D3.6 Lessons learned and recommendations to the EC for future FIWARE smart applications

Author(s): Marc Ribalta (EUT), Panagiotis Kossieris (NTUA), Alex van der Helm (WNT), Marcel Zandvoort (WNT), Ben Ward (SWW), Stéphane Deveughèle (3S)

Co-Authors: Lluis Echeverria (EUT), Dirk Vries (KWR), Siddharth Seshan (KWR), Josh Pocock (SWW)

31 May 2022

This project has received funding from the European Union's
Horizon 2020 research and innovation programme under
Grant agreement No. 821036.

Disclaimer

Intellectual Property Rights

This document contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

This document is the property of the Fiware4Water consortium members. No copying or distributing, in any form or by any means, is allowed without the prior written agreement of the owner of the property rights. In addition to such written permission, the source must be clearly referenced.

Project Consortium

Executive Summary

This document contains the lessons learned and recommendations from the different technical activities related to the four demonstration cases of the Fiware4Water project (Greece, France, Netherlands, United Kingdom) developed within Work Package 3 (WP3). The four demo cases are related to various aspects and parts of the water cycle: Raw water supply, water distribution networks, wastewater treatment, and citizen engagement and water use.

The document reports the technical conclusions in four actions, as implemented in each demo case: (i) The sensor configuration, explaining the planning, installation setup, sensor network design, and how is the data being stored; (ii) The data analytics and modelling-related activities, explaining the motivation behind the analytics, how the data are used to create the models, and which actions were taken into account when implementing these kind of solutions; (iii) The visualization activities, which are different in each demo case, ranging from a single visualization interface to more complex solutions like a whole phone application. The chapter explains, in each case, the technology being used and how is the information presented to the end-user; (iv) All the smart applications developed within WP3 have been deployed using components from the FIWARE ecosystem, improving their legacy system into a fully operational data management platform. This last chapter explains which components are integrated and the followed data management strategy.

The European Added Value (EAV) of this report comes from different perspectives. Each demo case implemented a set of actions not only in a different part of the water cycle but also in another European country, where the water use differs for different reasons (social behaviour, climate, and more). These actions offer an overall vision in how to implement a set of smart applications for the overall water cycle in different locations. Furthermore, the steps needed for the adaptation of a legacy system to interact with new incoming data, intelligent predictive models, and data visualization interfaces, using the components from the FIWARE ecosystem, are explained in the four cases where the legacy systems have different architectures. Thusa good amount of examples is provided, on how to prepare already existing platforms (legacy systems) for the integration of innovative smart solutions.

Related Deliverables

D3.1, D3.2, D3.3, D3.4, D3.5 – Describing the development of the smart applications for the different demo cases.

D4.1, D4.2, D4.3, D4.4, D4.5 – Describing the deployment of the smart applications and their final use.

Document Information

Programme	H2020 – SC0511-2018
Project Acronym	**Fiware4Water**
Project full name	FIWARE for the Next Generation Internet Services for the WATER sector
Deliverable	**D3.6: Lessons learned and recommendations to the EC for future FIWARE smart applications**
Work Package	**WP3: Smart Applications and Devices**
Lead Beneficiary	Eurecat (EUT)
Author(s)	Marc Ribalta (EUT), Panagiotis Kossieris (NTUA), Alex van der Helm (WNT), Marcel Zandvoort (WNT), Ben Ward (SWW), Stéphane Deveughèle (3S)
Contributor(s)	Lluis Echeverria (EUT), Dirk Vries (KWR), Siddharth Seshan (KWR), Josh Pocock (SWW)
Quality check	Lydia Vamvakeridou-Lyroudia (UNEXE/KWR)
Planned Delivery Date	31/05/2022
Actual Delivery Date	30/05/2022
Dissemination Level	Public

Revision history

Version	Date	Author(s)/Contributor(s)	Notes
Draft	25/05/2022	Marc Ribalta (EUT), Panagiotis Kossieris (NTUA), Alex van der Helm (WNT), Marcel Zandvoort (WNT), Ben Ward (SWW), Stéphane Deveughèle (3S), Lluis Echeverria (EUT), Dirk Vries (KWR), Siddharth Seshan (KWR), Josh Pocock (SWW)	First draft of the document, ready for internal review
Draft	26/05/2022	Lydia Vamvakeridou-Lyroudia (KWR)	Internal review
Final	27/05/2022	Marc Ribalta (EUT)	Addressed the suggestions indicated in the internal review.
FinalV2	19/07/2022	Marc Ribalta (EUT)	Adressed the comments received during the final review of the project.

Table of content

List of figures

List of Acronyms/Glossary

F4W	Fiware4Water project
NGI	Next Generation Internet *The Next Generation Internet (NGI) initiative, launched by the European Commission in the autumn of 2016, aims to shape the future internet as an interoperable platform ecosystem that embodies the values that Europe holds dear: openness, inclusivity, transparency, privacy, cooperation, and protection of data.*
WPL	Work Packages Leaders
WP	Work Package
CB	Context Broker
UI	User Interface
TF	Technical Functionality

Introduction

The digitalisation of the water sector is a fact that has been extending for years. However, the data revolution involves a change in computer systems in order to integrate information in a homogeneous and accessible way. In addition, in the last years, the use of data to extract analytics that improve internal processes and the generation of mathematical or data-driven models to create predictive solutions has become one of the main points of innovation for any company. In WP3 of the Fiware4Water project, the aim was to develop a set of smart applications easy to integrate with the FIWARE4Water architecture, parallelly developed during WP2.

To implement the smart applications, a previous step needs to be done, which is the sensor deployment and configuration to gather enough data to be used afterwards. This procedure is explained in WP4 deliverables, while in this document, the recommendations for the implementation for each demo case are explained.

The generated data needs to be stored, and during the execution of WP4, a set of connectors was developed to improve the existing legacy systems and adapt them to the FIWARE ecosystem. The technical difficulties and future recommendations to adapt legacy systems are described in this document and for each demo case.

Smart applications are mostly divided into two parts, the smart solution and the visual interface. Smart solutions can range from data analytics to predictive models, aiming at using the previously gathered data from the deployed sensors and profit from it, optimizing and automating internal processes, predicting anomalous events, anticipating demand issues, or implementing recommendation systems. This report explains the lessons learned from implementing them, mentioning which changes in the algorithm selection or which type of data is better to use in each case.

The output of the smart solutions needs to be visualized by the final user, and the user profile changes within each demo case, meaning the visual interface needs to be adapted to each user in each case. The user profile, technologies used to develop the interface, and the content shown in each case are explained within the document.

The technical development and deployment of the demo cases is explained in all the past deliverables of WP3 and WP4. The project includes four demo cases in Greece, France, the Netherlands and the UK respectively. The Greek demo case is explained in D3.1 “FIWARE-enabled applications for Raw Water Supply” and D4.1 “FIWARE4_Raw water supply system real-time operational management”. The French demo case is explained in D3.2 “FIWARE-enabled applications for Water Distribution”, D3.5 “FIWARE-enabled Water Quality Sensors”, D4.2 “FIWARE4_Leakage Management” and D4.3 “FIWARE4_Water Quality Monitoring and Pollution Response”. The Netherlands demo case is explained in D3.3 “FIWARE-enabled applications for Waste Water Treatment” and D4.4 “FIWARE4_Intelligent Control for Wastewater Treatment”. Finally, the United Kingdom demo case is explained in D3.2, D3.4 “FIWARE-enabled applications for Customers” and D4.5 “FIWARE4_Smart Metering and Citizen Engagement”.

