Data Platform for the Connection of Cognitive Ports

Title:	Document Version:
D5.5 Use Cases and Applications Evaluation Report	1.0

Project Number:	Project Acronym:	Project Title:
H2020-871493	DataPorts	A Data Platform for the Cognitive Ports of the Future

Contractual Delivery Date:	Actual Delivery Date:	Deliverable Type*-Security*:
M39 (March 2023)	M39 (March 2023)	R-PU

*Type: P: Prototype; R: Report; D: Demonstrator; O: Other; ORDP: Open Research Data Pilot; E: Ethics.

**Security Class: PU: Public; PP: Restricted to other programme participants (including the Commission); RE: Restricted to a group defined by the consortium (including the Commission); CO: Confidential, only for members of the consortium (including the Commission).

Responsible:	Organisation:	Contributing WP:
Paolo Calciati	ITI	WP5

Authors (organisation):	
Paolo Calciati (ITI)	Santiago Cáceres (ITI)
Andreu Belsa (UPV)	Miguel Bravo (ITI)
Héctor Iturria Sánchez (PRO)	Alvaro Martinez Romero (PRO)
Salvador Cuñat (UPV)	Pablo Giménez (VPF)
Tasos Nikolakopoulos (ICCS)	Alexandros Zerzelidis (CERTH)
Eirini Tserga (ThPA)	Vasilios Siopidis (CERTH)
Alexandros Tremopoulos (CERTH)	Stavros Devrelis (CERTH)

Abstract:

This document reports about the technical and business assessment of the different pilots, including aspects such as software performance, usage metrics, usability, revenue generated or number of users reached.

The first scope focuses on requirements evaluation, ensuring that the project respects and takes in consideration the requirements that have been defined.

The second scope deals with the evaluation of the use cases: the first part aims to ensure that the defined scenarios are correctly implemented, while the second part is based on the Technology Acceptance Model and aims to evaluate the perceived usefulness and ease of use of the project.

Finally, the document reports evaluation metrics related to software performance and usage metrics.

Keywords:

Evaluation, plan, pilot, assessment, survey

Revision History

Revision	Date	Description	Author (Organisation)
V0.1	01.11.2022	Initial draft	Paolo Calciati (ITI)
V0.2	06.02.2023	Added description of requirements evaluation	Paolo Calciati (ITI)
V0.3	13.02.2023	Added description of use case evaluation, integrating content received from different use cases leaders	Paolo Calciati (ITI)
V0.4	17.03.2023	Added Evaluation Metrics	Andreu Belsa (UPV)
V0.5	23.03.2023	Feedback from first evaluator	Kristina Weimer (FHG)
V0.6	26.03.2023	Feedback from second internal evaluator	Fabiana Fournier (IBM)
V0.7	30.03.2023	Consolidation of feedbacks	Paolo Calciati (ITI)
V1.0	31.03.2023	Final version	Paolo Calciati (ITI)

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement № 871493.

More information available at https://DataPorts-project.eu

Copyright Statement

INDEX

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

1.1 DATAPORTS PROJECT OVERVIEW

DataPorts is a project funded by the European Commission as part of the H2020 Big Data Value PPP programme, and coordinated by the Technological Institute of Informatics (ITI). DataPorts rely on the participation of 13 partners from five different nationalities. The project involves the design and implementation of a data platform, its deployment in two relevant European seaports connecting to their existing digital infrastructures and addressing specific local constraints. Furthermore, a global use case involving these two ports and other actors and targeting inter-port objectives, and all the actions to foster the adoption of the platform at European level.

Hundreds of different European seaports collaborate with each other, exchanging different digital data from several data sources. However, to achieve efficient collaboration and benefit from AI-based technology, a new integrating environment is needed. To this end, DataPorts project is designing and implementing an Industrial Data Platform.

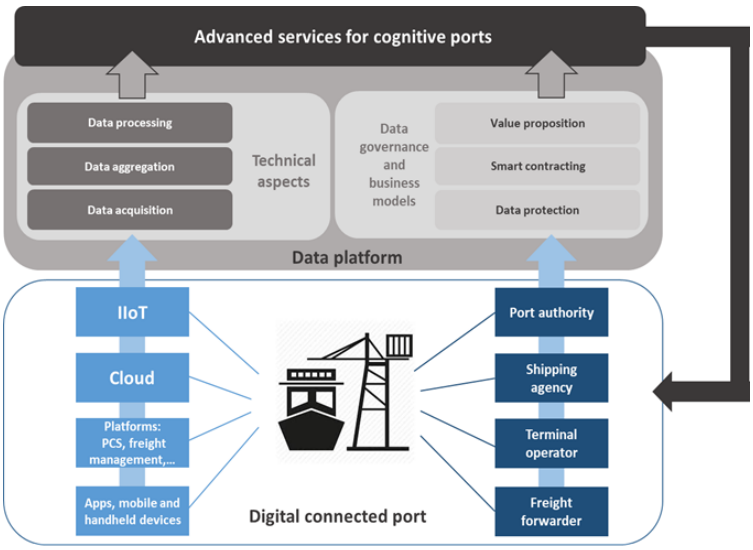

The DataPorts Platform aim is to connect to the different digital infrastructures currently existing in digital seaports, enabling the interconnection of a wide variety of systems into a tightly integrated ecosystem. In addition, to set the policies for a trusted and reliable data sharing and trading based on data owners' rules and offering a clear value proposition. Finally, to leverage on the data collected to provide advanced Data Analytic services based on which the different actors in the port value chain could develop novel AI and cognitive applications.

DataPorts will allow establish a future Data Space unique for all maritime ports of Europe and contribute to the EC global objective of creating a Common European Data Space.

1.2 DELIVERABLE PURPOSE AND SCOPE

Specifically, the DoA states the following regarding this deliverable:

This document will report about the technical and business assessment of the different pilots, including aspects such as software performance, usage metrics, usability, revenue generated and number of users reached.

This document reports the results of the evaluation, whose process is detailed in Deliverable D5.2 – Evaluation Plan, which comprises two different scopes: requirements and use cases. The description from D5.2 for the evaluation process is reported in the following three paragraphs:

- The first scope includes the evaluation of all functional and non-functional requirements defined in the project. The requirements must be verified one by one, assigning acceptance criteria for both mandatory and optional requirements in the form of a percentage of fulfilled requirements.
- The latter scope includes the evaluation of test cases and a user acceptance survey for each of the four use cases of DataPorts, as presented in Deliverable D5.4: Valencia Port, Thessaloniki Port, Smart Containers, and Port Integration. Each use case defines multiple scenarios, and each scenario specifies a series of test cases to be executed to verify its successful implementation. Like for

requirements, each scenario will be evaluated stand-alone, and the percentage of successful test cases will be checked against a pre-defined threshold to determine whether the evaluation is successful. The perceived usefulness and ease of use of the use cases are evaluated through a set of surveys which will be handed out to the final users after testing the DataPorts platform. The surveys will be based on the Technology Acceptance Model [1] and contain both closed (Likert scale) and open questions.

Conceptually, both scopes provide different levels of evaluating the project: on one hand, evaluating requirements allows to objectively check that the software was developed in line with the project specifications; on the other hand, evaluating the use cases ensures that the software can deliver the functionalities as promised. Finally, the user acceptance survey gives an idea of the perceived usefulness and ease of use of DataPorts.

D5.5 presents the results of the evaluation process detailing insights in each of the scopes previously described, discussing the conclusions regarding the achievements and benefits, but also addressing issues and new challenges to be addressed by the DataPorts platform.

Finally, besides the evaluation process previously described, D5.5 also explores different evaluation metrics, such as software performance and usage metrics, providing details regarding how the platform is deployed and how does it work by analysing the specifications of the machines on which the platform is deployed and how they perform under different load conditions.

1.3 DELIVERABLE CONTEXT

This document's relationship to other documents is as follows:

Primary Preceding documents:

- Consortium Agreement (CA): Deals with legal aspects between partners.
- Description of Action (DoA): Provides the foundation for the actual research and technological content of DataPorts. Importantly, the Description of Action includes a description of the overall project work plan.
- D5.2 – Evaluation plan: Contains the description of requirements and use case evaluation techniques.
- D5.4 – Use case oriented pilots final version: contains the description of pilots

Primary Dependant documents:

There is no dependant document, as this deliverable contains the final evaluation report for the project.

1.4 DOCUMENT STRUCTURE

This document is structured as follows:

- Section 1 contains a high-level overview of the project, the definition of the purpose and scope of this document, its context and structure.
- Section 2 presents the results of the evaluation of the requirements of the project.
- Section 3 presents the results of the use cases evaluation of DataPorts. This Section is divided into two different parts: the first (Section 3.1) describes the evaluation of the test cases; the second (Section 3.2) describes the evaluation of the project using the User Acceptance Survey.
- Section 4 discusses the remaining evaluation metrics described in the DoA: software performance, usage metrics, revenue generated, and number of users reached.
- Finally, Section 5 contains the conclusions of the project evaluation, as well as the next steps to address issues and new challenges that could arise in the future.

2 REQUIREMENTS EVALUATION

This section describes the methodology and results of the evaluation of the requirements of the DataPorts project, carried out as specified in the Deliverable D5.2.

2.1 METHODOLOGY

This section presents a summary of the evaluation methodology for the project requirements, which is fully described in Deliverable D5.2.

Each project requirement was evaluated on its own. The partner responsible for the requirement assigned it a value of PASS, MINOR REWORK, or FAIL, which was then confirmed by the corresponding reviewer.

Since DataPorts had both optional and mandatory requirements, the different subsets had a different threshold for considering the evaluation successful:

- 90% for MUST and WON'T requirements
- 75% for SHOULD requirements
- 33% for COULD requirements

The evaluation of the project was considered successful if the following conditions apply:

- DataPorts fulfilled the *PASS* threshold for both mandatory and optional requirements.
- Considering requirements with both *PASS* and *MINOR REWORK* status, *MINOR REWORK* should have been at most 25% of the total.

2.2 RESULTS

The results of the requirements evaluation are analysed by considering requirements of each work package separately. Table 1 contains the results of the requirements evaluation for all work packages (WPs), dividing them by the different typology of requirements (see Section 2.1), and reporting the status of the evaluation for each of them.

WP	Req type	Req Count	Pass Count	Pass %	Status
2	MUST	6	6	100	PASS
2	SHOULD	1	1	100	PASS
3	MUST	31	30	97	PASS
3	SHOULD	12	12	100	PASS
3	COULD	4	3	75	PASS
4	MUST	13	13	100	PASS
4	SHOULD	3	3	100	PASS
4	COULD	1	1	100	PASS
5	MUST	17	16	94	PASS
5	SHOULD	3	3	100	PASS
5	COULD	3	3	100	PASS

Table 1 - Requirements Evaluation Results

2.2.1 Work Package 2

WP2 contained a total of 7 requirements, all of which are in PASS status, so the Work Package evaluation is considered to be successful.

2.2.2 Work Package 3

WP3 has two requirements that did not pass the evaluation: one of them in the COULD category, and the other one in the MUST category. Due to the high number of requirements in this Work Package (47 in total) the evaluation is still successful.

The first failed requirement is R3.3, which is of priority COULD: "*Each data source that is going to be connected to the DataPorts infrastructure through an API could provide version management for its API, in order to allow a proper identification of the API and the data formats.*". The control version is not implemented for the API, so the requirement could not be considered completed.

The second failed requirement is R3.8, with priority MUST: "*The Data Access component must provide the proper connectors to communicate with the available data sources*". Not all the implemented agents have a template available, although there are templates for different types of basic agents which can be adapted for specific uses. Nonetheless, this condition is not sufficient for considering the requirement as fulfilled.

2.2.3 Work Package 4

WP4 contains a total of 17 requirements, all of which are in PASS status: the Work Package evaluation is successful.

2.2.4 Work Package 5

WP5 contains a total of 23 requirements. Out of those, one MUST requirement (out of 17) failed, which makes the evaluation for this Work Package still successful: "*The DataPorts Platform must provide an agent to provide Smart Containers data through DataPorts*". The reason behind the failure of this requirement is that there is no specific development to expose Smart Containers data from DataPorts.

3 USE CASES EVALUATION

This section describes the methodology and results of the evaluation of the use cases of the DataPorts project, carried out as specified in the Deliverable D5.2.

3.1 TEST CASES EVALUATION

3.1.1 Methodology

As described in Deliverable D5.2, the test cases of the different DataPorts scenarios are separately defined by each use case and collected in the Annex A (Section 7).

Each test case has a PASS or FAIL value, and the evaluation of a use case is considered to have a positive output if at least 90% of all test cases of each scenario have a PASS value. This ensures that the described scenarios are successfully implemented.

The test cases are executed by people that are familiar with the use case scenario, possibly working in the ports, or at least familiar with the port environment. The target number of evaluators for each of the use cases is five. Besides the result value of each test case, the user executing the test cases can provide a comment to clarify why the test case fails, or to address any doubt or clarification. The comment is mandatory if a test case fails, and optional otherwise.

3.1.2 Results

This section describes the results of the test cases evaluation. The results are divided into four different subsections - one for each use case.

The Excel files with the use cases evaluation can be found in the Annex A of this Document (see Section 7).

3.1.2.1 Port of Valencia

Table 2 presents the results for the test cases evaluation of the Port of Valencia. All test cases have a PASS status, resulting in a successful outcome for the test cases evaluation of the use case.

Scenario	Result	Comments
Tracking of Transport Operations	PASS	-
Verified Gross Margin	PASS	-
Digital Consignment Note	PASS	-
Port Authority Data Sharing & Analytics Scenario	PASS	-

Table 2 – Port of Valencia Test Cases Results

While for the first three scenarios of the Port of Valencia there were no comments from the evaluators, the last scenario, Port Authority Data Sharing & Analytics Scenario reported two issues from the evaluators, which did not prevent the test case from being successful, but should be addressed in the upcoming updates to the application:

- One user reported that the providers displayed when picking a dataset in the Training Wizard are not the ones mentioned in the description.
- While a second user reported that the option to select the granularity in the configuration of the Training Wizard is missing.

3.1.2.2 Port of Thessaloniki

Table 3 presents the test cases results for the Port of Thessaloniki use case. The Port of Thessaloniki defined eight test case scenarios, all of which were successfully executed by the evaluators, resulting in a successful

outcome for the test cases evaluation.

Scenario	Result	Comments
New COREOR Request	PASS	-
View Booking	PASS	-
Statistics (Vessel Visualisation)	PASS	-
Statistics (Vehicle Traffic)	PASS	-
Statistics (Vehicle Time At Port)	PASS	-
Queues Prediction (Vessel)	PASS	-
Queues Prediction (Custom)	PASS	-
Passenger Covid Analytics	PASS	1 test case out of 15 failed

Table 3 – Port of Thessaloniki Test Cases Results

The evaluators of the Port of Thessaloniki test cases produced the following comments with relation to the tests, which should be addressed when planning any update to the platform:

- In the “Statistics (Vehicle Traffic)” one user was in doubt on how to clear the node after selecting it. He was able to do it after being told how to do it, but this should be mentioned in the user manual and taken into consideration to improve the usability of the application.
- Another user complained that the web application was very slow to visualize in the “Statistics (Vehicle Time at Port)”.
- The “Queues Prediction (Custom)” test case produced a high number of comments, which is a possible indicator of some issues with the usability of the application. Users need to know the exact specifics of the vessel in order to use the application. More experienced users which tested this application did not have any issue with this, however less experienced users found it difficult to correctly populate all necessary fields.
    - A user complained that he should be only filling in name and date, not type of container (40'' or 20'') or weight.
    - One user complained that the list of predictions is not in alphabetical order.
    - Another user was unsure whether selecting the “40” option, the time would be predicted for 40 minutes only, showing that it is not clear what the “40” option is for.
    - Finally, two users mentioned that after clicking the “Clear” button to clear the form, the predictions table was still filled. In this case the behaviour is correct because the button refers to the form only.
- The “Passenger Covid Analytics” test case also produced an elevated number of comments from the evaluators:
    - Some evaluators complained about the slow visualization of the page at different points in the scenario.
    - All users reported that the clickable options for the “mobility” dropdown menu mentioned in the test description do not appear in the drop-down menu, which caused the test case to fail. The drop-down menu provides different options than the ones mentioned in the test case description. This was due to the fact that when creating the test cases the front end was not completely developed, and this caused a misalignment between test case and reality.

- Two different evaluators reported that when selecting a path from the “travel times” dropdown option, some names appear in Greek. The drop-down menu presents a list of street names, some of which are in Greek.

All those user comments were taken into consideration when drafting the user manual, which is included as Annex in Deliverable D5.4.

3.1.2.3 Smart Containers

The results for the Test Case evaluation of Smart Containers are shown in Table 4.

All Test Case scenarios for the Smart Container did pass the evaluation: a couple of test cases did fail, however the amount of fail does not exceed the 10% threshold, so the evaluation is considered to be successful.

Scenario	Result	Comments
Geofencing	PASS	1 test case out of 10 failed
Door Opening/Closing	PASS	1 test case out of 10 failed
Geolocation	PASS	-

Table 4 – Smart Containers Test Cases Results

In the geofencing test case all testers reported the step of identifying the Geofencing IN event as FAIL: the geofencing event was not present in the platform.

In the Door Opening/Closing test case, all testers reported the last step, which involves identifying the relevant door closing event as FAIL, specifying that no door closing event are present, despite being able to verify the door opening event in the previous step.

Both test cases that failed, despite not causing the general test case to fail the evaluation, should be addressed as soon as possible, as they could hinder the usability of the application.

Finally, the integration between the front end of the Port of Valencia and Traxens’ platform only registers the latest event generated in each of the three scenarios (geolocation, door open/close notification, and geofencing). This is enough to pass the defined test cases, however in the next release of the platform it should be considered to integrate the full history of events received by Traxens’ platform into the DataPorts’ front end.

3.1.2.4 Port Management System Integration

Table 5 reports the results for test case evaluation of the Port Management System Integration use case. All test cases passed the evaluation, which is considered successful.

Scenario	Result	Comments
Create destination for alerts	PASS	-
Subscribe for Change of ETA	PASS	-
Perform an ETA change	PASS	-
Check alerts	PASS	-
Subscribe for Water requests	PASS	-
Perform a Water Supply request	PASS	-
Disable subscription	PASS	-

Scenario	Result	Comments
Perform an ETA change (II)	PASS	-

Table 5 – Port Management System Integration Test Cases Results

The eight scenarios of the Port Management System Integration contain various comments from the evaluators, each issued only by a single evaluator:

- When logging into Posidonia Notifications' main page, one user reported that if a previous access was in cache and the access token is expired, the web shows an error and does not redirect to login. This issue happened because the provided URL required login to access but the redirection to the login webpage was not working. Writing the main URL of the application solves the issue.
- One user reported that the table visualized after creating a subscription for both test cases "Subscribe for Change of ETA" and "Subscribe for Water requests" does not have pagination controls and there are already more elements than can be shown in one page. To see the new subscription, the user has to sort by descending time of creation.
- The apply button for disabling a subscription in the "Disable Subscription" test case is rendered almost out of the window. There is a problem with the position of the button when the window is maximized on some given screen resolutions. The button is partially out of the window but can be clicked. It is a visualization error, but the action can be done.

All of these comments highlight minor bugs or issues that should be addressed in the future, none of the issues found in the tests prevented from completing the tasks.

3.2 USER ACCEPTANCE SURVEY

The user acceptance survey evaluates the usability of the software and how it will help users to perform their daily tasks. Each of the use cases of DataPorts will be evaluated with a survey based on the Technology Acceptance Model (TAM) [1] composed of both Likert closed questions (questions that use a four point scale) and open questions.

3.2.1 Methodology

As described in Deliverable D5.2, the survey is filled by all the users which participate in testing the use cases, as described in Section 3.1.

The questionnaire, described in detail in Section 3.2.2, contains a series of 4-points Likert questions with answers in the range "Strongly Disagree" to "Strongly Agree", with the addition of the "Not Applicable" value, shown in the questionnaire as blank. A value going from 1 to 4 (1 for Strongly Disagree, 2 for Disagree, 3 for Agree, 4 for Strongly Agree) is assigned to all the Likert questions present in the questionnaire. Not applicable values are not considered for the evaluation. The questionnaire of a use case is considered to have a positive outcome if the following conditions are met:

- The average of answered questions is 3 or above.
- At least 80% of the questions have an average of 3 or above.

