



**AGENT-BASED  
SUPPORT TOOL FOR  
THE DEVELOPMENT  
OF AGRICULTURE POLICIES**

## **D7.5 Report on AGRICORE testing platform evaluation**



Deliverable Number	D7.5
Lead Beneficiary	IDE
Authors	IDENER
Work package	WP7
Delivery Date	M50
Dissemination Level	Public

[www.agricore-project.eu](http://www.agricore-project.eu)



The Agricore project has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement No. 816078





## Document Information

Project title	Agent-based support tool for the development of agriculture policies
Project acronym	AGRICORE
Project call	H2020-RUR-04-2018-2019
Grant number	816078
Project duration	1.09.2019-31.8.2023 (48 months)

## Version History

Version	Description	Organisation	Date
0.1	Deliverable ToC	IDE	12-Dic-2023
0.2	Initial Content	IDE	7-feb-2024
0.3	Extended descriptions	IDE	25-feb-2024
0.4	Update to reflect software updates	IDE	15-may-2024
1.0	Final version	IDE	12-sep-2024

## Disclaimer

All the contributors to this deliverable declare that they:

- Are aware that plagiarism and/or literal utilisation (copy) of materials and texts from other Projects, works and deliverables must be avoided and may be subject to disciplinary actions against the related partners and/or the Project consortium by the EU.
- Confirm that all their individual contributions to this deliverable are genuine and their own work or the work of their teams working in the Project, except where is explicitly indicated otherwise.
- Have followed the required conventions in referencing the thoughts, ideas and texts made outside the Project.

## **Executive Summary**

D7.5 reports on the work done within Task 7.4 - AGRICORE testing platform evaluation. This task represented the starting point of the demonstration actions of the project, as it commissioned the AGRICORE platform as a whole. However, some elements of the AGRICORE platform were already deployed earlier (e.g. ARDIT), as they did not influence the use cases execution. This document delves into the different stages of the platform evaluation, from the preparation of the testing environment up to the analysis of the results achieved.

## Abbreviations

Abbreviation	Full name
ABM	Agent Based Model
AGRICORE	Agent-based support tool for the development of agriculture policies
EU	European Union
DWH	Data WareHouse
I/O	Input/Output

## List of Figures

Figure 1: AGRICORE IT architecture .....	8
Figure 2: Software modules composing the final Agricore Suite.....	8
Figure 3 : AGRICORE Software deployment for testing environment.....	9

## List of Tables

Table 1: AGRICORE HW Testing infrastructure.....	9
--	---

## Table of Contents

<b>1</b>	<b>Introduction .....</b>	<b>7</b>
<b>2</b>	<b>Preparation of the testing environment .....</b>	<b>8</b>
<b>3</b>	<b>Testing of the prepared environment and analysis of the results .....</b>	<b>11</b>
<b>4</b>	<b>Conclusions .....</b>	<b>12</b>

# 1 Introduction

This deliverable is part of the results of WP7 of the AGRICORE project. WP7's main objective is testing and validating the technologies developed within AGRICORE in real applications. As such, WP7 covered most of the aspects related to the preparation and analysis of the use cases, as well as the setup and evaluation of the required testing infrastructure. This WP is strongly related to all WPs in the project, as it includes all developments (WP1-WP5) and the integrated AGRICORE Suite (WP6).

Specifically, this deliverable is the result of the execution of Task T7.4 AGRICORE testing platform evaluation, which objective is to commission a testing infrastructure for the AGRICORE platform as a whole, testing its suitability for the execution of the use cases.

The structure of this deliverable starts with a prepared testing environment, including the cloud services used for this and the required networking and permission setups. It later delves into the actual testing of the testing environment and in the analysis of the testing results, as well as a set of corrective actions that were required to properly prepare the system for the execution of the use cases.

## 2 Preparation of the testing environment

The testing environment should allocate all the elements of the AGRICORE suite, which architecture is depicted in next **Error! Reference source not found.** However, this logical blocks, established within the proposal stage of the project, evolved in a set of specific software services in the final integrated version of the AGRICORE suite, which are depicted in Figure 2.