Therefore, the structure of this document, for each demonstration case:

- Sensor configuration: The decisions are taken to deploy the sensor network.
- Data related activities: The processes implemented and the key points to consider.
- Result visualization: The technology, the final user, and the visualization implemented.
- FIWARE integration: Technical lessons on how to integrate the legacy system with the FIWARE ecosystem.

I. Raw water supply optimisation (Greece)

The Athens Water Supply and Sewerage Company (EYDAP S.A.), the largest company of its kind in Greece, operates the external raw-water supply system that serves the metropolitan area of Athens (Greece; 5.000.000 inhabitants). The system transfers raw water mainly from surface water resources (Marathonas, Yliki, Mornos, Evinos), along with a series of boreholes, to the four water treatment plants of Attica, via an extensive system of aqueducts of a total length of 495 km. In the framework of F4W, we focused on the Mornos aqueduct, and specifically the "Giona – Dafnoula" channel, which conveys the vast percentage of total water supply (e.g., more than 90% of total raw water supplied within 2020) via a gravity conveyance system with a total length of 188 km. In the system, the hydraulic (e.g., flow, water depths) and quality (e.g., turbidity, temperature, conductivity) parameters of raw water are monitored via a network of sensors installed across the channel.

To ensure high reliability of operations, a key target of EYDAP, within F4W, is to upgrade the real-time operational management of the raw-water supply system by integrating data sources from different sensors, installed by different vendors over the years, into a common information system, taking advantage of the data portability and integration functionalities, provided by FIWARE. Taking advantage of the integrated source of data, a new platform, along with a series of smart applications, has been developed to enable the operational staff of EYDAP to monitor the system on a real-time basis, and get advice on the optimal management of the large conveyance system. The new platform provides access to real-time hydraulic and quality measurements, and allows the operators to access, process and analyse historical time series, giving also access to the 3 smart applications that developed within the project. These applications concern: (a) advice provision for optimal sluice-gate (flow regulation structures) operation; (b) early warning for high turbidity events and forecast of the level of turbidity at the downstream part of the systems, and (c) analysis and forecast of the water volumes conveyed by the system on daily basis.

I.1. Sensor configuration and integration

The hydraulic and quality parameters of raw water in the demo case are monitored via a network of sensors. In more detail, there are 18 quality sensors, measuring in real-time, temperature, turbidity and conductivity at 6 key points of the conveyance system, 20 water level meters, 9 water flow meters and 8 sluice gate opening valves. To support the development of smart applications, historical datasets from these sensors were processed and analysed concerning the quality and credibility of measurements. Further to desk studies, valuable information and hidden insights were obtained during the field visits at key points of the conveyance system. These activities proved especially valuable for the development of the application that provides advice on optimal sluice gate openings, and its implementation in an operational context. In this context, the analysis of time series from the existing water level meters showed that the data is characterized by a high percentage of missing values due to connection problems, while the field visits, and in situ comparisons of transmitted data with local measurements obtained manually, revealed that some sensors overestimate or underestimate the water depth in the channel. To improve the accuracy and validity of the model that simulates the opening of sluice gates, five new water level meters were installed at key points of the system. Specifically, 4 water level meters have been installed just upstream and downstream of two regulation structures (close to already existing water level meters), while one meter has been installed at a point between two sluice gates. It is worth mentioning that the data from the latter meter allowed the estimation of a critical hydraulic parameter for the part of the conveyance system studied, i.e., the Manning roughness coefficient, which is of high interest for EYDAP and practical importance for future hydraulic modelling activities. The data from the existing sensors, as well as from the 5 new ones installed in the framework of F4W, has been integrated, with the Fiware-enabled architecture, using FIWARE standardisation protocols. The integrated sensors continuously feed the relevant smart applications with data, providing decision support on a real-time basis. As discussed in Deliverable D4.1, the project delivers a fully functional Fiware-enabled solution to EYDAP, which allows the seamless integration of other sensors existing in the system, as well as services, in a straightforward and interoperable way.

I.2. Web platform

Within the F4W project, a new web platform, along with a new graphical user interface (dashboard), was designed and developed. This platform is based on the Nessie platform (a Web Server and Data Analysis & Archiving Engine), developed in NTUA, for the collection and analysis of high-resolution data from sensors. Nessie has been developed throughout EU-funded projects, such as the iWIDGET (FP7-GA 318272), Dessin (FP7 - GA 619039) and SUBSOL (H2020 – GA 642228) project. Nessie integrates into a unique dashboard real-time data from the existing sensors and gives access to the outcomes of smart applications, in order to provide operational support to the operational staff of EYDAP. Being part of the Fiware-enabled architecture, Nessie has been customised to interact with the Orion-LD CB, exchanging data according to FIWARE smart data models. The dashboard has been developed with the use of open-source analytics and libraries which are easily customisable and extendable to cover future requirements and applications, depending on the needs of the operational staff of EYDAP. It is worth highlighting that the dashboard allows the end-users to add new metering devices, also configuring the relevant static information (name, code, position, coordinates, city, installation time etc.). Furthermore, it allows to modify subscriptions to the Orion-LD CB, configure NGSI-LD data endpoints and create alternative instances of dashboards. These functionalities increase substantially the upscaling potential and interoperable character of the developed platform, decreasing the effort and time required for further extensions and customisations in the future.

I.3. Modelling

In the Greek demo case (Athens-DC1), 3 smart applications have been developed to support the operational staff of EYDAP in decision making. These are: a) the optimal regulation of flow in a system to convey timely specific water volumes, b) the forecast of water volumes that should be delivered on daily basis to serve the needs of customers, and c) the provision of early-warning for unusual (high-turbidity) quality events and forecasts for the level of turbidity at the downstream parts of the system given the quality conditions upstream. To provide operational and of practical interest and use applications, EYDAP and NTUA have worked in close collaboration. This allowed the exchange of complementary knowledge both on the way that such a complex system operates under different conditions, as well as on the modelling approaches that can be adopted, depending on the requirements posed by the operators and the available data sources. During the development of smart applications, large datasets of observations for different key qualitative and quantitative aspects of raw-water were collected, analysed and pre-processed. Much of the experience and insights gained during this process consist valuable prior knowledge that can directly inform the development of similar services across Europe. It is worth to highlight that this knowledge does not concern only the development of accurate models, but also on their implementation to a live operational environment, aiming to bridge science to practice.

The activities around the development of the model for the simulation of flow through and over sluice gates provided valuable lessons and insights on the implementation of physics-based hydraulic approaches, omnipresent in every relevant engineering handbook, in a real-world large and complex case study, such as that of the conveyance system of Athens. Specifically, the analysis of two key hydraulic parameters, i.e., the Manning roughness coefficient and Sluice gate discharge coefficient, suggested that the calibrated values, obtained based on real data, are quite different from the corresponding values proposed in the literature, which have derived from theory or ideal laboratory conditions. These findings lead us to follow a “grey-box” approach that combines hydraulic equations that pay respect to physical laws to model the flow through sluice gates and over spillways, with data-driven techniques for the estimation of the key parameters of the equations, which usually assumed constant and equal to a default value without further investigation.

Valuable insights have been also obtained from the development of the smart application that provides a forecast of turbidity events. The analysis of available time series from 6 measuring stations across the channel showed that turbidity of raw water exhibits a high noise behaviour, with sudden spikes (i.e., individual high-turbidity peaks) that appear for a few time intervals locally, without downstream propagation. Such peaks do not consist of high turbidity events with a permanent behaviour, but erroneous measurements which are attributed to sensor malfunction or bad sensor readings due to external factors (for instance, if a leaf stands upon the sensor or the sensor is very close to the surface of the water in the channel). This poses extra difficulties in the development of a model that forecasts the value of turbidity per se with high accuracy and validity.

