly separating the operating system layer from the node agent layer (local management software) within the "infrastructure layer" of the edge node.

E4S initially grouped the OS and the agent in the same layer, but the design proposed in this work separates them into distinct sequential layers (OS below, Node Agent above), which improves clarity and allows for a better definition of responsibilities and interfaces for each. This separation effectively introduces a new interface between the OS and the node agent (proposed as Interface 6 in our diagram), which would cover low-level interaction.

In addition, the definition of Interface 7, that defines the interaction between the node agent and the applications and Interface 8, that defines the interaction between different applications (not yet specified by E4S) was explored. Fig. 2 shows the proposed model, including the defined E4S interfaces and the new or modified ones.

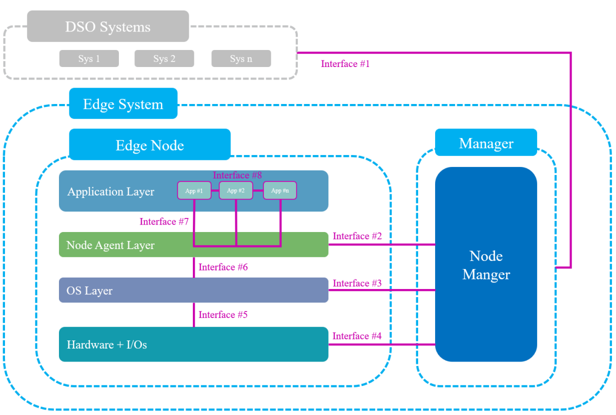

Fig. 2: Schematic representation of the proposed architecture

Three critical interfaces resulting from this improved architecture have been analyzed:

1) Interface 6: OS Layer - Node Agent: This interface enables communication between the Linux-based OS and the Node Agent, exchanging telemetry data such as CPU/memory usage, device status, events from the OS, and control commands or configuration requests from the Node Agent. Three communication protocols were evaluated:

- **Netlink sockets:** Efficient kernel-to-user-space communication with structured messages. Despite being Linux-specific, it is highly efficient and well-supported but may require custom kernel modules.
- **Unix domain sockets (UDS):** Effective user-space IPC method, offering high throughput and simplicity. However, it lacks predefined message framing, requiring custom schemas.
- **D-Bus:** Structured IPC message bus with built-in support for synchronous and asynchronous signals, widely adopted in Linux environments. Nevertheless, it incurs higher latency and moderate overhead.

Recommendation: Netlink sockets are recommended due to their efficiency, low overhead, and strong kernel integration.

2) Interface 7: Node Agent - Application Layer: This interface is critical for application lifecycle management, providing runtime services (logging, configuration distribution, health monitoring), and container control. Evaluated protocols include:

- **REST/gRPC APIs:** Structured and secure, language-agnostic, but introduce complexity in applications.
- **D-Bus:** Efficient local messaging but less suitable due to language constraints, scalability issues, and debugging complexity.
- **MQTT Broker:** Lightweight, robust asynchronous communication that clearly decouples applications from the agent but requires careful management of security and topics.
- **Container Orchestrator APIs (Docker):** Leverages mature Docker APIs, simplifying implementation. However, it restricts direct runtime interactions and real-time configuration changes.

Recommendation: MQTT combined with Docker APIs is proposed due to its simplicity, resilience, and strong decoupling.

3) Interface 8: Application - Application: This interface ensures semantic interoperability among applications, standardizing how data like measurements, control commands, and alarms are exchanged. Two primary approaches were analyzed:

- **Sparkplug B + CIM:** Efficient binary protocol with structured payloads and MQTT topics, widely adopted but requires predefined semantic models.
- **Web of Things (WoT) + CIM:** Provides rich semantic interoperability and dynamic discovery via JSON-LD Thing Descriptions, offering flexibility but increasing complexity and overhead.

Recommendation: Given the complexity and customization expected in secondary substations, WoT + CIM is recommended for its flexible semantic integration capabilities.

In summary, this modular architecture enables a DSO to change the "concentrator agent" (VDC) or add other containers (e.g., a low-voltage network control agent) without changing the underlying platform, as long as standard interfaces are respected. This enables multiple applications to coexist on the same edge node and allows different manufacturers to provide interchangeable components (for example, a VDC from another vendor) in compliance with interoperability standards.

D. Functional Validation, OLTC Use Case, and Scalability Tests

Once the infrastructure was deployed (Barbara on-premise platform and edge nodes with VDC), its operation was thoroughly validated through laboratory tests and field pilots. Three validation areas were considered.The first consisted of functional laboratory tests of the VDC;The second involved the development of an advanced OLTC use case implemented in a real secondary substation;And the third was a scalability test by connecting the VDC to multiple field substations simultaneously.The following sections describe these tests in detail, along with the connection architecture used in each case.

