Environment Administration

City of Gothenburg

The LoV-IoT project: Air and water monitoring with Internet of Things

Data platform

Report number R2020:20

Preface

The project Air and water monitoring with Internet of Things, LoV-IoT is an innovation- and development project which has examined the possibilities in using sensors and Internet of Things (IoT) to develop the environmental monitoring of air and water within cities. One aim of the project was to develop an effective system for gathering information on air and water quality in cities to contribute to better health among the citizens.

The project was running for three years, between autumn 2017 until autumn 2020 and it was financed by the strategic innovation program IoT Sverige, as a part of their work within IoT for societal benefits.

This report will describe the work done within work package 6, which worked with data platforms within the project. The report is written by Hagström consulting, together with IVL, RISE and the City of Gothenburg.

The LoV-IoT project: Air and water monitoring with Internet of Things
Data platform

City of Gothenburg, Environment Administration

Author: Author: Björn Hagström Hagström Consulting AB, Joel Wanemark IVL Svenska Miljöinstitutet AB, Lars Samuelsson Miljöförvaltningen, Hampus Bok Miljöförvaltningen, Ågot Watne Miljöförvaltningen

Photos/figures: Produced by the project

ISBN No: 1401-2448

If you want to use text or images from this report, please quote:
The City of Gothenburg, R2020:20 The LoV-IoT project: Air and water monitoring with Internet of Things Data platform

This is a report in the environmental administration's report series. The full series of reports can be found on https://goteborg.se/mfrapporter

Summary

Work Package 6 (WP6) of the LoV-IoT project focused on the central technical platform that WP4 and 5 could use to store, manage, visualize and analyze data. WP6 worked closely with WP4 and 5 in order to make sure that data from the measuring stations were sent in a way that could be received and standardized. The platform used was Fiware that also supplied the data models.

WP6 had 3 deliverables, L6.1 to L6.3. All deliverables were completed within the project. Deliverable 1 and 2 are detailed in this report and deliverable 3 can be found in an attached report.

Actors taking part in WP6 were The City of Gothenburg, IVL, Vinnter, Ericsson, Talkpool, IMCG, The City of Uppsala and the work was being led by Hagström Consulting.

All three project deliverables were completed within the project. The work itself was conducted in close cooperation with project partners. In the beginning a wider selection of partners were involved while a smaller group were responsible to get Fiware up and running and data flowing. WP6 mapped users and needs and used that as a basis for the development.

The technical platform used was Fiware that is open source, which was an important prerequisite for the City of Gothenburg. It also fit well with the project that had no budget for purchasing licenses. In the instances where Fiware were not enough we used additional open source products to achieve our goals.

Data was transformed to conform to Fiware data models and we used two APIs for accessing the data. A third way to access data was through a visualisation created in Grafana.

The diagnostics platform was an important deliverable and it was implemented to monitor sensor behavior and warn about unexpected behavior as well as creating one sensor specific and one general algorithm to try to increase the quality of the data.

There was interest in using our data from several parties. Since we did not publish it publicly no live usage happened but the fact that there is clear interest raised the city's awareness of the importance of sharing data. As a result the city now shares data from 3 high quality measuring stations, up from only 1 before the LoV-IoT project started.

Sammanfattning

Arbetspaket 6 (AP6) inom projektet LoV-IoT fokuserade på den centrala tekniska plattform som AP4 och 5 kunde använda för att lagra, hantera, visualisera och analysera data. AP6 arbetade i nära samarbete med AP4 och AP5 för att säkerställa att data från mätningarna skickades på ett sätt som gick att ta emot och standardisera. Plattformen som användes heter Fiware, och den försåg även dataplattformarna som användes.

AP6 hade tre leverabler; L6.1 till L6.3. Alla leverabler genomfördes inom arbetspaketet. Leverablerna ett och två är presenterade i denna rapport och leverabel tre är bifogad denna rapport.

Organisationer som deltog i arbetspaketet var Göteborgs Stad, IVL, Vinnter, Ericsson, Talkpool, IMCG, Uppsala stad. Arbetet i arbetspaketet leddes av Hagström Consulting.

Alla tre leverabler genomfördes inom arbetspaketet. Själva arbetet gjordes i nära samarbete med projektpartners. I början var en större mängd deltagare involverade, medan en mindre grupp var ansvariga för att få Fiware att fungera. AP6 kartlade användarnas behov och använda det som en grund för utveckling.

Den tekniska plattformen som användes inom projektet var Fiware, som har en så kallad ”open source”, något som var viktigt för Göteborgs Stad. Det passade även väl eftersom projektet inte hade en budget för att köpa in licenser. I de fall där Fiware inte var tillräckligt lades ytterligare ”open source”-produkter till för att uppnå målen.

Data transformerades för att överensstämma med Fiwares datamodeller och vi använde två API:er för att få tillgång till datan. Ett tredje sätt att få tillgång till data var genom en visualisering skapad i Grafana.

Diagnostikplattformen var en viktig leverabel och den var implementerad för att övervaka sensorbeteende och för att varna för oväntat beteende. Diagnostikplattformen hade en sensorspecifik algoritm och en mera generell algoritm för att försöka uppnå en högre datakvalitet.