At the same time, the findings from the analysis of raw-water turbidity series, along with oscillations observed in the series of other key parameters of the conveyance system (e.g., sluice gate openings), indicate a necessity for the development of a data validation application that will pre-process the raw data (e.g., detection and correction of faulty measurements, imputation of missing values and smoothing of series) before their use as inputs to a model. This is of high importance for future developments and applications since it will decrease substantially the time and effort required to pre-process the series.

I.4. FIWARE integration

The sensors available in the demo case and the new smart applications were integrated according to Fiware-enabled system architecture, around Orion-LD CB. To accomplish the integration, specific connectors were developed for the legacy system of EYDAP and the Nessie platform. The generic FIWARE legacy connector establishes communication between the Orion-LD CB and the Data Warehouse of EYDAP, which hosts the databases of the utility. This ensures the further upscaling of the developed solutions since every data source available in the Data Warehouse can be integrated straightforwardly using the developed connectors, with minimum reprogramming needs. On the other hand, any new smart application, service or model can be integrated with minimum effort, taking advantage of the implemented and configured Orion-LD as well as the Nessie dashboard that has been developed tailored for EYDAP. The new applications can also take advantage of the components installed and configured to store data persistently (retrieval of historical), which are currently operational for EYDAP.

The lessons learned based on the feedback of the Greek demo case about implementing smart applications in the F4W architecture are:

- The Context Broker should allow subscribers to receive past notifications, current notifications, or both. So, an archiving module (like Draco) should subscribe to both types of notifications and a real-time application only to the most recent one. This is crucial for cases where the sensor fails to send measurements but keeps measurements in memory and flushes them when the problem is resolved.

- More rules regarding Context Broker subscriptions are needed. For example, a consumer could choose to receive notifications only if the measured value is within a certain range or receive notifications only during the daily hours, or on weekdays.
- During development we had to browse through various websites to read documentation, find relevant tools, browse through data models, etc, anything regarding the different aspects of the FIWARE integration with and development or configuration of tools. It is thought that it would have been better if there was one portal with documentation, best-case scenarios, worst-case scenarios, helper tools, examples, data models, best practices, FAQs, etc.
- The Context Broker should have an option to notify subscribers if it has not received any measurement after a certain period. For example, the consumer could configure the Context Broker to send a "warning message" when it hasn't received values for a certain device.
- Permanent storage modules (like FIWARE Draco) should offer "UPSERTing" functionality by inserting new values and updating existing values depending on certain fields (e.g. the observed timestamp of the received value).

I.5. Recommendations for future applications

The work conducted in the framework of F4W project provided new insights and evidence on both EYDAP (water utility leading the Athens demo case) and NTUA (the scientific partner involved), also delivering a batch of lessons which are valuable for future developments and applications. The key to these lessons and recommendations can be summarized as:

- The implementation of scientific models should not be conducted in isolation, but in a close collaboration between data scientists/developers and operators of the system, at all stages of design, development and implementation of tools/models so as the latter ones to provide reliable outcomes, which will also support end-users at their real operational needs. Having said this, it is important to keep human factor in the modelling loop, so as to assess model outcomes on the basis of experience of operators and tune the models, accordingly. The close collaboration between developers and operators is closely associated with financial aspects, since it can reduce the amount of redundant effort as well as the expenditures.
- Redundant effort can be also avoided by first investigating the suitability/performance of more parsimony and less time-consuming modelling approaches (e.g., autoregressive approaches for forecast) for a task at hand, before deploying more complex and demanding machine learning approaches (e.g., deep neural networks). Evidently, the selection of the more suitable model can be conducted by exploiting the balance between parsimony and efficiency of alternative approaches.
- The collection, analysis, pre-processing, validation and reconciliation of raw data, coming from field sensors is a time-consuming process that should not be underestimated or neglected. In this vein, the development of data validation layers for data pre-processing (e.g., detection and correction of faulty measurements, imputation of missing values and smoothing of series) is of high importance, especially when this data comes from sensors of large open water-related infrastructures, such as raw-water conveyance systems. This is associated directly with financial aspects, since the existence of such a data validation layer at the utility's side, instead of using ad hoc and custom techniques, would reduce substantially the effort required to develop a model since the data would be ready for operational use a priory.
- The project delivers a series of operational analytics that unfold several benefits for EYDAP associated with financial and environmental aspects. The analysis and forecast of total water supply volumes enhances the operational management of hydrosystem, improves the level of preparedness of operational staff and contributes to the conveyance of proper quantities of

raw-water depending on the day of the year. Complementary to this, the estimation of optimal sluice gate settings depending on the flow of desire contributes to the reduction of redundant movements of sluice gates as well as to the reduction of possible overflows. Furthermore, the warnings for high turbidity events contributes to the improvement of the preparedness level of operators.

II. Water supply system management (France)

The French Demo Case (DC2), of the Fiware4Water (F4W) project, is the drinking water supply system of SICASIL (*Syndicat Mixte des Communes Alimentées par Les Canaux de la Siagne et du Loup*, in English, Mixed Water Union of Municipalities Supplied by the Siagne and Loup Canals). SICASIL has delegated the operation of its facilities to SUEZ under a Public Service Delegation contract.

The French Demo Case is located in the Cannes basin, in the south of France, on the shores of the Mediterranean Sea; it is, therefore, a region subject to high temperatures during the summer period. On the other hand, it is a tourist region whose population increases three-fold during the Cannes Film Festival or the summer holiday period. The combination of high heat and high population density requires fine-tuned and optimised management of water resources and drinking water supply system facilities. The operators of the French Demo Case of the F4W project have been confronted with these situations for several years.

Four business issues (BI01 to BI04) drive the French Demo Case:

- BI01 Forecast water resources availability
- BI02 Forecast water demand
- BI03 Detect water leaks
- BI04 Detect abnormal water quality events

II.1. Sensor configuration

According to the project proposal, the NANOsensor multiparameter probe (previously called PROTEUS), designed and developed by CNRS, was to be evaluated within the F4W project.

At the beginning of the F4W project, 3S decided to install, in parallel, 4 multiparameter probes other than the NANOsensor because no multiparameter probe was installed in the French Demo Case distribution network.

CIRSEE (*Centre International de Recherche Sur l'Eau et l'Environnement*, in English, *International Centre for Research on Water and Environment*), the most important research centre of the SUEZ Group and a third party of 3S for the F4W project, graciously lent four multiparameter probes for the duration of the project: 4 nano::stations (manufacturer: s::can) measuring 7 parameters: temperature, conductivity, UV254, TOC, free chlorine, pH and turbidity.

Two reasons motivated the installation of 4 nano::stations:

1. Monitoring the water quality in part of the French Demo Case distribution network as soon as possible
2. Collecting water quality measurements as soon as possible to build up the most exhaustive data history possible for the development, by 3S and TZW, of scientific models dedicated to the business problem BI04 "Detect abnormal water quality events".

These two reasons were also intended to secure the implementation of the French Demo Case in case of problems with the NANOsensor. On the other hand, the installation of the NANOsensor, in the third year of the project, on the distribution network of the French Demo Case would have allowed, in the best of cases, to have only one year of measurements, or even less in case of delay, which would have questioned the exhaustiveness of the measurements history thus constituted.

In the end, this decision paid off, given the technical problems encountered with the NANOsensor, described in deliverable D3.5 "FIWARE-enabled Water Quality Sensors".

The installation of the 4 nano::stations in the French Demo Case distribution network is detailed in deliverable D4.3 "FIWARE4_Water Quality Monitoring and Pollution Response". This description also includes installation, configuration, and testing of the whole communication chain set up to transmit the acquired measurements: Sensors ⇨ Data logger ⇨ SCADA TOPKAPI ⇨ AQUADVANCED® Water Networks "Distribution" (software product published by SUEZ Smart Solutions, used for the French Demo Case). The measurement analysis is detailed in deliverable D3.2 "FIWARE-enabled applications for Water Distribution" and deliverable D4.3.