Functional tests in the laboratory: The initial objective was to verify that the virtual VDC can perform all the basic functions of a traditional physical concentrator. To this end, an edge node with the VDC was connected to an i-DE laboratory environment composed of a PRIME base node (a physical Circutor concentrator used only as a PLC modem) and several smart meters. Fig. 3 shows this configuration: the VDC (on the edge node) communicated via Ethernet with the base node, and in turn with 9 low-voltage meters installed in the laboratory, while the Barbara Panel in the DMZ managed the node.

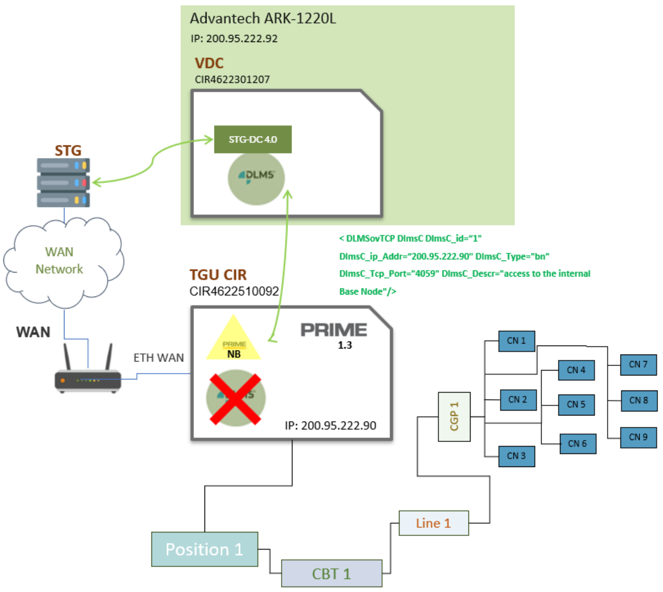

Fig. 3: Diagram of the VDC connections in the laboratory environment.

A sequence of three tests was defined:

1) *Connectivity and configuration test*: consisting of requesting an S12 report from the VDC via the central system to verify its factory configuration, and then sending a B07 command for initial parameterization (setting the IPs of the STG and FTP, as would be done with a new concentrator). Prior to this, a B07 was sent to the physical concentrator to disable its internal service and use it as a simple base node.
2) *Measurement reading test*: consisting of verifying that the VDC can periodically read data from meters. Reading commands were simulated (e.g., S24/G24 for device discovery on the PRIME network) and measurement queries (S21 reports for load profiles) to confirm that the VDC receives data from the meters through the base node and forwards it to the remote management system (STG) or stores it in FTP, as appropriate.
3) *Scheduled tasks test*: consisting of verifying that the VDC autonomously executes periodic tasks (as a physical concentrator would when collecting daily readings, billing closures, etc.). Typical tasks were scheduled, and it was verified through event logs that the VDC executed them on time and correctly interacted with the meters.

OLTC field use case: To demonstrate the potential of running advanced applications at the edge, an On-Load Tap Changer (OLTC) voltage regulation use case was developed on a smart transformer. The test was carried out in a real Secondary Substation in northern Spain equipped with a transformer with OLTC. In this substation, the existing physical concentrator (a ZIV) was used as a PLC base node, its concentrator functionality was disabled, and the laboratory VDC was remotely connected to the base node of that substation. Fig. 4 shows this architecture, where the edge node with the VDC remained in i-DE's laboratory (Madrid), communicating through i-DE's network with the remote substation. The purpose of this first scenario was to implement locally at the edge an algorithm that made tap change decisions on the transformer based on voltage measurements from the meters, without relying on the cloud or human intervention, achieving near real-time control.

The MV/LV substation used for the development of the use case had 175 meters, from which a preselection of 5 representative meters was made for monitoring. The selection criteria included: choosing only three-phase meters (400 V) since they are associated with larger loads or act as supervisors of several single-phase meters; selecting at least one meter from each of the 5 LV outputs of the transformer to cover the entire area; excluding meters with unreliable communication; and prioritizing those supplying many users.