Moreover, all questions with an average value of 2 or below for a given use case should be analysed more in detail, as they indicate a potential problem in the software.

TAM defines the following two main indicators: **Perceived Usefulness** and **Perceived Ease of Use**, to which it adds the dependent variable **Behavioural Intention**. Some TAM implementations also assess the intention to use indicator, however for this specific project it is not necessary, as DataPorts would either introduce new functionalities, or replace the systems in use, and not be an alternative to them.

Perceived usefulness is defined as "the extent to which a particular person believes that using a particular system would enhance his or her job performance". Perceived ease of use is defined as "the degree to which

a person believes that using a particular system would be free of effort". Behavioural intention defines individual intention to engage in specific behaviour, in this specific case it would be to use a particular technology: the DataPorts system.

The analysis of the questionnaire's answers, presented in Section 3.2.3, will first consider the three indicators (Perceived Usefulness, Perceived Ease of Use and Behavioural Intention) separately, and then focus on the results as a whole, by averaging the results for each of them in a final score for each use case.

The open questions planned in the User Acceptance Survey will be discussed in Section 3.2.3.

3.2.2 Questionnaire

The following is a general proposal for a questionnaire to evaluate the DataPorts use cases. Each use case responsible reviewed and modified the questions to best capture the essence of its use case: the four questionnaires, together with the answers received during the evaluation, are added to this deliverable in the Annex A in Section 7.

Each questionnaire is divided into three different parts, each collecting the questions for the corresponding indicator: Usefulness, ease of use, and behavioural intention.

Perceived Usefulness

- I think using the platform would help me perform my daily activities.
- I think the platform is better than the one offered by the current solution.
- Using the platform would allow me to work with a higher degree of security with respect to the current solution.
- Using the platform would allow me to access a higher volume of data with respect to the current solution.
- Using the platform would allow me to have access to more data sources with respect to the current solution.
- Using the platform would allow me to access more services with respect to the current solution.
- Using the platform would allow me to better manage data and information with respect to the current solution.
- Using the platform would allow me to share a larger amount of data with other partners with respect to the current solution.

Perceived Ease of Use

- The interface of the platform is clear and intuitive.
- I found it easy to perform the required tasks using the platform.
- It is easy to find the information needed in the platform.
- The data and analysis results presented by the platform are easy to understand.
- It would be easy for me to become skilful at using the platform.

Behavioural Intention

- I like the idea of using the platform.
- I believe it is a good idea to use the data platform of the platform to manage operations.
- I believe it is a good idea to use the data analytics of the platform to manage and better planning operations.

3.2.3 Results

Table 6 presents an overview of the user acceptance survey results. The table showcases the survey outcomes for each of the four use cases, presenting the three key indicators assessed - usefulness, ease of use, and behavioural intention. Additionally, the last column of the table ("Score") displays the average score for the three indicators.

Use Case	Perceived Usefulness	Perceived Ease of Use	Behavioural Intention	Score
Port of Valencia	3.69	3.3	4	**3.66**
Port of Thessaloniki	3.63	3.76	3.47	**3.62**
Smart Containers	3.5	2.7	3.8	**3.3**
Port Management System	3.1	3.36	3.65	**3.37**

Table 6 - User acceptance survey results

All four use cases have successfully passed the user acceptance survey evaluation, surpassing the minimum score average of 3 by a significant margin. Overall, the results have been exceptional across all indicators, with the perceived ease of use for the Smart Containers being the lowest recorded score, also being the only one not exceeding the threshold of 3.

An excellent result was obtained by the Port of Valencia in the behavioural intention indicator: by obtaining the maximum possible score, it shows that the application developed by the port will bring high value to its future users.

The following subsections collect answers to the open questions posed to all users who tested the platform. Each subsection collects the results for a different use case.

3.2.3.1 Open Questions Port of Valencia

The open questions presented in the Port of Valencia Use Case did generate two open questions in the evaluators.

The first question was "*How much would it cost to integrate a new data source? Which should be the process?*". Some project deliverables provide an answer related to the cost and the way to integrate a new data source:

- From an architectural point of view, Deliverable D2.4 describes the architecture and specification of the Data Access Component. In addition, the role of this component in the use cases is described in Deliverables D5.3 and D5.4.
- From a technical point of view, Deliverables D3.1 and D3.5 provide a complete overview of the data acquisition process in DataPorts.
- From the developer's point of view, this component has a specific section in the documentation portal, courses, and videos on the YouTube channel. In addition, the code, deployments, management user interfaces, and agent templates are provided in the git repository.
- From a development cost point of view, the Key Performance Indicator (KPI) 1.6.3.1 of Deliverable D6.9 focuses on clearly explaining the development effort required to integrate a new data source.
- Finally, from a business plan point of view, Deliverable D7.5 provides an overview of the maintenance cost of integrating a new data source when a third party performs this process.

Regarding the second question, "*Is not clear how can I revoke the access to a data source*", consumer access to datasets can be revoked by the provider in one of two ways:

1. Either as a unilateral action by the provider
2. Or by the provider accepting a revocation request made by the consumer

The first of these is explained in D3.4, Section 4 "Data Governance Services". More specifically, subsection 4.4.5 describes the frontend and Figure 12 shows the "Revoke" button available to the provider. Organisational-level access revocation by the provider is described in D3.7, subsection 2.2, where Figure 13 shows the "Revoke" button. In general, the functionality is available to the provider once they select the "Internal Datasets" menu item and then select the dataset in question. Once inside the dataset details, the

provider must click on the “Permissions” tab and then on the appropriate sub-tab, according to the level of permission they would like to revoke (either “Personal”, “Organisational” or “Public”). Inside each sub-tab the provider can remove any unwanted permissions and then click on the “Save Changes” button.

The second way to revoke consumer access to a dataset happens when the consumer elects to request an access revocation (either user-level or organisational-level) on a specific external dataset. This procedure, as it pertains to user-level access, is described in D3.7, subsection 2.4. Requesting the revocation of organisational-level access to a dataset can only be done by a data governance administrator of the consumer organisation and this is described in subsection 2.2 of the same deliverable. The general procedure is for the consumer user to select the “External Datasets” menu item, then find the datasets of the desired access level by choosing the appropriate sub-tab (either Personal, Organisational or Public) and then select the dataset in question. Once inside the dataset details, the consumer must click on the “Revoke Permissions” button. This creates a notification for the provider, who then needs to open the notification and accept the revocation request.

3.2.3.2 Open Questions Port of Thessaloniki

For the Thessaloniki Port use cases a few comments were recorded regarding any issues found while using the platform or any recommendations the users might have. In general, users replied that the DataPorts platform is easy to navigate, having a user-friendly interface. As for recommendations, the following answers were collected: Translation of the platform to Greek is required, to increase the number of end users. Although there is the choice through a button, the UI is still in English.

- Drop down list of streets of the mobility cases, contain street names, some of which are in English while others in Greek.
- In the public interface the visualisation of the map for the mobility cases is rather slow.

No major issues were recorded; only minor suggestions the platform’s improvement, all of which should be considered for improving the platform in the next iteration.

3.2.3.3 Open Questions Smart Containers

For the Smart Containers use case there were multiple answers to the two open questions proposed. The first question, “Did you have any issue while using DataPorts? Did you find anything not clear?”, tried to explore issues found while using the platform, uncovering the following items:

- Still some issues with data acquisition in DataPorts at the time of testing.
- Sometimes long loading time from login.
- Find a container with a tracking device not knowing its number and trying to discover it in the list of Transport Units.
- Difficult to find all the containers with tracking devices.
- Difficult to understand origin and destination of each container in transit.

The second question is connected to recommendations for changes to the platform: “Would you recommend any change to the platform?”. It collected the following answers:

- No visibility of the geozones shapes for geofencing.
- Having a positions history for a container on a map.
- Add notifications.
- Add a way to extract the data from a list in a spreadsheet file.
- Export data in files.

In general, the open questions for this use case did not reveal any major issue with the platform. The answers received highlight some minor issue with the platform or suggest some improvement, all of which should be considered for improving the platform in the next iteration.

3.2.3.4 Open Questions Port Management System Integration

The Port Management System Integration Questionnaire included three open questions.

The first one is related to the value brought by the integration of DataPorts with the Posidonia applications. “Do you think the integration of DataPorts with the Posidonia applications would help to improve the efficiency at the ports?”

- I definitely believe that the use of Dataports will help to improve efficiency.
- Yes. It will help to decrease times between interactions and speed up processes.
- Yes, as data from Posidonia apps can then be reused by port legacy apps and even be used to create more advanced apps fusing these data with other available data sources.
- Yes, having access to data in real time of what’s happening definitely improves the time between tasks in several ways. For instance, by detecting anomalies.
- Every action which aims to carry out some normalization/standardization with the use of the data among the different stakeholders in a Port Community will improve efficiency while reduce bottlenecks, idling and waste of resources. So, to sum up I would say that yes. DataPorts is a positive action as a transversal action to boost not only ports but also terminals and shipping lines.

All testers agree that the integration of DataPorts with the Posidonia applications would be beneficial, highlighting different advantages.

The second question is related to the benefits of integrating more Posidonia applications with DataPorts: “Would you find useful to integrate more Posidonia applications with DataPorts?”

- Yes, I think the more compatible it is, the better product it will become.
- Absolutely. We have different products depending on the client, so having more of them will increase the sources to detect the events.
- Depends on the needs of each port, but the more info you have available, the more insights you can obtain from data.
- Yes. Having more data available and ready to use is always good to have for any use case.
- The more, the better. IMHO, the only way to starting up offering advanced services (e.g., machine learning, IA, Big Data, IoT, 5G, etc.) and applications is having a solid base of mechanisms to exchange data among the different partners involved in the whole supply chain and a reliable communications infrastructure/network. So, yes. I would find useful to carry out more integrations on the suite of Posidonia which will enable the availability of more data to take better decisions with less uncertainty.

Again, the second question has a positive response: it reveals that the further integration of Posidonia applications with DataPorts would bring value to the platform’s users.

Finally, the third question, “In your opinion, what feature would be interesting to add to the alerting application?”, received a variety of responses: when asked about what additional features should be added to the alerting applications different users provided different suggestions, which should all be taken in consideration when further enhancing the platform.

- Increase the channels through which to notify, as well as make template creation better.
- Maybe to provide an API to manage the subscriptions to topics so they can be managed also from the source applications.
- Maybe a mobile app to receive and browse these notifications.
- A template editor would be good to have. Right now, it’s just a text box, and for example features like variables autocomplete, formatter, and special functions would be some nice features to implement.
- IMHO, the most important functionality with regards Posidonia Notifications (i.e., alerting application) is the capacity to alert/notify staff without being too much verbose (i.e., send a huge number of emails and messages) so that the notifications will be useless since the staff will have a

lot of information to manage in their inbox. Then, the next step should be to have certain intelligence on the core of the application which not only notify/remind one concrete event but also suggest smart actions to support the user who receives the alert as a way to alleviate the load.

The open questions for the Port Management System Integration use case confirmed the output of the questionnaire, revealing that the integration of the Posidonia applications with DataPorts brought value to the users, thus confirming the validity of the pilot. After some tests, users are very interested in using both DataPorts and Posidonia notifications together to help them to achieve better results managing their processes and saving time and let them focus on other activities while they wait for notifications.

4 FURTHER EVALUATION METRICS

In Deliverable D3.5, Section 3 defines the components and the information flows in the platform. Deliverable 5.4 defines the use of the described components to build the different use cases.

Section 4.1 introduces the infrastructure used for deploying the platform, while the following sections describe the evaluation metrics related to software performance and usage metrics.

Finally, the revenue generated is discussed in the Business Plan Deliverable (Section 6 – Revenue Streams, D7.5), while the number of users reached is included inside the General Impact KPIs described in Section 2.2.1 of Deliverable D6.9, which details the results of WP3, 4 and 5, focusing on specific KPIs defined in the scope of the project.

To assess performance metrics in different components, the following open-source tools were used:

- Autocannon[1]: a CLI-based tool for API stress testing that executes thousands of simultaneous petitions to a REST endpoint within a time frame. It supports multi-threading and can simulate multiple clients accessing at the same time.
- Prometheus[2]: an open-source monitoring solution for collecting and aggregating metrics as time series data. The monitored application must collect the metrics and serve them using a REST endpoint. That is commonly known as an "exporter" in Prometheus.
- Grafana[3]: a free and open-source visualization tool that can be used on top of a variety of different data stores but is most commonly used together with Graphite, InfluxDB, Prometheus, and Elasticsearch. It helps users to easily create and edit dashboards.

4.1 PORTS INFRASTRUCTURE

The infrastructure that covers the pilots of Port of Valencia, Smart Containers and Port Management System Integration use cases, and the DataPorts components required to run them, is deployed on the premises of Valencia Port Foundation (VPF) as presented in Section 4.1.1.

The infrastructure for the Port of Thessaloniki use case is described Section 4.1.2.

Finally, Section 4.1.3 reports the infrastructure used by Prodevelop for the Posidonia Port Management System for the global use cases and for providing Posidonia operations data to the Ports of Valencia and Thessaloniki.

For the Smart Containers use case there is no specific hardware deployed, as the pilot is running on the premises of Port of Valencia and Port of Thessaloniki, as described in Sections 2 and 3.

4.1.1 Port of Valencia

Table 7 reports the specifications of deployed hardware and the corresponding applications running on them.

Specifications	Use
4 vCPU, 8 GB RAM, 200 GB SSD	Kubernetes (Rancher) Blockchain Tracking application VGM application DCN application
4 vCPU, 8 GB RAM, 200 GB SSD	
6 vCPU, 16 GB RAM, 400 GB SSD	

[1] https://github.com/mcollina/autocannon

[2] https://prometheus.io/

[3] https://grafana.com/

Specifications	Use
6 vCPU, 16 GB RAM, 400 GB SSD	IAM (NTT Data) & Data Governance
8 vCPU, 30 GB RAM, 800 GB SSD	Data Processing Services (PRO+UPV+ICCS)
6 vCPU, 16 GB RAM, 400 GB SSD	Data Analytics (ITI) + Application (ITI)
4 vCPU, 8 GB RAM, 200 GB SSD	IDS (ITI)
6 vCPU, 16 GB RAM, 400 GB SSD	Process-based Analytics (UDE)

Table 7 - Port of Valencia infrastructure

4.1.2 Port of Thessaloniki

For the THPA pilot, two virtual machines were deployed. Both virtual machines are in THPA's VMWare infrastructure on a VxRAIL cluster. The cluster is comprised of 8 physical hosts (4 in the primary computer room and 4 in the secondary computer room, in an active-active configuration) and managed by vSphere version 7.0.3 The first VM (named vm-certh) has 4 CPUs, 16GB of RAM and 100Gb hard disk. It runs Ubuntu Linux 18.04. It is used for the Blockchain infrastructure and the analytics. The second VM (named vm-dpagents) has 4 CPUs, 16GB of RAM and 80Gb hard disk. It runs Ubuntu Linux 20.04. It is used for the agents.

4.1.3 Port Management Use Case

In this scenario we differentiate the applications that existed before DataPorts, and the new components implemented in this project.

In the first group we identify Posidonia Management, Posidonia PCS and Posidonia Operations. The required infrastructure for running these applications is out of the scope of this document.

On the other hand, we find the Data Access Component (DAC), The Semantic Interoperability Component and Posidonia Notifications. All these components have been deployed under Prodevelop provided hardware. The first two run in a VM while the application Posidonia Notifications is running in a docker container in a 5-node cluster managed with Swarm. Table 8 contains the specifications of the servers and VMs involved in the global pilot.

Specifications	Use
4 vCPU, 8 GB RAM, 100 GB HDD	Data Access Component (DAC)
44 vCPU, 250 GB RAM	Posidonia Notifications

Table 8 - Infrastructure for the Port Management System Integration use case

The application Posidonia Notifications is built from three projects: background, foreground and the service that sends the messages. The cluster where it is deployed is shared with other applications.

4.2 DATA ACQUISITION

The Data Access Component is built up using several services:

- A web UI to facilitate users some basic operation of the overall acquisition component.
- A REST API with a total of 53 endpoints used either by the web interface and by agents to import data into the system.
- 28 agents (16 "publish-subscribe" and 12 "on-demand") fulfilling all data import needs from all use cases.

The API pushes data to 2 sibling components:

- Publish-subscribe agents send data to the semantic interoperability component, more specifically to the “Orion Context Broker” by invoking its own REST API
- On-demand agents send the data to the data abstraction and virtualization component, more specifically to “Cygnus” using by invoking its own REST API

In the next subsections we will analyse the performance of the DAC’s API and a couple of agents which perform the heaviest data import duties. The web interface is not considered in this analysis since it is executed at the user’s web browser and therefore, does not add workload to the server.

4.2.1 DAC’s API testing

As commented above, the API is composed of 53 endpoints, but only a few (5) are used by agents to import data. We will run a stress test with autocannon with the following parameters:

- 10 parallel threads (equivalent to 10 agents connecting concurrently)
- 50 connection (5 per worker)
- 5000 requests (500 per worker)

The endpoints of interest are:

- GET /info/<<agent_id>>: this endpoint provides metadata about the “docker container” inside which the agent runs

In the test, we force a 404 error by querying a non-existing container. The command issued is:

```
autocannon -c 50 -a 5000 -w 10  -H Authorization="Bearer %TOKEN%" -H "Content-Type"="application/json" -H "Accept"="application/json" http://localhost:3000/info/fake_container
```

As a result, we observe that the endpoint can ingest 5000 requests from 50 clients with a latency of 148ms.

```
Running 5000 requests test @ http://localhost:3000/info/fake_container
50 connections
10 workers
```

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	94 ms	138 ms	247 ms	355 ms	148.51 ms	55.28 ms	690 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	150	150	337	422	332.34	62.8	150
Bytes/Sec	44.6 kB	44.6 kB	100 kB	125 kB	98.7 kB	18.6 kB	44.5 kB

```
Req/Bytes counts sampled once per second.
# of samples: 155

5k requests in 15.23s, 1.49 MB read
```

Figure 1 - Stress test for agent info request endpoint

- DELETE /info/<<agent_id>>: this endpoint is used by on-demand agents to signal the DAC that the container has finished running. If the agent is not found, it returns a 404 status code.

We replicate the previous test but attempting a deletion of a non-existing container. The command issued is in this case:

```
autocannon -c 50 -a 5000 -w 10 -H Authorization="Bearer %TOKEN%" -H "Content-Type"="application/json" -H "Accept"="application/json" -m DELETE  http://localhost:3000/info/fake_container
```

Results are similar, but in this case the latency is higher (190ms)

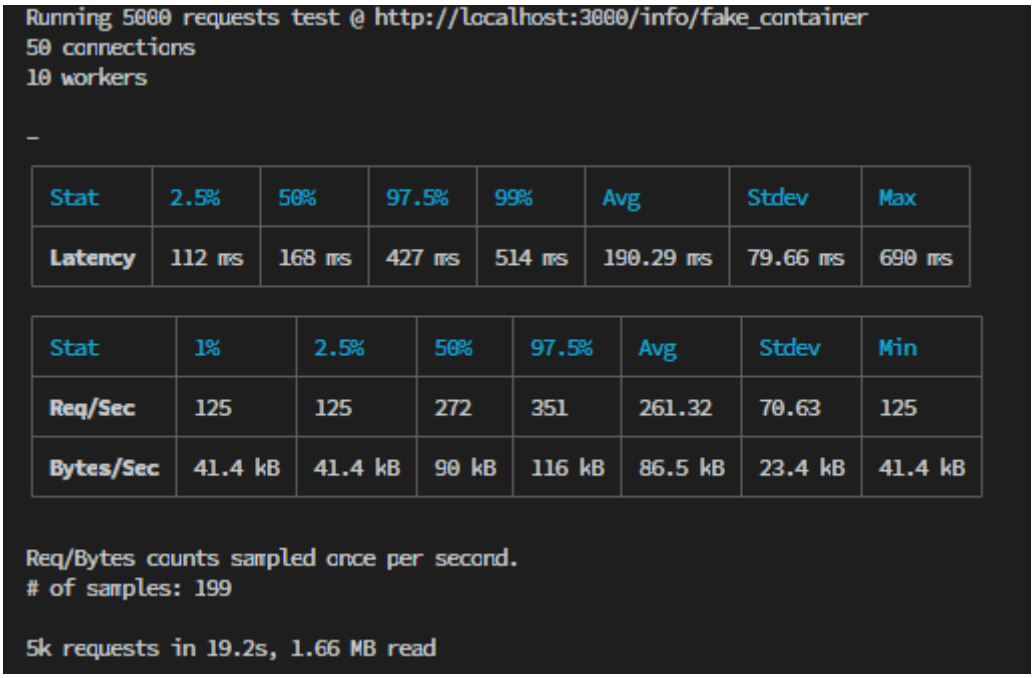

Figure 2 - Stress test for agent deletion request endpoint

- POST /notification: this endpoint is used by the agent to log events that are of interest to the DAC's users. For example, it can log when an agent starts/stops executing, when an error occurs (with stack traces information for debugging purposes, etc).