Det fanns ett intresse från flera aktörer att använda sig av datan som togs fram inom projektet. Eftersom vi inte publicerade datan publikt kunde ingen realtids-användning ta plats, men till följd av aktörerna intresse för realtidsdata fick Göteborgs Stad en ökad medvetenhet kring hur viktigt det är att dela data. Som ett resultat delar nu staden sin data från tre av sina ordinarie mätstationer för luft, jämfört med att data från endast en station delades innan projektet LoV-IoT startade.

Content

1 Project deliverables

All formal requirements for work package 6 have been met.

1.1 Deliverable L6.1

The deliverable is “En plattform för luft- och vattenkvalitetsdata, som är sökbar för M2M baserat på internationella standarder” that translates to “A platform for air- and water quality data that is searchable by M2M based on international standards”.

This has been achieved by Fiware and use of their standard data models.

1.2 Deliverable L6.2

The deliverable is “Diagnostikplattform för validering av data som hanteras i plattformen” that translates to “A diagnostic platform for validating data handled by the platform”.

This has been achieved by using a python based docker image.

1.3 Deliverable L6.3

The deliverable is “Rapport som beskriver koncept för att använda sensordata för att validera spridningsmodeller” that translates to ”R2020:21 Koncept för att använda sensordata för att validera spridningsmodeller”.

This has been achieved in a separate report found in the environmental administrations report series.

2 Start and planning

2.1 How the work was organized

The LoV-IoT project started off with discussions across work package borders where we as a team created an overview of needs, goals, possibilities, limitations and work processes. With WP6 we started off without preconceptions by creating diagrams of information flows and technical components and how they were related to each other.

After the initial process we limited ourselves to the participants in WP6 and a few invited contacts from other work packages for coordination purposes. During this phase we talked a lot about technical choices that needed to be made. The fact that the City of Gothenburg has open source as the first hand choice guided us, and as the city was already working with Fiware in several other projects it made it a natural choice. The LoV-IoT project did not have any budget for development or procurement, which made open source solutions even more appealing.

During the implementation of Fiware WP6 scaled down to 4–5 core participants necessary for the completion. We had biweekly meetings to make sure that the work progressed as needed.

During the entire project regular meetings with all WP leads have been held to solve any issues that arose.

2.2 Target groups and needs

To understand the needs concerning information sharing and its impact on information gathering we created an impact map to map target groups and needs.

The impact map was extended as new users and needs were discovered until the implementation of Fiware was started. The impact map was a valuable tool in keeping us focused towards the right goals.

The entire LoV-IoT project was by its very nature an exploratory project. This focus meant that it was not always the target groups' needs that were guiding us as learning and developing could be a goal in and of itself. These kinds of needs were not represented in the impact map.

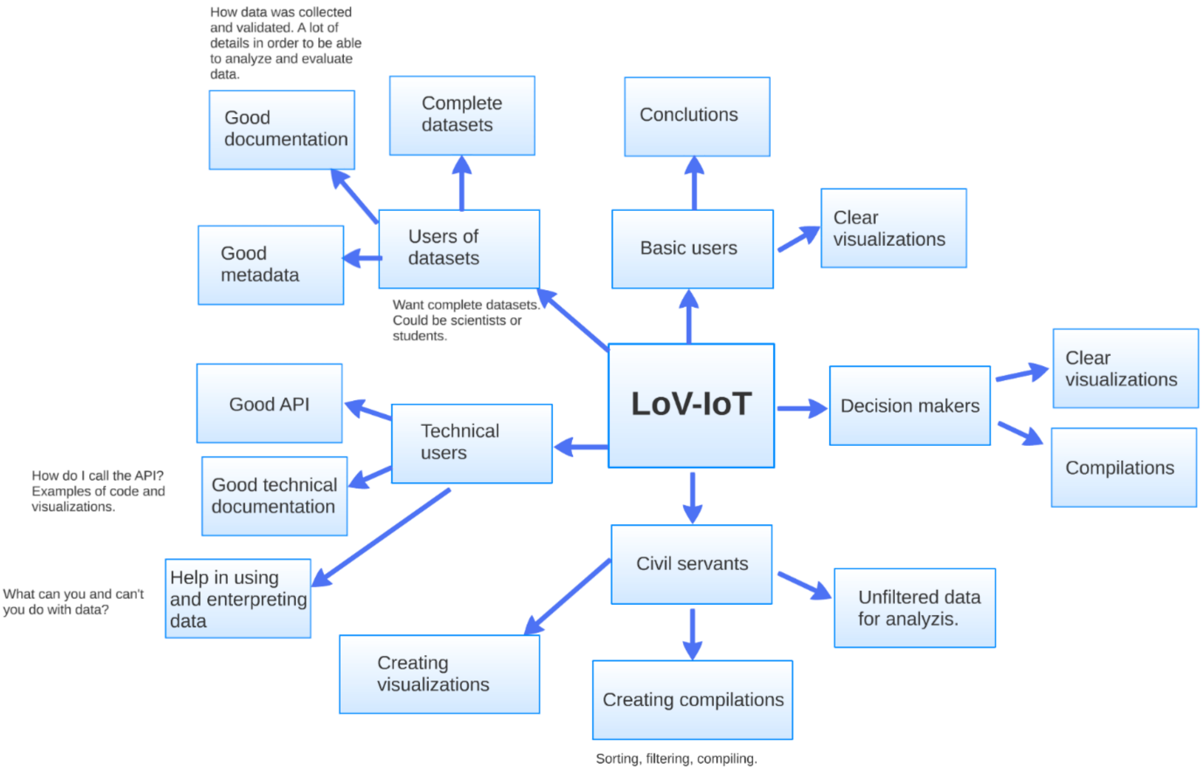

Figure 1. Impact map of target groups and their needs.