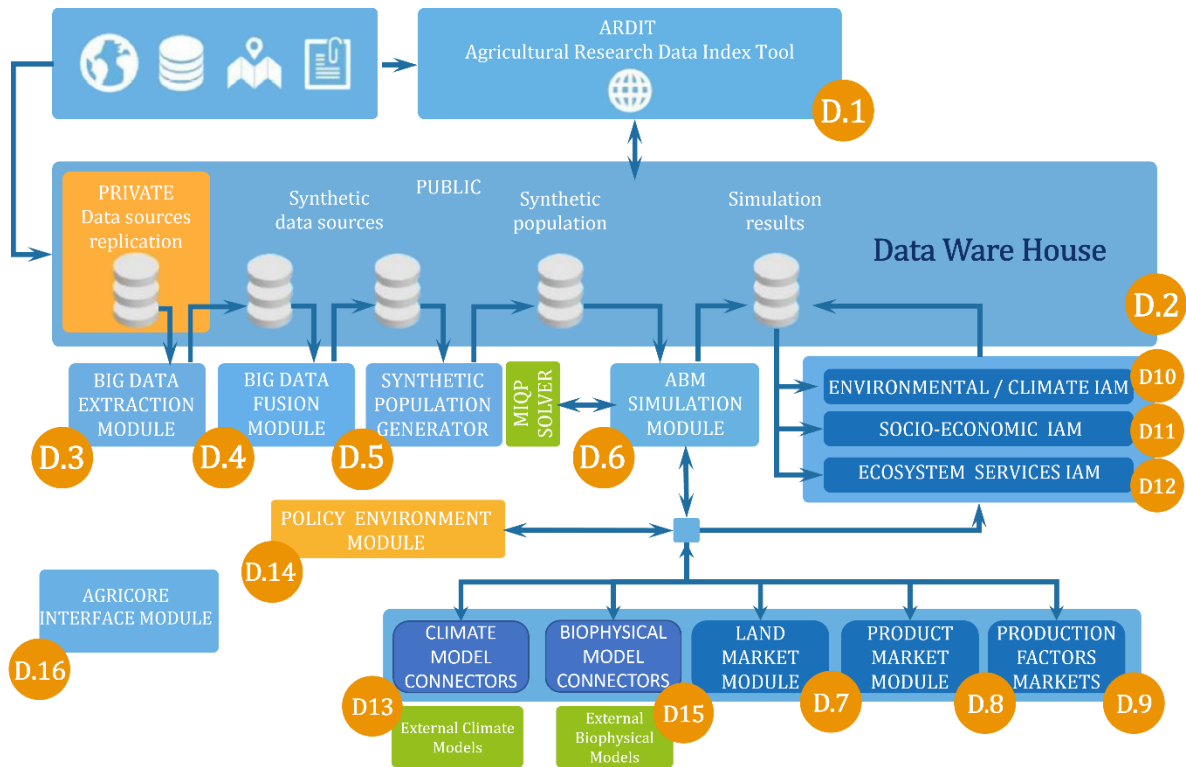


Figure 1: AGRICORE IT architecture

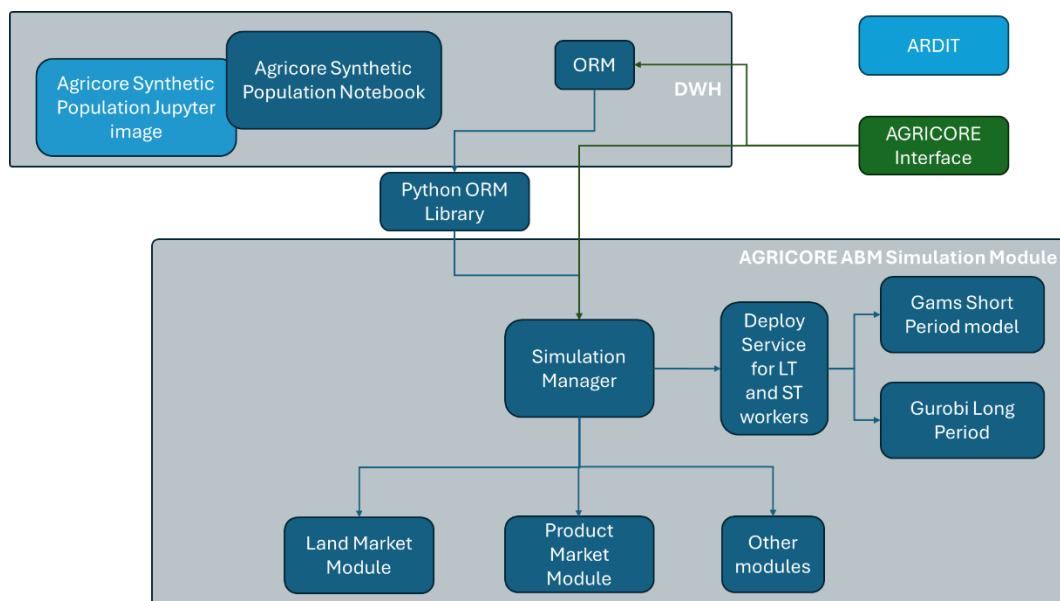


Figure 2: Software modules composing the final Agricore Suite

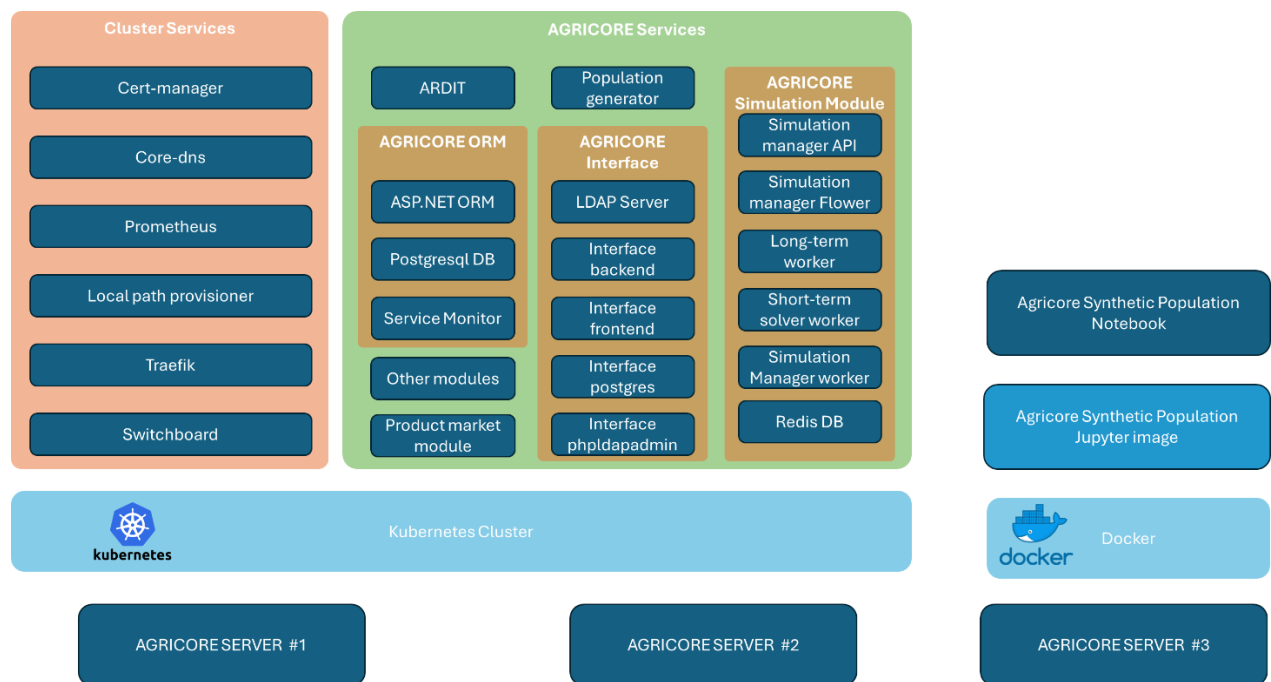


To allocate all these services, three servers were rented to one of the main computing infrastructure providers in Germany, HETZNER. The characteristic of these 3 servers is described in next Table 1.

**Table 1: AGRICORE HW Testing infrastructure**

Agricore Server #1	Agricore Server #2	Agricore Server #3
AMD Ryzen 9 3900	AMD EPYC 7401P	AMD EPYC 7401P
12 x 3.1 Ghz; 24 threads	24 x 2 Ghz; 48 threads	24 x 2 Ghz; 48 threads
128GB DDR 4 ECC RAM	128GB DDR 4 ECC RAM	128GB DDR 4 ECC RAM
2x 1.92 TB Nvme	2x 960 GB Nvme	2x 960 GB Nvme
1GB NIC Network	1GB NIC Network	1GB NIC Network

For the deployment of the software infrastructure of the AGRICORE suite, a Kubernetes cluster was setup on top of Server #1 and Server #2 as shown in Figure 3, while the third server was allocated only for the generation of the synthetic populations which, although lighter than the actual simulation process, also requires a relatively high level of computing capabilities.