Benefits induced by the F4W project: the operator SUEZ and its public client SICASIL expressed their satisfaction with the installation of the 4 nano::stations in the drinking water distribution network and the possible monitoring of the water quality. Independently but because of the F4W project, SICASIL has asked SUEZ to install, before the end of 2022, 10 other multiparameter probes (probably nano::stations) to have a more exhaustive vision of the water quality of the drinking network.

II.2. Data-driven modelling

Many software applications or services are now deployed on operation sites: data historian, monitoring, weather forecasts, diagnostic tools, hydraulic models (more rarely), interventions monitoring, users' complaints, etc. Twenty years ago, a great number of these applications functioned in a silo, i.e. with little or no exchange of information with other applications on the site. Over the past decade, there has been a strong need to exchange information between applications on the same operation site in order to obtain a more comprehensive view of the water supply system. More recently, this need has been extended to applications outside the operation site, mainly driven by service or data providers outside the operation site's IT environment; for example, weather data providers. As the IT environments and communication protocols are very heterogeneous, the technical challenge today is the ability to exchange data in a simple, standardised, reliable, high-performance and cyber-secure manner between the various actors of the water sector, whether they are organisations, individuals or legal entities (local authorities, water utilities, companies, the general public, etc.) or technical objects (IT applications, sensors, actuators, etc.).

Interoperability of systems based on FIWARE technology is a major theme addressed by the five projects constituting the DigitalWater2020 synergy group (DW2020), through the Task Force 1 "FIWARE and ontology": Fiware4Water, aqua3S, DigitalWater.City, NAIADES and SCOREwater. 3S is a partner in the two H2020 projects Fiware4Water and aqua3S.

Seven technical functionalities (TF01 to TF07) have been developed as part of the F4W project and detailed in deliverable D4.2 "FIWARE4_Leakage Management":

TF01: Network Notebook

The users of the AQUADVANCED® Water Networks "Distribution" applications can be local operators at the agency level or central operators at the regional level; they need to exchange information relevant to the management and operation of the network. In particular, they need:

- To quickly add a comment on a hydraulic zone or on a piece of equipment (valve, pump, sensor, etc.)
- A dedicated view to see all historical comments

TF02: Evolution of the AMR display

Currently, AQUADVANCED® Water Networks "Distribution" uses AMR data to display several indicators and curves. These indicators and curves are used by the operators to:

- Analyse the real consumption of their networks
- Compare it with the volume delivered in each DMA (District Metered Area)
- Calculate water losses

In DMAs where customers are not equipped with AMR meters, water losses are calculated using the night flow or the minimum flow. AQUADVANCED® Water Networks "Distribution" can use two methods (one based on AMR data and the other based on night flow data) to calculate the water losses, but before the new release, it only displayed one of the two.

TF03: Improvement of the communication between two systems (AQUADVANCED® Water Networks "Distribution" and the intervention system) using a geocoding service

AQUADVANCED® Water Networks "Distribution" has a function that allows users to send intervention requests in the intervention system. This function is very useful for the operators, as it is a way to communicate between the one who analyses the event and those who have to act on the field to repair the leaks.

AQUADVANCED® Water Networks "Distribution" provides a geolocation of the intervention area, but the intervention system needs a postal address because it does not manage the geolocation. This obliged the user to manually enter the postal address of the intervention corresponding to the geolocation provided by AQUADVANCED® Water Networks "Distribution".

Therefore, TF03 improves the communication between the user of the AQUADVANCED® Water Networks "Distribution" application and the person in charge of the intervention in the field and makes the intervention area reliable by assigning it a postal address provided automatically by a geocoding service.

TF04: Generation of events from acoustic loggers

AQUADVANCED® Water Networks "Distribution" retrieves data from several suppliers of acoustic loggers. It displays the noise data, together with the associated thresholds configured by the operators.

The thresholds help the operator to know if the noise is significant (which could mean that there is a leak in the network). The application associates a colour to the sound logger on the map and on the board to identify alerts. But no events are generated, unlike all the other thresholds configured in the application that generate events.

Some operators would like to be alerted by an event that they could manage in the events-view, while other operators do not want to generate events at all as they prefer to handle them separately.

We, therefore, propose first to allow the operator to generate the events he wants via a specific configuration, then to allow the system to create events according to the thresholds that the acoustic loggers generate. The interest will be to correlate the events between them (future developments planned for 2022) such as an acoustic logger on alert, an increase in the night-time flow and an increase in the volume of loss.

TF05: Online integration

Scientific models have been developed by 3S, EGM, EUT and TZW for the four French business issues (BI01 to BI04). The functional architecture, designed and implemented for the French Demo Case, is a strong and nice illustration of the theme of systems interoperability in two ways:

1. Online bidirectional data exchanges between two IT applications:
   - AQUADVANCED® Water Networks (software product published by 3S).
   - The scientific platform, based on the BentoML open-source component, is in charge of hosting and executing the scientific models, as a computation server.
   - Data exchanges are carried out through the F4W platform acting as a data gateway between both IT applications. IT connectors have been developed to exchange data under the FIWARE technology, based on the NGSI-LD standard.
2. A common functional architecture addressing jointly the French Demo Cases of both Fiware4water and aqua3S projects.

TF06: Big Data models integration

The deployment of the different Machine Learning models (generated by EUT) using the Spark Big Data tool is done using FIWARE components. It consists of showing how fast the predictions of "Forecast water demand" and "Detect water leaks" can be done under heavy data circumstances.

The main features are:

- Orion Context Broker to receive water sensor information
- Cygnus historic context to keep track of changes made in the data entities
- API storing the Big Data models
- Grafana visual interface to visualize and interact with the predicted results, and model's performance

TF07: Workforce tool integration

This functionality shows the interoperability and deployment of a workforce tool (generated by EUT) that optimally assigns and schedules operators' tasks according to maintenance operations to be performed within the water distribution system. It hosts a SAT solver that receives the input from the FIWARE component Orion Context Broker and returns to the user the actions taken out by the operators. The tool is deployed using the same FIWARE environment that is used in TF06, deploying together the Big Data models and the workforce tool within a single platform.

The main features are:

- Integration with FIWARE Orion and Cygnus
- Creation of a workforce tool composed of manager, planner, and solver modules, providing complete self-management of planification queues
- Customizable constraints for created tasks

II.3. Results visualization

TF01 to TF04: they have been developed in the AQUADVANCED® Water Networks "Distribution" software product, in connection with the F4W project; they enrich the functional scope of this product and will be available in the next commercial releases before the end of 2022. The entire user experience (UX) – application browsing, data configuration, data input, and results visualization – is performed through the user interface (UI) of this product. Therefore, for TF01 to TF04, no results are visualized by the user in the graphical interface of the F4W platform.

TF05: it illustrates systems interoperability, based on the following use case:

1. The input measurements required by a scientific model, stored in AQUADVANCED® Water Networks, are sent to the scientific platform, via the F4W platform
2. The F4W platform triggers the execution of the scientific model it hosts
3. The results calculated by the scientific model are sent to AQUADVANCED® Water Networks, via the F4W platform
4. The operator visualizes the calculated results through the graphical interface of AQUADVANCED® Water Networks

Therefore, for TF05, as the F4W platform acts as a data gateway, no results are visualized by the user in the graphical interface of the F4W platform.