2.3 Choice of platform and the reasoning behind it

LoV-IoT was a follow up project to the predecessor Miljöväder. The intent was for the LoV-IoT project to build on the technological platform used in Miljöväder, but we quickly realized that the value of doing so was too low. A contributing factor for the decision was the fact that the license for the technological platform had expired and we needed to go through the process of procuring a new platform and the City of Gothenburg had a policy of always evaluating open source software. That would have required money and resources we did not have. The fact that the project did not have a budget for hardware or software made this route impossible for us. We chose to discard the old platform in favor of something new, but we took a lot of the lessons learned from Miljöväder with us into the LoV-IoT project.

When the decision to move on to a new platform was taken a new exploratory phase with open discussions about the road ahead started. Different techniques were put forward and discussed. We identified information flows that needed to be supported and used that as a basis for our discussions. Several of the participating parties had their own platforms that they were interested in using in the project. The fact that the City of Gothenburg had a policy of evaluating open source systems as an alternative was something we took to heart. The City of Gothenburg was also involved in several other projects focusing on Fiware and that steered us in that direction. We saw the possibility of pooling resources and learning from each other and to be a part of an existing development within the city which would make our work a lot easier. IVL that would contribute the most resources for implementing the platform were very interested in learning more about Fiware too. The EU were quite clearly pushing for Fiware to be used by the European countries and that was also a factor we took into consideration.

One important point of discussion concerned if we were trying to create a platform for the project or a platform that could be used by all municipalities in time. We came to the conclusion that this is an exploratory project and a broader implementation can learn from our progress but will need to create its own platform. We did not have the resources for that kind of scaling.

Understanding Fiware and what it is and what it is not was a bit challenging. Something that made it harder for us was that Fiware itself was going through a cleansing process where many of the old, often unsupported, components were discarded and a new focus on the core components was set. It is hard to find where the discussion about Fiware is online and we are still missing this. There are several mailing lists used for different aspects of Fiware but many of them are inactive. We came to the conclusion that Fiware is three different things:

1. The data models
2. The context broker and core components
3. The api to connect different systems to the context broker

We were most interested in Fiware's data models, more than in the technological platform itself although we used all three parts of Fiware in the project.

3 Exploration of technical choices

3.1 Choosing Fiware components

When choosing which Fiware components that would be part of the IoT platform, the focus was more on the IoT than the air and water quality monitoring. The project had identified a value in learning more about the different IoT components of Fiware as these could have other applications as well. The setup would for example need to handle both large time series of data, secure communication, and auto-provisioning of sensors so that each sensor didn't have to be registered in the platform manually. The choice was also based on the maturity of components and their respective level of document-tation as of June 2018. Certain setups and configurations have been updated in the official Fiware guides during the project duration, but these updates have not been implemented in the project setup.

3.2 External components

Where necessary IoT components were non-existent or Fiware components were seen as unsuitable or immature, external open source components were used instead. These included components for visualization (Grafana), MQTT broker (mosquitto), and data transformation layers. The data transformation layers were custom built docker images1 for the LoV-IoT application.

3.3 Setup

An overview of the system setup and the component ecosystem is shown in the figure below and details are provided in the following sections. As previously described, the system was hosted on a Linux server within the IntraService department of Gothenburg City. There was also an unsuccessful test using VMware vSphere, details of that can be found in Appendix A.

[1] Available as open repositories at docker hub user *joelwanemark*.

