# FI.ICT-2011.1.8FINESCE

## D7.8 Version 1.0

## *Consolidated Trial Results*

**Contractual Date of Delivery to the CEC:** *September 30, 2015*
**Actual Date of Delivery to the CEC:** 16.09.2015

**Author(s):** Padraic McKeever
**Participant(s):** *RWTH Aachen*

**Workpackage:** *WP7– FINESCE Project Management*
**Estimated person months:** 1.0
**Security:** **PU**
**Nature:** **R** = Report
**Version:** 1.0
**Total number of pages:** **56**

**Abstract:**

This report is a consolidated final report on the results of the FINESCE trials.

**Keyword list:**
Generic Enabler, FIWARE, Smart Energy, API

**Disclaimer:**
All information provided reflects the current status of the trial site testbeds at the time of writing and may be subject to change.

# Executive Summary

This is a report on the results of the FINESCE trials.

FINESCE is the Smart Energy project in Phase 2 of the FI-PPP. FINESCE has performed field trials of the use of FIWARE Generic Enablers (GEs) in the smart energy sector and provides Domain Specific Enablers (DSEs) and an Application Programming Interface (API) which allow third-party clients (such as SMEs involved in the FIWARE Accelerator Programme) to develop applications which access the trial infrastructures, thus supporting the creation of a network of Smart Energy application developers.

The FINESCE trials comprise seven trial sites developed by five vertical work packages (WP) covering Smart Energy areas where Information and Communications Technology (ICT) can beneficially be applied.

This report gives an overview of the results of the FINESCE trials. The aspects covered are:

- Results related to FIWARE and GEs are summarised[1]. FINESCE has performed extensive work with FIWARE and GEs, following on from its precursor project FINSENY, which ran from 2011-2013, so that many FINESCE partners have effectively been working with FIWARE and GEs over a four-year period now. This report presents the lessons learned with FIWARE and the GEs with the objective of providing the Phase III participants with useful information and first-hand experience and to feedback the issues and problems to FIWARE.
- The results of the Smart Energy use case experiments performed at the trial sites are summarised.
- The results of simulation activities performed in FINESCE are summarised.
- Results which are published and available for use by third parties (FINESCE DSEs, FINESCE API).
- Results on the further development of the FINESCE trials based on an open platform approach.

[1] More details of results on FIWARE and GEs are available in the six "*Trial Results*" reports of the individual WPs (with separate reports for WP2 Horsens and WP2 Madrid)', and in D7.5 "*Consolidated Analysis of Generic Enablers and Domain Specific Enablers Integration*", available on http://www.finesce.eu/Results.html

# Authors


| Partner | Name | e-mail |
|---|---|---|
| **WP1** | | |
| E.ON | David Lillienberg | david.lillienberg@eon.se |
| | Lars Norrman | lars.norrman@eon.se |
| XLAB | Matej Artač | matej.artac@xlab.si |
| RWTH | Michael Diekerhof | MDiekerhof@eonerc.rwth-aachen.de |
| RWTH | Hassan Harb | hharb@eonerc.rwth-aachen.de |
| **WP2** | | |
| ACCIONA | José Luis Burón | joseluis.buron.martinez@acciona.com |
| ACCIONA | Ismael Rodriguez | ismael.rodriguez.cabado@acciona.com |
| ACCIONA | Lidia Godoy | lidia.godoy.parra@acciona.com |
| Insero Software | Thomas Hune | tsh@inserosoftware.dk |
| Insero | Andy Drysdale | adr@insero.dk |
| Insero | Christina Ernst Jørgensen | chjo@insero.dk |
| Insero | Munna Hoffmann | mhof@insero.dk |
| RWTH | Michael Diekerhof | MDiekerhof@eonerc.rwth-aachen.de |
| **WP3** | | |
| FIR | Julian Krenge | Julian.Krenge@fir.rwth-aachen.de |
| | Marco Roscher | Marco.Roscher@fir.rwth-aachen.de |
| RWTH | Adam Malik | A.Malik@wzl.rwth-aachen.de |
| | Marija Stevic | MStevic@eonerc.rwth-aachen.de |
| | Nicolas Berr | NBerr@eonerc.rwth-aachen.de |
| QSC | Ulrich Hacker | ulrich.hacker@qsc.de |
| | Mirja Steels | Mirja.Steels@qsc.de |
| | Adelheid Weinert | adelheid.weinert@qsc.de |
| XLAB | Matej Artač | matej.artac@xlab.si |
| | Marko Kuder | marko.kuder@xlab.si |
| SOPTIM | Heiner Halbach | heiner.halbach@soptim.de |
| | Frank Fiedler | frank.fiedler@soptim.de |
| **WP4** | | |
| Engineering Ingegneria Informatica S.p.A. | | |
| | Pasquale Andriani | pasquale.andriani@eng.it |
| | Luigi Briguglio | luigi.briguglio@eng.it |
| | Leandro Lombardo | leandro.lombardo@eng.it |
| | Massimiliano Nigrelli | massimiliano.nigrelli@eng.it |
| | Dario Pellegrino | dario.pellegrino@eng.it |
| Synelixis Solutions Ltd. | ArtemisVoulkidis | voulkidis@synelixis.com |
| | Theodore Zahariadis | zahariad@synelixis.com |
| ISMB Istituto Superiore Mario Boella | | |
| | Mikhail Simonov | simonov@ismb.it |
| | Federico Rizzo | rizzo@ismb.it |
| TeamWare Srl | Gianluca Zanetto | gianluca.zanetto@teamware.it |
| Devolo AG | Markus Waechter | markus.waechter@devolo.de |
| | Anil Mengi | anil.mengi@devolo.de |
| Grenoble INP | José Sanchez Torres | jose.sanchez@g2elab.grenoble-inp.fr |
| | Raphaël Caire | raphael.caire@g2elab.grenoble-inp.fr |
| ASM Terni S.p.A. | Massimo Cresta | massimo.cresta@asmterni.it |
| | Marco Paulucci | marco.paulucci@asmterni.it |
| **WP5** | | |
| WIT | Miguel Ponce de Leon | miguelpdl@tssg.org |
| | David Kirwan | dkirwan@tssg.org |
| | Jason Whelan | jwhelan@tssg.org |
| | Hisain ElShaafi | helshaafi@tssg.org |
| ESB | John Howard | john.howard2@esb.ie |
| Alcatel-Lucent Deutschland AG | Matthias Sund | Matthias.Sund@alcatel-lucent.com |
| | Detlef Hartmann | detlef.hartmann@alcatel-lucent.com |
| FUNITEC | Agustín Zaballos | zaballos@salleurl.edu |
| | Ramon Martin de Pozuelo | ramonmdp@salleurl.edu |
| | Joan Navarro | jnavarro@salleurl.edu |
| | Alan Briones | abriones@salleurl.edu |
| ALSTOMGrid | Amir Fazeli | amir.fazeli@alstom.com |
| RWTH | Mohsen Ferdowsi | MFerdowsi@eonerc.rwth-aachen.de |
| Orange Polska S.A. | Rafał Artych | rafal.artych@orange.com |
| **WP7 Technical Management** | | |
| RWTH | Padraic McKeever | pmckeever@eonerc.rwth-aachen.de |
| | Bettina Schäfer | bschaefer@eonerc.rwth-aachen.de |
| | Antonello Monti | amonti@eonerc.rwth-aachen.de |

# Table of Contents

1. Introduction ..... 5
2. Trial Results on Usage of Generic Enablers and FIWARE ..... 6
3. Trial Results from Technical Experimentation ..... 14
4. Trial Results from Simulation ..... 18
5. Conclusion ..... 20
ANNEX 1 Comparison of GEs to Alternative Technologies ..... 21
ANNEX 2 WP1 Trial Results ..... 37

# 1. Introduction

This is a report on the results of the FINESCE trials.

FINESCE is the Smart Energy project in Phase 2 of the FI-PPP. FINESCE has performed field trials of the use of FIWARE Generic Enablers (GEs) in the smart energy sector and provides Domain Specific Enablers (DSEs) and an Application Programming Interface (API) which allow third-party clients (such as SMEs involved in the FIWARE Accelerator Programme) to develop applications which access the trial infrastructures, thus supporting the creation of a network of Smart Energy application developers.

The FINESCE trials comprise seven trial sites developed by five vertical work packages (WP) covering Smart Energy areas where Information and Communications Technology (ICT) can beneficially be applied:

- WP1: development of demand side response and demand-side management solutions for mixed-use buildings in a city district in Malmö, Sweden. The trial has tested an integrated approach of energy carriers in order to demonstrate Demand Side Management and Demand Side Response tests based on either price or energy mix ($CO_2$) for both heat and electrical loads.
- WP2 Horsens: the trial takes place in a community of single family households in a small town near Horsens, Denmark. The houses are located outside the collective (district heating) heating supply area and originally had individual oil or gas-fired boilers. For the trial, the houses have had both energy producing and consuming equipment and components installed. The trial is studying efficient grid utilisation through demand-side management of prosumers.
- WP2 Madrid: this trial tests different energy services in an office building which has a Building Management System (BMS), PV energy production, energy storage in batteries, simulated wind power, a microgrid in the building and an extensive wireless sensor network.
- WP3: studies the interwork of an industrial demand side response with a cross-border Virtual Power Plant (VPP). The industrial trial site is in a smart factory in Aachen, Germany. The VPP trial site is located in Cologne, Germany and the VPP comprises about ten Distributed Energy Resources (DER) located in Belgium and Germany.
- WP4: this trial is based a branch network of the MV/LV grid in Terni, Italy. It develops an energy marketplace to provide demand side response to varying energy production from Distributed Energy Resources (DER).
- WP5 controlling electrical vehicle charging to balance DER supply and improved utility ICT infrastructure.

This report gives an overview of the results of the FINESCE trials. The aspects covered are:

- Results related to FIWARE and GEs are summarised[2]. FINESCE has performed extensive work with FIWARE and GEs, following on from its precursor project FINSENY, which ran from 2011-2013, so that many FINESCE partners have effectively been working with FIWARE and GEs over a four-year period now. This report presents the lessons learned with FIWARE and the GEs with the objective of providing the Phase III participants with useful information and first-hand experience and to feedback the issues and problems to FIWARE.
- The results of the Smart Energy use case experiments performed at the trial sites are summarised.
- The results of simulation activities performed in FINESCE are summarised.
- Results which are published and available for use by third parties (FINESCE DSEs, FINESCE API).
- Results on the further development of the FINESCE trials based on an open platform approach.

[2] More details of results on FIWARE and GEs are available in the six "*Trial Results*" reports of the individual WPs (with separate reports for WP2 Horsens and WP2 Madrid)', and in D7.5 "*Consolidated Analysis of Generic Enablers and Domain Specific Enablers Integration*", available on http://www.finesce.eu/Results.html

# 2. Trial Results on Usage of Generic Enablers and FIWARE

The purpose of this chapter is to present how FINESCE has used GEs and FIWARE and what lessons have been learned.

## 2.1 How has FINESCE used GEs?

FINESCE has built seven separate and independent trial sites based on FIWARE. FIWARE GEs have been integrated in the FINESCE trial sites along with other software and hardware components, some of which have been defined by FINESCE as DSEs. Although the trial sites are independent, there are broad similarities in their use of GEs. The predominant pattern, as shown in Figure 1, is that the sites gather data from remote equipment in buildings or vehicles, process the data and make it available over a WP-specific API. These WP APIs are used by the FINESCE API mediator and also, in some cases, directly by FINESCE partners or internally in the trial site.

In FINESCE, GEs from the FIWARE IoT chapter are typically used for data gathering, GEs from the Data/Context Management chapter for data handling and GEs from the Security chapter for controlling access via the WP's API.

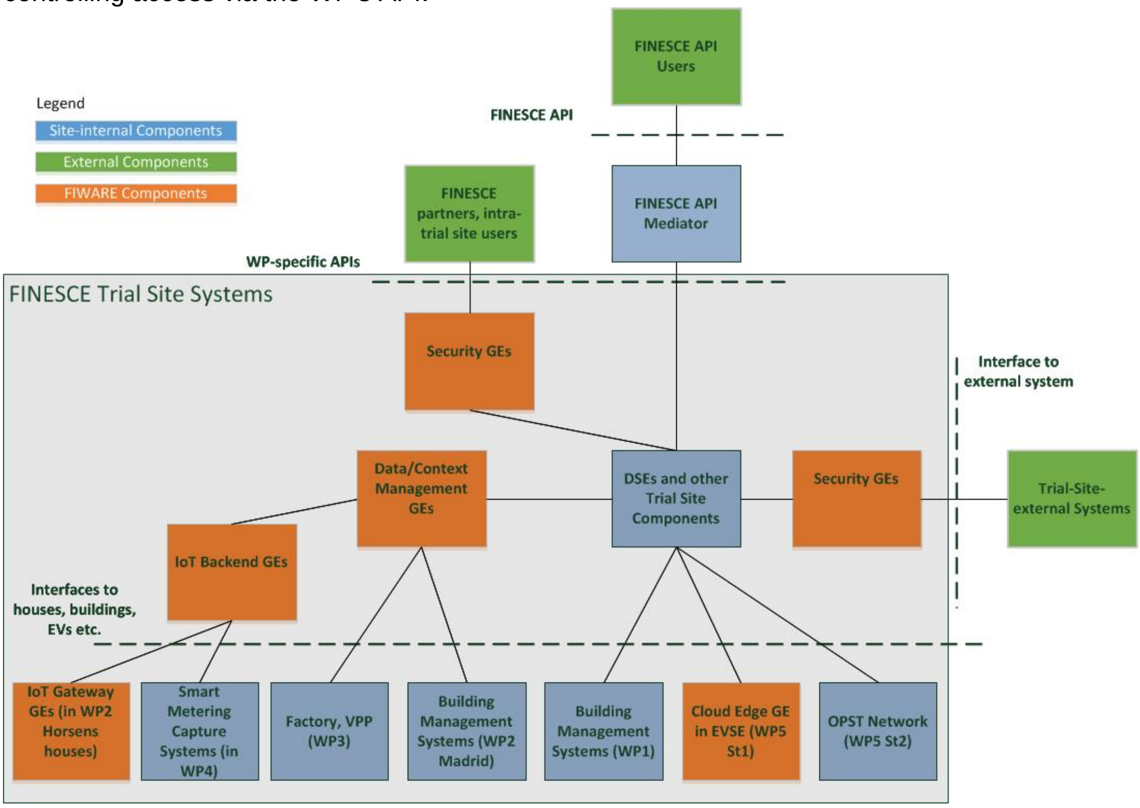


**Figure 1 – Broad Pattern of GE Usage across FINESCE Trial Sites**

A detailed overview of the FINESCE trial sites' integration architectures, including which GEs have been integrated is given in D7.5.

## 2.2 What GEs have been used?

The number of GE instantiations (GEis) that are integrated into FINESCE trial infrastructures is shown in Table 1, on a per FIWARE chapter basis. Forty-nine different instantiations of different GE implementations have been integrated into trial sites. It can be seen that the GEs of the Data/Context Management and Security chapters have been those most used: this reflects the FINESCE trials' emphasis on gathering and distributing Smart Energy data securely.

| FIWARE GE Chapter | GEis Integrated |
|---|---|
| Cloud | 5 |
| Data/Context Management | 17 |
| Apps | 6 |
| IoT | 7 |
| Security | 14 |
| I2ND | 0 |
| | 49 |

**Table 1:- GEis Integrated per FIWARE Chapter**

Table 2 gives a detailed per-WP breakdown of which GEis have been integrated. Not surprisingly, the GEis that have been integrated most often in FINESCE are the same ones which are given the most positive comments in GE evaluations, see Chapter 2.4 below. This is because these GEs provide useful functionality and good quality.

| FIWARE GEs | GEi Name | WP1 | WP2 Horsens | WP2 Madrid | WP3 Factory | WP3 VPP | WP3 Other[3] | WP4 | WP5 Str.1 | WP5 Str .2 | Totals |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Cloud Chapter | | | | | | | | | | | 5 |
| IaaS Data Center Resource Mgmt. | - | | | | | | | X | | | 1 |
| Self Service Interfaces | Cloud Portal | | | | | | | X | | | 1 |
| Object Storage | - | | | | | | | X | | X | 2 |
| Monitoring | TID | | | | | | | X | | | 1 |
| Data/Context Management Chapter | | | | | | | | | | | 17 |
| Publish/Subscribe Broker | Orion Context Broker | X | X | X | X | X | 2X | X | | | 8 |
| Complex Event Processing (CEP) | IBM (PROTON) | | | | X | X | | X | X | | 4 |
| BigData Analysis | COSMOS | X | | X | X | | | X | | | 4 |
| Publish/Subscribe Broker | Context Awareness Platform | | | X | | | | | | | 1 |
| Apps Chapter | | | | | | | | | | | 6 |
| Application Mashup | WireCloud | | | | X | | X | X | | | 3 |
| Store | WStore | | | | | | | X | | | 1 |
| Marketplace | UPM | | | | | | | X | | | 1 |
| Repository | UPM | | | | | | | X | | | 1 |
| IoT Chapter | | | | | | | | | | | 7 |
| (Backend) Device Management | IDAS DCA - TID | | | | | | | X | | | 1 |
| (Backend) IoT Broker | IoT Broker – NEC | X | X | | | | | | | | 2 |
| (Backend) Configuration Management | Orion Context Broker | | X | | | | | | | | 1 |
| (Gateway) Data Handling | EspR4FastData | | X | | X | | | | | | 2 |
| (Gateway) Protocol Adapter | ZPA | | X | | | | | | | | 1 |
| Security Chapter | | | | | | | | | | | 14 |
| Identity Management | KeyRock | X | | | | | X | X | | X | 2 |
| Authorization PDP | / AuthZForce | X | | | | | X | | | | 2 |
| PEP Proxy | / Wilma | X | | | | | 2X | | | | 3 |
| Identity Management | GCP | | X | | | | | | X | | 2 |

[3] The WP3 infrastructure comprises five parts, all of which separately use GEs. The column heading "WP3 Other" in Table 2 refers to the parts other than the Smart Factory and the VPP, i.e. the Future Internet Smart Factory Energy Planning System (FISFEPS), FINESCE Presentation Layer (FPL) and Simulation parts of the WP3 infrastructure.

| FIWARE GEs | GEi Name | WP1 | WP2 Horsens | WP2 Madrid | WP3 Factory | WP3 VPP | WP3 Other[3] | WP4 | WP5 Str.1 | WP5 Str .2 | Totals |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Data Handling | PPL | | | | | | | | X | | 1 |
| DB Anonimyzer | DBA | | | | | | | | X | | 1 |
| Content-based Security | CBS | | | | | | | | X | | 1 |
| Totals | | 6 | 6 | 3 | 5 | 2 | 7 | 13 | 5 | 2 | 49 |

**Table 2: Overview of GE Integration**

## 2.3 Why have these GEs been chosen?

FINESCE has performed an extensive process of selection and evaluation of GEs. An evaluation process has been developed and a detailed formal evaluation has been performed of the GEs which have been integrated into the trial sites. See D7.5 [4] for full details.

The basic organisation of FINESCE into independent trial sites carried over into the GE selection process, with each WP being responsible for selecting its own used GEs. The selection of GEs has, therefore, been a continuing, iterative process because the GEs themselves have continued to be developed and thus the available GE implementations, their level of maturity, the quality of the support offered by their developers and the level of user experience with them have evolved also.

WPs have used the following selection criteria for selecting GEs:

- the GE's fit to a role in the trial site's functional architecture and within the FINESCE partners' future plans. This involved firstly a study of the GEs on the level of their technical chapters, then on the level of the descriptions provided for the individual GEs (including its terms and conditions), then on the level of how the GE can fit to the trial site architecture. If the GE was included in the trial site architecture then it subsequently underwent a process of integration into the trial site, during which their functionality continued to be evaluated and more deeply understood as they were tested and debugged. GEs from some FIWARE chapters (such as Data/Context Management, Security, Internet of Things) have been of most interest to the FINESCE trial sites in the Smart Energy domain.
- the GE's terms and conditions and availability from FIWARE (in either the FIWARE Testbed, FIWARE Lab or as a downloadable product, as per the trial site's specific needs);
- whether the GE was included in the FIWARE Catalogue;
- whether the GE's documentation was of sufficient quality to allow the GE to be studied and, later, to be integrated into the trial site;
- whether there was sufficiently good support of the GE by its developers.

