

D4.10 COGNIT Serverless Platform - Software Source - e

Version 2.0

1 March 2026

Abstract

COGNIT is an open-source serverless framework for the cognitive cloud-edge continuum that enables the seamless, transparent, and trustworthy integration of data processing resources from providers and on-premises data centers in the cloud-edge continuum. Automatic and intelligent adaptation to where and how data is being processed according to application requirements can be achieved by enabling an AI-Enabled orchestrator. Further, the orchestrator is able to optimise the placement, acts when changing in application demands and behaviour, and the operation of the infrastructure in terms of the main environmental sustainability metrics. This standalone document provides the catalogue of open source software components developed in WP4 “AI-Enabled Distributed Serverless Platform and Workload Orchestration” during the project as part of the implementation of several of the main components of the COGNIT Framework (i.e. Cloud-Edge Manager and AI-Enabled Orchestrator).

Deliverable Metadata

Project Title:	A Cognitive Serverless Framework for the Cloud-Edge Continuum
Project Acronym:	SovereignEdge.Cognit
Call:	HORIZON-CL4-2022-DATA-01-02
Grant Agreement:	101092711
WP number and Title:	WP4. AI-Enabled Distributed Serverless Platform and Workload Orchestration
Nature:	R: Report
Dissemination Level:	PU: Public
Version:	2.0
Contractual Date of Delivery:	30/09/2025
Actual Date of Delivery:	01/03/2026
Lead Author:	Monowar Bhuyan (UMU) & Paul Townend (UMU)
Authors:	Yashwant Singh Patel (UMU), Adil Bin Bhutto (UMU), Rohail Gulbaz (UMU), Malik Bouhou (CETIC), Javad Forough (UMU), Aritz Brosa (Ikerlan), Idoia de la Iglesia (Ikerlan), Sébastien Dupont (CETIC), Fátima Fernández (Ikerlan), Joan Iglesias (ACISA), Tomasz Korniluk (Phoenix), Johan Kristiansson (RISE), Antonio Lalaguna (ACISA), Marco Mancini (OpenNebula), Philippe Massonet (CETIC), Nikolaos Matskanis (CETIC), Aitor Garciandia (Ikerlan), Mikel Irazola (Ikerlan), Álvaro Puente (Ikerlan), Juan José Ruiz (ACISA), Kaja Swat (Phoenix), Thomas Ohlson Timoudas (RISE), Iván Valdés (Ikerlan), Víctor Palma (OpenNebula), Michal Opala (OpenNebula), Jakub Walczak (OpenNebula), Carlos Moral (OpenNebula)
Status:	Submitted

Document History

Version	Issue Date	Status ¹	Content and changes
0.1	01/12/2025	Draft	Initial Draft
0.2	22/12/2025	Peer-Reviewed	Reviewed Draft
1.0	31/12/2025	Submitted	Final Version
1.1	27/02/2026	Draft	Initial Draft
1.2	28/02/2026	Peer-Reviewed	Reviewed Draft
2.0	01/03/2026	Submitted	Final Version

Peer Review History

Version	Peer Review Date	Reviewed By
0.2	29/12/2025	Antonio Álvarez (OpenNebula)
0.2	28/02/2026	Marco Mancini (Ikerlan)

Summary of Changes from Previous Versions

Second Version of Deliverable D4.10

¹ A deliverable can be in one of these stages: Draft, Peer-Reviewed, Submitted, and Approved.

Executive Summary

This is the fifth “COGNIT FaaS Model - Scientific Report” that has been produced in WP4 “AI-Enabled Distributed Serverless Platform and Workload Orchestration”. As this is the final deliverable for this work package, it provides a short description, licence information, version details, code repository and user guide, as well as the design, testing, and verification references of each of the Software Requirements that has been developed throughout the project (M4-M33) under WP4:

Cloud-Edge Manager

- **SR4.1** Provider Catalogue

Implement a backend to persist information about the available providers that can be used to extend the capacity of the COGNIT infrastructure

- **SR4.2** Edge Cluster Provisioning

The Cloud-Edge Manager must be able to provision Edge Clusters as a set of software-defined compute, network, storage on any cloud/edge location

- **SR4.3** Serverless Runtime Deployment

The Cloud-Edge Manager must be able to deploy Serverless Runtimes as Virtualized Workloads within an Edge Cluster.

- **SR4.4** Metrics, Monitoring and Auditing

Edge-Clusters monitoring, Serverless Runtimes metrics collection and continuous security assessment.

- **SR4.5** Authentication and authorization

Authentication and authorization mechanisms for accessing cloud-edge infrastructure resources by the devices for offloading workloads.

- **SR4.6** Plan Executor:

Plan Executor is responsible for the execution of plans produced by the AI-Enabled Orchestrator related to the placement and migration of Serverless Runtimes within an Edge Cluster.

AI-Enabled Orchestrator

- **SR5.1** Building Learning Models:

Provide AI/ML models based on collected metrics from the Cloud-Edge Manager monitoring service related to Edge Clusters, Serverless Runtimes, and infrastructure usage.

- **SR5.2** Smart management of Cloud-Edge resources:

AI-Enabled Orchestrator is responsible for managing and optimizing the lifecycle of Edge Clusters and serverless runtimes within Edge Clusters according to the application requirements, infrastructure and virtual resource usage, and energy-aware policies.