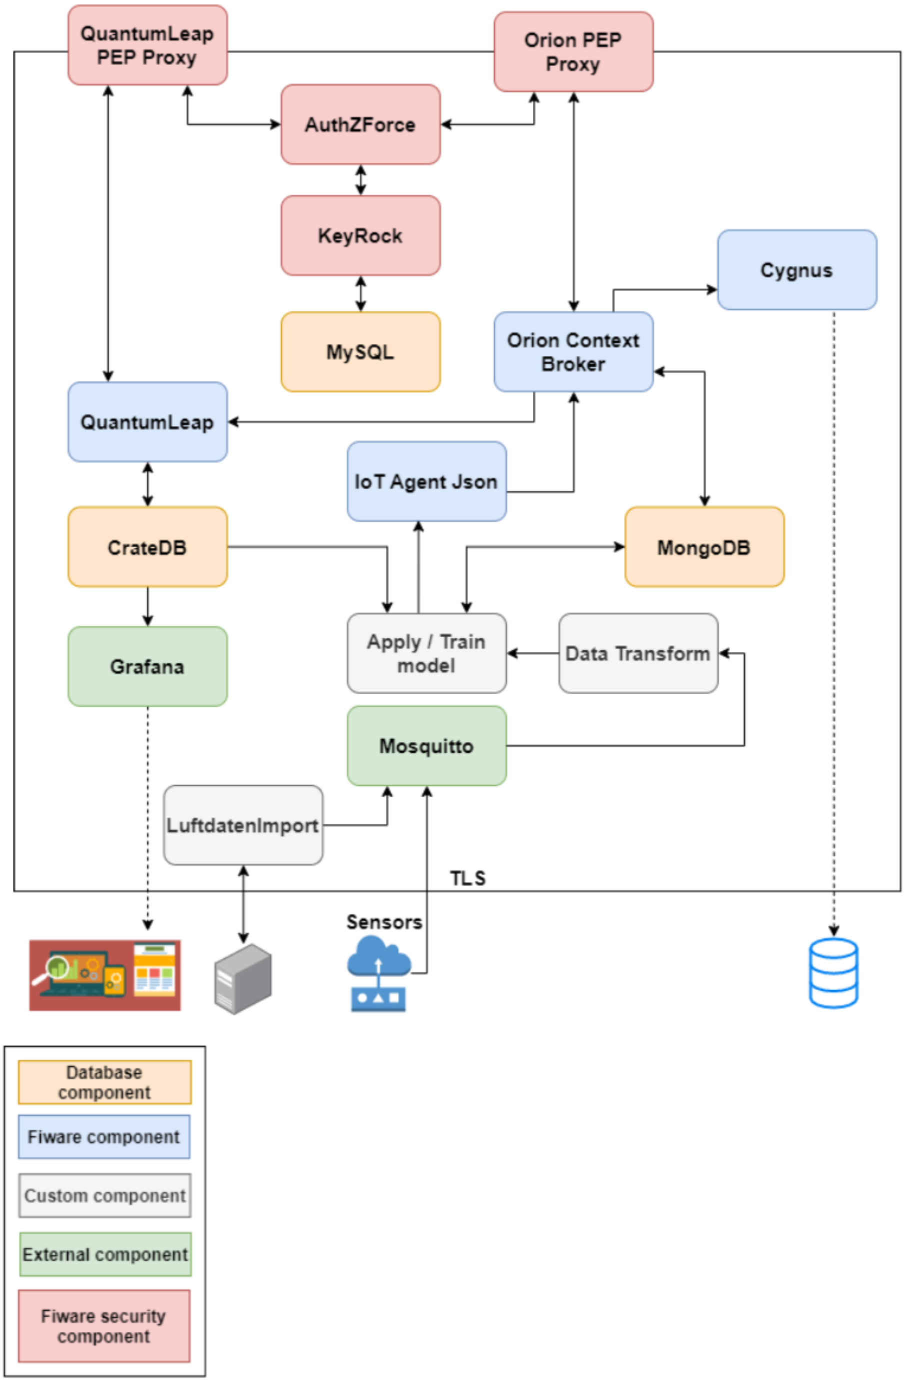

Figure 2. System and component overview. The arrows indicate data flow.

Table 1: Used components and their respective purpose.

Name	Purpose	Supporting backend
Orion Context Broker	Core functionality handling entities, metadata and holds current values.	MongoDB
Iot Agent JSON	Holds IoT device information. Updates the entities in the context broker with data from Mosquitto	MongoDB
QuantumLeap	Stores updates in the context broker to a time series database. Enables access to historical data through API.	CrateDB
Cygnus	Enables integration towards external databases	
AuthZForce	Policy decision point. Handles the policies set in KeyRock	
QuantumLeap PEP Proxy	Serves as a gatekeeper for access to the historical data API	
Orion PEP Proxy	Serves as a gatekeeper for access to the Context Broker API	
Keyrock	User identity management and permission settings	MySQL
Grafana	Visualization tool	CrateDB, MongoDB
Mosquitto	MQTT Broker	
Luftdaten Import	Imports air quality data from luftdaten.info	
AirData Transform	Transform air quality sensor data to match the IoT agent configuration	
WaterData Transform	Transform water quality sensor data to match the IoT agent configuration	
ApplyModel	Described in detail in *Diagnostics* section	MongoDB
TrainModel	Described in detail in *Diagnostics* section	CrateDB, MongoDB
Portainer	Manages docker containers	

3.4 Data input

The project had four different data input streams, of which three were air quality data and one water quality data. As the platform work package did not have responsibility for the data transmittance from the sensors to the platform, this data flow is explained briefly below.

For the air quality sensor boxes, measurement data in json-format was sent over 3G/4G to an AWS-server hosted by Vinnter who were responsible for the sensor box communication. The data was then published on a MQTT topic to which the platform broker subscribed. The AWS server was set up as a temporary solution when we had the sensor boxes ready to measure but no IoT platform to take care of the data, but we then used it as an intermediate backup storage once the platform was set up. As the data format coming from the sensors did not match the platform data models as configured in the IoT Agent, because of nearly a year's difference between the format decisions, a transformative layer was created as seen in figure 2.

The second air quality data stream came from the city's official measurement stations. The city's environmental division used the Fiware NGSI API to communicate the data directly from the measurement system to the Orion Context Broker through the PEP Proxy.

