CloudOpting

Migration of Open Public Services to the Cloud and Provision of Cloud Infrastructures

Deliverable D4.3 Evaluation report and replication plan

Due delivery date: 28/02/2017
Lead beneficiary: Worldline

Reference	
Authors	Xavier Cases Camats, Gonzalo Cabezas, Genís Mauri, Miquel Bonet, Gustavo Macias, Peter Nilsson, Eleanor Barley, Ciprian Pavel, Miquel Bonet, Saverio Ghiotti, Luca Gioppo, Lisa Cosmi, Saverino Reale, Sabine Delaitre, Jorge Cerpa, Junjo López.
Partner	Worldline
Work Package	WP4
Tasks	T4.3
Status-version	1.0
Dissemination	Public

Version List

Version	Date	Reason	Author	Review
0.1	20/01/2017	Table of Contents	Worldline	All Partners
0.2	03/02/2017	Contributions	WL, IMI, SDG, Netport	Worldline
0.3	07/02/2017	Contributions	Electric Corby	Worldline
0.4	17/02/2017	Contributions	Teamnet, SDG	Worldline
0.5	24/02/2017	Contributions	WL, Teamnet, Netport	Worldline
0.6	27/02/2017	Contributions	IMI, WL, CSI	Worldline
0.7	06/03/2017	Contributions	Corby, Teamnet, CSI, Netport, Wellness, WL	All Partners
0.8	08/03/2017	Contributions	SmartPartners, Netport, Corby, Wellness, Worldline	Worldline
0.9	10/03/2017	Contributions	Wellness	Netport, IMI
0.10	14/03/2017	Ammended after comments	Worldline	IMI
1.0	31/03/2017	Submission	IMI	Juanjo López Fumanal

Acronyms & Abbreviations

Concept	Definition
DoW	Description of Work of the project (Annex I of the Grant agreement no. 621146)
EC	European Commission
FP7	The Seventh Framework Programme for Research and Technological Development (2007-2013)
KPI	Key performance indicator
QoS	Quality of Service
VM	Virtual Machine
UI	User interface
MTR	Master Template Report
MTP	Master Test Plan
AR	Anomaly Report
IG	Installation Guide
RS	Release Notes
EC	European Commission

Table of Contents

Executive Summary

This deliverable is the third deliverable within Work Package 4 (WP4) and it corresponds to Task 4.3: Reproduction and validation plans. This task have been concerned with the evaluation of the services migrated and with the replication plans. An evaluation model based on criteria and metrics derived from the scenarios and requirements specified in Task 4.1 have been created. Subsequently, it has been used to assess the migration of services in the pilot deployed in Task 4.2. The replication plan has been used to reproduce these services migration in other locations.

During year 3, the four pilots have been deployed in four different cities; Barcelona (Spain), Torino (Italy), Corby (UK) and Karlshamn (Sweden). Each of these pilot cities comprises a service, or a set of services, that have been deployed sequentially and operated for 12 months during the monitoring period. This document will present both the evaluation of the platform and the pilots together with its services. It is focused on the platform and the four pilot cities represented by its most representative services.

This deliverable also includes the evaluation questionnaire, the tests performed to the platform and a list of improvements in the backlog of the platform.

The deployment and operation of the platform and the pilots can be found in deliverable *4.2 Pilots Deployment Report* within the scope of this work package.

Structure of the document

The document is divided in three main sections; the first one describes the evaluation overview about the methodology applied, the second will show and explain in detail the evaulation of each services and the platform depending on its KPIs, and the last will explain how to replicate a service and the platform with a successful use case.

1. Evaluation overview: In this section is described the evaluation methodology that has been used including the amendments completed during the evaluation period. Some of the points of the methodology have been changed in order to be adapted to the reality and to be more accurate and useful.
2. Evaluation of the platform and the pilots: This section contains the evaluation of the platform and the pilots by following the methodology described in the previous point, taking into account all KPIs defined.
3. The replication plan: In this section, a successful service replication use case will be explained in order to explain the methodology on how the platform and services can be replicated from machine to machine or cloud to cloud.

1. Introduction

To achieve the objectives of the CloudOpting project four pilots were designed and deployed in four different cities. According to the DoW, each of the pilots consist of a location (Piemonte, Corby, Barcelona, Karlshanm) where a set of services have been maintained for a period of 12-months.

The main objective of this document focuses on the evaluation of the platform and the pilots and selected services for each city after this period.

The platform and services implemented in the pilots have been monitored for 12 months. This document presents a detailed evaluation of the four more representative services of each pilot and a more generic evaluation for the other services. The four services selected for a more accurate evaluation (one per pilot city) were Clearò in Piemonte region, CKAN deployed in Barcelona, Bus Portal from the Corby pilot and '*Where can I charge my car?*" which belongs to the Karlshamn pilot.

The CloudOpting platform has been assessed through the operation of the services. For this purpose, particular metrics have been considered in deliverable D4.1 in order to measure the system as a whole, taking into account also all the services and their performance. These metrics and associated KPIs have been analysed to understand the impact of the CloudOpting platform and its operation.

Depending on the objectives of the pilot, particular metrics and KPIs, and threshold to compare with these KPIs, will be highlighted in order to evaluate the success of the service during and after the operational time of the project.

2. Evaluation Methodology

In this section the evaluation methodology of the CloudOpting platform and services is addressed.

It describes the most important aspects of the methodology described in deliverable D4.1 and the changes that have been made in order to better understanding of the platform and services.

The methodology explained in deliverable D4.1 does not takes into account the peculiarities of each service, so an update on the metrics may be applied in some services.

2.1. Objectives

The objective of the evaluation is to determine the value and quality of the platform and services.

The outcomes of the evaluation will assess:

- Technical reliability of the platform and services
- Usability of the platform from the final users point of view
- Business interest in the platform and the services

2.2. Methodology

The methodology used for the evaluation is based in the collection of metrics and information from the activities performed during this year and its subsequent evaluation. This metrics have been used to asses the key performance indicators of both the services and the platform. Assessment techniques provide the mechanism for measuring and evaluating the defined factors and to evaluate progress and impact.

The selection of metrics has been based in a mix of quantitative and qualitative methods. This compilation includes data extracted from different monitoring tools and data gathered from end user's feedback.

Quantitative methods

Each service used different KPI to measure success based on specific business goals and targets. Selecting the right KPIs depends on the service and which part of it is needed to be tracked. Each service presents its quantitative KPIs in the evaluation section of this document.

The quantitative metrics have been extracted from different monitoring tools. These tools include both the ones integrated in the platform and the ones used by some services (Google Analytics, other in-service tools, etc.).

Some of the quantitative KPIs that have been selected for the services are: CPU usage, memory usage, stability, availability, number of visits, datasets uploaded, etc. Notice that this KPIs may not apply to all the services, so each service has its own specific KPIs.

The following table lists some of the KPIs that have been used for the assessment of the services and the platform, its description and the associated metrics.

KPI	KPI Description	Associated Metrics
CPU usage	Average % CPU usage	CPU %
Memory usage	Averaged % Memory usage	Memory %
Stability	How stable is the service	Number of downs
Availability	Uptime of the service	Number of downs Service uptime
Access Time	How much time need the service to attend a request	Access Time
Service visits	Number of users using the service	Visits sessions
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded
Maximum Concurrent users	How many simultaneous users will use the service during a very short period	Concurrent users
Datasheet uploaded	Datasheet uploaded	Number of datasheet added

Table 1: Quantitative KPI sample.

Each quantitative KPI has been defined with an *Acceptance range of values to* decide at what point achievement of the objective begins. This range or threshole may be different depending on the service to evaluate and so forth the evaluation has been done accordingly.

CloudOpting have also used some testing tools to quantitative measure the performance of the platform. These tools are designed to extract analytical data, to test flows and integration, and to evaluate the performance and robustness of the platform.

Unit tests validated that a given piece of code, encapsulated in a function, worked as expected. Integration tests evaluated the communication among different components. At a higher level of abstraction, system tests were performed from the user interface. With this interaction, some flows were created and therefore evaluated.

Other type of tests executed were the performance & stress tests. This tests helped to know if the platform performed correctly and to measure the maximum concurrent userd accepted.

Qualitative methods

The qualitative methods have been focused in gathering information from two stakeholders: Public Administration and End Users. For this, a series of actions have been performed:

- Creation of surveys: two surveys were created, one for each stakeholder. These surveys can be seen at Annex III.
- A series of Hackathons were organised in different cities. The Hackathons included both the technical and business perspective.
- Face to face meetings were set to listen to specific requirements and feedback from end users.

The questionnaires provided valuable information that have been classified into a set of qualitative KPIs. The "Service and platform evaluation questionaire" has been answered by service providers and has been served to evaluate both, the services and the platform.

The qualitative KPIs that have been extracted from "Service and platform evaluation questionaire" are the following.

KPI	KPI Description
Service usability	The perception of the usability of the service for the subscriber
Customisation	Degree of customisation of the service
Support	The perception of the support received
Satisfaction	The perception of the satisfaction of the service for the subscriber
Monitoring usability	The perception of the usability of the monitoring for the subscriber
Willingness to renew	Willingness to renew the service

Table 2: Qualitative KPIs definition for the service subscriber.

The "End user service evaluation questionnaire" has been answered by end users and has been served to evaluate the quality of the services from their perspective and perception.

The qualitative KPIs that have been extracted from "End user service evaluation questionnaire" are the following.

KPI	KPI Description
Service usability	The perception of the usability of the service for the end user
Support	The perception of the support received
Rate	Rating of the services compared to others
Satisfaction	The perception of the satisfaction of the service for the end user

Table 3: Qualitative KPIs for the end user.

All the questions of the surveys were answered with a "1 to 5 stars" rating, or with a "5 word" answer, from low to high values. An average of all the answers have been performed and translated to a one to five star aproach in order to easy its evaluation. With all these information, and a defined range of acceptance, the services and platform have been evaluated.

The CloudOpting consortium also organised events and meetings in order to get a valuable feedback from the Public Administrations and End Users of key countries in the consortium.

One of the most relevant events were the hackathons driven in differents countries. During the hackathon, the platform and its architecture were presented. Atendees had the chance to play with the platform and the process to migrate a service. During the business day, participants had the chance to debate about the business model behind CloudOpting.

In addition to the hackathon, a workshop was done to present the CloudOpting platform in Barcelona. This workshop was also taken as an opportunity to meet relevant actors and assess the interest on the platform. More information about the Barcelona workshop can be found in deliverable D6.3

Apart from public events, some private meeting have been done with some interested Public Administrations.

The results and feeback gathered from these qualitative methods has been agregated and processed to understand the needs of the users and to draw the path to where the CloudOpting platform has to evolve. All the conclusions are described in the final Conclusions chapter.

3. Evaluation

The CloudOpting platform and the pilots will be evaluated in this section. For this, the collected data from the key performance indicators (KPI) will be used in order to have a base for interpretation.

Although all services will be evaluated, a more in-depth evaluation will be made for the four more representative services chosen in "D.4.1 Specification of pilot scenarios and evaluation criteria". The services selected for the action and evaluation plan (one per pilot city) were Clearò in Piemonte region, CKAN deployed in Barcelona, Bus Portal from the Corby pilot and '*Where can I charge my car?*" which belongs to the Karlshamn pilot.

3.1. Platform

The platform validation has been done by monitoring the platform server for the main metrics and KPIs along with the services. It was also validated by performing various tests organised around several levels: each one focused on validating a specific aspect of the platform.

The tests performed for the technical evaluation can be seen at Annex II *Platform Tests*. Tests have been classified in two main categories: functional and non-functional.

The functional tests are focused on the evaluation of the software characteristics that set up how the system behaves and what the system does. On the other hand, the non-functional tests evaluate other kind of characteristics such as performance or usability.

In CloudOpting project the following testing have been determined as a functional level; the unit tests, the integration tests, the system test and acceptant test. At non-functional level, the tests performed were performance and stress tests.

The platform validation involves not only the evaluation of quantitative aspects, but also qualitative ones. For qualitative metrics, a questionnaire for service owners (See Annex I) have been created. The values extracted from the questionnaire answers have been used to fulfil the Table 5 and see if they are what expected.

This kind of test also are called acceptance test. These tests enable the engagement of end users during the platform validation process. To consider the point of view of the final user reinforces the commercial vision of the product. This test main difference from others is that is not automated, i.e. it is done under consideration.

The objective is not only to know if the software is working as expected, but also to have qualitative opinions about it.

End user opinions have been considered during the validation process of the platform, as well as in its development, which allows the improvement of it as a product and fosters a better preparation of the market penetration and acceptance.

Quantitative metrics and KPIs

The quantitative metrics gathered during the Y3 period are listed in Table 4.

Metrics	Metric Description	Acceptance limit	Value
Total System CPU usage	The sum of total CPU usage	<70%	5%[1]
Memory used	Measures the amount of memory used in the system	<60%	55%
Number of downs	The number of downs in the CO platform	<4	2+1[2]
Number of user of the platform	Total sum of user using the platform of CloudOpting	>12	12
Number of users of pilots	Total sum of user using any service of CloudOpting	>1600	>36000
Number of Cities using CloudOpting	Sum of Cities involved in CloudOpting Services	>4	4
Number of Instances	Total sum of instances deployed using CloudOpting platform	>16	12
Number of services in CloudOpting	Total sum of services migrated to CloudOpting platform	>8	10
Number of cross city services	Total sum of services deployed physically in a city providing service to a different city	>8	6

Table 4: Quantitative KPI results.

Looking into the results from the table above, it can be concluded that almost all metrics have been accomplished as expected.

The value shown in the Table 4 for the CPU usage is very low as could be expected from an efficient website interacting with a simple database. The need of a powerful CPU comes when a service is requested for deployment. Then, the process consumes the full capacity of the CPU along its execution.

At the moment, the number of cities involved in CloudOpting services remains the same as in the beginning of the project, but some companies outside the project and outside the countries of the consortium have been

[1] The averaged CPU for a full year.
[2] Planned down for release.

requesting more information and meetings in order to know how to be involved and contribute to enhance the catalogue.
The number of users using the pilots has been higher than targeted in first place, meaning that the interest on the platform may grow depending on the services CloudOpting shows in its catalogue.

Qualitative variables

The qualitative metrics gathered during the Y3 period are listed in Table 5.

Metrics	Metric Description	Acceptance limit (over 5 Stars)	Values
Replicability	Could the same application provide service to different cities	Yes	Yes
Interoperability	This KPI measure the capacity of the system to be independent of the physical platform in which the application can be deployed.	Yes	Yes
Usability	How a user evaluates CloudOpting usability of the User Interface of CloudOpting	3 Stars	3.4 Stars
Customisation	Degree of service customisation	3 Stars	3.3 Stars
Technical support	How was the technical support given	3 Stars	3.9 Stars
Downtime	Downtime appreciated by Service Providers.	3 Stars	4 Stars
Satisfaction	Degree of satisfaction	3 Stars	3.8 Stars
Monitoring	How a user evaluate CloudOpting monitoring system	3 Stars	3.1 Stars
Satisfaction over other possibilities	Satisfaction over other possibilities	3 Stars	3.5 Stars

Table 5: Qualitative KPI results.

The table shows that the replicability and the interoperability have been a success. The platform can deploy applications to provide service to different cities (i.e. the CKAN Barcelona and Corby instances) and also has demonstrated its interoperability capabilities by deploying services in different clouds such as CloudStack and DigitalOcean.

To go more in detail, the following figure shows the answers of the question "How was the downtime of the services?".

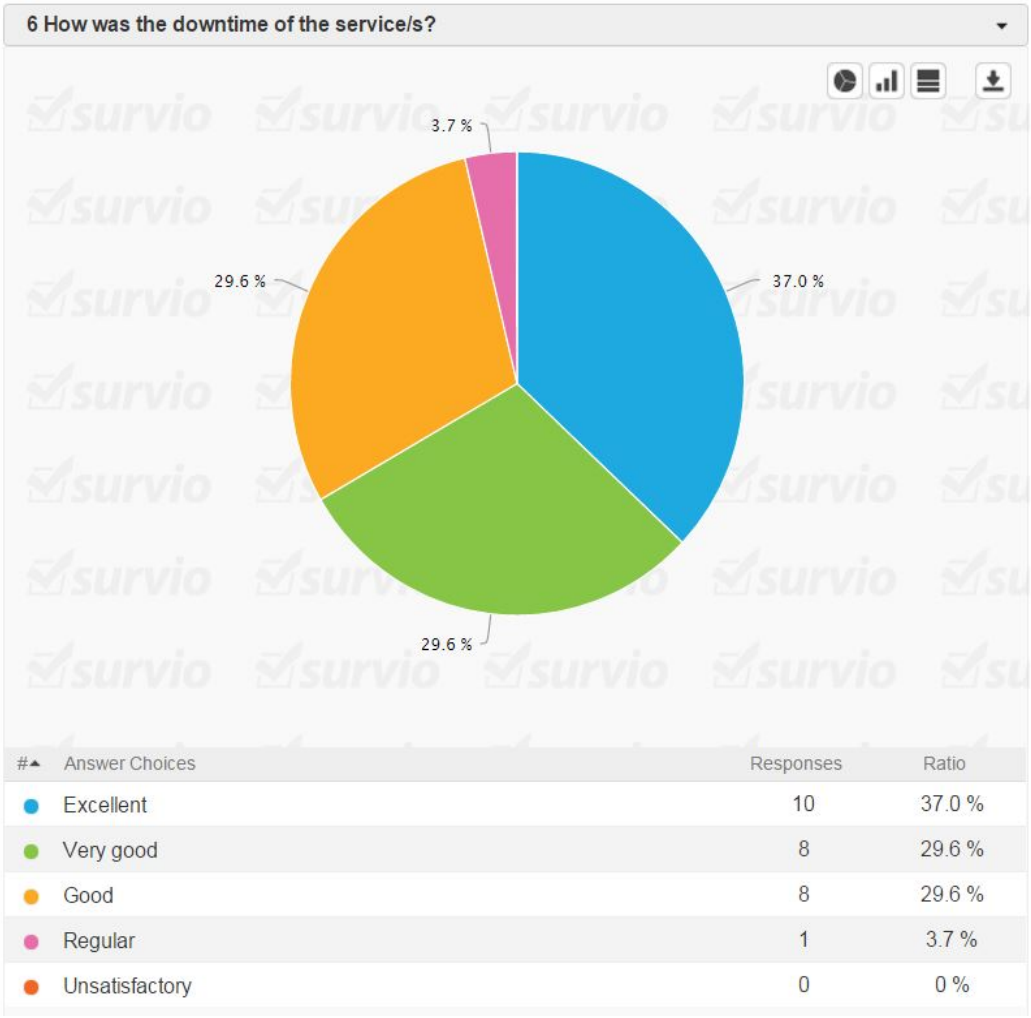

#	Answer Choices	Responses	Ratio
1	Excellent	10	37.0 %
2	Very good	8	29.6 %
3	Good	8	29.6 %
4	Regular	1	3.7 %
5	Unsatisfactory	0	0 %

Figure 1: Questionnaire results downtime of the services.

The figure above shows that a third of the users (66,6%) are very satisfied with the downtime perception of the services. All services have been deployed using strong methodologies and robust software solutions.

Additionally, thanks to the technologies that were selected to deploy and configure the services (Docker), the availability and the recovery time in case of failure is at the milliseconds order. Using these technologies will help to improve and incorporate new features that will lead to a great positive impact.

The following figure shows the answers of the question "How was the technical support given?"

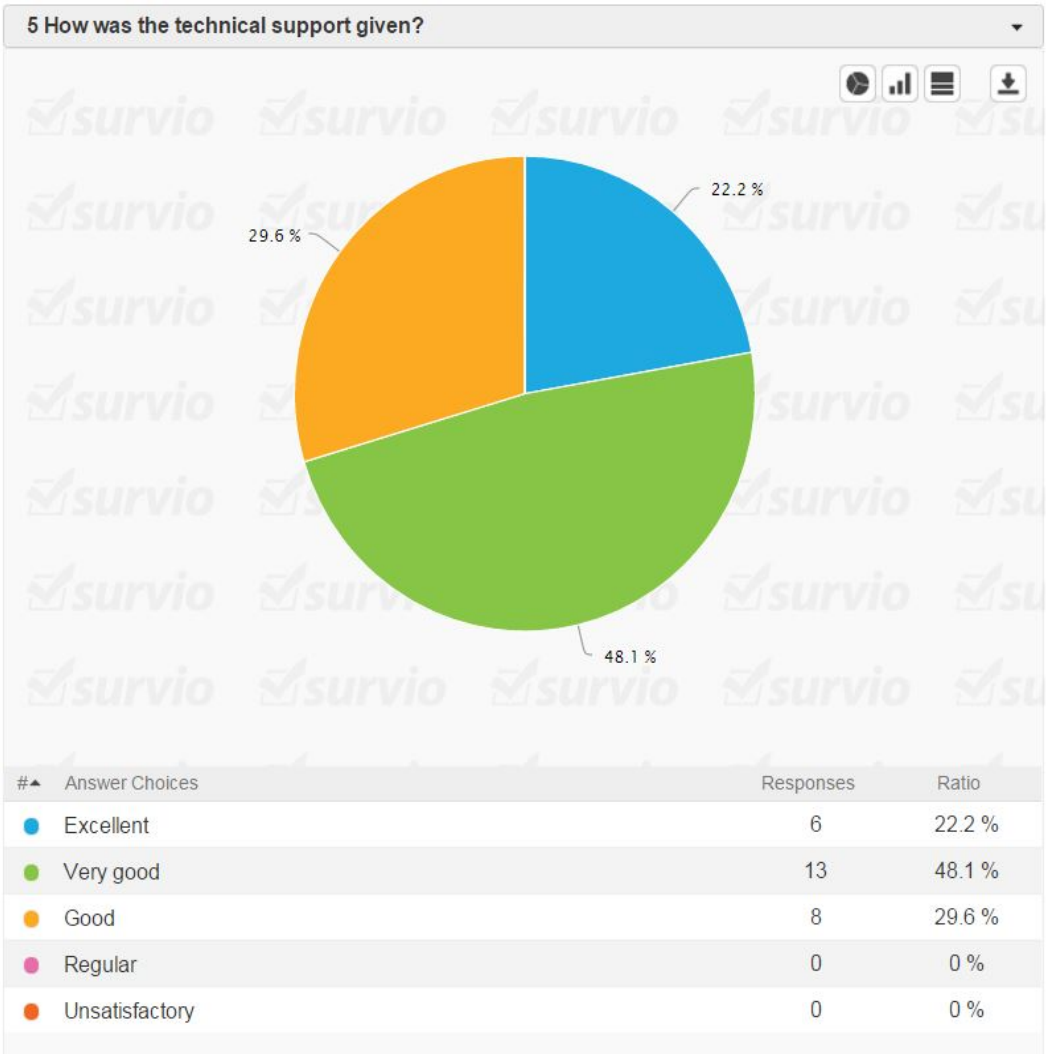

#	Answer Choices	Responses	Ratio
	Excellent	6	22.2 %
	Very good	13	48.1 %
	Good	8	29.6 %
	Regular	0	0 %
	Unsatisfactory	0	0 %

Figure 2: Questionnaire results of the technical support given.

The answers depicted in the graph shows that users were not unsatisfied at all and that technical problems were solvet with diligence and made the user feel that the technical support for the service was Good, Very Good, or Excellent.

The tracking of the issues and incidents of the platform has been done by means of the GitHub Issues system. That helped to bring users and developers together and to reduce the resolution time of the issues.

The following figure shows the answers of the question "How would you rate the monitoring system usability?"

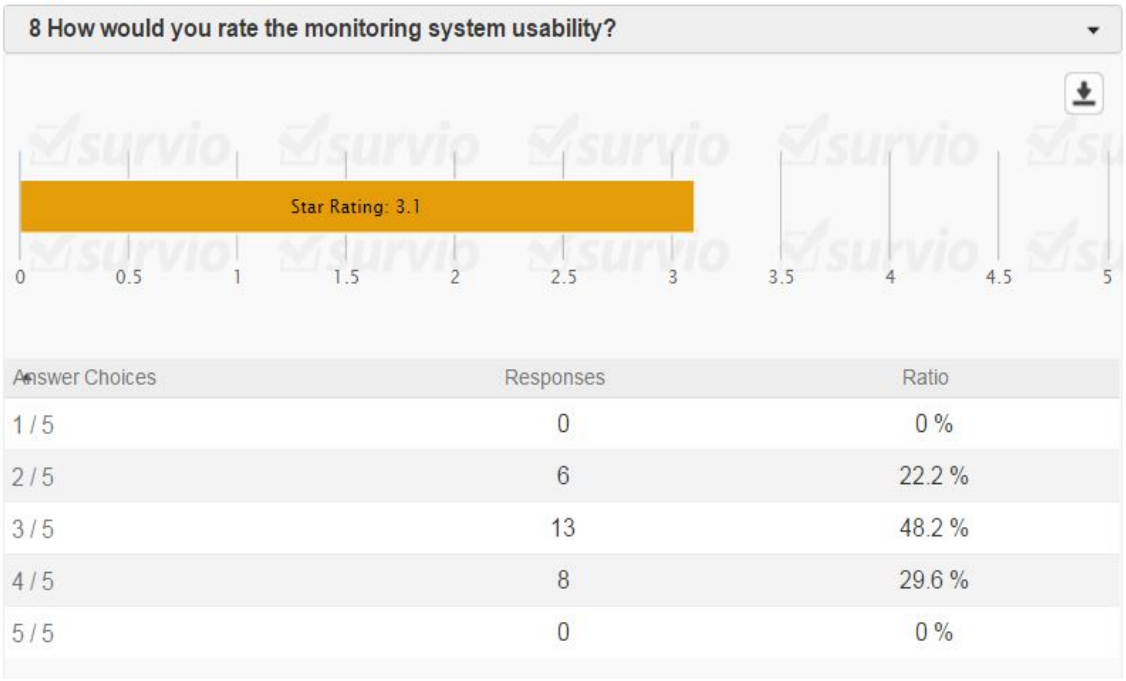

Answer Choices	Responses	Ratio
1 / 5	0	0 %
2 / 5	6	22.2 %
3 / 5	13	48.2 %
4 / 5	8	29.6 %
5 / 5	0	0 %

Figure 3: Questionnaire answer of monitoring system usability.

The rating average of this question is one of the lower. It states that some improvements can be performed to satisfy user's needs and expectations. As an example, the time to retrieve information to draw some graphs is too high and the number of metrics offered is big, so the user could be overwhelmed.

Figure 4 below, shows the answers of the question "Which was the degree of customisation of the services?"

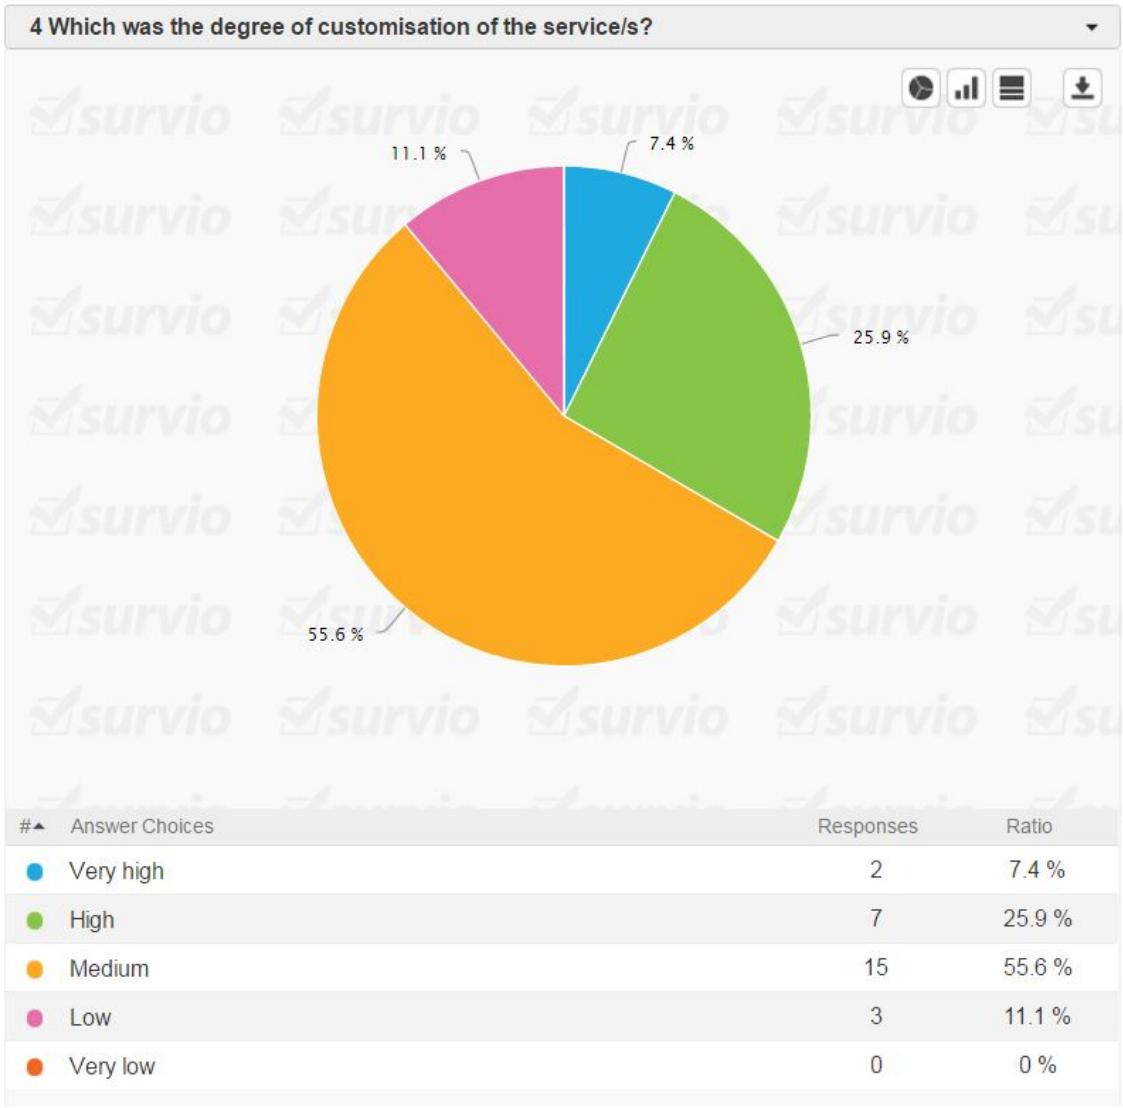

#	Answer Choices	Responses	Ratio
●	Very high	2	7.4 %
●	High	7	25.9 %
●	Medium	15	55.6 %
●	Low	3	11.1 %
●	Very low	0	0 %

Figure 4: Questionnaire results of customisation degree.

This is a complex aspect that each service should evaluate separately. The platform gives the possibility to Service Providers to choose what parameters the subscribers will be able to customise. Therefore, this is not only a platform problem but a global problem.

For the Service Provider, the customisation of a service is done during the migration process by using the tools the platform provides. If they do not feel comfortable because of their technical skills, the TOSCA IDE may be used instead.

In case of the Service Subscriber, the customisation depends on the decisions taken in the subscription process. They have to include the

information required by Service Providers by filling the blanks at subscription time that will be used during the deployment.

Evaluation and conclusions

From the end user perpective can be concluded that the platform has been a success, even though some improvements can be made in order to increase the overall satisfaction.

Replicability and Interoperability have been acomplished successfully. An extra interoperability level has been added thanks to the docker technology. It provides a more extense compatibility over different systems and also allows to replicate services reduce its startup time to seconds.

The unit tests, focused on the most critical functionalities of the platform, were passed successfully (See Annex II). Besides, the integration tests, with the two fundamental blocks of CloudOpting (Manager and Crane), returned a positive feedback, meaning that the interaction between them works expected.

The system tests, used to test the platform automatically by simulating the interaction from a user perspective, also were passed successfully.

As a summary, the platform worked properly at the unit level. Consequently, this allows the evaluation of the platform from a higher level as follows.

For the Hackathons, valuable information was extracted. More than 80 participants attended the Hackathons in the 2 cities. All the target groups were present: service providers, service subscribers and cloud providers. If including those who connected online via streaming, there were over 160 attendants.

As mentioned, a Hackathon and a workshop were done, providing very useful information from potential customers and users. For more information go to deliverable D6.3

Face to face meetings have been very fruitful to qualitative evaluate and complete the expectations from end user point of view. Meetings were useful to know what is expected from the platform in terms of user interaction, hierarchy of the platform users and organisations, and what they would like to see/have in the platform.

3.2. Pilots

The four more representatives services implemented in the pilots have been monitored for a year. In the following sections you can find the evaluation of all of them.

3.2.1. Barcelona

The Barcelona pilot is located in the city of Barcelona. It uses the IMI's cloud for the deployment of services listed below:

- CKAN
- Sentilo
- City Agenda
- Next to Me
- FixThis

Apart from these services, IMI's cloud has been used as a logical cloud for a service of the Corby pilot (CKAN Corby), and the services from Karlshamn (Where can I charge my car and POI and Next2Me).

CKAN has been considered the most representative service of the Barcelona pilot. Because of that, its extended evaluation is described in the next section while the other Barcelona services are included in the *Other Services* chapter.

4.2.1.1 CKAN

Overview

The CKAN instance of Barcelona has been operated for 12 months during the Y3 of the project. This section describes the evaluation of the service during this period.

According to deliverable D4.1, the pilot is considered a success if the CKAN instance deployed in the IMI's cloud is capable of providing a reliable service to any CKAN user. In order to evaluate the reliability of the service, a set of KPIs were defined. Their corresponding metrics have been recorded to be evaluated at the end of the pilot phase.

During the deployment and the first week of operation the number of concurrent users was modified from a range [100 to 1000] to a more realistic range [50 to 100] taking into account the projection and needs of the service.

Quantitative metrics and KPIs

The quantitative KPI and metrics gathered during the monitoring period are listed in Table 6. The associated data is the average figures taken from the reporting templates of the evaluation period.

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	15%
Memory usage	Averaged % Memory usage	Memory %	< 80	25%
Stability	How stable is the service	Number of downs	[0 – 20] per year	2
Availability	Uptime of the service	Number of downs Service uptime	90%	99%
Access Time	How much time need the service to attend a request	Access Time	[20- 3000]ms	1900ms
Service visits	Number of users using the service	Visits sessions	[100 – 10000]	776
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	Not monitored
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[50 – 100]	>75
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[20 – 200]	62

Table 6: Service metrics and KPIs for CKAN.

Qualitative variables

The qualitative KPIs were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end users. These surveys can be found in Annex I. The link of the one focused on the end users was published on the CKAN instance while the service subscribers one was answered by the administrators of the portal.

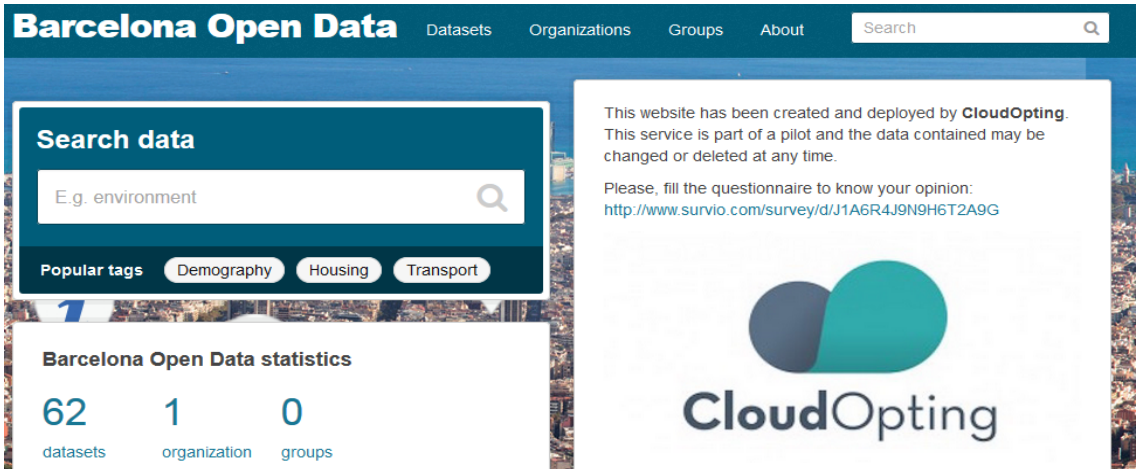

Figure 5: CKAN front panel including the link of the questionnaire.

All the qualitative metrics were measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service subscriber satisfaction	The perception of the quality of the service for the subscriber	Service satisfaction	> 3 Stars	**4 Stars**
Support	The perception of the quality of the support	Quality of Technical Support	> 3 Stars	**3.5 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 7: Qualitative KPIs for the service subscriber.

The following chart represents the service subscriber satisfaction gathered from the questionnaires delivered to the CKAN managing team on the Barcelona City Council.

Figure 6: Service Subscriber satisfaction

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
End users satisfaction	Overall end user satisfaction	Satisfaction rating	> 3 Stars	**3.25 Stars**
Service usability	The perception of end user about the usability of the service	Usability rating	> 3 Stars	**3.50 Stars**
Support	The perception of end user about the support received	Support rating	> 3 Stars	**2.85 Stars**
Service comparison	End user service comparison rating	Comparison rating	> 3 Stars	**3.42 Stars**

Table 8: Qualitative KPIs for the end user

Next figure represents the end user satisfaction.