TF06: the predictive model's visualization is implemented using two major components, **PostgreSQL** as a relational database for storing the model's information such as performance metrics, training date, features and predictions, alongside **Grafana**, which is an open-source analytics and monitoring solution that works with a huge number of different data sources. With those two components, the visualization is achieved by creating two different dashboards, one for each Machine Learning model, which case contains relevant graphics and information.

TF07: the visualization of the workforce tool's output is currently done via console, where it is possible to check all provided solutions for a given proposed planning.

II.4. FIWARE integration

The integration of components / services based on FIWARE technology concerns TF05 to TF07.

For TF05: the IT connectors developed for data exchanges between AQUADVANCED® Water Networks and the scientific platform, through the F4W platform, did not present any particular technical difficulties. The Stellio context broker, the core of the F4W platform, was enhanced by EGM for the French Demo Case for two reasons:

1. Transmit time series whose start or end dates can be in the past, at the current time or in the future
2. Trigger the execution of a model at a fixed period or conditionally to the update of certain input variables required by the model

TF05 has been deployed in SUEZ Microsoft Azure cloud where 4 virtual machines have been allocated.

For TF06: the deployment of the different Machine Learning models using the Spark Big Data tool is done using FIWARE components: (i) Orion Context Broker to store the current state of entities, (ii) Cygnus responsible for the historic context.

With the development of the Big Data Machine Learning models, many benefits were obtained:

- Fast large-scale predictions and output visualization
- Data consistency and history
- Scalability and integrability

For TF07: the workforce tool was developed using the OptaPlanner as a solver, which is responsible for optimizing plans, schedules and more based on pre-defined constraints.

The developed workforce tool has a focus on being self-manageable, fast, and fully integrated with FIWARE technologies, to keep and store the data entities consistently and to be responsive when new planner requisitions are received through the system. Among its main features some key points can be showcased:

- System interoperability
- Self-management
- Scalability
- Highly customizable

The FIWARE technology is of great importance for the European Union (EU), a major topic being its system interoperability, i.e. the ability for systems to exchange data in a simple, standardised, reliable, efficient and cyber-secure way between the various actors in a given sector.

While the water sector had little experience with FIWARE technology, the Fiware4Water project developed open-source components and services (Generic Enablers) specifically for the water sector, compliant with FIWARE technology, i.e. compliant with the NGSI-LD standard. These components enrich the library of open-source Generic Enablers of the FIWARE framework available to developers and integrators. Beyond these IT developments, the French Demo Case demonstrated the speed and ease of use of these Generic Enablers to build smart services.

Moreover, the work done for the French Demo Case would not have been possible without the close and extensive collaboration between the following Fiware4Water partners from different countries: EGM (France), Eurecat (Spain), TZW (Germany) and 3S (France).

II.5. Recommendations for future applications

The developments within the French Demo Case have provided the different partners with an additional learning experience in smart application building. These lessons learned are of great use for the water community to reduce the cost, effort and time to produce similar applications in the future. Below, these lessons are presented as recommendations for further development in water distribution digital solutions:

- Sensor installation is one of the most important actions to develop smart applications, since having data as soon as possible benefits the overall outcome of a project. The SUEZ operators and their public client SICASIL expressed their satisfaction with the installation of the 4 nano::stations in the drinking water distribution network and the possible monitoring of the water quality. Independently but because of the F4W project, SICASIL has asked SUEZ to install, before the end of 2022, 10 other multiparameter probes (probably nano::stations) in order to have a more exhaustive vision of the water quality of the drinking network. These nano::stations can be installed in any water distribution system, and this project has demonstrated it.
- The IT connectors developed for data exchanges between AQUADVANCED® Water Networks and scientific platform, through the F4W platform, did not present any particular technical difficulties. These demonstrates the potential of FIWARE and its adaptability to legacy systems.
- The development of a scientific model for a given area of interest does not guarantee its transposability to other areas of interest in the same territory, and even less so on a national or even international scale. Adapting a model from one area of interest to several others is usually a long, technical and costly exercise, whose conclusion may underline the technical and economic non-viability of a deployment on several areas of interest because the operation would be too long and/or too costly.

- As time passes, data storage and data processes become bigger and complex. With the development of the Big Data Machine Learning models, general benefits can be obtained, enabling smart applications to scale either vertically and horizontally, easily adapting to more powerful hardware, and produce predictions and outputs in a fast large-scale.

III. Intelligent control for wastewater treatment (Netherlands)

The wastewater treatment plant (WWTP) Amsterdam West demo case (DC#3) is situated in Amsterdam, the Netherlands. WWTP Amsterdam West consists of 7 treatment lanes. Each treatment lane has its own state and is controlled individually. One of the treatment lanes (lane 2) is dedicated as a full-scale research lane. For this demo case, additional sensors are deployed in the research lane. Furthermore, AI smart applications are developed, these are described in Deliverable 3.3. The F4W architecture is integrated into the WNT legacy system with AI smart applications that are tested in practice, this is described in Deliverable 4.4.

III.1. Sensor configuration

Currently, a limited number of sensors are used to control and monitor the wastewater treatment process. In the research lane, several additional sensors were placed at different locations in the process. Some were placed to improve the process control while others were placed to gain more insight into the treatment process and/or to check the effluent quality (compared to the other treatment lanes). Construction of the on-site electrical and data infrastructure which are required for the operation of the new sensors has been completed during this project. At several locations of the research lane, new power and ethernet connection points and control cabinets for the sensors were introduced. Cable trays for data, electrical wiring and an optical fibre data line of several hundreds of meters were installed at the WWTP. The data of all sensors and analysers were made available in the legacy Data Historian process information management system (PIMS). The WNT legacy Data Historian PIMS is integrated with the F4W architecture by making use of the interoperable properties of FIWARE. The F4W architecture runs continuously on a WNT virtual machine (VM) processing the latest data of multiple sensors and actuators while running several F4W AI smart applications that are developed in Deliverable 3.3 using F4W data models developed in Deliverable 2.3. By integrating all sensors (existing and new), actuators and setpoints from the WNT legacy system in F4W in the same way, infrastructure and code can be reused which results in decreased development time and increased upscaling potential.

III.2. Data-driven modelling

WWTPs are used throughout Europe and the world. They can differ in size, type of treatment processes, level of automation, type of sensors used to monitor and control the process, the way they are operated etc. So, although the potential for WWTP AI smart applications is enormous, not all developed WWTP smart applications can be used directly at any WWTP in Europe. For instance, the influent soft sensor per lane, described in Deliverable 3.3 is very specific for the plant layout and sensor setup of WWTP Amsterdam West. However, the AI data validation model described in Deliverable 3.3 is highly generic and can be trained for other WWTPs and other sensors. Also, the AI influent flow prediction using rain meter data can be trained for WWTPs at other locations. The control model is applied at the research lane, one of the seven treatment lanes of WWTP Amsterdam West, and can also be implemented in the other six treatment lanes. The applied approach works, so the AI model architectures can be used for other (WNT) WWTPs with some modifications (e.g. including or excluding other sensors) and model training.

The F4W Amsterdam demo case demonstrates the possibilities for the use of FIWARE in the (waste)water sector. The interoperable properties of the FIWARE-enabled F4W architecture together with the developed F4W data models, enable replication and upscaling. The methodologies, approaches, and developed technologies in the Amsterdam demo case present a successful baseline to guide other water utilities for future digitalization processes. By closely working together in developing the AI models, integrating the AI models and deploying the AI models at WWTP Amsterdam West; WNT, KWR and EUT collaborated extensively on a transnational level. The F4W project provided the opportunity to bring together complementary knowledge on wastewater treatment, sensors, AI and model deployment from a research and practice perspective. This collaboration made the development of the smart applications, the integration of the F4W architecture in the legacy system and the use of the smart applications for WWTPs efficient and effective. The gained insights on AI in the water sector can be used across Europe.