Taking these selection criteria into account, the 5 optimal devices were identified, and a containerized Python application was developed on the edge node to implement the OLTC control logic. This application was designed to periodically request S21 load profile readings from the VDC for the five meters (simulating a periodic order from the central system), obtain the measured phase voltages, and calculate the average three-phase voltage of the monitored points. If the average deviates from the nominal value (230 V) beyond a tolerance band of $\pm 3\%$, the algorithm decides a tap adjustment by determining whether the tap should be increased, maintained, or decreased. This decision is recorded as a tap up/down/no-change command.

The algorithm runs every 5 seconds, and each cycle stores both the voltage values of each meter and the calculated tap command in a time-series database (InfluxDB) installed on the edge node. In parallel, a container with Grafana was deployed on the edge, which queries the Influx database and visualizes in real time the voltages and tap decisions.

Fig. 4 schematically shows how the virtual VDC reads data from the meters (via base node), the control script processes these readings and decides the OLTC action, the data is stored in InfluxDB, and Grafana displays the voltage trends and tap commands—all within the edge node. This distributed architecture enables real-time monitoring and control directly at the MV/LV substation, without relying on the latency of communicating each reading to the central system.

Scalability test: In this scenario, the VDC's ability to manage multiple substations in parallel was evaluated, pushing it to its

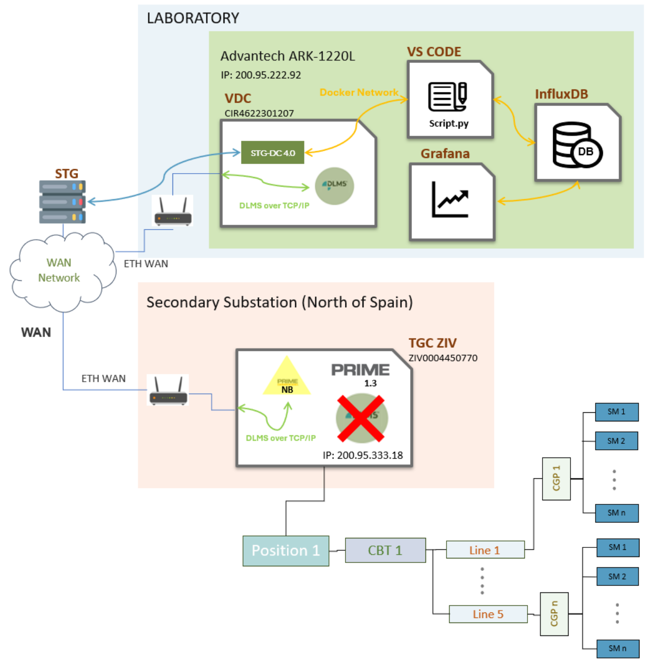

Fig. 4: Logical connections involved in the VDC use case scenario, showing the interaction between components.

theoretical limit. In the test, 12 real secondary substations from the urban area of Madrid were connected, selected among those with the highest number of subscribers (900 meters each). All these MV/LV substations were simultaneously linked to the VDC instance running on an edge node in the laboratory. Fig. 5 represents this configuration, where the VDC acts as a multi-substation virtual concentrator.