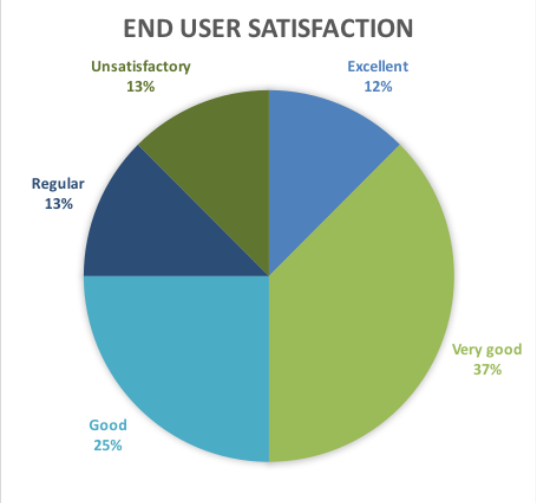

Figure 7: End User satisfaction

Evaluation and conclusions

- Deployment and Customisation:

The deployment of the service created a fully working instance, but some customisation improvements could be included in the following software versions. The CKAN configuration *ini* file template included in the installation process can be easily modified, but it was considered to leave it with a high security level and to modify it after the deployment of the instance. The modification of the *ini* file can be directly edited on the host while the process of reloading the instance is a matter of seconds thanks to the containerisation (Docker) of the service.

- Operation and KPI:

The operation of the instance has been performed with a reduced number of support staff due to the reliability of the service. Only a few issues regarding the cloud management affected the service and required the action of a technical operator. Most of the effort was then done by the publishers of the datasets.

Regarding the technical indicators, by the end of the period, the service included 62 datasets published by the operators of the service at IMI required less than 15% of CPU usage and 25% of memory. Thus, considering that the performance KPI successfully covered the expected values, the VM size may be decreased to reduce costs in the following deployments.

The containerisation of CKAN has been demonstrated to be very efficient in terms of use of resources and reliability reporting very few services down, most of them related to the cloud maintenance rather than the service itself.

In the area of Barcelona, where the service has been running during this year, no other cities host data portals with more than 50 datasets (Table 9). Thus, the technical configuration that has been evaluated could be implemented in most of the regional cities in this area.

City	Datasets
Sant Cugat	9
Tarragona	9
Viladecans	19
El prat	23
Cornella	28
Sant Feliu	40

Table 9: Number of datasets in relevant cities.

The number of concurrent users successfully tested (>75) is excellent considering the number of visits of the original Barcelona Open Data portal.

- User experience:

Although the number of questionnaires answered is low, the results reported a satisfactory use of the CKAN instance by end users and service subscribers. The support was rated below expectations, but contacts made for issues were very low. The usability of the instance and the overall satisfaction has reached the acceptance values indicating a good performance of the instance and the piloting phase.

The user experience with the CKAN instance deployed with CloudOpting has been demonstrated to be transparent and very similar to the one when using standard technologies. Thus, it has been proved the viability to use Docker as core technology for the CloudOpting platform.

As a conclusion of the overall evaluation, the deployment and performance of the service included in the platform during the piloting phase has successfully achieved the expectations.

3.2.2. Corby

4.2.2.1 Bus Portal

The Corby Bus Portal is an online website that allows users to find their nearest bus stop and the bus they need to take for where they need to travel. The Bus Portal was given the domain name 'Catch a Bus', and can be found at www.catchabus.co.uk and www.catchabus.eu

Overview

The Corby Bus Portal (www.catchabus.co.uk) has been operating for 12 months during the Y3 of the project. The service is part of the Corby pilot which includes a set of other services. This section will describe the evaluation of the service during this period.

The Corby Bus Portal allows user to search for their nearest bus stop and find the bus they need.

According to deliverable D4.1, the Bus Portal instance would be a success if it is capable of providing a reliable and cost effective service to users and stakeholders. The Bus Portal was a new service that would test whether such an instance was able to perform in a production environment, and it was important to test whether the Bus Portal would be a success if it was capable of supporting a standard traffic volume of users.

Metrics for the Service

- Results / KPIs

Quantitative metrics:

- Technical variables
  - CPU
  - Memory
  - Server Failures
  - Number of downs
  - Downtime after a failure

- Business variables
  - Effort (Staff-hours)
  - Time to identify bus route
  - Number of interactions between users and bus companies
  - Use of other local applications
  - Issues related to pilot plan execution

- Service analytics
  - Audience: Visits (Sessions) and Unique Visitors (Users)
  - Number of participants who reviewed a journey
  - Datasets added
  - Bus routes feedback
  - Bus stations feedback

Qualitative variables:

- Service subscriber perception about the quality of service received
- Service subscriber perception about the technical support
- Service subscriber willingness to renew the subscription service
- End user perception about the quality of service received:
- End user perception about the response time
- End user perception about the usability of the service
- End user opinion about comparison with other similar services.

Quantitative metrics and KPIs

The quantitative metrics gathered during the monitoring period are listed in Table 9 below. The associated data are the average figures for one month taken from the reporting templates across all of the evaluation periods (except where specified).

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	25%
Memory usage	Averaged % Memory usage	Memory %	< 80	95%
Stability	How stable is the service	Number of downs	[0 – 20] per year	0
Availability	Uptime of the service	Number of downs Service uptime	90%	100%
Average Number of interactions by users	How many interactions did users have with the service per month	Number of user interactions	[20- 4000]	1720
Service visits	Average number of users using the service per month	Visits sessions	[20 – 200]	134
Unique and Returning Visitors	Average number of users that were unique to the service per month	Unique Users	[20 – 200]	88.75
Time to Identify Bus Route	The time taken in seconds to identify the bus route	Manual evaluation of time taken to identify route	[1 – 5]min	5s

Journey reviews	Total number of participants who reviewed a journey	Number of journey reviews	[0-100]	0
Total Datasets added	Total number of datasets added	Datasets added	[50-2000]	1400
Bus Routes Feedback	Total number of feedback received on bus routes	Bus routes feedback	[0-100]	0
Bus Stations Feedback	Total number of feedback received on bus stations	Bus stations feedback	[0-100]	1

Table 10: Service metrics and KPIs for the Bus Portal.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focussed on the administration of the instance and the other one focused on the end users. These questionnaires can be found in Annex I.

All the qualitative metrics were measured through questionnaires distributed to users of the Corby Services through the Service websites and on the CloudOpting pages on the Electric Corby website.

The qualitative metrics were measured using satisfaction values from 1 to 5. Table 11 below summarises the qualitative KPI results for the Corby Service Subscribers.

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service subscriber satisfaction	The perception of the quality of the service for the subscriber	Service satisfaction	> 3 Stars	**4.5 Stars**
Support	The perception of the quality of the support	Quality of Technical Support	> 3 Stars	**4.5 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**4.5 Stars**

Table 11: Qualitative KPIs for the service subscriber.

The qualitative KPIs taken from the End User Questionnaires can be found below in Table 12.

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
End users satisfaction	Overall end user satisfaction	Satisfaction rating	> 3 Stars	**3.8 Stars**
Service usability	The perception of end user about the usability of the service	Usability rating	> 3 Stars	**4.3 Stars**
Support	The perception of end user about the support received	Support rating	> 3 Stars	**4.3 Stars**
Service comparison	End user service comparison rating	Comparison rating	> 3 Stars	**3.6 Stars**

Table 12: Qualitative KPIs for the end user.

The end user satisfaction can be seen in further detail in the chart below. It indicates that a high amount of end users were satisfied with the Corby services.

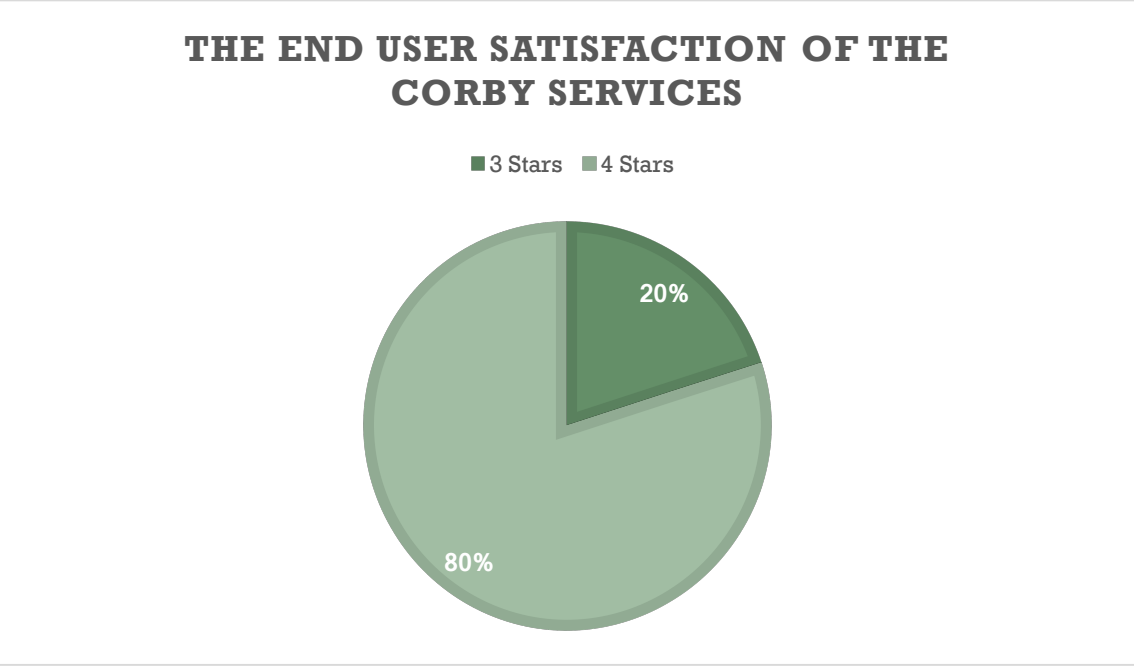

Figure 8: End User Satisfaction of the Corby Services.

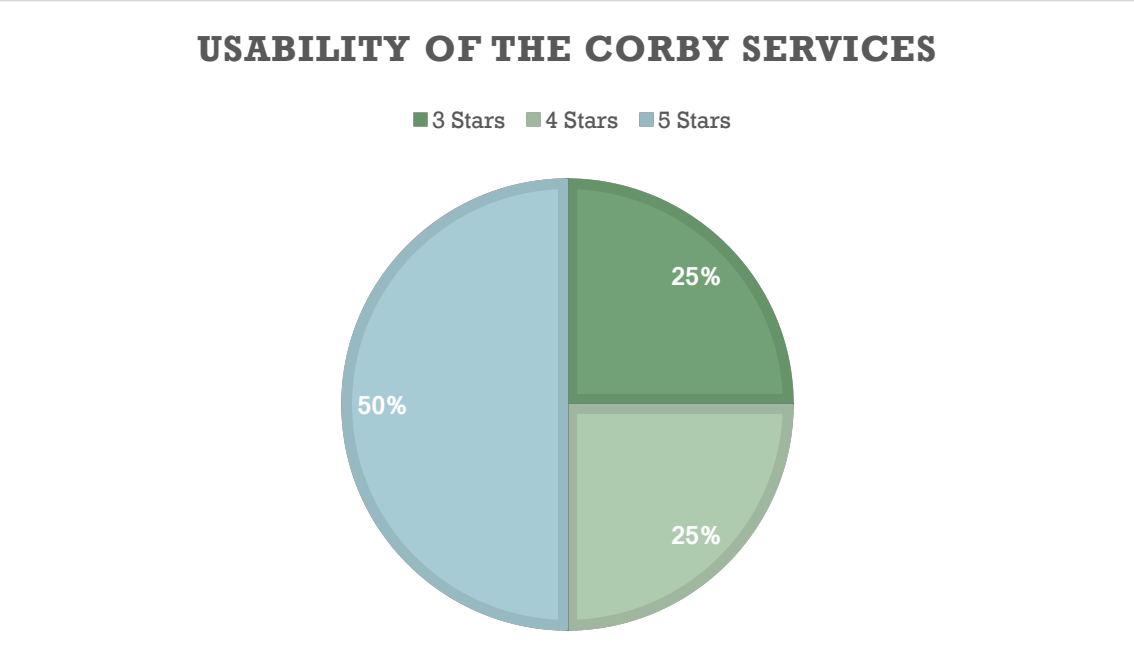

Figure 9: Chart displaying the users' perceived usability of the service.

One comment that an end user left for the bus portal was that they would like to see route planning based on GPS, but that they found the service simple and well delivered.

Evaluation and conclusions

- Deployment and Customisation

As discussed in D4.2, the Bus Portal was deployed in week 1 after the Deployment of the catalogue, with the domain name www.catchabus.co.uk and www.catchabus.eu.

The Bus Portal was deployed and hosted in CSI's cloud. In order to operate correctly, the software 'Open Street Map' and data from Traveline, was maintained by Teamnet.

After the deployment of the Bus Portal, some basic variables were customised, that were not previously configured through the CloudOpting platform, related to the aesthetic of the website. These included:

- Website title
- Theme colour selection
- Background colour
- Welcome message
- About message

For this instance, the Bus Portal was customised to feature the Electric Corby logo and brand details, and information about Electric Corby. The website was also customised so that the map used for the bus portal showed the local area of Corby, and the local Bus Stops and Bus Routes.

The level of customisation that could be achieved was good, and the Bus Portal allowed for both aesthetic customisation and more technical customisation to be of use to the local area.

- Operation and KPI

According to the results seen in the Quantitative KPIs seen in Table 10 , we can conclude that the Bus Portal provided a successful service to users from the CloudOpting platform. It achieved a reliable service as evidenced through the number of downs observed (0, see Stability) and the Service Uptime (100%, see Availability).

The Bus Portal was also able to support a standard traffic of users, as evidenced by receiving on average 152 visits per month, with 2231 interactions. In total, the Bus Portal received 20645 interactions, 1606 visits, and 1065 unique and returning visitors.

- User Experience

The User Experience for the service was good, as evidenced by the number of users of the service as well as the overall satisfaction of the Corby Services.

The Corby Services as a whole have achieved very good qualitative results from respondents of the End User Questionnaire and the Service Subscriber Questionnaire. This good results from the End Users rating include their satisfaction with the Corby Services (with 3.8 stars out of five), and the support provided and service usability with a very good overall evaluation .

3.2.3. Piedmont

4.2.3.1 Clearò

Overview

The service Clearò was already operating before Y3, as previously stated. During Y3, the service was specifically migrated to CloudOpting for 12 months for the scope of the project and then assessed. According to deliverables D4.1 and D4.2, the pilot would be considered a success if a new instance of Clearò deployed on the CloudOpting catalogue is sufficiently mature to move to a production environment.

This section describes the results of the monitoring and assessment of the service during this period with support of a set of specific KPIs and other qualitative data of the service.

Metrics for the service

Here below the quantitative and qualitative metrics (KPI) applied are listed.

Quantitative metrics:

Technical variables.

- CPU usage
- Availability
- Stability
- Access Time

Service analytics:

- User audience
- Public Audience
- Files Uploads
- Files downloads
- Use of Dynamics Data List (DDL)

Qualitative variables:

All qualitative data are based on a commonly agreed online Questionnaire, investigating the following items.

- Service subscriber perception about the quality of service received
- Service subscriber perception about the technical support
- Service subscriber willingness to renew the subscription service
- End user perception about the quality of service received:
- End user perception about the response time
- End user perception about the usability of the service
- End user opinion about comparison with other similar services.

All these qualitative metrics have three possible values:

- Excellent
- Good
- Unsatisfactory

Quantitative metrics and KPIs

The next table identifies the finally elicited KPIs for the service. Some amendments from the Pilot Plan for tests have been operated.

KPI name	KPI Description	Associated Metrics
CPU usage	Average % CPU usage	CPU %
Availability	Uptime	Server failures
Stability	How stable is the service	Number of downs
Access Time	How much time the service waits for a request	Access Time
User audience	How many back-office users are accessing to the service	Number of back-office users' accesses
Public Audience	How many visits and visitors the service has	Number of visits (Sessions) Number of unique Visitors (Users)
Files Uploads	How many files are being uploaded by the back-office users and their size	Number of uploaded file Size of uploaded files
Files downloads	How many files have been downloaded by public users	Number of file downloads
Use of Dynamics Data List (DDL)	How much the DDL has been modified by back-office users	Number of updated rows

Table 13: Service metrics and KPIs for Clearò.

The Qualitative Variables result: (Excellent, Good or Unsatisfactory)

Description
Service subscriber perception about the quality of service received
Service subscriber perception about the technical support
Service subscriber willingness to renew the subscription service
End user perception about the quality of service received:
End user perception about the response time
End user perception about the usability of the service
End user opinion about comparison with other similar services.

Table 14: Qualitative variables for Clearò.

All statistics on pilot performance are based on three main sources:

- Google Analytics;
- Symon (enterprise dashboard for service monitoring implemented by CSI);
- Logs from Clearò platform.

From Symon we extracted data on KPI Apdex, availability and user satisfaction. Especially, **Apdex** (Application Performance Index) is an open standard developed by an alliance of companies that defines a standardized method to report, benchmark, and track application performance (http://www.apdex.org/overview.html).

The index is based on three dimensions of application's responsiveness:

- **satisfaction**: users are fully productive. This dimension represents the time value (T seconds) below which users are not impeded by application response time;
- **tolerance**: users notice performance lagging within responses greater than T, but they continue the process;
- **frustration**: performance with a response time greater than F seconds is unacceptable, and users may abandon the process.

Apdex Rating	Apdex
Excellent	0.94–1.00
Good	0.85–0.93
Fair	0.70–0.84
Poor	0.50–0.69
Unacceptable	0.00–0.49

Figure 10: APdex Table.

The graphs below report data referred to period 01/01/2016 – 31/12/2016.

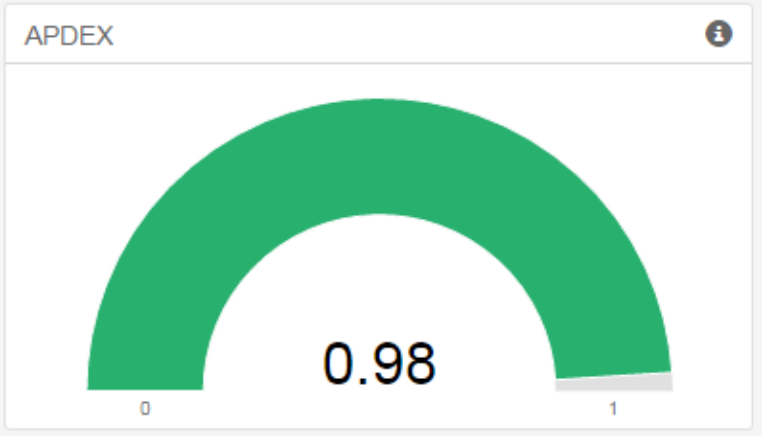

Figure 11: APDEX, Clearò's performance 2016. (Source: Symon)

Figure 11 reports the value of APDEX, that results highly satisfied. Also Availability, see Figure 12, is quite good as it is very close to 100%.

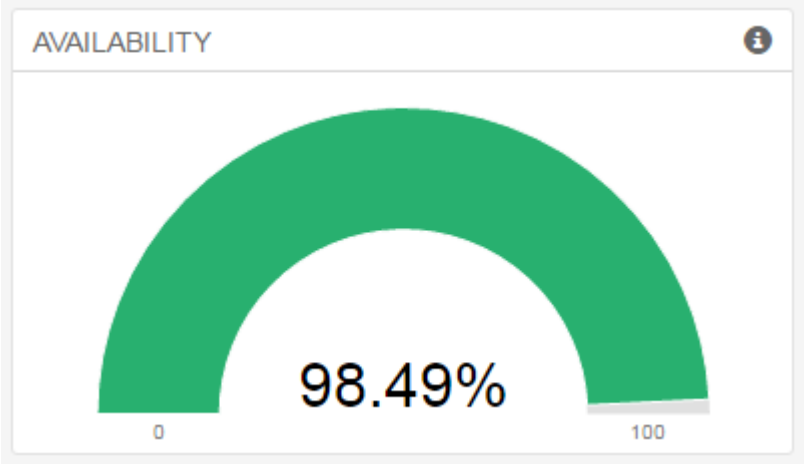

Figure 12: Availability, Clearò's performance 2016. (Source: Symon)

Next graph shows results in terms of User Satisfaction, see Figure 13.

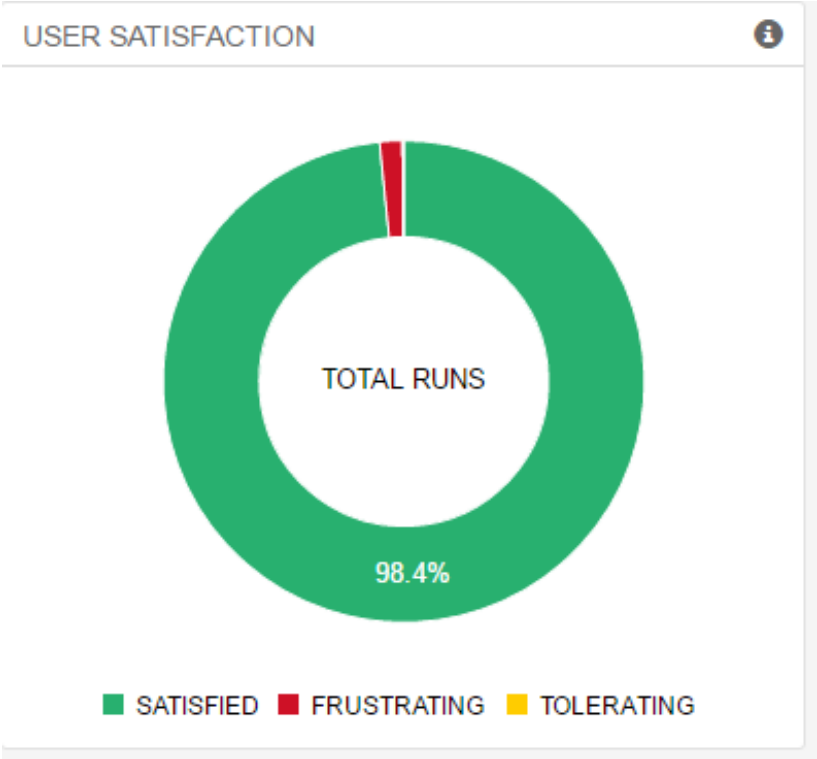

Figure 13: User satisfaction, Clearò's performance 2016. (Source: Symon)

The 1.5% of "frustrated" is a result of the benchmark "Response time" as defined in Symon (for all CSI's applications). Throughout the monitoring period, this benchmark was always outdone at first access because Clearò needs a daily restart (for the backup of the database that involves an upload in the cache of all statics resources and thus is generally time-consuming).

Next Figure 14 clearly represents this phenomenon. Herein the peak in the daily access times for the first access is clearly visible. In order to have a more robust result for this variable, one should then exclude the times for first access.

Figure 14: Performance – Access time per day. (Source: Symon)

The next table finally reports the values tracked during 2016 of all KPIs for Clearò.

KPI	KPI Description	Associated Metrics	Acceptable range of values	Result	KPI achievement
CPU usage	Average % CPU usage	CPU %	< 1% (***)	20%	YES
Availability	Uptime	Server failures	98%	98.49%	YES
Stability	How stable is the service	Number of downs	3 per year	2	YES
Access Time	How much time the service waits for a request	Access Time	<1 second	7.62 second (*)	NO
User audience	How many back-office users are accessing to the service	Number of back-office users' accesses	> of current value (**)	134 users' accesses (0.37 users/day) (0.53 user/working days)	NA
Public Audience	How many visits and visitors the service has	Number of visits (Sessions) Number of unique Visitors (Users)	> of current value (**)	6,042 visits 3,630 unique visitors	NA

Files Uploads	How many files are being uploaded by the back-office users and their size	Number of uploaded file Size of uploaded files	> of current value (**)	493 upload	NA
Files downloads	How many files have been downloaded by public users	Number of file downloads	> of current value (**)	64,132 download	NA
Use of Dynamics Data List (DDL)	How much the DDL has been modified by back-office users	Number of updated rows	> of current value (**)	3,486	NA

Table 15: Quantitative metrics and benchmark for Clearò evaluation, 2016.

(*) "Access Time" is higher than the threshold because one decided to monitor the overall availability of the service rather than the access of the final user (citizen). Consequently, only the time for login and download of the home page is considered.

(**) A number of KPIs defined as "Acceptable" range the actual value "> of current value". Unfortunately, initial values (t0) for the service are not available. The overall monthly trend has been analysed to elicit out potential increases in the rate of uses/downloads, with some biases: in the first phase of the project, a number of massive uploads for data porting on the system have been done, so as to mislead the trend. Moreover, transparency has in general a seasonal trend so that monthly accesses' rate are not so significant for trend analysis.

(***) As far as CPU is concerned, please refer to graphs below. All data derive from the server for the monitoring of CloudOpting (the name of the machine is defined by CSI's standard and refer to all 2016).
CPU consumed (Figure 11) is very low because the hosting virtual machine is of high level (8 GB RAM and 4 CPU). Band's consume highlights a constant use by part of the CPU. Figure 12 demonstrate that the service did not experience any failure on 2016 (CSI's server is *FEAPI-opendata.odsp.csi.it)*.

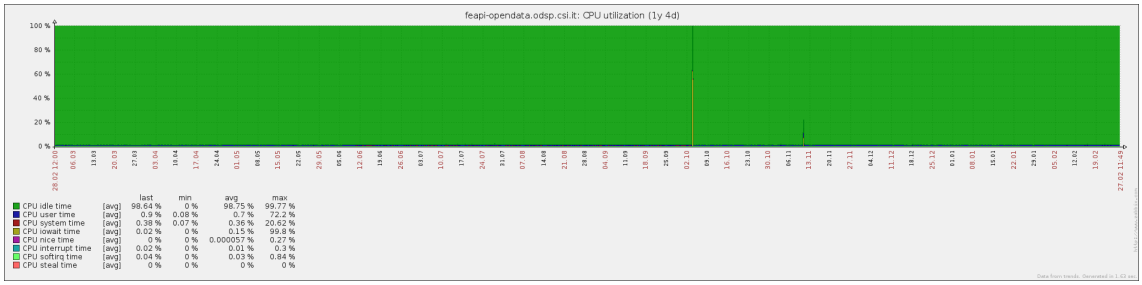

Figure 15: CPU usage 2016, Clearò. (Source: Symon)

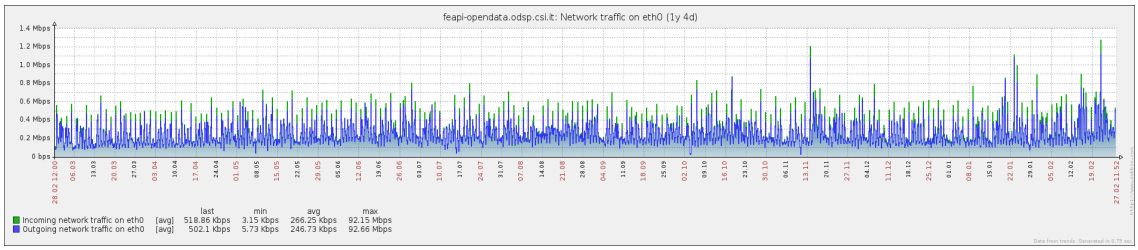

Figure 16: Band network usage, 2016 Clearò. (Source: Symon)

Qualitative data

A number of overall 12 questionnaires (on 26 total users) collected from **service subscribers** were collected in 1 week. Here below a summary of data emerged.

Citizens (final users) questionnaires are too few to give a robust analysis (4 questionnaires), so that we excluded them.

As for usability, the overall response of the questionnaires is positive: 75% of users say Clearò's usability is good to excellent (8%). Unsatisfactory is only 8% (1 answer).

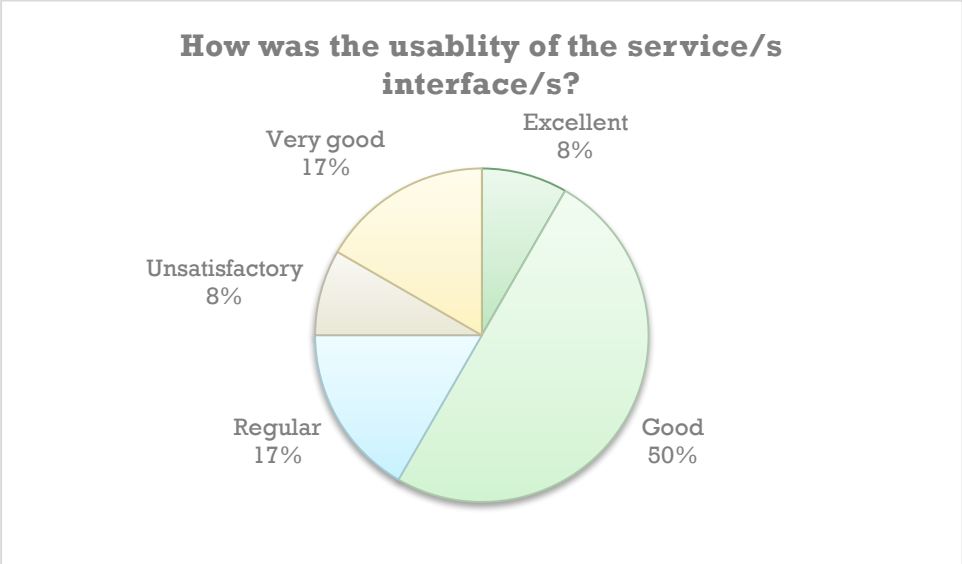

Figure 17: Usability of Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

As for Customisation, once again, the overall response of the questionnaires is positive: 42% of users say it is high to excellent (8%). Unsatisfactory is only 8% (1 answer).

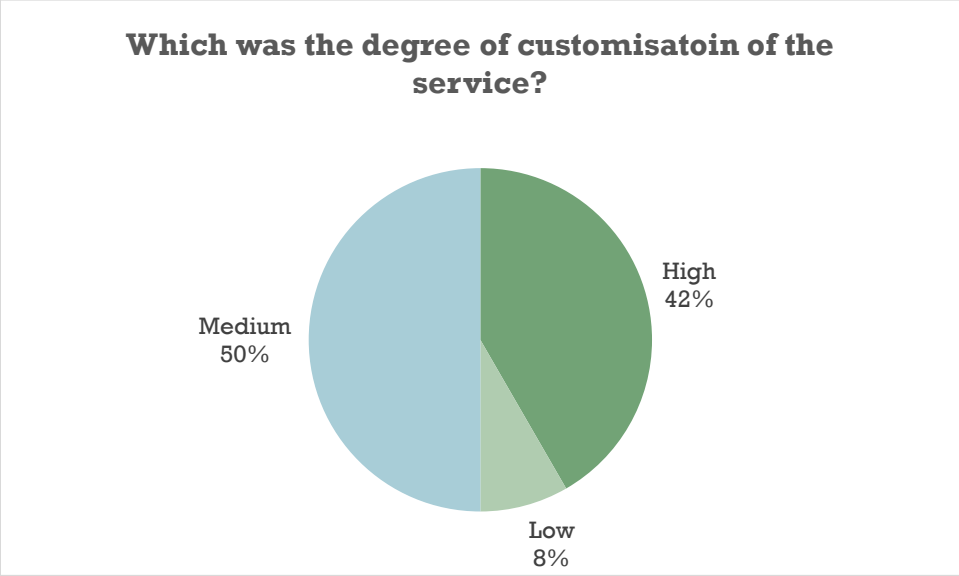

Figure 18: Customisation of Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

As far as the technical support given is concerned questionnaires collected only positive response, with "Excellent" rated at 25%.

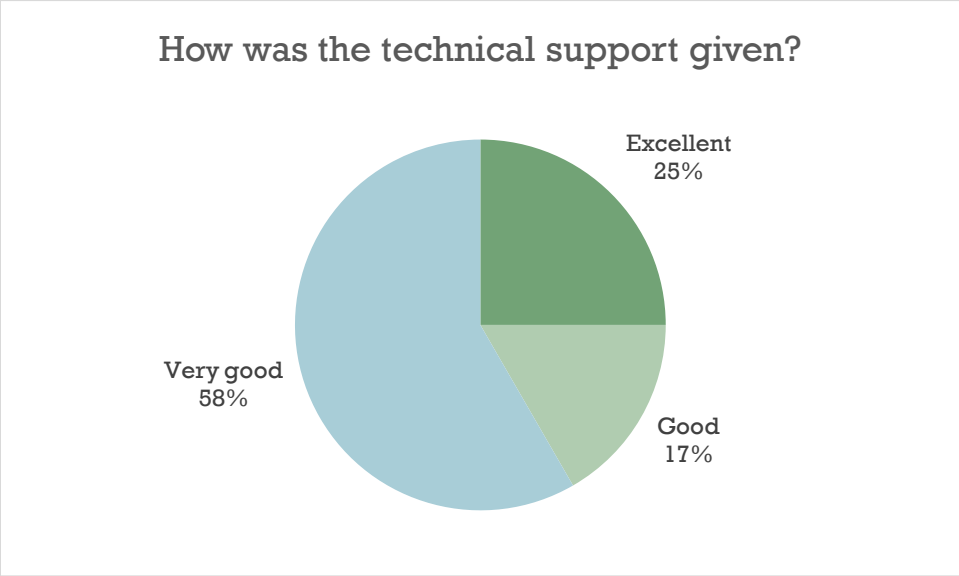

Figure 19: Technical support of Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

The regular response for downtime only was chosen by one. Only positive responses in general, with "Excellent" rated 25%.

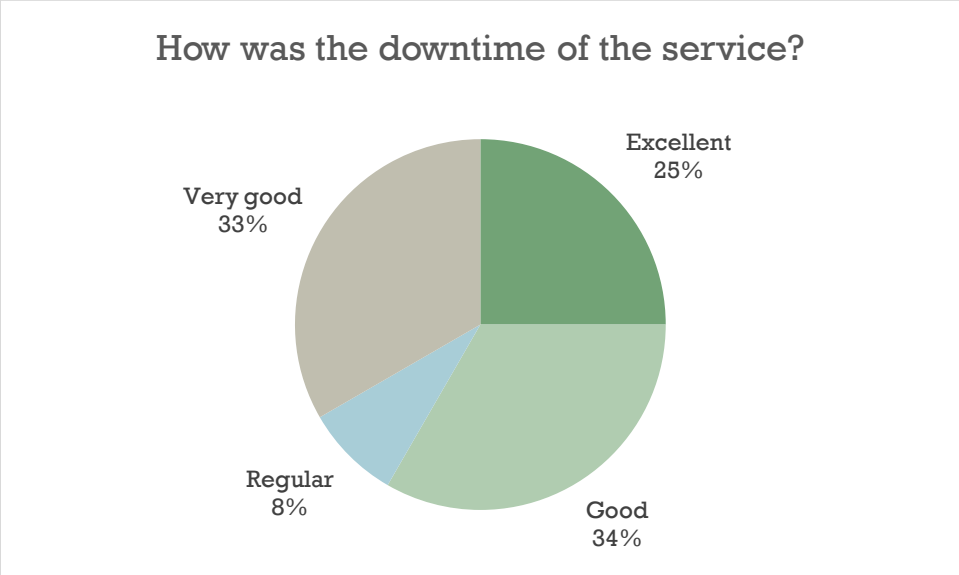

Figure 20: Downtime of Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

Next, the overall satisfaction is rated 3 or more (on a 5 scale degrees) on 91%.

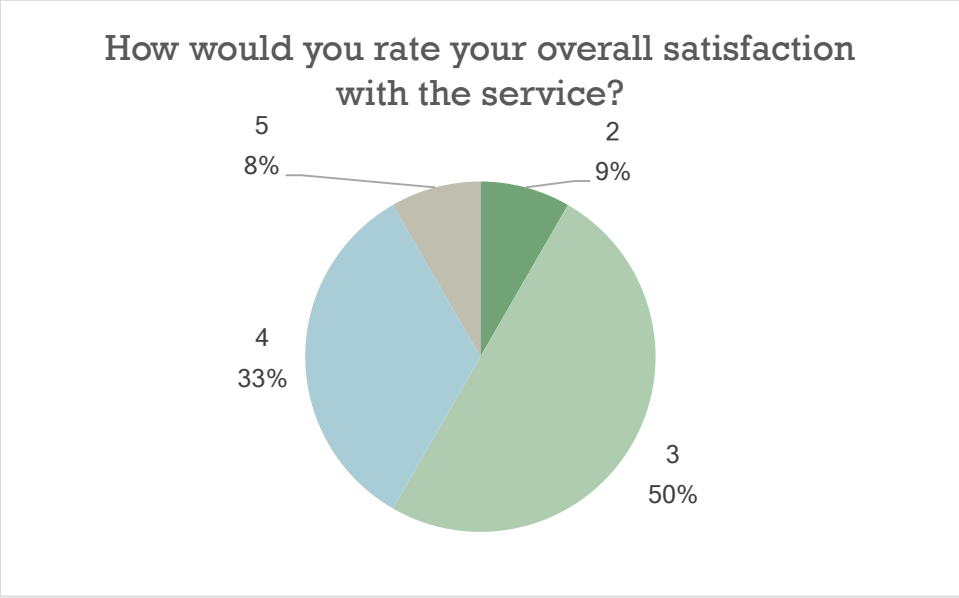

Figure 21: Overall satisfaction on Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

This question was somehow more difficult to be answered; in general, the response from the questionnaires is positive (3 or more) in 83% of answers.