The European added value (EAV) encompasses not only boosting excellence in AI for the water sector but also includes insights which were gained regarding nitrous oxide emissions (N2O) reduction in practice. The F4W project has contributed to process knowledge and has boosted research about the production and reduction of N2O emissions from WWTPs. N2O emissions data sets are growing in size, enriched, and are being stored for future use. Consequently, this demo case contributes to the sustainability and greenhouse gas emissions reduction ambitions of the EU. The demo case Amsterdam West WWTP also contributes directly to the acceleration of the EU twin – green and digital – transition, which is regarded as a necessity to reach the climate goals by 2030.

III.3. Result visualization

Result visualization is realised through the user interfaces (UIs) on which the users of applications can see results and/or interact with the applications. In the F4W Amsterdam demo case, the starting point is to use existing UI systems designed for operating personnel and wastewater technologists as much as possible. The reasoning behind this is that the identity and access management (IAM) is then secured via existing IAM processes, hence, no additional administrative processes are necessary for IAM. Furthermore, using, where possible, existing UIs and infrastructure will lead to more robust and secure implementation and fast adoption.

Therefore, the use of the control model is integrated in the distributed control system (DCS), used by the plant operators to control the WWTP. The results of other AI smart applications are brought into the legacy systems and can be visualized e.g. by wastewater technologists, using the legacy visualisation tools they are used to working with.

Next to using existing UIs, the AI smart applications also make use of advanced visualisation via dashboards developed with the use of the open-source analytics and interactive visualization web application Grafana. Grafana provides a dashboard with charts, graphs, and alerts when connected to supported data sources. In the Amsterdam demo case, Grafana is implemented in the F4W architecture and connected to the PostgreSQL database in F4W, which holds all sensor data, and data from the smart AI components. Building visualizations in Grafana takes minimal time to learn, the user interface is simple and functional, it is well designed and offers a good user experience and interaction, i.e. it works in an intuitive way.

III.4. FIWARE integration

Next to F4W AI smart applications, also specific dedicated connectors for AI smart applications with the WNT legacy system and a generic FIWARE smart legacy connector is developed for the WNT legacy system. The generic FIWARE smart legacy connector is able to ingest real-time data of all sensors and actuators from WWTP Amsterdam West and present the data for direct use in F4W smart applications. With the FIWARE smart legacy connector the full upscaling potential of F4W smart applications is ensured. The F4W architecture is deployed and integrated in the legacy system with (near) real-time (ranging from no delay up to maximum of 2 minutes delay) AI smart applications developed in Deliverable 3.3, using F4W data models developed in Deliverable 2.3.

The configuration options and choice of components in FIWARE are ample. This gives a large degree of freedom in setting up a F4W architecture for specific AI smart applications and uses cases fitting for integration with different legacy systems. However, the time needed to get acquainted and to configure components is substantial, due to the complexity involved and the concise, but limited documentation. This has led to a longer than expected lead time for the implementation and integration of the F4W architecture in the legacy systems.

The lessons learned based on the feedback of the Amsterdam demo case with regard to implementing of AI smart applications in the F4W architecture are:

- extending FIWARE to insert missing data through Orion and Cygnus into PostgreSQL (knowing that there is missing data is something different than not knowing if there is data at all);
- adding subscriptions with conditions that are comprised of multiple sensors at a timestamp / time-window simultaneously;
- on a subscription notification, to add the ability to add historical values of a sensor. To ingest historical data which is sometimes needed with a new data point for AI-driven applications which utilise e.g. GRU and LSTM models;
- for new users of FIWARE, time for deployment can be decreased by providing tutorials on adding components to an existing FIWARE architecture;
- FIWARE features are needed to deal with time-critical responses, for example, timeouts, ordering of events / subscription notifications, etc.;
- reducing lead time on opened pull requests;
- allowing components like Cygnus with a structured database to handle updates if the entity / device changes in the context broker;
- with regard to error handling, current subscriptions are based on the concept of fire-and-forget, while it would be nice if there are also error handling features.

III.5. Recommendations for future applications

Research towards mitigating N_2O emissions from wastewater treatment plants has significantly increased over the last years. However, knowledge on how to structurally minimize these emissions in practice is still at its infancy. The current control model implemented in the Amsterdam demo case is a first version and attempt to control the N_2O emission by AI in a full-scale treatment plant application. The control model should now be evaluated and improved in iterative cycles. Following this the control model can be extended and improved by adding more control parameters. Furthermore, influent parameters that determine the incoming waste load should be added to the digital twin as these are expected to further improve the AI control model. If the model-based control proves to be successful in optimizing N_2O emission and energy use it will contribute to the global effort to reduce greenhouse gas emissions.

To extrapolate this approach to other wastewater treatment plants a set of key parameters that enable AI control that reduces N2O and energy use should be deduced, this will make the control more reliable and more likely to be applied elsewhere. Finally, other approaches such as hybrid modelling may make the AI control models better and might require less training time.

AI techniques will inevitably become an integrated part of the water sector. It is therefore of vital importance for the water sector that water utilities start to learn about AI solutions and implement them in their work processes. Water utilities need to get hands on experience and knowledge and all levels of the workforce need to be prepared for adoption and implementation. Learning by deploying and making AI part of the daily processes will allow the water utilities to get experience and confidence in AI solutions and to be able to evaluate market solutions.

General lessons learned for the demo case project are:

- installing and maintenance of sensors with high accuracy is necessary, time consuming and should not be underestimated, the data is the basis of the AI smart applications;
- developing AI smart applications leads to increasing data quality and gives lots of new insight in the processes;
- close interaction between data scientists and wastewater technologists is essential for effective data-driven models;
- data science teams should consist of data scientists and machine learning (ML) engineers;
- close interaction between ML engineers, IT and process automation experts is key in implementing AI smart applications in practice;
- in the current labour market, it is a challenge for water utilities to ensure a stable data science team for the duration of H2020 project time spans of 3 years;
- the water sector is just at the start to explore the power of implementing AI.

IV. Smart metering and citizen engagement (United Kingdom)

South West Water provides water and sewerage services to c.1.7million people in Devon, Cornwall and the Isle of Scilly. Water only services are also provided in the regions of Bournemouth and Bristol, as depicted in Figure 1.