The command to run the test in autocannon is:

```
autocannon -c 50 -a 5000 -w 10 -m POST -H Authorization="Bearer %TOKEN%" -i ./tests/notification.json http://localhost:3000/notification
```

We are using as payload a notification that contains a 1.3KB json file with an error notification. The stress test gives the results shown in Figure 3 and Figure 4:

```
/
```

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	100 ms	179 ms	478 ms	614 ms	207.18 ms	103.43 ms	911 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	93	93	243	370	238.1	81.08	93
Bytes/Sec	38.5 kB	38.5 kB	101 kB	153 kB	98.6 kB	33.6 kB	38.5 kB

```
Req/Bytes counts sampled once per second.
# of samples: 210

5k requests in 21.22s, 2.07 MB read
```

Figure 3 - Stress test for the notification endpoint

Figure 4 - API Metrics during notification endpoint stress testing

- POST /cygnus: this endpoint is invoked by on-demand agents to create entities in cygnus. This endpoint accepts an array of entities which will be stored in the same collection in a single transaction. The array size varies between 25 and 75, depending on the agent and the size of the data, since Cygnus limits the payload to 2MB per call.

In this test, due to size of the request's payload (102kB), we have reduced the number of connections to 10 and to a total number of requests. The command issued is:

autocannon -c 10 -a 500 -w 10 -m POST -H "Content-Type"="application/json" -H "Accept"="application/json" -i .\tests\cygnus.json http://localhost:3000/cygnus

The latency of the requests is high (1.22 seconds on average, as shown in Figure 5) since we are sending 30 entities on each request and cygnus internally decomposes that array in 30 atomic data insertions and there are 10 concurrent requests at all times. When running a similar analysis with just 1 concurrent request, the latency slightly decreases to 1.06 seconds, meaning that most of the time is used by cygnus to process the data.

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	730 ms	1189 ms	1977 ms	2173 ms	1224.55 ms	313.1 ms	2418 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	0	4	8	11	8.05	1.97	4
Bytes/Sec	0 B	1.18 kB	2.36 kB	3.25 kB	2.38 kB	579 B	1.18 kB

```
Req/Bytes counts sampled once per second.
# of samples: 617

500 requests in 61.26s, 148 kB read
```

Figure 5 - Stress test for the cygnus ingestion endpoint

When looking at the API metrics (Figure 6), we observe that there is an increase in the use of memory and CPU, that is freed just after the stress test finishes.

Figure 6 - API Metrics during cygnus ingestion endpoint stress testing

- POST /cygnusInformation: this endpoint is used also by on-demand agents to notify the DAV's component that a new dataset is available in cygnus. The payload carries the name of the collection created and the size of the collection. This endpoint was not stress-tested since the DAV component was not available in the testing machine. However, this endpoint is rarely used and could not pose a performance problem in the DAC.

4.2.2 Agents testing

Agents created are able to import data from different types of data sources (e.g., polling an external REST API, importing data from CSV files, listening to message queues, etc.) and map the source data to compliant DataPorts models (see https://egitlab.iti.es/dataports/data_processing/datamodel).

The test is done on developer's computer running all the stack (apart from other desktop apps to stress the machine's memory and processing capabilities). The results obtained in this test should be considered as a baseline for a worst-case scenario, since servers' setup in the different pilots have larger resources and therefore will perform better).

On-Demand

The most intensive agent in terms of processing requirements is the "Customs" on-demand agent, which imports publicly available customs information from the Spanish Tributary Agency [2]. Each year contains around 100 million registries and 32 mapped fields per registry.

To test the performance of this agent, we ran an import job for just 1 month (December 2022), which will import 9.9 million registries into cygnus.

The assessment was done by collecting metrics using a Prometheus and charting them with Grafana. In more detail, we prepared two dashboards to assess both the agent and the API performance during the import job execution:

The first dashboard (Figure 7) contains metrics related to memory and CPU use and network traffic (input and output) of the docker container running the agent. The memory charts show a consistent use of memory during the execution meaning that once the agent starts and Docker creates the container, it won't require more resources to run the import job. The spike observed in the graph 1.5h after the age starts was caused due to a SSH connection to the agent's container to validate data files had been correctly downloaded.

The CPU use is also consistent to an average of 10% of 1 single processor's computing capacity. The IO traffic also shows that the agent downloads a high volume of data at the beginning (it essentially downloads all source data from the Spanish Tributary Agency's web site) and then the output traffic is steady at around 1MB/s, meaning that the agent is sending data to cygnus as fast as it can ingest those data.

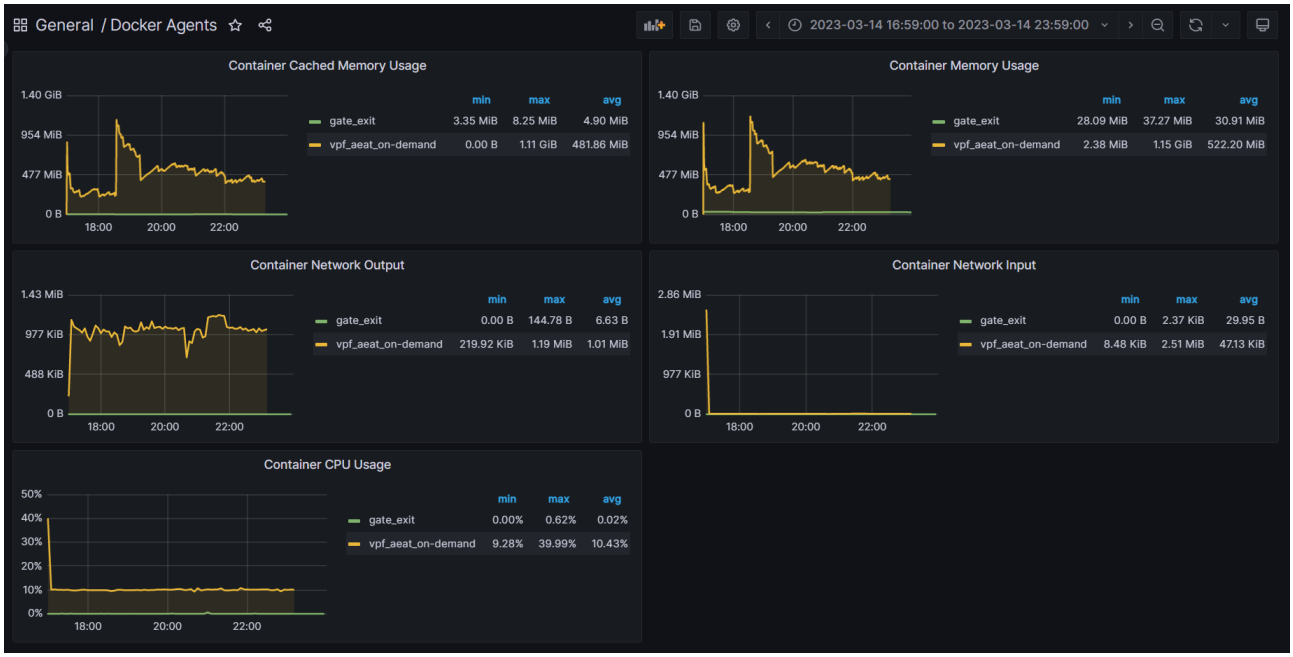

Figure 7 - agent performance dashboard

The second dashboard (Figure 8) shows the API metrics during the same period as in Figure 7.

Figure 8 – DAC's API performance dashboard

On the upper row we can see some instant values for the CPU and memory used by the API, as well as a histogram of running agents (int this case, we are running 2 agents: "customs" on-demand, and a publish-subscribe that has low latency and does not interfere with the analysis).

The middle row shows the average number of calls to API endpoints per second and the CPU used by the API itself. As it can be observed, the agent operates at full capacity by calling to the cygnus endpoint around 15 times/second and consuming around 50% of the capacity of 1 processor in the machine.

The lower row shows the memory use of the API, where it is observed that the use is steady at around 80 MB with no memory leaks appreciated.

In Figure 9 an overview of the imported data by Cygnus (and persisted in a MongoDB collection) is shown. The agent creates a new collection for each month of data imported (in this case "col_/_urn:ngsi-ld:ITI:Customs:2022:12_Customs") which contains the predicted 9.9 million entities with a raw size of 7.6GB.

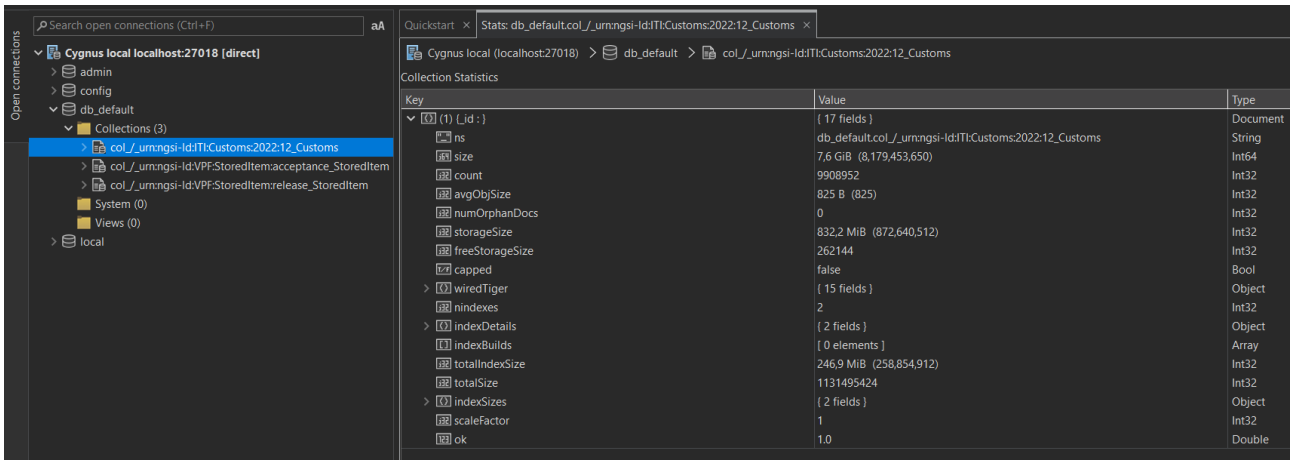

Figure 9 - imported data in Cygnus and stored in MongoDB

In summary, the execution of the agent took 6h15m, being the download time of raw data negligible. That gives an overall performance of 524.3 records imported per second.

Publish-Subscribe

The most demanding publish-subscribe agent is the PCS agent which imports port calls information for the current month and stores/updates the information in Orion, which was extensively introduced in Deliverable D3.5, Section 5.2.1. This agent does not persist data through the API, but connects directly to the Orion instance, meaning that the agent performs all the workload.

The agent is programmed to run every minute; it downloads from the VPF systems all the port calls for the current month (minimum granularity allowed by the port's API) and processes the results sequentially, mapping the information to the DataPorts data model and sending each value to Orion.

In Figure 10 we can observe 20 executions of the agent. All charts show steady resources use, meaning that no memory leaks nor increased processing power is needed between executions. Each execution lasts for 6 seconds on average, being 99% of that time the processing and storage of the data in Orion.

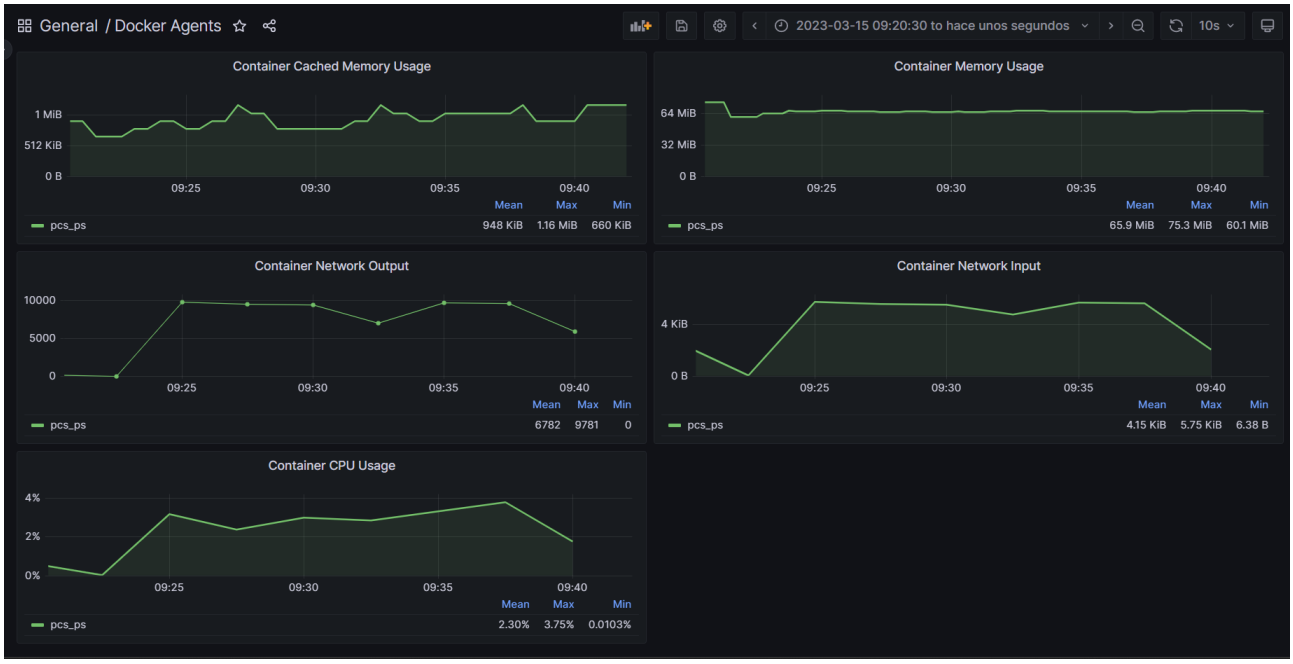

Figure 10 - PCS publish-subscribe agent metrics

In Figure 11 there is an extract of one agent run, showing the time it takes to download data from the port's API (orange area, ~600ms) and the latency to map and store each data point (green area, ~15.5ms).

```
2023-03-15 09:36:49 2023-03-15 08:36:49.089 DEBUG __init__ - run: Running job Job(interval=1, unit=minutes, do=main, args=(), kwargs={})
2023-03-15 09:36:49 2023-03-15 08:36:49.089 INFO script_PCS_ITI - get_vessel_calls_file: Requesting vessel calls for year=2023 month=3
2023-03-15 09:36:49 2023-03-15 08:36:49.090 DEBUG connectionpool - _new_conn: Starting new HTTPS connection (1): www.valenciaportpcs.net:443
2023-03-15 09:36:49 2023-03-15 08:36:49.289 DEBUG connectionpool - _make_request: https://www.valenciaportpcs.net:443 "POST /portcalls/search/ExportBerths HTTP/1.1" 200 123392
2023-03-15 09:36:49 2023-03-15 08:36:49.362 INFO script_PCS_ITI - get_vessel_calls_file: Creating file "./data/2023_03.xls"
2023-03-15 09:36:49 2023-03-15 08:36:49.362 DEBUG script_PCS_ITI - main: Reading files...
2023-03-15 09:36:49 Getting file:  2023_03.xls
2023-03-15 09:36:49 2023-03-15 08:36:49.721 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.736 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.750 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.767 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.784 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.801 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.822 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.836 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.849 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.865 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.879 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.892 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
2023-03-15 09:36:49 2023-03-15 08:36:49.908 DEBUG connectionpool - _make_request: http://host.docker.internal:1026 "POST /v2/entities?options=upsert HTTP/1.1" 204 0
```

Figure 11 - PCS agent run logs

4.3 SEMANTIC INTEROPERABILITY

In this section tests about Orion's and KrakenD's performance in the platform and the performance impact the publish / subscribe agents interacting with them are reported. To this end, Orion has prepared some load tests on their service, providing results for a machine of similar specifications as the machine in the Valencia

Port pilot. However, these tests are for a machine that is exclusively running Orion, the project has multiple agents as well as the Data Access Manager, the DAV component and the API gateway that ties everything together. To present a more realistic image of the component, tests were executed using the default docker tools as well as the third-party tool "Autocannon".

According to the official Fiware load test repository Orion can take a much heavier load in a similar position, using a machine with similar specs, although their tests assume running Orion almost exclusively. With this in mind, logic follows that Orion is going to have a similar behaviour in the pilots that have a lesser network load than what Fiware tested with Gatling, but a higher system load due to the agents being run in it. The number of petitions shown is due to the amount of publish/subscribe agents used in the pilots and their activity, since no agent is neither continuously sending petitions nor activating at the same time as other agents.

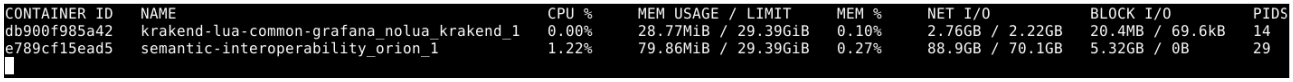

CONTAINER ID	NAME	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
db900f985a42	krakend-lua-common-grafana_nolua_krakend_1	0.00%	28.77MiB / 29.39GiB	0.10%	2.76GB / 2.22GB	20.4MB / 69.6kB	14
e789cf15ead5	semantic-interoperability_orion_1	1.22%	79.86MiB / 29.39GiB	0.27%	88.9GB / 70.1GB	5.32GB / 0B	29

Figure 12 - Docker stats for both the KrakenD and Orion containers

It is worth mentioning how little system memory these components used throughout all the real deployment tests. This will prove to be very useful when compared with some of the more resource intensive elements of the deployment.

The test starts multiple connections to retrieve all the entities in Orion at the same time, which is quite intensive for Orion. If several thousand petitions are launched directly to the endpoint /v2/entities ignoring the gateway a spike in the CPU usage can be seen, but that's nothing Orion is not prepared to handle.