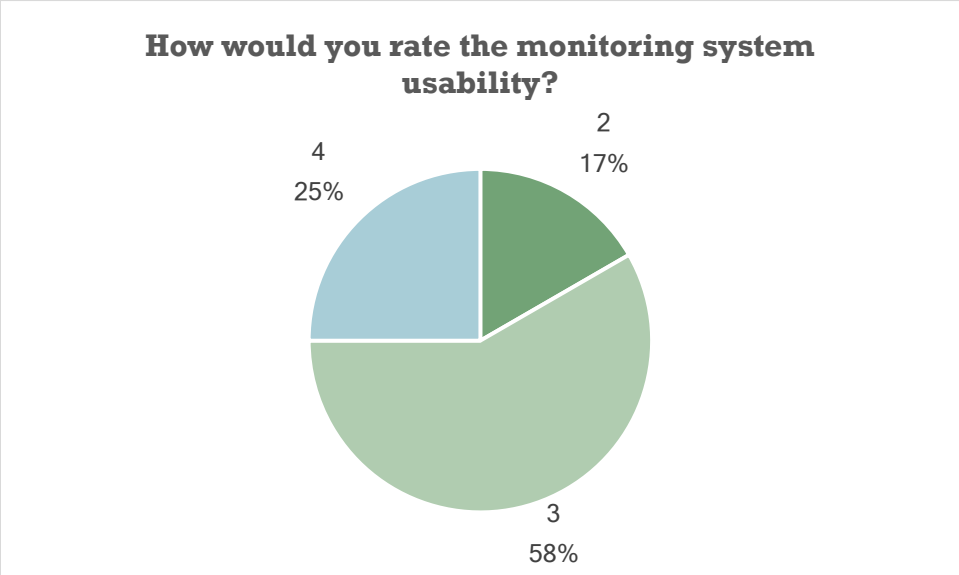

Figure 22: Monitoring system of Clearò in the opinion of service's subscribers. (Source: Online questionnaire)

The last question asked for a declaration of future renewal of the service. In this case, a less positive framework emerges, with "Definitely yes" and "Probably yes" rated for "only" 50%, and "I don't know" rated for 42%.

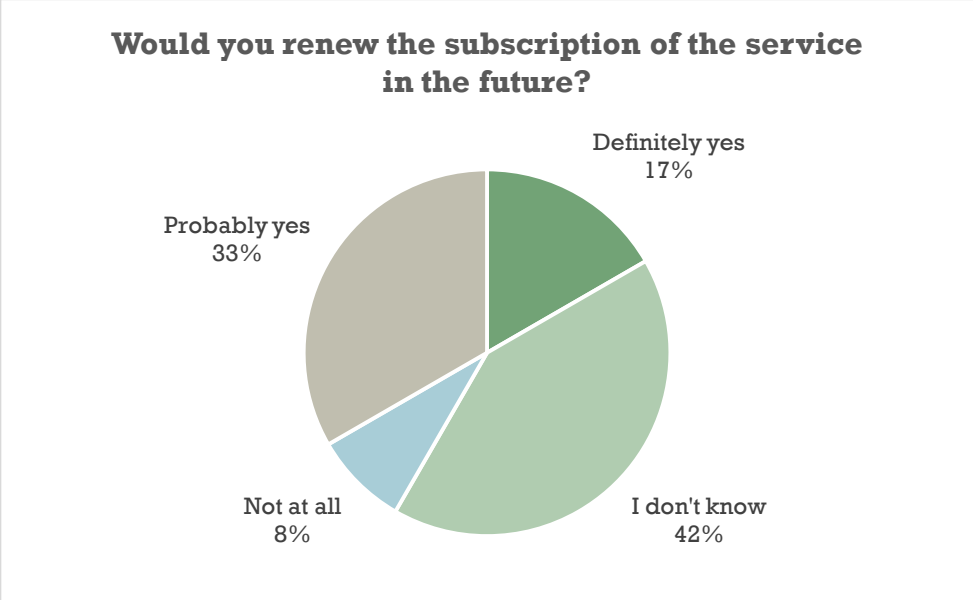

Figure 23: Future renewal of service's declarations by service's subscribers. (Source: Online questionnaire)

Evaluation and conclusions

The overall evaluation of the pilot for 2016 is abundantly positive. This is evident both from a technical and qualitative perspective (results from the KPIs look&feel evaluation, general users' experience).

The technical performance was very good (see also deliverable D4.2). Clearò did not experience any downs of the service (incidents), except for two cases of planned downs due to service maintenance. KPIs are positive in general, even if no actual (initial) references can be documented for this first year of service run on CloudOpting's platform. In 2017 the assessment will be verified again with higher robustness, thanks to the availability of historical series of data. For the assessment of the service by means of enterprise dashboards like Symon, some elements should be carefully elicited, in particular the impact of daily restart of the service and other variables explained above. The final, overall, rating is excellent.

From a qualitative perspective it is important to highlight the low number of answers collected (12 on 26 users) and the fact that Clearò is an institutional service offered by CSI itself to a number of elicited employees. Nonetheless the general feeling is positive. Assessment should be improved for 2017 with higher involvement of final users (citizens), for instance by means of an online questionnaire embedded in the service itself. In general terms, (internal) service subscribers declared that Clearò is a user friendly and high quality service. The technical support is also very good in the overall rating. The yearly trends of use are impacted by the institutional activities more than other variables (public tenders and other calls' publication), however the internal users appear loyal to the service.

3.2.4. Karlshamn

The Karlshamn pilot comprises two different services, hosted in the IMI cloud in Barcelona:

- Where can I charge my car (also known as Take Charge)
- POI/Next2Me

"Where can I charge my car" is the most representative service of the Karlshamn pilot. The "Where can I charge my car" evaluation is deeply described in next section while the POI/Next2Me service is included in the *Other Services* section.

4.2.4.1 Where can I charge my car?

Overview

The "Where can I charge my car" instance has been operating for 12 months during Y3 of the project. This section will describe the evaluation of the service during this period.
According to deliverable D4.1, the pilot will be considered a success if the design and development of the "Where can I charge my car" service is effectively migrated to the CloudOpting platform. The following visual criteria for the application should be met:

- Show the charging stations on the map
- Show the surrounding points of interest on the map
- Point on the map to get routing

In order to evaluate the reliability of the service, a set of KPIs were defined in deliverable 4.1. The corresponding metrics have been recorded to be evaluated at the end of the pilot phase.

Quantitative metrics and KPIs

The quantitative KPI and metrics gathered during the monitoring period are listed in Table 16. The associated data is the average figures, unless otherwise stated, taken from the reporting templates from all the evaluation period.

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 20%	2.5%
Memory usage	Averaged % Memory usage	Memory %	< 80	Not monitored
Stability	How stable is the service	Number of downs	2 per year	1
Availability	Uptime of the service	Number of downs Service uptime	90%	92%

Access Time	How much time need the service to attend a request	Access Time	[20- 800]ms	Not monitored
Service visits	Average number of users using the service per month	Visits sessions	[100 – 10000]	10,3
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	NA
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[20 – 50]	Not monitored
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[3 – 5]	NA

Table 16: Service metrics and KPI's for Where can I charge my car.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end users. These questionnaires can be found in Annex I.

All the qualitative metrics will be measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values [1-5]
Service usability	The perception of the usability of the service for the subscriber	Service usability	> 3 Stars	**3.25 Stars**
Customisation	Degree of customisation of the service	Service customisation	> 3 Stars	**2.75 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.25 Stars**
Satisfaction	The perception of the satisfaction of the service for the subscriber	Service satisfaction	> 3 Stars	**3.75 Starts**
Monitoring usability	The perception of the usability of the monitoring for the subscriber	Monitoring satisfaction	> 3 Stars	**3.75 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**3.75 Stars**

Table 17: Qualitative KPIs for the service subscriber.

The following chart represents the perception of the usability of the service for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team.

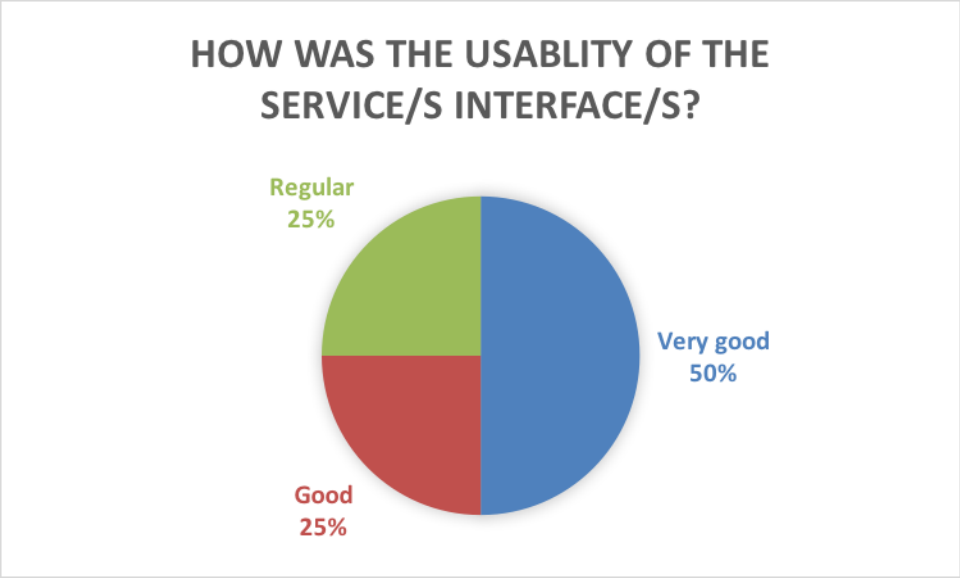

Figure 24: Service Subscriber usability.

The chart in figure 25 represents the degree of customisation of the service for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team.

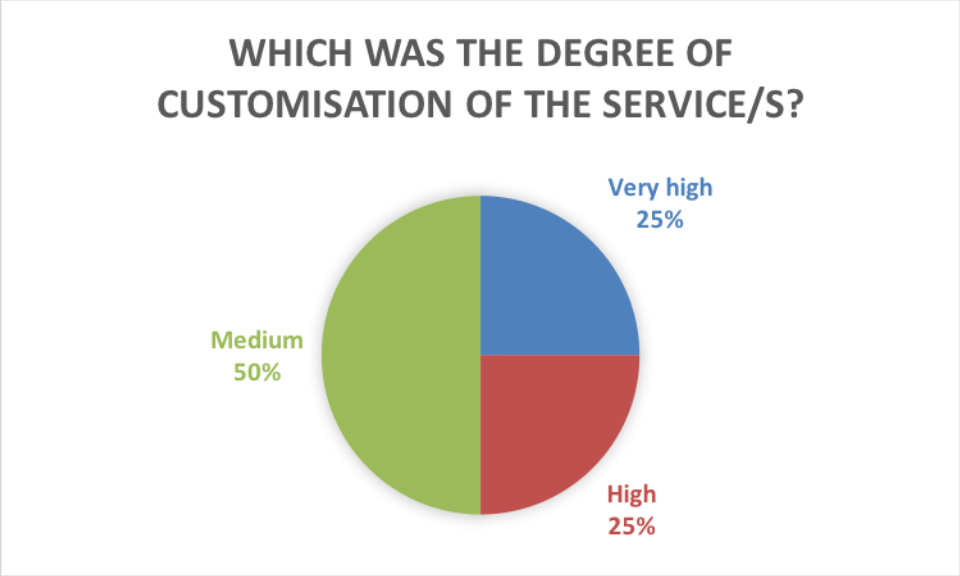

Figure 25: Service Subscriber customisation.

Figure 26 shows the perception of the support received for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team.

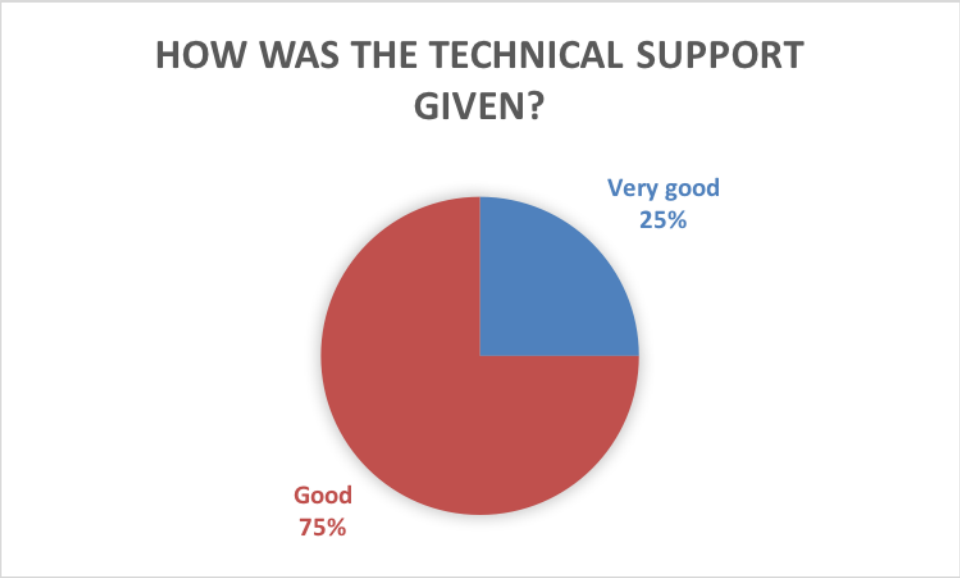

Figure 26: Service Subscriber support given.

The following chart represents the perception of the system downtime received for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team.

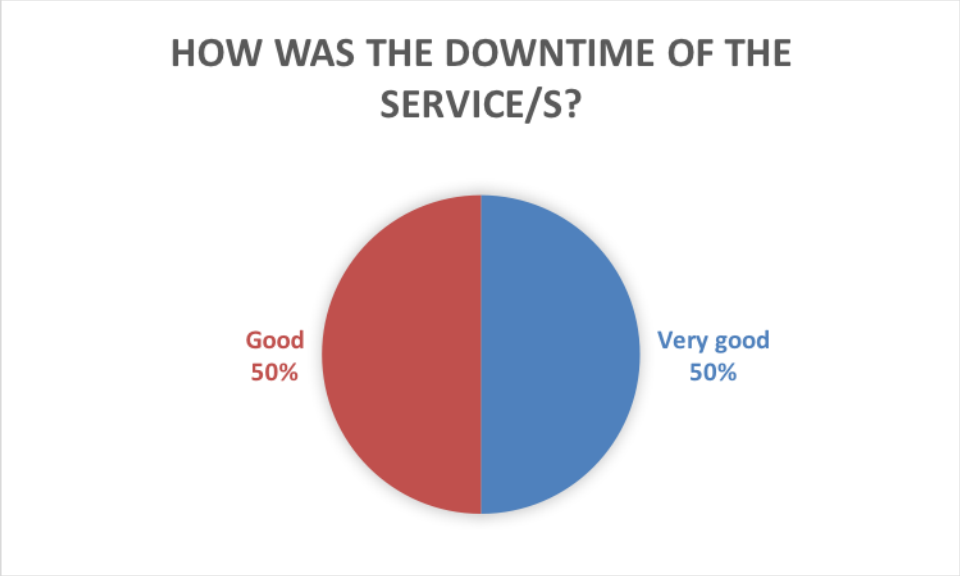

Figure 27: Service Subscriber downtime.

Figure 28 represents the perception of the satisfaction of the service for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team.

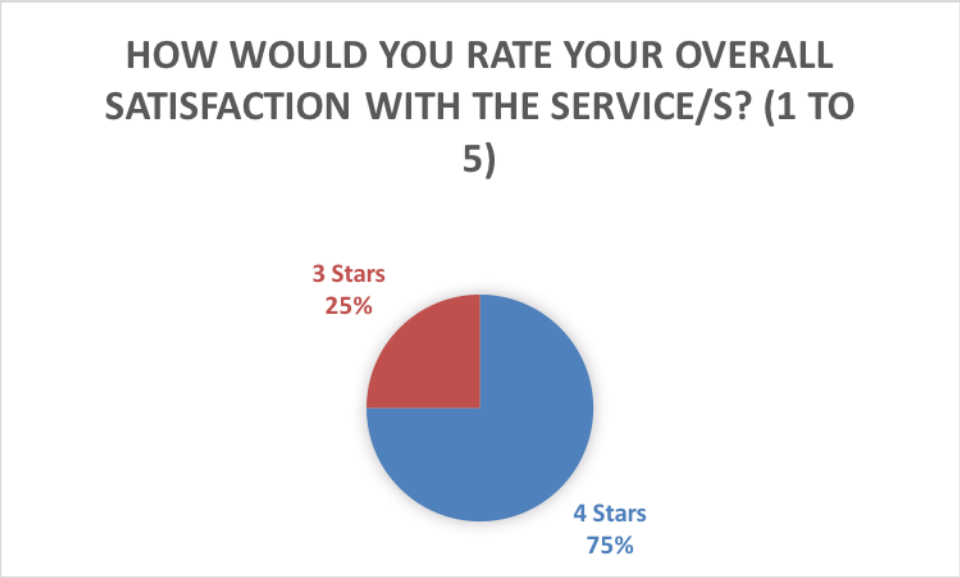

Figure 28: Service Subscriber satisfaction.

The perception of the usability of the monitoring for the subscriber gathered from the questionnaires delivered to the TakeCharge Service managing team is shown in Figure 29 below.

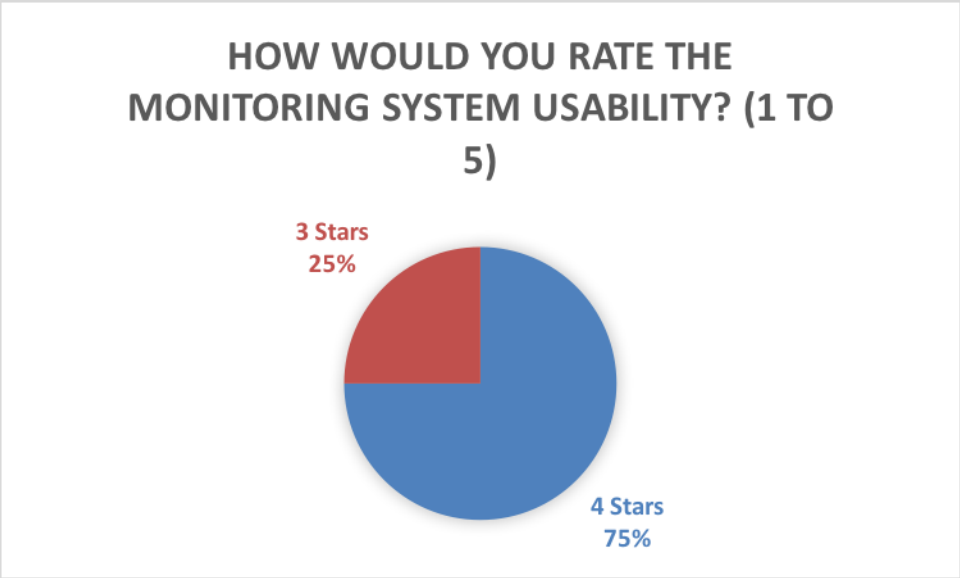

Figure 29: Service Subscriber monitoring.

The last question asked for a declaration of future renewal of this service from a service subscriber perspective.

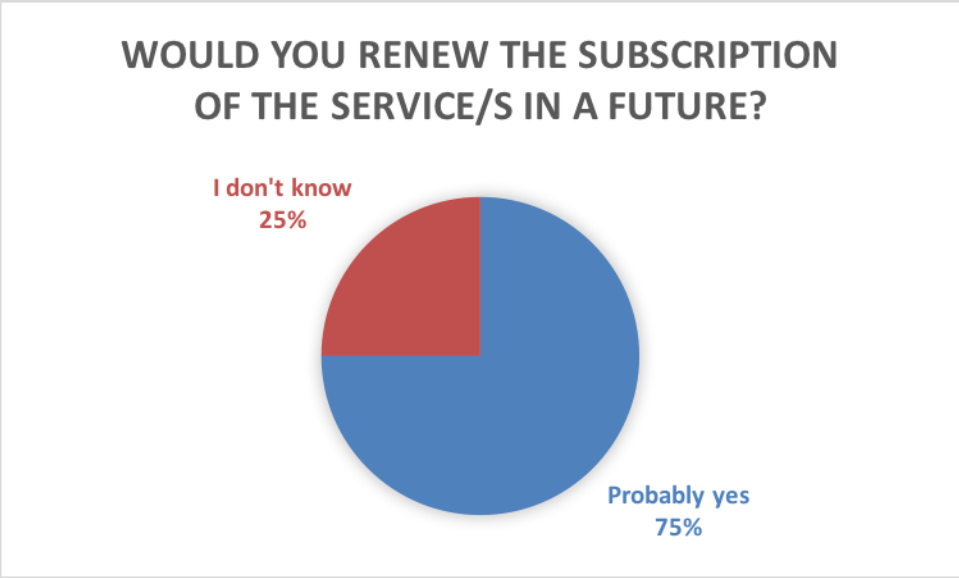

Figure 30: Service Subscriber renewal.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values [1-5]
Service usability	The perception of the usability of the service for the end user	Service usability	> 3 Stars	**3.4 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.6 Stars**
Rate	Rating of the services compared to others	Rating	> 3 Stars	**3.75 Stars**
Satisfaction	The perception of the satisfaction of the service for the end user	Service satisfaction	> 3 Stars	**3.5 Starts**

Table 18: Qualitative KPIs for the end user.

The chart in figure 31 represents the perception of the usability of the service for the end user gathered from the questionnaires.

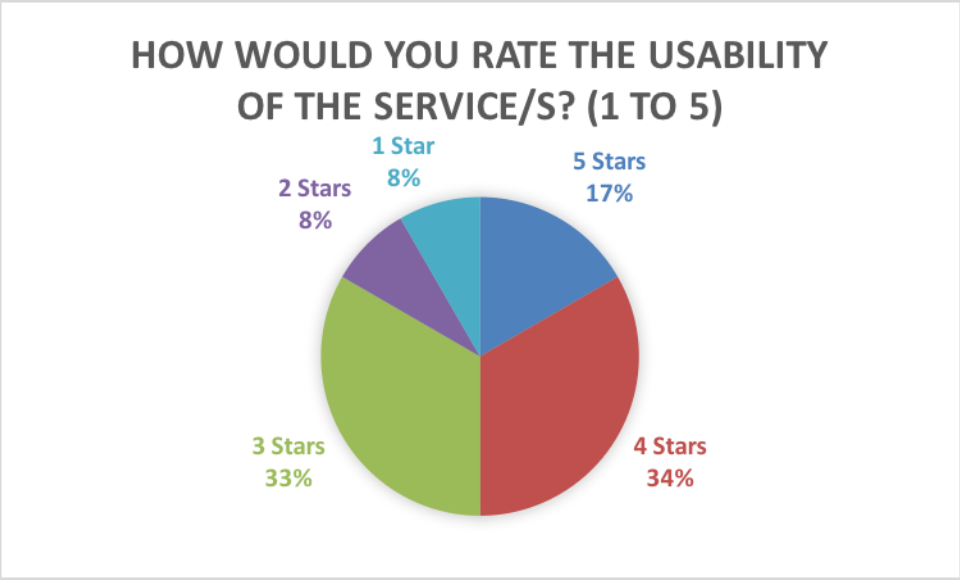

Figure 31: End User usability.

The following chart represents the perception of the support received for the end user gathered from the questionnaires.

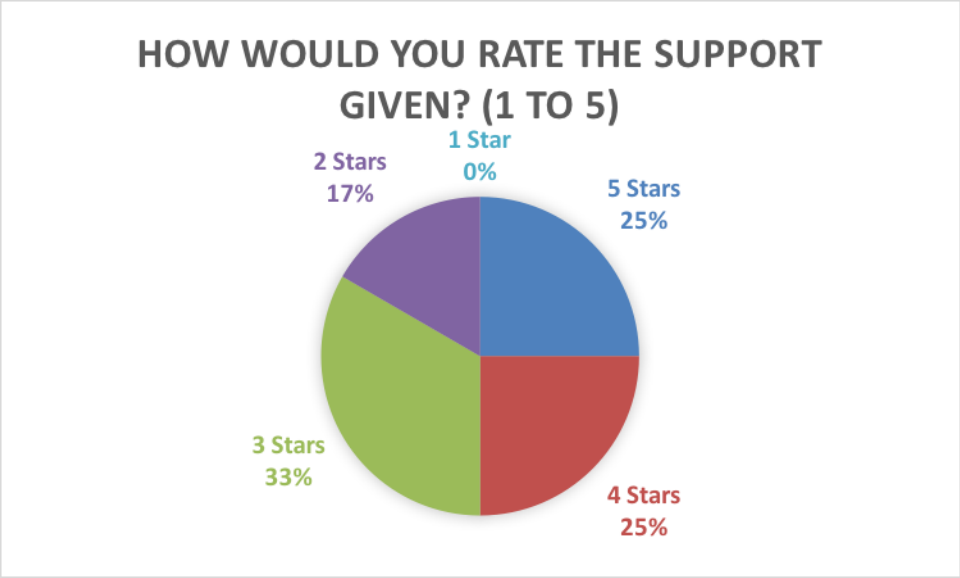

Figure 32: End User support.

Figure 33 represents the rating of the services compared to others for the end user gathered from the questionnaires.

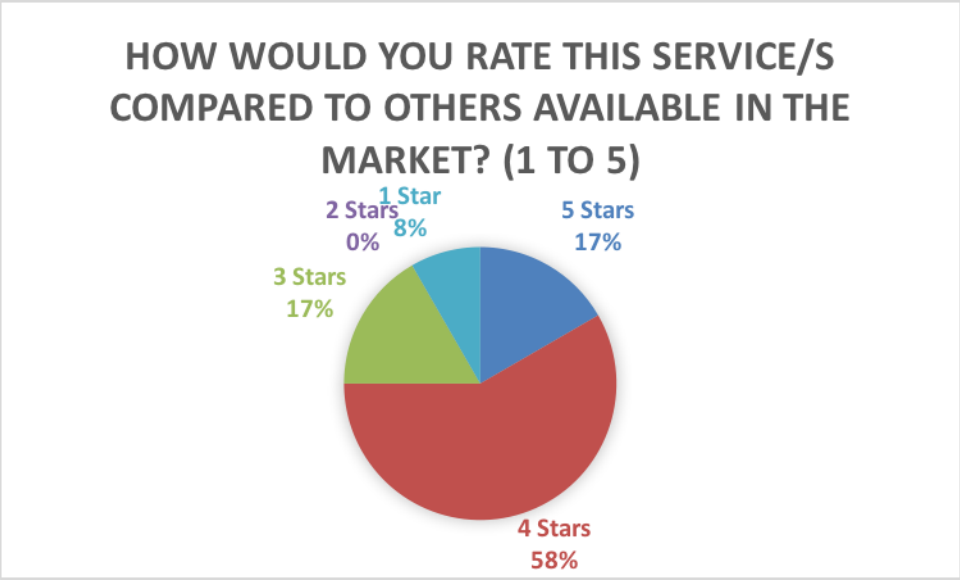

Figure 33: End User rate against others.

The following chart represents the perception of the satisfaction of the service for the end user gathered from the questionnaires.

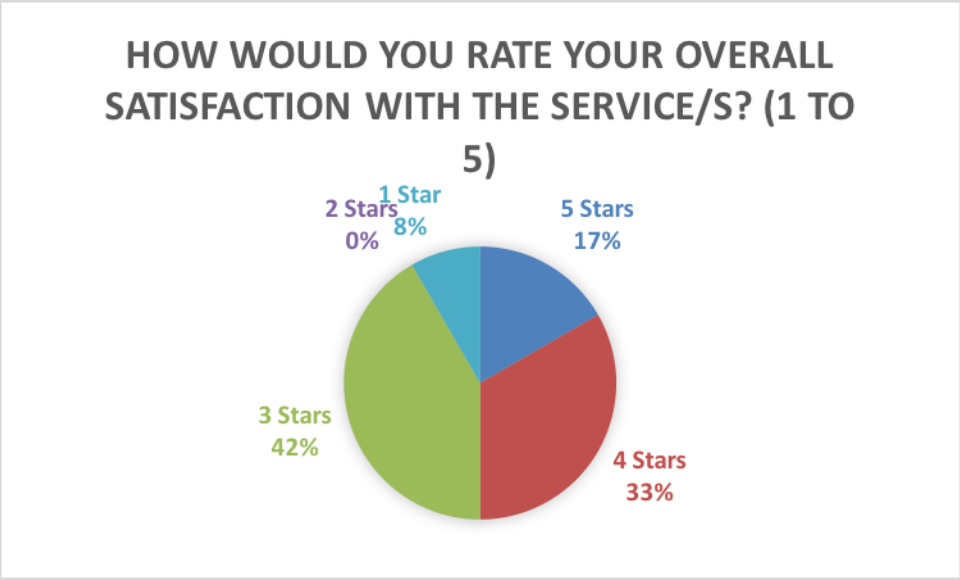

Figure 34: End User satisfaction.

Evaluation and conclusions

According to the tables above, our conclusion is that all KPIs have been achieved but customisation. The usage of the service did not reach a level where platform limitations could be expected or tested. However, considering that Karlshamn is a small city (about 30 000 citizens), not many of them having an electric car yet, the usage figures should be considered realistic anyway. The customisation level has been one of the identified items to improve the service. Considering the low CPU usage (2.5%) the size of the virtual machine instance may be decreased to reduce cost, without affecting end user experience.

The conclusion of evaluating the technical and quality metrics for the service, is that end user experience is well within expected goals. Regarding system administration, some basic knowledge of Docker is required by the system administrator for the maintenance procedures such as dataset updates. This could have been avoided or improved by including an administrator interface to the service application, but on the other hand this also demonstrates the flexibility of using Docker technology in the CloudOpting platform. Another area of improvement is the data source intergration, whereas current solution is based on administrator procedures to import new datasets, but the preferred solution should be to have much more tight data source integration using automtic procedures. Improving the data quality should also improve the end user attraction to the service.

The end users appreciated the application concept providing route planning for electrical vehicles, combined with things to do while recharging. This concept currently seems to be fairly unique, but as the market for electrical vehicle now grows extensively, similar services a very likely to pop up.

The final conclusion of the overall evaluation, is that the deployment and behaviour of the service running on the CloudOpting platform has been successful, only a few improvements to the service itself regarding the end user experience and data source handling as elaborated above could be suggested.

3.3. Other services

Apart from the representatives services aforementioned, the rest of services deployed in the pilots have been tested and monitored to enrich the knowledge base with new lessons learned about interaction between applications and the platform.

3.3.1. Sentilo

Sentilo is an open source sensor and actuator platform. During the deployment, a dedicated Sentilo instance was created for the CloudOpting project.

Overview

Sentilo has been operated since the first month of the monitoring period. Together with CKAN, the service has been fully operated from Barcelona and its operation department. During this period, an agreement with the Sentilo operator and the *Servei Meteorologic de Catalunya* has been reached to create a subscription in order to publish the meteorological data from their sensing stations. This agreement delayed the publication of data for 2 months, time used to manage users, sensors and configure the instance final look and feel.

The general KPIs described in D4.1 were adapted to better describe the performance of the service during this period and are reported in the next section. Thus, *Requests processed* and *Active Sensors* have been included to the general table.

- *Requests processed*: Read and Write requests to sensors data.
- *Active Sensors*: Number of sensors registered.

Quantitative metrics and KPIs

The accumulated KPI's reported during this period are shown in the table below:

KPI	Description	Associated Metrics	**Reported Value**
CPU usage	Average % CPU usage	CPU %	**<20%**
Memory usage	Averaged % Memory usage	Memory %	**30%(4GB)**
Stability	How stable is the service	Number of downs	**1**
Availability	Uptime of the service	Service uptime	**99.5%**
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	**>25**

Access Time	How much time need the service to attend a request	Access Time	**<2000ms**
Requests processed	Number of requests processed since service deployment	Requests processed	**5861**
Active sensors	Number of active sensors	Number of sensors	**34**

Table 19: Service metrics and KPIs.

The following image shows the Barcelona map and a detail of the X4 component sensors including data humidity, wind speed and rain among others.

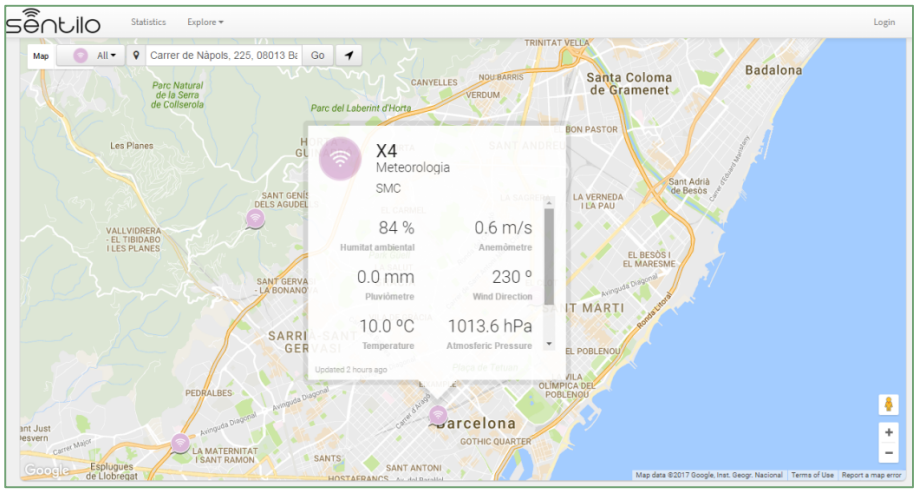

Figure 35: X4 detail component on the Barcelona map.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service subscriber satisfaction	The perception of the quality of the service for the subscriber	Service satisfaction	> 3 Stars	**4.5 Stars**
Support	The perception of the quality of the support	Quality of Technical Support	> 3 Stars	**3.75 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 20: Qualitative KPIs for the service subscriber.

The following chart shows the service subscriber satisfaction gathered from the questionnaires delivered to the CKAN managing team on the Barcelona City Council.

Figure 36: Service Subscriber satisfaction.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
End users satisfaction	Overall end user satisfaction	Satisfaction rating	> 3 Stars	**4.25 Stars**
Service usability	The perception of end user about the usability of the service	Usability rating	> 3 Stars	**4.25 Stars**
Support	The perception of end user about the support received	Support rating	> 3 Stars	**3.5 Stars**
Service comparison	End user service comparison rating	Comparison rating	> 3 Stars	**3.75 Stars**

Table 21: Qualitative KPIs for the end user.

The next figure shows the end user satisfaction results.

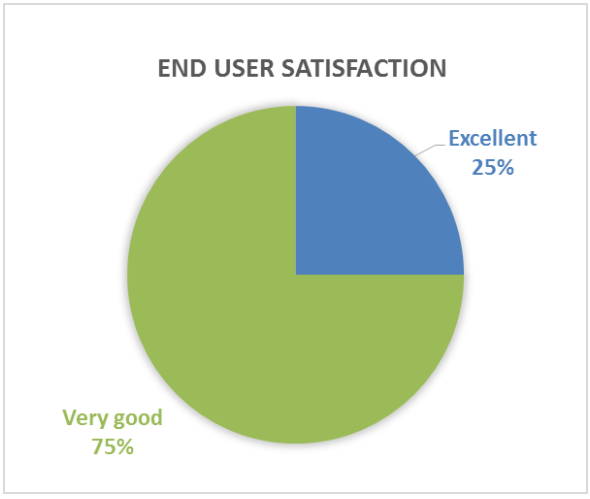

Figure 37: End User satisfaction.

Evaluation and conclusions

- Deployment and Customisation:

The deployment of the service created a fully working instance, but some customisation improvements could be included in the following software versions (by default, the displayed map is centred in Barcelona).

- Operation and KPI:

The management of the instance became simpler than expected during this period. Only a few issues regarding the cloud management affected the service. Thus, thanks to using Docker containers it was easy to restore the service after it's down.

Regarding the technical indicators, the service supported the more than 6000 request for 34 sensors during this period. Taking into account that the CPU and memory use has been below 50% during all period, the VM size may be decreased to reduce costs. Docker containerisation has been demonstrated to be very efficient in terms of use of resources.

- User experience:

The user experience with the Sentilo instance deployed with CloudOpting has been transparent and identical to the one when using standard technologies.

On the other hand, the system administrators require having some basic knowledge about the Docker basics in order to administrate the system.

Thus, it is proved the viability to use Docker as core technology for the CloudOpting platform.

As a conclusion of the overall evaluation, the deployment and behaviour of the service included in the platform has been very positive.

3.3.2. Next2Me

Next2Me is a mobile service that allows citizens to find services and events occurring around him. The objective was to develop a mobile solution that looks for nearby services and public facilities, allowing the user to filter them by category of the service/event or to look for a specific one. The application provides a list of services and facilities, geo-location of the user and a map of the services and facilities.

Overview

Next2Me has been operated for 12 months during the Y3 of the project. The service is part of the Barcelona pilot which includes a set of other services. This section will describe the evaluation of the service during this period.

Next2Me has a scheduled job that checks every night for new data in the open cloud of Barcelona, called Barcelona Open Data, and retrieves this new datasets. After getting the data it is transformed and filtered by the platform, then it is loaded to the application database, and finally it is indexed by Apache Solr.