The third air quality data stream was set up using the Luftdaten REST API to collect data every 5 minutes for all the active sensors within a range of 15 km from Gothenburg central station which was around 50 sensors. A docker image was created to fetch the data and subsequently publish it to the platform MQTT broker.

We set up the water quality data streams in a similar way as the air quality sensor boxes. Talkpool were responsible for the sensor deployment and communication and they had a MQTT server to which the data was sent, and the platform MQTT broker then subscribed to that. As for the air quality data, a transformative layer was needed as the sensor data model did not match the data model in the IoT Agent. Because of the nature of the water quality measurements, as they measured only when the water flows were high, and the low number of sensors the number of water quality data points was significantly lower than for air quality.

3.5 Data output and visualization

The data was published in three different ways: two different types of API:s and through a visualization website.

Grafana was used as a visualization tool to make the data available to all project partners. Grafana was considered suitable as it is a simple tool that enables visualizations for the project's basic needs. A publically accessible website was created at http://loviot.miljoforvaltningen.goteborg.se:8003/ (not published openly), and a dashboard including a map based visualization and time series graphs was created. An example of this is shown below. This was created for both users who had an interest in getting a simple visualization and overview of the current air quality status, and data users who needed a more graphical access point to all data other than through API:s. Even though the website was publically accessible the dashboards were however password protected so in practice it was only available to project participants. This was done as the project decided that the data should not be publically accessible before we could evaluate the data quality.

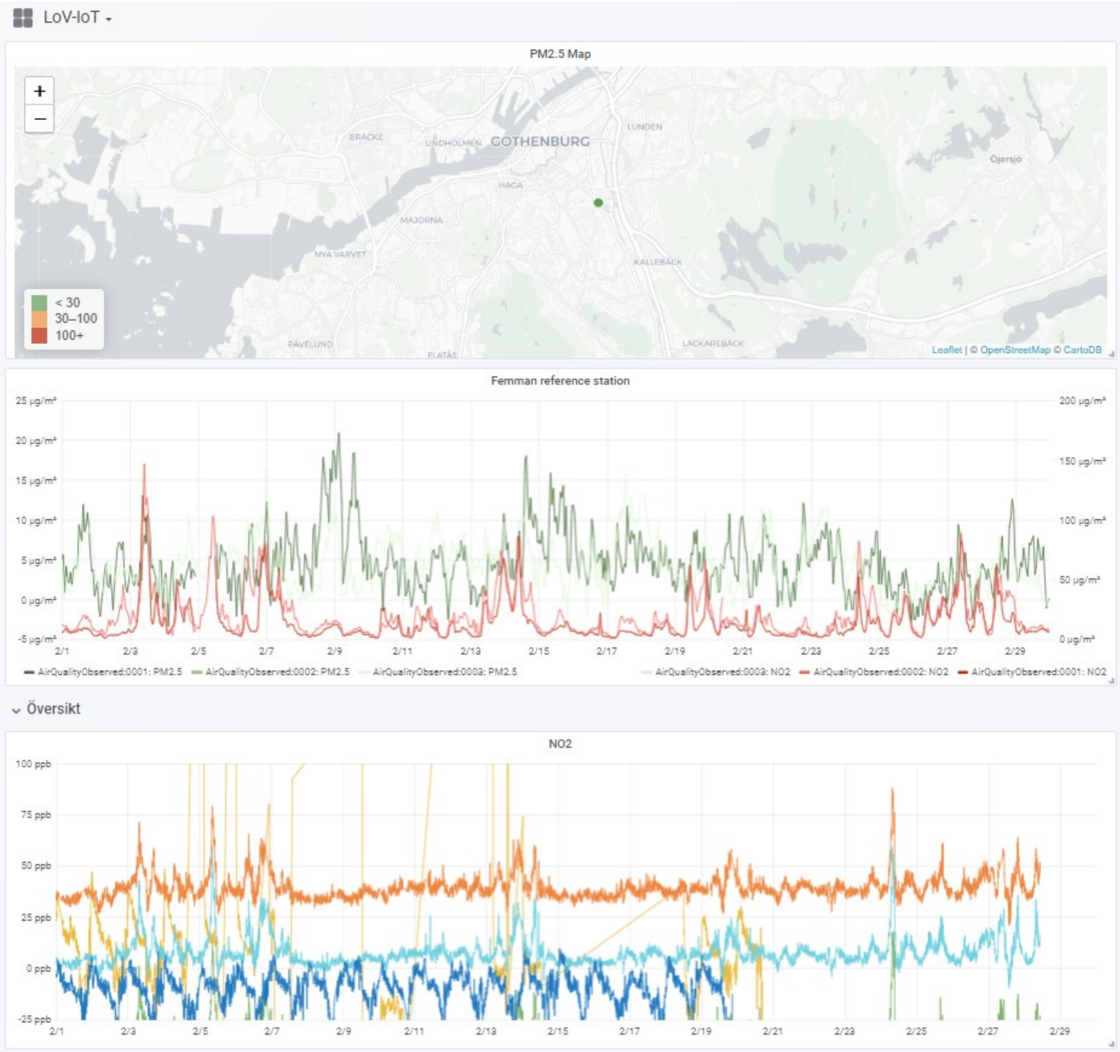

Figure 3. An example of the visualization of data.