The formal evaluation of the GEs was made by scoring the GE's performance on a set of criteria developed from the ISO/IEC 25010:2011 specification ("Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE). The criteria headings are: Functional suitability, Performance Efficiency, Compatibility, Usability, Reliability, Security, Maintainability and Portability. Under each of these headings, a set of sub-criteria has been developed, FINESCE's requirements on the GE have been defined per sub-criterion and scores given on how well the GE meets the requirements. The result of the formal evaluation was an overall score for each GE. Full details of the results are given in FINESCE deliverable D7.5.

## 2.4 What GEs are good/bad?

FINESCE has found the following six GEs to be really useful, although with varying technology readiness levels (TRLs)[5]. Please note that the TRL assessment given here is from FINESCE.

[4] See deliverable D7.5, "*Consolidated Analysis of Generic Enablers and Domain Specific Enablers Integration*", on http://www.finesce.eu/Results.html

[5] See https://en.wikipedia.org/wiki/Technology_readiness_level

1. **Orion:** very mature as software although FIWARE Lab unreliable, Question whether its terms and conditions clear for its use long term (TRL 8)
2. **Cosmos;** good but "R' support should be included (TRL 6)
3. **Complex Event Processing**, good software (TRL 6)
4. **Identity Management KeyRock:** reliable as an entry point to the cloud, although unannounced changes in its API have reduced trust in it(TRL 5-6)
5. **WireCloud:** useful and good (TRL 6)
6. **PEP Proxy:** only recently available, so still in testing – it needs changes (TRL 4)

Table 3 gives some brief "PROs and CONs" of the different GEs which have been heavily used in the project. Deliverable D7.5 should be consulted for more detailed comments.

| FIWARE GEs | PROs (☺)and CONs (☹) | |
|---|---|---|
| Cloud GEs | +<br>+<br>-<br>- | Creating and managing VMs (and VMs with a GE image aboard) is easy.<br>Cloud portal is easy to use.<br>Too many maintenance episodes and a couple of unexpected "crashes".<br>From the beginning of 2015, no chance to increase the number of VMs (and public IPs) assigned to a region |
| COSMOS Big Data Analysis GE | +<br>+<br>-<br>-<br>- | Easy storage and retrieval of large datasets via HIVE SQL-like queries.<br>Stable, well documented.<br>Some unexpected "crashes" of the data retrieval service.<br>Reverse subscription for returning data from Cosmos to Orion needed.<br>Support for R needed. |
| WIRECLOUD Application Mashup GE | +<br>+<br>- | Excellent performance and stability;<br>Documentation is precise and complete.<br>Environment could be more user-friendly (especially for code debugging). |
| ORION Context Broker GE and Backend Configuration Manager GE | +<br>+<br>+<br>+<br>+<br>-<br>- | Plays the role it has been designed for. Key component in several trials.<br>Good support from the developers.<br>Frequent updates with advance notice.<br>GE has improved continuously. Documentation and performance good.<br>FI-LAB image makes the installation process very easy.<br>VMs with pre-installed ORION images come with a "working" partition whose size is <u>ONLY</u> 5 GBs. In a scenario of extreme logging activities (due to a huge amount of incoming data) logs may definitely saturate the partition thus blocking ORION (although it would not seem to throw/notify issues/exceptions).<br>Images have not been entirely stable (FI Lab issue, not GE issue). |
| PROTON Complex Event Processing GE | +<br>+<br>- | GE shows the expected behaviour.<br>Documentation helpful and complete.<br>UI for loading rules for evaluating incoming events is not very intuitive. |
| IDAS (Backend) Device Management GE | +<br>-<br>- | It works as expected.<br>Compatibility with the SensorML (language for representing measurement transformation of observations) limited to the version "1.0.1".<br>Debugging tools are limited only to a set of CURLs. During the verification and testing, there were no GUI tools that could be used |
| Security GEs Identity Manager KeyRock, Authorization PDP, PEP Proxy Wilma | +<br>-<br>-<br>- | Identity Manager KeyRock for Single Sign-On works well with Authorization PDP for authorisation and PEP Proxy Wilma as a gate keeper.<br>Keyrock does not support custom attributes for users.<br>Keyrock installation slow due to sparse documentation.<br>PEP Proxy behaves like prototype, version |
| Identity Management (GCP). | + | Many features, generally stable, good documentation and support. |

| FIWARE GEs | PROs (☺)and CONs (☹) | |
|---|---|---|
| IoT Service Enablement GEs | - | Difficult integration, incompatibility of interfaces between GEs. EspR4FastData not recommended for critical components yet. |

**Table 3- "PROs and CONs" of FIWARE GEs**

## 2.5 What lessons have been learned from using FIWARE and GEs?

**GEs work as a basis for building new services quickly** but must be more reliable and predictable for use in a commercial and long term context. Usage of FIWARE shortens development times through its provision of both a good set of integrable components (the GEs) using open standard interfaces and a good hosting and development environment (FIWARE Lab) where generic components can easily be deployed and new instances can be set up. The generic functionality of the GEs provides ready-made building blocks for needed basic functionality. Having these building blocks allows the developers to focus on implementing the energy application and so reduces the effort and time for development from concept to production-level solutions. For example, the GEs used to implement security aspects such as authentication have enabled a FINESCE development team that did not have prior experience with security in a cloud based application to quickly allow third parties access to trial data in a secure manner.

The **openness** and the availability of the source code of the GEs means that there is no vendor lock-in. Developers were able to choose the GEs appropriate for the solution and could adapt them to their needs. The GEs' open interfaces enable the development of interoperable applications.

The **documentation of the GEs** was mostly good and supported the developers in setting up the application. Questions and problems occurring during the development were generally answered by the developers of the GEs and by the helpdesk of FIWARE very competently, although the response time for issues with the IT infrastructure (FI-LAB) where GEs are instantiated must be substantially improved in order to be able to compete with commercial alternatives.

The **information about what is available in FIWARE** and where to use it for different purposes (test, large scale demonstration, commercial use) needs to be improved, if FIWARE is to compete with other cloud hosting solutions providing much of the same functionality. Communication between the FIWARE organisation and its users is haphazard and unreliable. There is a lack of a reliable communications channel to inform users of upcoming events such as GE downtime or maintenance. Some GEs (such as COSMOS) provide positive exceptions to this lack of communication, showing that the FIWARE communication policy and methodology towards users is not uniform. A particularly bad area is the FIWARE Catalogue, where there is absolutely no communication of actual or planned changes. If any GE is going to be deprecated or removed, it should be communicated well in advance to the customers using it, so they can adapt to the change.

FINESCE has put a lot of effort into the evaluation, testing and integration of the GEs. We have learned that the **technology is not the issue in moving to renewable energy** – the trials show that the FIWARE technology works – the issues are in the business processes and business offering to the market as services or products and in regulatory issues. The GEs selected for integration into the trial infrastructures have generally fulfilled their roles in the trial infrastructures well, particularly GEs from the Cloud, Data/Context Management and Security chapters.

There are many good ideas in FIWARE but the overall maturity has not yet reached a level suitable for a critical infrastructure or for commercial operation of solutions based on FIWARE. It is possible to deploy a set of scalable infrastructure components but much work is needed to combine these into complete solutions. The GEs tend to be presented by FIWARE as individual components rather than also presenting combinations of GEs that can work together to perform typical tasks. There are cases where GEs do not work well together: examples are that

Wirecloud GE does not work with Identity Management GE and that integration of GEs from the IoT Chapter failed due to lack of interoperability. This indicates the need for FIWARE to adopt and enforce architectural concepts supporting better interoperability between the GEs. One possibility would be to focus on a core set of GEs and hosting, and increase maturity for this selection rather than try to make a very broad set of GEs.

There were frequent availability/stability issues with the FI-LAB platform, which questions the maturity and readiness of this platform for already hosting commercial applications and services. However, the stability problems seem not to be directly related to the GEs themselves, but rather to the capacity of the underlying IT infrastructure. Therefore it could be expected that once these infrastructure limitations get solved, FIWARE will be much closer to being able to offer the performance and reliability that commercial energy services will demand.

The FINESCE trials are generally small scale. Scalability testing of the trials has generally been performed in simulation, where the GEs are not tested.

## 2.6 What are the most valuable FIWARE Technologies Developed by FINESCE?

FINESCE's architectural approach is based on building independent trial sites to perform multiple tests of FIWARE technology. The FINESCE trial sites represent an architecture where the individual use cases are implemented on separate vertically-integrated silos. The link between the trial sites is provided by the FINESCE API[6], which provides live **Data as a Service** for use by third parties. Historical datasets are available through the FI Lab Data platform[7].

During the course of FINESCE, the FINESCE partners' thinking about how to incorporate ICT into Smart Energy use cases has changed. The FINESCE partners have recognised that a **common, open-source platform architecture** is better suited to the Smart Energy domain and that FIWARE can act as the basis of this platform Such an **open Smart Energy Platform** concept supports the concept of a Utility company driven by services, called **Utility 4.0** because it represents a revolutionary development of the traditional business model of utilities. The energy services that can be built on top of such a platform are not limited to utilities, however, but can be developed by third parties to be offered to the end users. Building an open platform means also that it is not owned by a single player and that the effort of building it and the right to use it can be shared between many partners. There will be **applications to H2020** from FINESCE partners for new projects to implement new use cases based on FIWARE technologies applied in an open energy platform. The way that the vertically-integrated FINESCE trial site architectures map to an open platform architecture is outlined in deliverables D7.13 and D7.14[8].

Some of the software components which FINESCE partners have developed in the course of building their trial site infrastructures are considered to be of potential further use. These have been packaged and made freely available for public download and use as open-source Domain Specific Enablers (DSEs)[9].

## 2.7 Why use FIWARE in Smart Energy?

The transformation of the energy supply in Europe is under way. The goals are to avoid the dependence on fossil energy sources and expensive energy imports by reducing energy consumption and by integrating the increasing number of renewable energy resources into the energy supply. To reach these goals presents great challenges for utilities in Europe. They need

[6] See deliverable D7.7 "*Consolidated FINESCE API and Handbook*" on http://www.finesce.eu/Results.html and http://www.finesce.eu/FINESCE_API.html

[7] https://data.lab.fiware.org/, see deliverable D7.12 "FINESCE Data Repository in FIWARE Lab Data.

[8] D7.12 "*Consolidation of FINESCE trial sites specific data models with a Unified Data Model for the Smart Energy Platform*", D7.13 "*Mapping of FINESCE Trial Site Architectures onto Smart Energy Platform Architecture*".

[9] See http://finesce.github.io/FINESCE_DSEs.html?option=2

to contribute to these climate goals by finding new business models and customer-relationship models to conform to the new technologies developing with the changing situation including the private distributed energy resources and the future needs and requirements of their customers. The FIWARE platform and the GEs provide a setup, which enables developers to implement, test and validate new applications to meet the needs of the utilities to deal with the greater diffusion of renewable energy sources, high load applications like electric vehicles and increased prosumer activity. This Future Internet approach of reusable software components like the GEs, the cloud based services promote the standardisation of interfaces and reduce the product development costs and the development times.

The FINESCE project shows that using FIWARE and the FIWARE GEs the challenges of the energy transition can be met and that FIWARE provides Future Internet technology for the energy area. They enabled the development team to focus on implementing the energy applications and reduced the effort and time of development. This leads to an early start of implemented applications which could then be tested and refined.

FINESCE has built seven trial site infrastructures based on FIWARE and GEs. It has shown that a number of GEs provide core functionality that is needed for Smart Energy systems and that using FIWARE technology enables Smart Energy systems to be built quickly.

The key result of FINESCE is the development of the Smart Energy Platform concept for a common, open, middleware platform which interfaces towards field equipment and sensors and allows value-added services to be implemented. FINESCE has shown that the functionality of the platform can be based on the GEs which provide the basic functional generic building blocks.

One major advantage is the openness and the availability of the source code of the GEs. Because GEs employ open standards, using FIWARE does not lead to any vendor lock-in. Developers can choose the GEs appropriate for the solution and can adapt them to their needs. FIWARE-based solutions can be adapted to work with other solutions mandated by the potential customers (e.g., a utility or a DSO) due to their strict administrative policies.

FIWARE comprises an advanced Open Stack cloud-based infrastructure as well as a library of GEs which offers developers a powerful set of resources. The GEs offer a wide range of functionality and several GE chapters are particularly useful in the Smart Energy domain.

The possibility of having a test-bed such as FI-LAB, where generic components can easily be deployed and new instances can be set up, is very attractive. FIWARE gives developers the opportunity to easily realise new functionality. This possibility is also attractive to established companies, such as utilities, and encourages fast development and experimentation.

## 2.8 Compared to using another technology, how did the GEs integrated perform in the trial?

The FINESCE partners have not built trial sites using GEs and also alternative trial sites using other components. Hence, a direct comparison between the performance of GEs and the performance of alternatives is not possible. However, FINESCE partners have performed reviews of alternatives and compared them to GEs. The results of these reviews are gathered in ANNEX 1.

Some of the generic advantages of using GEs compared to commercial competition are:

- its flexibility in data management, allowing local or cloud data storage;
- its trustworthiness compared to large Cloud platform supported by large US corporations which have significant trust problems in the EU, particularly in the energy domain;
- its openness: there is no vendor lock in;
- its adaptability: you can take a GE and make your own DSE;
- access to data through FIWARE;
- the extensive support and coaching provided;
- that it is a one-stop shop from the developer perspective (functionalities and hosting).

## 2.9 Other FINESCE Results

**FINESCE API**

The FINESCE architectural concept is to build several separate trial sites based on FIWARE technology, each of which controls its own physical equipment, but to make live Smart Energy data from the trial sites available through a unified portal, the FINESCE API[10].

The FINESCE API data is available as historical data through the FIWARE Lab Data Portal https://data.lab.fiware.org/. All available FINESCE datasets can be found by searching under the tag "FINESCE" in the FIWARE Lab Data portal. These datasets are made available as open data; any sensitive information available through the FINESCE API which is not suitable to be published as open data has been anonymised or left out of the datasets. Full details are available in D7.11[11].

**FINESCE DSEs**

A DSE is an enabler, like a GE, but covers some re-usable functionality inside a domain, in FINESCE's case the Smart Energy domain. That is, FINESCE DSEs are software components that are considered potentially useful to other developers in the Smart Energy domain. FINESCE DSEs consist of a reference implementation inside one of the trial sites and downloadable specifications and code are miscellaneous. Fourteen DSEs are available for download at http://finesce.github.io/FINESCE_DSEs.html?option=2.

FINESCE has published the specifications of some other components at http://finesce.github.io/FINESCE_DSEs.html?option=1.

[10] See http://www.finesce.eu/FINESCE_API.html and Deliverable 7.7 "*Consolidated FINESCE API and Handbook*" on http://www.finesce.eu/Results.html.

[11] See Deliverable 7.11 "*FINESCE Data Repository in FIWARE Lab Data*" on http://www.finesce.eu/Results.html

# 3. Trial Results from Technical Experimentation

This chapter gives an overview of the results of the technical experimentation carried out by the FINESCE WPs on their trial infrastructure. Full details are available in the WPs' Dx.7 "*Trial Results*" reports, available on available on http://www.finesce.eu/Results.html. This chapter covers the scientific results of the individual trials. The FIWARE/GE results are in Chapter 2.

## 3.1 Technical Experimentation Results of WP1

The technical experimentation concerned:

- Proof of concept and evaluation of solution based on a distributed energy management capability and a centralised portfolio management capability;
- Proof of concept regarding cost optimisation on price signals for heat and electricity based on different business model(s);
- Increased knowledge on future potential for Demand Side Management and Demand Side Response as well as ideas on customer's potential to act as balancing power;
- Evaluation of the thermal load shifting potential of different heating systems, e.g. underfloor heating and radiators, while leveraging the building's thermal inertia.

The Big Data GE in the WP1 trial infrastructure was used to perform analysis of a dataset consisting of hourly sensor data collected from buildings of 350 district heating customers in a Swedish city over the course of one year. The results obtained concerned various different aspects of customers' energy consumption.

In the area of energy optimisation, the developed trial infrastructure has proven to be a very flexible system able to deliver benefits both on a local level, optimisation in the building, and at the same time on a global level, system optimisation. The infrastructure can handle all types of energy carriers, and is controlling electrical loads, district heating loads and district cooling loads. Five buildings are connected to the infrastructure. All originally identified use cases have been implemented, investigated and proven. These use cases include amongst others tests for load curtailment and load shifting. Some of the use cases have been found to have a commercial potential and plans are made to progress towards a commercial phase, post FINESCE.

The potential for shifting loads without significant impact on the customer's comfort has been shown to be bigger than initially expected. Therefore E.ON is now further exploring how the flexibility can be used to enable system optimisation of district heating and district cooling grids. For example, provided that desired flexibility is available, that could enable avoidance of firing up peak production units which usually have higher operational costs and $CO_2$ emissions. Thus, rolling out the infrastructure to the wider Malmö (here 5 buildings would not be enough, 50+ are required) could enable benefits to the whole City of Malmö.

Tests performed on dynamic district heating prices (optimisation of heating complete buildings) indicated savings of 5%.

## 3.2 Technical Experimentation Results of WP2 Horsens

The technical experimentation concerned:

- Access to and collection of data – near real-time and through the FINESCE API;
- Demonstration of shifting energy consumption
- Creating incentives to shift energy consumption away from peak periods

In the area of involvement of the users, i.e. the people living in the houses in the trial, four focus areas have been studied: Technology and social behaviour, Energy visualisation, External control of energy resources and Prosumers in the smart grid. Across these four focus areas, desirability, usability, flexibility and acceptance emerged as crucial topics. The study showed that nearly all participants actively engage in "hands-on" optimising guided by common sense, that the .families adjusted towards "more rewarding" consumption practices and that the sense of achieving something is essential for the users' attitude towards external control. The users emphasise the importance of issues like transparency, level of security of energy supply and shared control when asked about a future partnership with an external control service supplier.

Insights from end-user co-operation have generated three recommendations that can be emphasised as having universal relevance for any end user engaging in a “green” energy transition, interested in converting their private energy resources into intelligent alternatives.

1. Well-designed support and instructions should accompany the delivery of intelligent, green technology – both hardware and software – to encourage and empower end-users in the transition and to prevent old habits from counteracting the full potential of their investments.
2. Users ask for easy means to interact with external control services, and system interfaces must invite the users to easy-accessible, shared control features with clear feedback on “who” has control, and for how long.
3. In the case of offering control as a service, the services must be open to accommodate clear local motivators. This means that intelligent control services must be flexible and adaptive in order to meet the present price and payment conditions, as in the case of profiting from free solar energy.

In the area of energy consumption and control, the trial has shown that it is possible to build a smart energy solution for a private home based on existing components that are not all designed for this purpose. However, it is also clear that equipment built for integration with a smart home setup is needed if cheap and easy to install solutions are to be built.

The trial included a demonstration of controlled charging of 19 EVs. The tests demonstrated that:

- it is possible to reduce the peak load on the 400 V power grid by controlling the charging of EVs without reducing the users’ comfort
- in addition, controlled charging can reduce CO2 footprint and cost of electricity for the individual user.

With current technologies a scalable EV charging system can be deployed that can use changes in electricity prices or tariffs, or provide a “green profile” for users. However, in a Danish setup the business case for such solutions does not exist yet. At the same time such solutions can provide support for levelling the grid load, or can be deployed solely with that purpose.

Eight houses in the Horsens trial were selected for participation in a demonstration of controlled operation of heat pumps. These eight houses had the same type of heat pump, and hence same interface and heat production profile. Initial testing of control of heat pump operation showed issues in the configuration performed during installation of the heat pumps. The configuration is now performed and technical control of the heat pumps is possible from the software platform. The optimisation of heat pump operation has not been performed because the season for using heating in Denmark has passed.