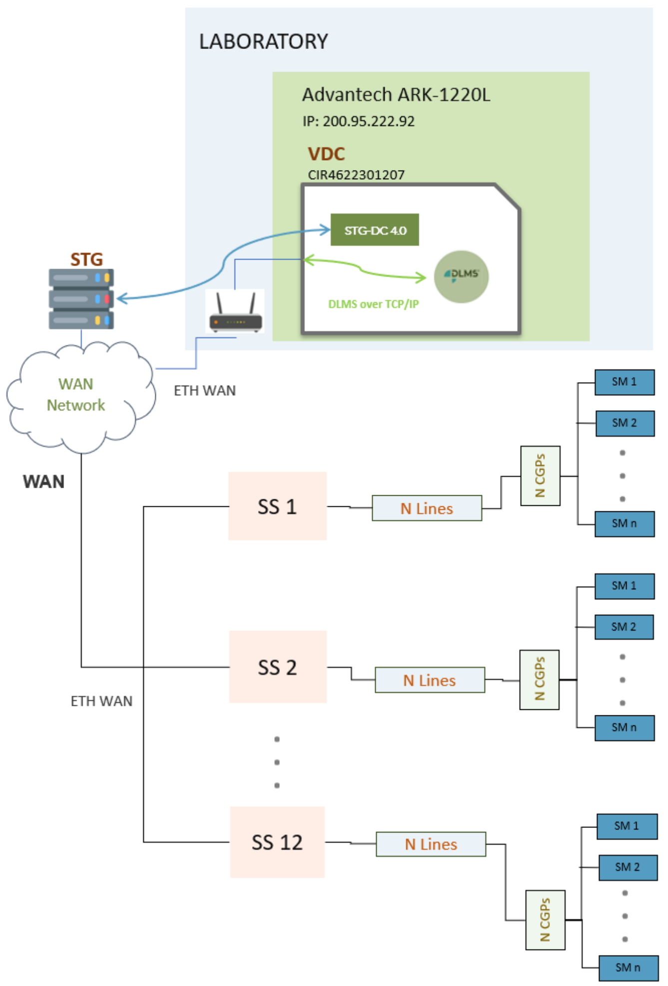

Fig. 5: Connection Architecture Used in the VDC Scalability Evaluation

V. Results

A. Functional validation of the VDC

Functional validation tests confirmed that the VDC deployed on the edge node correctly reproduces all the essential functions of a physical concentrator, with minor exceptions that are solvable. Table I summarizes the main tests carried out and their results.

TABLE I: Summary of Validation Tests

Category	Test	OK	Reported
Config.	S12 – Factory config report	✓	✗ (Missing `ipCom2`)
	B07 – Disable DLMS	✓	–
	B07 – Initial config	✓	–
	S24/G24 – Meter discovery	✓	✓
Meas.	S21 – Meter 1 (no load)	✓	✓
	S21 – Meter 7 (60 W)	✓	✓
Tasks	S05 – Daily billing	✓	✗ (STG issue)
	S02 – Hourly load profile	✓	✓
	S04 – Monthly closure	✓	✓ (STG issue)
	S09 – Meter events	✓	✓
	T01 – Reboot	✓	✓
	T07 – Time sync	✓	✓
	S17 – DC event log	✓	✓
	S24/S11 – SM list / BN stat	✓	✗ (no BN)
	S14 – Voltage/current prof.	✓	✗ (no LVS)
	G01/02/12 – Stats/history	✓	✓

In summary, the VDC was able to receive and respond to configuration commands (S12) and parameterization orders (B07), discover meters on the PRIME network (S24/G24), query hourly measurements (S21), and enter data into the central system, as well as execute scheduled daily, hourly, and occasional tasks autonomously. All test cases yielded satisfactory results in terms of execution in the VDC, and in most of them the report reached the remote management or target storage system as expected. For example, it was verified that the VDC, after an S24 order, correctly detected all meters connected to the base node, and that in response to S21 reading queries, the values were stored on the FTP server. Likewise, in periodic tasks it was observed that the VDC promptly initiated each request and transmitted the responses from the meters.

Some minor incidents were identified during these tests, but they were not attributable to VDC failures, rather to limitations of the test environment. For example, the initial S12 report was rejected by the RMS due to missing mandatory fields (ipCom2, ipMask2) in the configuration, which could be easily resolved by adjusting the specification so that these fields are not required or by filling them with default values.

Also, in the scheduled tasks, certain reports such as S05 and S17 were not stored in the development STG due to known issues in that environment, which would not occur in the updated production environment. In any case, no fundamental technical impediments were found; the VDC proved to be technically viable as a replacement for physical concentrators, requiring only minor adaptations to the existing specification. This is a crucial result, as it clears doubts about the basic functionality of the VDC within the i-DE infrastructure. In conclusion, the laboratory tests validated the feasibility of deploying VDCs as replacements for physical DCs, fulfilling the expected functionalities once configuration details are resolved.

B. OLTC Use Case Results

In the OLTC pilot, the VDC's ability to enable real-time control at the edge was evaluated. During the field test (CT with i-Trafo), the control algorithm at the edge node was active for several hours, continuously monitoring the voltages of the 5 selected meters and making decisions regarding the transformer's tap. Through the Grafana interface, both the three-phase voltages measured at each point and the tap adjustment decision calculated in each iteration were observed in real time. Fig. 6 shows how initially the transformer was in the nominal tap position, with voltages within range. Then, in response to a voltage increase in certain phases beyond the ±3% tolerance, the algorithm detected the deviation and began issuing commands to lower the tap in order to reduce the supplied voltage. The graphs show how, after the upper threshold was exceeded, the algorithm repeatedly recommended lowering the tap by one step.

It is important to note that, for operational and safety reasons, the commands were not physically applied to the transformer's actual OLTC during the pilot. That is, the system did not execute the tap change on the device, but only simulated it by recording it. Due to this lack of closure in the control loop, the measured voltages remained high and the algorithm continued issuing the same recommendation in each cycle. Despite this, the pilot demonstrates that the VDC and the edge application are capable of obtaining real-time measurements from the meters, processing them locally with a custom algorithm, and generating control actions with very low latency. The entire process occurred at the edge node, confirming the feasibility of implementing distributed control in the CT using the VDC infrastructure. The results visualized in Grafana showed how the system would dynamically react to voltage conditions. This use case sets a precedent for advanced applications at the edge, which can be added on top of the modular platform.