Furthermore, access to the data was given through the Context Broker and QuantumLeap open API:s. These API:s enabled more advanced users to access the data programmatically rendering possibilities for further data usage. The Context Broker API gave the user the ability to see lists of entities and access current data, whereas the QuantumLeap API gives access to historical data. Access to these two different API:s was handled by the three security components KeyRock, AuthZForce and Wilma PEP Proxy. Each use case within the project was given a separate access key with specific roles and permissions attached to the key.

3.6 Diagnostics

As one of the goals of the project was to increase the quality and value of the data produced by the sensors to meet the needs of the data users and stake-holders, a diagnostics tool was developed. The tool was a Python based docker image designed to do two things:

1. Create a variable indicating if the sensor does not behave as expected, here done by calculating the correlation with the sensor's closest three neighbors.
2. Train and apply a correction algorithm based on reference measure-ments. Before deployment, most sensors were placed next to a reference instrument to create a sensor specific correction algorithm. This algorithm was then updated weekly if the sensor still was nearby a reference instrument. Based on the collected data a generic model was also created for each sensor type, so that deployment of new sensors could be done without having a sensor specific model.

A schematic overview of the diagnostics is shown in figure 4, split on what was done in real time and what was done on a weekly basis. The time window of one week was set based on discussions with stakeholders.

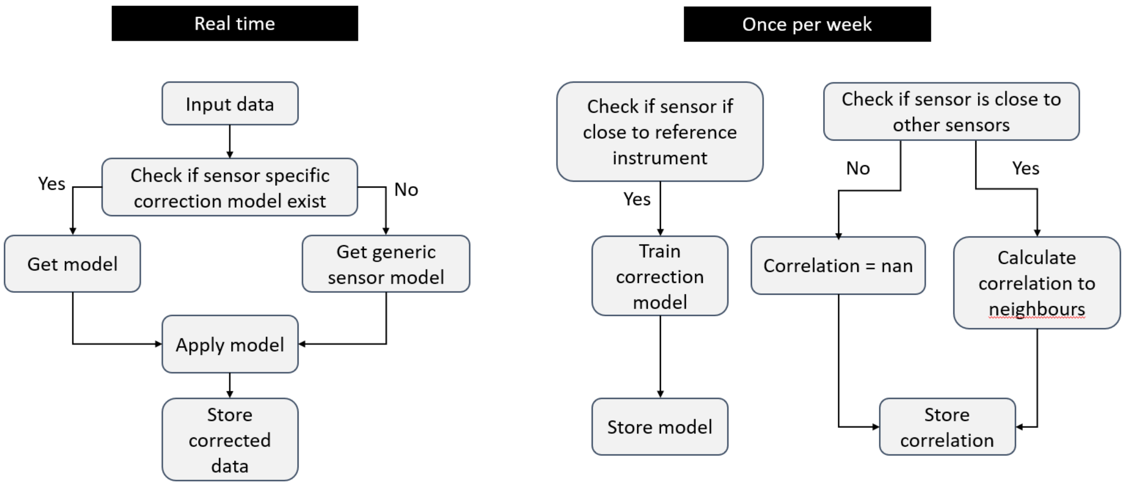

Figure 4. Diagnostics overview.

Furthermore, was a digital testbed developed that automatically rendered statistical evaluations of the different sensors when situated in proximity to a reference station. This enabled a more objective comparison between different sensor's performance and stability over time, which provided both sensor manufacturers and city officials with valuable information. Detailed information on the testbed can be found in the LoV-IoT WP4 final project report.

3.7 Standards and data models

A major part of the discussions in the platform work package have been around the different standards and data models in and around IoT and smart city applications. Our general view was that the project would only use existing open standards and data models, and since the platform was built on the Fiware architecture we used the existing Fiware data models. As previously explained, the Fiware documentation and developed data models in the beginning of the project were not as complete as today but the project settled on using the *AirQualityObserved* and *WaterQualityObserved* data models as a starting point in our discussions of data models. Quickly we realized that these models would need some modifications to fit the needs of the project, mainly focused on two things:

1. the amount of metadata regarding sensor type, sensor quality, sampling time etc. not being handled by the IoT Agent, and
2. the sensors provided many more signals that were also interesting than just a simple pollution attribute.

There were also some technical issues in following the standard nomenclature as for example the time series database would not accept "." as an attribute name, as in *PM2.5*. Thus, what we did was that we stored these sensors as new *types*, not as *AirQualityObserved* entities. However, the city's official reference data was stored and handled according to the *AirQualityObserved* data model and the water quality sensors did also fit the *WaterQualityObserved* data model. Since the NGSI API used to retrieve information from the Orion Context Broker uses entity types as the primary resource endpoint, not all data could be retrieved using the same endpoint. A detailed description of the data models you find in Appendix B.