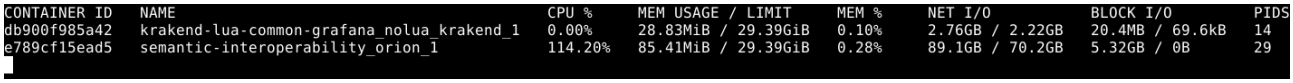

CONTAINER ID	NAME	CPU %	MEM USAGE / LIMIT	MEM %	NET I/O	BLOCK I/O	PIDS
db900f985a42	krakend-lua-common-grafana_nolua_krakend_1	0.00%	28.83MiB / 29.39GiB	0.10%	2.76GB / 2.22GB	20.4MB / 69.6kB	14
e789cf15ead5	semantic-interoperability_orion_1	114.20%	85.41MiB / 29.39GiB	0.28%	89.1GB / 70.2GB	5.32GB / 0B	29

Figure 13 - Docker stats during the load test for Orion

Using the third-party tool "Autocannon" more detailed results of the performance tests can be seen. In this case the tests ran 10 connections for 10 seconds, and almost 100MB of raw data with an average of 366 petitions per second was seen, which is more than enough for the pilots and use cases of the DataPorts Project. As a reference, according to the data provided by the Prodevelop team, no pilot is going to have more than 10 publish/subscribe agents running at the same time and the agents do not need to send data at the same time either, this means that even if the machine is getting overloaded, it will be able to continue to work as it is.

```
Running 10s test @ http://localhost:1026/v2/entities
10 connections

┌─────────┬──────┬──────┬────────┬──────┬──────────┬──────────┬──────────┐
│ Stat    │ 2.5% │ 50%  │ 97.5%  │ 99%  │ Avg      │ Stdev    │ Max      │
├─────────┼──────┼──────┼────────┼──────┼──────────┼──────────┼──────────┤
│ Latency │ 4 ms │ 26 ms│ 51 ms  │ 56 ms│ 26.74 ms │ 9.46 ms  │ 86 ms    │
└─────────┴──────┴──────┴────────┴──────┴──────────┴──────────┴──────────┘

┌──────────┬──────────┬──────────┬──────────┬──────────┬──────────┬──────────┬──────────┐
│ Stat     │ 1%       │ 2.5%     │ 50%      │ 97.5%    │ Avg      │ Stdev    │ Min      │
├──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
│ Req/Sec  │ 336      │ 336      │ 360      │ 406      │ 366.2    │ 21.28    │ 336      │
├──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
│ Bytes/Sec│ 8.94 MB  │ 8.94 MB  │ 9.58 MB  │ 10.8 MB  │ 9.74 MB  │ 565 kB   │ 8.93 MB  │
└──────────┴──────────┴──────────┴──────────┴──────────┴──────────┴──────────┴──────────┘

Req/Bytes counts sampled once per second.
# of samples: 10

4k requests in 10.03s, 97.4 MB read
```

Figure 14 - Autocannon output for the Orion load test

Alternatively, when running the same tests against the KrakenD gateway Orion never actually reaches the same CPU spikes as before, as the petitions are passing through KrakenD.

```
NAME                                      CPU %     MEM USAGE / LIMIT     MEM %     NET I/O           BLOCK I/O         PIDS
krakend-lua-common-grafana_nolua_krakend_1 20.63%   38.87MiB / 29.39GiB   0.13%     3.4GB / 2.34GB    20.6MB / 69.6kB   16
semantic-interoperability_orion_1         13.38%    80.57MiB / 29.39GiB   0.27%     89.2GB / 70.4GB   5.33GB / 0B       20
```

Figure 15 - Docker stats for the KrakenD load test

When looking at the Autocannon stats what can be seen is that there is an average of 32 petitions per second being accepted by KrakenD, this is due mainly to the fact that KrakenD is verifying every petition with Keycloak, which is in a different machine in a completely different network. Regardless of this, KrakenD will be able to handle the incoming requests from the pilots. It would work even better in a scenario where the identity manager is behind Keycloak.

```
Running 10s test @ https://localhost:9000/orion/v2/entities
10 connections

┌─────────┬──────┬──────┬────────┬──────┬──────────┬──────────┬──────────┐
│ Stat    │ 2.5% │ 50%  │ 97.5%  │ 99%  │ Avg      │ Stdev    │ Max      │
├─────────┼──────┼──────┼────────┼──────┼──────────┼──────────┼──────────┤
│ Latency │276 ms│292 ms│ 398 ms │424 ms│ 301.55 ms│ 33.6 ms  │ 473 ms   │
└─────────┴──────┴──────┴────────┴──────┴──────────┴──────────┴──────────┘

┌──────────┬──────────┬──────────┬──────────┬──────────┬──────────┬──────────┬──────────┐
│ Stat     │ 1%       │ 2.5%     │ 50%      │ 97.5%    │ Avg      │ Stdev    │ Min      │
├──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
│ Req/Sec  │ 28       │ 28       │ 33       │ 36       │ 32.6     │ 2.38     │ 28       │
├──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┼──────────┤
│ Bytes/Sec│ 746 kB   │ 746 kB   │ 880 kB   │ 960 kB   │ 869 kB   │ 63.3 kB  │ 746 kB   │
└──────────┴──────────┴──────────┴──────────┴──────────┴──────────┴──────────┴──────────┘

Req/Bytes counts sampled once per second.
# of samples: 10

336 requests in 10.08s, 8.69 MB read
```

Figure 16 - Autocannon output for the KrakenD load test

4.4 DATA ABSTRACTION AND VIRTUALIZATION

The Data Abstraction & Virtualization (DAV) component is designed to efficiently pre-process, filter & clean incoming datasets, and then store them to a data lake, DAV's Virtual Data Repository. The component has been designed in a way that allows the completion of the whole flow in a small-time window, which (of course) depends on the incoming datasets' volume. Regarding the dataset used for testing, it is the "urn:ngsi-ld:VPF:PortCall", which contains about 31200 documents. It cannot be considered as a "big data" dataset, but DAV's efficiency can be highlighted by showing its speed while processing & storing it. The resources used (along with the results) are the same in both Port virtual machines (VPF and THPA). Both resources and results can be seen through DAV's Apache Spark Cluster WebUI since DAV is hosted in an Apache Spark Cluster environment. The image below shows the Spark driver's ID, the timestamp and the resources used (by the driver) while running the Pre-processing & Filtering Software (PaFS), DAV's first sub-component for processing, filtering, cleaning and storing the dataset. It is shown in Figure 17.

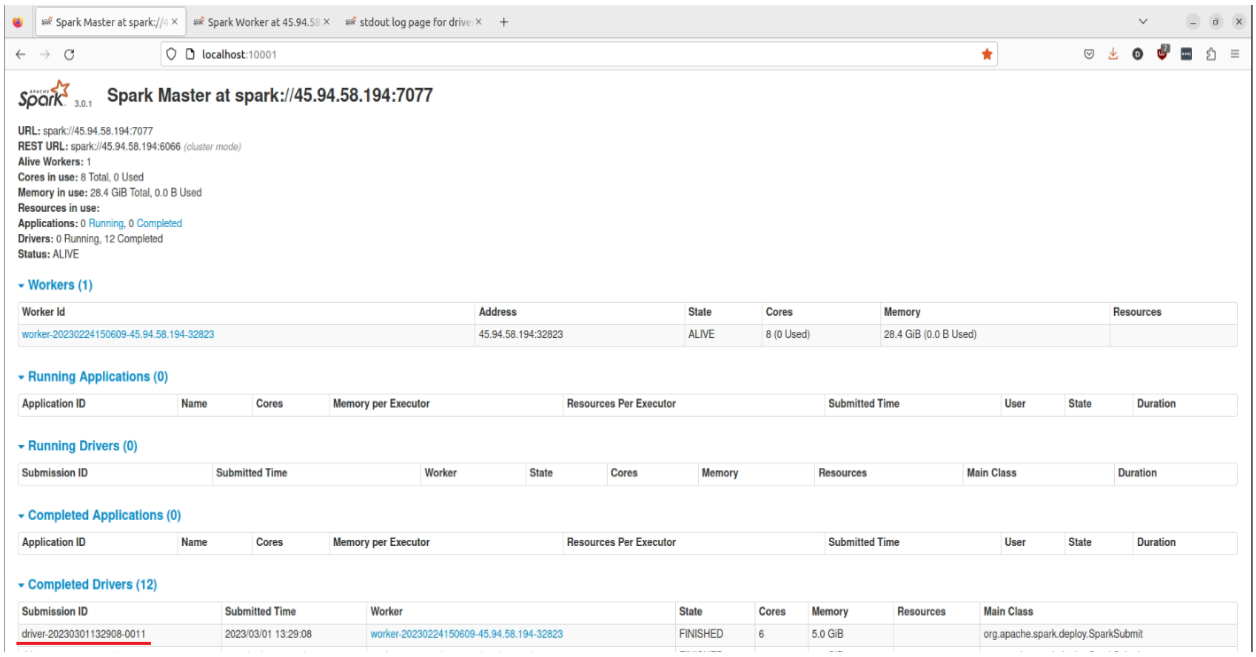

Figure 17 - Screenshot from DAV's Apache Spark Cluster WebUI

As seen in Figure 17, driver with ID "20230301132908-0011" was used in order to load PaFS and process, filter, clean, and store the dataset. This driver has been triggered by the On-demand agent which is responsible for loading the dataset. This means that the whole flow is being done automatically, with each sub-component (or software tool) triggering the next. Our Spark cluster driver used 6 CPU cores and 5 GBs of RAM in order to run PaFS. The time it took to complete the whole process can be seen in Figure 18, inside the driver's logs.

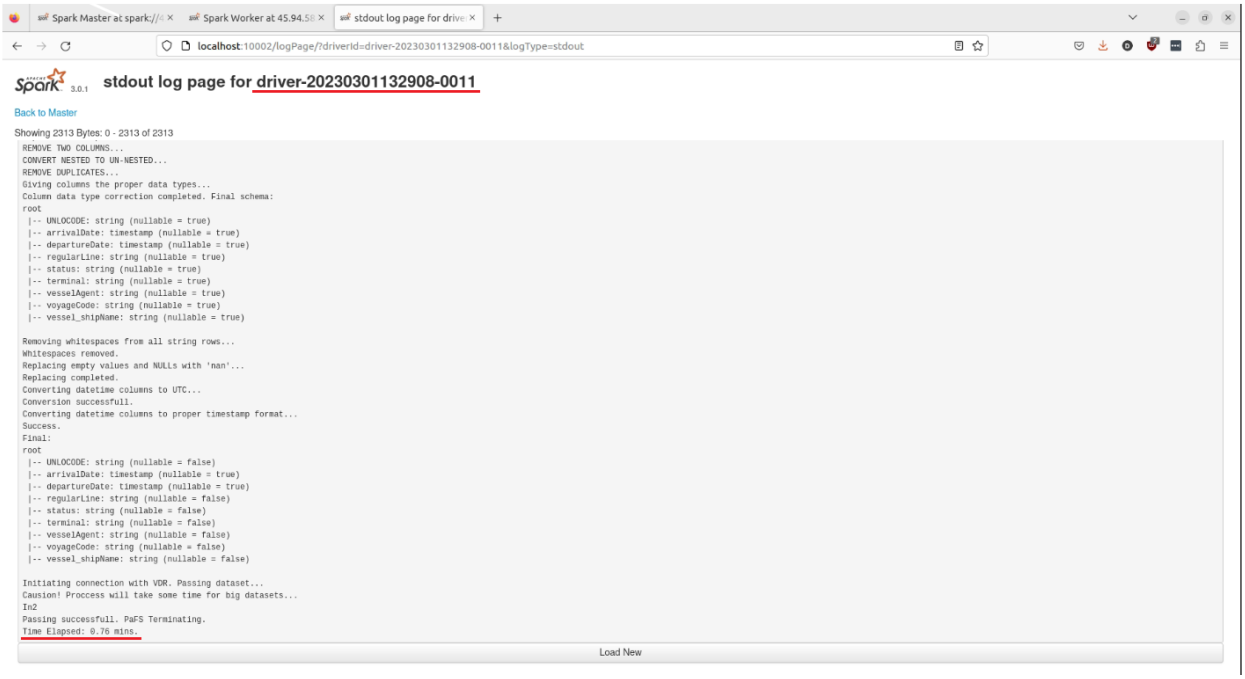

Figure 18 - Screenshot from Spark cluster driver's logs

The log's title confirms that we are inside the correct driver's information. In the end of the log file, one can read the total time it took for PaFS to process, filter clean, and store the selected dataset. The time it took was 0.76 minutes, which translates roughly to 45 seconds. This is a very short period of time, given the series of tasks PaFS applies to the dataset. In the end, the dataset is being stored clean and safe in the Virtual Data Repository, ready for retrieval by a potential data recipient.

4.5 DATA GOVERNANCE

The performance for the Data Governance Hyperledger Fabric backend was tested on a Virtualbox machine running Ubuntu 18.04 with 8GB of RAM, a 60GB disk, and 2 CPUs. To monitor performance we utilized Prometheus, a tool for monitoring system metrics. Node Exporter, designed for collecting hardware and OS metrics, was used in conjunction with Prometheus. For visualization, we relied on the Grafana platform.

The tests focused on the creation of a new external organisation. An external organisation is an organisation that participates in the Data Governance ecosystem without participating in the blockchain network, i.e., the organisation does not contribute to the network with its own peer machines, but is rather a logical entity within Data Governance. A major reason external organisations were implemented in this way was to minimise their impact on the performance of Data Governance as a whole. Another was to also simplify their use of Data Governance, removing the need for them to create their own peer and connect it to the Data Governance network. Still, creating an external organisation is one of the more intensive functionalities in Data Governance, since, along with the organisation itself (the logical entity), we also create its admin user. Therefore, it is important to see that this functionality, although intense compared to others, has indeed been kept to low levels of CPU and memory consumption, as was the initial intent.

During the procedure, the organisation's logical entity is created and added to the blockchain. Next, the admin is registered with the certificate authority and the necessary certificates are produced for the user. Figure 19 & Figure 20 depict the RAM and CPU usage during request processing, revealing that this particular process does not consume a significant amount of machine resources.

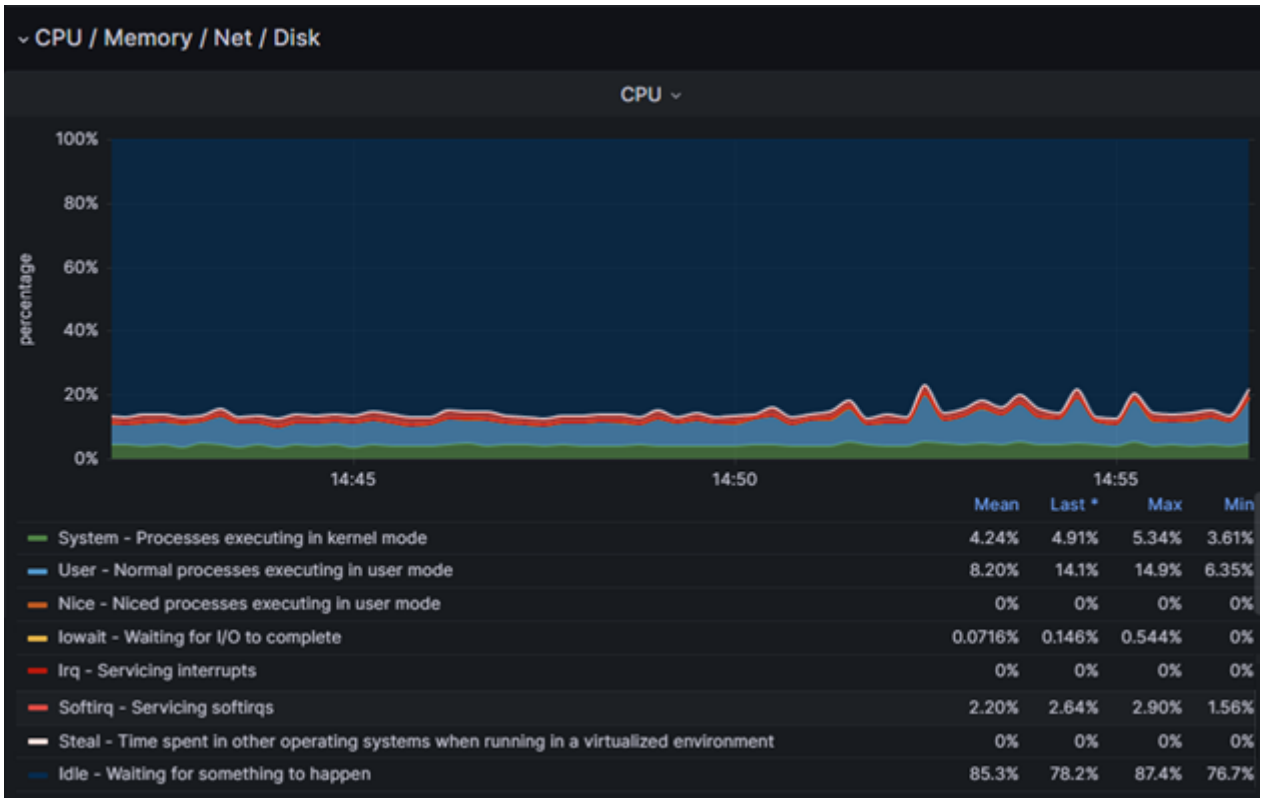

Figure 19 - Data Governance CPU utilisation

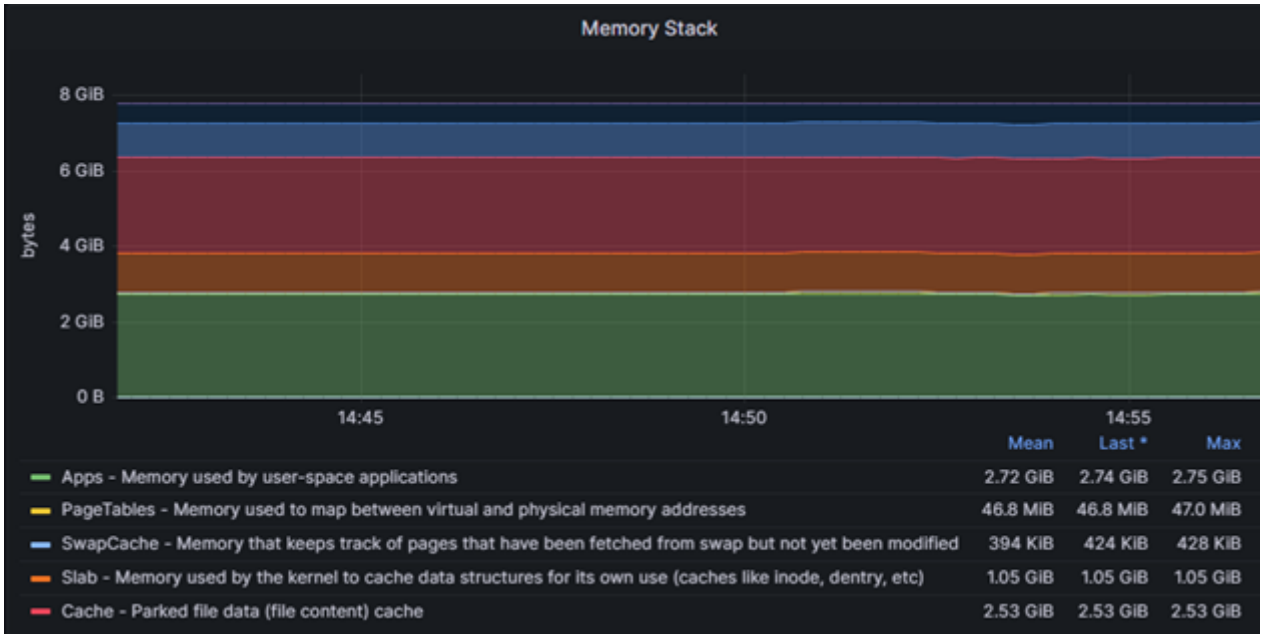

Figure 20 - Data Governance memory utilisation

Using Postman, we have also timed the two most important endpoints Data Governance offers to other components of the platform. Their average measured execution times are shown below:

- `/getPermissionsByDataSourceID => 434.8 ms`
- `/queryMetadataByDataSourceID => 314.2 ms`

These times are considered acceptable for the needs of the platform.

4.6 DATA ANALYTICS

Regarding the Data Analytics services, a set of evaluation metrics were excerpted to fully appraise the performance of the components.

4.6.1 Automatic Models Training Engine (AMTE)

To perform the evaluation process, the component was tested in an infrastructure provided by Valencia Port where the component is deployed and which provides the following computing parameters: 6 vCPU, 16 GB RAM, 400 GB SSD.

4.6.1.1 Cognitive Services Metrics

A total of **8 different cognitive services** were created to test the feasibility of all the different data mining approaches to extract insights for port-business KPIs. That implied the training and cross validation of a total of **920 Machine Learning algorithms** by following a series of data pipelines designed to optimize the data processing phase and thereby maximize the quality of the predictive models obtained.

Name	Task	Type	Description	Last trained date	Status		
Average Vessel Berth Time	Time Series Forecasting	Average Vessel Berth Time	Saturation of the port of Valencia in the next 3 weeks	18/11/2022 11:14:23	RUNNING	Stop	•••
Customs Trade Volume	Time Series Forecasting	Customs Trade Volume	Estimation of the tons of animal products imported to Valencia in the next 3 months	18/11/2022 11:17:12	RUNNING	Stop	•••
Missing Origin/Destination Identification	Values Imputation	Missing Origin/Destination Identification	Prediction of the unknown district from the historical data of PCS traceability	18/11/2022 11:20:23	RUNNING	Stop	•••
Machinery forcasting 7 days	Time Series Forecasting	Container Goods Volume	Number of TEUS of machinery imported to Valencia in the next 7 days	18/11/2022 11:30:52	RUNNING	Stop	•••
Glass products imported to Valencia in 3 months.	Time Series Forecasting	Container Goods Volume	Amount of TEUs of glass products expected to arrive to Valencia in the next 3 months.	23/01/2023 10:56:19	RUNNING	Stop	•••
Valencia Ports congestion in the next 3 weeks	Time Series Forecasting	Vessels Port Calls Calculator	Number of port calls that are expected at the Port of Valencia in the next 3 weeks	08/03/2023 07:08:33	RUNNING	Stop	•••
Congestion of the port of Valencia in the next 7 days	Time Series Forecasting	Vessels Port Calls Calculator	Congestion of the port of Valencia in the next 7 days	08/03/2023 09:02:32	RUNNING	Stop	•••
Vessels Time of Departure Estimator ..	Time Series Forecasting	Vessel Time of Departure Estimator	Time Series Forecasting Vessel Time of Departure Estimator	09/03/2023 10:24:23	RUNNING	Stop	•••

Figure 21 – Cognitive services created

Consequently, to train and evaluate the whole set of models a total of **1607,86 GB of data** were utilized by following a cross-validation approach.