It has been demonstrated that a system based on existing standard components and internet communication can be used to shift energy consumption. The energy consumption can be moved in time to where the consumption is “best” according to both individual constraints (such as price) and global constraints (such as overall load of the grid).

## 3.3 Technical Experimentation Results of WP2 Madrid

Office, industrial and tertiary buildings give little freedom to end users to control the building equipment generating, storing or consuming energy, control of which, is normally centralised and implemented through a BMS which integrates the operation of the different building subsystems. However, even if the larger portion of the energy savings will usually come from more optimized global control strategies, the influence of end users on the final consumption cannot be completely neglected, as they still retain certain degree of control over some of the energy consuming devices of the building. Therefore the promotion of energy awareness and energy efficient habits among end users, which has been deeply researched in the residential buildings domain, needs to be further developed in non-residential buildings in order to underpin the savings achieved through centralized control strategies.

Also, the increasing integration of renewable energy technologies in buildings call for more advanced control systems and algorithms which can optimise the matching of the energy demand of the building with own energy generation and storage capabilities, taking into account the dynamic prices and any demand management signal coming from the electricity grid. This needs to be addressed by Energy Service Companies (ESCOs), utilities or other stakeholders who will need to be provided with the necessary technological foundations for supporting the delivery of these services to multiple buildings.

The Madrid trial uses an office building equipped with a BMS, which can be remotely operated through a software platform called Building Control Centre (BCC), this platform can control the BMS from multiple remote buildings. The building also has PV panels, simulated wind power generation, and a battery for electricity storage, and a smart metering and sensing infrastructure, and a weather forecasting service has been developed to feed the building energy control algorithms.

The objective of the FINESCE Madrid trial has been to assess how the FIWARE-based infrastructure can meet the requirements for data handling, processing and exchange in the context of the provision of smart energy services for industrial and tertiary buildings.

Different energy management and energy efficiency service prototypes have been implemented and tested in the trial, including demand management services, energy auditing/monitoring/verification services, and energy generation/storage management services.

In the area of involvement of the users and stakeholders, awareness about global energy consumption has been increased through the provision of feedback to end users about the global consumption and comfort of the building, presenting also the detailed breakdown of the different building areas, in order to change behaviour to save energy in the portion of the building energy consumption over which they can have a direct influence. Users can subscribe to data and get customised alerts about the different energy consuming devices of the building.

## 3.4 Technical Experimentation Results of WP3

The technical experimentation concerned demonstrating a FIWARE-based infrastructure supporting plan-driven balancing of energy demand and supply through the mutual exchange of Smart Factory energy consumption and VPP production forecasts and whose event-oriented architecture allows for fast response times to intra-day changes and deviations from those forecasts.

The impact of the VPP operation on the electrical grid and in particular the effectiveness of the energy balancing method when distributed generation is present in the electrical grid has been demonstrated using real-time simulations. It is shown that the proposed balancing approach has a positive impact on electrical grid operation compared to a “fit&forget” principle of integrating DERs.

A visualisation application, the FINESCE Presentation Layer (FPL), has been developed for visualising the WP3 Smart Factory, VPP and simulation data.

## 3.5 Technical Experimentation Results of WP4

The technical experimentation concerned realising and testing a local energy marketplace with the specific purpose of reducing power losses due to the reverse power flow effect by aligning in near-real-time the energy demand with the energy production.

Two new models of smart meters have been installed in the field, the first one is based on GPRS communication, the second one uses G3 PLC communication and are read every 5 minutes. Both methods are working and do not interfere with one another.

A near-real time simulation infrastructure is included in the trial site architecture and performs power production and consumption predictions which matched actual measured values with good accuracy.

## 3.6 Technical Experimentation Results of WP5 Stream 1

The technical experimentation concerned electricity balancing, showing how both temporal and geographical supply-demand imbalances, resulting from the ever increasing use of renewable power in electricity grids globally, can be addressed by having users' energy demand (in the form of when they may charge electric vehicles) track the supply of energy from renewable sources, the opposite of the conventional approach, in a fully operational charging optimisation system (COS), serving real customers. Two use cases have been studied:

1. Grid emergency: this is where a fault results in a major drop in power generation or supply, and emergency action to reduce electric vehicle charging load in order to avoid blackouts. COS has been designed to ensure that a malicious user cannot initiate a grid emergency event. The critical parameter in this case is the charging optimisation system's speed of response, the faster the response the greater the economic value to grid operators. Tests showed that average response times of unencrypted and encrypted event messages were 462.2 ms and 166.8 ms respectively, which fulfils the requirements set by the DSO/TSO.
2. Supply-demand balancing: tests have been performed on the how quickly electric vehicle charging can be controlled. Response times of 300ms on LTE and 640ms on WiMAX were measured towards the real electric vehicles in the trial. Large-scale tests performed towards 200, 1000, 5000, and 20000 simulated electric vehicles showed that COS processed the data streams from the electric vehicles in 1.07, 3.12, 3.60, and 11.83 seconds, respectively. COS's distributed architecture means that it can scale to support more electric vehicles.

The latency and availability of the communications between COS and the electric vehicles was measured. The latency was below 1 second, which is very satisfactory. The overall communications availability was found to be 98.5% meaning that about 1.5% of potential interruptible load cannot be interrupted on average due to communications problems.

## 3.7 Technical Experimentation Results of WP5 Stream 2

The WP5 Stream II trial uses a network of FIDEV (FINESCE Devices) platforms as a distributed storage system that provides high-availability and reduces the latency in acquiring data from the local sites of the utility while offering a secure solution for sharing data with external stakeholders.

The experimentation scenario concerns data storage and replication between nodes located in two different cloud environments, public (and remote) and private (and distributed around the utility facilities), creating a hybrid cloud. These nodes can replicate information between themselves with the goal of storing information in several storages with authorised access from anywhere.

Security requirements for a hybrid cloud managing Smart Grid data has been defined jointly by utility experts and academia. Security issues have been analysed and solutions for known issues incorporated into the trial infrastructure design. Some vulnerabilities that are found in the cloud scenario deployed using GEs are considered as highly critical if GEs such as Object Storage or Identity Management Keyrock are wanted to be used for commercial reasons. However, the solutions found and suggested to fix these issues could be easily incorporated by the GE developers.

The use of the TRILL layer 2 protocol has been tested, validating this protocol as a possible solution to provide a data communication network with redundancy, high availability, and fast convergence times in case of data link faults. TRILL provides a data network with a minimum configuration of the nodes, allowing easy deployment and interconnection of new nodes, which are automatically discovered These features could help to achieve the requirements of some Smart Grid communications use cases, such as electrical network fault detection, teleprotection, or selfhealing. Difficulties in large scale deployment have been found, depending on the intermediate network hardware.

# 4. Trial Results from Simulation

This chapter gives an overview of the results of the simulation work carried out by the FINESCE WP. Full explanations and details are available in the WPs' Dx.7 "*Trial Results*" reports, available on available on http://www.finesce.eu/Results.html.

## 4.1 Simulation Results from WP1

**Study of Building Thermal Flexibility**

The aim of this work is to develop a method for the identification of the building thermal flexibility, described by the parameters of a simplified building model with a clear physical interpretation. The determined thermal flexibility will describe the capability of the building to act as short term heat storage and will therefore represent its load shifting potential for district heating networks. The development of the method aims its applicability on different building types without significant adjustments.

This work has presented a method for assessing the building thermal characteristics based on collected measurement data and basic building information. Test parameter estimation was done for different building models to find which model was simple enough to be parametrised through fitting of measured data but complex enough to reproduce the building thermal response accurately enough for the use in Demand Side Management (DSM) measures. After a process of filtering and completion of the input data, an optimisation algorithm was applied to fit the simplified building model to the available measurement data. The model parameters estimated by the optimisation algorithm allow for the direct derivation of the building thermal flexibility, the building time constant and therewith the potential of the building for DSM measures. The wide applicability of the building model structure and the parameter constraints for the optimisation, developed within the method, allows for an implementation on different buildings types. The qualitative analysis of the model temperature prediction and the evaluation of the estimated model parameters revealed that only the two capacity building model with an additional consideration of the indoor air as a massless node (I-E-A-model) combines an accurate qualitative reproduction of the indoor temperature fluctuations and a clear physical interpretation of the estimated parameters.

An evaluation of the building thermal flexibility using the I-E-A-model determined that the building can maintain a comfortable indoor temperature (less than 1 K temperature decrease) over 20 hours for an average outdoor temperature of 3°C and a heating reduction of 70%. The high thermal flexibility of the building was confirmed by step tests performed at the real building.

**Modelling a $CO_2$-steering signal for Demand Side Management in district heating grids**

Usually the operation of combined heat and power plants for district heating grids follows the demand. Peak units, which have higher $CO_2$ emissions and are most cost intensive, are only used in high demand periods. Demand Side Management can make use of the thermal flexibility of buildings, given by their thermal inertia or thermal water storages, to enable optimised operation of the heating plants so that firing up the peak units can be avoided.

This work derived a $CO_2$ steering signal based on generation data from the heating plants plus hourly outdoor temperature in Malmö of the years 2011 and 2012. The derived signal can be applied to the customers to give incentives for shifting demand into times of lower $CO_2$ emissions.

## 4.2 Simulation Results from WP2 Horsens

The houses in the Horsens trial have electric vehicle charging, heat pump consumption and photovoltaic generation. The houses are in the Stenderup village whose LV/MV grid is a typical rural grid with a radial distribution system. The simulation investigated scalability by analysing what would happen in the Stenderup grid if further buildings were similarly equipped as the trial houses by performing a power flow analysis. Two scenarios were analysed: a very sunny day with of high global irradiation and therefore a lot of PV energy production and a cold winter's day where the heat pumps are heavily used.

The results for the PV-generation scenario show that in particular around noon, where the PV generation is the highest, the transformer loading in some parts of the area is overloaded up to 120 %. The voltage profiles at the nodes are increasing significantly during high PV generation but are never exceeding the limits of 1.1 pu.

The results for the heat-pump scenario show that the increased consumption of heat pumps and electric vehicles will cause transformer overload and huge voltage drops along the lines. The limit of 0.36 kV respectively 0.9 pu will be violated several times per day.

## 4.3 Simulation Results from WP3

The scalability of the energy balancing mechanism has been studied through two simulation scenarios.

In the first scenario the VPP operates DERs and industrial loads located in two different areas from an electrical grid perspective but with the DERs close to the loads. The local balancing in each area is then the same as in the actual trial (with a single industrial load). The simulation results showed that the global energy balancing, i.e. considering both areas, is improved by having two areas, showing an improved voltage profile and reduced import of external power. This is a positive indication of the scalability of the energy balancing mechanism.

In the second scenario more DERs and more industrial loads were added into a single area. The simulations showed that having too few loads results in a worse voltage profile and more need to import external power.

## 4.4 Simulation Results from WP4

The simulator has been developed as part of a DSE which is an integral part of the trial site infrastructure. Its near-real-time results are crucial to the functioning of the energy marketplace application. The simulator's tasks are to: *a)* compute power losses and voltage drops of an unbalanced distribution network with the insertion of distributed energy sources, and *b)* Predict the power consumption and production of every user in the Trial Site, 1h, 3h, 6h, 12h and 24h ahead. The simulation highlighted the importance of computing the power losses for a distribution network as decision factor for the Market Place.

## 4.5 Simulation Results from WP5 Stream 1

The WP5 Stream 1 trial site has been implemented in such a way that EV charging schedules proposed by the Charging Optimisation System (COS) are automatically approved, i.e. no charging authorisation system that checks the effect of the proposed schedule on grid conditions and approves or rejects it has been implemented.

A data-driven approach to implementing a charging authorisation system has been designed, based on Artificial Neural Networks and has been tested in a simulated grid environment with good results.

Use of LTE for communications to electric vehicles and smart meters has been studied in an environment containing a simulation of the electrical grid combined with real LTE communications. The performance of two versions of LTE has been compared.

# 5. Conclusion

FINESCE has built seven separate trial sites, each implementing individual Smart Energy use cases in a trial infrastructure based on FIWARE and GEs. The FINESCE partners have developed extensive experience of FIWARE and GEs by working with them as they have developed. Finally FINESCE integrated a total of 49 GE instantiations. The integrated GEs represent those GEs whose functionality is most suited to the Smart Energy domain and which have the best quality in terms of functionality, usability and support.

The FIWARE platform offers a development environment and software building blocks which together allow the speedy development of applications. Some of the GEs are considerably more mature and usable than others. It would be a useful improvement if all GEs could adopt the level of support offered by a high quality GE such as Orion Context Broker.

FINESCE partners have gained a positive experience with FIWARE, particularly in the Data/Content Management and Security chapters. The overall impression of the GEs is positive, but the overall maturity of the GEs and the FIWARE platform has not yet reached a level suitable for delivery of commercial services. Further work is still needed to build completely reliable solutions.

Live data from the FINESCE trials is available over the FINESCE API and open historical data over the FIWARE Lab Data.

A set of components developed by the FINESCE partners especially for their trial infrastructures, but considered to be useful for third parties, has been published as open-source DSEs.

FINESCE has developed the concept of a common, open-source FIWARE-based platform architecture for the Smart Energy domain. The intention is to develop the platform in future projects.

FINESCE partners have carried out a set of experimentation based on the use cases trialled in their infrastructures. Simulation has been used extensively in order to test scalability, to study particular scenarios or as a component in the live trial site. Detailed results of the trial experimentation and simulation results have been published in separate reports.

# ANNEX 1 Comparison of GEs to Alternative Technologies

## 1.1 WP1 Comparison of GEs to Alternative Technologies

WP1 has compared *Big Data* GE with Hortonworks using the FINESCE GE evaluation criteria (see Deliverable D7.5 for details).

| **Category/Criteria** | **FI-WARE Big Data GE Score** | **Hortonworks Score** | **Comment** |
|---|---|---|---|
| | | | |
| **Functional suitability** | | | |
| Functional completeness | 15 | 20 | Better support from Hortonworks, cloudera et.al. |
| Functional correctness | 25 | 25 | Better support from Hortonworks, cloudera et.al. |
| Terms and Conditions | 25 | 25 | Would probably go for Hortonworks instead |
| **Performance efficiency** | | | |
| Time behaviour | 15 | 20 | Unknown during load, however the initial mock data seems to be performant |
| Resource utilization | 15 | 20 | Hard to determine as this stage |
| Capacity | 15 | 25 | Hard to determine as this stage |
| **Compatibility** | | | |
| Co-existence | 12 | 12 | |
| Interoperability | 3 | 3 | Integration with the Context broker is broken! Std. Interfaces from the Hadoop community seems not been hidden or otherwise broken therefore the 4... |
| Standards compliance | 12 | 15 | In an open Hadoop platform we can e.g. have scoop to speak to other databases out of the box etc. |
| **Usability** | | | |
| Appropriateness recognisability | 3 | 15 | Three wiki pages + a few presentations is not documentation. Have a look at http://docs.hortonworks.com/ and https://ilearning.seertechsolutions.com/lmt/clmsCatalogSummary.prMain?site=hw to compare |
| Learnability | 0 | 15 | We have, almost exclusively, looked on all other documentation (on the net) instead of the ”non” documentation. |
| Quality of Documentation | 5 | 25 | See above comments – this is a driver of high consumption of time spent (in the area of ”lock-down” into FILAB way) |
| Operability | 6 | 12 | |
| Support for Implementation | 15 | 20 | |
| **Reliability** | | | |

| | | | |
|---|---|---|---|
| Maturity | 15 | 20 | Unknown |
| Availability | 15 | 20 | Unknown |
| Fault tolerance | 9 | 12 | Unknown |
| Recoverability | 9 | 15 | Unknown |
| **Security** | | | |
| Confidentiality | 15 | 15 | In our understanding we have our own data nodes but the environment is completely unknown |
| Integrity | 15 | 20 | Unknown |
| Non-repudiation | 9 | 9 | Unknown |
| Accountability | 9 | 12 | Unknown |
| Authenticity | 15 | 20 | Unknown |
| **Maintainability** | | | |
| Modularity | 4 | 4 | |
| Reusability | 3 | 4 | |
| Testability | 4 | 5 | |
| **Portability** | | | |
| Adaptability | 4 | 4 | |
| Installability | 15 | 20 | |
| Replaceability | 15 | 25 | Unknown |
| | | | |
| **Total Score** | **322** | **457** | |
| | out of | out of | |
| | **545** | **545** | |
| | 59,08% | 83,85% | |

## 1.2 WP2 Horsens Comparison of GEs to Alternative Technologies

### 1.2.1 Identity Management: GCP

Because security is difficult and because it is crucial to most applications, WP2 being no exception, delegating it to an external subsystem (the GCP) has some excellent advantages. However, it also comes with a host of issues that would not exist if identity management was handled by the application itself, using libraries designed specifically with the same purpose in mind rather than delegating it externally. This is the alternative we will compare the GCP to.

Specifically, because the applications written in WP2 are created in the language C#, we will compare it to the newly developed ASP.NET Identity (https://aspnetidentity.codeplex.com/), hosted in an OWIN environment (http://katanaproject.codeplex.com/) (ASP.NET Identity+Katana). These are libraries designed to make hosting and implementation of identity management easy. It covers identity issues such as secure storage of identities, various ways of authenticating users, authorizing access to resources and handling creation of cryptographically secure tokens, which on many points, is the same identity issues that the GCP is specified to solve.

In general, the GCP is providing the functionality promised based on the tests we have performed. However, the needs in the Horsens trial are very limited compared to the very broad functionality provided by the GCP. Therefore, the benefit of integrating the GCP, compared to the cost of integrating a complex third party component, may not be big enough, when something like ASP.NET Identity+Katana is integrated directly into the framework used in the development of the Horsens trial.

Functional suitability

On the point of functional correctness/completeness, both the GCP and ASP.NET Identity+Katana have support for everything WP2 requires: OAuth2, management of identities etc. On this point, the GCP will handle everything and expose the functionality through a REST API. The only negative point that can be said about this is the support for the OAuth2 client credentials flow (described in the integration section).

ASP.NET Identity+Katana on the other hand does not handle everything out of the box, but it does allow one to implement, with relative ease, everything needed through simple call-backs you can hook yourself into. It is a pluggable system that allows you to decide the implementation, while the framework takes care of the protocols (for example OAuth2). This means that you are partially responsible for the functional correctness and completeness of your OAuth2 implementation. It also means that you must have some knowledge of the security aspects of the protocol you are implementing call-backs for, in order for it to be secure and complete.

The GCP can be used for free as part of the FI-PPP programme by partners, but is hosted by DT, whereas ASP.NET Identity+Katana is completely free and integrates directly into the product being implemented.

Performance efficiency

When it comes to performance ASP.NET Identity+Katana wins by default because it is not an external subsystem.

One of the primary issues with integrating with the GCP is that, it alone can authenticate the users. What this means is that every time a request is received, we must ask the GCP for the identity of the user, because we are unable to do this locally on our own resource servers. The issues that this causes are: potential bottleneck in the system and additional latency to every request that has to be serviced.

Compatibility

ASP.NET Identity+Katana lives only in a .NET or Mono environment. It is meant to be a library as part of applications you make in .NET and nowhere else. Therefore, if the application is not implemented in .NET/Mono, ASP.NET Identity+Katana is not an option.

The GCP, on the other hand, is an external subsystem, so it is language agnostic.

In terms of interoperability and standard compliance both score high because both of them comply with all standards and protocols used by WP2.

Usability

The GCP is relatively easy to learn to use as it has a well-documented REST API, that supports OAuth2 (client & server), OpenID, identity management, email and much more.

ASP.NET Identity+Katana has a .NET API that you integrate with, so the effort required to integrate is difficult to compare. It also supports OAuth2 (client & server), OpenID, identity management, email and much more. Its documentation is also good, and many people are adopting it, so it is very easy to find solutions to common problems.

It is worth mentioning here that if your platform of choice is .NET the effort required to use ASP.NET Identity+Katana is almost non-existent. This stands in sharp contrast to integrating with an external subsystem such as the GCP, which takes a lot more effort, by the virtue of it being external and the fact that you must somehow correlate external identities to resources that you store on your own servers. This adds an additional layer of complexity if using the GCP in comparison to just ASP.NET Identity+Katana.