Table of Contents

Abbreviations and Acronyms	6
1. Cloud-Edge Manager	7
2. AI-Enabled Orchestrator	10

Abbreviations and Acronyms

AI	Artificial Intelligence
API	Application Programming Interface
CPU	Central Processing Unit
FaaS	Function as a Service
MOGA	Multi-Objective Genetic Algorithm
ML	Machine Learning
NSGA-II	Non-dominated Sorting Genetic Algorithm II
REST	Representational State Transfer
SPEA2	Strength Pareto Evolutionary Algorithm 2
SR	Serverless Runtimes
VM	Virtual Machine

1. Cloud-Edge Manager

[SR4.1] Provider Catalog

Description The Provider Catalogue provides a unified abstraction layer for managing multi-provider infrastructure enabling seamless integration of resources from public cloud providers, on-premise infrastructure and edge locations. The Provider Catalogue manages the registration and lifecycle of infrastructure providers.

Licence Apache 2.0

Version 7.1.80

Design D4.5 → SR4.1 Provider Catalogue

Code [Public Repository](#)

User Guide [Documentation](#)

Testing D5.6 → 9.4 Cloud-Edge Manager

Verification D5.6 → 9.4 Cloud-Edge Manager

[SR4.2] Edge Cluster Provisioning

Description Edge Cluster Provisioning provides automated provisioning of complete Edge Clusters using declarative infrastructure-as-code approaches deploying compute, storage, and networking resources across heterogeneous providers.

Licence Apache 2.0

Version 7.1.80

Design D4.5 → SR4.2 Edge Cluster Provisioning

Code [Public Repository](#)

User Guide [Documentation](#)

Testing D5.6 → 9.4 Cloud-Edge Manager

Verification D5.6 → 9.4 Cloud-Edge Manager

[SR4.3] Serverless Runtime Deployment

Description Serverless Runtime Deployment addresses the challenge of managing the execution environment for offloaded functions within the COGNIT Framework. A Serverless Runtime is a lightweight, specialized virtual machine instance that executes user functions submitted by Device Clients through the Edge Cluster Frontend. Serverless Runtimes are dynamically provisioned, scaled, and terminated based on application demand, embodying the serverless computing paradigm where infrastructure management is abstracted from application developers.

Licence Apache 2.0

Version 7.1.80

Design D4.5 → SR4.3 Serverless Runtime Deployment

Code [Public Repository](#)

User Guide [Documentation](#)

Testing D5.6 → 9.4 Cloud-Edge Manager

Verification D5.6 → 9.4 Cloud-Edge Manager

[SR4.4] Metrics, Monitoring, Auditing

Description Comprehensive monitoring system providing visibility into infrastructure health, resource utilization, application performance, and power consumption.

Licence Apache 2.0

Version 7.1.80

Design D4.5 → SR4.4 Metrics, Monitoring, Auditing

Code [Public Repository](#)

User Guide [Repository README](#)

Testing D5.6 → 9.4 Cloud-Edge Manager

Verification D5.6 → 9.4 Cloud-Edge Manager

[SR4.5] Authentication & Authorization

Description	Token-based authentication and authorization system using Biscuit tokens, providing secure, decentralized verification and multi-tenant isolation across the distributed cloud-edge continuum.
Licence	Apache 2.0
Version	7.1.80
Design	D4.5 → SR4.5 Authentication & Authorization
Code	Public Repository
User Guide	Repository README
Testing	D5.6 → 9.4 Cloud-Edge Manager
Verification	D5.6 → 9.4 Cloud-Edge Manager

[SR4.6] Plan Executor

Description	Plan Executor component is responsible for the execution of plans produced by the AI-Enabled Orchestrator related to the placement and migration of Serverless Runtimes within an Edge Cluster.
Licence	Apache 2.0
Version	7.1.80
Design	D4.5 → SR4.6 Plan Executor
Code	Public Repository
User Guide	Documentation
Testing	D5.6 → 9.4 Cloud-Edge Manager
Verification	D5.6 → 9.4 Cloud-Edge Manager

2. AI-Enabled Orchestrator

[SR5.1] Building Learning Models

Description SR5.1 provides AI/ML models based on collected metrics from the Cloud-Edge Manager monitoring service related to Edge Clusters, Serverless Runtimes, and infrastructure usage. These models enable predictive analytics for proactive resource management, allowing the COGNIT Framework to anticipate workload changes and optimize resource allocation before performance degradation occurs.

Licence Apache 2.0

Version 3.0

Design D4.5 → SR5.1 Building Learning Models

Code [Public Repository](#)

User Guide [Repository README](#)

Testing D5.6 → 9.5 AI-Enabled Orchestrator

Verification D5.6 → 9.5 AI-Enabled Orchestrator

[SR5.2] Smart management of Cloud-Edge resources

Description SR5.2 manages and optimizes the lifecycle of Edge Clusters and Serverless Runtimes within Edge Clusters according to application requirements, infrastructure and virtual resource usage, and energy-aware policies. This capability enables COGNIT to dynamically adapt resource allocation across the cloud-edge continuum, balancing performance, energy consumption, and cost objectives.

Licence Apache 2.0

Version 7.1.80

Design D4.5 → SR5.2 Smart Management of Cloud-Edge Resources

Code [Public Repository](#)

User Guide [Repository README](#)

Testing D5.6 → 9.5 AI-Enabled Orchestrator

Verification D5.6 → 9.5 AI-Enabled Orchestrator