The KPIs described for "Other Services" in the D4.1 were adapted in order to describe in a better way the performance and the operation. A new KPI has been added:

- New data loaded: Amount of data downloaded from the Barcelona Open Cloud in MB.

Quantitative metrics and KPIs

In order to be able to evaluate this service we show the accumulated KPI's of all months:

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	0.8%
Memory usage	Averaged % Memory usage	Memory %	< 80	50%
Stability	How stable is the service	Number of downs	[0 – 20] per year	4
Availability	Uptime of the service	Number of downs Service uptime	90%	94%
Access Time	How much time need the service to attend a request	Access Time	[20- 4000]ms	20 ms
Service visits	Number of users using the service	Visits sessions	[100 – 10000]	<100

Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	Not monitored
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[50 – 100]	>100
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[20 – 200]	1 per dey running

Table 22: Service metrics and KPIs.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end uses. These questionnaires can be found in Annex I.

All these qualitative metrics will be measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the subscriber	Service usability	> 3 Stars	**4.33 Stars**
Customisation	Degree of customisation of the service	Service customisation	> 3 Stars	**2.33 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.66 Stars**
Satisfaction	The perception of the satisfaction of the service for the subscriber	Service satisfaction	> 3 Stars	**4.5 Starts**
Monitoring usability	The perception of the usability of the monitoring for the subscriber	Monitoring satisfaction	> 3 Stars	**3.33 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 23: Qualitative KPIs for the service subscriber.

The following chart represents the perception of the usability of the service for the subscriber gathered from the questionnaires delivered to the Next2Me Service managing team.

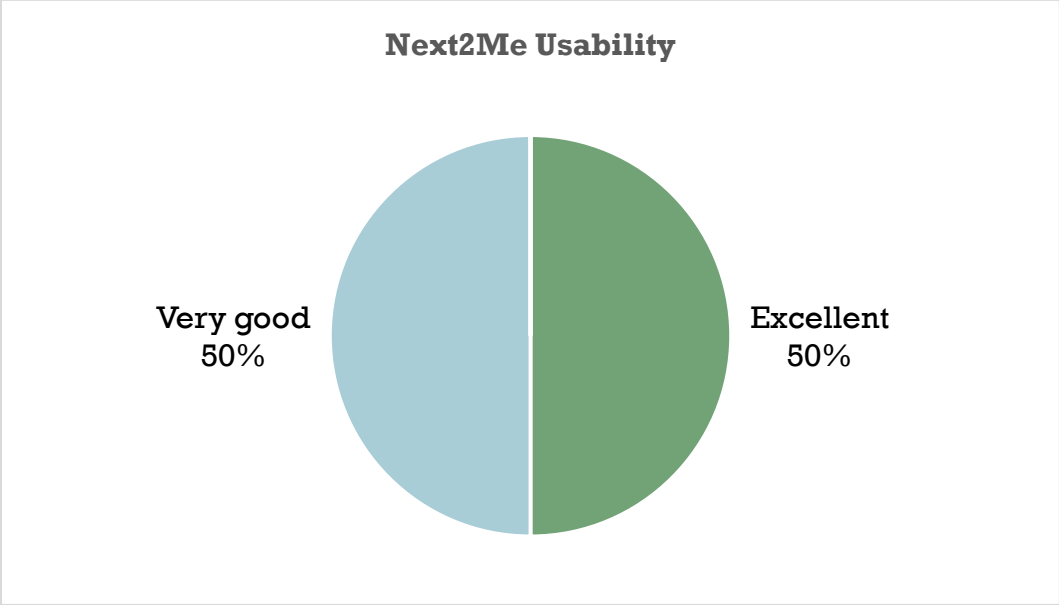

Figure 38: Next2Me – Service Subscriber Usability.

Next chart represents the degree of customisation of the service for the subscriber gathered from the questionnaires delivered to the Next2Me Service managing team. As depicted in figure 39, the satisfaction in customisation is below expectations.

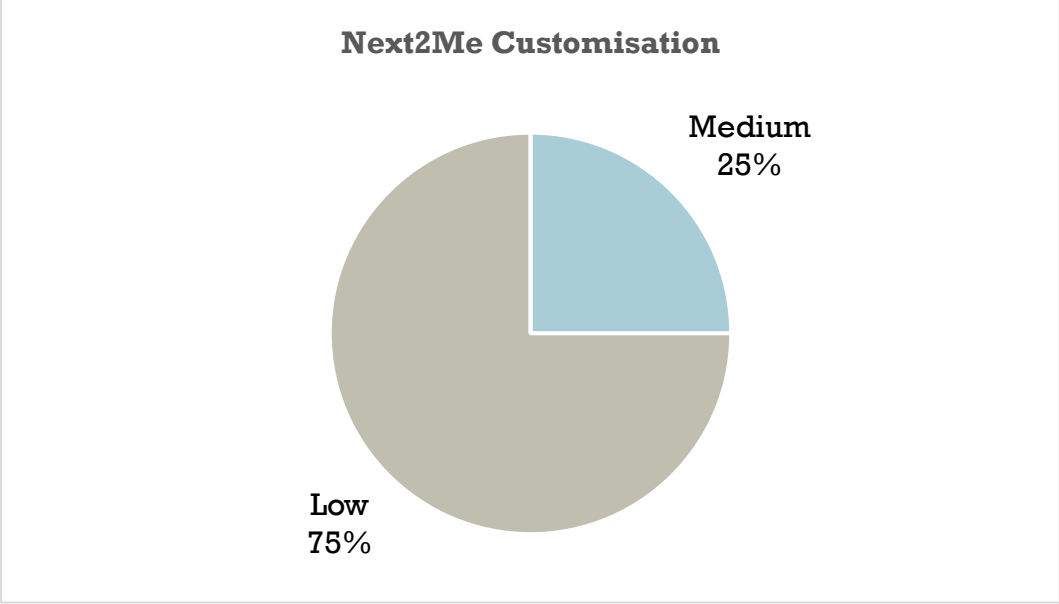

Figure 39: Next2Me – Service Subscriber customisation.

Chart in figure 40 represents the perception of the support received for the subscriber gathered from the questionnaires delivered to the Next2Me Service managing team.

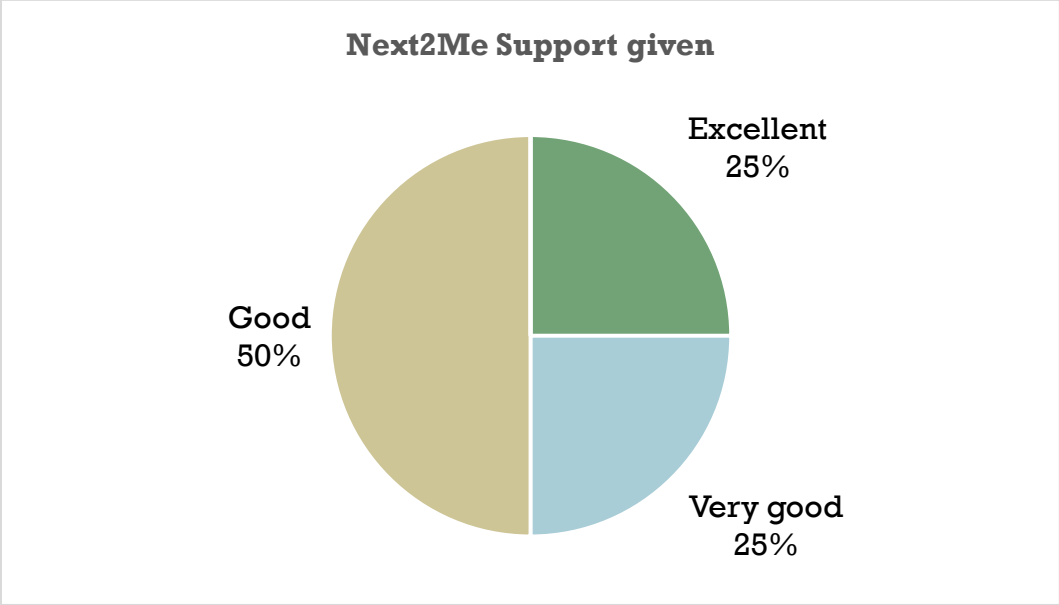

Figure 40: Next2Me – Service Subscriber support given.

The following chart represents the perception of the satisfaction of the service for the subscriber gathered from the questionnaires delivered to the Next2Me Service managing team.

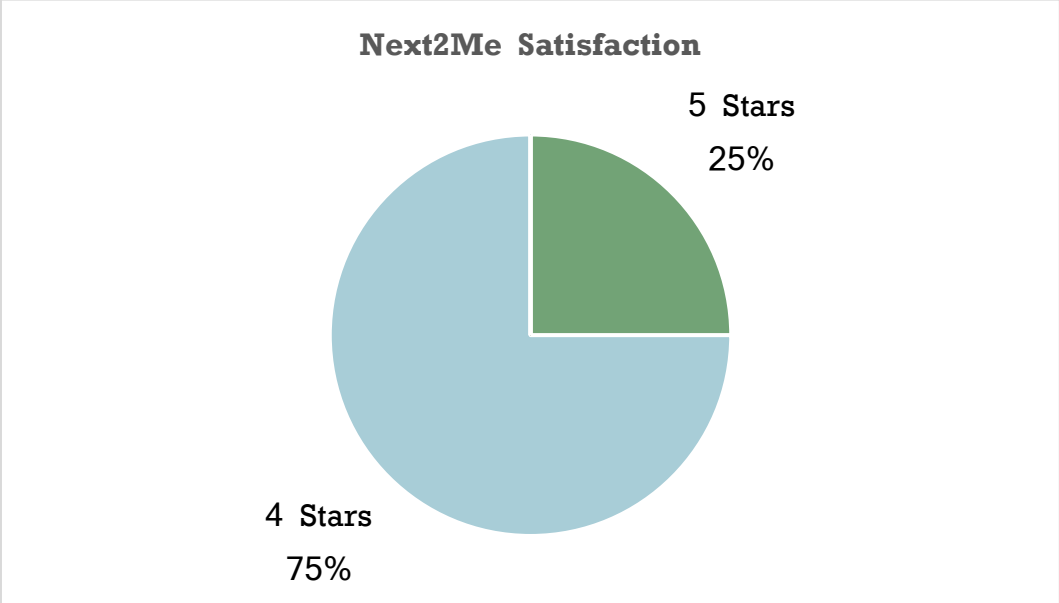

Figure 41: Next2Me – Service Subscriber satisfaction.

Next chart represents the perception of the usability of the monitoring for the subscriber gathered from the questionnaires delivered to the Next2Me Service managing team.

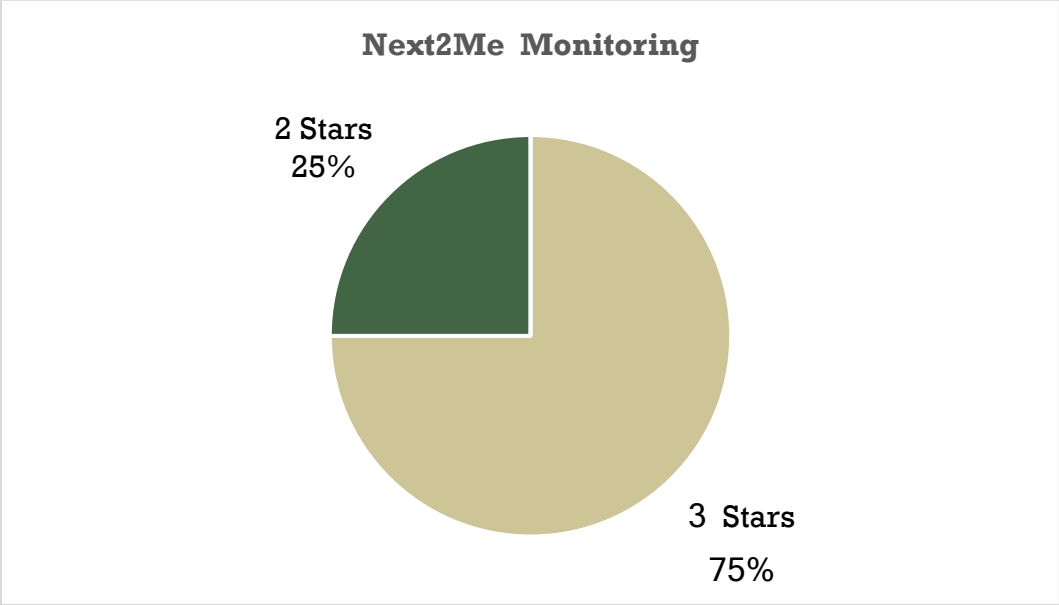

Figure 42: Next2Me – Service Subscriber monitoring.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the end user	Service usability	> 3 Stars	**4.09 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.2 Stars**
Rate	Rating of the services compared to others	Rating	> 3 Stars	**4.45 Stars**
Satisfaction	The perception of the satisfaction of the service for the end user	Service satisfaction	> 3 Stars	**4.18 Starts**

Table 24: Qualitative KPIs for the end user.

Chart in figure 43 represents the perception of the usability of the service for the end user gathered from the questionnaires delivered to the Next2Me Service managing team.

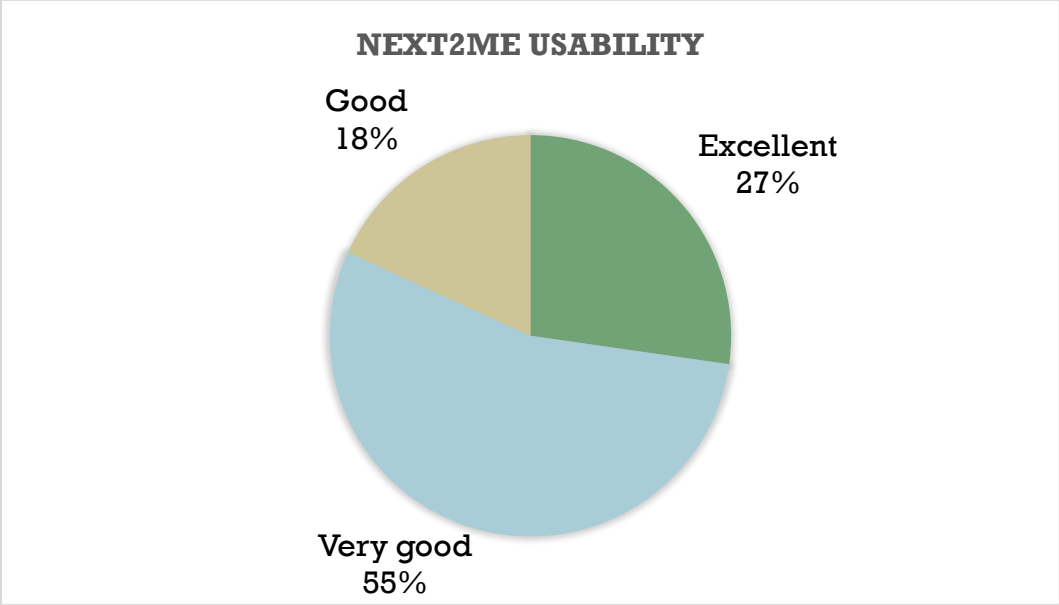

Figure 43: Next2Me – End User usability.

The perception of the support received for the end user gathered from the questionnaires delivered to the Next2Me Service managing team is shown below.

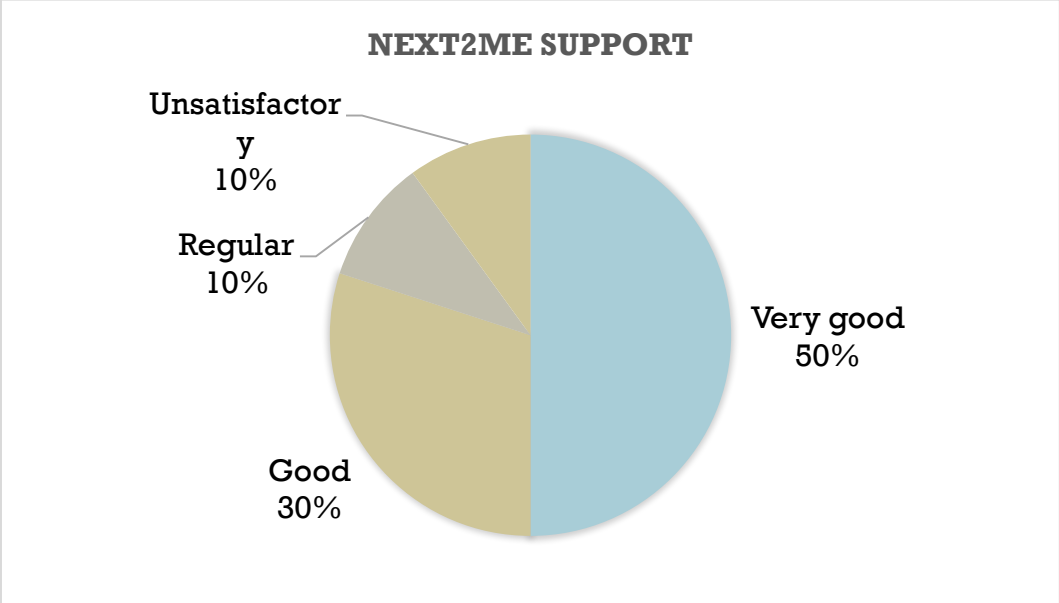

Figure 44: Next2Me – End User support.

The following chart represents the rating of the services compared to others for the end user gathered from the questionnaires delivered to the Next2Me Service managing team.

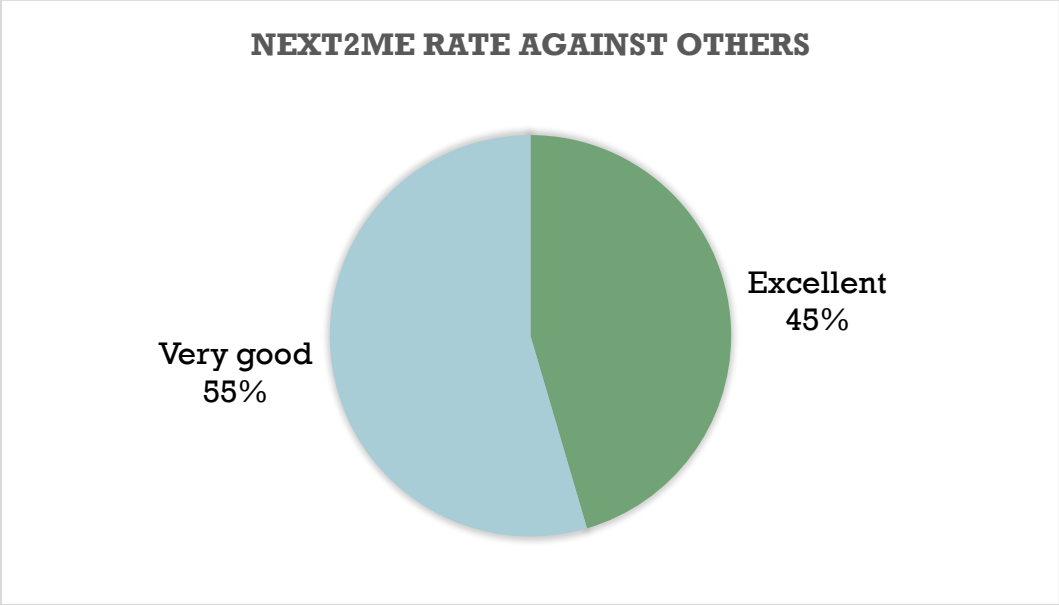

Figure 45: Next2Me – End User rate against others.

Below, a chart shows the perception of the satisfaction of the service for the end user gathered from the questionnaires delivered to the Next2Me Service managing team.

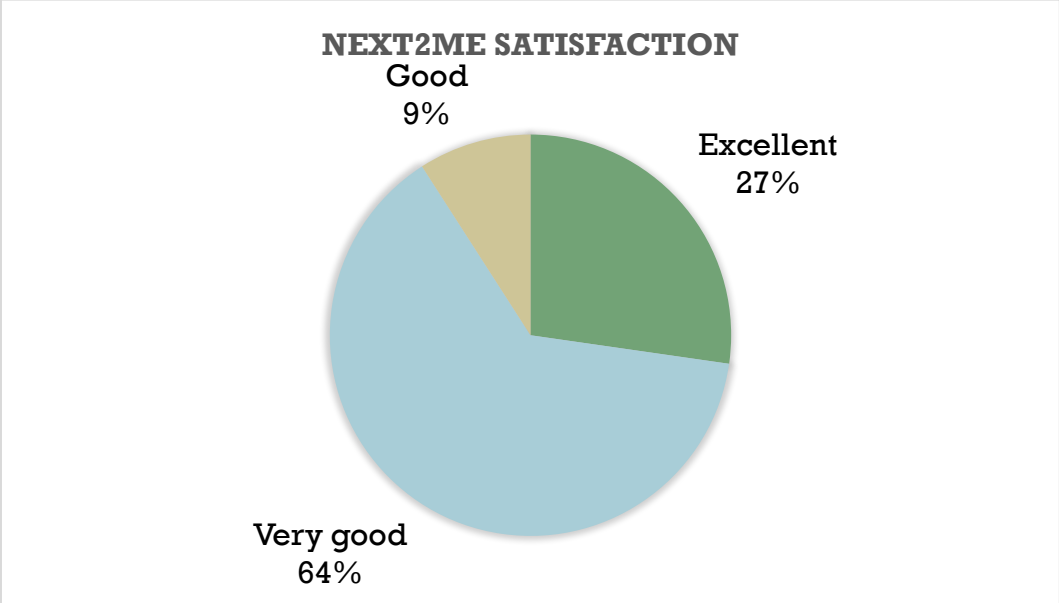

Figure 46: Next2Me – End User satisfaction.

Evaluation and conclusions

- Deployment and Customisation:

The deployment of the service created a fully working instance, but some customisation improvements could be included in the following software versions as the users stated in the questionnaire. The need to change or add new data sources, change colors and logos in order to fit better the city are some aspects that are on top of the table for imprvement.

- Operation and KPIs

The operation has been easy to perform, thanks to the Docker containers technology and the developed job to load data automatically to the database.

Although some downs were reported, only one of them was due to a technical issue of the application. It ran out of space because of the amout of data loaded from the Barcelona Open Data.

According to the tables from the "Results / KPIs", all KPIs have been achieved but the number of visits, impacting in the number of POI's consulted.

The new data loaded into the platform does not have a POI associated, so there are not a quantitative way to show its impact. To calculate the amount of new data loaded the difference of initial and final data volume in database has been performed, doubling the information from the begining.

- User experience:

All graphs from the user experience questionnaires demonstrate that the user experience for Next2Me has been very positive. At least the 91% of the people that answered the quetionnaire says it is very good or excelent.

With all this information it can be concluded that the results, as not compile what expected in some parameters, was sucessful and apreciated by the users.

3.3.3. Fix This

FixThis is a mobile application for easing the interaction among citizens and the city council. With FixThis citizens can report geo-located incidences with their smartphones, which are automatically sent to the incidents systems of the Public Entity for reporting, and then take corrective actions.

The application provides:

- Aa list of possible incidences.
- The geo-location of the user to record of incidences sent by the user.
- Status of the incidences.

Overview

Fix This has been operated for 12 months during the Y3 of the project. The service is part of the Barcelona pilot which includes a set of other services. This section will describe the evaluation of the service during this period.

Quantitative metrics and KPIs

In order to be able to evaluate this service we show the accumulated KPI's of all months:

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	0.8%
Memory usage	Averaged % Memory usage	Memory %	< 80	5%
Stability	How stable is the service	Number of downs	[0 – 20] per year	2
Availability	Uptime of the service	Number of downs Service uptime	90%	98%
Access Time	How much time need the service to attend a request	Access Time	[20- 4000]ms	20 ms
Service visits	Number of users using the service	Visits sessions	[100 – 10000]	<100
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	Not monitored
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[50 – 100]	>100
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[20 – 200]	N/A

Table 25: Service metrics and KPIs.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end uses. These questionnaires can be found in Annex I.

All these qualitative metrics will be measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the subscriber	Service usability	> 3 Stars	**4.2 Stars**
Customisation	Degree of customisation of the service	Service customisation	> 3 Stars	**2.2 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.4 Stars**
Satisfaction	The perception of the satisfaction of the service for the subscriber	Service satisfaction	> 3 Stars	**4.6 Starts**
Monitoring usability	The perception of the usability of the monitoring for the subscriber	Monitoring satisfaction	> 3 Stars	**2.6 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 26: Qualitative KPIs for the service subscriber.

The following chart represents the perception of the usability of the service for the subscriber gathered from the questionnaires delivered to the FixThis Service managing team.

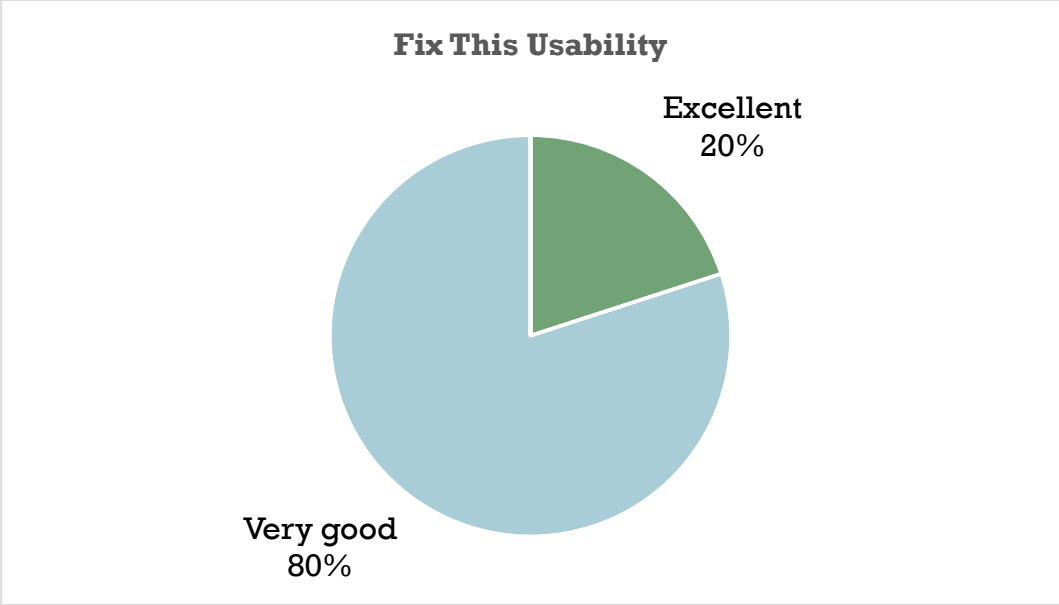

Figure 47: FixThis – Service Subscribers usability.

Next chart represents the degree of customisation of the service for the subscriber gathered from the questionnaires delivered to the FixThis Service managing team. The customisation satisfaction does not fit the desired values.

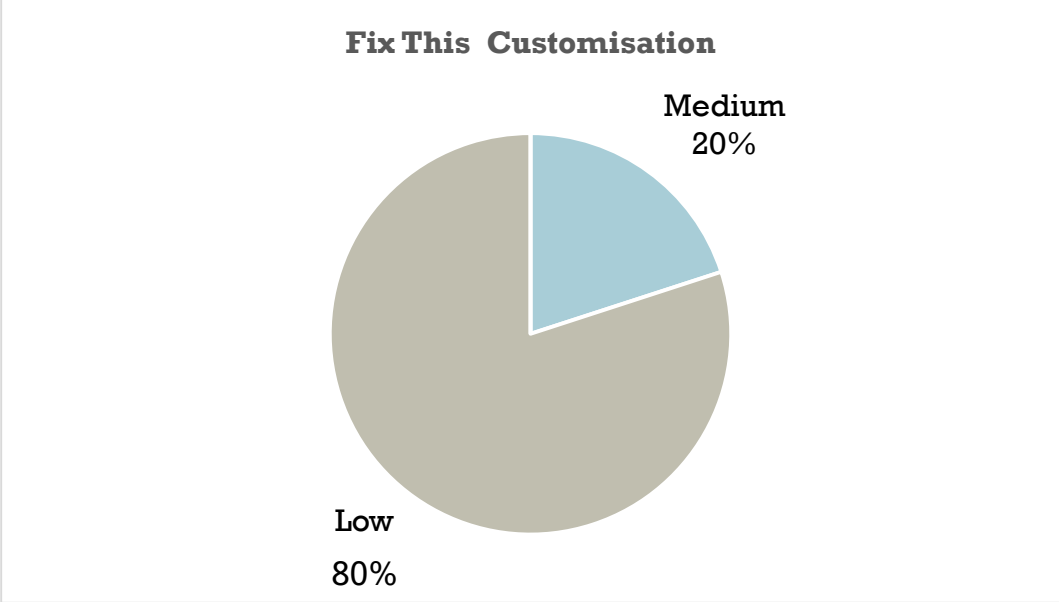

Figure 48: FixThis – Service Subscribers customisation.

Below, a chart represents the perception of the support received for the subscriber gathered from the questionnaires delivered to the FixThis Service managing team.

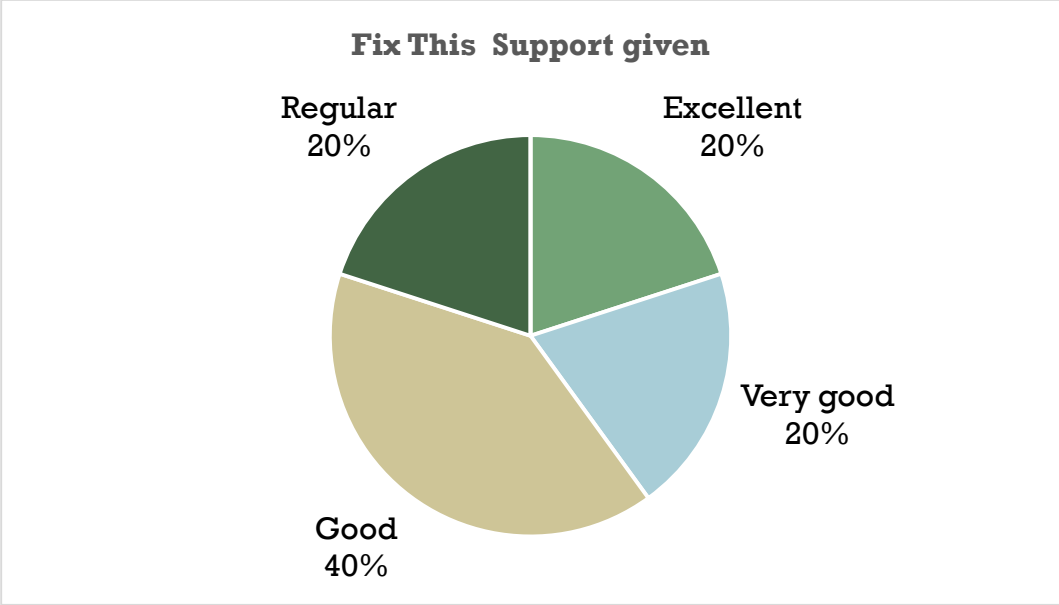

Figure 49: FixThis – Service Subscribers support given.

In Figure 50, a chart represents the perception of the satisfaction of the service for the subscriber gathered from the questionnaires delivered to the FixThis Service managing team.

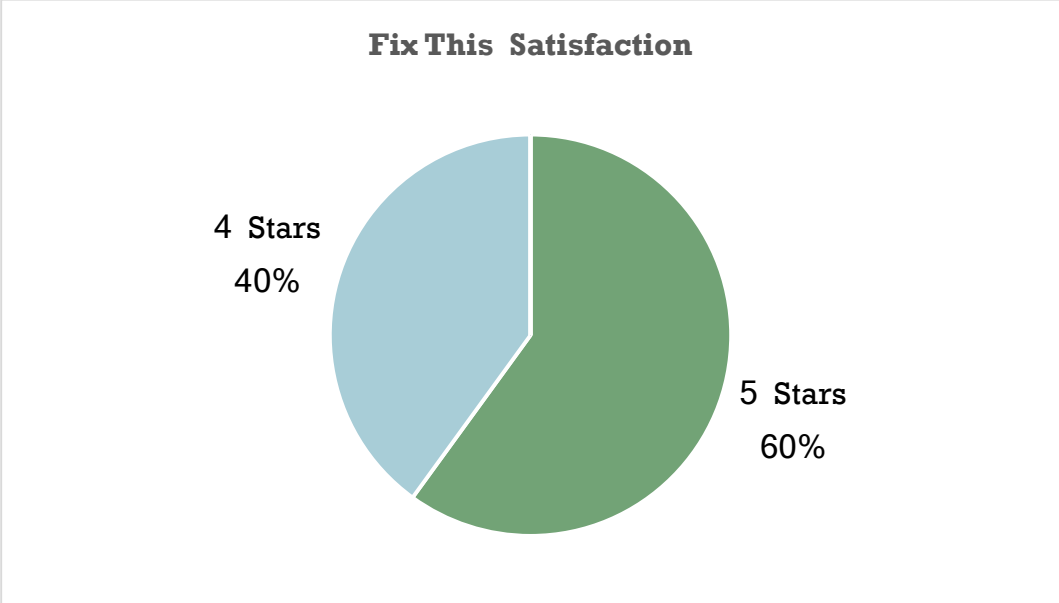

Figure 50: FixThis – Service Subscribers satisfaction.

The following chart represents the perception of the usability of the monitoring for the subscriber gathered from the questionnaires delivered to the FixThis Service managing team.

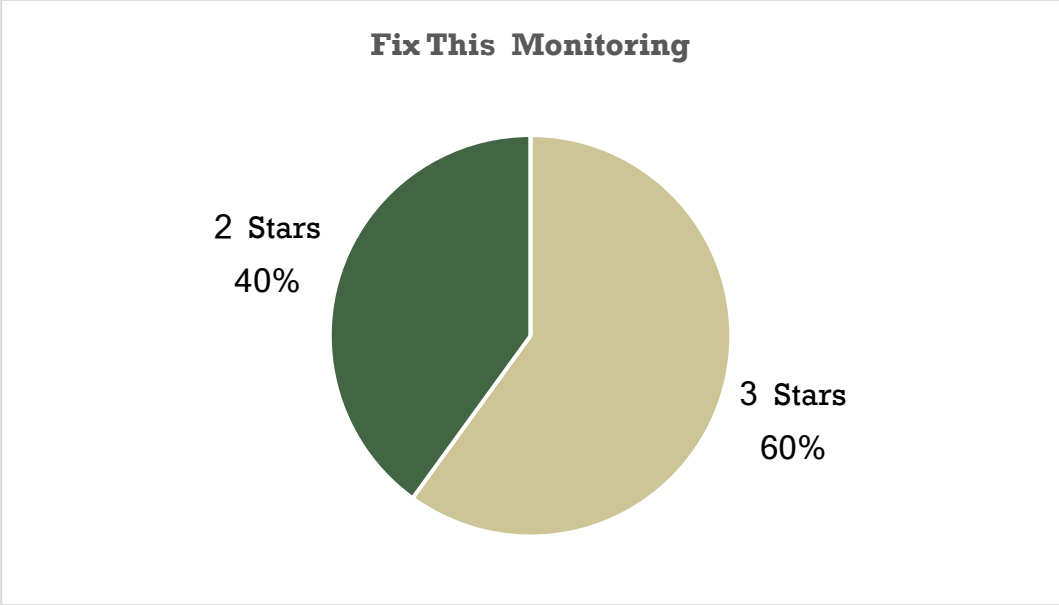

Figure 51: FixThis – Service Subscribers monitoring.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the end user	Service usability	> 3 Stars	**4.12 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.33 Stars**
Rate	Rating of the services compared to others	Rating	> 3 Stars	**4.14 Stars**
Satisfaction	The perception of the satisfaction of the service for the end user	Service satisfaction	> 3 Stars	**4 Starts**

Table 27: Qualitative KPIs for the end user.

The chart in Figure 52 below represents the perception of the usability of the service for the end user gathered from the questionnaires delivered to the FixThis Service managing team.

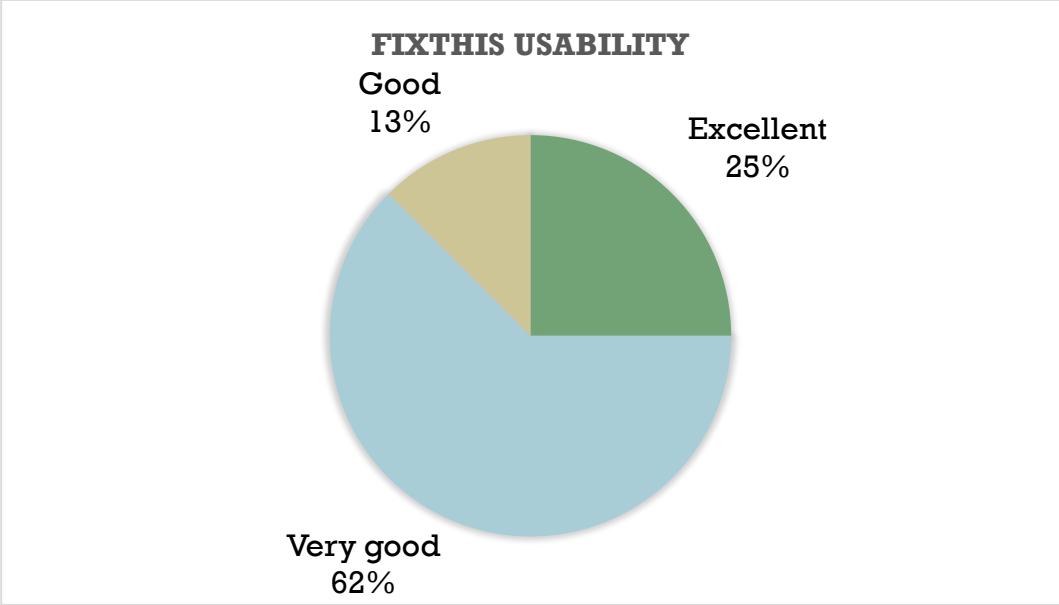

Figure 52: FixThis – End User usability.