It is also worth mentioning that ASP.NET Identity+Katana and the GCP are not mutually exclusive. You can easily use ASP.NET Identity+Katana to authenticate against the GCP, which is also, partially, what we are allowing in the API of our resource servers.

Reliability

In terms of maturity, ASP.NET Identity+Katana is very new and has yet to be proven. The GCP on the other hand is mature and proven.

In terms of availability, fault tolerance and recoverability, we have not experienced any issues with either of the two systems, so both score high.

Security
The primary concern of both the GCP and the ASP.NET Identity+Katana is authentication, that is ensuring the identity of the user (at least for WP2), and they both do this well. Accordingly, they both score high in authentication. However, for the most part, neither of the two deals with confidentiality, integrity, non-repudiation and accountability. This is instead a primary concern for resource servers when resources are being accessed.

Maintainability
The GCP is separated into distinct modules that each serves a purpose that can be turned on or off for a specific service. Many of these modules are not required by WP2, so this is a good feature. Because it is an external system, it also allows you to use the same identities across multiple different systems by default.

ASP.NET Identity+Katana is also very modular. This is on a library level though. For example, if you need to implement an OAuth2 backend (which WP2 needs) then you simply roll in the "OAuth2" middleware, if you want your users to authenticate against external OAuth2 providers, such as google, then you simply roll the "google" authorization middleware. If you cannot find a package for a specific purpose, you can with relative ease roll out your own. ASP.NET Identity+Katana is also storage agnostic, meaning you can choose to store the identity data wherever you please, without having to deal with all the difficult security logic, such as hashing of passwords, generation of security tokens, etc.

Portability
The GCP is offered in a SaaS fashion, so it cannot be scored on these points.

ASP.NET Identity+Katana offers good backwards compatibility on updates, so your code will not require updates after upgrading. It is also very easy to install and update through the nuget package manager, which also easily allows you to discover, when there are updates available.

### 1.2.2 Big Data Analysis – Cosmos

In the case of the Madrid trial, a comparative assessment has been done between one of the main FIWARE blocks, which have been integrated, the Big Data Analysis – Cosmos GE, with one of the main alternatives available in the market, which is the BigQuery product belonging to the Google Cloud Platform. It should be noted that the Orion Context Broker GE is also involved in the comparison, since it has been necessary to use it for integrating the Big Data GE.

In order to understand the comparison, it is necessary to recall the process followed in order to use the Big Data GE in the Madrid trial. The aim of this process is to enable a programmed delivery of data in XML format from the different trial subsystems to the GE. For doing so it is necessary to correctly configure at terminal level both the Orion Context Broker GE and the BigData GE. This configuration consists of the following steps:

- Creation of entities in the Orion Context Broker GE with the necessary parameters.
- Creation of subscriptions to the created entities for the required parameters.
- Creation of the dataset and tables in Cosmos, where data will be stored.
- Launch a process (which execute continuously) for parsing the XML files sent from the trial to be injected in Cosmos.
- Program the data output from the trial subsystems, through Java programmed tasks which send the most recent data through REST requests.

Once the data loading process is completed, it should be noted that in order to perform the analysis, there is no interface provided for making tests, so the analysis are launched directly in HiveQL on the datasets which have been created. In order to check these analyses, a Java client program has been developed in order to insert them into a database.

If we wanted to perform a similar task using Google Cloud Platform instead of FIWARE, the process for loading the data to BigQuery would consist of the following steps:

- Creation of the necessary datasets through an interface provided by the platform.
- Building the data structure that will be sent.
- Use of client program (for instance a Java application), with OAuth2 authentication, which will make use of a REST API (https://cloud.google.com/bigquery/docs/reference/v2/) for programming data deliveries.
- Building in JSON the data bundle which needs to be sent, and send the request through the aforementioned Java application.

Once the data are loaded, in the case of the Google Cloud Platform, a web interface is provided which allows launching the analyses on the datasets. Test datasets are also available for checking beforehand the effectiveness of the analyses. The use of this web interface is quite intuitive.

After having outlined the process of data loading and data analysis in both FIWARE and Google Cloud Platforms, the following subsections provide more detailed comparative results grouped under different analysis criteria.

Functional suitability

Functional suitability has to be assessed by checking the capabilities of both FIWARE and Google Cloud Platforms to perform the following tasks:
- Support periodical delivery of data through web services, and injection of this information into a big data cluster.
- Perform periodically a set of analyses on the injected data.
- Store the results of these analyses in a database external to the big data cluster, so they can be shared or visualized.

All of these tasks can be accomplished with both platforms in a similar way, with some particularities. For instance, in the case of FIWARE, it is necessary to use the Orion Context Broker GE in order to parse the information before inserting it into the big data cluster. Therefore each platform has some differences in the way data are loaded into the cluster and in the way the data analyses are performed, but in general terms both platforms cope with the functionalities which are required from them.

Performance efficiency

Both platforms use Hadoop file system and MapReduce technology for performing queries to huge numbers of files in an efficient manner, so the tests done over datasets with FIWARE and Google Cloud Platform rendered similar results, although the latter proved to be slightly faster.

Compatibility

Considering that both the BigData GE from FIWARE and BigQuery from Google Cloud Platform are external subsystems, we can make abstraction of the way in which they are programmed internally, and focus on the way the communication with them is carried out. Both platforms make use of a standard REST API, so the compatibility offered is quite similar.

Usability

In this area some important differences have been spotted. In the case of the Google Cloud Platform, the configuration overhead is quite low because it provides an intuitive platform for supporting the configuration of all the necessary blocks.

On the other hand, the BigData GE is lacking such interface, which makes it necessary to carry out these configurations by accessing directly all the machine instances which have been created in the cloud cluster, and performing all the steps which have been described at the beginning of this chapter for the Orion Context Broker GE and Big Data GE configuration process.

For these configuration steps, which are rather complex, the development team of Telefónica I+D provides support in case there is any problem, but a configuration interface similar to the one provided by Google Cloud Platform would be much appreciated.

Reliability

Regarding the reliability issues, it has to be considered that the BigData GE is an environment under continuous development, and this has compromised its reliability during the testing periods, due to maintenance tasks and successive upgrades applied to different underlying components, such as the platform for the creation of new virtual machine instances, the data cluster or the programs that are executed continuously for sending data.

The Google Cloud Platform has not been tested as exhaustively as the BigData GE, so despite no problems of service outage have been encountered, from this it cannot be concluded that Google Cloud Platform is more reliable than the FIWARE platform, apart from the fact that we are comparing a system which has been under development to another one which is currently stable.

Security

In terms of security, both platforms are equipped with similar protection mechanisms, which include IP filtering of the machines that try to access for performing certain actions (e.g. restriction for allowing the storage of contents from a specified IP in the public API), as well as the use of OAuth2 authentication for the client applications.

Maintainability

Since both FIWARE and Google Cloud are external platforms, any change or upgrade in the way of injecting data or extracting the results of the analyses applied would impact in a similar way the applications which make use of these platforms, for instance, the data parsing procedures could need an upgrade, but no significant difference in this respect has been detected between both platforms.

Portability

Both platforms are provided as SaaS, so no difference can be spotted with regards to portability.

As general conclusion from the comparison, it can be said that FIWARE Big Data GE (together with the Orion Context Broker GE) provide a similar functionality as BigQuery from Google Cloud Platform, and can receive a similar score under almost all evaluation criteria, if we can abstract from the fact that we are comparing a platform under constant development with another one already stable. The only evaluation criterion under which FIWARE lags clearly behind Google Cloud Platform is usability, due to the lack of a web interface for supporting the configuration and analysis processes. Therefore further work in FIWARE in the direction of enabling adequate interfaces would be desirable, in order to ensure a wider uptake of the platform.

## 1.3 WP2 Madrid Comparison of GEs to Alternative Technologies

During the evaluation of FIWARE Generic Enablers, an assessment was carried out for comparing the performance of one of the GEs integrated in the Madrid trial, the Big Data – Cosmos GE, and a commercial alternative, the BigQuery component of the Google Cloud Platform. The detailed description of this comparison process and of its results was included in deliverable 2.3.2 “Midterm Analysis of GEs and DSEs”. The comparison was made according to different common criteria agreed with the rest of FINESCE trials, namely: functional suitability, performance efficiency, compatibility, reliability, security, maintainability and portability.

It is remarkable that even at the early stage when this comparison was done, the evaluated FIWARE GE did quite well when compared with a commercial alternative from a top competitor like Google. Nevertheless some issues were detected where Big Data – Cosmos GE lagged behind, particularly in terms of usability and ease of configuration. Other issues were more related to the fact that FIWARE is a platform still under development, while the evaluated alternative is an already stable commercial solution, a fact which occasionally originated availability and maintainability problems. Nevertheless, it can be expected that these will be mostly solved once the FIWARE development is stabilized and its supporting IT infrastructure is dimensioned according to the demand that it will receive from multiple users.

The comparison which was performed with the Big Data – Cosmos GE could be further extended by comparing it with additional commercial alternatives from other vendors or by comparing other integrated FIWARE GEs with any commercial counterpart. However, this work has been already carried out by other trials (i.e. all the trials together give an overall picture of the performance of the different FIWARE GEs when compared with commercially available solutions), and therefore it is more interesting to look at the general advantages provided by the own concept of FIWARE platform, than to the individual technical performance or functional suitability of each of its components, namely:

- FIWARE platform relies on an open specification which can be implemented and maintained by any company/organization, which helps to avoid dependence on a single provider, one of the key features when evaluating the acquisition of a platform of this type.
- As a consequence of the previous point, very large organizations such as the main energy utilities may consider the option to deploy and maintain their own FIWARE instantiation in a private cloud.
- The fact that FIWARE is the result of a European research effort gives more confidence to customers in Europe and other world regions for storing and processing their energy data within this platform, rather than other platforms with which they may have privacy concerns.
- Furthermore, the FIWARE Accelerator Programme which is currently running supported by European funding is another opportunity to fine tuning the performance of FIWARE platform and its GEs, as it will have to cope with the demanding requirements of an extensive base of entrepreneurs who want to use FIWARE as the technological foundation of their businesses.

## 1.4 WP3 Comparison of GEs to Alternative Technologies

### 1.4.1 Alternatives to the GE Publish/ Subscribe Context Broker

The Orion Context Broker is an implementation of a Publish/Subscribe Context Broker GE belonging to the Data Management and Context Management chapter. To our knowledge, there is no single solution comparable to what Publish/Subscribe Context Broker - Orion has to offer: a combination of a key-value store and a publish/subscribe messaging service. The implementation of the NGSI-9 standard provides the functionality of a service (data source) directory, which in general is a functionality of key-value stores such as Memcached[12], Riak[13], Scalaris[14] and others. Instead of being a general solution, however, Publish/Subscribe Context Broker - Orion has a focus on metadata describing data providers and consumers. The data provider who creates an entity in a Publish/Subscribe Context Broker - Orion instance declares a key - a name of the entity - and provides the value - the address of the service providing entity's data.
In the NGSI-10, the implementation also provides brokerage of the actual data emitted by each registered entity. This, again, can be likened to a key-value store that uses a prescribed structure of the values. The key here is the name of the entity, while the value is a collection of attributes and attribute metadata as well as additional information on the entity prescribed by the NGSI-10 standard. Updating the entity's attribute values in the GE overrides any previously set values of the attribute, while querying for the entity's attributes retrieved the most recently set values. This is a normal behaviour of the key-value stores.
In both standards, the functionality of having data at rest at the broker is extended with the ability to push the changes to the registered services given that the conditions for triggering the push are met. This is a functionality that could be provided with solutions such as the Java

[12] Memcached: http://memcached.org/
[13] Riak: http://basho.com/riak/
[14] Scalaris: https://code.google.com/p/scalaris/

Message Service (JMS)[15] or one of the Advanced Message Queuing Protocol (AMQP)[16] (e.g., RabbitMQ[17]).

Similarly to the message bus implementations, the Publish/Subscribe Context Broker - Orion GE offers a loosely coupled communication between a data producer (e.g., a power meter) and data consumers (e.g., a status display in a web application). The message buses use topic names to name the channels that the data producers send the messages to. The data consumers subscribe to the topic names to receive the messages. In the Publish/Subscribe Context Broker - Orion GE implementation, the topic name is the same name of the entity as it was used to create or update the entity's context. This mostly caters for the IoT applications, where entities emit more or less regularly data on their attributes.

As demonstrated in this trial, however, the developers are free to involve this GE in other, more complex use cases, such as notifying the Energy Balancing services of new or deviated plans, so that it then returns the optimal result through the same channel. All it takes is that administrators deploy the GE, making it readily available to accept entities and context updates. This is in contrast to competitor solutions, which would require custom implementations before they could be used in the suggested way.

The use of open standards, namely the FIWARE NGSI-9 and NGSI-10 has further advantages. Rather than using a binary protocol that requires its specific set of libraries to communicate with the message bus, any client capable of making HTTP calls over the network can be used as either the data producer or the data consumer. The data exchanged is humanly readable, which is important during the development and debugging of an application. By complying with the open standards, any part of the architecture is also fully interchangeable, and this includes the Publish/Subscribe Context Broker - Orion GE which may be exchanged to another possibly more efficient GE implementation. The services can be deployed across networks, as long as their endpoints are reachable from each client end, which is not always possible without additional message bus bridges when using JMS or RabbitMQ.

The choice of using the GE Publish/Subscribe Context Broker in FINESCE has proven valuable at the stage of involving the Open Call SMEs. The new partners were able to quickly adopt the GE Publish/Subscribe Context Broker - Orion instances set up by the original partners to be used as an effective web service interface, being able to choose from passively receiving of the data or actively polling the instance. Normally, the publish/subscribe message buses do not enable accessing the last data in a topic to be at rest and available for newly arrived message consumers. Instead, they have to first subscribe to a topic, and only receive the values with the next update. Further, by properly organising the topology of the Publish/Subscribe Context Broker - Orion instances it is possible to create web services which at different access points provide different entity attribute availability. Also, having a common well-defined standard enabled fast integration. If we had decided for any of the alternate technologies, we could have expected steeper learning curves and longer integration times.

### 1.4.2 Alternatives to the GE Complex Event Processing (CEP) and GE Gateway Data Handling

Complex Event Processing software, or engines as they are usually called, is a domain widely researched and a quite well addressed market. There are several advances from research side led by universities such as

- Stanford University[18]
- Cornell University[19]
- UC Berkeley[20]
- Karlsruhe Institute of Technology[21]

They all offer prototypes or ready to use software packages for the described tasks of analysing incoming events in real time.

On the commercial market, also different vendors offer their solutions, e.g.:

- Microsoft StreamInsight[22]

[15] Java Message Service: http://www.oracle.com/technetwork/java/index-jsp-142945.html
[16] Advanced Message Queuing Protocol: http://www.oracle.com/technetwork/java/index-jsp-142945.html
[17] RabbitMQ: http://www.rabbitmq.com/
[18] Stanford Stream Data Manager: http://infolab.stanford.edu/stream/
[19] Cayuga: http://www.cs.cornell.edu/bigreddata/cayuga/
[20] TelegraphCQ: http://telegraph.cs.berkeley.edu/
[21] SpoVNet: http://www.spovnet.de/

- JBoss Drools[23]
- SAP Event Stream Processor[24]
- IBM Active Middleware Technology[25]
- Software AG Apama Analytics & Decisions Platform[26]
- Esper[27]

The last software also builds the foundation for the GE Gateway Data Handling which shows alternative software to the GE Complex Event Processing (CEP), the GE Gateway Data Handling itself. Both offer comparable functionality in terms of analysis of data streams with pre-defined patterns. The GE CEP however targets larger data streams and facilitates cloud based architecture, whereas the Gateway Data Handling is meant for on-premises deployment and a pre-filtering of data. However, the FIWARE documentation gives not a clear point when to switch from one solution to another, though this might not be necessary.

Both event processing GEs, the CEP and Gateway Data Handling, offer very good functionality, configurability, and scalability. There is no loss against existing commercial and open source solutions, interoperability of the software allows for an easy exchange of the software.

In general, CEP engines are rather easy to substitute and do not offer any specific advantages in terms of functionality. But, the distinct advantage of both GEs is their embeddedness into FIWARE and the surrounding software tools such as the Publish/ Subscribe Context Broker. This allows for a much easier and more seamless integration of the CEP engine into existing environments. This creates a convincing value proposition for FIWARE.

## 1.5 WP4 Comparison of GEs to Alternative Technologies

This chapter presents a brief comparison between some FIWARE GEs and products available on the market offering the same functionality.

For each GE, we have divided this activity into two tasks:

- Selecting the “alternatives”; e.g. alternative products/services available on the market (either Open Source or COTS);
- Defining a series of extra-features that a user may consider "an advantage to gain" when comparing the FIWARE GEs to the “alternatives”.

### 1.5.1 Alternatives on the Market to FIWARE GEs

The “alternatives” have been identified among those offering the same functionalities/services as the ones provided by the FIWARE GEs.

The following table shows the results of this analysis:

| FIWARE GEs | “Alternatives” products/services |
|---|---|
| Cloud GEs | Cloud services from TELCO and IT companies (e.g. Amazon AWS, Google Cloud, Microsoft Azure) |
| COSMOS Big Data Analysis GE | Cloudera’s Hadoop, HortonWorks’s Hadoop, MapR’s Hadoop, EMC-spinoff PIVOTAL, IBM InfoSphere BigInsights |
| WIRECLOUD Application Mashup GE | Mashup[28] platforms such as: iMashup, iGoogle, Apache Shindig, Apache Rave, Apache |

[22] Microsoft StreamInsight: http://technet.microsoft.com/de-de/library/ee362541(v=sql.111).aspx
[23] JBoss Drools: http://www.drools.org/
[24] SAP Event Stream Processor: http://scn.sap.com/community/developer-center/esp
[25] IBM AMT: http://www.research.ibm.com/haifa/dept/services/papers/amt_fact_sheet.pdf
[26] Software AG Apama: http://www.softwareag.com/corporate/products/bigdata/apama_analytics/overview/
[27] Esper: http://esper.codehaus.org/
[28] Mashup - http://en.wikipedia.org/wiki/Mashup_(web_application_hybrid)

| | |
|---|---|
| | Wookie & Cordova |
| ORION Context Broker GE | Message Brokers - Distributed publish-subscribe Messaging System[29] such as:<br>Redis, RabbitMQ, Apache Kafka, Apache ActiveMQ, and Kestrel |
| PROTON Complex Event Processing GE | Event Processing Software[30] such as:<br>Oracle Event Processing, Tibco Streambase, Esper, Drools and IBM Infosphere |
| OBJECT STORAGE GE | Cloud storage products/services based on CDMI such as those listed in [31] |
| KEYROCK Identity Management GE | Different implementation of the OAuth2 standard[32] |
| IDAS (Backend) Device Management GE | IoT Device Management products such as:<br>Oracle's Internet of Things platform, Axeda Ready M2M, Device Cloud by Etherion and Wind River® Intelligent Device Platform XT |

**Table 4 - FIWARE GEs vs. Alternative Products/Services**

### 1.5.2 Advantages of using FIWARE GEs compared to selected "Alternatives"

For filling in the following table we have used an approach which is based on defining a series of advantages that a user may consider appealing when choosing to go for a FIWARE GE instead of an alternative product/service available on the market (either Open Source or COTS).