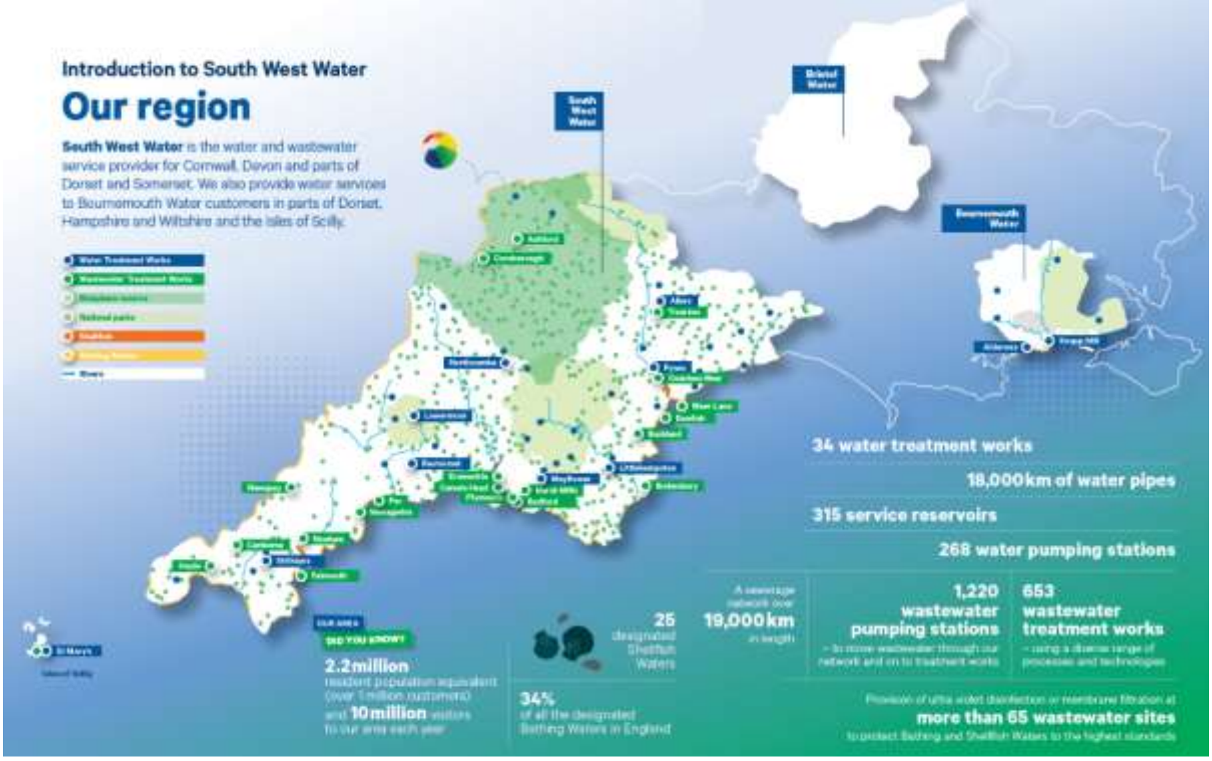

Figure 1 South West Water actions and information

The goal of this demo case (DC#4) was to implement a FIWARE enabled pipeline to retrieve consumption data from smart meters and provide analytics to customers and the utility itself via:

A. A customer smartphone application to drive positive changes in water use behaviour, reduce consumption, and reduce the customers' water bill
B. A utility web-based application to help identify customer side leakage by viewing and interrogating consumption/leakage data received from the water meters

IV.1. Sensor configuration

Both the communication infrastructure and the physical installation of smart meters need careful consideration and planning. In this case study, the communication infrastructure selected was SigFox. This was due to the network's compatibility with South West Water's smart meters and due to the long-range/coverage achievable via such technology.

Before the Fiware4Water project, that was zero wireless connectivity that was suitable. For reasons of power consumption and inability to penetrate underground metal lids, WiFi and 4G were immediately ruled out. Instead, we chose to build a Low-Power, Wide Area Network (LPWAN) from scratch, in-order to connect Great Torrington to FIWARE. SIGFOX was selected over LoRaWAN and NB-IoT for its superior battery lifespan and transmission prowess through metal lids.

The first stage was to survey the GT area for suitable locations to site the SIGFOX equipment. Radio planning is described as 80% science-20% art. Based on elevation, urban clustering and coverage models, we selected two sites to provide the coverage across the pilot area of c.6,000 population in Great Torrington. The first, and primary mast, is located at one of South West Waters nearby service reservoirs. It is a powered mast located approximately 1.6km from the population centre of Great

Torrington. A second (repeater) radio mast was also installed at the local council's main car park. Between the two masts, the area of Great Torrington is provided with full SigFox coverage.

Lessons learned and recommendations:

- Ideally, the mast should be positioned at an elevated vantage point across the desired coverage area. We have learnt that water service reservoirs are often ideal sites for positioning the primary communication mast. By their nature, the service reservoir is usually positioned within reasonable proximity of the population that they serve and are typically located on hillside such that their associated supply is by gravity. Service reservoirs normally offer suitable power supplies, access and space for the mast installations.
- Where necessary, repeater stations were better suited to being installed within closer physical proximity to the properties being served but often at the expense of an aerial advantage. Sites with these characteristics are typically not owned by the utility provider and therefore establishing partnerships with local councils (who are more likely to have infrastructure of this nature) is recommended. In this pilot study, we used the local councils public toilets located at the main town centre car park as the location for the repeater mast installation.

Each smart meter, when installed on a non-metered supply, requires the construction of a meter chamber typically in the public footpath at the boundary of the property.

Depending on the properties configuration it was not always possible to install an external water meter and chamber. Primarily, there were three types of meter installation:

1. **External meter (71%) –** traditional meter chamber installation completed externally of the property
2. **Meter insitu (5%) –** traditional meter chamber already present, requiring the water meter to be fitted within the existing chamber
3. **Internal meter (11%) –** insufficient space or practicalities to install an external chamber. The smart meter was therefore installed within the property, i.e., beneath the kitchen sink.
0. **Infeasible (12%) –** insufficient space or practicalities to install either an internal or external water meter.

Further details on the customer engagement process undertaken before meter installation are covered in D4.5.

IV.2. Data analytics

As described in D4.5, five primary pieces of information are utilised in both the consumer and utility application:

1. the date/time of the reading (observedAt)
2. the unique identifier for the reading (observedBy)
3. the consumption reading which is a cumulative total, i.e., the total volume of water passed recorded through the meter (waterConsumption);
4. alarm/alert and their associated statuses, mainly the continuous flow alarm for leak detection (alarmFlowPersistence);
5. min and maximum flows related to the alarms (minFlow) and (maxFlow).

The utility application is primarily concerned with the 'alarmFlowPersistence' status. This status is set to 1, where the meter has not recorded a period of zero flow, i.e., the flow in the property has a background level of continuous flow over the 24 hour period. This continuous flow status provides the best estimate that leakage is occuring and its magnitude can then be estimated from the 'minFlow' data point. This data point records the lowest non zero flow volume during the 24 hour period which is likely to background level of leakage.

The smart meter data is collected on a rolling daily basis (with several meters being processed every hour, rather than processing all meters at a given point in the day) and stored in Stellio using the FIWARE smart water WaterConsumption entity. Whilst this was not a huge issue for the system, as consumption is measured as a total since meter installation, rather than absolute daily consumption, it did make visualization and data processing over given periods awkward due to the potential for arbitrary gaps in data.

To address this, functionality was developed to automate the production of 'infill' data that would create estimated consumption values between known start and end points. This approach was particularly useful to visualise historical consumption data for the customer app. The infilling of daily consumption data was identified as a key requirement by the customers testing the application. They acknowledged that as although it may not be an accurate daily consumption reading, it served to provide a common language/comparator.

The key lesson learnt from working with smart meter data is to expect to have data loss and, therefore, to build data management about dealing with gaps in data.

IV.3. Result visualization

The utility application is not built on hard and fast rules because the user is an informed user of the information and therefore, we have provided the flexibility for them to customise the rules that are applied. To allow this to happen, the application provides the user with easy to use controls which are set up as filters. For example, the user can select the number of consecutive days that the 'alarmFlowPersistence' records a continuous flow before a leak alarm is triggered. Similarly, the user can filter by 'minFlow' values to filter out low volume leaks.

When dealing with this type of data across a large number of customers, customisation is an important necessity to allows the user to filter down and focus on a smaller sub-set of customers, i.e., where leaks are not being repaired by the customer (long duration) and/or leaks of high flow rates. The trigger levels are likely to be unique to each organisation using the application and will be dependent on the company resources and investment programmes dedicated to supporting the resolution of customer side leakage.

For the customer application, which is aimed at less informed users, the controls are removed. Instead, the information about whether or not the leak alarm is active is presented to the user alongside the estimated leakage flow volumes. The customer application does include feedback/interactivity which is limited to the visualisation of consumption data and specifically how this compares to other users. Predefined average consumption values were used based on occupancy rates, however, there is an opportunity for these to be based on observed data from the smart meters once a sufficient sample is developed. As a recommendation, it will be important that upper/lower limits are set to prevent outliers from influencing the comparative consumption data whilst a sufficiently sized dataset is established.