3.8 Security

The end-to-end security has been part of the project, although split over the different work packages. Each sensor development work package, for air and water sensors, has had responsibility to secure the data transfer from the sensors to their respective storage server. Thus, we have had trust in them delivering untampered data to the platform. The platform work package has dealt with securing the communication to and from the platform, with SSL security for all the inbound data streams and SSL and OAuth2 handling the outgoing data and API access. The security components that have been used are developed by the Fiware community and are considered secure, but the system as a whole has however not been professionally pen-tested.

In this proof of concept environment the security aspect might not be of utmost importance, but as the project has an aim for scalability and replicability these are questions that need to be handled with more time and resources, which has also been identified by IoT Sverige. When it comes to IoT, security needs to be a clear part of the work from the start.

4 Data usage

4.1 Nordic Way 2, Task 8 – Smart routing based on infrastructure policy

The project, where the City of Gothenburg is an active participant, examined the possibility of using real time data of air quality monitoring in order to adapt the drive of hybrid vehicles when the air quality is bad. The project demonstrated the concept that hybrid vehicles entering a zone with bad air quality could automatically change to electric drive, to help improve the air quality. The goal was originally to use the data from real measurements. The Fiware installation from LoV IoT was examined as a source but was not used primarily because the project needed data that changed in order to demonstrate the use cases. Therefore, live data was unsuitable. The project does show that there is demand for this kind of service.

4.2 ElectriCity

ElectriCity equips the electric buses for line number 55 with particle measuring stations for air quality. Ericsson connects the stations to the ElectriCity Innovation Platform in order to examine air quality. One problem encountered was to evaluate how accurate the measurements were compared to the more advanced stations employed by the City of Gothenburg for official measurements. ElectriCity has access to API:s from the LoV-IoT project and evaluated if it was suitable to use our data. However, to the best of our knowledge data from the LoV-IoT project was not used.

4.3 Within the City of Gothenburg

The City of Gothenburg collects its data from their measuring stations to the Air Quality Management system Airviro, Airviro is located on a Linux server. Within the project, we have created a solution for uploading data from the Airviro system. The described data model AirQualityObserved was used. The uploading of data to Fiware is done with a script, using a curl-command. The script runs once an hour to upload near real time data.

At the present, data are uploaded from three stations:

- Haga south (PM10 and PM2.5)
- Haga north (NO2)
- Femman (NO2, O3, PM10 and PM2.5)

This makes it easier to compare data between the low-cost sensors and the reference instruments. As a direct consequence of the LoV-IoT project the city has started publishing information openly from more measuring stations from one earlier to three now.

5 Impact

5.1 City as a Platform

The project City as a Platform (CaaP), a part of “Viable cities” is run by RISE and started 2018 and is a direct result of the needs identified in the first four IoT projects being run as part of the Vinnova (Sweden’s Innovation Agency) program IoT Sverige.

Several of the partners connected to the LoV-IoT project are also involved with CaaP. Through CaaP the LoV-IoT project reached several new partners, among others 16 municipalities, KTH, LTU and SiS.

CaaP focuses on practical matters related to technical platforms, standards and how to organize the work with IoT within municipal organizations. The LoV-IoT Project has contributed with input, expertise and questions to CaaP.

LoV-IoT saw it as a natural part of the project to scale technical aspects to other municipalities, an effort that has been partially successful even though we wished for even more followers. Of the 16 CaaP municipalities we have established closer contact with six of them in order for them to test our solutions.

Both the LoV-IoT project and CaaP have contributed to the fact that the City of Gothenburg has identified the need for an overarching strategy and access to infrastructure within the IoT domain. The city wants to use technical platforms for as many uses as possible while maintaining demands concerning security, architecture and manageability.

5.2 The City of Gothenburg is a frontrunner within Fiware

During 2018 the City of Gothenburg's work with Fiware within the LoV-IoT project and other projects were acknowledged by the Fiware community. In November 2018 City of Gothenburg joined the Fiware Foundation as a front-runner. The Fiware foundation frontrunner is a collaboration program to support the adoption of a reference architecture and compatible common data models that underpin a digital market of interoperable and replicable solutions for smart cities.

Since then the city has entered an intense period of work in order to create and implement an architectural framework for working with IoT. The goal is to establish technical, legal and organizational solutions for handling IoT that can be used by the entire municipality. Parts of Fiware will most likely be part of this solution going forward with building the smart city.

In november 2018 City of Gothenburg joined the Fiware Foundation as a front-runner. The Fiware foundation frontrunner is a collaboration program to support the adoption of a reference architecture and compatible common data models that underpin a digital market of interoperable and replicable solutions for smart cities.

Reference: https://www.fiware.org/news/fiware-foundation-and-tm-forum-launch-front-runner-smart-cities-program/

5.3 The Swedish environmental protection agency

Through collaboration with the government assignment "Smarter environmental information" being run by the EPA the city has brought forth issues concerning principles for sharing information, standardization for API:s and data models from Fiware for reporting air quality measurements between municipalities, the EPA and the EU.

It has come to light that there exists a gap between the demands from legislatures and the Fiware data models. The EPA has deemed that the Fiware data models do not meet the requirements set by the EU. This is an important finding and something to learn from going forward and an issue that needs to be remedied. Therefore, it is important that we continue the work within Sweden and together with the EU in order to bring change and not just state differences.