[29] Apache Kafka - http://www.infoq.com/articles/apache-kafka
Exploring Message Brokers: RabbitMQ, Kafka, ActiveMQ, and Kestrel - http://java.dzone.com/articles/exploring-message-brokers
RabbitMQ vs Kafka - http://www.quora.com/RabbitMQ/RabbitMQ-vs-Kafka-which-one-for-durable-messaging-with-good-query-features

[30] An Overview of Event Processing Software - http://www.complexevents.com/2014/08/25/an-overview-of-event-processing-software/

[31] CDMI Server Implementations - http://www.snia.org/technology-communities/cloud-storage-initiative/snia-cloud-technology-community/list-cdmi-server-imp

[32] OAuth open standard to authorization - http://en.wikipedia.org/wiki/OAuth

| ADVANTAGES \ GENERIC ENABLERS | Cloud GEs | COSMOS Big Data Analysis GE | WIRECLOUD Application Mashup GE | ORION Context Broker GE | PROTON Complex Event Processing GE | OBJECT STORAGE GE | KEYROCK Identity Management GE | IDAS (Backend) Device Management GE |
|---|---|---|---|---|---|---|---|---|
| free support (contact person available via email) | ● | ● | ● | ● | ● | ● | ● | ● |
| no hardware to be purchased | ● | ● | | | | | | |
| free usage (no costs per hour / no licensing costs) | ● | ● | ● | ● | ● | ● | ● | ● |
| cloud based (available as SaaS) + no hardware to be purchased + no installation activities to be carried forward (GE Global instance) | | ● | | ● | | ● | | ● |
| openness: no vendor lock in | | ● | ● | ● | ● | ● | ● | ● |
| no need of extreme hardware configuration | | ● | | | | ● | | ● |
| comes with native integration with other GEs, thus being part of an entire ecosystem | | ● | ● | ● | ● | | ● | ● |
| cloud based: installable from image available on a catalogue (coming with a default configuration) on a cloud-based VM (GE private instance) | | | ● | ● | ● | | | |
| natively integrated in the FIWARE cloud infrastructure | | | | | | ● | ● | |

**Table 5 - "Advantages" of using FIWARE GEs**

From the above table, a couple of considerations can be made:

- the key factor of choosing a GE would definitely be the availability of a contact person (the "GE owner") who can eventually help a user in sorting out issues during both set-up and operation;
- set-up an entire infrastructure would not imply costs of hardware, hosting and licensing in the context of FI-PPP programme;
- most of the GEs are natively integrated each with the other thus representing a Future Internet "ecosystem" through which data can be acquired, stored, processed/analysed and, finally, exposed;
- GEs are based on open standards and so can easily work in conjunction with other products both open source and COTS.

## 1.6 WP5 Comparison of GEs to Alternative Technologies

### 1.6.1 Data Handling GE

A comparison was undertaken between the Data Handling GE and the OpenPDS [deMontjoye] [33].

Functional suitability

For the COS, we tried to find existing solutions that help ensure privacy of Personally Identifiable Information (PII) stored or handled by the system. The scientific background of the Data Handling GE is sound. Some published work explains the theory behind the GE such as in [deMontjoye]. OpenPDS is a personal metadata management framework that allows individuals to collect, store, and give fine-grained access to their metadata to third parties. OpenPDS is also a research prototype. No commercial product appears to offer

[33] de Montjoye Y-A, Shmueli E, Wang SS, Pentland AS (2014) openPDS: Protecting the Privacy of Metadata through SafeAnswers. PLoS ONE 9(7): e98790. doi: 10.1371/journal.pone.0098790

required functions. In terms of the required privacy functions, Data Handling is better suited to the COS as PII is stored within the system and privacy obligations allow control and detailed notifications of third party access to the data. However, Data handling GE appears as a work in progress with several bugs such as linking individual sticky policies to their respective files as well as the lack of appropriate documentation on the supported functions.

Performance

Privacy protection using DH GE as well as OpenPDS is expected to introduce some effect on performance. DH GE involves event handling and processing of privacy obligations. DH GE is responsive and no performance issues are identified. The GE is hosted at the TSSG data centre. As described by [deMontjoye], OpenPDS introduces significant performance overhead due to storage of data at the data subjects location and processing of individual requests involving security and privacy mechanisms.

Compatibility

Both DH GE and OpenPDS provide RESTful interfaces allowing integration with other systems and components regardless of their implementation language. DH GE is more compatible to COS given the requirements and system constraints such as internal storage of data.

Usability

DH GE has a simple Web-based user interface. However, interactions with COS are performed through the GE's REST API. Similarly, OpenPDS provides a Web interface for users. Researchers performed user evaluation of the tool and received good results regarding usability [deMontjoye]. OpenPDS provides better documentation.

Reliability

Even considering the current maturity level of DH GE, it would be expected that the GE provide available functions with no major bugs. However, reliability problems are frequent such as returned errors from the API in the form of Java and Hibernate stack traces which is inappropriate practice for a Web based API. Other errors exist in the documentation of the GE such as specifying wrong parameters. OpenPDS is also under development with several privacy functions are promised for the future. However, field studies and user feedback indicate better reliability for OpenPDS. OpenPDS's usage documentation with moderate detail is available on GitHub.

Security

The function of DH GE is to enforce the access control rules imposed by the data owner and execute the obligations on the usage of the data. The focus of OpenPDS is more on data protection than privacy obligations. OpenPDS can be subject to a number of attacks such as collusion between client apps and vulnerabilities caused by reliance on personal metadata.

Maintainability

Both tools exhibit good testability and modularity. DH GE provides better reusability.

Portability

DH GE runs in a Web server. Therefore, it is flexible in terms of the installation environment and usage scenarios.

### 1.6.2 DB Anonymizer GE

A comparison was undertaken between the DB Anonymizer GE and the ARX Data Anonymization Tool (http://arx.deidentifier.org/).

Functional suitability

DB Anonymizer GE facilitates protection of privacy during data disclosure through data anonymisation and helps to improve anonymisation policies. It receives a raw dataset together with the disclosure policy. It then analyses the policy and evaluates its effectiveness in ensuring privacy protection. At the end of the evaluation process, it returns a percentage of the original dataset that an attacker can reconstruct. The higher percentage the more possible to reconstruct the anonymized data and hence, it is necessary re-evaluate the anonymisation policy with more restrictions.

Despite its name, the GE does not currently perform anonymisation of data. However, the GE developers reported that they are working on adding this function in future versions. Therefore, the name of the GE is a bit misleading. Although, the GE does not include performing anonymisation of data, it fits the COS requirement for evaluating anonymisation policies.

ARX provides a multifunctional anonymisation tool. It implements several anonymisation techniques including k-Anonymity, l-Diversity, t-Closeness and d-Presence. It also supports added features, such as generalization hierarchies (structured grouping of entities sharing common attributes), exploration of the solution space and techniques to compare transformed datasets to original dataset. Therefore, ARX wins in terms of the functional suitability.

Performance
During the evaluation of this GE, it requires an unanonymised dataset to provide as an input and a specific anonymisation policy which will be used as dataset disclosure policy. The unanonymised dataset will be any SQL database and anonymisation policy will include as a XML policy file. These two inputs are necessary to evaluate the effectiveness of the dataset disclosure policy. This might cause performance and scalability problems in case of big data sets.

ARX uses optimized search algorithms and provides benchmarks on the performance and efficiency of those algorithms. It was noticed that ARX becomes unresponsive for a time when importing datasets.

Compatibility
DB Anonymizer provides RESTful interfaces allowing integration with other systems and components regardless of their implementation language.
ARX project does not only include GUI based application, it also offers free Java library that provides data anonymisation functionalities to other software systems. As ARX is implemented using Java, it supports multiple platforms including Windows, OSX and Linux. ARX provides Java API documentation for developers. ARX also supports multiple formats of datasets including CSV, XLS and SQL formats.

Usability
GE developers provide moderate level of documentation on the supported functionality and integration. However, discussion is provided regarding the scientific basis of the anonymisation evaluation and scoring except a published article[34]. This may cause uncertainty regarding the validity and reliability of the resulting scores without feedback from the security community. Possibly, additional measures are required to further ensure anonymity.

The online service (https://dbanon.lab.fi-ware.eu/) allows uploading datasets and policies for anonymity evaluation. The services frequently experience errors and unavailability problems.
ARX project provides elaborate documentation on API, usage and anonymisation technologies. The GUI application is user friendly and easy to learn.

Reliability
ARX tool wins on reliability as current version of DB Anonymizer features faults and errors. GE developers are dealing with those issues. Regarding reliability of anonymisation results, ARX also wins due to usage of multiple well-known technologies and high quality research publications in Journals and conferences that detail the scientific background of ARX.

Security

[34] Trabelsi, S, Salzgeber, V, Bezzi, M, Montagnon, G, (2009) "Data disclosure risk evaluation," 4th International Conference on Risks and Security of Internet and Systems (CRiSIS), pp.35-72, DOI:10.1109/CRISIS.2009.5411979

The functions of DB Anonymizer and ARX are to support privacy and confidentiality of sensitive data. ARX supports more privacy features and uses proven anonymisation methods.

Maintainability
DB Anonymizer GE currently provides single function using an anonymisation scoring algorithm and RESTful interface. No evidence of modularity. GE developers provide support for testing and evaluating the GE.

ARX project appears well maintained with source and documentation available online. ARX developers provide datasets and guidelines to assure evaluators of reproducibility of their results. The tool appears well structured and modularised.

Portability
DB Anonymizer GE runs in a Web server and provides REST API. Therefore, it is flexible in terms of the installation environment and usage scenarios.

ARX implementation in Java allows portability and flexibility regarding the installation environment. In addition, ARX can be used a Java library integrated into software systems.

### 1.6.3 Identity Management GCP

A comparison was undertaken between the main feature of interest within this GE (OAuth implementation) and the Spring Security OAuth extension.

Functional suitability
The Identity Management GCP GE provides customer administration as well as identity management, authentication and authorization services. The GE implements security standards such as OAuth and OpenId. The GE has been selected to provide authentication to clients who require to access data in the COS DSE system. The IdM GCP GE is used in the Charging Optimisation System as a means for authenticating its Web API users. As the API provides access to data stored within the COS, security is a major factor and as such it is of upmost importance that persons allowed to access this data can be verified in a controlled and reliable manner. The IdM GCP GE provides suitable documentation regarding the integration and usage of the GE. The GE supports other features that are not being considered at this stage by COS DSE such as customer management.

Spring Security includes OAuth support for providers and consumers. The library also supports a range of other authentication and authorisation mechanisms and features that can be of use to the COS system in future versions.

Performance
IdM GCP GE is only available as an online service which could have scalability implications. Current business usage scenario of the service is aimed towards authenticating limited number of SME clients.

Spring Security has no performance issues. Spring provides detailed documentations on guidelines and best practices in implementing scalable systems using Spring Framework components including Spring Security.

Compatibility
IdM GCP GE provides RESTful API allowing interoperability with systems built in any programming language.
Spring Security provides security mechanisms for J2EE-based enterprise software applications. It particularly supports projects built using Spring Framework. This limits its usefulness for the COS system as it is not built in Java.
Therefore, IdM GCP GE scores higher than Spring Security on compatibility.

Usability
Developers of the GE provided technical support for the integration and created admin access to COS DSE developers in order to configure the IdM GCP online service. Its

documentation is of good quality and comprehensively tackles all functionality that is provided. The documentation also covers most if not all errors which could be experienced during interaction with the GE. Another reason why no major issues were experienced is because all functionality that is described in the documentation does actually work correctly in the implementation.

Spring Security is mature framework with excellent supporting documentation and learning material. It has community support as well as formal approach to reporting and tracking bugs and enhancement requests. However, Spring only provides commercial support.

Reliability

Both are mature components with high reliability. No faults or errors are experienced with either.

Security

IdM GCP GE supports authentication and authorization of customers using common security protocols including OAuth, OpenID and SAML. As an online service with no access to the internal structure or code there needs to be an element of trust regarding its security. Spring Security supports a range of authentication and authorization mechanisms and features. Spring Security also supports testing using common techniques and tools such as integration testing using Junit. The framework is widely tested and evaluated by a large user community. This can be an advantage in detecting and reporting potential vulnerabilities in the security mechanisms.

Maintainability

IdM GCP GE is an online service with no view of the modularity of its internal structure. However, the online administration interface is well structured and allows enabling and disabling features as required.
Spring Security is highly modular and it forms part of Spring Framework known for its modularity and adherence to best practices and patterns in software development. As noted above, Spring Security supports unit and integration testing using common testing tools such as JUnit. In terms of reusability it is limited to applications using Java and Spring Framework.

Portability

As an external service, IdM GCP GE has no limitation as to installation or usage environment. Spring Security can be installed and used in different platforms and environments as part of Spring Framework. However, there are often issues regarding upgrade between versions of Security and other Spring components requiring matching between compatible versions.

## 1.6.4 Content Based Security

A comparison was undertaken between the CBS GE and an open source alternative FileSender (https://www.assembla.com/spaces/file_sender/wiki).

Functional suitability

Content-Based Security (CBS) GE protects data and its metadata at its source and integrates access control to the data. The data is protected by encrypting or signing at the time of its generation. Access to the encrypted data is controlled by restricting access to the cryptographic keys needed to remove protection from the data. This function is required by COS to ensure secure communication of sensitive grid emergency data.
FileSender allows secure transfer of files from source to destination through authentication of sender and receiver of data. FileSender does not currently offer file encryption. Encryption of the transferred file can be performed separately using a tools such as AESCrypt. However, there is no secure key exchange mechanisms offered by AESCrypt or FileSender. Therefore, CBS GE better matches COS requirements.

Performance

No performance issues are identified in either of the tools in terms of time or capacity utilisation. The GE requires limited memory and disk space. Resource consumption highly depends on the load i.e. number of concurrent requests.

Compatibility
CBS GE uses standard encryption algorithms. It has dependency on Access Control GE and IdM GE. The GE provides RESTful API allowing interaction with other systems with restrictions. It also runs in the multiplatform Apache Tomcat Web server.
FileSender server requires Linux and can be installed in a Web server e.g. Apache2. It supports standard SAML based authentication. It supports multiple standards such as HTML5, LDAP and RADIUS.

Usability
CBS GE provides documentation on the installation as well as on the architecture of the GE. It also provides description of unit testing of the GE security features. The documentation quality and level of detail can possibly be improved. GE developers do respond to enquiries regarding its installation and usage.

FileSender on the other hand provides detailed user and developer documentation, server and client requirements, known issues, mailinglists, etc. It is also easy to learn and operate. However, this may also be attributed to the fact that it provides less security features compared to the CBS GE. FileSender also supports multiple international languages.

Reliability
It is hard to claim high level of maturity of the GE given the available documentation and usage experience. We also experienced errors during decryption of the data. The dependency on Access Control GE and IdM GE may reduce fault tolerance and add complexity to its usage and exposure to errors in theses GEs.
FileSender is more mature. It provides detailed information and updates on current and upcoming releases. It is deployed at several academic institutions such as Waterford Institute of Technology. No errors were experienced during the usage of FileSender.

Security
CBS GE aims to ensure access control and protection of sensitive data through encryption and digital signature. This helps protect confidentiality and integrity of data. Authentication relies on the Access Control GE. Assurance regarding those security functions requires wider user community and feedback regarding potential vulnerabilities.
FileSender supports exchange of files in a moderately secure way. Support of cryptography is necessary in order to enhance its level of security.

Maintainability
CBS GE provides details on testing of its features. Testing is vital to assuring the effectiveness and robustness of the GE security functions. It can be reused in multiple scenarios and different environments. Its dependence on specific components i.e. GEs, may hamper its reuse. CBS consists of multiple modules i.e. consumer, producer and broker. However, further component modularity cannot be confirmed.

FileSender is currently well supported with testing and usage in multiple production environments. Information and guidance on various testing procedures are available on the Website. FileSender uses Hudson continuous integration tool for automated building and testing. It supports several functional features but no description available regarding its level of modularity.

Portability
As with some of the other GEs, CBS GE runs in a Java Web server and is flexible in terms of the installation environment and usage scenarios. No major issues regarding replaceability and adaptability are identified.

FileSender client supports multiple browsers and the server can be installed on any Web server environment preferably Apache2. FileSender project provides upgrade instructions for moving to current releases.

# ANNEX 2 WP1 Trial Results

This Annex contains an edited version of the confidential D1.7 “Trial Results” report of WP1, omitting some chapters. The other Dx.7 reports of the other WPs are public.

## Executive Summary

The WP1 Trial Results is WP1’s seventh deliverable in the FINESCE project. The purpose of this deliverable is amongst other to describe and report WP1’s final trial results within the FINESCE project.

The deliverable is split into the following results sections: usage of Generic Enablers and FIWARE, Energy optimization, Simulations, and Other.

## 2.1 Introduction

The purpose of this deliverable is amongst other to describe and report WP1’s final trial results within the FINESCE project.

The scope of the trial has been to execute Demand Side Management and Demand Side Response tests with external buildings in, Malmö, Sweden. The solution should be capable of testing activities of an integrated approach of energy carriers in order to demonstrate Demand Side Management and Demand Side Response tests based on either price or energy mix (CO2) for both heat and electrical loads.

The desired outcomes are stated here below.

- Understanding of how Future Internet technologies can contribute to an efficient and robust Demand Side Management system
- Proof of concept and evaluation on solution which architecture is based on distributed energy management capability and centralized portfolio management capability
- Proof of concept regarding cost optimization on price signals for heat and electricity based on different business model(s)
- Increased knowledge on future potential for Demand Side Management and Demand Side Response as well as ideas on customer’s potential to act as balancing power
- Evaluation of the thermal load shifting potential by different heating systems, e.g. under floor heating and radiators, while leveraging the building’s thermal inertia
- Definition of a scale-up strategy for the trial, e.g. ability for other towns, regions or business sectors to use the results and functionality

All of the desired outcomes for WP1 have been met and documented in different FINESCE deliverables.

The deliverable is split into the following results sections: usage of Generic Enablers and FIWARE, Energy optimization, Simulations, and Other.

Different results and conclusions have also been reported in previous deliverables. See for example the below deliverables for more information.

- WP1 Analysis of Generic and Specific Enablers Integration (D1.4)
- WP1 Trial Demonstration (D1.5)
- WP1 Innovation and Business Report (D1.8)

### 2.1.1 Trial architecture

The below figure illustrates the architecture developed in WP1.

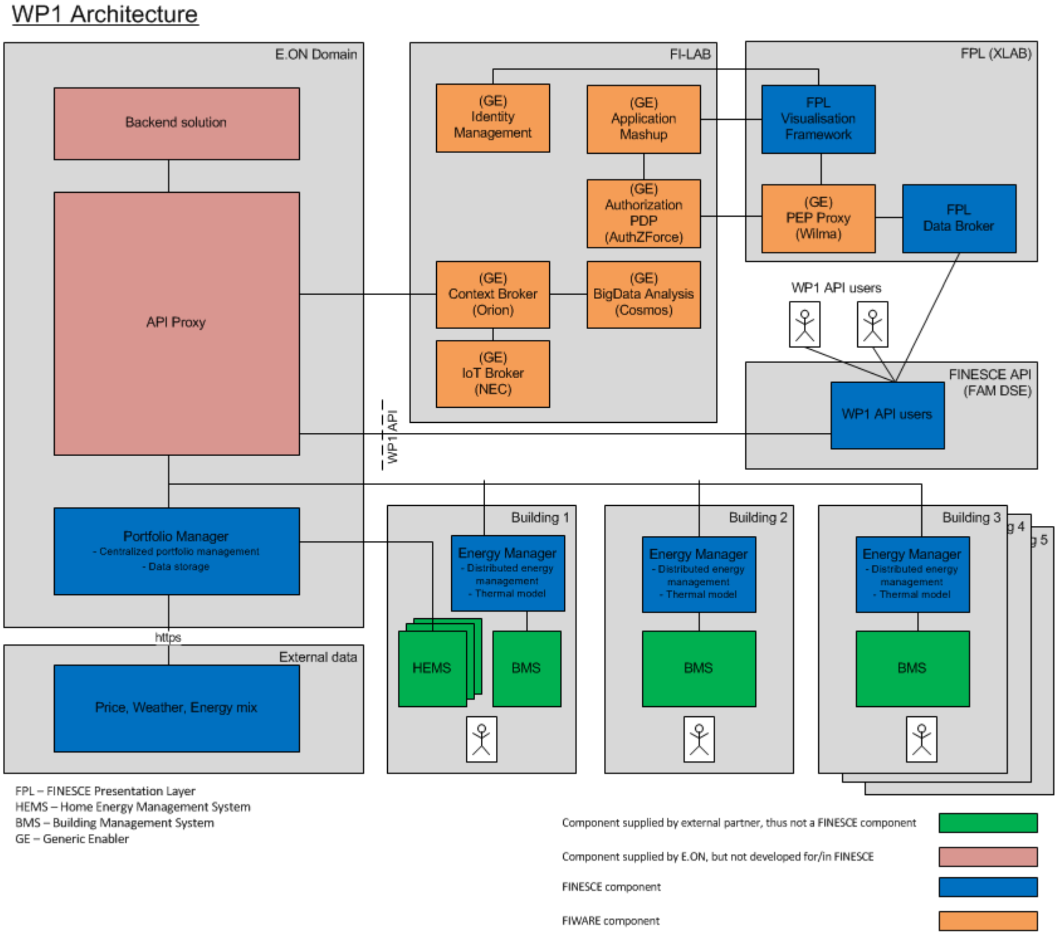


**Figure 5-1 WP1 trial architecture**

## 2.2 Trial results on usage of Generic Enablers and FIWARE

## 2.3 XLAB’s view on the FIWARE platform and Future Internet capabilities

The FIWARE proved highly beneficial in the trial development both in terms of the offered infrastructure and the GE’s readily available for the integration. From the perspective of XLAB as a SME in the consortium, this means less time and effort spent on technology, and more attention given to the actual content of the bigger Smart Energy projects.