Below, the perception of the support received for the end user gathered from the questionnaires delivered to the FixThis Service managing team is shown (Figure 53).

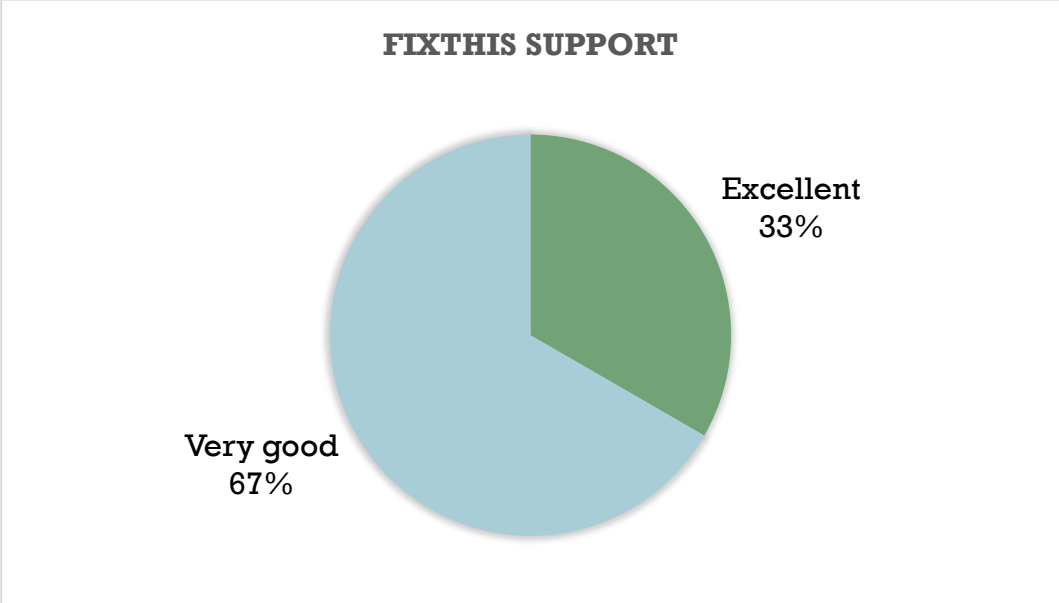

Figure 53: FixThis – End User support.

The following chart represents the rating of the services compared to others for the end user gathered from the questionnaires delivered to the FixThis Service managing team.

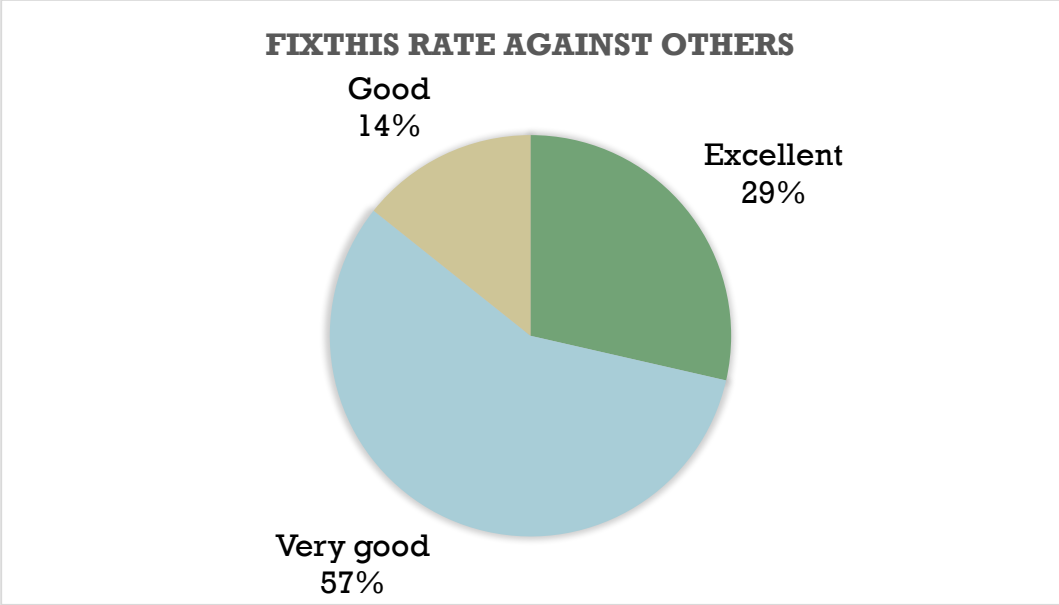

Figure 54: FixThis – End User rate against others.

Next chart represents the perception of the satisfaction of the service for the end user gathered from the questionnaires delivered to the FixThis Service managing team.

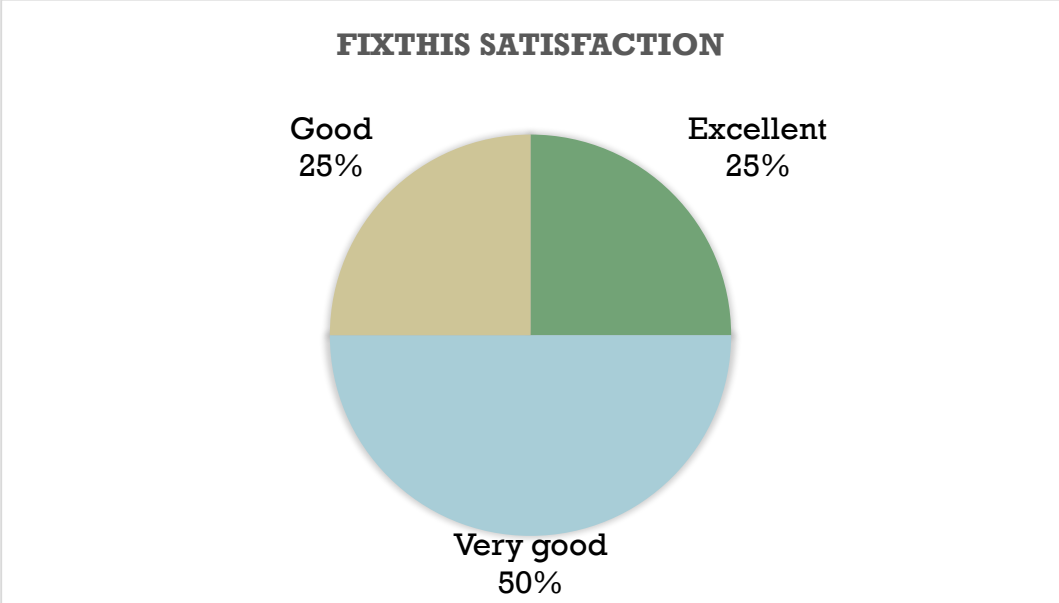

Figure 55: FixThis – End User satisfaction.

Evaluation and conclusions

- Deployment and Customisation:

The deployment of the service created a fully working instance, but some customisation improvements will be included in the following software versions to satisfy user expectatives. The need to change colors and logos in order to fit better the city are some aspects that are on top of the table for imprvement.

- Operation and KPIs

The operation has been easy to perform, thanks to the Docker containers technology that allows to start/stop/restart a very easy thing to manage.

According to the tables from the "Results / KPIs", all KPIs have been achieved but the number of visits, impacting in the number of POI's consulted.

- User experience:

All graphs from the questionnaires demonstrate that the user experience for FixThis has been very positive. At least the 75% of the people that answered the quetionnaire says it is very good or excelent.

With all this information, it can be concluded that the results, even taking into account that some parameters are lower than exected, was a sucessful service and so it was appreciated by the users.

3.3.4. City Agenda

City Agenda is a mobile service that updates the citizen with the schedule of events occurring around him. This service is a mobile service that aims to increase the citizen participation by keeping them updated with the schedule of the events. The objective is to develop a mobile solution that allows the users to look for cultural and leisure events –by theme– happening in Barcelona. The application provides a list of events occurring within an established distance, geo-location of the user, map of events, search tool for events and possibility of highlighting events.

Overview

City Agenda has been operated for 12 months during the Y3 of the project. The service is part of the Barcelona pilot which includes a set of other services. This section will describe the evaluation of the service during this period.

Quantitative metrics and KPIs

In order to be able to evaluate this service we show the accumulated KPI's of all months:

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	0.8%
Memory usage	Averaged % Memory usage	Memory %	< 80	50%
Stability	How stable is the service	Number of downs	[0 – 20] per year	3
Availability	Uptime of the service	Number of downs Service uptime	90%	95%
Access Time	How much time need the service to attend a request	Access Time	[20- 4000]ms	20 ms
Service visits	Number of users using the service	Visits sessions	[100 – 10000]	<100
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	Not monitored
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[50 – 100]	>100
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[20 – 200]	1 per dey running

Table 28: Service metrics and KPIs.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end uses. These questionnaires can be found in Annex I.

All these qualitative metrics will be measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the subscriber	Service usability	> 3 Stars	**4.5 Stars**
Customisation	Degree of customisation of the service	Service customisation	> 3 Stars	**2.25 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**4 Stars**
Satisfaction	The perception of the satisfaction of the service for the subscriber	Service satisfaction	> 3 Stars	**4.25 Starts**
Monitoring usability	The perception of the usability of the monitoring for the subscriber	Monitoring satisfaction	> 3 Stars	**2.75 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 29: Qualitative KPIs for the service subscriber.

The following chart represents the perception of the usability of the service for the subscriber gathered from the questionnaires delivered to the City Agenda Service managing team.

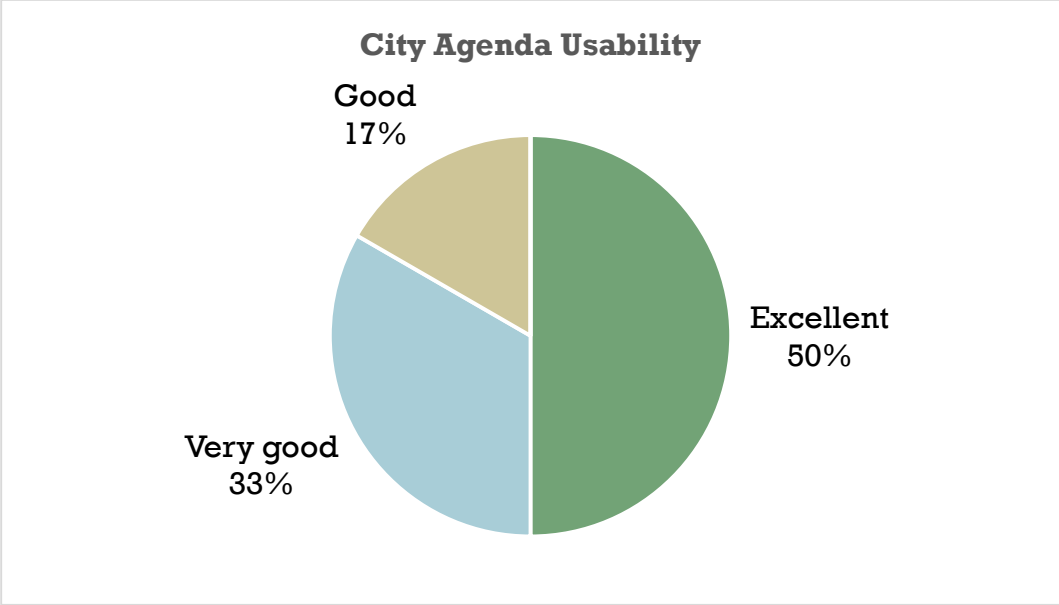

Figure 56: City Agenda – Service Subscribers usability.

Next chart represents the degree of customisation of the service for the subscriber gathered from the questionnaires delivered to the City Agenda Service managing team. As shown in figure 57, the answers have a low range of satisfaction.

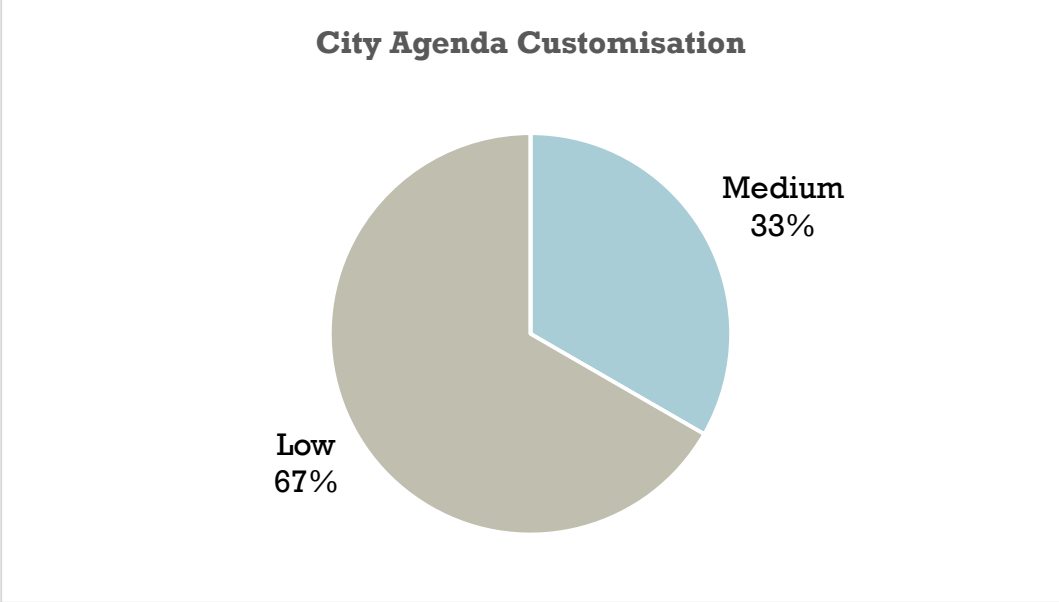

Figure 57: City Agenda – Service Subscribers customisation.

Chart in Figure 58 represents the perception of the support received for the subscriber gathered from the questionnaires delivered to the City Agenda Service managing team.

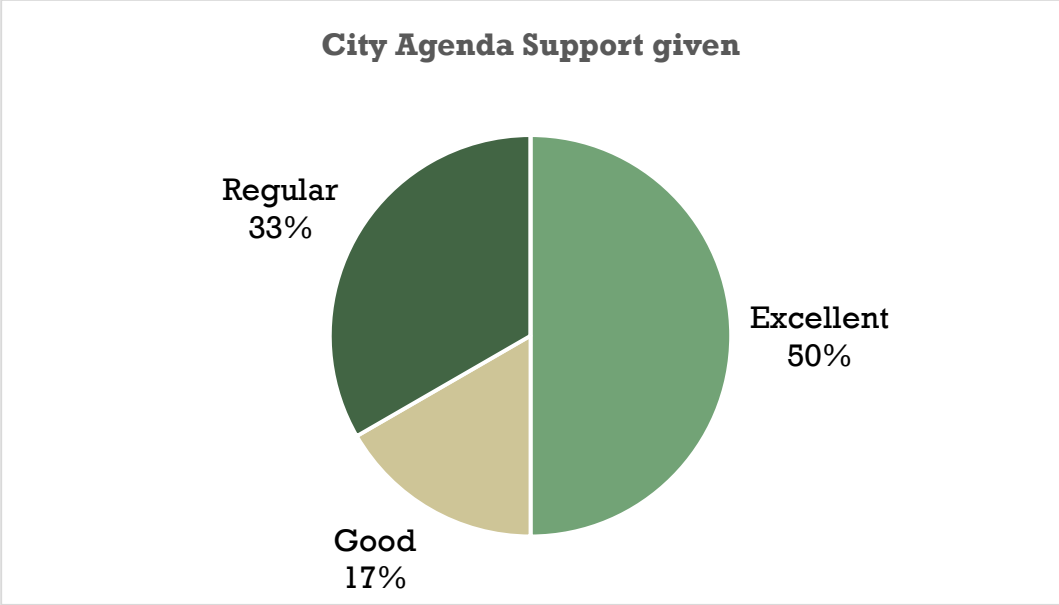

Figure 58: City Agenda – Service Subscribers support given.

The chart below represents the perception of the satisfaction of the service for the subscriber gathered from the questionnaires delivered to the City Agenda Service managing team.

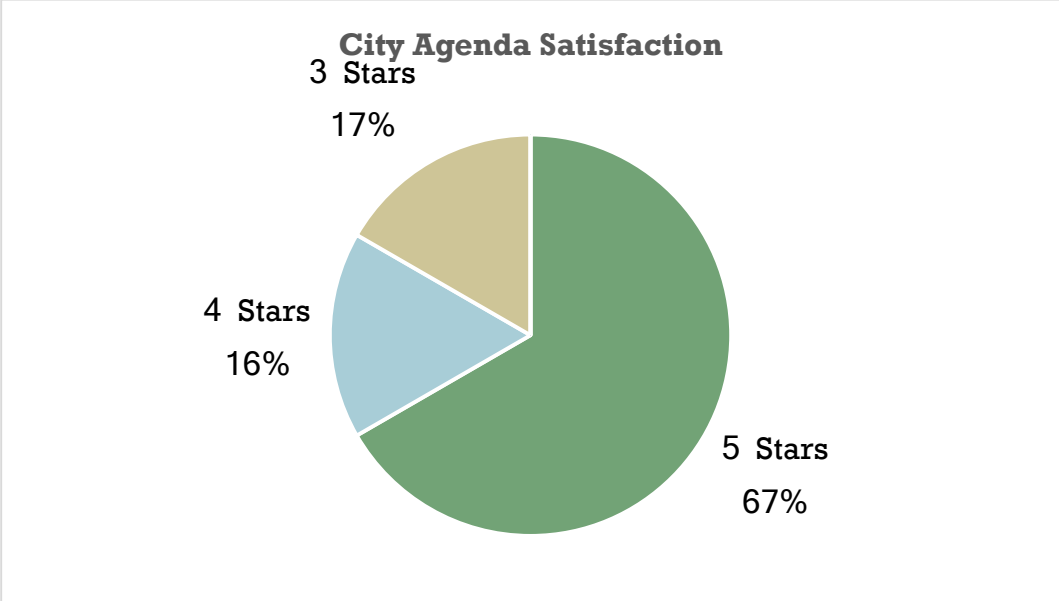

Figure 59: City Agenda – Service Subscribers satisfaction.

Chart in Figure 60, below, represents the perception of the usability of the monitoring for the subscriber gathered from the questionnaires delivered to the City Agenda Service managing team.

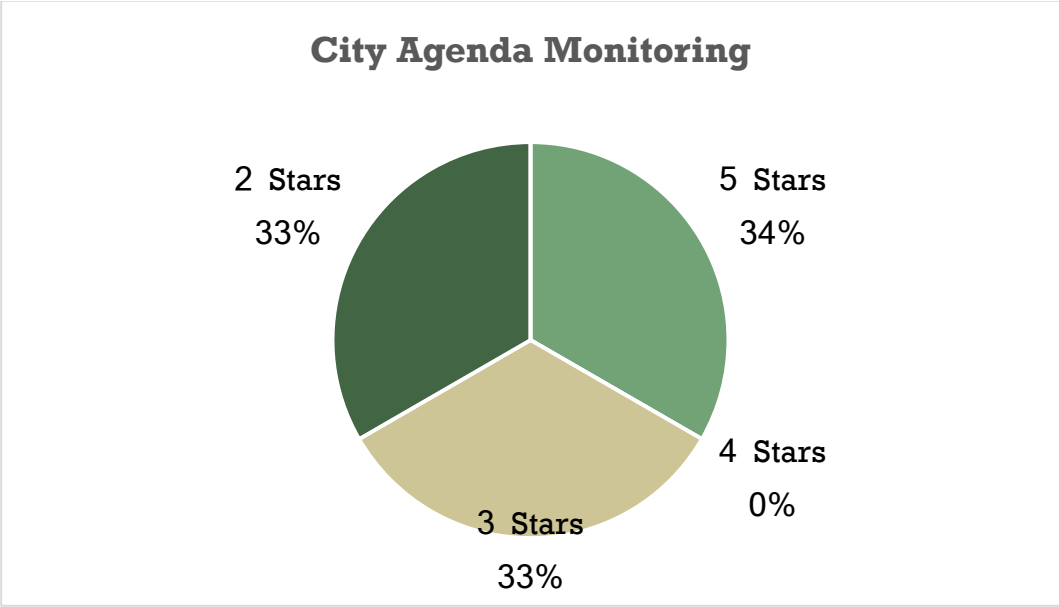

Figure 60: City Agenda – Service Subscribers monitoring.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the end user	Service usability	> 3 Stars	**4.4 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.75 Stars**
Rate	Rating of the services compared to others	Rating	> 3 Stars	**3.75 Stars**
Satisfaction	The perception of the satisfaction of the service for the end user	Service satisfaction	> 3 Stars	**4.4 Stars**

Table 30: Qualitative KPIs for the end user.

The following chart represents the perception of the usability of the service for the end user gathered from the questionnaires delivered to the City Agenda Service managing team.

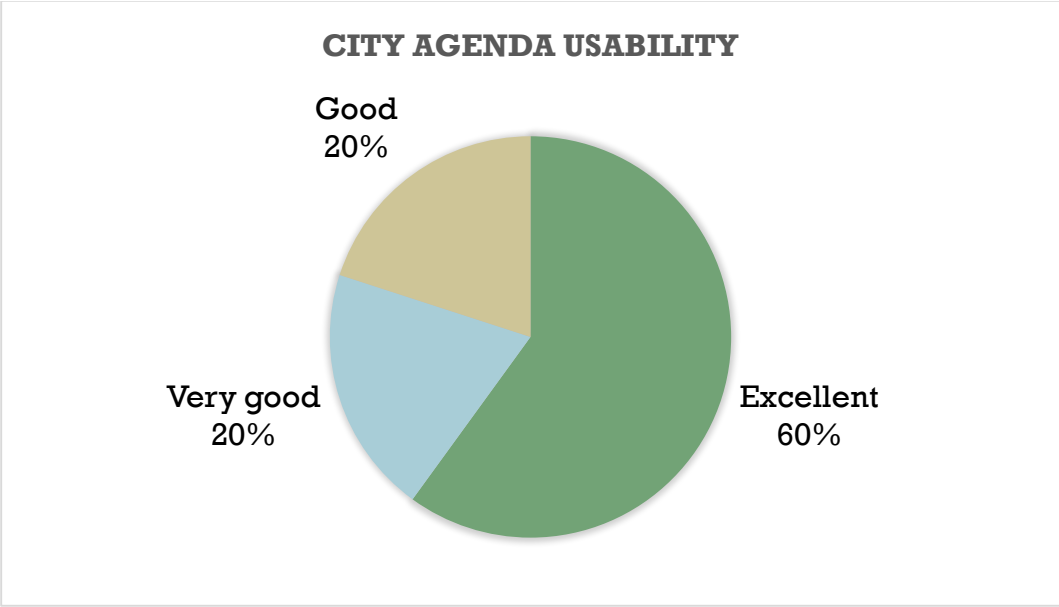

Figure 61: City Agenda –End User usability.

Next chart represents the perception of the support received for the end user gathered from the questionnaires delivered to the City Agenda Service managing team.

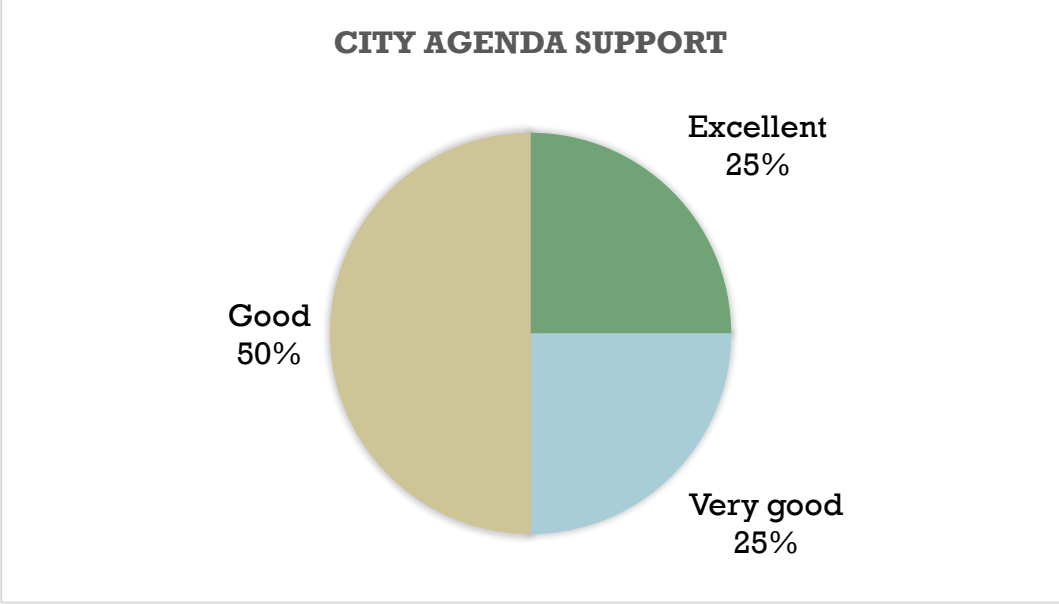

Figure 62: City Agenda – End User support.

Figure 62 below shows the rating of the services compared to others for the end user gathered from the questionnaires delivered to the City Agenda Service managing team.

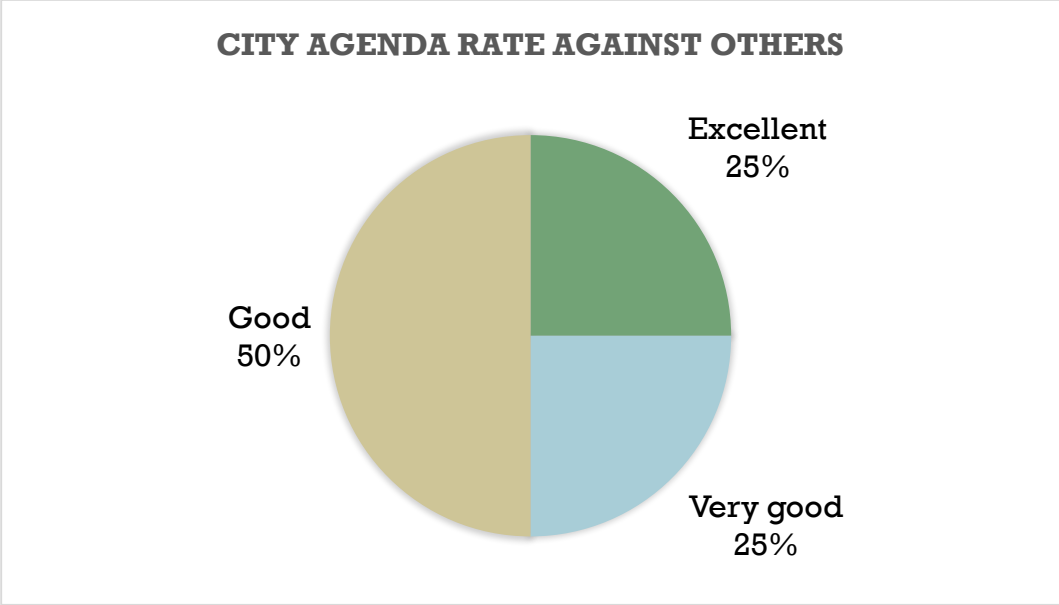

Figure 63: City Agenda – End User rate against others.

The following chart represents the perception of the satisfaction of the service for the end user gathered from the questionnaires delivered to the City Agenda Service managing team.

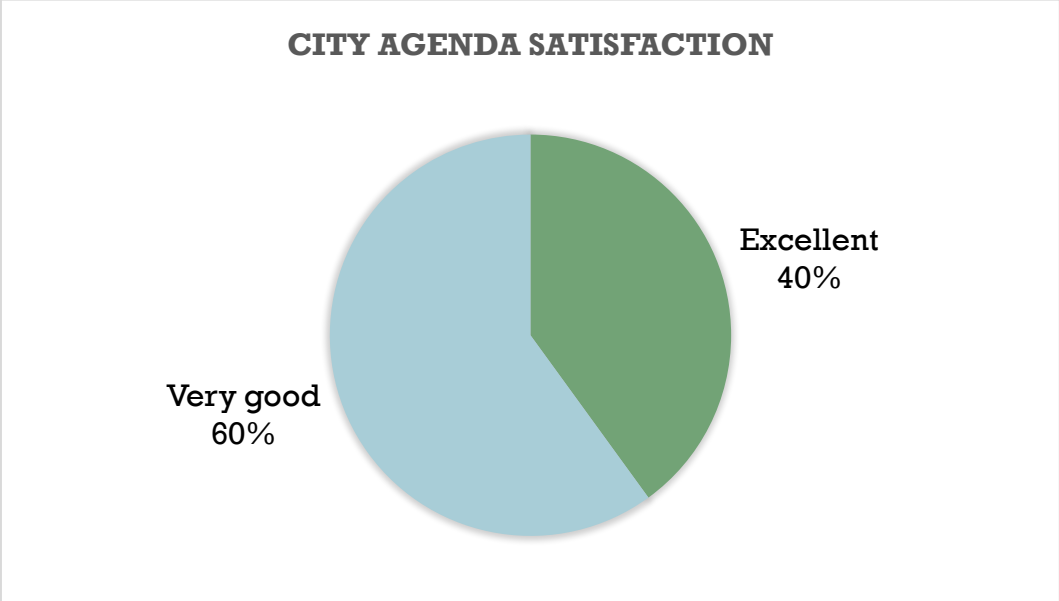

Figure 64: City Agenda – End User satisfaction.

Evaluation and conclusions

- Deployment and Customisation:

The deployment of the service created a fully working instance, but some customisation improvements have to be included in the following software versions in order to enhance user's satisfaction. The need to change or add new data sources, change colors and logos in order to fit better the city are some aspects that are on top of the table for imprvement.

- Operation and KPIs

The operation has been easy to perform, thanks to the Docker containers technology and the developed job to load data automatically to the database.

Although some downs were reported, only one of them was because of the application, that ran out of space due to the amout of data loaded from the Barcelona Open Data.

According to the tables from the "Results / KPIs", all KPIs have been achieved but the number of visits, impacting in the number of POI's consulted.

The new data loaded into the platform does not have a POI associated, so there are not a quantitative way to show its impact. To calculate the amount of new data loaded the difference of initial and final data volume in database has been performed, doubling the information from the begining.

- User experience:

All graphs from the user experience questionnaire demonstrates that City Agenda experience has been very positive, and the 100% of the people that answered the quetionnaire says it is very good or excelent.

With all of this information it can be concluded that the results, as not compile what expected in some parameters, were sucessful and apreciated by the users.

3.3.5. CKAN Corby

CKAN run a number of open data registries around the world, including the UK government open data website data.gov.uk. CKAN allows for greater openness and transparency of public data, making it easier to find, share, and reuse open content and data.

The CKAN instance in Corby contains datasets from across Corby and Northamptonshire, bringing local data together in an easy to use and visible format.

Overview

The CKAN instance for Corby (www.corbydata.co.uk) has been operated for 12 months during the Y3 of the project. The service is part of the Corby pilot, which includes a set of other services, and is hosted by IMI. This section will describe the evaluation of the service during this period.

Quantitative metrics and KPIs:

- Technical variables
  - CPU Usage
  - Memory Usage
  - Stability
  - Availability
  - Maximum concurrent users
  - Access Time

- Business variables
  - Effort (Staff-hours)
  - Service visits
  - Service pageviews
  - Datasheets uploaded

Quantitative metrics and KPIs

The metrics to evaluate the service are shown in the following table:

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	<20%
Memory usage	Averaged % Memory usage	Memory %	< 80	1GB
Stability	How stable is the service	Number of downs	[0 – 20] per year	2
Availability	Uptime of the service	Number of downs Service uptime	90%	98.52%
Access Time	How much time need	Access Time	[20- 3000]ms	<2000ms

	the service to attend a request			
Service visits	Number of users using the service	Visits sessions	[250 – 400]	178 sessions
Service pageviews	Number of visits to pages of the service	Pageviews	[300 – 500]	802 pageviews
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[1 – 50]	>25[3]
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[20 – 100]	41[4]
Effort (Staff-Hours)	Number of staff hours for project execution	Staff Hours	150h	135h

Table 31: Service metrics and KPIs.

Qualitative variables

The following chart represents the satisfaction of the users of the gathered from the questionnaires delivered to the CKAN Corby Service managing team.

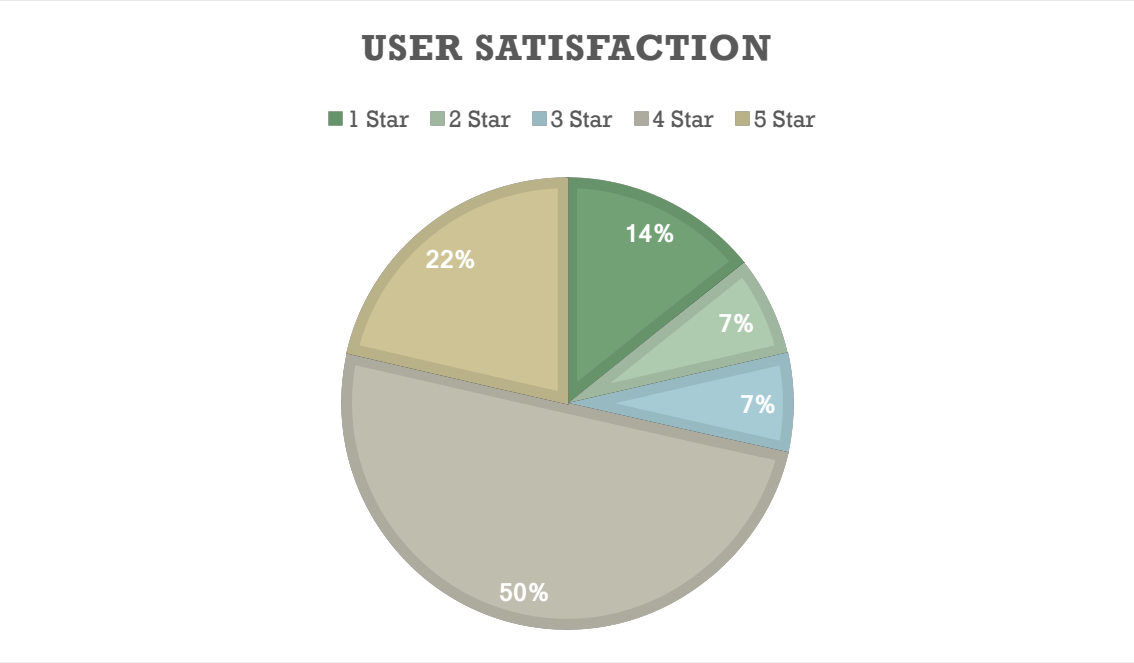

Figure 65: The overall user satisfaction of the CKAN services, including the Corby instance.

[3] This figure has been obtained by means of a test tool using virtual users accessing from a host located at Dublin, Ireland.
[4] At present, the service has 26 datasets which may contain one or more data files.

The following chart represents the user rates about usability gathered from the questionnaires delivered to the CKAN Corby Service managing team.

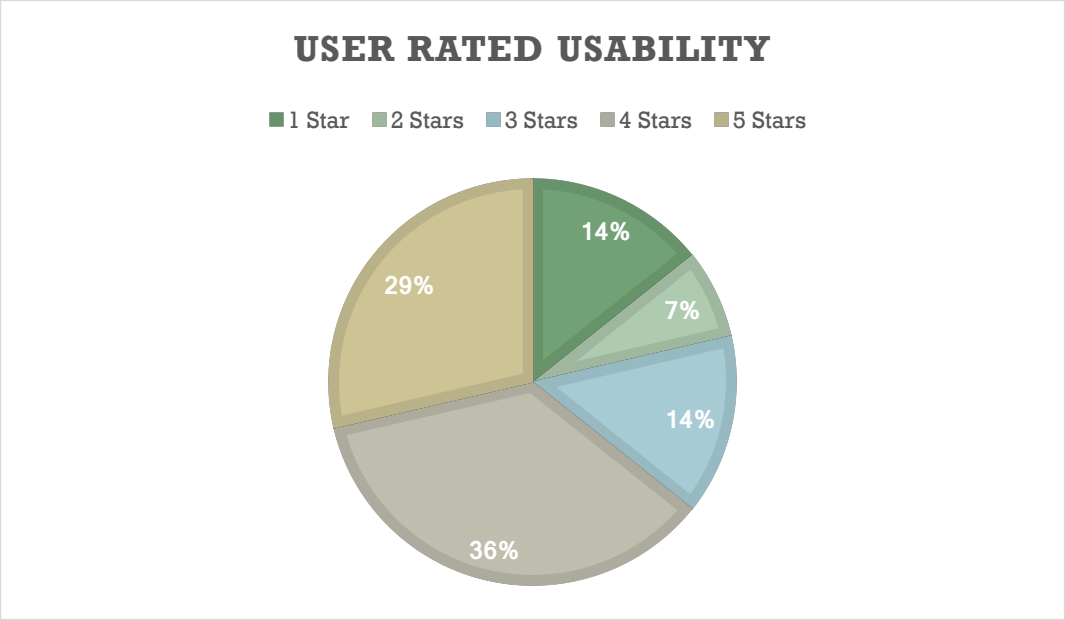

Figure 66: The usability of the CKAN services rated by the users in the End User Service Questionnaire.