**Figure 3 : AGRICORE Software deployment for testing environment**

The K8 cluster was setup using a standard version of [K3s](#), a lightweight flavoured Kubernetes distribution that automatically setup some of the core elements to quickly having a running cluster.

Due to the number of servers used in the deployment, High Availability was not targeted, as the required number of nodes for such setup is 3 (and only 2 were included in the cluster). The default K3s distribution was tweaked to provide proper local storage capabilities, as all the loads that the consortium deployed in the cluster had specific affinity for one of the servers, which suited the purpose of a testing environment but that is not recommended for a production one.

On top of K3s, different services were deployed to provide core functionality, including cert-manager (to manage required SSL certificates in an automated way), Prometheus (to extract load metrics from the server to monitor their usage), traefik (used a single-entry point for all the

services deployed in the cluster) and switchboard (which automatizes the interaction between traefik, cert-manager and the dns manager).

On the AGRICORE side, all the required modules were deployed. ARDIT and a running engine for population generation (which did not have much usage as this item was also present in the server outside the cluster) were deployed. Some minor extra services were also added. Then, the 3 main deployments were installed in the server: The AGRICORE interface, the AGRICORE simulation module and the AGRICORE ORM).

Finally, [Lens](#) was used at the administrators computers to monitor and interact with the cluster deployment.

On the other side, Server #3 used an Ubuntu LTS Server distribution on top of which Docker was installed. The AGRICORE Synthetic Population Jupyter Image was used to deploy a container in which exposes a Jupyter server. This was used to further develop and deploy the diverse synthetic population generation notebooks that were used during the generation of the population for the use cases.

### 3 Testing of the prepared environment and analysis of the results

Once the infrastructure was ready, the AGRICORE development team started making the required test to verify that the system was ready for the execution of the use case.

For such testing, the team decided to test the complete process against a simple, smaller population. For this purpose, a non-synthetic population was created by extracting data (1000 farms) from the RICA dataset that was available for the execution of the Italian Use case.

At the early stages of the testing, several issues were identified and accordingly a wide set of request to the development team was done. This covered, initially, issues with the communication between services, unclarity on the units of measures of specific variables, hangs of the simulation process and inaccuracy of some of the results (due to long-term and short-term model misalignment). All these issues were properly addressed and solved in the last versions of the affected modules. Moreover, the development team realised that the system was somehow limited on the level of information it offered for debugging purposes, which trigger a considerable amount of effort towards proper reporting not only at user level, but also at developer one.

On a second stage, once all runtime errors were corrected, the simulation of the testing scenario was completed successfully. However, new issues on the generated data were identified, which clearly points to specific bugs in the code of several modules (from software services to the mathematical models) which were also addressed consequently.

Finally, after the second round of tests, the system was considered ready for the execution of the use cases.

## 4 Conclusions

The work reported in this deliverable covers the setup of the testing environment which will be used for the execution of the use cases. When this took place, the team properly demonstrated that the system was indeed ready for such campaign.

However, as this deliverable was delivered later than expected, it is relevant to report further discoveries that happened later in the project. Specifically, and despite the fact that the testing campaign demonstrated the operability of the platform, the tests were not properly designed to account for all the use cases casuistic. Specifically, the decision to use a small data subset for the testing campaign avoid to early identify several issues related to modules performance and hardware resources exhausting. Once the simulation of the first use cases began, these issues were identified and required many more efforts on corrective actions. The Andaluca’s case encompass the simulation of >100.000 farms across 4 years. Given the simulation detail and granularity of the AGRICORE approach, this require a considerable amount of computing resources to be done. Being AGRICORE a pioneer on this kind of high-level detail simulations, the team realised that a new level of optimisation both in the software implementation and the algorithms involved was required. As a result of this effort, the AGRICORE suite evolved to the version described in D6.5 and shared in the open-source repositories. This version properly exploits all available parallelisation capabilities and is able to properly exhaust the available resources in the cluster where it is deployed.

Finally, the AGRICORE team acknowledges that several new improvements can be done at architectural level to increase the overall reliability of the overall platform, supporting more concurrent simulations and avoiding any potential issues due to excessive resource usage. These optimisations should be carried out in future activities that exploit AGRICORE results.

For preparing this report, the following deliverables have been taken into consideration:

<b>Deliverable Number</b>	<b>Deliverable Title</b>	<b>Lead beneficiary</b>	<b>Type</b>	<b>Dissemination Level</b>	<b>Due date</b>
D6.4	ABM Simulation Module	IDE	0	Public	M52
D6.5	AGRICORE Suite	IDE	0	Public	M54