

D3.7 COGNIT FaaS Model - Software Source - b

Version 1.0

30 April 2024

Abstract

COGNIT is an AI-enabled Adaptive Serverless Framework for the Cognitive Cloud-Edge Continuum that enables the seamless, transparent, and trustworthy integration of data processing resources from public providers and on-premises data centers in the cloud-edge continuum. The main goal of this project is the automatic and intelligent adaptation of those resources to optimise where and how data is processed according to application requirements, changes in application demands and behaviour, and the operation of the infrastructure in terms of the main environmental sustainability metrics. This document offers a catalogue of those open source software resources developed in WP3 “Distributed FaaS Model for Edge Application Development” during the Second Research & Innovation Cycle (M10-M15) as part of the implementation of several of the main components of the COGNIT Framework (i.e. Device Client, Serverless Runtime, and Provisioning Engine).



Copyright © 2023 SovereignEdge.Cognit. All rights reserved.



This project is funded by the European Union’s Horizon Europe research and innovation programme under Grant Agreement 101092711 – SovereignEdge.Cognit



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

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:	WP3. Distributed FaaS Model for Edge Application Development
Nature:	R: Report
Dissemination Level:	PU: Public
Version:	1.0
Contractual Date of Delivery:	31/03/2024
Actual Date of Delivery:	30/04/2024
Lead Author:	Idoia de la Iglesia (Ikerlan)
Authors:	Monowar Bhuyan (UMU), Malik Bouhou (CETIC), Aritz Brosa (Ikerlan), Sébastien Dupont (CETIC), Aitor Garciandia (Ikerlan), Torsten Hallmann (SUSE), Johan Kristiansson (RISE), Martxel Lasa (Ikerlan), Marco Mancini (OpenNebula), Alberto P. Martí (OpenNebula), Philippe Massonet (CETIC), Nikolaos Matskanis (CETIC), Daniel Olsson (RISE), Goiuri Peralta (Ikerlan), Samuel Pérez (Ikerlan), Thomas Ohlson Timoudas (RISE), Paul Townend (UMU), Iván Valdés (Ikerlan), Constantino Vázquez (OpenNebula), Jorge Lobo (OpenNebula), Michal Opala (OpenNebula).
Status:	Submitted

Document History

Version	Issue Date	Status ¹	Content and changes
0.1	22/04/2024	Draft	Initial Draft
0.2	25/04/2024	Peer-Reviewed	Reviewed Draft
1.0	30/04/2024	Submitted	Final Version

Peer Review History

Version	Peer Review Date	Reviewed By
0.1	25/04/2024	Per-Olov Östberg (UMU)
0.1	24/04/2024	Antonio Álvarez (OpenNebula)

Summary of Changes from Previous Versions

First Version of Deliverable D3.7

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

Executive Summary

This is the second “COGNIT FaaS Model - Software Source” report that has been produced in WP3 “Distributed FaaS Model for Edge Application Development”. It provides a short description, licence, version, code repository and user guide, as well as the design, testing, and verification references of each of the software requirements that have had active development tasks during the Second Research & Innovation Cycle (M10-M15) in connection with these main components of the COGNIT Framework:

Device Client

- **SR1.1** Interface with Provisioning Engine:
Implementation of the communication with the Provisioning Engine.
- **SR1.2** Interface with Serverless Runtime:
Implementation of the communication with the Serverless Runtime.
- **SR1.3** Programming languages:
Support for different programming languages (Python and C).

Serverless Runtime

- **SR2.1** Secure and Trusted FaaS Runtimes:
Automated building of secure and trusted OS images (vulnerability scans, security assessment) related to different flavours of FaaS Runtimes.

Provisioning Engine

- **SR3.1** Provisioning Interface towards the Device:
Provide an interface to the Device asking for a Serverless Runtime to offload functions and data transfer on any resource of the cloud-edge continuum.

This deliverable has been released at the end of the Second Research & Innovation Cycle (M15), and will be updated with incremental releases at the end of each research and innovation cycle in M21, M27, and M33.