IV.4. FIWARE integration

The system architecture is based around the Stellio context broker, FIWARE Generic Enabler implementing the NGSI-LD specification. The context broker is mainly in charge of managing the information context, storing the historical data of the water smart meters and providing subscription / notification mechanisms to allow consumer applications to be notified of events of interest.

On the sensors side, the smart meters send their data to a Sigfox back-end managed by a third-party. The Sigfox back-end then transmits every new measure to the Sigfox IoT Agent deployed inside the FIWARE architecture. The Sigfox IoT Agent checks the authenticity of the sender, decodes the data, and then sends it to the context broker, using the NGSI-LD API exposed by Stellio.

To ensure the security of the communications, all exchanges from or to components external to the FIWARE platform are made through a secure HTTPS channel. Also, an authentication and authorization server (namely Keycloak) has been integrated into the architecture. It is used to authenticate exchanges from the Sigfox IoT Agent to the context broker and is ready to be used soon to authenticate calls from the user application. Finally, it is also ready, in synergy with the context broker, to manage necessary access rights to the information stored in the context broker.

Complementary to this, two utility scripts have been developed to ease the setup and enhance the usability of the platform: the first one is used to provision the water-smart meters installed by SWW, and the second one is used to inject past historical data recorded for the sensors (i.e. data recorded before the deployment of the platform).

Security was a particularly concerning issue, given the sensitivity of data. Although the Stellio service was provided with OAuth authentication, there was no granularity to security control, i.e. once a user was authenticated, they had complete control over the full data lifecycle. To mitigate this, a Django server was installed to serve smart meter application requests and pass them to Stellio. This removed update / delete type requests from the customer-side application and gave appropriate levels of control to the app.

Stellio performance tended to be dependent on HTTP requests, particularly in functionality that required several smart meter records to be read, e.g., determining periodic consumption for the entire population, where naïve solutions resulted in long delays in presenting daily or historic consumption data, given the requirement to present a user’s consumption in the context of the entire population. Work was undertaken to reduce this impact by threading multiple request calls, rather than running them purely sequentially and this did help, though the ultimate solution was to process users daily and store results on a separate broker which could be queried on-demand with little cost.

IV.5. Recommendations for future applications

FIWARE has been an incredibly useful technology for this project, providing a straightforward and generally easy-to-use environment for managing data, in particular time series data, without the need to resort to SQL/no SQL databases. However, the FIWARE paradigm can be difficult to follow. It is recommended that developers that are new to FIWARE should look to develop their familiarity with FIWARE through small learning activities prior to engaging in project-based activities. The Stellio ‘beehive’ API walkthrough (https://stellio.readthedocs.io/en/latest/API_walkthrough.html) is a good starting point and the CIM NSGI-LD specifications (https://www.etsi.org/deliver/etsi_gs/CIM/001_099/009/01.05.01_60/gs_CIM009v010501p.pdf), though it is worth noting that not all context broker implementations follow the specification exactly.

There was an initial assumption within the project that data collection would be completely reliable, though this was soon abandoned given gaps in data collection, typically caused by vehicles parking over ground-based meters blocking data transfer. For this project, it was not a huge issue as the key data, ‘total consumption’, would eventually be transmitted to the FIWARE broker and the nature of the data was such that missing consumption data could be ‘infilled’ along a straight-line estimate. However, the same could not be said for the continuous flow attributes which would limit the performance of leak detection.

The resulting recommendation is that consideration needs to be given to the data being collected, the impact of data loss and the likelihood and approaches, if appropriate, for data reconstruction and interpolation.

Much of the initial development in this project was geared around applications interacting with the FIWARE broker. Whilst this worked well for early development activities, it did become increasingly difficult to work in testing, particularly with edge case generation and detection as there was a strong desire to keep testing data away from the live data in the broker.

It is recommended that developers working with FIWARE should look to implement a ‘devops’ style environment, such that functionality can be developed, tested, and operated in separate, but functionally and structurally equivalent environments. We found that moving to a containerised environment with Docker greatly reduced installation-style issues, i.e. going from development to operational environments, and removed a lot of ‘operations only’ application failures. It is also recommended for developers that are looking at containerised environments should look to develop their Docker expertise before moving onto a live development project and avoid using a current project as a learning environment.

Also, developing in a purely live environment raises clear issues about data security, in that if errors are made to the live FIWARE broker, the result could be the loss of live data. This creates a broad recommendation that consideration should be given to data security, to ensure that live data is not lost. In addition, a data policy should be developed to determine how data can be duplicated from a live server into test environments and how data can be backed up, achieved, and recovered as necessary.

Conclusions and perspectives

Each chapter contains a conclusion and reflexion on the work done during WP3 and WP4 for each technological advancement, and the respective deliverables from WP3 and WP4 contain an extensive conclusion and a focused explanation on each. Therefore, technical conclusions are explained in these other chapters/deliverables.

The Fiware4Water project WP3 and WP4 have resulted into great demonstrator cases in how to build smart applications and improve legacy systems to enable powerful data-driven solutions. These demonstrators show solutions around the whole water cycle that can be interpreted, understood and reimplemented by other organizations in any other location. This document offers a set of reflections, experiences, lessons and recommendations, which are useful for any organization that intends to improve their digital systems, deploy smart devices, integrate legacy systems into the FIWARE ecosystem, and develop ready-to-deploy smart applications.

EU added value and upscaling

The water use cycle is a critical process present in every country in Europe and the rest of the world. It involves multiple processes, with complex human and material resource management, and risk prevention that needs to be prioritised in order to reduce unexpected additional economic costs and environmental problems. Therefore, support systems provide additional information and results that can be considered during the management of operations.

The developments of the project involve the implementation of several solutions in four demonstration cases that go through the whole water cycle: i) Raw water supply in Athens, Greece, ii) Urban water distribution in Cannes, France, iii) Water use management, Great Torrington, United Kingdom, iv) Wastewater treatment, Amsterdam, Netherlands. This document offers several lessons learned and recommendations to implement smart applications in different parts of the water cycle, with the added value of showcasing these implementations in different regions of Europe, where there are different climate, social and economic conditions.

Water utilities throughout Europe may differ in size, type of treatment processes, level of automation, operations, sensors, etc. So, although the potential of AI-powered smart applications is big, the produced predictive models cannot be directly duplicated. However, the data modelling, mathematical formulas, and train-test-validate methodology is of great value for any water utility and can be reused to directly produce their own AI models. Furthermore, some utilities may have really similar conditions to the ones demonstrated within the project, and the models can be almost duplicated with very few changes.

Another important aspect of the project is the adaptation of legacy systems to integrate smart solutions and deploy innovative platforms that are reliable and scalable to new and future needs. The four demonstration cases show how the FIWARE ecosystem can be used to improve legacy systems and easily integrate new smart solutions. Each case had a different legacy system, with complex architectures that needed to address new issues such as multiple sensor data gathering, intensive data handling, predictive models, and result visualization. The recommendations of this report are based on four successful adaptations of legacy systems, and the issues addressed are vital for any organization.

Finally, the development of each demonstration case involves the close and extensive collaboration of the partners. The partners are split into architecture development, mathematical/data-driven modelling experts, software development, research teams, and the water utilities working in different parts of the water cycle. This union provides advanced knowledge and expertise in critical water and digital solutions, obtained through the collaboration of partners located in different regions of Europe.

This document will assist in bridging science to practice and science to policy across Europe.