Of the GE’s available and integrated into the trial, we found the greatest use in the ones from the Security chapter of the FIWARE Catalogue. The GE Identity Manager KeyRock has been serving as the Single Sign-On solution for the visualisation applications. It also serves as the go-to place for authenticating users and machine clients, who access the data handling and serving layer of the visualisation services in the trial. To implement the authorization functionality, we use the GE Authorization PDP AuthZForce, and to complement the two services, we also use the GE PEP Proxy Wilma. By deploying the services in a topology where the PEP Proxy serves as a gate keeper to the data service, we have a stack where data is protected from unauthorized accesses.

The benefit of the FIWARE’s whole package of the two of the three As (Authentication, Authorization, Accounting) is that the trials and their applications can take advantage of the accounts of the stakeholders already registered in the FIWARE. This means that both the Phase II and the Phase III projects’ participants can use a unified and safe approach in presenting their user credentials. At the same time, the GEs employ open standards (e.g., the XACML for the Authorization PDP), so using FIWARE does not represent any vendor lock-in.

Thanks to this open design, the solutions currently working in FIWARE could be adapted to work with other solutions mandated by the potential customers (e.g., a utility or a DSO) due to their strict administrative policies. Of course this customisation requires a minor amount of effort and some overhead.

### 2.3.1 Lessons learned with FIWARE and GE’s

Details concerning lessons learned with FIWARE and GE’s are documented in the below deliverables. Those deliverables include example of concrete feedback given to GE developers as well as a comparison between the Generic Enabler BigData Analysis (Cosmos) and a similar service (Hortonworks).

- WP1 Mid-term Analysis of Generic and Specific Enabler Integration and Trial Impact (D1.3.2)
- WP1 Analysis of Generic and Specific Enablers Integration (D1.4)

## 2.4 Trial results on energy optimization

The developed trial infrastructure has proven to be a very flexible system with regard to handling different use cases and business models. One of the infrastructure's many strengths is the ability to deliver benefits both on a local level, optimization in the building, and at the same time on a global level, system optimization.

The infrastructure can handle all types of energy carriers, and we have been controlling electrical loads, district heating loads and district cooling loads.

Currently five buildings are connected to the infrastructure. All originally identified use cases have been implemented, investigated and proven. These use cases include amongst others tests for load curtailment and load shifting. Even additional use cases have been added. New use cases are found as spinoff when developing and investigating, resulting in new future opportunities. See previous deliverables (D1.1 and D1.2) for more information concerning energy optimization.

Further, some of the use cases have been found to have a commercial potential and plans are made to progress towards a commercial phase, post the FINESCE project. This concerns for example system optimization of district heating and district cooling. As for the previously reported district cooling case in the Western harbour, Malmö, the current status is that a rollout of the WP1 infrastructure to circa 25 buildings in 2016-2017 (i.e. post the FINESCE project) would have a positive NPV compared to a "conventional" case. The "conventional" case includes investments in additional production capacity (MW). This rather expensive production capacity would not be required to the very same extend thanks to the load curtailment use case.

### 2.4.1 Results

The infrastructure is capable to shift load according to defined use cases. The potential for shifting loads without significant impact on the customer's comfort has been shown to be bigger than initially expected. All this is of course very positive as it indicates good opportunities for leveraging the loads' flexibility.

Therefore E.ON is now further exploring how the flexibility can be used to enable system optimization of district heating and district cooling grids. For example, provided that desired flexibility is available, that could enable avoidance of firing up peak production units which usually have higher operational costs and $CO_2$ emissions. Thus, rolling out the infrastructure to the wider Malmö (here 5 buildings would not be enough, 50+ are required) could enable benefits to the whole City of Malmö.

In order to quantify the potential, different analyses and simulations have been activated, in addition to planned FINESCE activities, to identify a potential return of investment given the costs to set up a commercial operation of the infrastructure and rollout of required technology (compared to today's pilot operation with 5 buildings). These promising aspects would not have been this far without E.ON's involvement in the FINESCE project.

#### 2.4.1.1 Dynamic district heating prices

E.ON has been testing so called dynamic district heating prices together with 3 of the buildings in the WP1 trial. The prices varied on hourly basis. The test is referring optimization of heating for the complete building, i.e. on building level, not necessarily individual apartment level.

There are many different possibilities to build dynamic price models for district heating with different advantages and drawbacks. The selected price model for this project was based on the Nord Pool electricity prices.

Moreover, Nord Pool is the Nordic electricity market, hence not at all district heating market. However, when it is cold outside, the electricity prices increase. The same concerns district heating. When it is cold outside, more district heating production is required. Hence there is a correlation between electricity prices and district heating production. Still there is not at all a

perfect correlation. This correlation is relevant over hours and days, however it is not as strong over longer periods of time, for example over months and years.
The tests of using dynamic district heating prices based on Nord Pool electricity prices have been successful in general, but also resulted in a number of learnings for the future.

As the building's heating system should not be turned off completely (to avoid physical stress in the building), a 50% maximum curtailment capacity was set. It means that when the BMS system in the building requires 60kW of heat power at a given moment, for example, the WP1 optimization algorithm reduced the heat consumption to 30kW during the that period. This 50% maximum curtailment cap limits as well the capacity of monetizing the price differences within the day by half.

The dynamic district heating price based on Nord Pool has periods with very low variations, and periods with higher variations. See below figure.

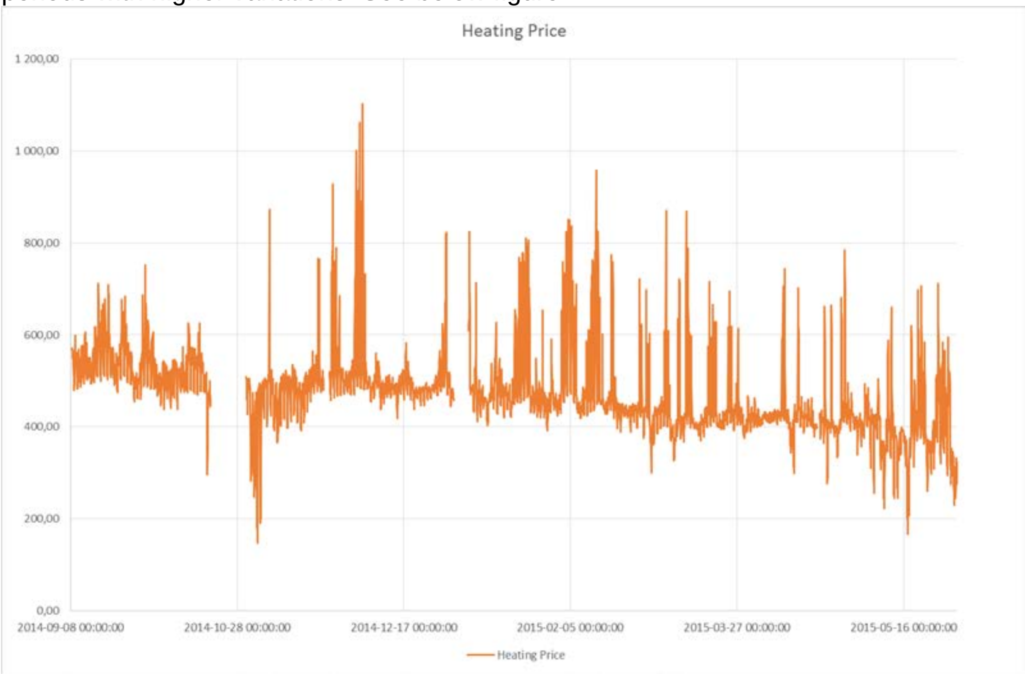


**Figure 5-2 Prices over 2014-2015**

It can be seen that there are periods of high variations and periods with low variations. Periods with high variations generally have peak price changes close to double the price of the lowest price of the day.

The uncertainty of these factors makes it very difficult to state an overall saving potential for the customer. However using the chosen price model, for a selected period of time, and estimated flexibility, calculations indicates savings around 5% (compared to doing no optimization at all).
The figures of 5% derive from a so called baselining process. Calculating the baseline is very complex. It is very difficult – if not impossible – to know for certain how the load behaviour would have looked like without optimization.

Moreover, it is also important to be aware of that some days have a very flat price profile, which indicates a low possible to lower the costs. On the other hand some days have a more extreme variation which indicates a high possible to lower the costs. That is one of the reasons why it is impossible to define an exact percentage concerning the reduction potential.

# 2.5 Trial results on simulations

## 2.5.1 Simulation

### 2.5.1.1 Introduction

In district heating grids, the conventional operation of heating plants is demand oriented. As a result, peak units that are generally characterized by a high cost and CO2-footprint, and specific generation costs are operated to ensure the security of supply in periods of large demand.

Residential and commercial buildings account for up to 38% of the total end energy consumption worldwide and hence provide a large potential for energy savings and Demand Side Management (DSM) [1]. The embedded thermal mass of a building may be actively used as structural thermal storage capacity. Therefore intelligent control strategies can be used to optimize the use of this capacity by taking into account the thermal characteristics of the building. This is realized by a dynamic control of the indoor temperature to flatten the buildings' heat demand profile maintaining or even improving thermal comfort. Such dynamic control strategy would preheat the building and activate the storage capacity by increasing the indoor temperature setpoint in times of low demand that correspond to low CO2 emissions.

Alternatively, the heating setpoint could be lowered in high load periods which induce the operation of peak units that result in high CO2 emissions. Consequently the thermal mass releases the stored energy thereby reducing the energy demand allowing for avoiding CO2 emissions.

The most crucial barrier of estimating the building inertia is the lack of knowledge about the building physical properties. A detailed investigation of these factors requires extensive monitoring and analysis which is, in practice, applicable only on a small fraction of the total existing building stock. The aim of this work is to develop a method for the identification of the building thermal flexibility, described by the parameters of a simplified building model with a clear physical interpretation. The parameter identification process involves optimization of the model fitting using a few input variables measured at the building and very basic information about the building. The determined thermal flexibility will describe the capability of the building to act as short term heat storage and will therefore represent its load shifting potential for district heating networks. The development of the method aims its applicability on different building types without significant adjustments.

### 2.5.1.2 Approach

This work is based on a systematic approach that will allow for a straightforward future application of the method on other buildings. The below figure gives an overview of the applied approach. First the input data measured at the building is filtered and completed in order to be usable for the model fitting. The model parameters are initialized and constrained based on basic building information combined with specifications from norms and standards. The estimation of the model parameters is performed by fitting the simulation output to the available measurement data. An optimization algorithm is used for the parameter approximation based on an interior point optimization method for solving linear and non-linear convex optimization problems.

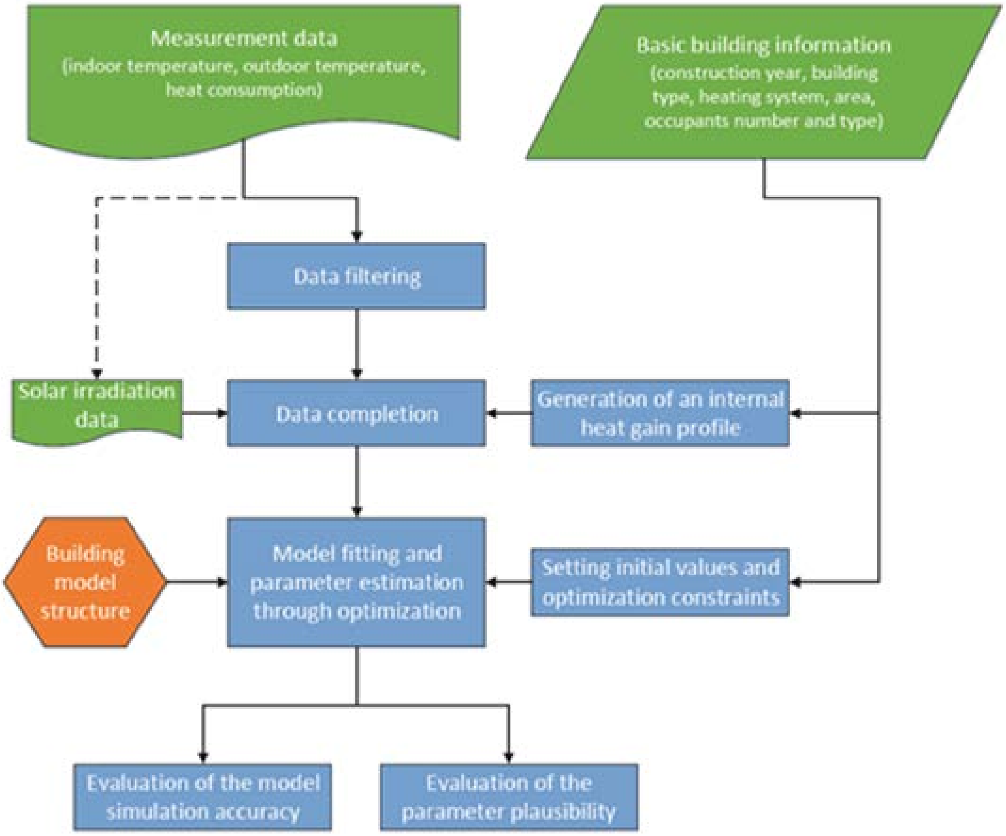


**Figure 5-3 Flow chart of the building modeling and input data estimation**

### 2.5.1.3 Models

The estimation of the model parameters through optimization of the model fitting is tested for different model structures. The six models tested in the study include different building components and vary regarding their complexity and accuracy to represent the real physical interrelationships in a building. The below table gives an overview of the presented building models and the physical effects they regard. The test parameter estimation approach on the different building models will reveal which ones are simple enough to be parametrised through fitting of measured data but complex enough to reproduce the building thermal response accurate enough for the use in DSM measures.

**Table 5-6 Overview on the building models considered in this work**

| Separate consideration of: | Building models | | | | | |
|---|---|---|---|---|---|---|
| | I | I-H | I-E | I-E-A | I-H-E | I-H-E-A |
| interior component | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| exterior component | ✗ | ✗ | ✓ | ✓ | ✓ | ✓ |
| heater component | ✗ | ✓ | ✗ | ✗ | ✓ | ✓ |
| indoor air temperature node | ✗ | ✗ | ✗ | ✓ | ✗ | ✓ |
| transmission heat losses | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| infiltration heat losses | ✗ | ✗ | ✓ | ✓ | ✓ | ✓ |
| radiative heat exchange between interior and exterior | ✗ | ✗ | ✓ | ✓ | ✓ | ✓ |
| convective heat exchange | ✗ | ✗ | ✗ | ✓ | ✗ | ✓ |
| convective and radiant contribution of the heat gains | ✗ | ✗ | ✗ | ✓ | ✗ | ✓ |
| heat gains on interior and exterior | ✗ | ✗ | ✓ | ✓ | ✓ | ✓ |

### 2.5.1.4 Case study

The presented approach was tested on a residential apartment building in Malmö, Sweden connected to a district heating grid operated by E.ON. The observed residential building includes 53 apartments and a total floor area of 4740 m². As shown in the below figure, the building integrates three blocks - two consisting of 5 floors and a third block with an additional penthouse floor. The construction year 2013 suggests a good insulated building envelope.

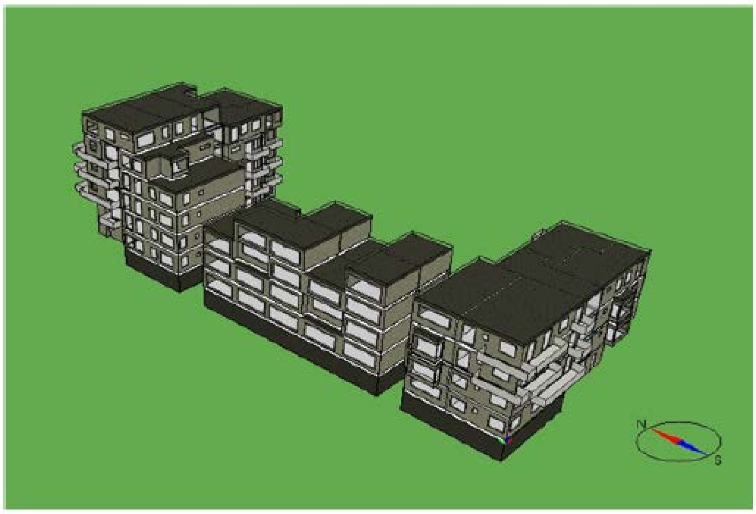

**Figure 5-4 of the test building in Malmö, Sweden**

The input data for the model simulation is presented by times series of measurements extracted from the building management system (BMS) and the thermostat system of the test building for the period between end of January and beginning of March 2014. The collected data includes hourly values of the indoor temperature (average value for the building), the heat consumption and the outdoor temperature. Additionally, values for the solar irradiation on horizontal surface for the area of Malmö could be obtained by the Swedish Meteorological and Hydrological Institute.

### 2.5.1.5 Results

The parameters of the simplified building models were identified using the monitored data of 14 consecutive days from the 14th to the 28th of February. Subsequently, the models were simulated for the whole available period from the 14th of February to the 3rd of March. The most suitable building model is identified by comparing the simulation output accuracy and the parameter plausibility of the different models.

**a. Quantitative analysis**

The quantitative accuracy of the model regarding their ability to represent the indoor temperature fluctuations are evaluated based on the root mean squared error (RMSE) of the residuals between the simulated temperature and the measured indoor temperature of the building. The results show that simulation error decreases slightly with the rising complexity of the models. Still, the absolute differences between the RMSE of every two models does not exceed 0,03 K - an insignificant value considering the sensitivity of the temperature sensors in the rooms.

**b. Qualitative analysis**

Considering the similar RMSE values of the different models, the qualitative analysis of the temperature prediction presents a better way to evaluate the accuracy of the model simulations. The models I-E-A, I-H-E and I-H-E-A give a better representation of the building cooling rate in the night hours, even if the average absolute residuals to the measured temperature do not differ significantly.

**c. Plausibility of the estimated parameters**

The below figure gives an overview over the estimated model parameters regarding their physical plausibility. Most of the models have several parameter assessed to their boundary values. Only the I-E-A-model provides a parameter set which is completely within the physically plausible range.