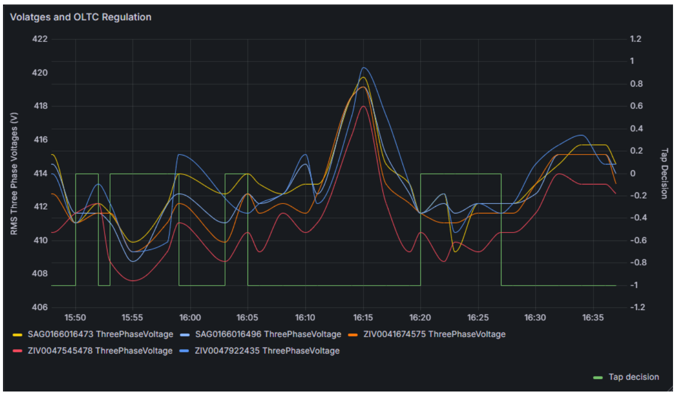

Fig. 6: Evolution of three-phase voltages under tap switching conditions.

C. Scalability Results

The scalability test confirmed that the proposed solution can handle a volume of devices far greater than that of current individual concentrator deployments. The virtual VDC managed to maintain simultaneous communication with 6000 meters distributed across 12 different MV/LV substations, which is equivalent to concentrating the work of 12 physical concentrators into a single edge node. This result suggests that in the future it may be possible to centralize several substations by zones or groupings into a single powerful edge node, reducing the number of devices needed in the field.

One finding was that the practical limit during the test was determined by meter security issues, not by VDC saturation. The resource usage statistics shown in Fig. 7 show that the edge node still had headroom in CPU, memory, and network while handling 5900 connections. Moreover, although data could only be obtained from unsecured meters, the VDC did send requests to all meters. That is, the software scaled up to that number of sessions without crashing or collapsing the network. Under real conditions (with an STG capable of managing the keys), it would be expected that the VDC would successfully serve all meters. In fact, it was verified that some broadcast tasks (such as a general S05 command) were sent but failed only due to the "active key pending" issue in many of them. This demonstrates the robustness of the VDC in scheduling and handling multiple threads of concurrent requests.

D. Economic Impact

The economic analysis carried out aims to compare two deployment scenarios for the modernization of Iberdrola i-DE's transformer centers (CT): the traditional scenario with the deployment of new physical TGUC concentrators versus the innovative scenario with edge nodes running a virtual concentrator (VDC). For this comparison, the initial investment (CAPEX) and operational costs (OPEX) assumptions shown in the corresponding tables, Table II and Table III, have been used.

TABLE II: CAPEX Breakdown by Device Type

Device	Cost	Unit
TGUC	450.00	€/Un
LVS	70.00	€/Un
Edge Node	300.00	€/Un
NB PRIME	100.00	€/Un
License VDC	250.00	€/Instance
License DPP	175.00	€/node
Equipment installation in SS	170.00	€/SS
DPP Deployment	445,000.00	€

TABLE III: OPEX Parameters

Concept	Value	Unit
Intervention in MV/LV SS	140.00	€/SS
O&M interventions in SS	3%	% of SS per year
Reduction in Edge interventions	30%	% of interventions per year

Considering the total number of affected MV/LV substations, specified in Table IV, a gradual and realistic deployment has been assumed, described by a Weibull-type curve (Fig. 8). Based on these assumptions, a detailed cost-benefit analysis has been carried out for each scenario over the project's time horizon (15 years), the results of which can be visually observed in the corresponding graphs (Fig. 9 and Fig. 10).

TABLE IV: Number of Secondary Substations by Area and Configuration

Category	SSs	SSs with $\geq$2 positions	Positions
Urban	55,819.00	22,314.00	78,133.00
Rural	26,185.00	0.00	26,185.00
Scattered Rural	20,050.00	0.00	20,050.00
Total	102,054.00	22,314.00	124,368.00