Evaluation and conclusions

- Deployment and Customisation

The Corby CKAN pilot was successfully deployed in Month 25 of the project, and so has been running successfully for 12 months.

In order to customise the CKAN instance for the Corby pilot, Electric Corby were registered to the website and authorised to publish on the website. This allowed Electric Corby to continue to publish datasets that would be useful and interesting for the local area.

Other customisations made to the CKAN instance were aesthetic, including the use of the Electric Corby logo and other brand details (including font and colour).

The range of customisation of the CKAN instance is very good, as it allows local authorities to customise the instance using their own local data.

- Operation and KPI

According to the table above, it can be concluded that all of the KPIs have been achieved, apart from the service visits KPI. The KPIs indicate that the Corby CKAN instance has been performing correctly and reliable, as its technical metrics remained well within the KPIs.

- User Experience

The User Experience across all of the Corby services has been previously evidenced as good, as seen in section 4.2.2.1 for the Corby Bus Portal.

Additionally, user satisfaction across the whole of the CKAN services, including the Corby CKAN instance, was also good. The chart of Figure 65 shows the overall user satisfaction with the CKAN services, which had an average star rating of 3.6. This was taken from the results from the CloudOpting End User Service Questionnaire.

The support given to the users was also rated well with an average of 3.4 stars, and the usability of the CKAN services also received a good rating, with an average of 3.6 stars. This can be seen broken down further in the chart on Figure 66.

3.3.6. Business Portal

Overview

The Business Portal was running at http://84.240.187.22/ , however the service is no longer in use. The Business Portal was originally developed as part of the Growing Corby project in Corby, acting to faciliate the exchange of information and communication between local businesses. This website template has been hosted successfully by the CloudOpting platform for 12 months, however the service is not in use since the Growing Corby programme ended and the website currently has no business funding to offer local businesses.

Quantitative metrics and KPIs

The KPIs were monitored between March and August 2016, and included:

Quantitative Metrics

- CPU Usage
- Memory Usage
- Stability
- Availability

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	15%
Memory usage	Averaged % Memory usage	Memory %	< 80	80%
Stability	How stable is the service	Number of downs	[0 – 20] per year	0
Availability	Uptime of the service	Number of downs Service uptime	90%	100%

Table 32: Service metrics and KPI's

Evaluation and conclusions

While there were no further KPIs measured from the Business Portal as a result of the Growing Corby Project's closure, the results from the KPIs above show that it succeeded as a test of the technical aspects of the CloudOpting platform. The Business Portal itself is a replicable webpage template that is easy to use, and thus can be an important application for the CloudOpting repertoire as a tool for local authorities to offer their own local business support.

3.3.7. Energy Dashboard

The Energy Dashboard allows users to monitor, track, and analyse their energy consumption. It also provides up-to-date local weather information, which in this case has been customised for Corby. This data includes solar radiation and current and historical temperature data.

Overview

The Energy Dashboard (yourenergydashboard.co.uk) has been operating for 8 months during the Y3 of the project. The service is part of the Corby pilot which includes a set of other services. This section will describe the evaluation of the Energy dashboard during this period.

Quantitative metrics and KPIs

The Quantitative metrics gathered during the monitoring periods are listed in the table below. The metrics listed are the average monthly figures taken from the reporting templates from all of the evaluation periods.

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	25%
Memory usage	Averaged % Memory usage	Memory %	< 80	95%
Stability	How stable is the service	Number of downs	[0 – 20] per year	3 (2 Planned)
Availability	Uptime of the service	Number of downs Service uptime	90%	100%
Page Views	Website traffic	Number of page views	[50- 100]	4801
Accounts Made	Amount of accounts created	Accounts Registered	[5– 8]	0
Monitors Installed	Monitors installed and reporting to system	Number of Monitors	[30 – 40]	0
Data Sheets Uploaded	Number of datasheets added to the platform	Number of datasheets added	[1000 – 2000]	0
User Energy Reduction	Reduction in users energy from historical data	Electrical Consumption Reduction %	[5-8]%	0
Monitor Samples	How many samples are currently in the database	Number of samples	[5000-10000]	12633

Table 33: Service metrics and KPIs for the Energy Dashboard.

Evaluation and conclusions

- Deployment and Customisation

The Energy Dashboard has been fully operating since Month 29 of the Project. The Energy Dashboard required more time to design the dashboard and specify its requirements before it could be deployed. Once the customisations and new features were fully completed and developed, this service successfully went live with the domain name www.yourenergydashboard.co.uk.
Customisations to the Energy Dashboard included minor look and feel changes, including the addition of Electric Corby logos, font, and brand details (including the colour scheme).

- Operation and KPI

According to the table above, most of the KPI's have been achieved. However, due to issues in collecting monitoring data and no accounts were made, we were not able to quantify the KPI 'User Energy Reduction'.

Nevertheless, the KPI results indicate that the service was running correctly and able to support a standard volume of traffic, and ran reliably on the CSI Cloud Platform.

- User Experience

As stated in the user experience evaluation in the Bus Portal section, the Corby services as a whole received good user feedback. While users were unable to show actual energy reductions due to issues in obtaining actual energy data, the weather dashboard provided useful local information with a large amount of data.

3.3.8. POI/Next2Me

The purpose of this service is to provide a multi city application for citizens to find interesting services, events and public facilities in the neighbourhood. This pilot service adds a dataset from Karlshamn municipality including point-of-interest (POI) data such as playgrounds, schools, parks etc. to be included and presented in the Next2Me application from Worldline. This, in fact, turns the Next2Me application into a multi city application serving both Barcelona and Karlshamn.

Overview

POI (Point Of Interest) has been operated for three months during the Y3 of the project. The service is part of the Karlshamn pilot which includes two services. This section will describe the evaluation of the service during this period.

Quantitative metrics and KPIs

In order to be able to evaluate this service we show the accumulated KPI's of all months:

KPI	KPI Description	Associated Metrics	Acceptable range	Measured values
CPU usage	Average % CPU usage	CPU %	< 75%	0.8%
Memory usage	Averaged % Memory usage	Memory %	< 80	50%
Stability	How stable is the service	Number of downs	[0 – 20] per year	4
Availability	Uptime of the service	Number of downs Service uptime	90%	94%
Access Time	How much time need the service to attend a request	Access Time	[20- 4000]ms	20 ms
Service visits	Number of users using the service	Visits sessions	[100 – 10000]	<100
Service Demand	No of datasets downloaded from the service.	No of datasets downloaded	[20 – 200]	Not monitored
Maximum Concurrent users	How many simultaneous users will use the service during a very short period of time	Concurrent users	[50 – 100]	>100
Datasheet uploaded	Datasheet uploaded	Number of datasheet added	[1-10]	1

Table 34: Service metrics and KPIs.

Qualitative variables

The qualitative KPI were measured by means of two questionnaires, one focused on the administrator of the instance and the other one focused on the end uses. These questionnaires can be found in Annex II.

All these qualitative metrics will be measured through questionnaires with five possible satisfaction values from 1 to 5 stars.

The qualitative KPIs results for the service subscribers:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the subscriber	Service usability	> 3 Stars	**4.33 Stars**
Customisation	Degree of customisation of the service	Service customisation	> 3 Stars	**2.33 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.66 Stars**
Satisfaction	The perception of the satisfaction of the service for the subscriber	Service satisfaction	> 3 Stars	**4.5 Starts**
Monitoring usability	The perception of the usability of the monitoring for the subscriber	Monitoring satisfaction	> 3 Stars	**3.33 Stars**
Willingness to renew	Willingness to renew the service	Willingness to renew the service	> 3 Stars	**NA**

Table 35: Qualitative KPIs for the service subscriber.

The following chart represents the perception of the usability of the service for the subscriber gathered from the questionnaires delivered to the POI/Next2Me Service managing team.

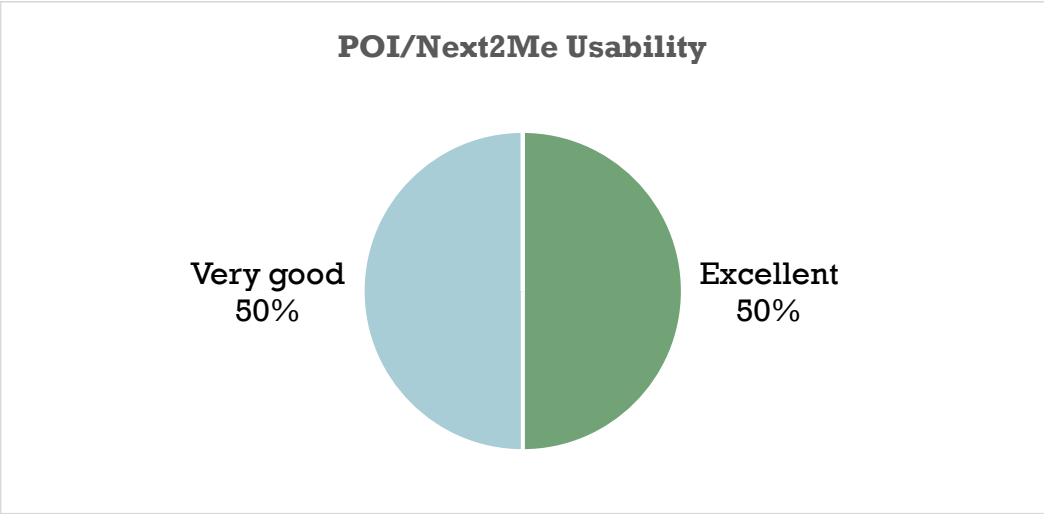

Figure 67: Next2Me – Service Subscriber Usability.

Next chart represents the degree of customisation of the service for the subscriber gathered from the questionnaires delivered to the POI/Next2Me Service managing team.

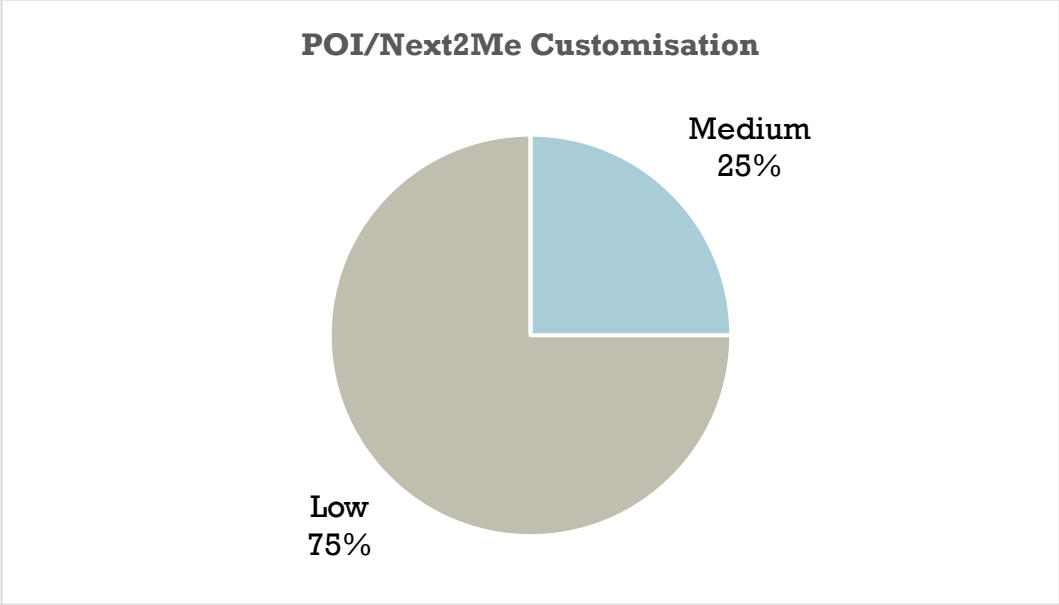

Figure 68: Next2Me – Service Subscriber customisation.

Next chart represents the perception of the support received for the subscriber gathered from the questionnaires delivered to the POI/Next2Me Service managing team.

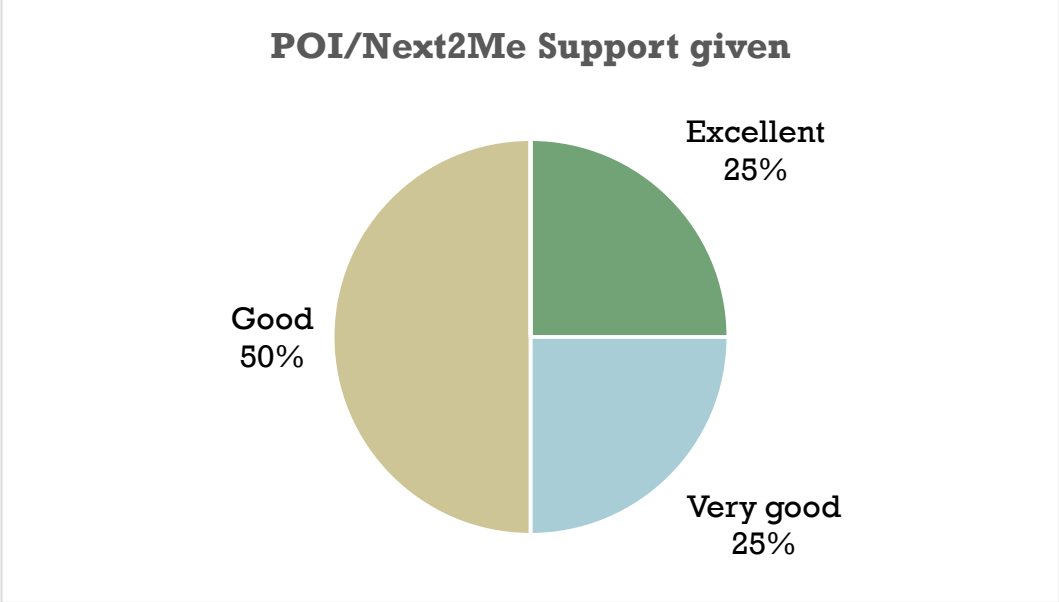

Figure 69: Next2Me – Service Subscriber support given.

Chart in Figure 70 below represents the perception of the satisfaction of the service for the subscriber gathered from the questionnaires delivered to the POI/Next2Me Service managing team.

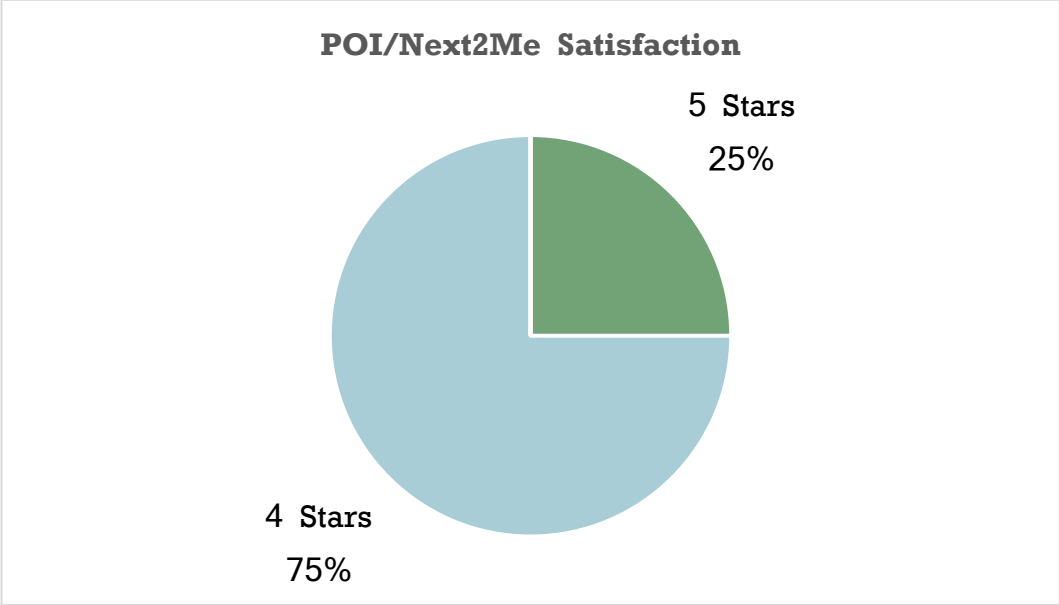

Figure 70: Next2Me – Service Subscriber satisfaction.

The following chart represents the perception of the usability of the monitoring for the subscriber gathered from the questionnaires delivered to the POI/Next2Me Service managing team.

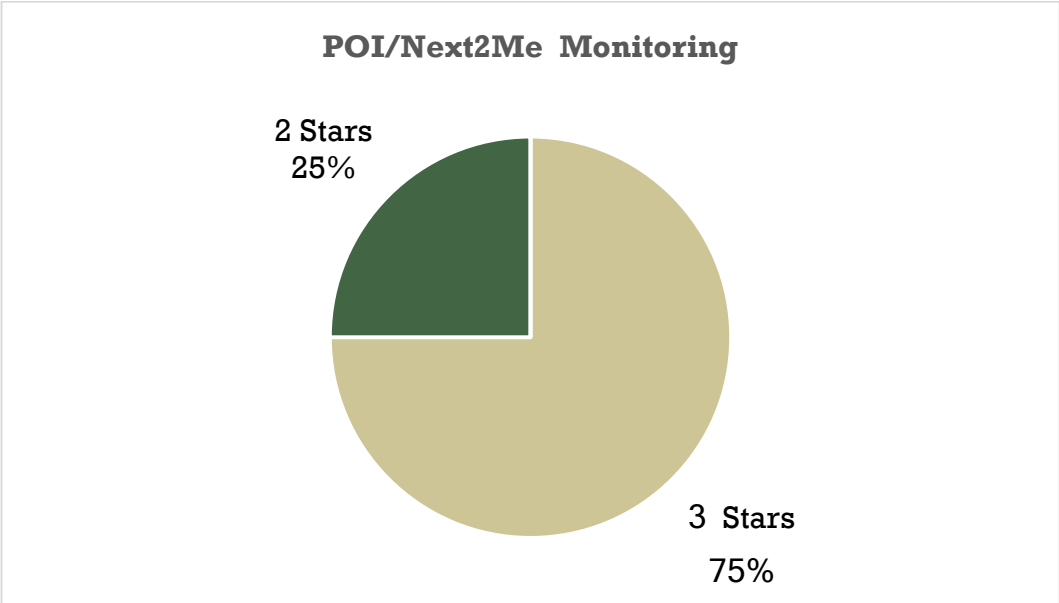

Figure 71: Next2Me – Service Subscriber monitoring.

The qualitative KPIs results for the end user:

KPI	KPI Description	Associated Qualitative Metrics	Acceptable range of values	Measured values
Service usability	The perception of the usability of the service for the end user	Service usability	> 3 Stars	**3.75 Stars**
Support	The perception of the support received	Service support	> 3 Stars	**3.5 Stars**
Rate	Rating of the services compared to others	Rating	> 3 Stars	**3.25 Stars**
Satisfaction	The perception of the satisfaction of the service for the end user	Service satisfaction	> 3 Stars	**3.75 Stars**

Table 36: Qualitative KPIs for the end user.

Next chart represents the perception of the usability of the service for the end user gathered from the questionnaires.

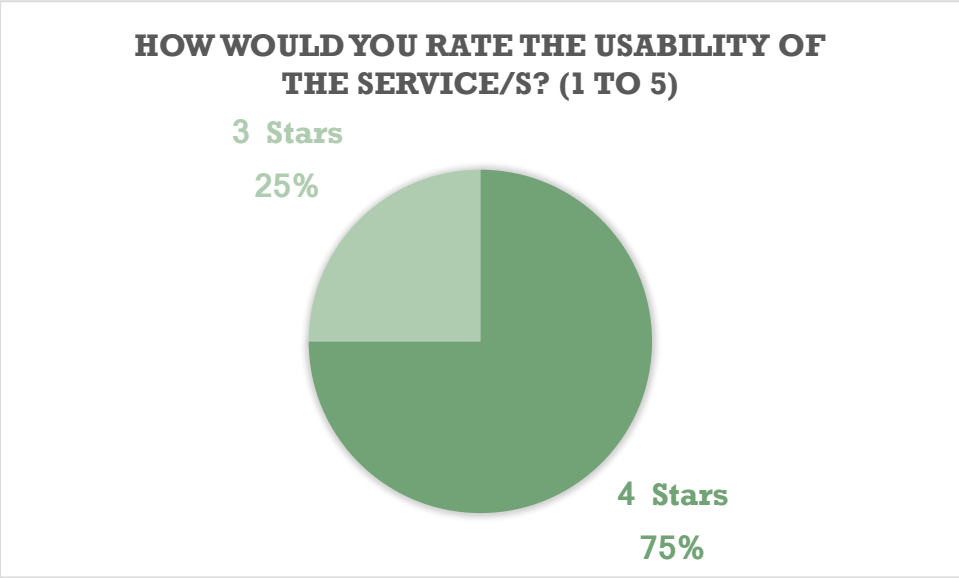

Figure 72: End User usability.

The following chart shows the perception of the support received for the end user gathered from the questionnaires.

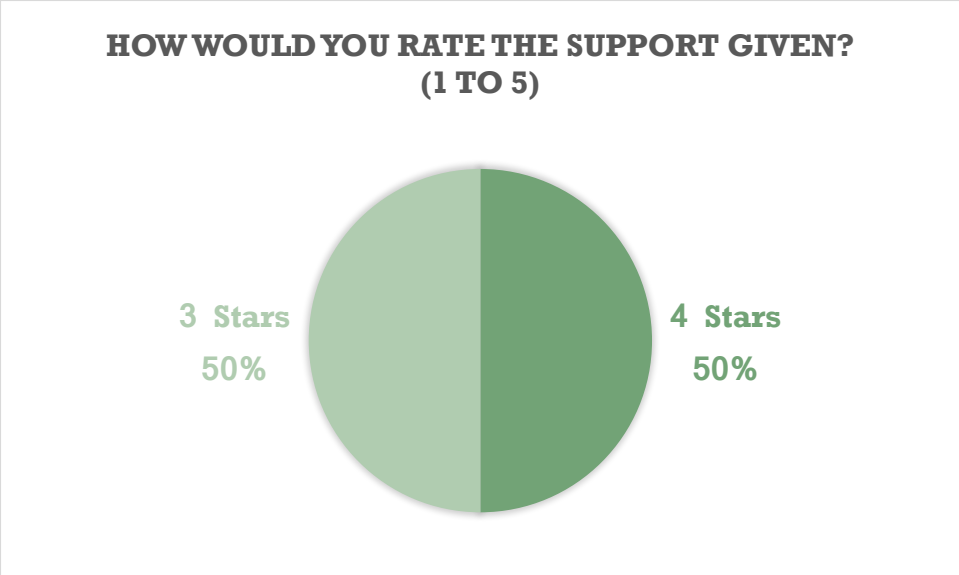

Figure 73: End User support.

In Figure 74, the chart represents the rating of the services compared to others for the end user gathered from the questionnaires.

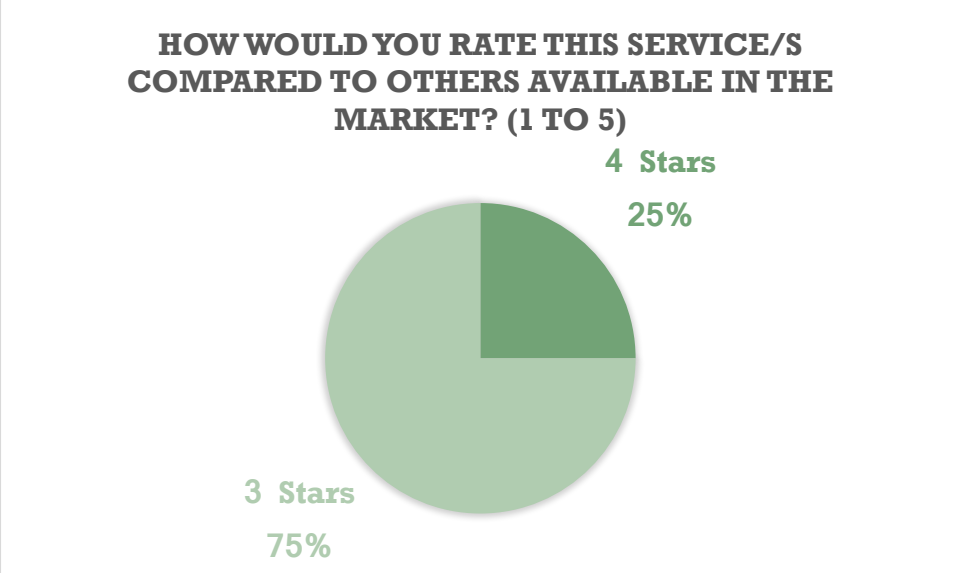

Figure 74: End User rate against others.

The following chart represents the perception of the satisfaction of the service for the end user gathered from the questionnaires.

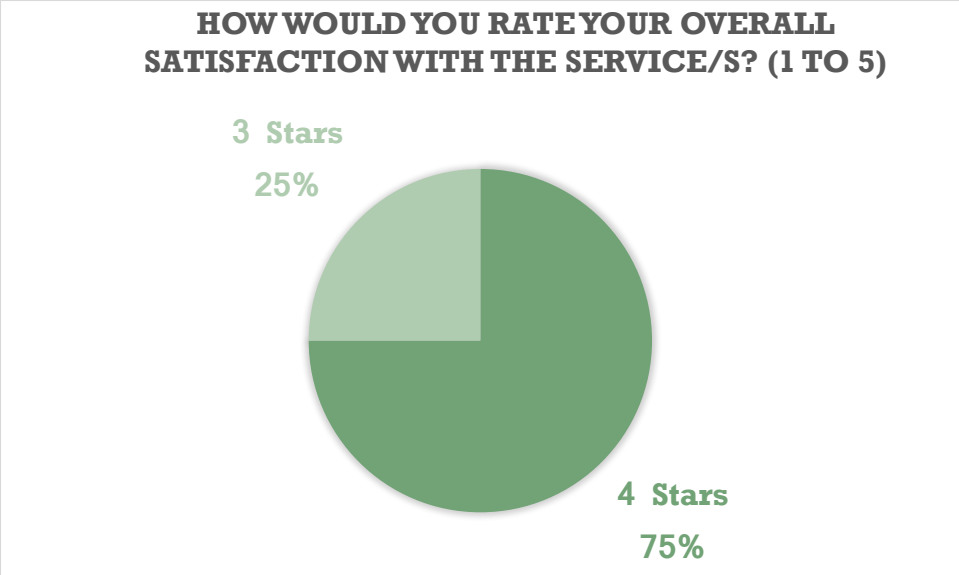

Figure 75: End User satisfaction.

Evaluation and conclusions

According to the facts above, our conclusion is that all expected KPIs have been achieved, except number of visits that was slightly lower than expected and the customisation level. The updating procedure of the POI dataset was not automatic, because the data was not available as open data from the municipality of Karlshamn, the data was only made available on demand. The lesson learned here is that to increase end user attraction, a service like this should implement automatic update procedures from open data sources to be as up to date as possible.

Another conclusion is that for a multi city service like this, the data sources available in the different cities will almost certainly differ regarding content, structure and openness. Thus, both the service data adaption procedures, as well as end user interfaces, must be very flexible and resilient to handle these differences.

One more experience was that a strategy for multiple language support should be in place for a multi city/country service like this, to effectively support both local and visiting end users in a region/city.

The final conclusion is that the Next2Me service performed well running on the CloudOpting platform, only a few improvements to the service itself regarding the end user experience and data source handling as elaborated above could be suggested.

4. Replication Plan

4.1. Introduction

The replication plan chapter describes the different scenarios and activities required to replicate the CloudOpting ecosystem in other locations. This ecosystem includes, essentially, the services, the CloudOpting platform and the monitoring module.

4.1. Scenarios

There are different scenarios to describe the replication plan of the CloudOpting platform. This scenarios are linked to specific business needs of the actors involved in the different process. Thus, while a service subscriber may be interested in deploying a service in its own cloud, others would prefer to deploy it in a cloud provided by a third party.

A representation of the CloudOpting ecosystem is shown in Figure 76. The representation includes the cloud, the services, the catalogue of services, the monitoring module, and the CloudOpting manager.

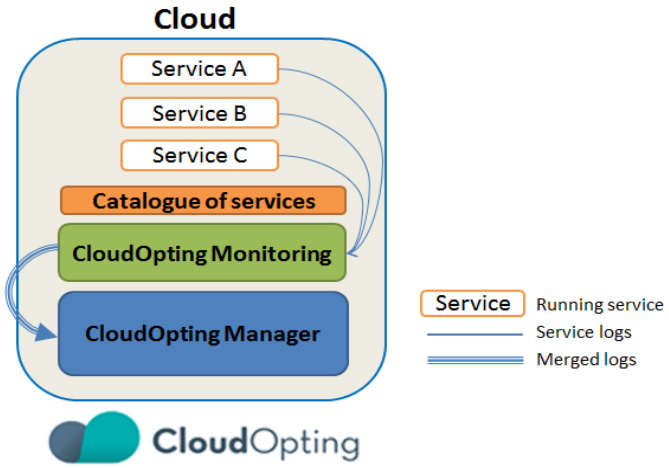

Figure 76: CloudOpting ecosystem.

The cloud hosts all the components of the platform and a series of running services. The monitoring module processes the logs gathered from the services, and the CloudOpting manager controls all the processes of the platform.

From the image above, several replication scenarios can be described, but the most relevant ones according to the acquired knowledge are:

- In-Cloud Service Replication
- New Cloud Service Replication
- Service and Monitoring Replication
- Platform Replication

4.1.1. Service Replication In-Cloud

This scenario describes the replication of a service from the catalogue in the same cloud where the platform is hosted. Service subscribes willing to subscribe to a service may not be interested in deploy it in a specific cloud so may just choose it from the ones available in the platform. A new isolated VM will then be created in the same cloud and linked to the platform for further monitoring.

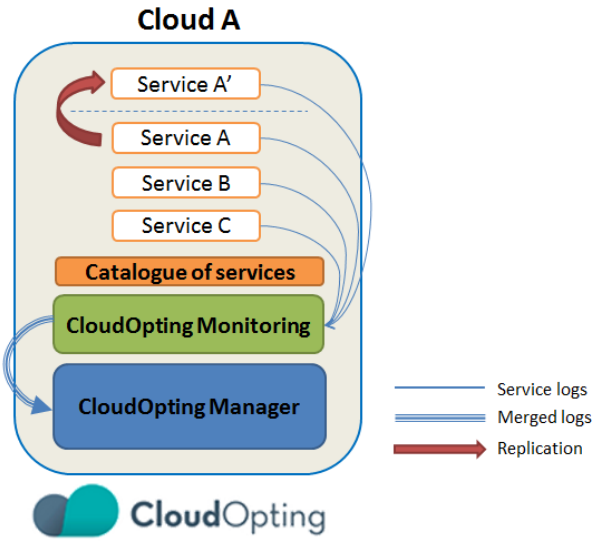

Figure 77: In-cloud service replication.

4.1.2. Service Migration New Cloud

This scenario is very similar to the previous one except by the cloud selected to host the new service. In this case, the service subscriber selects a different cloud from the ones available in the platform to deploy the service. This cloud may be the subscriber's private cloud or may be supplied by a different cloud provider. The new VM created has to be linked to the platform to send the logs to the monitoring component.

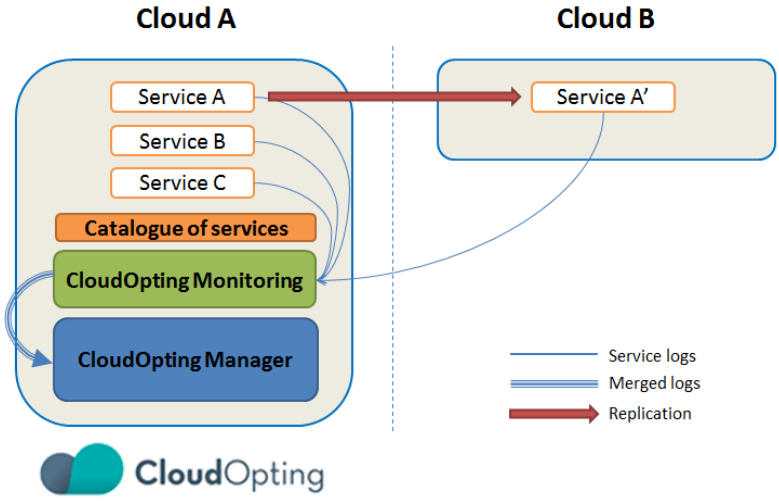

Figure 78: New cloud service replication.

The monitoring traffic generated by the new service instance is sent from cloud B to cloud A. This traffic between clouds requires connecting clouds to each other and to the internet. When the new instance is created, it also creates the firewall rules to grant access to specific ports, and creates static routes to forward traffic to the manager.

4.1.3. Service and Monitoring Replication

In this case, both a service and the monitoring component are replicated in a different cloud. This scenario may be of interest to service subscriber willing to keep their logs locally and not sending them to an external administrator. The new service instances are linked to the monitoring component that has been replicated and the merged logs are sent to the manager.

This scenario describes the piloting year of the project where the monitoring component was hosted in CSI's cloud (Italy) and the manager was kept in IMI's cloud (Barcelona).

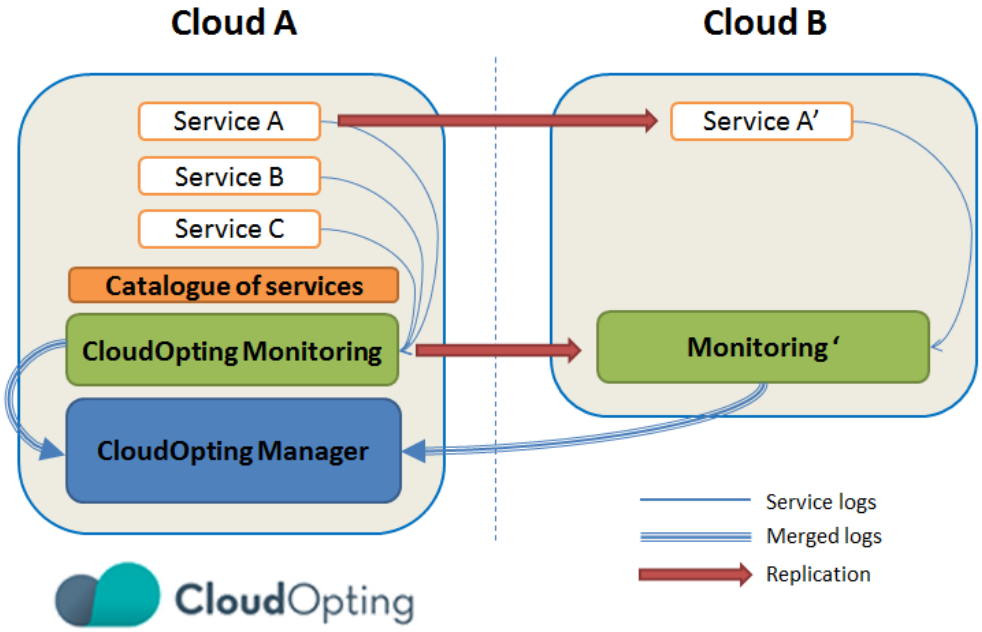

Figure 79: Platform replication.

In this scenario, the service provider has 2 services deployed in 2 different clouds. The monitoring traffic generated by the service is sent to the monitoring component hosted in the same cloud as the service. If more instances are deployed, all the generated logs will be sent to the monitoring component. Therefore, the main traffic is generated inside the cloud and the aggregation of data in the new cloud is sent to the CloudOpting manager in the original platform.

4.1.4. Platform Replication

The last replication scenario is related to the replication of the platform itself. This scenario is of interest to platform providers willing to replicate the whole platform in a different location. Hence, all the components included in the platform are replicated. This replication of the platform starts, essentially, with the installation of the platform (*D4.2 Pilot Deployment Report, Annex II: Platform Installation Guide)*.