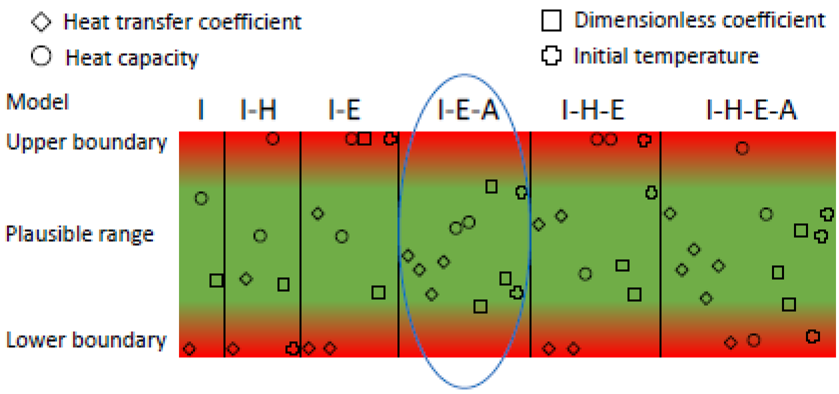


**Figure 5-5 Physical plausibility of the estimated model parameters**

In conclusion, the only building model that presents a whole parameter set with reasonable physical interpretation and reproduces accurately the indoor temperature dynamics is the I-E-A-model. This model is used further for the estimation of the thermal flexibility of the test building.

### 2.5.1.6 I-E-A-model

The below figure presents the model structure of the I-E-A model. All interior and exterior building components are summarized in one respective capacity. Additionally the indoor air is observed as massless temperature node. The model distinguishes between infiltration heat losses, connecting the indoor air directly to the environment, and transformation heat losses, which transfer the heat first to the exterior and then to the ambience.

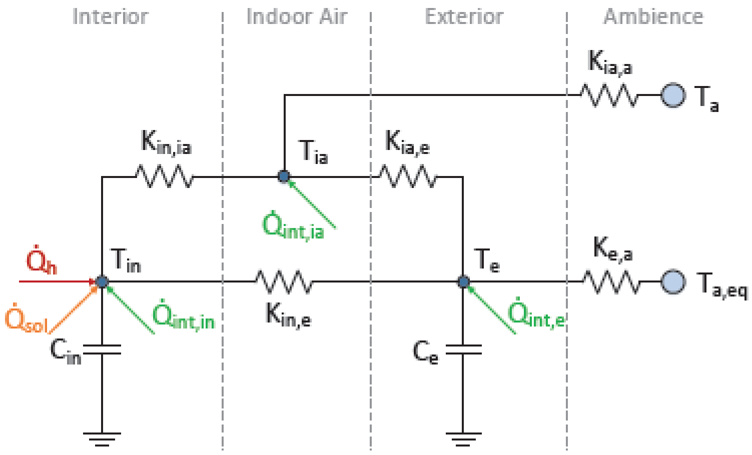


**Figure 5-6 Simplified building model: I-E-A**

As presented in the below figure, the model fitting gives a good match of simulated and measured temperatures. It must be noted that the recorded temperature drops are reproduced by the model simulation in a very accurate way.

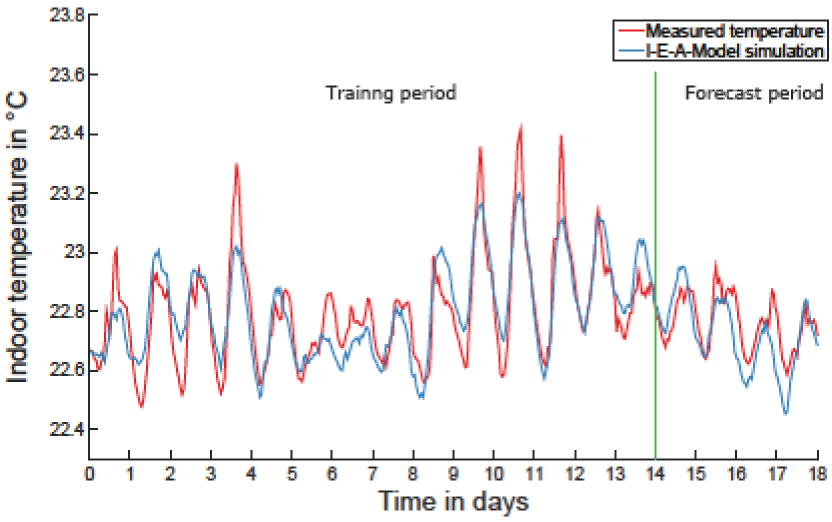


**Figure 5-7 Simulation results of the I-E-A model**

### 2.5.1.7 Thermal flexibility estimation

District heating network operators need information about the thermal flexibility of the heat consumers in order to plan demand side management measures. Therefore, the practical application of the developed method involves the derivation of the building thermal flexibility from the estimated model parameters. In this sense, the potential of a building for demand side management is defined by its thermal flexibility.

The below figure presents an estimation of the thermal flexibility of the case study building as a function of the permitted indoor temperature decrease and the heat load reduction performed using the I-E-A-model.

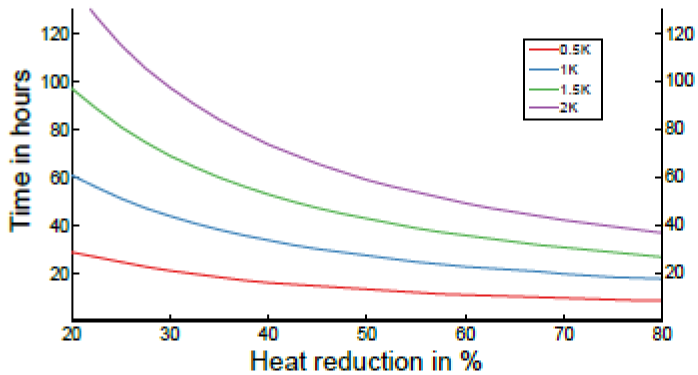


**Figure 5-8 Maximum time of building wall mass heat storage discharge for preset indoor temperature drops**

Each curve corresponds to the allowed temperature drop in the building. High offsets of the building heating reduce the influence of the various indoor temperature drop constraints on the maximum cool down times. For the observed building, the heat stored in the building mass is sufficient to keep the indoor temperature in a comfortable range (less than 1 K temperature drop) for 20 hours with a heating reduction of 70%. These results are observed for an average outdoor temperature of 3°C, present at the period of the simulated heat reduction.

A step test was performed by E.ON at the real building in Malmö to verify the calculated results. On the 12th of December at 1:00 am in the morning the heating system operation of the test building was reduced by 30% for 14 hours. As a result, the average indoor temperature of the building decreased by less than 0,2 K. This confirms the high thermal flexibility of the building assessed by the building model simulation.

#### 2.5.1.8 Conclusion

This work has presented a method for assessing the building thermal characteristics based on collected measurement data and basic building information. For this purpose, the building was represented through a lumped capacity model using a grey-box modelling approach. After a process of filtering and completion of the input data, an optimization algorithm was applied to fit the simplified building model to the available measurement data. The initial parameter values and the optimization constraints were derived from norm and standard specifications as well as basic principles of the building physics and were defined to be generally valid for different building types. The model parameters estimated by the optimization algorithm allow for the direct derivation of the building thermal flexibility, the building time constant and therewith the potential of the building for Demand Side Management measures. The wide applicability of the building model structure and the parameter constraints for the optimization, developed within the method, allows for an implementation on different buildings types. The qualitative analysis of the model temperature prediction and the evaluation of the estimated model parameters revealed that only the two capacity building model with an additional consideration of the indoor air as a massless node (I-E-A-model) combines an accurate qualitative reproduction of the indoor temperature fluctuations and a clear physical interpretation of the estimated parameters.

An evaluation of the building thermal flexibility using the I-E-A-model determined that the building can maintain a comfortable indoor temperature (less than 1 K temperature decrease) over 20 hours for an average outdoor temperature of 3°C and a heating reduction of 70%. The high thermal flexibility of the building was confirmed by step tests performed at the real building.

### 2.5.2 Modelling a CO2-steering signal for Demand Side Management in district heating grids

#### 2.5.2.1 Introduction

Usually the operation of combined heat and power plants for district heating grids follows the demand. Peak units are used to cover the demand in times of higher demand periods. In general those units have higher CO2 emissions and in particular are most cost intensive. Hence, Demand Side Management or Demand Side Response concepts for district heating grids can be sufficient solution for an optimized operation of the heating plants. Firing up the peak units could be avoided.

Such concepts could make use of the thermal flexibility on the building side given by the thermal inertia of the building or thermal water storages. The flexibility could shift the demand from times with high CO2 emissions into times of lower CO2 emission to support a more sustainable operation of the heating grid.

This work derives a CO2 steering signal based on generation data from the district heating plants in Malmö of the years 2011 and 2012. The CO2 steering signal is the same as the CO2 emissions prediction. This data plus hourly outdoor temperature and the corresponding CO2 emission factor per power plant is provided by the district heating grid department. The derived signal is applied to the customers to give incentives for shifting demand.

Both an artificial neural network (ANN) and a regression based method are applied to the data for modelling the CO2 signal.

#### 2.5.2.2 Modeling

First, the provided data is used to calculate the CO2 emissions for all heating plants and the total amount in every hour of the years 2011 and 2012 according to the outdoor temperature. We filter the data according to the outdoor temperature with a resolution of 0,5°C and the corresponding hour of the day. This results in information of the CO2 emissions for each hour of a characteristic day depending on a certain outdoor temperature. For taking into account inaccurate information an exponential robust fitting is applied to each hour. This results in 24 different fitting curves offering the opportunity to calculate the emissions at a certain hour depending on the temperature during this hour h.

$E_{CO2}(h)=a(h)\cdot e^{b\cdot T_{out}(h)}+c(h)$ Eq. 1

Eq. 1 applied to each hour of a day is then used to derive a 24hour CO2-emission forecast implied that the outdoor temperature forecast is available. Due to the fact that if the daily outdoor temperature is higher than 15°C there is no heating required within the building we set the CO2 emissions to zero during those time periods. An example for the robust fitting of one hour in the one-day model is pictured in the below figure.

For covering the disadvantage that those days can appear also in winter and transition periods, we extend the model to a three-season model. This represents, using the VDI 4655, the different weather periods, such as winter, transition and summer period. Studies showed that a linear fitting approach, given in equation 2 fits best here.

$E_{CO2}(d_s,h)=a(d_s,h)\cdot T_{out}(d_s,h)+c(d_s,h)$ Eq. 2

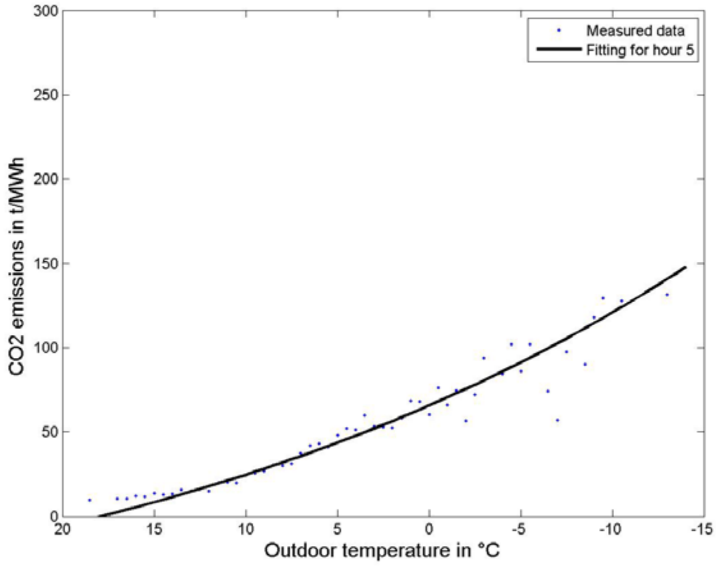


**Figure 5-9 Example of the robust fitting approach for hour 5 using the one-representative day model**

The below figure shows an overview on the work flow for the modelling of the three-day representative model. The distinction between these periods is based on the average day temperature. If the average temperature is less than 5°C it is considered to be a winter day, whereas an average temperature higher than 15°C refers to summer period. Other average temperatures lead to transition period.

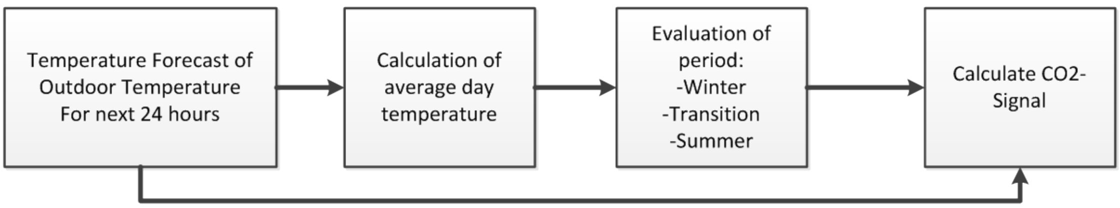


**Figure 5-10 Work Flow of CO2-Signal Modelling**

### 2.5.2.3 Results

For the evaluation of the model both the one representative and the three-seasonal model are analysed. The fitting results for the measurements of two years data are considered. In addition, we also apply a nonlinear autoregressive network with exogenous inputs (NARX) artificial neuronal network model. It contains 20 neurons and 2 layers. The data of 2011 and 2012 is used to forecast the CO2 emissions of the year 2013.

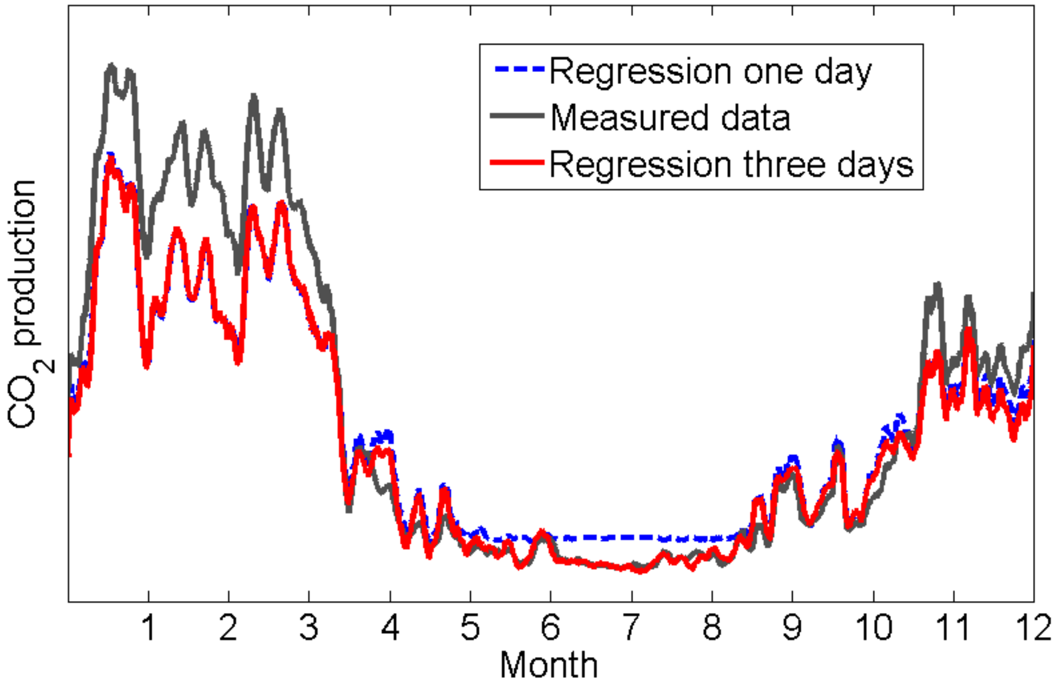


**Figure 5-11 Comparison of results taken from**

As evaluation criteria we mainly focus on the sum of squared error (SSE). The analysis shows that the ANN model has the biggest SSE and is hence neglected for further comparison. Looking into the two other models shows that the three-seasonal model has a slightly smaller SSE than the one-day model and hence gives a better performance. In general both models show that the dynamic behaviour of the CO2-emission profile can be captured by both approaches. The below table provides an overview on the SSE per model. In general a lower SSE is related to a higher performance.

**Table 5-7 Overview on SSE per Model Approach**

| Model Approach | Sum of Squared Error (SSE) * $10^5$ |
|---|---|
| Artificial Neuronal Network | 5.00 |
| One-day Approach | 3.98 |
| Three-day Approach (VDI 4655) | 3.75 |

### 2.5.2.4 Conclusion

This analysis showed that the model trained with historic production from two years can be used to evaluate a CO2 signal. This signal is able to be used for DSM concepts in district heating grids to give incentives to move consumption into times of lower CO2 emissions. In general the three-seasonal model according to VDI 4655 should be preferred due to a lower SSE meaning a better performance. However, further training data is desired to reduce the SSE.

## 2.6 Other Trial Results

Results linked to WP1 customer engagement and partner interaction was presented in deliverable WP1 Trial Demonstration (D1.5).

### 2.6.1 Visualization

The Demand Side Management trial involves the tenants of the buildings and the staff working in the commercial buildings. The amount and the dynamics of heating directly affects the comfort levels of the people. At the same time, the consumption of energy is associated with cost, which the consumers wish to minimise. One of the better ways of controlling the energy consumption is by monitoring it in a Smart Energy graphical application.

In FINESCE, XLAB used the FINESCE Presentation Layer (FPL), developed in the WP3, to create an application, which visualises the energy consumption of the smart building. The architecture of the FPL is presented in a greater detail in D3.7, but in short, it is composed of the Data Broker for collecting and enriching the Smart Energy data, a visualisation framework and visualisation widgets. Using these components, end-user applications are created. The aim in the design of these applications was to present the data to the users with a pleasing appearance, but at the same time also to offer crucial visual information on the current and past energy consumption.

The application is suitable for a number of use cases. Principally it is aimed at the owners of the building, and the operators of the heating and electricity service, such as E.ON. However, it is also possible to extend its use to include individual tenants in the use cases where the energy consumption meters are installed at the individual consumer’s apartment.

A single installation of the application supports all three use cases. The users will see the data and functionality suitable for their assigned role. The role system also enables various granularities of the data accessible to the users, protecting the data from the unauthorised views and operations.

The design of the application follows the principles of the responsive interfaces. This assures that the visualisation is suitable for a wide range of displays, starting at the largest panels used in the operating centre to the small screens of the smart phones. The choice of the underlying technology also enables portability of the application, and assures compatibility with the great majority of modern and popular web browsers.

The FPL is also designed to respond quickly to the requests, making sure that the users do not perceive the delay between the requesting a view and receiving it on screen. This includes the views, which display the Smart Energy data on a wider time scale. The quick responses are possible because FPL’s Data Broker aggregates the data at various levels and stores the aggregations to be quickly displayed.

#### 2.6.1.1 Energy Service Providers

The first use case of the Smart Building application assumes that the energy providers or the DSO’s need to monitor the ongoing energy consumption by their customers. For these users, the administrator must have created an account, and has also created the users in the application, assigning then the role of the providers. When a user with this role logs into the application, they first receive the Overview view. The below figure shows an example of the view, which contains the following elements:

- the latest readings of the aggregated electricity and heat consumption,
- a chart showing the history of the energy consumption readings starting at the beginning of the day, both for the heating and the electricity as the energy source,
- the history of the outside temperature readings from the beginning of the day,
- the history of the energy prices, both for electricity and heating.

The purpose of this view is for the operator to have it always open and visible, thus being able to perform continuous monitoring of the network status. By showing the live data and the chart that periodically updates to display the history between the beginning of the current day and the current time, it enables a view that requires no input parameters.