One challenge is that both government agencies, county administrative boards and municipalities are independent and there is a lack of common ground and direction as well as steering in order to achieve interoperability concerning data management.

5.4 An established platform within the City of Gothenburg

One impact from the LoV-IoT project is that the city, since 2019, has accelerated its work with developing, establishing and implementing a technical platform (infrastructure, security etc), organization (processes) and standardization concerning smart cities and IoT within the city's architectural framework and technical environment. The work is being conducted in close cooperation with national actors such as Swedish municipalities and regions (SALAR), Inera (SALARS it-company), Swedish Standards Institute (SiS), The agency for digitalization (DIGG) as well as other municipalities.

The city has among other things been working on adapting LAN and WLAN, establishing processes, catalogues for master- and metadata, mapped needs and demands as well as creating an organizational wide architectural framework for IoT.

During 2020 work has been started with establishing a City information platform as a foundation for the sustainable smart city.

5.5 The City of Gothenburg’s policy and plan for digitization

During 2019-2020 the city has created a policy for digitalization and IT. The policy describes the guiding principles for the city's work with organizational development connected to digitalization and its basic enabler, IT.

Findings from the LoV-IoT project have clearly contributed to the work and content of the new policy that focuses on issues concerning information management and supply, information sharing, standardization, common infrastructure and frameworks for architecture and information security.

6 Collaboration

The LoV-IoT project has collaborated with several other projects and initiatives within different domains. Through collaboration we have had an active exchange with other actors. The exchange of ideas and findings has encompassed how to work with IoT, project findings and in some cases results and techniques.

IRIS

IRIS (https://irissmartcities.eu) is a Smart Cities-project run by the three lighthouse cities Gothenburg, Utrecht and Nice.

The goal is in part to create a replicable smart cities platform and in part to create a data market platform. The cities goal is to learn from the other cities in order to use the best ideas and findings when developing the cities own IoT-platform.

For the LoV-IoT project this has, among other things, meant that the platform the city uses for publishing open data (Entryscape) has been developed by the supplier in order to handle dataflows from Fiware and that the issues of a common data market has become part of the cities strategies.

SCORE

SCORE (https://northsearegion.eu/score/) is a Smart Cities project focusing on spreading and reusing open data and open source code. Nine cities within the EU have been part of the project.

The collaboration with the LoV-IoT project has focused on sharing findings around open data and standards.

As a spin off we packaged and replicated work being done with citizen science together with the Science fair. This has also been shared with the city of Bradford in the UK that adopted the entire concept.

The science fair

In collaboration with the science fair and the association Luftdata (luftdata.se) we have built and distributed cheap air quality monitoring stations to the public in order to get them to share their data with us. All data from this type of measuring stations has been added to our platform and are available for use by the city through API:s and visualizations.

Red Hat

In May of 2020 it was announced that Red Hat were to become new members of the Fiware foundation. This was in part a result of a dialogue between Red Hat, that is one of the city's procured suppliers, and the city concerning prerequisites for working with and implementing strategies for open IoT-platforms.

(https://www.fiware.org/news/fiware-foundation-announces-red-hat-as-platinum-member).

SCOREWater

SCOREwater will be a start for cities to collaborate with the goal to be water resilient using ICT. Open Source and Open Data will make it easier for cities to share development, data and computational models and we will together explore the possibilities this can have for the different stakeholders as well as for the economy in this new digital ecosystem ScoreWater (https://www.scorewater.eu/).

The LoV-IoT project shares findings connected to organisation, IT and open platforms with the project.

Appendix A: Technical implementation

During the LoV- IoT project, IntraService first set up an environment using VMware vSphere to host the application in their enterprise environment. This was done as a test to evaluate the possibility of creating and maintaining an IoT platform in a production like environment. The vSphere version did however not offer full support to all Docker commands and features, and several issues remained unsolved when the decision was made to instead host the platform on a regular linuxLinux server. The main issues regarded network and version compatibility as communications between the different Docker components did not work as well as between the components and external networks.

When setting up a Fiware based platform as a containerized application using Docker, knowledge of how Docker and Docker Compose work is crucial. Many problems that arose within the LoV-IoT project can be avoided by having a deeper understanding of docker networks, ports and versions.

Appendix B: Data models

The data models used in the platform were based on the type of sensor that was providing the data. Two examples are shown below, showing both type *AirQualityObserved* and *PMS5003* and exemplifying how the original AirQualityObserved data model was extended:

Table 1: AirQualityObserved

Attribute name	Type
dateObserved	string
entityId	string
entityType	string
fiwareServicePath	string
location	text
no2	float
o3	float
pm10	float
pm2_5	float
reliability	float
source	string

Table 2: PMS5003

Attribute name	Type
altitude	float
db10	float
db2_5	float
deviceId	string
entityId	string
entityType	string
fiwareServicePath	string
latitude	float
longitude	float
pm10	float
pm10_atm	float
pm2_5	float
pm2_5_atm	float
sensorId	string
sensorType	string
timestamp	string

Environment Administration, City of Gothenburg
Box 7012, SE-402 31 Gothenburg
Telephone: + 46 (0)31 365 00 00
E-mail: miljoforvaltningen@miljo.goteborg.se