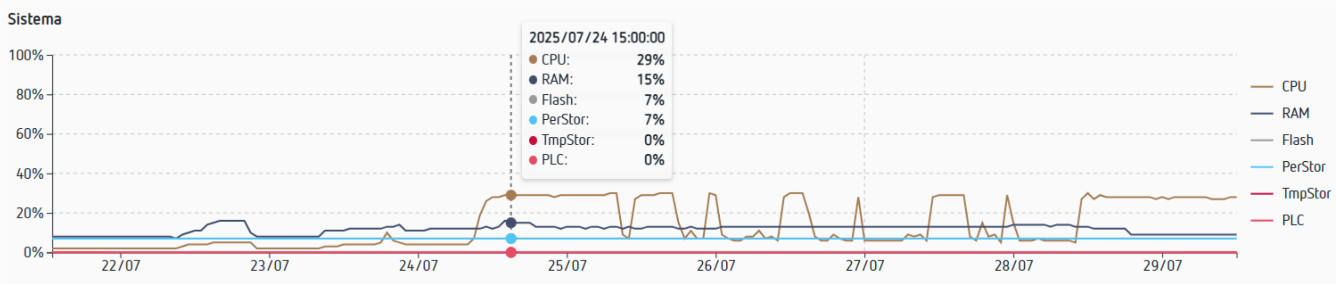

Fig. 7: Resource consumption metrics of the VDC application during scalability tests, showing CPU, RAM, Flash, PerStor, TmpStor, and PLC usage over time.

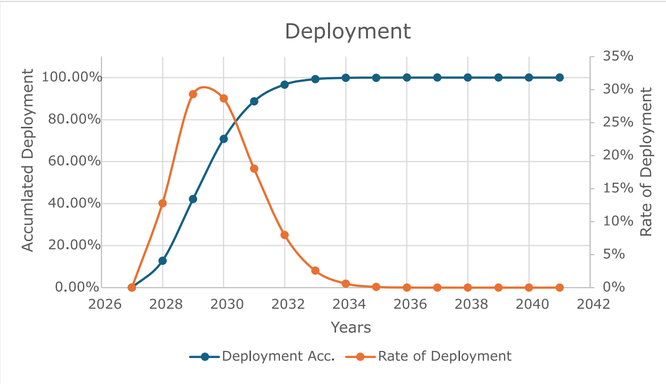

Fig. 8: Estimated deployment of Data Concentrators by i-DE for the next regulatory period.

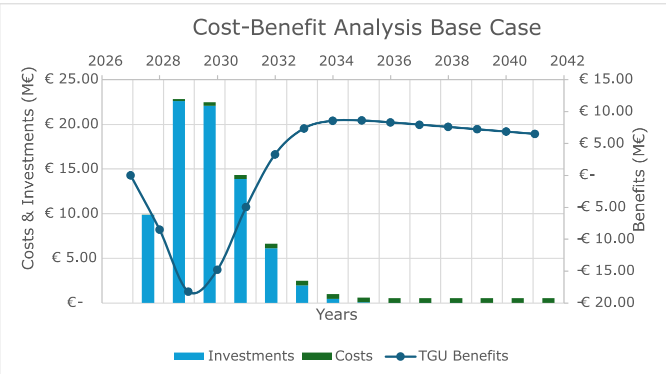

Fig. 9: Cost-Benefit Analysis of the Base Case

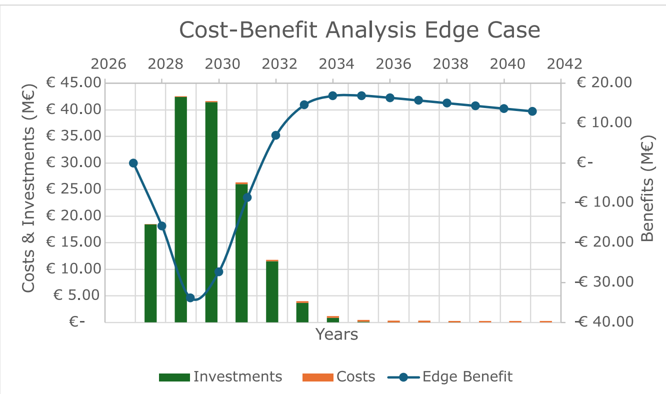

Fig. 10: Cost-Benefit Analysis of the Edge Case

The key economic indicators calculated for each scenario include the Net Present Value (NPV) and the Internal Rate of Return (IRR). These indicators have been evaluated under three different assumptions regarding the regulatory remuneration of software costs: full remuneration (100%), no remuneration (0%), and the minimum remuneration required to make the deployment of edge nodes more advantageous compared to the TGUC base case. The summarized results of these calculations are presented in Table V.