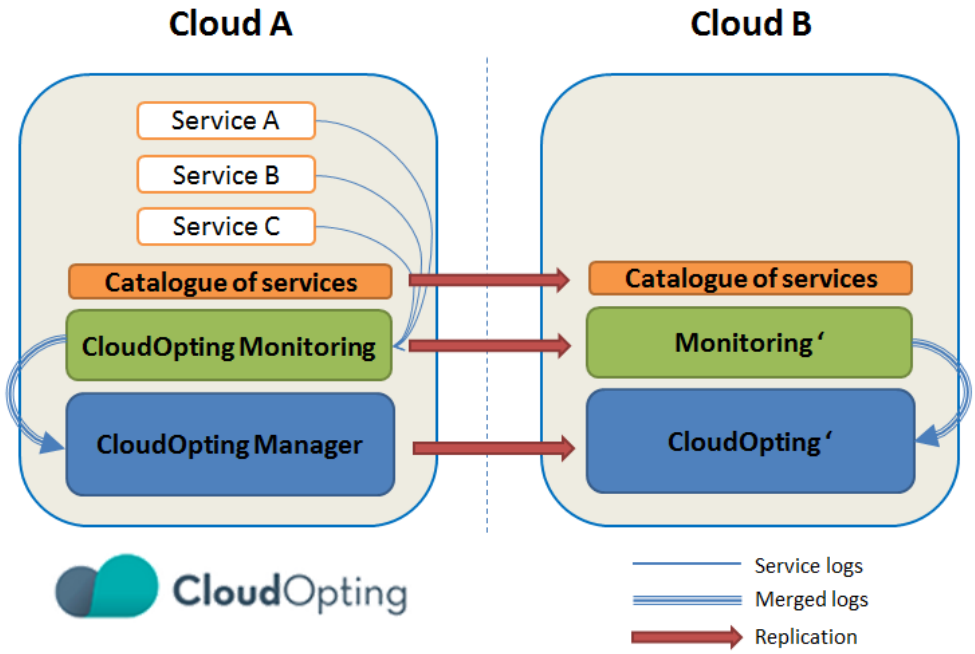

Figure 80: Platform replication.

By default, new instances of the platform are totally independent among them. This means that there is no communication of any kind among instances or clouds.

The replication of the services catalogue requires a special mention here. When installed, a new platform is created with an empty catalogue of services. Platform providers willing to include services to the new catalogue have to come to new agreements with the service providers. This situation may lead to service providers providing services to different platforms.

4.2. Service replication

Service replication is one of the main goals of the CloudOpting Platform. Once a service is published into the CloudOpting Catalogue, the service is available to all of the CloudOpting users to be subscribed. The process of subscribing to a service is defined and designed in deliverable D2.3 in section *1.2.5 Subscribe a Service*.

Essentially, the process of service replication is displayed in 3 main steps included in the following diagram:

Figure 81: High Level Replication process.

Next sections will go into a detailed description of this process.

4.2.1. Service publication

Before a service is made available in the CloudOpting Platform it has to be published by its creator, the Service Publisher. Once the service is published (Figure 82) it can be instantiated by CloudOpting users with Service Subscriber roles.

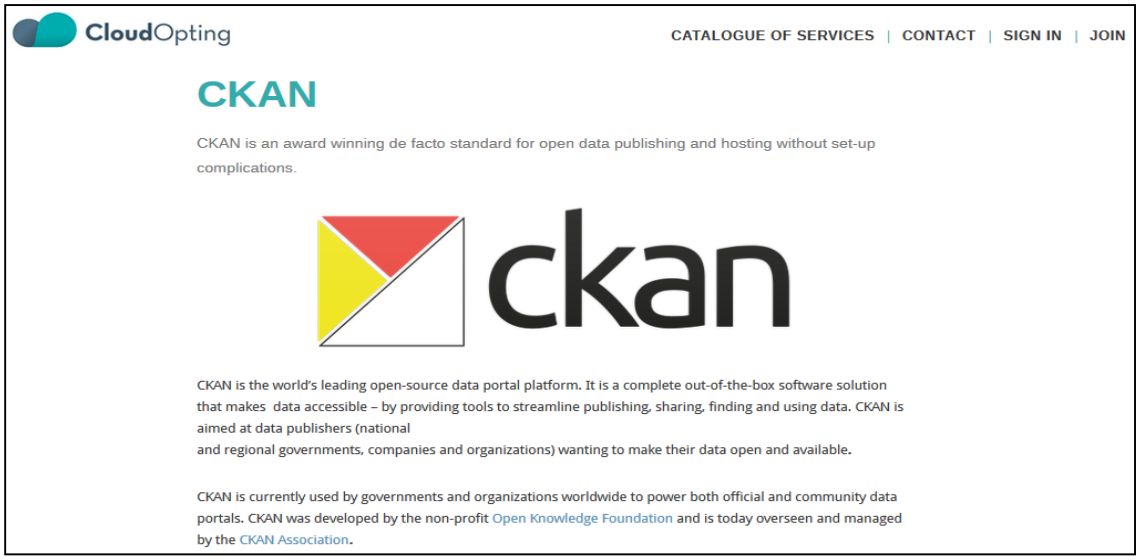

Figure 82: CKAN published information screenshot.

The platform includes an administrator section to manage all the services that have been published or are pending to be published. After the service has been formally published its status is listed in the Services List as *Published*.

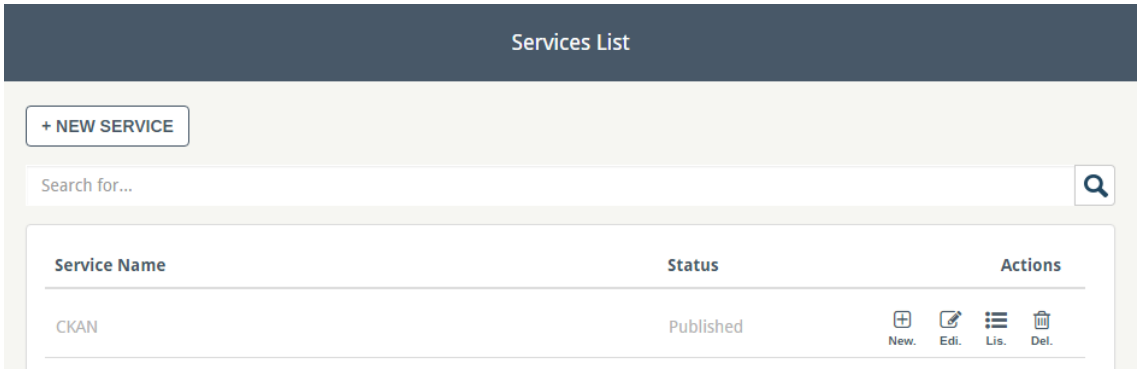

Figure 83: Service list screenshot.

4.2.2. Service Subscription

When the service is available in the CloudOpting Catalogue, any CloudOpting user can subscribe to it. As mentioned before, subscription process details can be found in D2.3. During the subscription process, the service subscriber may customise some deployment aspects of the future service instance. Furthermore, the service subscriber has to select the cloud infrastructure where the instance will be hosted. The following image shows the cloud selection panel.

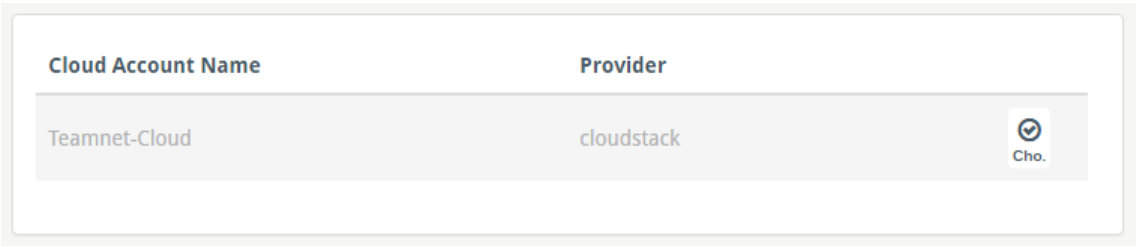

Figure 84: CloudOpting Platform Provider selection screenshot.

When users subscribe to a service, the Instance List shows the status of the new instance for that particular service. All the service instances are available to be managed by the Service Publisher. As depicted in the next image, before the actual deployment of an instance, new deployments have to be accepted by service providers so the deployment status remains *Requested* until its aproval by the service provider:

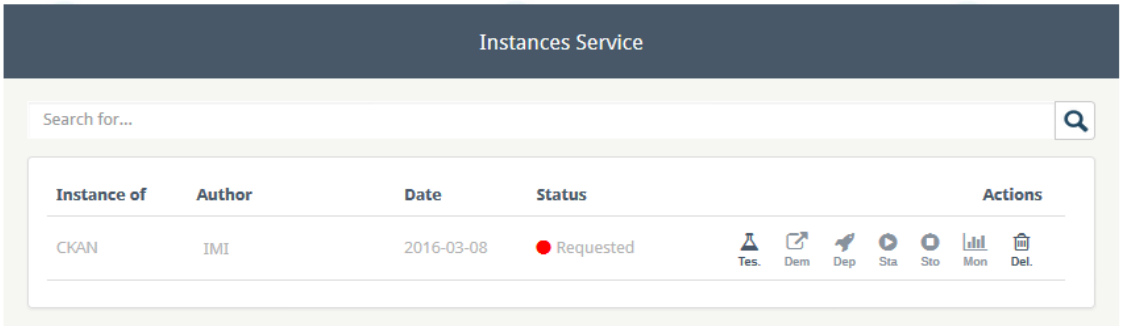

Figure 85: Service Instance list; Requested service.

4.2.3. Service Deployemnt

When the Service Provider aproves the deployment of a new instance of the service, the status of the deployment changes to *Running* (Figure 86).

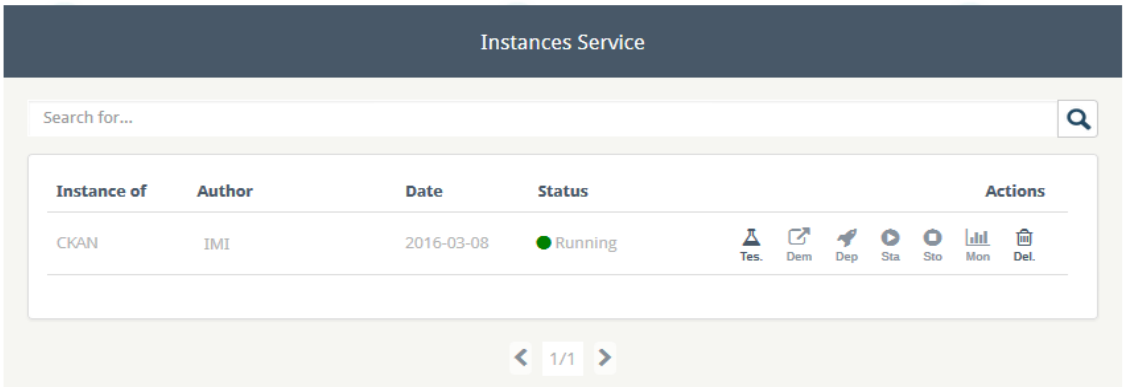

Figure 86: Service Instance list; Running service.

Then, the service instance is available to the end users. This step is the final step of the replication process:

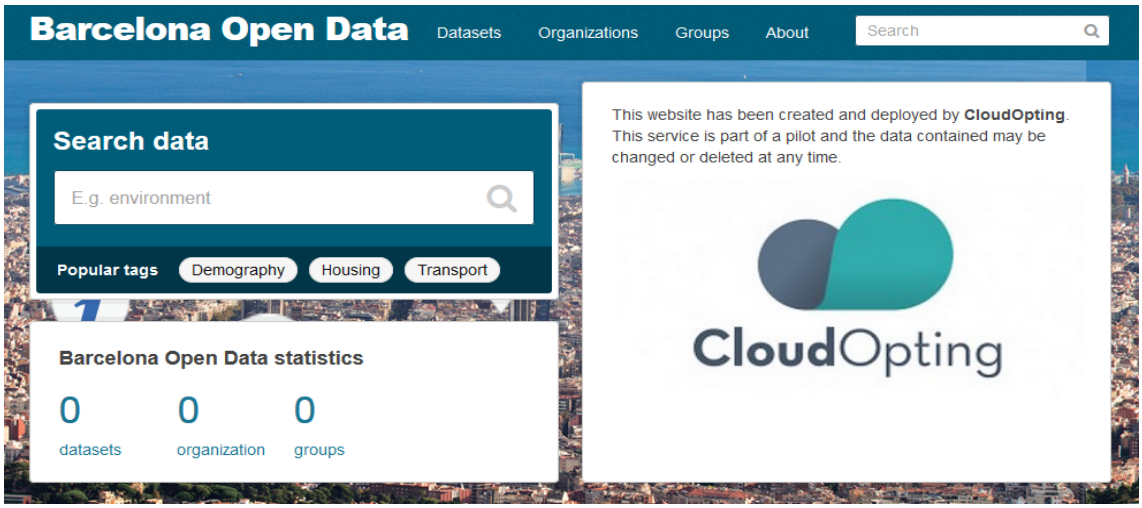

Figure 87: Service is available.

4.3. Services replication example in Y3: CKAN

During the third year of CloudOpting, the CKAN service from Barcelona has been offered to Barcelona and Corby by means of CloudOpting. There were two instances of the same service for two different cities running in the same cloud.

The process of replicating the CKAN service from Barcelona to Corby was the following:

1. The CKAN service was published by IMI as Service Provider. This has lead to the result of having the CKAN service available in the CloudOpting Service List:

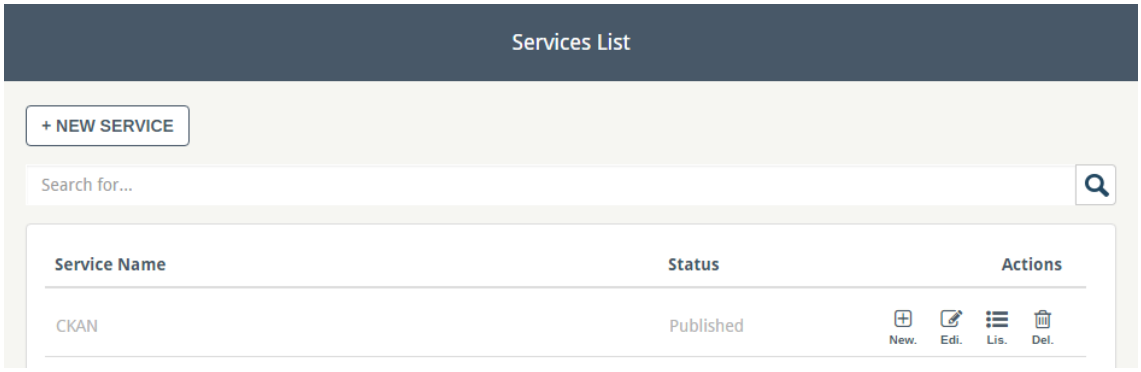

Figure 88: CKAN Service publication.

2. After the service was succesfully published, a CKAN instance was deployed by IMI as Service Subscribe to offer this service for the Barcelona Municipality. This instance used IMI's infrastructure and, therefore, IMI assumed the role as Cloud Provider.

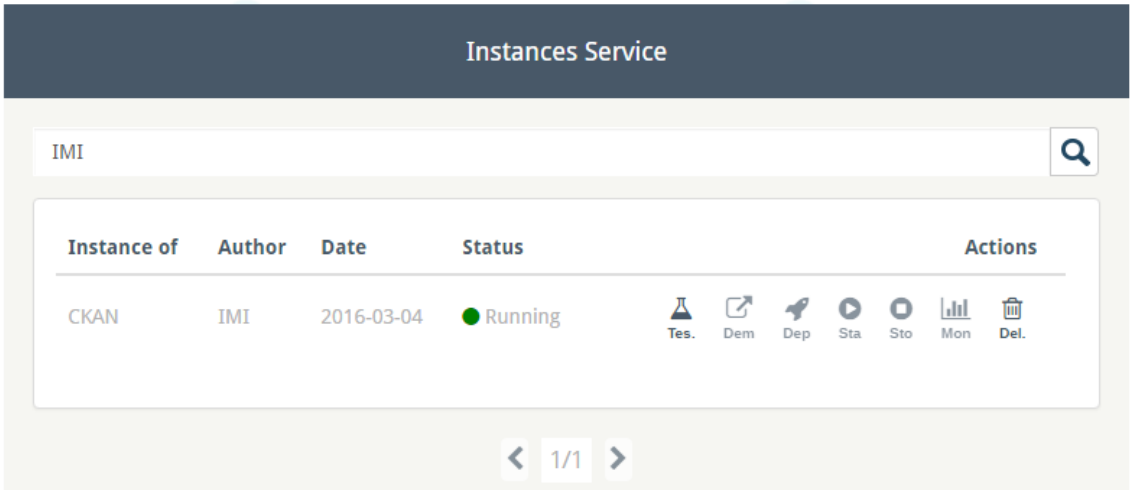

Figure 89: CKAN Barcelona instance.

3. Corby (as service subscriber) subscribed to the service creating a new instance of the CKAN. The second isntance was deployed by Corby in the same cloud as the Barcelona CKAN one. The Instance list shows then both active instances (Figure 90).

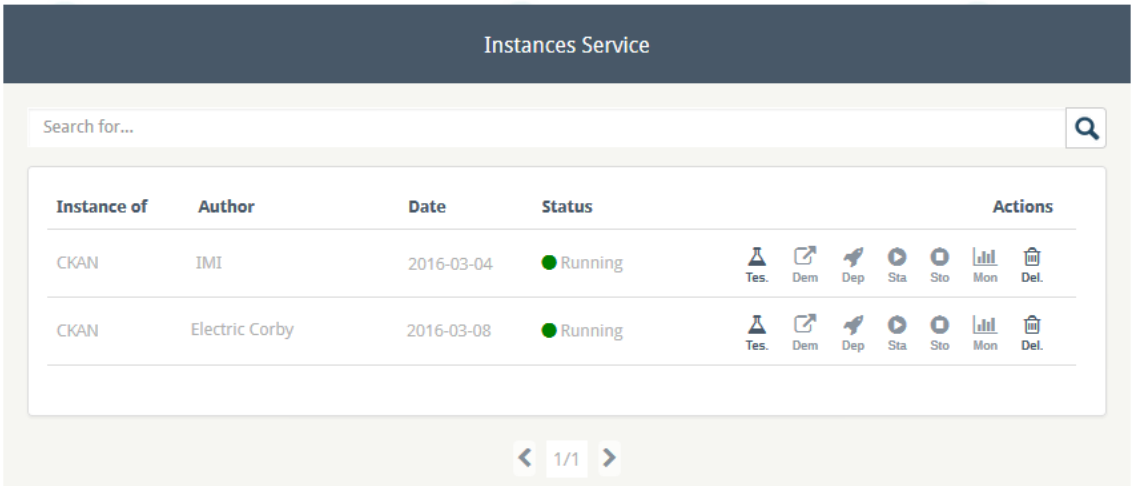

Figure 90: CKAN Corby instance.

When the instances were active and running, both portals were ready to be finalised and start managing users and datasets.

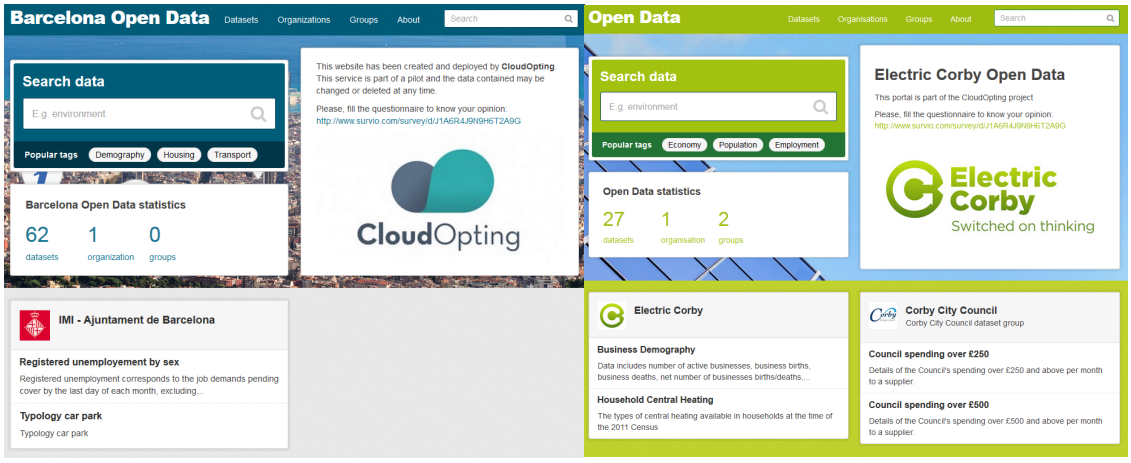

Figure 91: CKAN Corby instance.

5. Conclusions

In this document we had defined a methodology for the evaluation of the platform, an evaluation of four different pilots in four different cities, an evaluation of a set of additional services, and the associated replication plans. This document also presented the results of the evaluation questionaires, the results of the platform tests and the platform backlog.

After all the testing and analysis of the results, it can be concluded that the platform is acceptable and ready for production. A set of improvments have been identified that will upgrade the platform for better user experience.

On the technical perspective, the **evaluation of the platform** has been very positive, and it is proven that the technical approach chosen at the beginning of the project was correct. The technologies chosen are complient with the standard quality expected, and the technical features of the platform are wide enough to be operative. As a result of that, the platform is able to host a complete marketplace of public services whithin its catalogue of services ready to use, and ready to be deployed with little technical effort in many cities around Europe.

As mentioned in deliverable 4.2, the operation of the platform is low consuming and highly approviate for host a cloud marketplace.From the results of the feedback methodology some improvments about the high availability and clutering were suggested and included in a future

The **evaluation of the pilots** has been divided into a technical evaluation, a users satisfaction evaluation and a business approach evaluation gathered from the results of the series of Hackathons and the Barcelona Workshop.

Technical Evaluation of the Pilots has been very positive as has been shown in the respective Key Performance Indicators, that are very reliable in terms of the stability of the platform.

Users had concers about the reliability of the platform at the beginning of the Y3, but after a one year running the platform has demonstrated that it is robust and reliable.

Further from technical and operational results, the feedback revieved from potential users states that the strength of the platform will come with a set of appealing services. Thus, the set of services is something that has to be considered; which services fits consumers better and how they can suggest new services and be heard.

Also, as part of the technical evaluation of the platform we carried out in the **Hackathon technical day,** an end user evaluation of the platform relevant feedback was collected regarding: the adoption of technologies like Docker, about the scalability of the platform, the operating systems that Service

providers tend to use (linux vs Windows), the privacy concerns of the end users, the use of the cloud provider of the service provider's election, and the multi-instance approach of the architecture vs the multitenancy approach debate.

The **Users Evaluation** also had been very positive in terms of the user satisfaction gathered by the questionaires fullfilled.

The **Business evaluation** is a cross-domain evaluation gathered within the scope of the business approach stated in WP5. The results gathered from the discussions arisen within the public entities that were participating in the series of Hackathons can be resumed into several statements coming from different types of Public Administrations and Service Providers.

The most important aspect for **regional governments** is that there is an opportunity for CloudOpting in the form of a homologated entity that has undergone a process of public concurrence. There, local administrations from the region can choose among its homologated services without having to undergo a complex procurement process.

Also, **Aggregators of Municipalities said that i**t is fair and understandable that SMEs join CloudOpting if they consider that it is a competitive advantage for them to be in it. Nevertheless, in order to ensure free competition public administrations cannot make it a requirement.

Municipalities present in the hackathons said that the three presented business models are suitable and could be adopted by the Public Administrations in every country depending on the size of the organisation and the contracting policies and legal framework. It also has to be taken into consideration that the current procurement model of aggregation of demand has been proven that it is not the most suitable solution and that CloudOpting would have more chances of success with a procurement model based on the aggregation of offer.

As far as **Service Providers** is concerned, there is a very low access of SMEs to the European Market because they lack the resources and the channel to spread beyond their local area of influence. CloudOpting could support them with the mechanisms to access the market and gain visibility at the European level.

To conclude, the acceptance of the potential users of the platform is good, but the success of the project is directly linked to the enrollment of appealing services and to find a suitable and simple contracting procedure. The technical aspects of the platform, although always can be improved, are robust and strong enough to support the cloud marketplace.
To finalise, the CloudOpting project has been a really interesting project to the eyes of Public Bodies as well as for the Service Providers to enhance and make the European Single Market a reality.

6. Annex I Evaluation Questionnaire

One of the key aspects when evaluating the platform is linked to the services that it is offering. Qualitatively, it is necessary to know the feedback of the users who have been testing and using such services during the whole project.

When assessing the feedback of the users, it is possible to divide them into two groups: the first group contains the users who are directly related to the subscription of the services (Public Administrations), whereas the second, contains the users who ultimately benefit from using the service (end users).

Therefore, two evaluation questionnaires have been created:

- **Service and platform evaluation questionnaire**: This survey has been answered by Public Administrations and has been served to evaluate both the services and indirectly the platform.

- **End user service evaluation questionnaire**: This survey has been answered by citizens and has been served to evaluate the quality of the services from their perception.

The SaaS tool "Survio" has been utilised to develop the questions and get the sharable link to the users. The questionnaires have a general profile meaning that they are reusable for all the services deployed in the platform, thus, it can be possible to have the same criteria if a cross-services evaluation should be done.

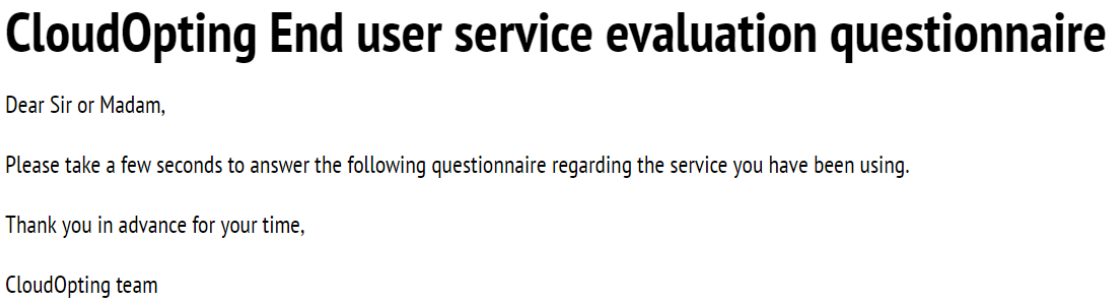
CloudOpting End user service evaluation questionnaire

Dear Sir or Madam,

Please take a few seconds to answer the following questionnaire regarding the service you have been using.

Thank you in advance for your time,

CloudOpting team

Figure 92: Look of the main page of the End user questionnaire

Platform and services evaluation questionnaire

The questionnaire for Public Administration follows the next structure, which has been implemented in order to be simple and appealing for the responders:

CloudOpting platform and services evaluation questionnaire

Dear Sir or Madam,

Please take a few seconds to answer the following questionnaire regarding the service you have been consuming on CloudOpting.

Thank you in advance for your time,

CloudOpting team

Figure 93: Look of the main page of the Service subscriber questionnaire

Which country is your organisation from?

- Italy
- Romania
- Spain
- Sweden
- United Kingdom
- Other

Which service have you been using?

- Business Portal
- Bus Portal
- CKAN
- City Agenda
- Clearó
- Energy Dashboard
- FixThis
- Next2Me
- Sentilo
- Take charge

How was the usability of the service/s interface/s?

- Excellent
- Very good
- Good
- Regular
- Unsatisfactory

Which was the degree of customisation of the service/s?

- Very high
- High
- Medium
- Low
- Very low

How was the technical support given?

- Excellent
- Very good
- Good
- Regular
- Unsatisfactory

How would you rate your overall satisfaction with the service/s?

★ ★ ★ ★ ★

How would you rate the monitoring system usability?

★ ★ ★ ★ ★

Would you renew the subscription of the service/s in a future?

- Definitely yes
- Probably yes
- I don't know
- Probably no
- Not at all

If you have a comment or suggestion regarding the service/s please let us know about it

End user service evaluation questionnaire

CloudOpting End user service evaluation questionnaire

Dear Sir or Madam,

Please take a few seconds to answer the following questionnaire regarding the service you have been using.

Thank you in advance for your time,

CloudOpting team

Which country are you from?

- Italy
- Romania
- Spain

- Sweden
- United Kingdom
- Other

Which kind of service have you been using?

- Mobile application
- Transparency portal
- Energy efficiency
- Other

How would you rate the usability of the service/s?

★ ★ ★ ★ ★

How would you rate the support given?

★ ★ ★ ★ ★

How would you rate this service/s compared to other available in the market?

★ ★ ★ ★ ★

How would you rate your overall satisfaction with the service/s?

★ ★ ★ ★ ★

If you have a comment or suggestion regarding the service/s please let us know about it

7. Annex II Platform Tests

This section presents the overall results obtained after the execution of the testing plan for the platform of the CloudOpting project.

7.1.1. Unit Tests

This testing level evaluates the main functions of a programming code. This is the lowest level of testing. The result obtained in each test has been classified as a Boolean, where "True" means the test has passed successfully, and "False" the test fails. In the latter case, it was also possible to indicate at which point it fails in order to make code debugging easier.

In CloudOpting project, these kinds of tests have been employed to evaluate the most critical functionalities, below listed, with the latest results obtained:

Test	Result
Crane alive	OK ✓
Docker status	OK ✓
Delete engine containers, images and crane data	OK ✓
Create build context	OK ✓
Create build context	OK ✓
Puppet module download	OK ✓
Docker images build	OK ✓
Docker images base build	OK ✓
Save Docker images in registry	OK ✓
Docker host and Docker cluster connect	OK ✓
Deploy containers	OK ✓

Table 37: Unit test results.

Currently, all of the above features have been tested and are properly working. As introduced before, the unit tests focused on the most critical functionalities of the platform. Since these features are critical, it is mandatory that all of them return a positive result in order to enable a correct operation of the platform. As shown in the Table 37, this objective has been successfully completed.

As a summary, the platform properly works at the unit level and this consequently allows the evaluation of the platform from a higher level as follows.

7.1.2. Integration Tests

This testing level evaluates the integration between the different blocks that compose the system. In CloudOpting, there are two main blocks: CloudOpting-Manager and CloudOpting-Crane.

These two blocks communicate through a component integrated into the Manager and the API REST of Crane.

To evaluate this integration, this component is tested at the code level, evaluating the main Crane functionalities from Manager.

Because these are code tests, like the tests of the unit level, the result of each test is a Boolean indicating whether the test has passed or failed.

For this level of test, it has been determined that the best way to test the integration is by simulating the different scenarios that could be presented. Each of these possible scenarios were composed for the use and the tests of the different necessary functionalities. They are not included for the reader's convenience.

Test	Result
Crane connection	OK ✓
Context creation with puppetfile and group name	OK ✓
Context creation with puppetfile but without group name	OK ✓
Context creation without puppetfile but with group name	OK ✓
Context creation without puppetfile and without group name	OK ✓
Check context by crane	OK ✓
Log context information	OK ✓
Log detailed context information	OK ✓
Context removal	OK ✓
Creation and check of Docker images	OK ✓
Building Stop of images	OK ✓
Docker images building without puppet	OK ✓
Local deployment of container	OK ✓
End to end creation and deployment	OK ✓

Table 38: Integration test results.

Currently, all the above features have been tested and are properly working and this means the two fundamental blocks of CloudOpting platform interact properly.

The positive results of the integration tests mean also Manager and Crane perfectly meet the requirements of all foreseen interaction scenarios. Since these scenarios make use of different functionalities, this leads to conclude the platform is running properly at functional level.

7.1.3. System Tests

The level system testing evaluates the system in its entirety as a black box in which the result, or a given outcome, is tested for a certain input.

For this kind of test, a software tool called Selenium has been used. Selenium is a software-testing framework for web applications and is able to automate the navigation of a web site and to evaluate if the offered results by this web correspond with expectations.

Thanks to this procedure, CloudOpting could be tested from its GUI, which is moreover the input and output interface that the user will employ.

Seven different tests have been designed to simulate the navigation of a user with seven different following objectives:

- Test service management
- Test search service
- Test user management
- Test instance list
- Test search organisation
- Test search user
- Test login

Each of these tests performs a complete navigation and allows the tests of different functionalities in the course of such navigation.

The results obtained in the last execution of these tests are the following:

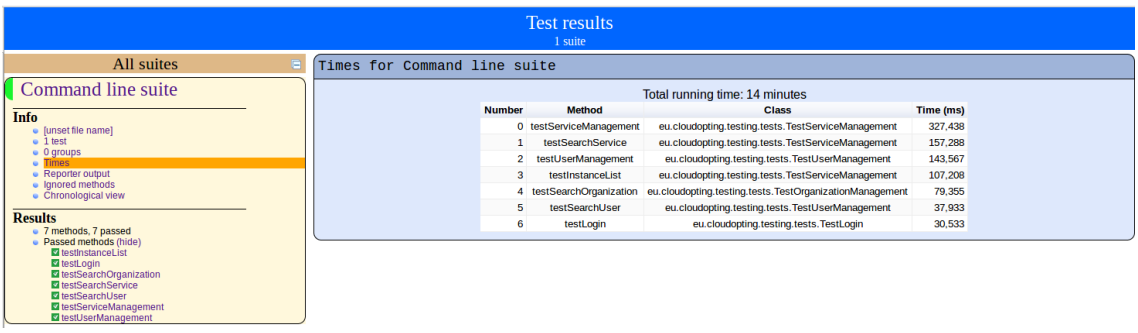

Figure 94: Results and execution time of the system tests.

Seven tests have passed satisfactorily as shown in the blue frame with the indication of the time taken to execute each test.

The examination of the obtained results shown in the above figure confirm that the platform behaves as expected for all targeted navigation scenarios and this is a clear guarantee of an adequate user interaction and the user-friendless of the platform GUI. This also implies that the platform internally responds correctly at the functional level.

7.1.4. Performance & Stress Tests

In order to test the performance of the platform, we used an online tool called "webperformance" that hits the pages you specify.

Here are the paràmetres that we used in order to perform the test:

- 5 minutes. This platforms perform test for 5 full minutes
- 500 concurrent users màximum. It starts with 125 and goes to 250, 375 and 500.

- 2 seconds goal per page màximum.

Here you can see a summary of the performance Goal Analysis:

User Level	Avg Dur	Average Wait Time	Failure Rate
125	✓	✓	✓
250	✓	✓	✓
375	✓	✓	✓
500	✓	✗	✓

Figure 95: Summary of the performance Goal Analysis.

Slowest Pages/Steps by Average Duration

The chart and table below show the slowest pages/steps, as determined by the page/step durations measured during the test. The table lists the pages/steps along with the time that the peak duration was recorded, the number of users running the testcase at that time and the base duration of the page/step (as recorded). Refer to the related sections for more detailed analysis of each listed page/step and testcase.

User Level Analysis

This section summarises the metrics by user level.

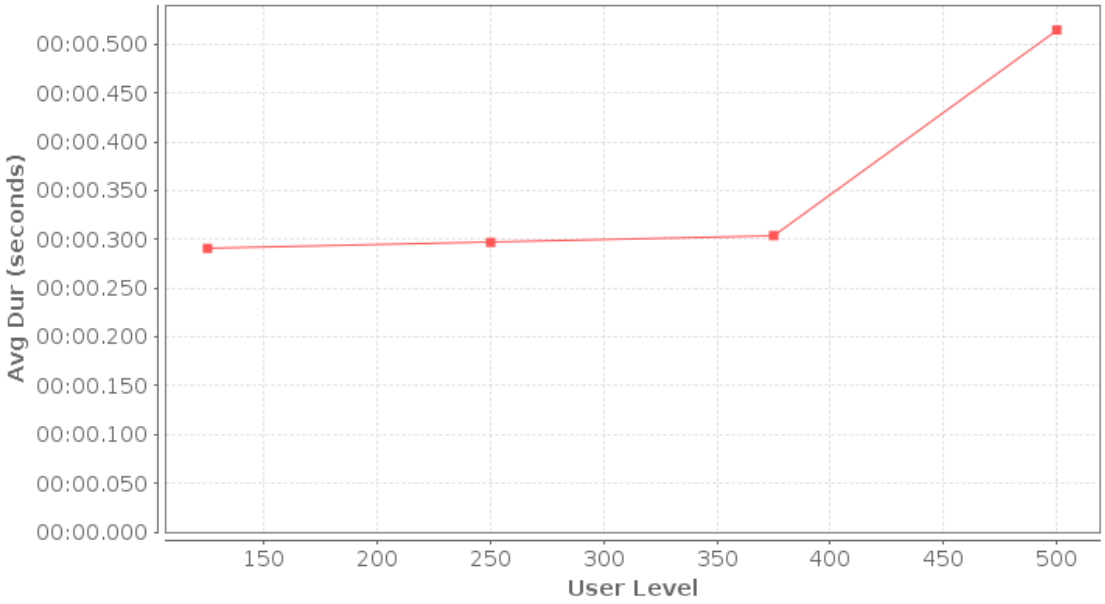

Figure 96: Serving time per page average per users.

Time-based Analysis

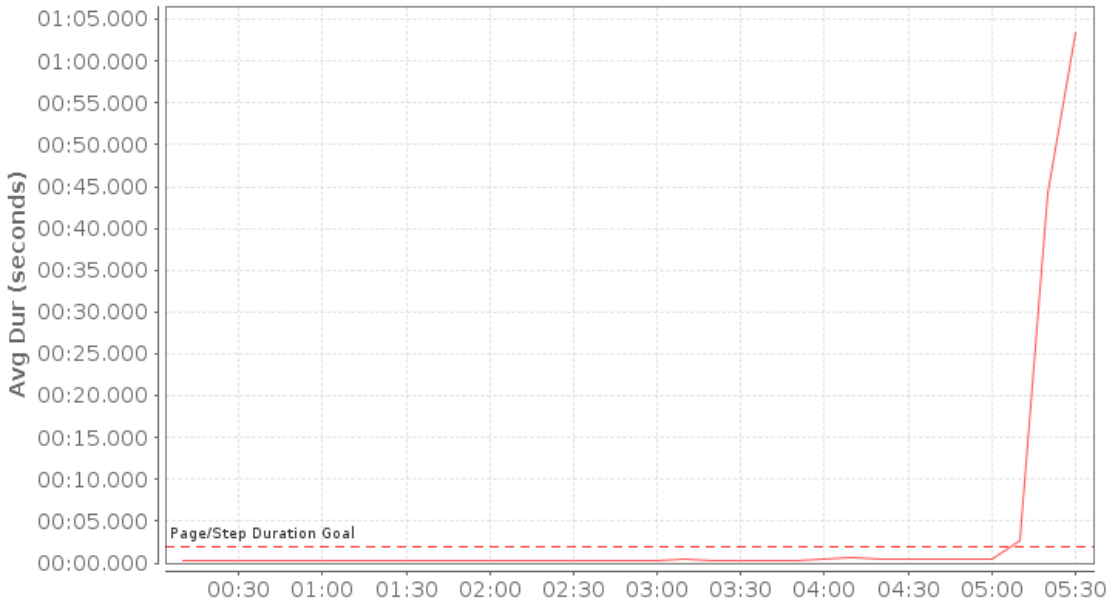

Figure 97: Page/Step duration over the testing period.