Service Type	Dataset	Dataset size (GB)	# of services	Total size (GB)
VTDE	PCS Traffic	0,009	1	0,009
VPCC	PCS Traffic	0,009	2	0,018
AVBT	PCS Traffic	0,009	1	0,009
CTV	AEAT	14,5	1	14,5
CGV	PCS traceability	531	2	1062
MODI	PCS traceability	531	1	531
			Total	1607,536

Table 9 – Amount of data utilized to train the set of cognitive services

The whole training collection took **54 mins of training**, making it an average training time of **6.75 mins per service**.

4.6.1.2 Parallel training of data pipelines

The training of data pipelines is performed by following a parallel computation as it can be observed in the following figures, **where the whole set of available CPUs are utilized reaching 100% of their capacity**.

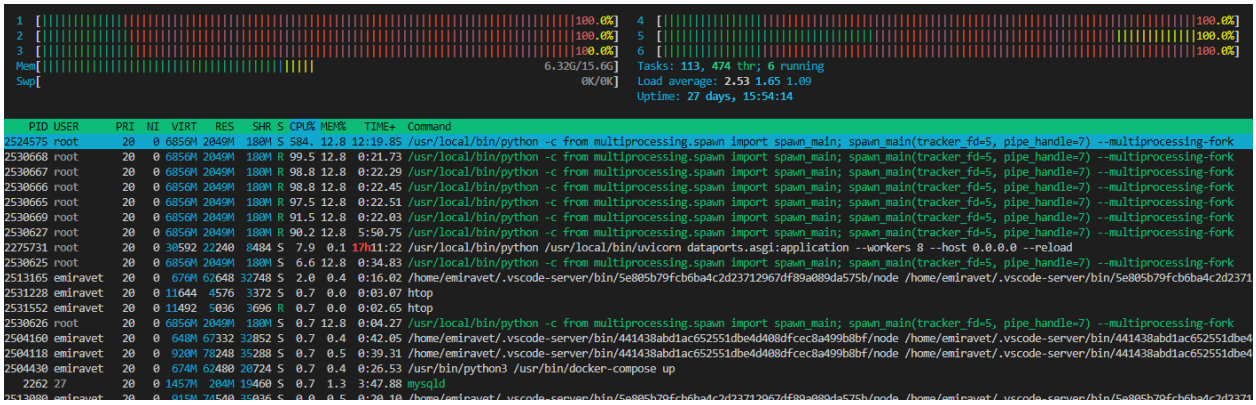

Figure 22 – Use of CPUs while training data pipelines in parallel

4.6.1.3 Comparison of AMTE predictions with the read ETD

To evaluate the performance of the predictions made by the Automatic Models Training Engine component, the cognitive service Vessels Time of Departure Estimator was selected due to its potential to compare the ETD predicted by the best model trained and the real ETD registered by the PCS calls system.

Figure 23 shows a boxplot representing the time difference between AMTE predictions and the real ETD, along with some statistics excerpted from it.

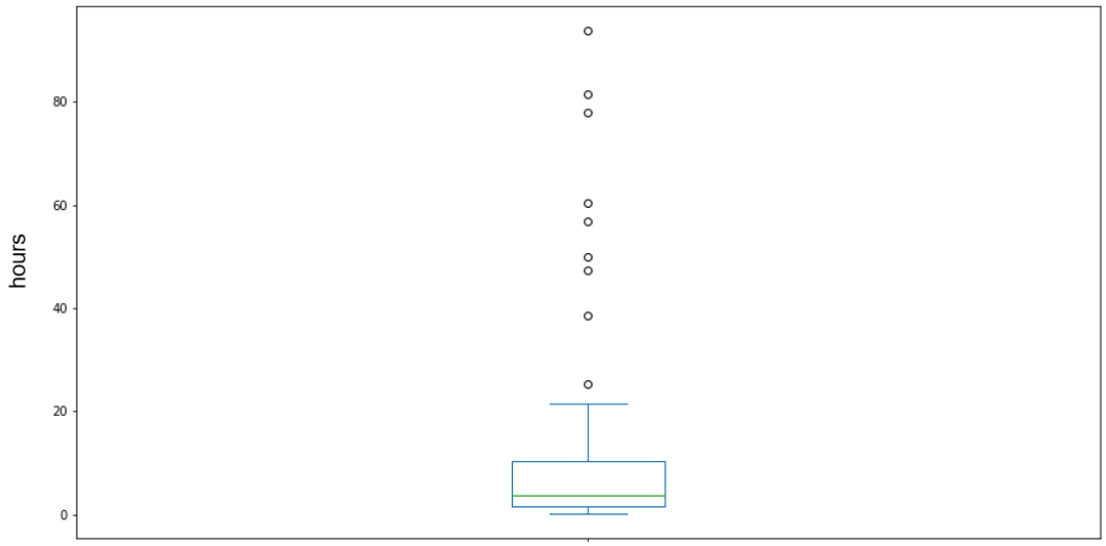

Figure 23 – Boxplot of the time difference between AMTE and reality

Median	count	mean	std	min	P25%	P50%	P75%	max
3,76 h	140,00	9,13 h	15,12 h	0,09 h	1,43 h	3,77 h	10,26 h	93,88 h

Table 10 - Statistics excerpted from the boxplot.

Figure 23 and Table 10 set out that the deviation from reality achieved by AMTE is highly promising as the minimum reached is close to zero and the median is quite low, achieving a value of 3.76 hours. The deviations are quite concentrated into the median as showed by the standard deviation (std), percentile P25 and percentile P75, meaning that the vast majority of predictions are below 10h of deviation. Nevertheless, there are some specific high deviations in certain vessels (small circles above P75) reaching a maximum value of 93.88h, that are considered as outliers and can be due to anomalous operations of vessels, namely unexpected delays or technical issues, or even a lack of additional external information that the model is not aware of.

4.6.1.4 Benchmarking of AMTE and PCS calls system

The PCS Calls system currently performs their own predictions on vessels with Estimated, Authorized and Operational status. Hence, to appraise AMTE against the existing predictive system, a set of comparisons was carried out. In the following figure it can be analysed the deviation from reality in hours of both predictive systems for each vessel status where each circle represents a vessel. The closer to 0 hours, the better the prediction was, and the median of each system is highlighted in a diamond shape of darker colour.

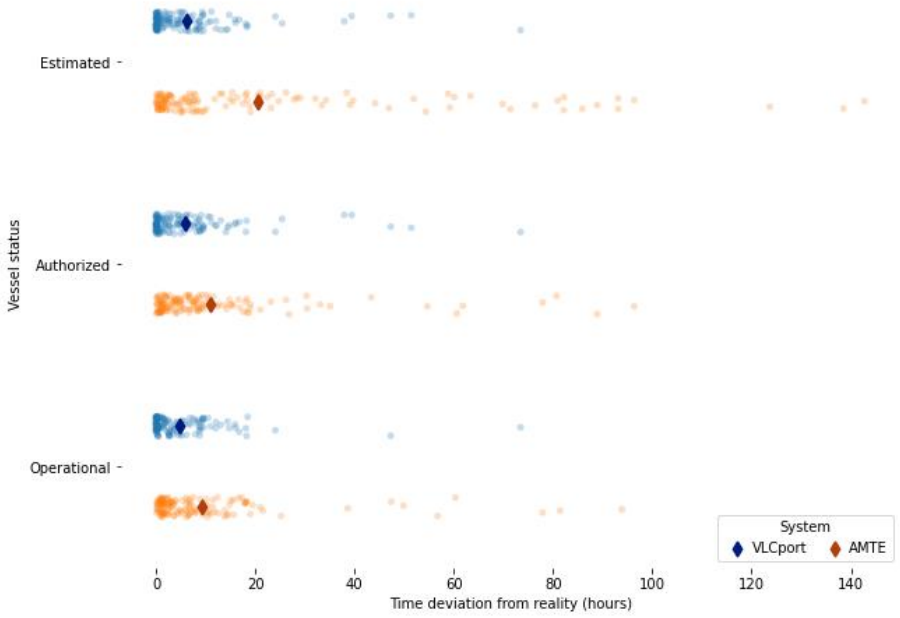

Figure 24 – Deviations of predictions from reality in hours, of AMTE and PCS calls system

Based on the results presented in Figure 24, the medians of the predictions performed by both predictive systems on all status are quite close to each other and near to zero, meaning that the vast majority of predictions are close to the reality as the concentration of circles is higher in that left end.

Yet, precise predictive system as AMTE seems, it shows some specific predictions that are far away from the reality compared to the predictions registered in PCS calls system. That behaviour can be explained by a lack of additional information that can be liable for AMTE to determine the ETD with more precision compared to the information PCS may handle. As far as it was informed, the predictions showed by the PCS calls system are obtained by the shipping companies themselves and not calculated by the Port Authority of Valencia. Therefore, the nature of the calculations is unknown, so the reasons of such differences cannot be fully explained.

Finally, it is worth considering that the pandemic had a profound impact on the normal operation of trading and ports functioning all around the globe, thereby machine learning models understanding of the data can be utterly influenced by this fact, eventually performing biased predictions.

4.6.2 Benchmark-based Evaluation of Process-based Analytics Component

Given the availability of extensive public (open access) real-world benchmark data sets to evaluate process-based analytics solutions, we decided to evaluate the Process-Based Analytics Component using these benchmark data sets to facilitate the comparability with the state of the art and foster replication of the DataPorts outcomes.

The Process-Based Analytics Component (developed by UDE) uses data available in the form of event logs (i.e., sequences of business process actions [3]) to deliver predictive, prescriptive, and explainable process monitoring capabilities. In this regards it is complementary to the Automatic Model Training Engine (developed by ITI), which uses data in the form of time series.

An important difference event log data and time series data is the kind of data used for the AI/ML components. While time series data represent values of single or multiple variables at different points in time (such as number of containers or trucks), event log data represent sequences of process events. In addition to a timestamp, a process event typically includes an event label (uniquely identifying the process step) and additional attributes of the event. As such, there is an important semantic difference between a data point in a time series and a process event [4]. Figure 25 illustrates the difference between these two main approaches towards analytics pursued in DataPorts.

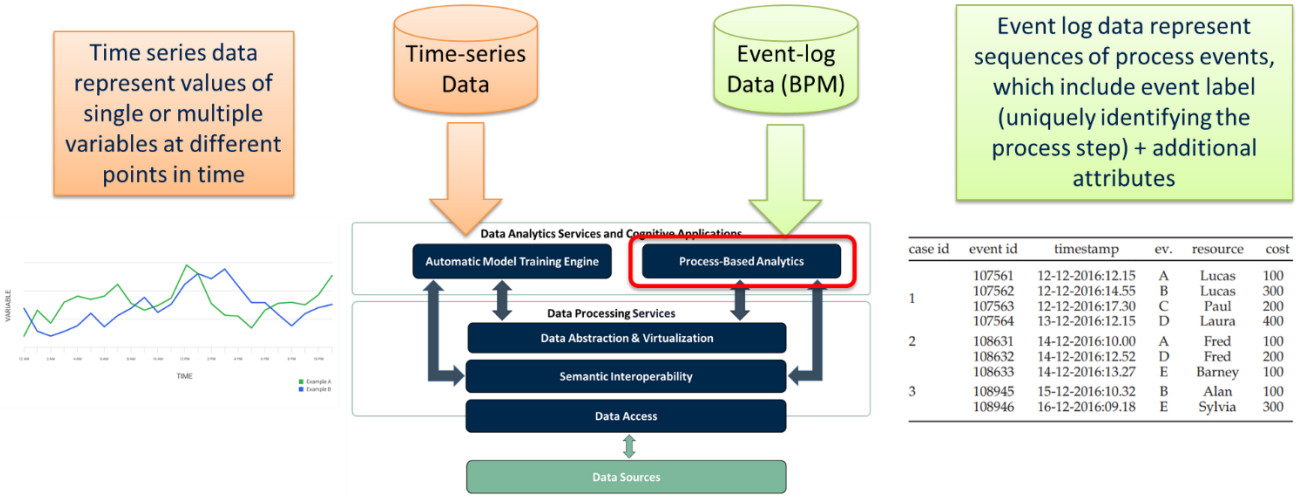

Figure 25 - Difference between time series data and event log data

The following two capabilities of the Process-Based Analytics Component provide the key innovations for process analytics as developed in DataPorts:

- **Prescriptive Process Monitoring / Analytics** [5]**:** Prescriptive business process monitoring provides decision support to process managers on when and how to adapt an ongoing business process to prevent or mitigate an undesired process outcome, and thereby constitutes an important next step from predictive analytics, which forecasts the future state or outcome of a process. The UDE solutions leverage Deep Online Reinforcement Learning to deliver prescriptive analytics capabilities, addressing as one major challenge the concept drift in the underlying data.
- **Explainable Predictive Process Monitoring / Analytics** [6]**:** Predictive business process monitoring increasingly leverages sophisticated prediction models. Although sophisticated models achieve consistently higher prediction accuracy than simple models, one major drawback is their lack of interpretability, which limits their adoption in practice. *Explainable* predictive business process monitoring aims to increase the interpretability of prediction models. The UDE solutions deliver counterfactual explanations, which provide insights into why a specific prediction was made instead of another possible prediction, addressing humans' preference for counterfactual explanations.

Table 11 depicts which kind of benchmark data sets were used for which of the capabilities of the Process-based Analytics component, which is then followed by the actual evaluation results.

Name	Type of Deviation	Rate of Deviations	Size of "Test" Data	Process Variants	Max. length	Used for Evaluating: Prescriptive Analytics	Used for Evaluating: Explainable Predictive Analytics
BPIC12 [7]	Unsuccessful loan application	25%	4,361	3,587	48	✓	
BPIC17 [8]	Unsuccessful loan application	41%	10,500	2,087	71	✓	✓
Traffic [9]	Unpaid traffic fines	58%	50,117	185	5	✓	
Cargo [10]	Delay in cargo delivery	31%	1,313	144	21	✓	

Table 11 - Data sets used in evaluation of process-based analytics

As can be seen, these data sets exhibit different characteristics and cover different application domains: finance (BPIC12 and BPIC17), government (Traffic), and transport (Cargo). This facilitates generalizability of the evaluation results.

Evaluation Results for Prescriptive Process Monitoring / Analytics

We evaluated the cost savings made possible by using prescriptive analytics by comparing it with the state-of-the-art baseline. Results are given in Table 12, indicating around 12% of cost savings on average.

a_{min} = 0.5	k = .1			k = .2			k = .3			k = .4			k = .5			
	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	***Avg***
Traffic (excerpt)	27,7	25,2	*9,1%*	37,5	41,4	*-10,5%*	47,2	57,6	*-22,0%*	56,9	73,8	*-29,6%*	66,7	89,9	*-34,9%*	***-17,6%***
BPIC 2012	17,8	17,8	*0,0%*	20,3	21,1	*-3,6%*	22,8	23,6	*-3,3%*	25,3	25,9	*-2,1%*	27,8	28,1	*-1,1%*	***-2,0%***
BPIC 2017	8,2	8,8	*-7,8%*	14,8	15,2	*-2,5%*	21,5	21,5	*0,2%*	28,2	27,6	*1,9%*	34,8	33,7	*3,3%*	***-1,0%***
Average			***0,4%***			***-5,5%***			***-8,4%***			***-9,9%***			***-10,9%***	***-6,9%***
Traffic (full)	17,6	17,4	*1,2%*	23,8	28,0	*-17,6%*	30,1	38,6	*-28,5%*	36,3	49,3	*-35,7%*	42,6	59,9	*-40,8%*	***-24,3%***
Average			***-2,2%***			***-7,9%***			***-10,6%***			***-12,0%***			***-12,9%***	***-9,1%***

a_{min} = 0.0	k = .1			k = .2			k = .3			k = .4			k = .5			
	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	RL	Thr	*Diff*	***Avg***
Traffic (excerpt)	28,4	27,8	*1,9%*	38,1	43,9	*-15,3%*	47,8	60,0	*-25,5%*	57,5	76,1	*-32,3%*	67,2	92,2	*-37,1%*	***-21,7%***
BPIC 2012	22,3	22,6	*-1,3%*	24,6	25,3	*-2,7%*	27,0	28,0	*-3,7%*	29,4	30,6	*-4,2%*	31,7	32,7	*-3,1%*	***-3,0%***
BPIC 2017	8,3	11,0	*-32,5%*	15,0	17,7	*-18,1%*	21,6	23,4	*-8,4%*	28,2	29,1	*-3,1%*	34,9	33,1	*5,1%*	***-11,4%***
Average			***-1,8%***			***-6,0%***			***-8,1%***			***-9,2%***			***-9,5%***	***-7,1%***
Traffic (full)	17,7	19,8	*-11,7%*	23,9	30,3	*-26,7%*	30,2	40,9	*-35,4%*	36,4	51,5	*-41,2%*	42,7	62,0	*-45,3%*	***-32,1%***
Average			***-15,2%***			***-15,8%***			***-15,9%***			***-16,2%***			***-14,4%***	***-15,5%***

Table 12 - Average cost savings by prescriptive analytics for different relative costs (k and a_{min} are cost parameters); all differences are significant with a p-value < 0.001

Evaluation Results for Explainable Predictive Process Monitoring / Analytics

To evaluate the quality of explanations, we used the *fidelity* metric from explainable AI to evaluate the faithfulness of explanations. This means, we evaluate how faithful the interpretable model *c* generated to provide interpretability of the predictions and thus the explanations generated from *c* are with respect to the black-box model *b* to be explained.

Table 13 reports the evaluation results. Overall, our approach reaches a very high fidelity of 97.69% across different prediction points. Fidelity remains stable across all prediction points with a variance of less than 0.5%, suggesting that our approach can generate good explanations independent of the accuracy of the underlying black-box model.

Prediction Point	Nbr. of processes	Fidelity
5	6,283	0.9754
10	6,277	0.9735
20	5,338	0.9849
30	3,454	0.9755
40	2,074	0.9730
	Total: 23,426	**Avg.: 0.9769**

Table 13 - Fidelity of explainable analytics

4.7 THESSALONIKI APPLICATIONS

This section describes the evaluation metrics for the applications developed for the Thessaloniki Use Case.

4.7.1 Analytics applications

This section describes the applications developed using the data obtained using the DataPorts Platform's components. The main difference with respect to the application described in Section 4.6 is that the Data Analytics application is a building-block of the DataPorts platform, while the ones described inside this section rely upon information obtained by building blocks provided by DataPorts.

In this section we will present the measurements taken from the analytics applications developed for the ThPA pilot. These are the Statistics for ThPA, Queues predictions, Passenger facilitation, and Covid statistics use cases. The way these use cases correspond to the various views of the private and public ThPA analytics pages is summarised in.

Use case	Vessel visualisation	Prediction	Custom prediction	Vehicle traffic	Vehicle time at port	Mobility analytics (public)	Traffic analytics (public)
Statistics for ThPA	✓			✓	✓		
Queues predictions		✓	✓	✓			
Passenger facilitation						✓	✓
Covid statistics						✓	

Table 14 - Correspondence of analytics use cases to frontend views

The following table lists the measured metrics. The metrics were measured server-side, i.e., connected to the up and running server.

	Vessel visuali-sation	Predic-tion	Custom predic-tion	vehicle traffic	vehicle traffic at port	real time vessel data	mobility monitor-ing	mobility statistics	traffic statistics	travel times	Park-ing spots
CPU max usage	2.5%	2.1%	1.5%	1.9%	0.7%	1.0%	32.3%	34.6%	2.0%	21.8%	5.7%
Max Memor y	138mb	138mb	137mb	141mb	137mb	128mb	3050mb	2980mb	1270mb	1406mb	1437mb

Avg page load time	304ms	<100ms	<100ms	242ms	11.5s	<100ms	23.33s	28.63s	700ms	1.3s	810ms

Table 15 - ThPA analytics metrics

Page loading times were taken using Chrome browser's devtools. Page loading time includes any iFrames the page might contain. On the other hand, CPU usage and memory were extracted using Ubuntu's System Monitor. We tried to "pressure" each tab with interactions, reloading and use some functionalities, so to capture max CPU usage.