TABLE V: Summary of Economic Impact (Cost & Benefit Analysis)

C&B Analysis	NPV (€)	IRR
Base case	5,234.11	6.462 %
Edge Case (50% Remuneration)	-12,947,412.67	4 %
Edge Case (0% Remuneration)	-31,080,257.63	-1 %
Edge Case (86% Remuneration)	5,234.11	6.462 %

As the main conclusion of the economic analysis, it is worth highlighting that, given the current regulatory framework in Spain, which favors investment in physical assets (CAPEX), it has been determined that the option of edge nodes with VDC is economically more advantageous than the baseline TGUC scenario only if the regulator remunerates the software cost at 86% or higher. This threshold represents the critical point from which the edge computing alternative becomes preferable from a financial profitability perspective.

VI. Conclusions

This work has successfully demonstrated the technical and practical feasibility of virtualizing the Data Concentrator (VDC) in distributed environments, thoroughly validating its operation through comprehensive laboratory and field tests. The comparative analysis between Barbara's edge computing platform, currently adopted by Iberdrola i-DE, and the previously assessed Minsait solution highlighted key architectural differences. Specifically, the evaluation indicated that Minsait aligns more closely with structured standardization processes and stringent cybersecurity frameworks, whereas Barbara offers superior flexibility for third-party integration and supports innovative application development.

The detailed documentation of the deployment and implementation process on industrial-grade hardware provides clear guidelines and insights into integrating containerized edge computing solutions within operational environments, demonstrating reproducibility and feasibility for future deployments.

A significant contribution of this research is the proposal of a modular edge computing architecture aligned with the E4S alliance. The proposed architecture addresses gaps identified in the current E4S reference model, specifically defining critical interfaces for improved interoperability. For Interface 6 (OS Layer – Node Agent), the use of Netlink sockets was identified as the optimal choice due to their efficiency in kernel-user space communications and low overhead. For Interface 7 (Node Agent – Applications), a MQTT-driven design potentially complemented by lightweight container orchestration APIs was recommended, given its modularity, resilience to application restarts, and minimal resource usage. Furthermore, Interface 8 (Applications – External Access or APIs) highlighted the adoption of a flexible, semantic-rich approach leveraging Web of Things (WoT) and Common Information Model (CIM) standards to accommodate diverse, vendor-specific requirements.

The extensive validation activities performed confirmed that the virtualized Data Concentrator effectively replicates essential functionalities traditionally handled by physical concentrators, including configuration management, connectivity assurance, data measurement collection from smart meters, and automated execution of scheduled tasks. These tests provide robust evidence supporting the technical viability of virtualization in critical operational environments.

Additionally, the development and testing of an edge-based use case for On-Load Tap Changer (OLTC) regulation demonstrated the practical utility of deploying intelligent edge applications. Field-deployed equipment successfully utilized real-time data from the VDC to autonomously perform voltage regulation calculations, underscoring the potential for distributed intelligence and autonomous local control in smart substations.

An economic analysis conducted as part of this research assessed the deployment viability of edge nodes within the context of the PRADA project, comparing the virtualized approach against traditional physical concentrators. The findings underscored the sensitivity of economic feasibility to regulatory recognition of software licensing costs as capital expenditure (CAPEX) and highlighted potential economic advantages gained through hosting multiple applications on a single edge node.

Finally, scalability tests performed on field-deployed hardware demonstrated that the VDC solution can manage significantly larger numbers of smart meters compared to typical current deployments. Stress testing involving thousands of devices confirmed stable resource consumption and reliable system operation, thus validating the VDC's readiness for large-scale deployment.

In conclusion, this project has successfully proven that the containerized virtualization of Data Concentrators is technically robust, economically feasible under supportive regulatory frameworks, and scalable enough to meet future grid requirements. The modular and interoperable architecture developed through this work establishes a solid foundation for the ongoing digital transformation and evolution towards intelligent, flexible, and resilient electric distribution networks.

VII. Future work

The outcomes achieved in this thesis open several promising avenues for future development, particularly aimed at enhancing security, interoperability, and scalability of the virtualized edge platform:

- **Secure Northbound Communications:** Implement HTTPS with mutual TLS authentication for communications between the VDC and the remote management system (RMS), pending support from the vendor (Circutor). This will require proper configuration and integration of the VDC's digital certificate within Iberdrola's Public Key Infrastructure (PKI) to ensure trusted, encrypted data exchange.
- **Virtualization of Additional Substation Agents:** Extend the containerization architecture to other substation devices such as RTUs (Remote Terminal Units) and LVS (Low Voltage Supervisors). This would enable the edge platform to host multiple smart grid applications concurrently on a single node, consolidating functionalities like data concentration, control, and monitoring within one hardware unit.
- **Deployment on High-Performance Hardware:** Evaluate the feasibility of running the virtualized platform on over-provisioned substation hardware (for instance, deploying the VDC on a high-capacity RTU device). This assessment should examine multitasking performance and resource sharing when multiple containerized functions operate in parallel on the same host, ensuring that critical applications remain compatible and responsive under combined workloads.
- **Decoupled Node Architecture:** Further separate the base operating system from the edge agent software on the node. This decoupling would facilitate more flexible, vendor-agnostic deployments by allowing the core OS platform to remain independent of any specific data concentrator or edge agent implementation, thereby easing future integration of different vendors' applications on standard hardware.