Page/Step Duration

The duration chart shows the range of durations recorded for the item. The minimum, maximum and average values are computed for items that completed within the sample periods summarised for each user level.
It shows that from 125 to 375 concurrent platform users pagesa re served with less than 2 seconds, but with 500 concurrent users, sometimes it takes almost 16 seconds.

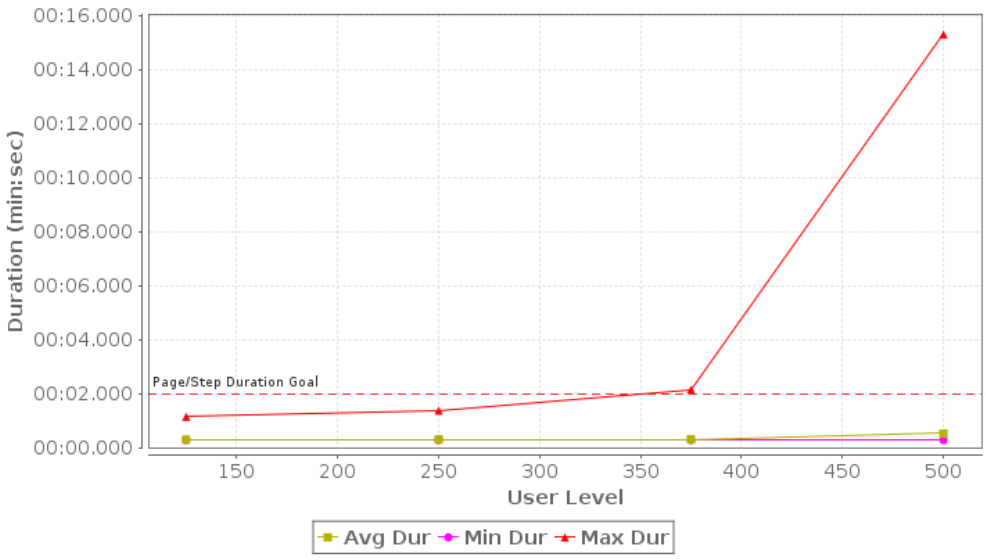

Figure 98: Page/step duration per number of users.

Failures

The failures chart shows the number of failures recorded for the item during the sample periods summarised for each user level and the rate of failures compared to the total attempts for this item. A failure of a page/step or transaction could indicate a wide variety of problems from an inability to establish a network connection to the server to the failure of a validator configured to check the returned content.

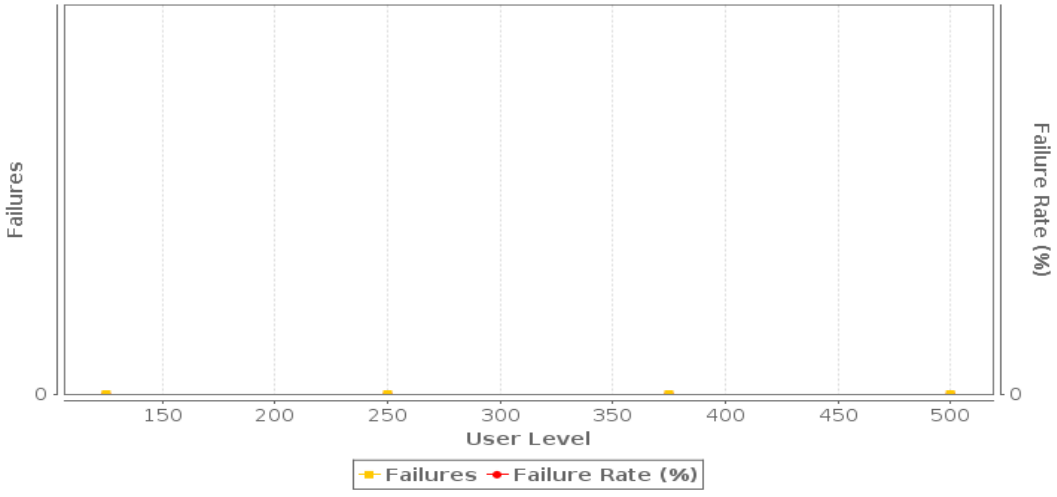

Figure 99: Failures/Failures Rate per number of users.

Waiting Users

The detailed durations chart plots each duration recorded for the item during the test. Gaps in the chart could indicate a period where an item was not being serviced by the application even though there were outstanding requests for the item. Used in conjunction with the Waiting Users chart, this can be used to identify stalling behavior on the server.

According to Figure 100, with 500 concurrent users the response time reaches almost 3 seconds, and this is out of the limits fixed before ("2 second goal" parameter).

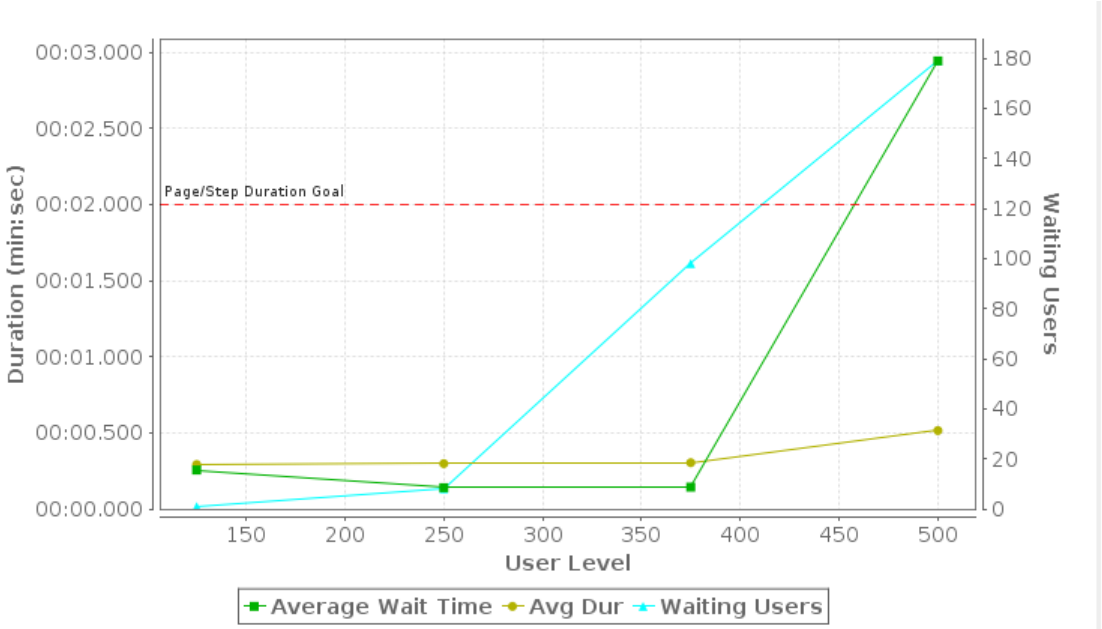

Figure 100: Duration/Waiting Users per number of users graph and table.

User Level	Successes	Failures	Failure Rate	Avg Speed	Avg Size	Min Dur	Avg Dur	Max Dur	Waiting Users	Average Wait Time	Dur Std Dev
125	3,870	0	0.00%	491.6 kbps	17.5 KB	00:00.277	00:00.291	00:01.139	1	00:00.256	00:00.067
250	7,593	0	0.00%	484.0 kbps	17.5 KB	00:00.275	00:00.297	00:01.368	8	00:00.140	00:00.079
375	11,294	0	0.00%	476.0 kbps	17.5 KB	00:00.273	00:00.304	00:02.145	98	00:00.143	00:00.157
500	22,631	0	0.00%	280.9 kbps	17.5 KB	00:00.274	00:00.515	00:15.317	179	00:02.943	00:00.969

Figure 101: Data from graph on figure 100.

Time-based Analysis

Page/Step Duration

The duration chart shows the range of durations recorded for the item. The minimum, maximum and average values are computed for items that completed within a sample period.

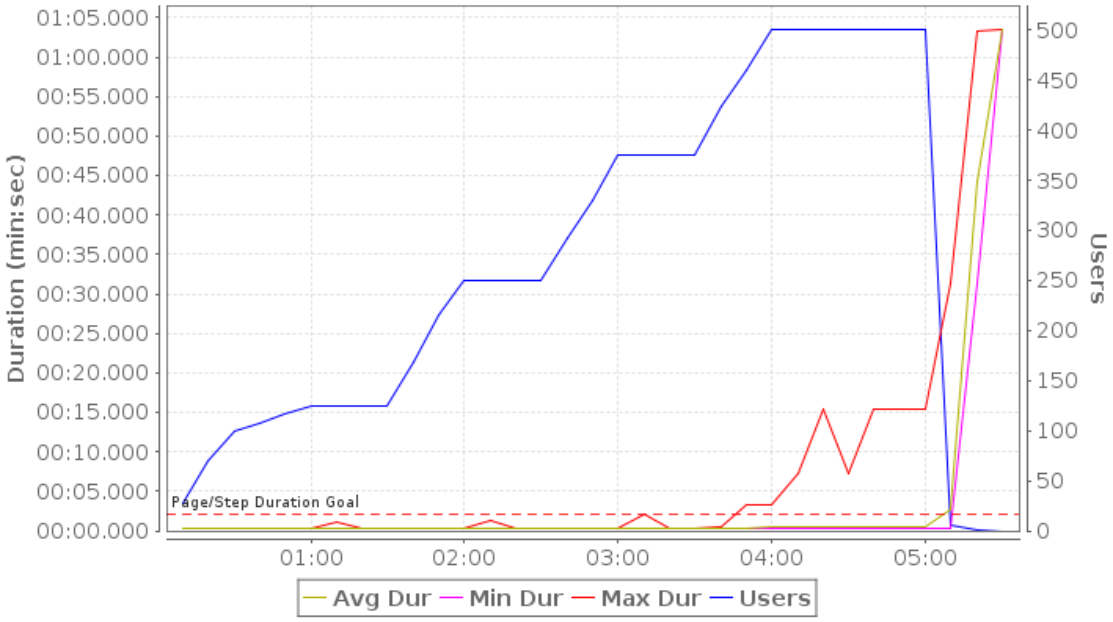

Figure 102: Duration/Users over the duration of the test.

Failures

The failures chart shows the number of failures recorded for the item during each sample period and the rate of failures compared to the total attempts for this item. A failure of a page/step or transaction could indicate a wide variety of problems from an inability to establish a network connection to the server to the failure of a validator configured to check the returned content.

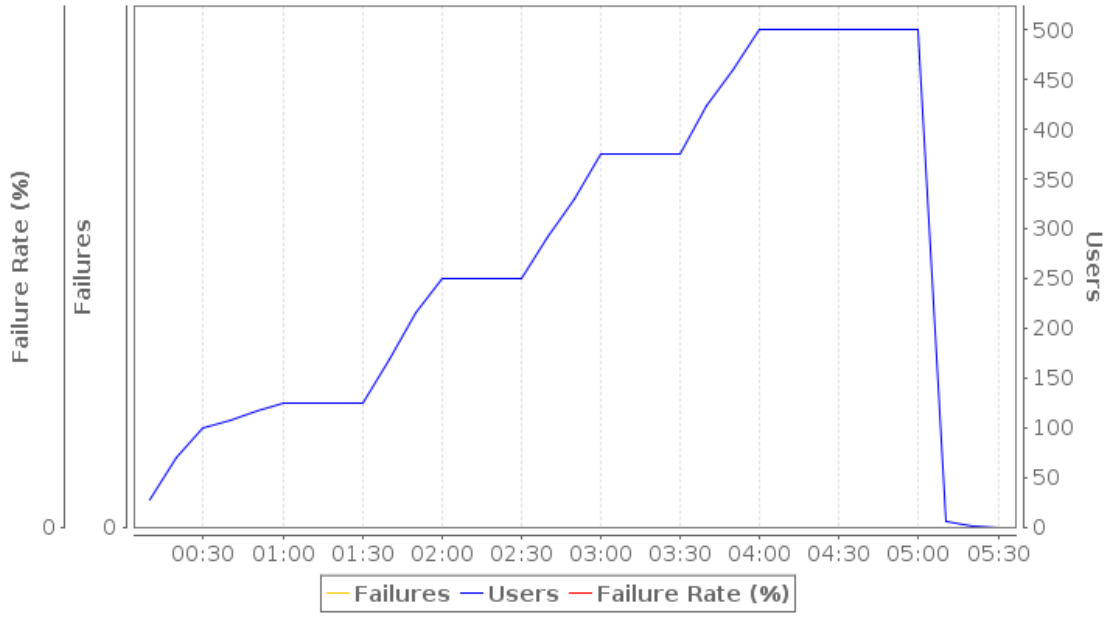

Figure 103: Failure/Users over the duration of the test.

Waiting Users

The *Waiting Users* and *Average Wait Time* metrics help diagnose certain types of performance problems. For example, they can help determine what pages/steps users have stopped on when a server becomes non-responsive. The *Waiting Users* metric counts the number of users waiting to complete the item at the end of the sample period. The 'Average Wait Time' describes the amount of time, on average, that each of those users has been waiting.

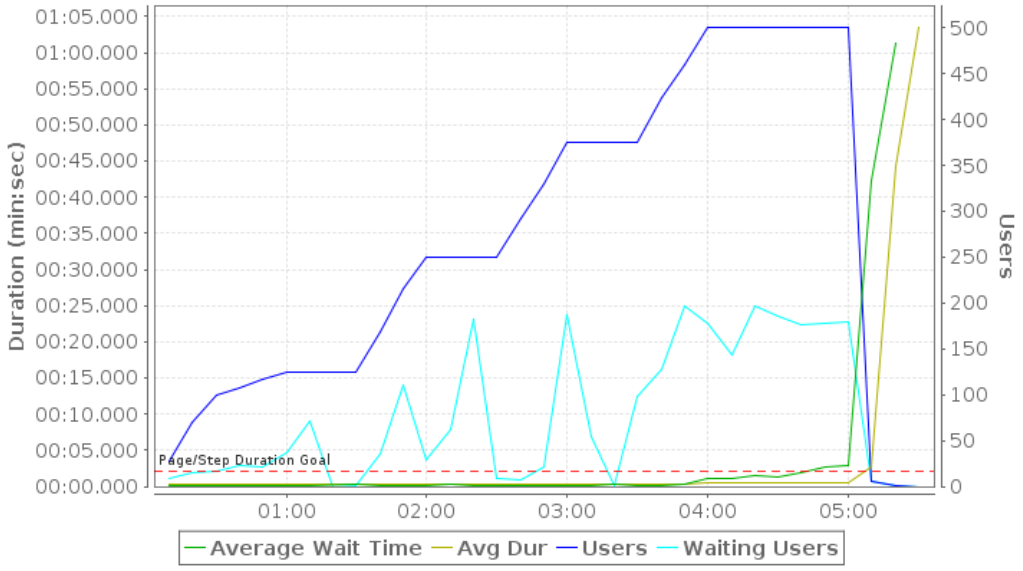

Figure 104: Duration/Users over the duration of the test.

Load Engines

The Load Engines section provides metrics about the performance of the load engines during the test.

Users

Number of users running on the engine(s)

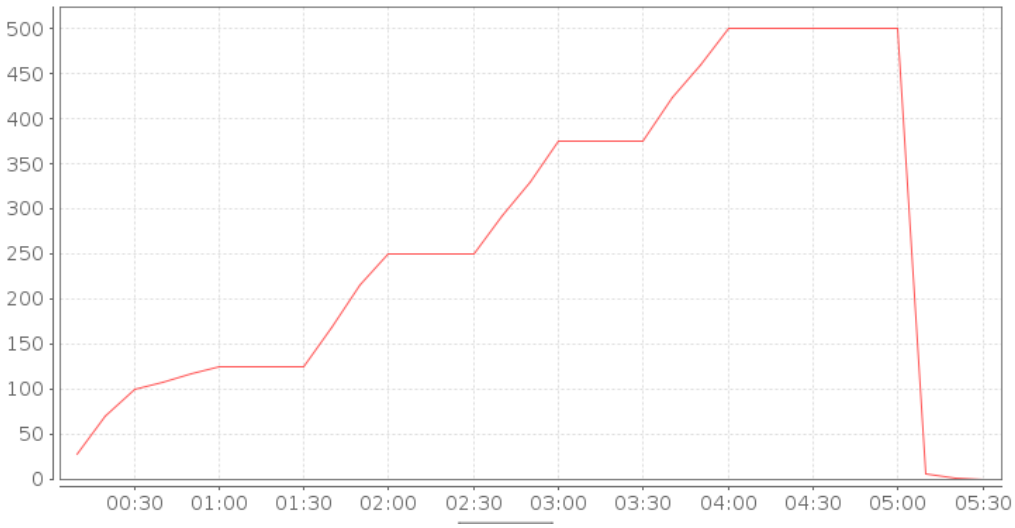

Figure 105: Users over the duration of the test.

Total Capacity

Estimated user simulation capacity of the engine(s)

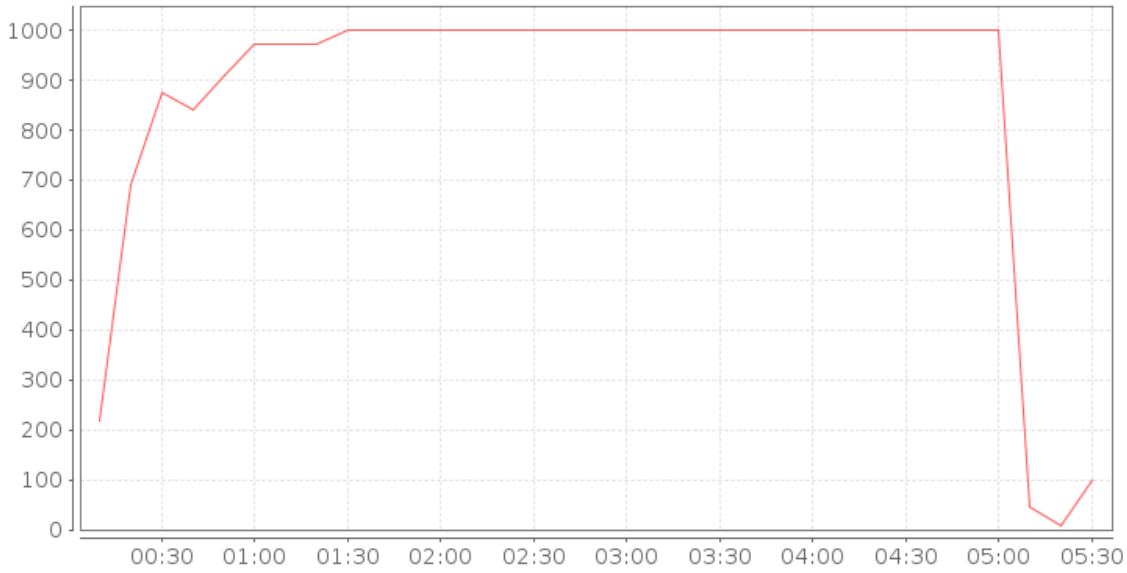

Figure 106: Estimated user capacity of the engine.

CPU %

CPU utilization % at end of sample period

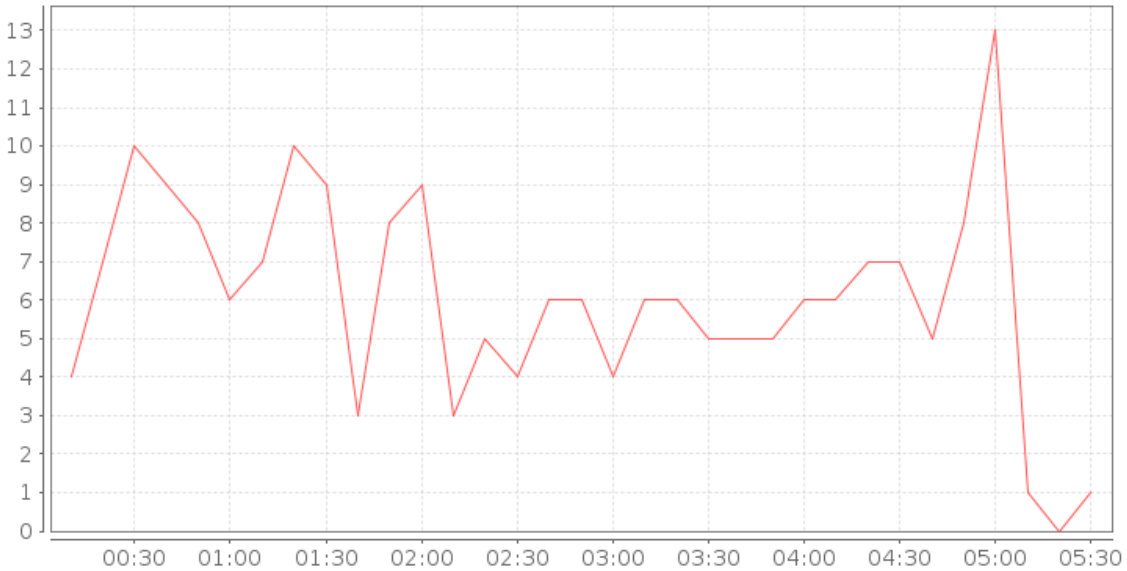

Figure 107: CPU consumption over the duration of the test.

VM Memory %

VM Memory utilization % at end of sample period

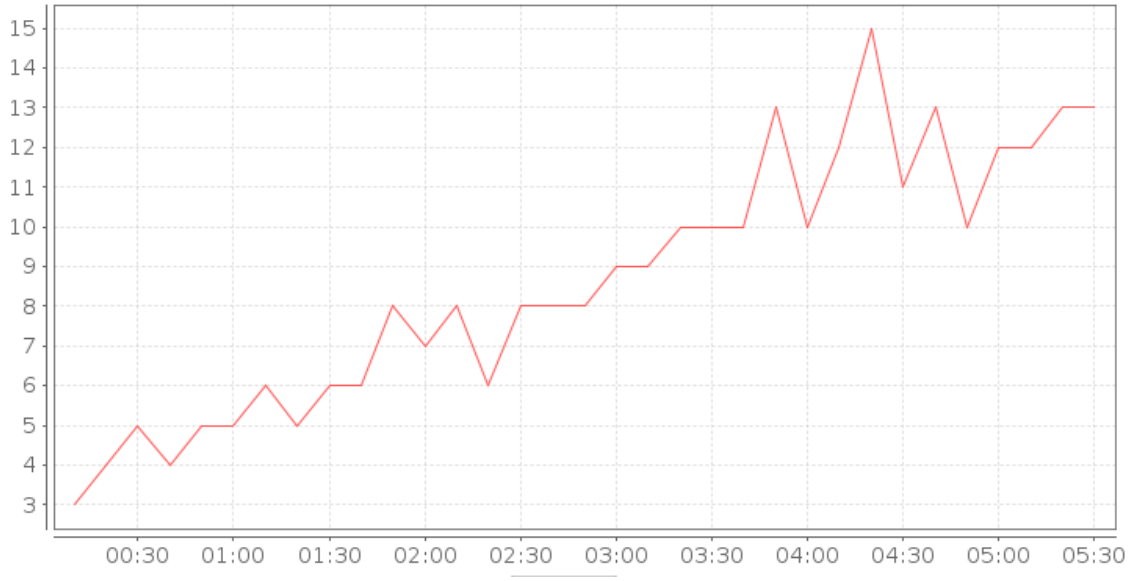

Figure 108: Virtual Memory usage over the duration of the test.

OS Memory %

OS Memory utilization % at end of sample period

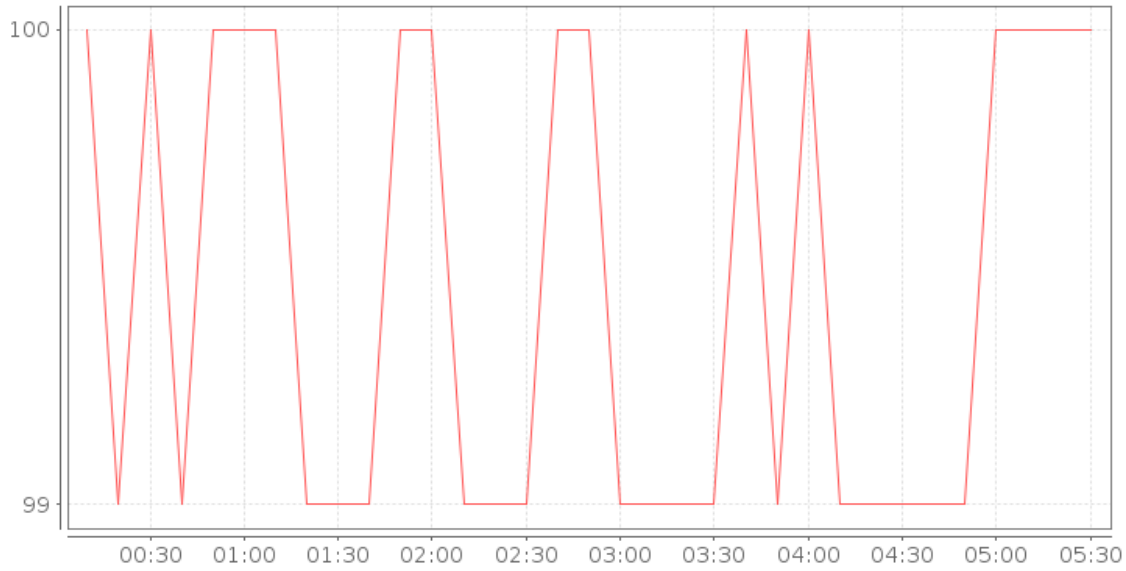

Figure 109: Operative System memory usage over the duration of the test.

Figure 109 shows a graph where a busy CPU can be seen during the testing, even with a low number of users. It may mean that the platform may need a bit more of CPU in order to go smoothly.

Bandwith In

Total data transfer rate (in to the load engine) during the sample period

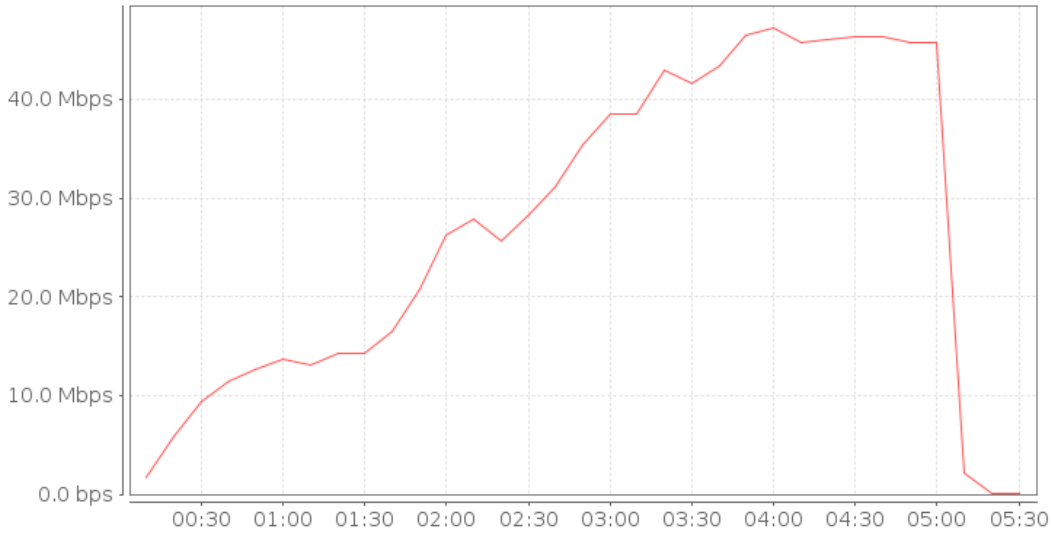

Figure 110: Bandwidth in usage over the duration of the test.

The amount of data being sent to the server is big, but also the number of concurrent users in the testing. If the number of users was, in average, more than 250, maybe the bandwidth of the server will have to be modified.

Bandwith Out

Total data transfer rate (out from the load engine) during the sample period

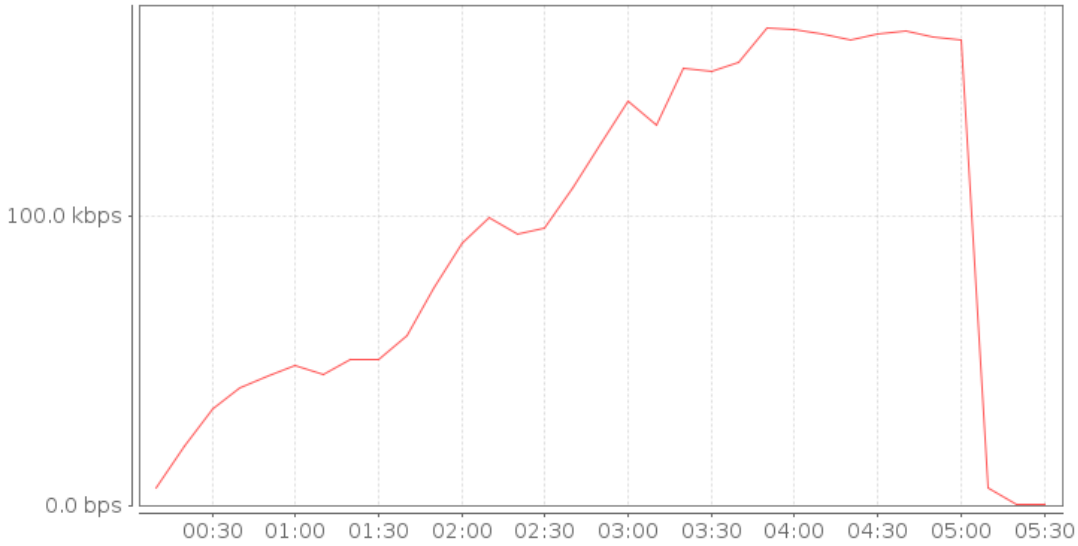

Figure 111: Bandwidth out usage over the duration of the test.

The figure above shows that the amount of information going out is quite huge, but should not be a problem taking into account that the number of concurrent users in the testing was quite larger that the initial expectations.

Steal %

The time that the guest CPU had something runnable, but the hypervisor chose to run something else instead (applies only to systems running on virtual machines)

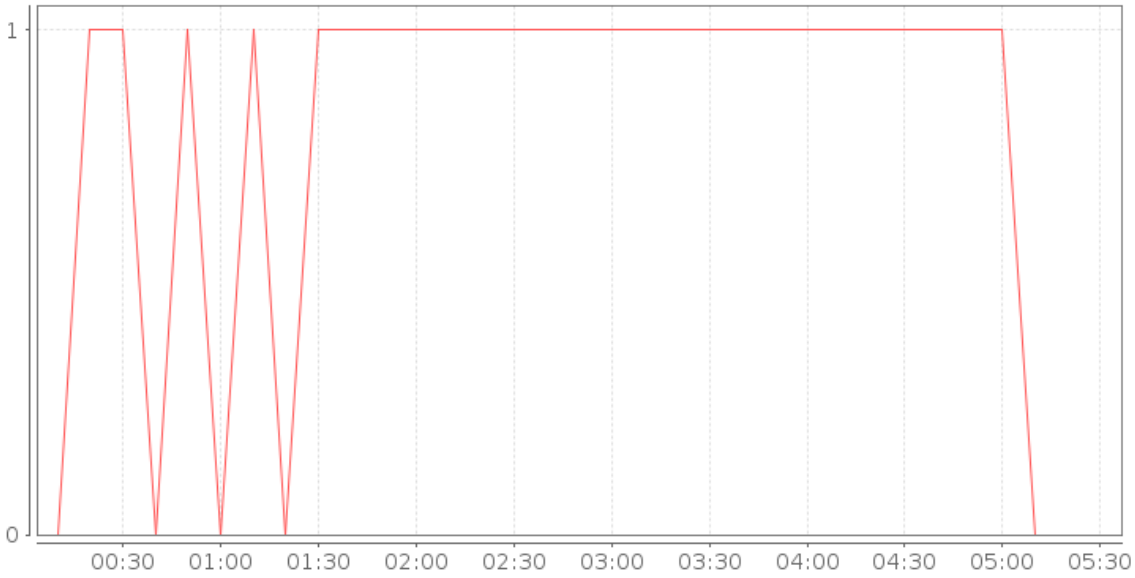

Figure 112: Steal % over the duration of the test.

The conclusion that can be extracted from these results is that the platform can hold a concurrency of 375 users without being afected. That does not mean that it could not handle more, but that the waiting time will not be satisfactory for our users.

8. Annex III Platform Backlog

8.1. Technical feedback

The following technical improvements have been identified during the 12 months of running the pilots:

- HW resources should be flexible for the user. The IaaS console should be offered to the user of the CO platform to be enhanced with this possibility. Billing process should take this into consideration. The use case is that either the user did not select the appropriate information when subscribing to the service or the usage of the service has increased

- Platform should enable the service publishers to easily upgrade the service running for the service subscribers. Having a small amount of service subscribers can have the update done manually, but for more subscribers this task may become impractical

- CloudOpting base image for web proxy should include out-of-the-box support for SSL communications. This will enable the service subscriber to use more features of the subscribed service in a more secured manner.

8.2. Business feedback

From the WP5 point of view, CloudOpting Platform will play a major role in making CloudOpting a successful European marketplace of digital services. Therefore, it is vital that the platform and its user interface adapt to the needs of its costumers while being appealing and offering all the functionalities that are expected nowadays from a marketplace such as this.

To this end, WP5 has been gathering constant feedback from different sources such as:

- **External** actors to the project (public administrations, app developers, cloud providers, …) which have been contacted through events, like the Mobile World Congress, taking into special consideration to the CloudOpting Hackathon, which was held in different locations in Europe and where representatives of all the stakeholders of the project where brought together to discuss CloudOpting.

- And **internal** continuous analyses of the market which pretend monitor the competitor's evolution and track their "movements" and detect interesting functionalities that could be added to CloudOpting.

Based on all the feedback gathered, it has been identified various aspects that the platform should incorporate in order to be successful. Afterwards, a series of issues were opened in Github so the developers of the project could implement them.

The issues can be divided into the following categories:

- **Aggregator of public administrations**, as it has been explained in deliverable D5.2, aggregators are one of the target customers of CloudOpting for several reasons. Therefore, the CloudOpting platform has to be able to incorporate this kind of customers and fulfil all their needs.

  An example could be that when a Service Customer subscribes to a service "A", the Aggregator didn't have the option to select/deselect public entities to which the service "A" was offered. So the issue #375 was developed and called "[Aggregators] select/deselect PA of a service".

- **User Interface** is an Interface in which users can navigate through the Service Catalogue. An easy navigation through the platform has to be a main point to achieve in order to create a more appealing to users interface. Besides, the platform has to incorporate added value functionalities such as the one covered by the issue #349, which allowed to a Service Subscriber to rate a service and make it visible for all users.

- **CloudOpting Manager**, is used for the deployment and customisation of the applications offered in the Service Catalogue. So it has to be as easy and practical as possible. An example of improving the easiness of it, is the issue #293 which created filter buttons for the vertical solutions in order to speed up the search of solutions.

8.3. Backlog

Here a list of the most relevant tasks extracted directly from the GitHub that have been detected in order to improve the platform and tagged to be done in the future (backlog):

- Include a search engine in the catalogue
- Service Catalogue - Create a screen for Service Providers to control subscriptions
- Service Catalogue - Create a screen in the private profile to enable users management
- Service Catalogue - Create a sheet in the private profile with the services subcribed/provided by the organisation

- Service Catalogue - Create a section in the private profile to show the status of the organisation
- Service Catalogue - Create sheet in the private profile to update details in the public profile
- Service Catalogue - Create a public profile for Service Subscribers
- [Aggregators] Monitoring comparison by service and public entity
- [Aggregators] Monitoring metrics in a global and/or particular views
- [Aggregators] Trial period of a service
- [Aggregators] select/deselect PA of a service
- Service Catalogue - Group sizes of a service in one service main page
- TOSCA IDE - Be able to add new kind of nodes
- 2 levels within an organisation (Aggregators)
- Write public comments for the services
- Create a public profile of the Service Providers
- When browsing a service any user should be able to see its rating
- A Service Subscriber should be able to rate a service
- Translate images to FRENCH
- Translate to FRENCH
- Filter buttons for the vertical solutions
- modify registration form to add also organisation
- add organisation admin role
- Type in organisation needs to be defined and managed
- REST API versioning