To get a feeling of how much the amount of available data affects our algorithms, we have measured the CPU max usage, memory usage and page loading time for the publicly available mobility analytics pages under four different amounts of available mobility data: 100MB, 500MB, 1GB, 1.5GB and 2GB. We have plotted the results and came up with the line charts shown in Figure 26, Figure 27 and Figure 28.

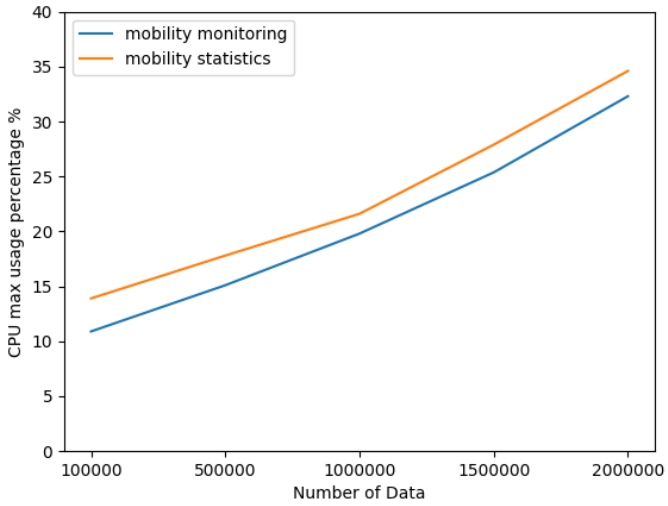

Figure 26 - Mobility analytics "CPU usage vs number of available data"

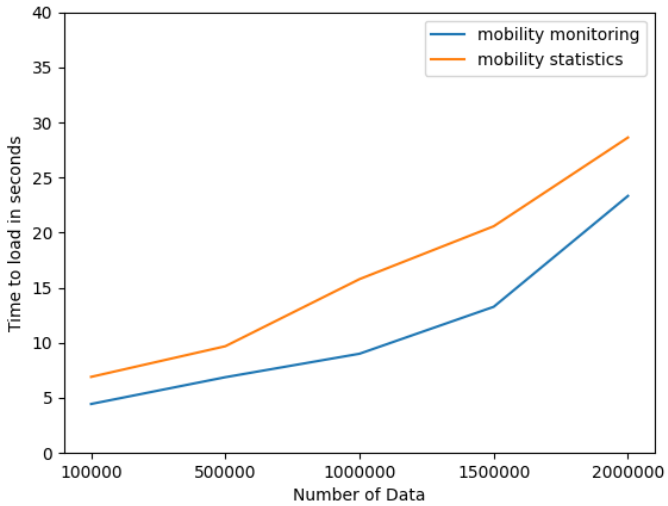

Figure 27 - Mobility analytics "page loading time vs number of available data"

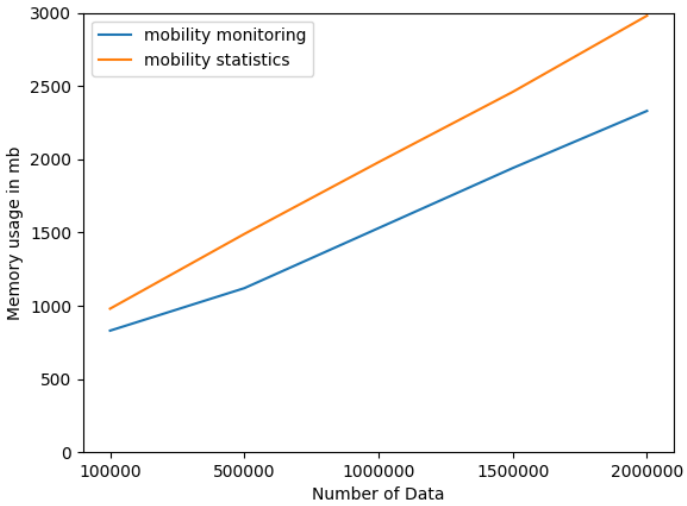

Figure 28 - Mobility analytics "memory usage vs number of available data"

By looking at the above charts, we can see that by increasing the number of available data our algorithm needs exponentially more execution time to run, which is to be expected given the nature of the algorithm. However, memory usage increases linearly, which is a positive outcome, since it means that there is no uncontrolled consumption of memory.

We have also calculated the execution time for our three most important analytics endpoints. These were calculated by taking the response time of each endpoint from Postman and then corroborating that with measurements taken from within the python code. Code measurements were taken using simple code instrumentation. The following simple construct was added to each endpoint:

```
start_time = time.time()
……   # endpoint code
end_time = time.time()
execution_time = end_time - start_time
```

The execution times calculated with both methods were almost identical. Their average is shown below (Table 16). The costliest endpoint by far is `/predict` for ship predictions, with an average execution of 6 seconds. This is due to the large amount of data available for this prediction, in contrast to traffic data, which are limited to just three months' worth of data.

	Ship predictions/date submit	Ship predictions/predict	vehicle traffic/predict
Execution time average	100ms	6000ms	100ms

Table 16 - Analytics endpoints average execution time

4.7.2 Container Pick-Up application

For the Container Pick-Up (CPU) application we utilised the aforementioned autocannon benchmarking tool to measure a characteristic endpoint of the CPU backend, the `/prelogin` endpoint, which retrieves from the blockchain and returns the list of all organisations present. This endpoint is called by the frontend when loading the login page. Autocannon was run with four different numbers of concurrent connections: 10, 20, 50 and 100. Each time the number of seconds before timing out and resetting a connection was set to 50 seconds. An example command is the following:

```
autocannon -c 100 -t 50 https://dataports.thpa.gr/backend/prelogin
```

The output of the above command can be seen in Figure 29. Taking the average from all four executions and plotting them onto a chart can be seen in Figure 30. Looking at the chart we can see that the latency experienced when calling the endpoint increases with the increase in concurrent users. However, the increase is not exponential and can be coped with by the system. Furthermore, we do not expect more than 100 concurrent users to be a realistic number for our application, where, normally, we would expect 1-2 users from each participating organisation/company to be logged in. That means that even with 50 organisations, we can expect logical latencies.

```
Running 10s test @ https://dataports.thpa.gr/backend/prelogin
100 connections


┌─────────┬──────┬──────┬──────┬──────┬──────┬──────┬──────┐
│ Stat    │ 2.5% │ 50%  │ 97.5%│ 99%  │ Avg  │ Stdev│ Max  │
├─────────┼──────┼──────┼──────┼──────┼──────┼──────┼──────┤
│ Latency │667 ms│1437 ms│7172 ms│10123 ms│1814.17 ms│1561.39 ms│10123 ms│
└─────────┴──────┴──────┴──────┴──────┴──────┴──────┴──────┘
┌─────────┬──────┬──────┬──────┬──────┬──────┬──────┬──────┐
│ Stat    │ 1%   │ 2.5% │ 50%  │ 97.5%│ Avg  │ Stdev│ Min  │
├─────────┼──────┼──────┼──────┼──────┼──────┼──────┼──────┤
│ Req/Sec │ 0    │ 0    │ 9    │ 15   │ 8    │ 5.18 │ 3    │
├─────────┼──────┼──────┼──────┼──────┼──────┼──────┼──────┤
│ Bytes/Sec│ 0 B  │ 0 B  │10.9 kB│18.1 kB│9.67 kB│6.26 kB│3.63 kB│
└─────────┴──────┴──────┴──────┴──────┴──────┴──────┴──────┘

Req/Bytes counts sampled once per second.
# of samples: 10

180 requests in 10.17s, 96.7 kB read
```

Figure 29 - Running autocannon on the `/prelogin` endpoint with 100 concurrent connections

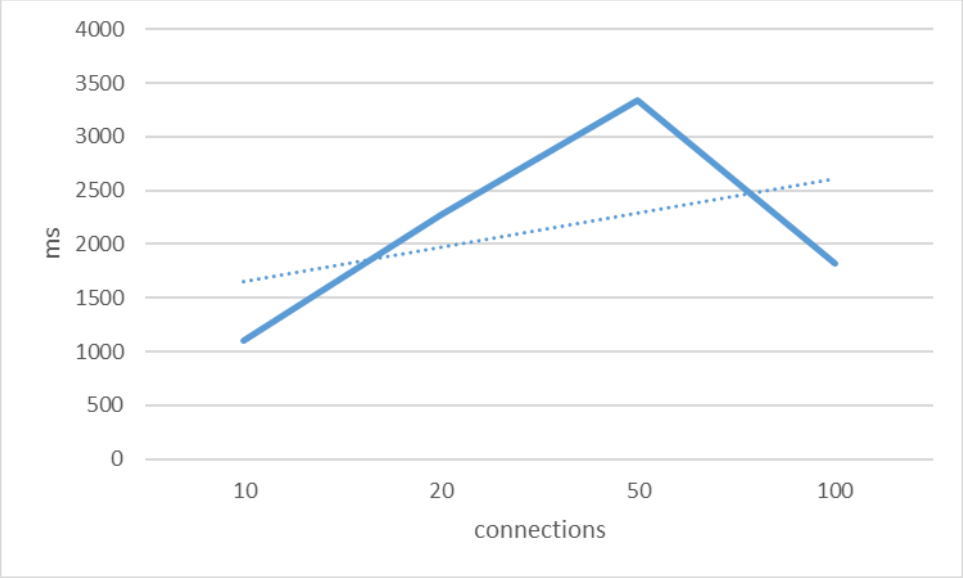

Figure 30 - Average request latency vs number of concurrent connections

The performance of Hyperledger Fabric deployment at ThPA was monitored by utilising Prometheus, a tool for monitoring system metrics. Node Exporter, designed for collecting hardware and OS metrics, was used in conjunction with Prometheus. For visualization, we relied on the Grafana platform.

The tests focused on the creation of a new Hyperledger Fabric organization (a shipping agent), which entailed configuring the organization, generating appropriate certificates and bringing up docker containers for peers

and CouchDBs. The network was then updated to include the new organization, and the new organization joined the channel. Smart contracts were also deployed in the new organization. Finally, the admin was registered with the certificate authority and the necessary certificates were produced for the user. Figure 31 and Figure 32 depict the RAM and CPU usage during request processing, revealing that this particular process consumes a significant amount of machine resources, with CPU usage reaching nearly 100%.

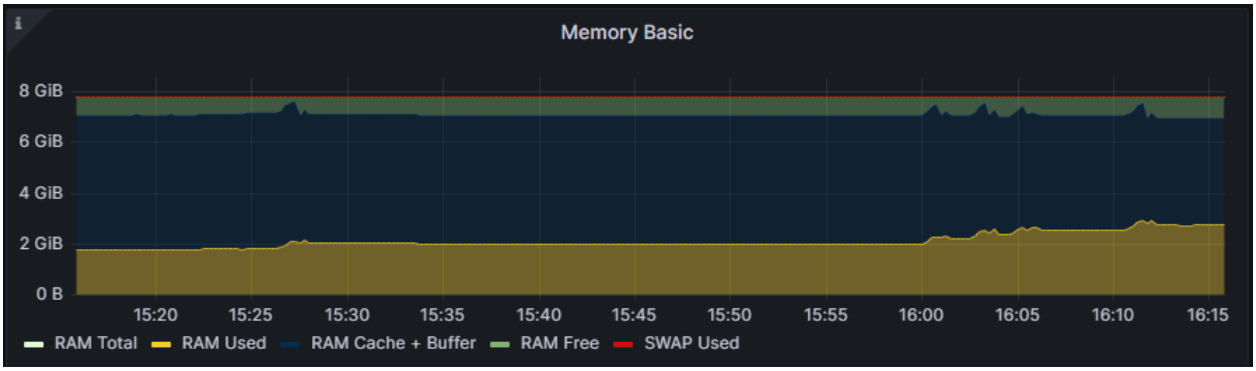

Figure 31 – Container Pick-Up memory utilisation

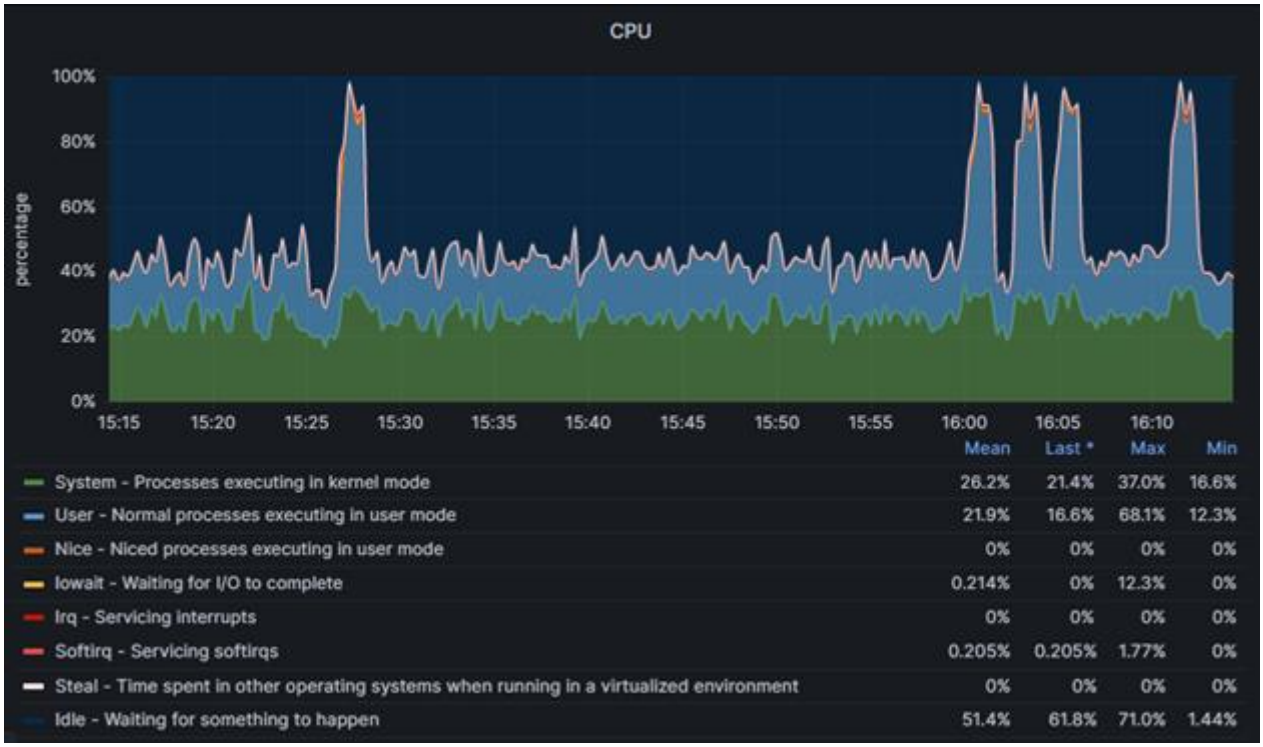

Figure 32 – Container Pick-Up CPU utilisation

The conclusion is that creating an organisation is, as expected, an intensive procedure, which, however, comes in relatively short CPU bursts, lasting about 1-1.5 minutes each, having to do with the different stages of configuring the organization, generating certificates and bringing up the various docker containers. Furthermore, memory usage is marginally affected, which is a positive outcome.

5 CONCLUSIONS

This document provides a comprehensive view of the evaluation tasks carried out by DataPorts, and the views of external end users on it.

The main finding uncovered during the evaluation is that users testing the platform did value the effort put into the project, reporting that the deployment of DataPorts would bring additional value and help them during their daily tasks.

Despite still having some issues to fix, the evaluation process uncovered a number of errors and problems with the platform, which have already been addressed, thus enabling the evaluation to be successful with respect to the methodology discussed in Deliverable D5.2.

Concerning the limitations of the evaluation process, the first thing that should be mentioned is that the number of evaluators needed to be increased. A higher number would have helped to have a complete view. Also, in some use cases, it was impossible to have the platform evaluated by its future final users, but this was mitigated by resorting to technical people not working as end users.

Finally, the multiple evaluation metrics reported in Section 4 give valuable insights regarding how the platform is deployed and how it works. It shows the pilots multiple machines' specifications and how they behave under heavy loads, as well as the actual performance under normal behaviour, using multiple tools and interfaces to show the data obtained from the evaluation metrics related to software performance. Section 4 is the finishing touch to completing the information reported in Deliverable D5.4. From here onwards, the next steps should be learning from these performance tests to improve the platform and the ability to scale the platform upwards.

6 REFERENCES AND ACRONYMS

6.1 REFERENCES

[1] F. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quarterly,* pp. 319-340, 1989.

[2] "https://sede.agenciatributaria.gob.es/Sede/estadisticas/estadisticas-comercio-exterior.html," [Online].

[3] *A. Metzger, P. Leitner, D. Ivanovic, E. Schmieders, R. Franklin, M. Carro, S. Dustdar, and K. Pohl, "Comparing and combining predictive business process monitoring techniques," IEEE Trans. on Systems Man Cybernetics: Systems, vol. 45, no. 2, pp. 276–290,* 2015.

[4] A. Metzger, A. Neubauer, P. Bohn, and K. Pohl, *"Proactive process adaptation using deep learning ensembles," in 31st Int'l Conference on Advanced Information Systems Engineering (CAiSE 2019),* Rome, Italy, June 3-7, 2019, ser. LNCS, P. Giorgini and B. Weber, Eds., vol. 11483. Springer, 2019, pp. 547–562.

[5] A. Metzger, T. Kley, and A. Palm, *"Triggering proactive business process adaptations via online reinforcement learning," in 18th Int'l Conference on Business Process Management (BPM 2020),* Sevilla, Spain (virtual), September 13-18, 2020, ser. LNCS, D. Fahland, C. Ghidini, J. Becker, and M. Dumas, Eds., vol. 12168. Springer, 2020.

[6] T. Huang, A. Metzger, and K. Pohl, *"Counterfactual Explanations for Predictive Business Process Monitoring," in 18th European Mediterranean & Middle Eastern Conference on Information Systems (EMCIS 2021),* December 8-9, 2021, Online, Springer LNBIP 437, pp. 399-413 – Best Theoretical Paper Award.

[7] "https://data.4tu.nl/articles/dataset/BPI_Challenge_2012/12689204," [Online].

[8] "https://data.4tu.nl/articles/dataset/BPI_Challenge_2017/12696884," [Online].

[9] "https://data.4tu.nl/articles/dataset/Road_Traffic_Fine_Management_Process/12683249," [Online].

[10] "https://archive.ics.uci.edu/ml/datasets/Cargo+2000+Freight+Tracking+and+Tracing," [Online].

6.2 ACRONYMS

Acronym List	
AI	Artificial Intelligence
AMTE	Automated Model Training Engine
CA	Consortium Agreement
CPU	Central Processing Unit
DAC	Data Access Component
DAV	Data Abstraction & Virtualization

DCN	Digital Consignment Note
DoA	Description of Action
ETD	Estimated Time of Departure
IAM	Identity and Access Manager
IDS	International Data Spaces
IO	Input / Output
ML	Machine Learning
OS	Operative System
PCS	Port Community System
RAM	Random Access Memory
SSH	Secure Shell
TAM	Technology Acceptance Model
UI	User Interface
VGM	Verified Gross Mass
VM	Virtual Machine
WP	Work Package

Table 17 – Acronyms

7 ANNEX A: REQUIREMENTS AND TEST CASES

This Annex Section contains the results of the test cases evaluation of each use case, as part of the evaluation process described in Section 3.1.