Appendix A
Alignment with the Sustainable Development Goals

This project aligns with five key Sustainable Development Goals (SDGs), out of the seventeen defined by the United Nations Organization in 2015 [13]. By modernizing the low-voltage grid infrastructure and increasing its efficiency through virtualization and edge computing, the project contributes to the following goals:

- **Affordable and Clean Energy (Goal 7):** The deployment of containerized Virtual Data Concentrators (VDCs) enhances the reliability, efficiency, and flexibility of the LV distribution network. This indirectly optimizes energy delivery, supports the integration of distributed resources, and reduces CAPEX by promoting the use of a single generic hardware platform instead of multiple specialized devices.
- **Industry, Innovation and Infrastructure (Goal 9):** The implementation of edge computing, containerization, and standardized protocols in secondary substations facilitates the transition toward more resilient and scalable utility infrastructures. This contributes to the United Nations' objective of building reliable and sustainable electrical infrastructure.
- **Sustainable Cities and Communities (Goal 11):** The project also contributes to more sustainable urban mobility by enabling more efficient integration of electric vehicles. This is achieved through the introduction of edge computing in secondary substations, which enhances LV network monitoring and control, and enables local smart charging algorithms to optimize EV charging and avoid electrical congestion in urban areas.
- **Responsible Consumption and Production (Goal 12):** The proposed transition from traditional hardware-based

data concentrators to software-based systems aims to minimize the need for single-purpose devices. This helps to avoid hardware obsolescence, promotes the reuse of general-purpose industrial computers, and reduces electronic waste, aligning with SDG 12's principles of sustainable resource usage and circular production models.

- **Climate Action (Goal 13):** The adoption of virtualization and edge computing in secondary substations also facilitates the integration of distributed energy resources (DERs), supports higher penetration of EVs, and ultimately contributes to reducing CO_2 emissions. This aligns with the SDG 13 objective of taking urgent action to combat climate change.

REFERENCES

[1] FUTURED, "Centro de transformacion inteligente: Fundamentos para distribucion de funcionalidades," Tech. Rep., 2021.

[2] B. I. SL, "Ficha de producto: Barbara, la plataforma edge industrial cibersegura," 2022.

[3] I. M. E. J. Segado, "Mide4s libro de trabajo ct inteligente funcionalidades de control, integración y de gestión," Tech. Rep., 2020.

[4] J. Li, G. C., X. Y., and L. F., "Edge-cloud computing systems for smart grid: State-of-the-art, architecture, and applications," *Journal of Modern Power Systems and Clean Energy*, vol. 10, no. 4, pp. 805–817, 2022.

[5] L. Guoqing and T. K. I. K. I. H. N. C. L. P. T. P. P. P., "The convergence of container and traditional virtualization: Strengths and limitations," *SN Computer Science*, vol. 4, 2023.

[6] V. H. Nguyen and H. A. D. R. H. R. T. T. C., "Containerized iec 61850-based protection in smart grid – lab demonstration and performance assessment for large scale deployment," 2024.

[7] N. R. Pérez, "Final master thesis: Analysis of an edge-computing-based solution for local data processing at secondary substations," Universidad Pontificia Comillas, Madrid, Tech. Rep., 2020.

[8] I. N. de Redes, "Distributed processing platform rfi evaluation report," Tech. Rep., 2025.

[9] E4S Alliance, "E4S Alliance — official website," 2025. [Online]. Available: https://www.e4salliance.com/

[10] OpenAPI, "Openapi specification v3.0.3," https://spec.openapis.org/oas/v3.0.3.html, 2020, accessed: 6 June 2025.

[11] W3C, "Web of things (wot) architecture 1.1," https://www.w3.org/TR/wot-architecture11/, 2023, accessed: 6 June 2025.

[12] E. Consortium, "E4s - dso requirements," Tech. Rep., 2025.

[13] U. Nations, "The 17 goals - sustainable development goals," https://sdgs.un.org/es/goals, accessed: 12 June 2025.