Table of Contents

Abbreviations and Acronyms	5
1. Device Client	6
2. Serverless Runtime	8
3. Provisioning Engine	9

Abbreviations and Acronyms

AI	Artificial Intelligence
API	Application Programming Interface
DaaS	Data as a Service
FaaS	Function as a Service
REST	Representational State Transfer
SDK	Software Development Kit

1. Device Client

SR1.1 Interface with Provisioning Engine

Description The Device Client is the component that enables the devices to communicate with the COGNIT Framework in order to perform offloading of tasks. This component communicates with the Provisioning Engine to request/delete/update a Serverless Runtime. It communicates with the provided Serverless Runtime to perform the offloading of functions and the uploading of content to the Data Service, if configured.

The device runtime is delivered as a library with implementations in Python and C which abstracts the user from the internal application protocol by offering a user-friendly API.

The interface with the Provisioning Engine establishes communication with the COGNIT Framework, allowing the device to access its permitted resources.

Licence Apache 2.0

Version [release-cognit-1.0](#)

Design D3.2 → [SR1.1] Interface with Provisioning Engine

Code [Public Repository](#)

User Guide [Repository README](#)

FAQ [Wiki documentation](#)

Testing D5.3 → 10.1 Device Client

Verification D5.3 → 10.1 Device Client

SR1.2 Interface with Serverless Runtime

Description The interface with the Serverless Engine allows the user to interact with the Serverless Runtime that has been assigned to it. Through the defined API, the Device Client is able to manage offloaded tasks at the convenience of the application it is running.

Licence Apache 2.0

Version [release-cognit-1.0](#)

Design	D3.2 → [SR1.2] Interface with Serverless Runtime
Code	Public Repository
User Guide	Repository README
FAQ	Wiki documentation
Testing	D5.3 → 10.1 Device Client
Verification	D5.3 → 10.1 Device Client

SR1.3 Programming languages

Description	In this development cycle an extended version of the Python Device Client has been implemented (representing interpreted languages), and the first version of the C Device Client (representing compiled languages).
Licence	Apache 2.0
Version	release-cognit-1.0 (Python version) release-cognit-1.0 (C version)
Design	D3.2 → [SR1.3] Programming languages
Code	Public Repository (Python version) Public Repository (C version)
User Guide	Repository README (Python version) Repository README (C version)
FAQ	Wiki documentation (Python version)
Testing	D5.3 → 10.1 Device Client
Verification	D5.3 → 10.1 Device Client

2. Serverless Runtime

SR2.1 Secure and Trusted FaaS Runtimes

Description The Serverless Runtime is the service deployed into the scheduled node that will be in charge of executing offloaded tasks. This service exposes the Serverless Runtime API to allow devices to upload functions along with the data needed to execute them.

There will be several flavours of Serverless Runtime to be deployed, depending on the function requirements. It will communicate through the defined RESTful API with the Device Client that is offloading the concerned function.

The Serverless Runtime's VM image will need to contain all the software requirements for the function to be executed.

This requirement focuses on the FaaS (Function as a Service), the actual environment where offloaded functions are executed.

Additionally, the Serverless Runtime exposes function execution information through a Prometheus exporter.

Licence Apache 2.0

Version [release-cognit-1.0](#)

Design D3.2 → [SR2.1] Secure and Trusted FaaS Runtimes

Code [Public repository](#)

User Guide [Repository README](#)

FAQ [Wiki documentation](#)

Testing D5.3 → 10.2 Serverless Runtime

Verification D5.3 → 10.2 Serverless Runtime

3. Provisioning Engine

SR3.1 Provisioning Interface for the Device to manage Serverless Runtimes

Description The Provisioning Engine is a software component that acts as the single point of contact for any device / application that requests access to a Serverless Runtime. It consists of a FaaS Runtime to offload computation through the FaaS paradigm, and/or a DaaS Runtime to offload data into the cloud-edge continuum. The update operation was implemented in this second development cycle

Licence Apache 2.0

Version [release-cognit-1.0](#)

Design D3.2 → [SR3.1] Provisioning Interface for the Device to manage Serverless Runtimes

Code [Public Repository](#)

User Guide [Repository Wiki](#)

Admin Guide [Repository Wiki](#)

Testing D5.3 → 10.3 Provisioning Engine

Verification D5.3 → 10.3 Provisioning Engine