7.1 TEST CASES VALENCIA

Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments
TTO 1	The user logins in the platform https://vgm.dataports.com.es	user: representante@prueba.app pwd: prueba123	The user access to the platform	Pass	
TTO 2	Shipping line gets the location of all its containers https://vgm.dataports.com.es/#/pages/transportUnits	-	The container events available and authorised are shown	Pass	
TTO 3	Export company gets the location of its goods https://vgm.dataports.com.es/#/pages/shipments	-	The shipments events available and authorised are shown	Pass	
TTO 4	Freight forwarder gets the status of the transport operations which is managing https://vgm.dataports.com.es/#/pages/trackableEvents	-	All the events available and authorised are shown	Pass	
TTO 5	User can filter and order in any of the previous steps	Filter shipments by "CARGADOR 2" Filter transportUnits by "TRLU123456" Filter events by "GateIn"	Specific results	Pass	

Figure 33 - Tracking of Transport Operation

Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments
VGM 1	The user logins in the platform https://vgm.dataports.com.es	The user needs to be registered in a company The truck need to be registered user: representante@prueba.app pwd: prueba123	The user access to the platform	Pass	
VGM 2	The freight forwarded or Shipper creates the VGM request https://vgm.dataports.com.es/#/pages/vgm The user clicks "New VGM request" and complete the data	Truck: 4535HLY Trailer: R2311BCR Boobking: 53550324 Container: SEGU5769764 Shipping-line: HAPAG-LLOYD	VGM request is created	Pass	
VGM 3	The scale operator updated the VGM request with the weight measure https://vgm.dataports.com.es/#/pages/vgm The user clicks "Modify VGM request" in the proper request	Serch for previous VGM request and add "Scale" value: 32000	VGM request is updated	Pass	
VGM 4	The freight forwarded or Shipper can access and review the data https://vgm.dataports.com.es/#/pages/vgm	Serch for previous VGM request	VGM details are shown	Pass	

Figure 34 - Verified Gross Margin

Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments
DCN 1	Consignment note data is received from ValenciaportPCS to DataPorts platform	-	Consignment note is registered as PDF/A	Pass	
DCN 2	The user logins in the platform https://vgm.dataports.com.es	user: representante@prueba.app pwd: prueba123	The user access to the platform	Pass	
DCN 3	Haulier company has access to the consignment note. Click in the pdf icon https://vgm.dataports.com.es/#/pages/shipments	-	Last version of the consignment note is downloaded	Pass	

Figure 35 - Digital Consignment Note

Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments
PADS 1.1	Go to AMTE's main page: https://amte.dataports-project.eu	Login: " amte@dataports.org" Pass: "dataports"	The web site is shown without any error	Pass	
PADS 1.2	Click on "Create" tab on the left panel	-	The Training Wizard screen is shown	Pass	
PADS 1.3	Training wizard - Task. Training wizard - Task. Fill up the boxes "Name" and "Description" and click on the service called "Vessel Time of Departure Estimator"	Any strings are valid	No errors occurred and the next step 'Dataset' is shown	Pass	
PADS 1.4	Training wizard - Dataset. Click on the dataset named "Hinterland traceability of Valencia's Ports". Then click "Next".	-	No errors occurred and the next step 'Configuration' is shown	Pass	There are some providers but not the one in the description
PADS 1.5	Training wizard - Configuration. Select the port information by choosing a port, a terminal, and a regular line	-	Dropdown buttons allow the selection of items	Pass	
PADS 1.6	Training wizard - Configuration. Select the Time granularity by choosing a specific date and time on the "Generate prediction for the following time period" box. Then, click on "Next"	-	The calendar appears and it allows the selection of a date and time. No errors occurred and the next step 'Strategy' is shown	Pass	No option to select granularity
PADS 1.7	Training wizard - Strategy. Click on any of the training strategies available.	-	No errors occurred and the next step 'Confirmation" is shown	Pass	
PADS 1.8	Training wizard - Confirmation. Click on the button "Train service"	-	No errors occurred and the screen of "Services" list is showed	Pass	
PADS 1.9	Observe that the new created service appears in the list with the status "Training"	-	The new created service is added to the list and shows the status "Training"	Pass	

Figure 36 - Port Authority Data Sharing

7.2 TEST CASES THESSALONIKI

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
1.1	The shipping agent admin logs in using the login page	login data username:"shippingAgent1_admin" pass:"12345"	The shipping agent dashboard is shown	Pass		Pass		Pass		Pass		Pass	
1.2	The shipping agent admin clicks on the "COREOR Requests" menu item on the left panel		"COREOR Requests" list view is shown	Pass		Pass		Pass		Pass		Pass	
1.3	The shipping agent admin clicks on new the "COREOR Request" button		"New COREOR Request" page is shown	Pass		Pass		Pass		Pass		Pass	
1.4	The shipping agent admin submits the form with all data required (all fields with asterisk are mandatory)	form data (any data, as long as mandatory fields have a value)	"COREOR Requests" list view is shown. The new request is also submitted to the Port's infrastructure through their API and can now be also seen on the port's systems.	Pass		Pass		Pass		Pass		Pass	
1.5	The port admin logs in using the login page	login data username:"thpa_admin" pass:"12345"	The port dashboard is shown	Pass		Pass		Pass		Pass		Pass	
1.6	The port admin clicks on the "COREOR Requests" menu item on the left panel		"COREOR requests" list view is shown. COREOR Requests view contains the new request, shown as pending.	Pass		Pass		Pass		Pass		Pass	
1.7	The port admin proceeds to accept the new request through the port's system		The port's system informs us through our API of the COREOR's acceptance	Pass		Pass		Pass		Pass		Pass	
1.8	The port admin clicks on the row corresponding to the new request		The COREOR request detail view is shown and the new request appears as accepted	Pass		Pass		Pass		Pass		Pass	
1.9	The shipping agent clicks on the "COREOR Requests" menu item on the left panel		"COREOR Requests" list view is shown. New COREOR request now appears as accepted.	Pass		Pass		Pass		Pass		Pass	
1.10	The trucking company admin logs in using the login page	login data username:"truckingCompany1_admin" pass:"12345"	The trucking company dashboard is shown The admin sees the notification bell displaying a red badge, signifying the existence of a new accepted COREOR request to act upon	Pass		Pass		Pass		Pass		Pass	
1.11	The trucking company admin proceeds to create a new booking request through the port's system		The port's system sends us the new booking request through our API	Pass		Pass		Pass		Pass		Pass	
1.12	The port admin proceeds to accept the new booking request through the port's system		The port's system informs us through our API of the booking's acceptance	Pass		Pass		Pass		Pass		Pass	
1.13	The trucking company admin clicks on the "Booking Requests" menu item on the left panel		New request now appears as accepted	Pass		Pass		Pass		Pass		Pass	
1.14	The trucking company admin clicks on the row corresponding to the accepted booking request		The booking request detail view is shown	Pass		Pass		Pass		Pass		Pass	
1.15	The trucking company admin clicks on the booking details tab		The QR code is available for use by the truck driver	Pass		Pass		Pass		Pass		Pass	

Figure 37 - New COREOR Request

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
2.1	The port/shipping agent/trucking company admin logs in using the login page	login data (see Test Case 1)	The corresponding organizational dashboard is shown	Pass		Pass		Pass		Pass		Pass	
2.2	The admin clicks on "Bookings" menu item on the left panel		"Bookings" page is shown	Pass		Pass		Pass		Pass		Pass	
2.3	The admin clicks on a record in order to view the booking		Booking details are shown	Pass		Pass		Pass		Pass		Pass	
			• For a booking where the booking date has not yet passed and the trucking company's truck has not exited the port's yard, the booking is marked "pending"	Pass		Pass		Pass		Pass		Pass	
			• For a booking where the trucking company's truck has exited the port's yard, the booking is marked "completed"	Pass		Pass		Pass		Pass		Pass	
			• For a booking where the booking date has passed and the trucking company's truck has not exited the port's yard, the booking is marked "failed"	Pass		Pass		Pass		Pass		Pass	

Figure 38 - View Booking

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
3.1	The user logs in using the login page.	Analytics user test credentials username:"thpa_analyticsUser" pass:"12345"	The port's dashboard is shown.	Pass		Pass		Pass		Pass		Pass	
3.2	The user clicks on the "Analytics" menu item on the left panel.		The "Analytics" page is shown.	Pass		Pass		Pass		Pass		Pass	
3.3	The user clicks on the "Vessel visualization" tab.		The contents of the "Vessel statistics" tab is shown: - A scatterplot of vessels at the top left (largest panel in the screen), with a special icon per vessel. - An area plot at the bottom left, showing the number of vessels at the port each day. - A vessel information panel at the top right, showing details for a selected vessel. - Four bar charts at the middle right, showing the number and average weight of vessels per vessel type and cargo type. - An area plot at the bottom right, showing the distribution of the work time ratio. - A button "Select outliers" above the work time ratio plot.	Pass		Pass		Pass		Pass		Pass	
3.4	In the vessel number per day plot (bottom left), the user selects with the mouse a date range, clicking at a point and dragging to the right or left.		- The selected date range is highlighted on top of the area plot. - The date limits of the main scatterplot (horizontal axis) are limited to the selected range. - The bar plots at the middle right are updated to cover only vessels handled within the selected date range. - The work time area plot at the bottom right is updated to cover only vessels handled within the selected date range.	Pass		Pass		Pass		Pass		Pass	
3.5	In the vessel number per day plot (bottom left), the user moves the selected date range to the left or right, by clicking and dragging it with the mouse.		The views of the scatterplot, bar charts and work time plot are updated, as above.	Pass		Pass		Pass		Pass		Pass	
3.6	In the vessel number per day plot (bottom left), the user clicks outside the selected date range.		- The previously selected date range is un-highlighted. - The views of the scatterplot, bar charts and work time plot are reset to cover all available data.	Pass		Pass		Pass		Pass		Pass	
3.7	In the main scatterplot (top left plot), the user clicks on the icon corresponding to a vessel.		The vessel information panel (top right) is filled with details about the selected ship.	Pass		Pass		Pass		Pass		Pass	
3.8	In the work time ratio plot (bottom right), the user selects a horizontal range, clicking with the mouse and dragging to the right or left.		- The selected value range is highlighted on the work time ratio plot. - The scatterplot and bar chart views are updated to cover only the vessels having a work time ratio within the selected value range.	Pass		Pass		Pass		Pass		Pass	
3.9	In the work time ratio plot (bottom right), the user moves the selected value range to the left or right, by clicking and dragging it with the mouse.		The scatterplot and bar chart views are updated to cover the vessels falling within the selected range, as above.	Pass		Pass		Pass		Pass		Pass	
3.1	In the work time ratio plot (bottom right), the user clicks outside the selected value range.		- The previously selected value range is un-highlighted. - The scatterplot and bar chart view are reset to cover all the available data.	Pass		Pass		Pass		Pass		Pass	
3.11	The user clicks on the "Select outliers" button, above the work time ratio plot (bottom right).		- In the work time ratio plot (bottom right), the range from 0 to 0.2 is highlighted. - The scatterplot and bar chart views are updated to cover only the vessels with a work time ratio less than 0.2 (idle vessels – outliers).	Pass		Pass		Pass		Pass		Pass	

Figure 39 - Statistics (Vessel Visualization)

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
4.1	The user logs in using the login page.	Administrator test credentials	The port's dashboard is shown.	Pass		Pass		Pass		Pass		Pass	
4.2	The user clicks on the "Analytics" menu item on the left panel.		The "Analytics" page is shown.	Pass		Pass		Pass		Pass		Pass	
4.3	The user clicks on the "Vehicle traffic" tab.		The contents of the "Vehicle traffic" tab is shown: - A map view (top left part) showing the positions of the port gates as circular marks, with embedded numbers of the total number of trucks that entered through each gate. - A set of 7 bar charts showing the temporal patterns of truck arrivals (top right). There is one bar chart per day of week. Each bar chart shows the number of trucks entering the port through all gates at the considered time period. - A timeline plot (bottom), in which we can observe some information from end of February in 2021 till end of August in 2021. In X-axis there are the days in this timeframe and in Y-axis there is the number of trucks in port. In the background there is a light blue line, that is the bookings of trucks that took place, while the darker blue is how many trucks where actually each day in port. A brown line at the right side shows the prediction of the truck number for the next 10 days. A separate chart below shows this prediction line in more detail.	Pass		Pass		Pass		Pass		Pass	
4.4	In the map view (top left), the user clicks on one of the circular marks (gate positions).		The bar charts (top right) are updated to show statistics for only the trucks that passed through the selected gate.	Pass		Pass		Pass		Pass		Pass	
4.5	In the map view (top left), the user holds Ctrl and clicks on another of the circular marks (gate positions).		The bar charts (top right) are updated to show statistics for the trucks that passed through both of the selected gates.	Pass		Pass		Pass		Pass		Pass	
4.6	The admin clicks on a node of the plot graph, which is a specific day.	node click	If we click a day (a node), we can see a red line with the prediction that our models made regarding how many trucks will be in port the next 10 days. In the last day of August which is the 'demo' current date, due to data used on those dates, the next 10 days that are predicted are the future days.	Pass		Pass	How do we clear the node?	Pass		Pass		Pass	

Figure 40 - Statistics (Vehicle Traffic)

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
5.1	The port admin logs in using the login page	login data	The port's dashboard is shown	Pass		Pass		Pass		Pass		Pass	
5.2	The admin clicks on the "Analytics" menu item on the left panel		"Analytics" page is shown	Pass		Pass		Pass		Pass		Pass	
5.3	The admin selects the "Vehicle time at port" tab from the menu		Shows the number of vehicles parked in the port premises into groups, according to the duration of stay, i.e. x cars stayed for 5 minutes, y cars stayed for 1/2 hour etc.	Pass		Pass		Pass	Very slow to visualise	Pass		Pass	
5.4	The admin selects month, day or hour		Shows number of vehicles parked in the port premises into groups given duration month, day or hour.	Pass		Pass		Pass		Pass		Pass	
5.5	The admin selects a stayed time category bar		Under the main graph, a new graph shows each type of vehicle i.e. car, motorcycle etc... that stayed in that period.	Pass		Pass		Pass		Pass		Pass	
5.6	The admin selects "no specific month/day/time chosen" button		Resets the process and shows starting interface	Pass		Pass		Pass		Pass		Pass	

Figure 41 - Statistics (Vehicle Time At Port)

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
6.1	The port admin logs in using the login page	login data	The port's dashboard is shown	Pass		Pass		Pass		Pass		Pass	
6.2	The admin clicks on the "Analytics" menu item on the left panel		"AnalyticsVessel visualization" page is shown tab is shown	Pass		Pass		Pass		Pass		Pass	
6.3	The admin can observe three four different tabs and clicks on the "Prediction" tab	button click	A "Date" field next to a submit button is shown along with three empty tables belowVisualizations are shown, updated in real-time, according to user input	Pass		Pass		Pass		Pass		Pass	
6.4	The admin selects a Date in the "Date" field and then clicks submit	Form data and button click	The results table that contains all the ships in port in the selected date is shown, along with information like Ship Name, Ship Type, Cargo, Arrival, Mooring and Work	Pass		Pass		Pass		Pass		Pass	
6.5	The admin selects a ship from the table	button click	• Two tables are shown, with information about containers of the selected ship and another with the time predictions for three procedures in port. • Below and left from the table named "Containers" is filled and it shows how many empty and non-empty containers along with their weight, each cargo of the selected ship contains. • The table on the right side called 'Predictions for "Ship Name" ' is filled. It contains time predictions for turnaround, mooring, and work procedures of the ship in average with a min-max range, and an estimated date if the start date of the procedure is known. In addition, if a procedure in the selected date is already done, predictions will not take place.	Pass		Pass		Pass		Pass		Pass	
6.6	The admin selects a tab number from the 'Containers' table to see the rest cargos of the ship (if any)		The 'Container' table turns page with the rest information of the cargos	Pass		Pass		Pass		Pass		Pass	

Figure 42 - Queues Prediction (Vessel)

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
7.1	The port admin logs in using the login page	login data	The port's dashboard is shown	Pass		Pass		Pass		Pass		Pass	
7.2	The admin clicks on the "Analytics" menu item on the left panel		"Vessel visualization" tab is shown	Pass		Pass		Pass		Pass		Pass	
7.3	The admin can observe four different tabs and clicks on the "Custom prediction" tab	button click	A form is displayed in the left side with 11 fields to be filled an in the right side an empty table called 'Predictions'.	Pass		Pass		Pass		Pass		Pass	
7.4	The admin fills the form data in which those with an '*' are necessary to proceed (6 in total) and then clicks submit below form	Form data and button click	The predictions table is filled with time predictions for the procedures for turnaround, mooring, and work procedures of the ship in average with a min-max range, and an estimated date if the start date of the procedure is known.	Pass	User should have to fill in name and date, not type of container (40" or 20") or weight	Pass		Pass	lists is not in alphabetical order	Pass	If click "40", predicts time only for 40'?	Pass	
7.5	The admin selects 'Clear' button below form	button click	The form is cleared.	Pass	Predictions table is still filled	Pass	Predictions table is still filled	Pass		Pass		Pass	

Figure 43 - Queues Prediction (Custom)

				Examiner #1		Examiner #2		Examiner #3		Examiner #4		Examiner #5	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
8.1	A visitor (be it a passenger, a professional or simply someone interested in visiting the port) accesses the THPA web site		The THPA site is displayed	Pass		Pass		Pass		Pass		Pass	
8.2	The visitor clicks on the "Analytics" link on the web page		"Analytics" page is shown, containing dropdown options for the two analytics of mobility and traffic use cases, facilitation of passengers etc." and Covid-related statistics	Pass	Slow visualisation	Pass		Pass	Slow visualisation	Pass		Pass	
8.3	The visitor can explore the analytics results shown		Visualizations are shown, updated in real-time, according to user input	Pass		Pass		Pass		Pass		Pass	
8.4	The visitor selects "mobility" dropdown option		Two more clickable options appear, "general" and "average"	Fail	"Mobility monitoring" & "Mobility Statistics" are the clickable actions in "mobility" drop down menu	Fail	No "general" or "average" in drop down list	Fail	No "general" or "average" in drop down list	Fail	No "general" or "average" in drop down list	Fail	No "general" or "average" in drop down list
8.5	The visitor selects "Mobility monitoring" option		User can see a map interface of mobility in Thessaloniki created with historical data. Visitor can choose to see mobility with absolute or relative numbers and is also capable of seeing COVID thresholds. By changing timeline slider, in graph of active users, visitor is able to see mobility per hour of day.	Pass		Pass		Pass		Pass		Pass	
8.6	The visitor selects "Mobility statistics" option.		A new map interface of Thessaloniki shows up. Visitor can choose to see map by average numbers or by COVID thresholds colours i.e. red too much users, orange medium mobility, green safe places. Visitor can also choose month and day to see average mobility in bar plot where he can choose specific hour to see in the map average mobility for every cell individually and predict whether or not visit a specific area.	Pass		Pass		Pass		Pass		Pass	
8.7	The visitor selects "Traffic" dropdown option.		Three more clickable options appear, "Traffic statistics", "travel times" and "parking spots"	Pass		Pass		Pass		Pass		Pass	
8.8	The visitor selects "Traffic statistics" option.		A map interface of Thessaloniki shows up.	Pass		Pass		Pass		Pass		Pass	
8.9	The visitor selects month or day dropdown		The visitor can choose month and day to see average traffic in map and in bar plot of 24 hours of day.	Pass	only 23 h of day	Pass		Pass		Pass		Pass	
8.10	The visitor clicks a bar from average users Bar Plot		The visitor can see average traffic for specific hour of day into the map and predict whether or not visit an area	Pass		Pass		Pass		Pass		Pass	
8.11	The visitor selects "travel times" option.		Shows map interface of Thessaloniki's road network	Pass		Pass	slow visualisation	Pass	slow visualisation	Pass		Pass	
8.12	The visitor selects a path from dropdown option.		Shows selected path with red colour and average duration in bar plot to traverse it.	Pass		Pass	Some names in Greek; requires a lot of zooming to disdinguish colours	Pass		Pass	Some names in Greek	Pass	
8.13	The visitor selects a month or day from dropdown options.		Shows for selected path average travel time for every hour during the selected month and day	Pass	Why only 3 months?	Pass		Pass		Pass		Pass	
8.14	The admin selects "parking spots" tab from the menu		Shows average number of visitors for both parking lots (or each one individually) for each day of the week and each hour every day. Also shows actual historical data in a graph	Pass		Pass		Pass		Pass		Pass	
8.15	The admin uses the subgraph under the main graph and adjusts it by cropping or moving it.		Shows better visualization of historical data with a selected timeline from the user	Pass		Pass		Pass		Pass		Pass	

Figure 44 - Passenger Covid Analytics

7.3 TEST CASES SMART CONTAINERS

Test Case #1				Evaluator 1 - TRX Lucas Moulin (Head of Product)		Evaluator 2 - TRX David Serrurier (Product Owner)	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments
SCO 1.1	Login to Traxens-hub application Main page : https://hub.traxens.com/traxui/#/	Login and password are individual and provided for each user	The web site is shown without any error	Pass		Pass	
SCO 1.2	Search for an equipment and click on it to see the details	MRSU4958035 for example, favor a diversity of containers for tests	The details page of the container is displayed	Pass	MRKU7166548	Pass	MRSU4418593
SCO 1.3	Click on "History" tab to see IOT events	-	The IOT events of the container are displayed	Pass		Pass	
SCO 1.4	Search for and click on the last geofencing IN event	-	The corresponding icon is displayed on the map. Memorize the geofencing of the container (timestamp, position and location).	Pass	2023-02-19 04:10:32 UTC TRANSITION=IN_to_OUT, ZOI_NAME=Port Tangier Mediterranee, ZOI_CODE=MAPTMTER01	Pass	2023-03-02 00:26:02 TRANSITION=IN_to_OUT, ZOI_NAME=Strasskirchen, ZOI_CODE=DESKI
SCO 1.5	Search for and click on the last geofencing OUT event	-	The corresponding icon is displayed on the map. Memorize the geofencing of the container (timestamp, position and location).	Pass	2023-02-19 04:10:32 UTC TRANSITION=IN_to_OUT, ZOI_NAME=Port Tangier Mediterranee, ZOI_CODE=MAPTMTER01	Pass	2023-03-02 00:26:02 TRANSITION=IN_to_OUT, ZOI_NAME=Strasskirchen, ZOI_CODE=DESKI
SCO 1.6	Login to DataPorts platform VPF main page: https://vgm.dataports.com.es/	Login: representante@prueba.app Pass: prueba123	The web site is shown without any error	Pass		Pass	
SCO 1.7	Click on "Logistics" tab on the left panel	-	The menu unfolds	Pass		Pass	
SCO 1.8	Click on "Events"	-	A list of events is shown	Pass		Pass	
SCO 1.9	Identify the relevant Geofencing IN event (search for geofencing events of the previous container)	Same container than previously	Container ID / Geofencing / arrived / port location / performed time. Data are the same beetween Traxens-hub and DataPorts applications.	Fail	No geofencing IN event present at all	Fail	No geofencing / arrived
SCO 1.10	Identify the relevant Geofencing OUT event (search for geofencing events of the previous container)	Same container than previously	Container ID / Geofencing / departed / port location / performed time. Data are the same beetween Traxens-hub and DataPorts applications.	Pass	Departure / Port Tanger Med / 19/02/23 5:10 local time	Pass	Departed / Strasskirchen / 02/03/2023 01:26:00

Figure 45 - Geofencing

Test Case #2				Evaluator 1 - TRX Lucas Moulin (Head of Product)		Evaluator 2 - TRX David Serrurier (Product Owner)	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments
SCO 2.1	Login to Traxens-hub application Main page : https://hub.traxens.com/traxui/#/	Login and password are individual and provided for each user	The web site is shown without any error	Pass		Pass	
SCO 2.2	Search for an equipment and click on it to see the details	MRSU4958035 for example, favor a diversity of containers for tests	The details page of the container is displayed	Pass	MRKU7166548	Pass	MRSU4418593
SCO 2.3	Click on "History" tab to see IOT events	-	The IOT events of the container are displayed	Pass		Pass	
SCO 2.4	Search for the last door opening event	-	Memorize the door opening of the container (timestamp and door status).	Pass	OPEN 2023-01-18 07:39 PM UTC	Pass	OPEN on 2023-02-21 1:58 PM
SCO 2.5	Search for the last door closing event	-	Memorize the door opening of the container (timestamp and door status).	Pass	CLOSED 2023-01-18 07:52 PM	Pass	CLOSED on 2023-03-01 10:03 AM
SCO 2.6	Login to DataPorts platform VPF main page: https://vgm.dataports.com.es/	Login: representante@prueba.app Pass: prueba123	The web site is shown without any error	Pass		Pass	
SCO 2.7	Click on "Logistics" tab on the left panel	-	The menu unfolds	Pass		Pass	
SCO 2.8	Click on "Events"	-	A list of events is shown	Pass		Pass	
SCO 2.9	Identify the relevant door opening events (search for door opening events of the previous container)	Same container than previously	Container ID / doorOpening / doorOpened / performed time. Data are the same beetween Traxens-hub and DataPorts applications.	Pass	OPEN 2023-01-18 08:39 PM local time	Pass	21/02/23 14:58 doorOpened
SCO 2.10	Identify the relevant door closing events (search for door closing events of the previous container)	Same container than previously	Container ID / doorClosed / performed time. Data are the same beetween Traxens-hub and DataPorts applications.	Fail	No door closing event present at all	Fail	No door closed

Figure 46 - Door Opening Closing

Test Case #3				Evaluator 1 - TRX Lucas Moulin (Head of Product)		Evaluator 2 - TRX David Serrurier (Product	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments
SCO 3.1	Login to Traxens-hub application Main page : https://hub.traxens.com/traxui/#/	Login and password are individual and provided for each user	The web site is shown without any error	Pass		Pass	
SCO 3.2	Search for an equipment and click on it to see the details	MRSU4958035 for example, favor a diversity of containers for tests	The details page of the container is displayed	Pass	MRKU7166548	Pass	MRSU4418593
SCO 3.3	Click on "History" tab to see IOT events	-	The IOT events of the container are displayed	Pass		Pass	
SCO 3.4	Search for and click on the last geolocation event of the container	-	The position icon is displayed on the map. Memorize the position of the container (timestamp and position).	Pass	18.900579, 72.997017 2023-03-06 02:58:31	Pass	2023-03-02 00:26:02 with coordinates 49.000042,12.138250
SCO 3.5	Login to DataPorts platform VPF main page: https://vgm.dataports.com.es/	Login: representante@prueba.app Pass: prueba123	The web site is shown without any error	Pass		Pass	
SCO 3.6	Click on "Logistics" tab on the left panel	-	The menu unfolds	Pass		Pass	
SCO 3.7	Click on "Events"	-	A list of events is shown	Pass		Pass	
SCO 3.8	Identify the relevant container location event (search for golocation events of the previous container)	Same container than previously	Container ID / geolocation / performed time. Data are the same beetween Traxens-hub and DataPorts applications.	Pass	06/03/2023 03:58:00 at 18.900579, 72.997017	Pass	02/03/23 01:26 at 49°00'00.2"N 12°08'17.7"E

Figure 47 - Geolocation

7.4 PORT MANAGEMENT SYSTEM INTEGRATION

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 1.1	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass	If a previous access was in cache and the access token is expired, the web shows an error and does not redirect to login	Pass		Pass	
PNO 1.2	Click on "User destinations" tab on the left panel	-	An empty table is shown	Pass		Pass		Pass		Pass		Pass	
PNO 1.3	Press the "+" button on the top of the table.	-	A new form is shown.	Pass		Pass		Pass		Pass		Pass	
PNO 1.4	Create a new destination. Select the communication channel and enter the direction. Press the save button at the bottom right of the page.	Channel: "Email" Destination: "baleares_vessel_consignee@prodevelop.es"	The table is shown and contains the just created destination.	Pass		Pass		Pass		Pass		Pass	

Figure 48 - Create destination for alerts

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 2.1	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 2.2	Click on "Subscriptions" tab on the left panel	-	A table with current subscriptions is shown	Pass		Pass		Pass		Pass		Pass	
PNO 2.3	Press the "+" button on the top of the table.	-	A new form is shown.	Pass		Pass		Pass		Pass		Pass	
PNO 2.4	Create a subscription. Select the topic to subscribe, the filter, the communication channel and the destination. Press the confirmation button at the bottom right of the page.	Topic: "Change of ETA" Select Filter: "Organization" Channel: "Email" Destination: baleares_vessel_consignee@prodevelop.es	The table is shown and contains the just created subscription.	Pass		Pass		Pass	Table does not have pagination controls and there are already more elements than can be shown in one page. To see y subscription I have to sort by descending	Pass		Pass	

Figure 49 - Subscribe for change of ETA

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 3.1	Go to Posidonia Management's main page: http://docker-baleares.prodevelop.es/Posidonia	Login: "dataports" Pass: "1"	The web site is shown without any error	Pass		Pass		Pass	On first login attempt, an error page was displayed. An F5 refreshed and displayed the home page	Pass		Pass	
PNO 3.2	Click on "Port Calls" tab on the left panel. Then click on "Port Calls"	-	The screen with the list of visits is visualized	Pass		Pass		Pass		Pass		Pass	
PNO 3.3	Search for the test visit. Filter the visits using the filter "Fast search" placed at the top right of the application.	P202300001	Only one visit is shown of the list	Pass		Pass		Pass		Pass		Pass	
PNO 3.4	Select the visit and choose the action "Edit ETA" from the action list on the top left.	-	The form for editing the date is shown.	Pass		Pass		Pass		Pass		Pass	
PNO 3.5	Set the new date and confirm.	Any future date or time that is before the ETD of the visit.	The window is closed and you get back to the visit list, that has the updated ETA.	Pass		Pass		Pass		Pass		Pass	
PNO 3.6	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 3.7	Click on "Notifications" tab on the left panel	-	A table with a notification for the topic "ETA changed" is shown.	Pass		Pass		Pass		Pass		Pass	

Figure 50 - Perform an ETA change

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 4.1	Go to the email incoming box: https://outlook.office.com	Login: "baleares_vessel_consignee@prodevelop.es" Pass: "E2qqViBnVM08Ed3YWAJ0"	The mailing box is shown and there is a mail alerting of an ETA change.	Pass		Pass		Pass		Pass		Pass	
PNO 4.2	Open the Posidonia Notifications received mail.	-	An email about the ETA change of a visit has been received.	Pass		Pass		Pass		Pass		Pass	

Figure 51 - Check alerts

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 5.1	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 5.2	Click on "Subscriptions" tab on the left panel	-	A table with current subscriptions is shown	Pass		Pass		Pass		Pass		Pass	
PNO 5.3	Press the "+" button on the top of the table.	-	A new form is shown.	Pass		Pass		Pass		Pass		Pass	
PNO 5.4	Create a subscription. Select the topic to subscribe, the filter, the communication channel and the destination. Press the confirmation button at the bottom right of the page.	Topic: "Water Supply Requested" Select Filter: "Stakeholder" Channel: "Email" Destination: "baleares_vessel_consignee@prodevelop.es"	The table is shown and contains the just created subscription.	Pass		Pass		Pass	Same consideration as PNO 2.4	Pass		Pass	

Figure 52 - Subscribe for water requests

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 6.1	Go to Posidonia PCS' main page: "http://docker-baleares.prodevelop.es/pcs"	Login: "consigbuqueall" Pass: "1234"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 6.2	Click on "Professional Services" tab on the left panel. Then click on "Water and Energy Supply"	-	The screen with the list of services is visualized	Pass		Pass		Pass		Pass		Pass	
PNO 6.3	Press the "+" button at the top of the table.	-	A form for requesting supplies is opened	Pass		Pass		Pass		Pass		Pass	
PNO 6.4	Fill the form and press the save (a disk image) button at the top right of the application.	Select a Port Call from the list through the "Port" and "Port of call" drop lists. Resquest of: "Water" Quantity: "100"	The form is closed and you get back to the list with the new request created.	Pass		Pass		Pass		Pass		Pass	
PNO 6.5	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notificati	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 6.6	Click on "Notifications" tab on the left panel	-	A table with a new notification for the topic "Water Supply Requested" is shown.	Pass		Pass		Pass		Pass		Pass	3x messages received. Same hour/content. This is expected because there are 3 active subscriptions

Figure 53 - Perform a water supply request

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 7.1	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 7.2	Click on "User destinations" tab on the left panel	-	A table with a destination for the email baleares_vessel_consignee@prodevelop.es is shown	Pass		Pass		Pass		Pass		Pass	
PNO 7.3	Delete a destination. Select the destination to delete and press the (red trash icon) delete button on the top of the table.	-	The destination disappears from the list	Pass		Pass		Pass		Pass		Pass	
PNO 7.4	Click on "Subscriptions" tab on the left panel	-	A table with a subscription to ETA change and a subscription for Water Supply Resquest are shown	Pass		Pass		Pass		Pass		Pass	
PNO 7.5	Disable the subscriptions. Select the subscriptions to disable and select the Disable action from the top left dropdown list. Accept the action.	-	The subscriptions are modified and are disabled.	Pass		Pass		Pass	The apply button is rendered almost out of the window	Pass		Pass	

Figure 54 - Disable subscription

				Evaluator 1 - PRO		Evaluator 2 - PRO		Evaluator 3 - PRO		Evaluator 4 - PRO		Evaluator 5 - PRO	
Test Step ID	Test Step Description	Test Data	Expected Results	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments	Pass / Fail	Comments
PNO 3.1	Go to Posidonia Management's main page: http://docker-	Login: "dataports" Pass: "1"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 3.2	Click on "Port Calls" tab on the left panel. Then click on "Port Calls"	-	The screen with the list of visits is visualized	Pass		Pass		Pass		Pass		Pass	
PNO 3.3	Search for the test visit. Filter the visits using the filter "Fast search" placed at the top right of the application.	P202300001	Only one visit is shown of the list	Pass		Pass		Pass		Pass		Pass	
PNO 3.4	Select the visit and choose the action "Edit ETA" from the action list on the top left.	-	The form for editing the date is shown.	Pass		Pass		Pass		Pass		Pass	
PNO 3.5	Set the new date and confirm.	Any future date or time that is before the ETD of the visit.	The window is closed and you get back to the visit list, that has the updated ETA.	Pass		Pass		Pass		Pass		Pass	
PNO 3.6	Go to Posidonia Notifications' main page: http://svldock04.prodevelop.es:11900/notifications	Login: "baleares_vessel_consignee" Pass: "dataports2022"	The web site is shown without any error	Pass		Pass		Pass		Pass		Pass	
PNO 3.7	Click on "Notifications" tab on the left panel	-	No new notification has been created for the event since there is not an active	Pass		Pass		Pass		Pass		Pass	

Figure 55 - Perform an ETA change (II)

8 ANNEX B: USE CASE QUESTIONNAIRES

This annex contains the four questionnaires used for evaluating the use cases, as described in Section 3.2.

Table 18, Table 19, Table 20, and Table 21 show the results of the use case survey. Each table presents the result for one of the four questionnaires. The questions are grouped by indicator (usefulness, ease of use, behavioural intention), the columns “E#1” through “E#5” report the score assigned by each of the evaluators involved in the survey, and finally the last column (“Score”) reports the average score for each question.

Ind.	Question	E#1	E#2	E#3	E#4	E#5	Score
Perceived Usefulness	I think using the platform would help me perform my daily activities.	4	4	4	4		4
	I think the platform is better than the one offered by the current solution.	3	4	4	4		3.75
	Using the platform would allow me to work with a higher degree of security with respect to the current solution.	3	3	3	4		3.25
	Using the platform would allow me to access a higher volume of data with respect to the current solution.	4	4	4	4		4
	Using the platform would allow me to have access to more data sources with respect to the current solution.	4	4	4	4		4
	Using the platform would allow me to access more services with respect to the current solution.	4	4	4	3		3.75
	Using the platform would allow me to better manage data and information with respect to the current solution.	3	4	3	3		3.25
	Using the platform would allow me to share a larger amount of data with other partners with respect to the current solution.	4	3	3	4		3.5
Perceived Ease of Use	The interface of the platform is clear and intuitive.	3	3	3	3		3
	I found it easy to perform the required tasks using the platform.	4	3	3	4		3.5
	It is easy to find the information needed in the platform.	2	3	3	3		2.75
	The data and analysis results presented by the platform are easy to understand.	4	4	3	4		3.75
	It would be easy for me to become skilful at using the platform.	3	3	4	4		3.5
Behavioural	I like the idea of using the platform.	4	4	4	4		4
	I believe it is a good idea to use the data platform to manage operations.	4	4	4	4		4
	I believe it is a good idea to use the data analytics of the platform to manage and better planning operations.	4	4	4	4		4

Table 18 - Questionnaire Valencia

Ind.	Question	E#1	E#2	E#3	E#4	E#5	Score
Perceived Usefulness	I think using the platform would help me perform my daily activities.	4	3	3	4	3	3.4
	I think the platform is better than the one offered by the current solution.	4	4	4	4	4	4
	Using the platform would allow me to work with a higher degree of security with respect to the current solution.	3	3	3	3	3	3
	Using the platform would allow me to access a higher volume of data with respect to the current solution.	4	3	4	4	3	3.6
	Using the platform would allow me to have access to more data sources with respect to the current solution.	4	3	4	3	3	3.4
	Using the platform would allow me to access more services with respect to the current solution.	4	4	4	4	4	4
	Using the platform would allow me to better manage data and information with respect to the current solution.	4	4	4	4	4	4
	Using the platform would allow me to share a larger amount of data with other partners with respect to the current solution.	4	4	4	3	3	3.6
Perceived Ease of Use	The interface of the platform is clear and intuitive.	4	4	4	4	3	3.8
	I found it easy to perform the required tasks using the platform.	4	4	4	3	3	3.6
	It is easy to find the information needed in the platform.	4	4	4	4	3	3.8
	The data and analysis results presented by the platform are easy to understand.	4	4	4	3	3	3.6
	It would be easy for me to become skilful at using the platform.	4	4	4	4	4	4
Behavioural	I like the idea of using the platform.	4	4	3	4	3	3.6
	I believe it is a good idea to use the data platform of the platform to manage operations.	3	3	3	3	2	2.8
	I believe it is a good idea to use the data analytics of the platform to manage and better planning operations.	4	4	4	4	4	4

Table 19 - Questionnaire ThPA

Ind.	Question	E#1	E#2	E#3	E#4	E#5	Score
Perceived Usefulness	I think having geofencing events from smart containers will improve my operations with respect to existing data.	3	3	4	3	3	3
	I think having door opening events from smart containers will improve my operations with respect to existing data.	4	4	4	4	4	4
	I think having geolocation events from smart containers will improve my operations with respect to existing data.	3	4	4	4	3	4
	I think having smart container events in DataPorts is helpful compared to using Traxens platform in parallel of DataPorts.	3	2	2	2	3	2
	I think smart container data will help provide better services to my customers.	4	4	4	4	4	4
Perceived Ease of Use	The interface of the platform is clear and intuitive.	3	3	2	3	3	3
	I found it easy to perform the required tasks using the platform.	4	2	2	2	4	2
	It is easy to find the information needed in the platform.	3	2	2	2	3	2
	The data and analysis results presented by the platform are easy to understand.	3	3	3	3	3	3
	It would be easy for me to become skillful at using the platform.	3	3	3	3	3	3
Behavioural	I like the idea of using the platform.	4	4	4	3	4	4
	I believe it is a good idea to use the data of the platform to manage operations.	4	4	4	4	4	4
	I believe it is a good idea to use the data analytics of the platform to manage and better planning operations.	3	4	4	4	3	4

Table 20 - Questionnaire Smart Containers

Ind.	Question	E#1	E#2	E#3	E#4	E#5	Score
Perceived Usefulness	I think using the platform would help me perform my daily activities.	3	3	3	2	2	2.6
	I think the platform is better than the one offered by the current solution.	4	4	3	3	2	3.2
	Using the platform would allow me to improve the response time dealing with requests with respect to the current solution.	4	4	3	3	3	3.4
	Using the platform would allow me to be more aware of the port activities with respect to the current solution.	4	4	2	3	3	3.2
	Using the platform would allow me to better manage data and information with respect to the current solution.	4	3	2	3	3	3
	Using the platform would allow me to share a larger amount of data with other partners with respect to the current solution.	3	3	4	3	3	3.2
Perceived Ease of Use	The interface of the platform is clear and intuitive.	4	4	3	3	3	3.4
	I found it easy to perform the required tasks using the platform.	3	3	3	4	3	3.2
	It is easy to find the information needed in the platform.	4	3	2	4	3	3.2
	The data presented by the platform is easy to understand.	3	4	2	4	4	3.4
	It would be easy for me to become skilful at using the platform.	4	4	3	4	3	3.6
Behavioural Intention	I like the idea of using the platform.	4	4	3	4	3	3.6
	I believe it is a good idea to use the data platform for the data integration of Posidonia products.	4	4	3	4	3	3.6
	I believe that using the platform will help me to be more efficient in the requests and services management.	4	4	4	4	3	3.8
	I believe that using the platform will help me to be more aware of the current status of the port activities.	4	4	3	4	3	3.6

Table 21 - Questionnaire Port Management