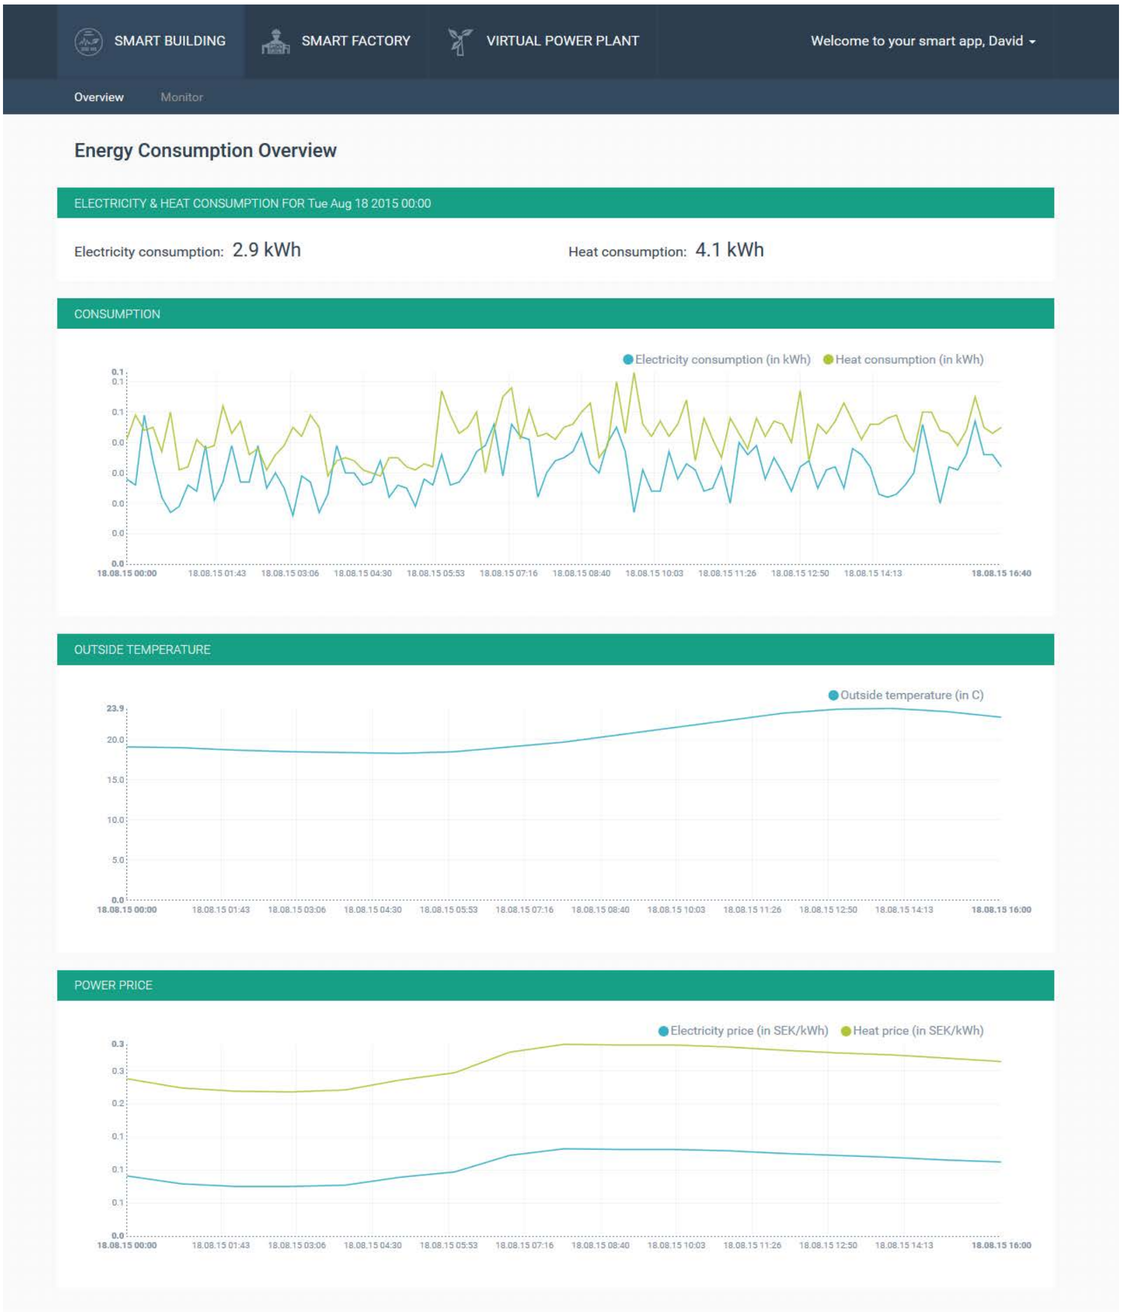


**Figure 5-12 The Smart Building application's Overview displays the data relevant for continuous monitoring of the Smart Factory's energy consumption**

To access a view where the user has much more control over the range of the data shown, the user can switch to the Monitor view. There, a provider user obtains a list of all the meters registered in the system, with each meter placed in a building and a region. The user can then select from this list a region, a house or a meter, and the widgets on the rest of the view will show the data applicable for the selection. Additionally, the user can adjust a start time and an end time of the range of data used in this view. The data shown on the visualisation widgets will therefore represent a relevant aggregation of the data if the user selects a building or a region. . The Figure 5-13 shows an example of this view.

The data display at the Monitoring view includes the following:

- a chart showing the history of the heating and the electricity power consumption

- an aggregation of the heating and electricity energy consumed in the selected time range
- a chart showing the history of cost of the energy consumed
- a chart showing the history of prices of the energy as it has been readjusted in time
- a chart showing the history of the outside temperature

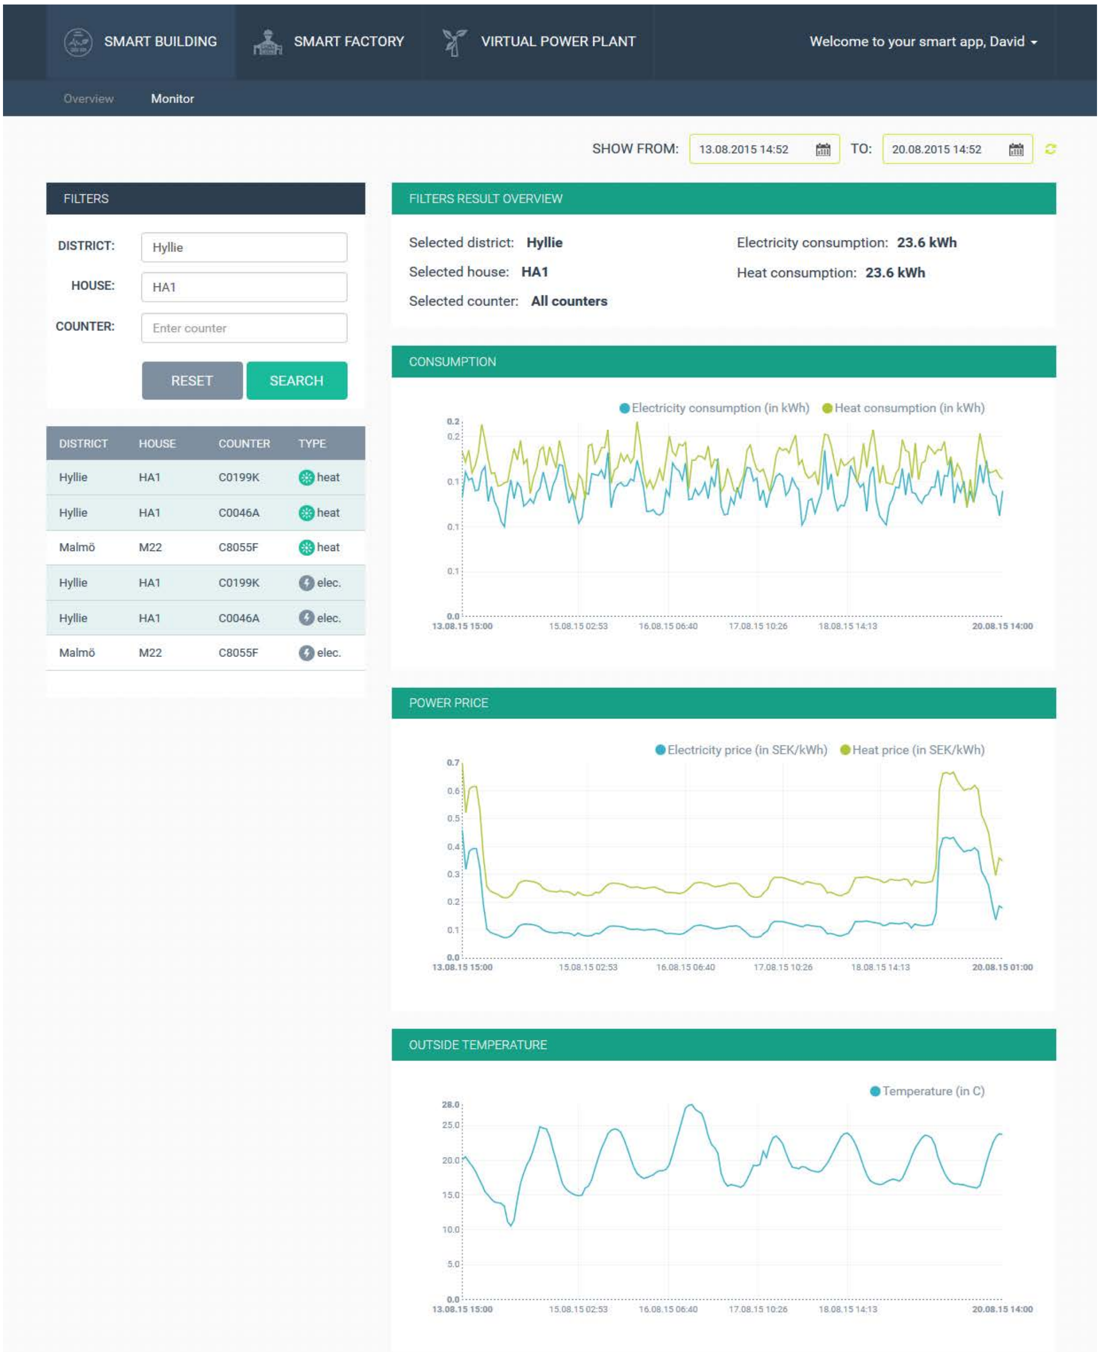


**Figure 5-13 The Smart Building application's Monitor view for the energy provider's operator shows data at various levels of aggregation**

### 2.6.1.2 Customers and consumers

The second major role in the application represents the recipients of the energy, who are the direct customers of the energy provider (e.g., the building owners) or the end consumers (e.g., tenants in the house or the building). We assume that they purchase the energy and pay at a monthly interval for the energy metered and consumed in the previous month. The prices are set by the market and possibly the provider, and they can be flat or can dynamically change throughout the day. The consumer's interest in the metered data is therefore the history of the energy consumption and the cost accumulated so far.

Following this rationale, the users upon logging in receive the **Overview** view (Figure 5-14), which visually summarises the following:

- live power currently consumed according to the most recent reading
- a chart showing the electricity and heating power consumption history since the beginning of this month (or since the start of the accounting month)
- accumulation of the energy consumed since the start of the month
- the cost of the energy consumed
- the grade of the energy consumption efficiency in this month as compared to the metered consumption in the previous month
- a chart showing the outside temperature since the beginning of the month
- the current weather

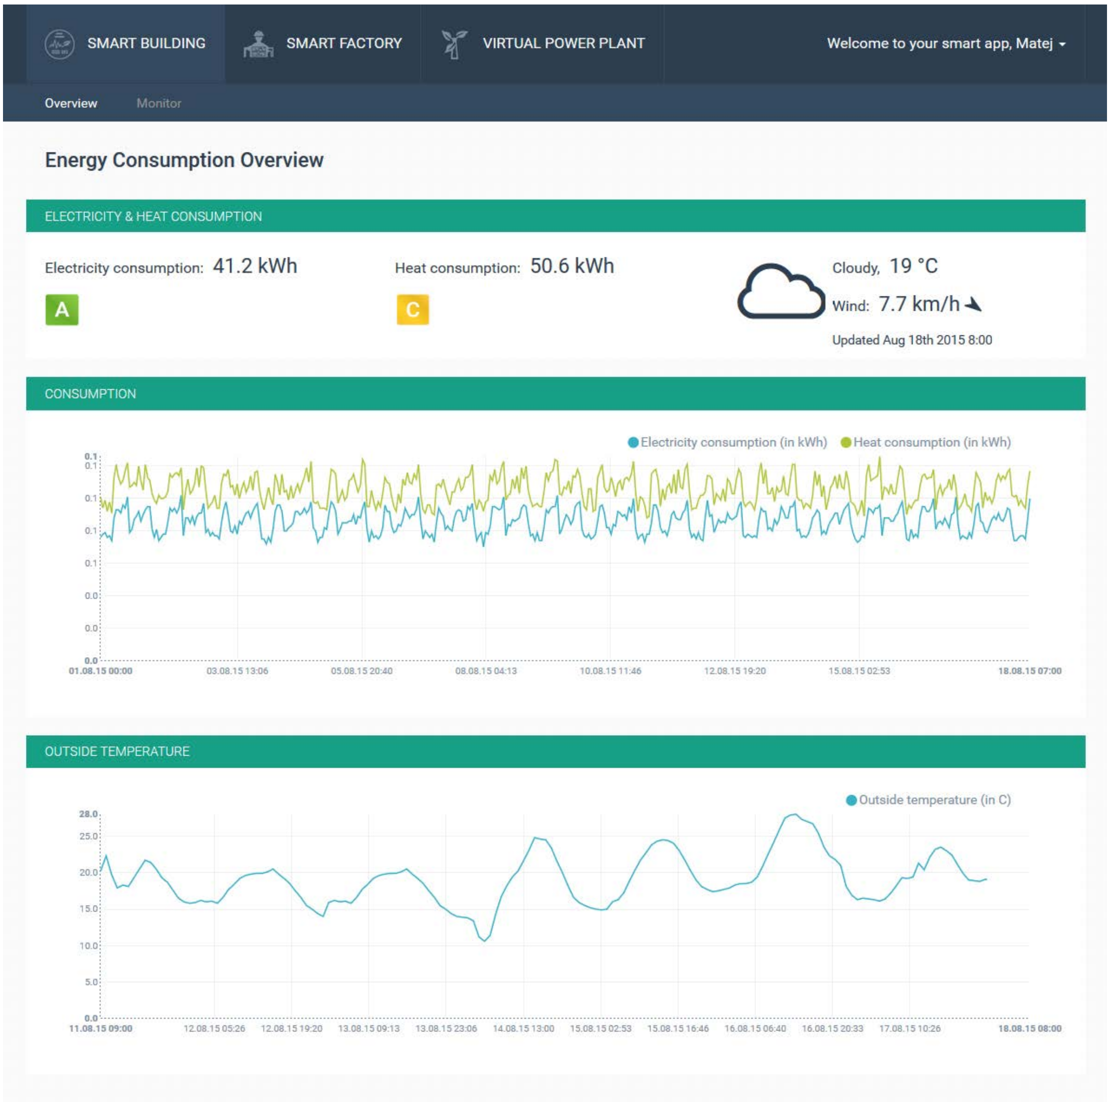


**Figure 5-14 The Smart Building application's Overview tab for the customers and consumers displays the data relevant for monitoring the consumption in the current**

The consumers may also switch to the Monitor view for the ability to select a custom time range for the visualised data. The Figure 5-15 shows an example, which consists of the following widgets:

- a chart showing the history of the heating and electricity power consumed during the selected time
- a chart showing the history of the cost of the energy consumed in the selected time range

- a chart showing the history of the outside temperature in the selected time range,
- a weather history

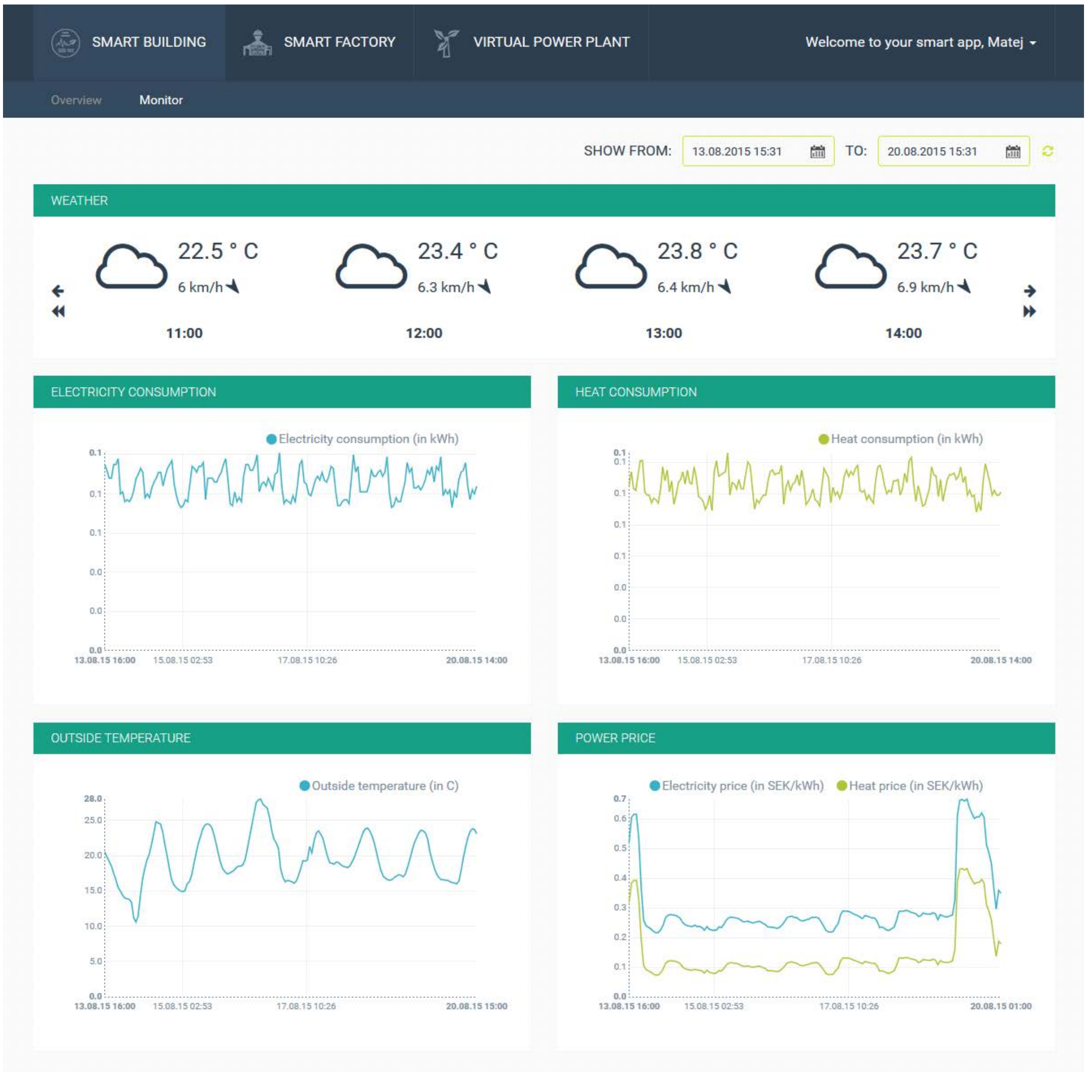


**Figure 5-15 The Monitor view for a consumer user shows the data related to a single building or apartment**

#### 2.6.1.3 Visualization security

The heat consumption data of a building are considered sensitive data. When produced by tenants, the data is private, because it could be used to discern a deeper insight into the habits of the consumers. For the whole buildings, they could represent a business secret. In either case, we wanted to build into the application the ability for the users to only access the data when they are logged in. Therefore we integrated with the application the use of the GE Identity Manager KeyRock. In practice this means that the users need to provide their FIWARE credentials (user name and password). The application supports Single-Sign-On protocols such as OAuth2, therefore the user's password goes to the Identity Manager directly, never to be intercepted by the application.

The application keeps the access policies encoded in an XACML format, compatible with the GE Authorization PDP AuthZForce. The GE PEP Proxy Wilma serves as a gatekeeper for the Data Broker service, permitting access only to the requests, which come from authorized and properly logged-in users, and which the AuthZForce grants the access to. The access policies base the conditions to grant access on the user's known attributes, which include the user's assigned role, the account (e.g., a household or a company) that the user belongs to, and a list of the meters owned or rented by the user. To support this functionality, we complemented the KeyRock's features by implementing our own **Attribute Manager**.

In the trial, the application's deployment relied on the FIWARE Lab's instances of KeyRock and AuthZForce. Thanks to the two GEs' source availability and well-written documentation, we are also able to deploy and use our own instances of the GEs. Additionally, our Attribute Manager can also be used in other contexts and applications, including other FINESCE use cases such as the ones from WP3.

## 2.7 Conclusion

The deliverable includes extended results for the following results sections: usage of Generic Enablers and FIWARE, Energy optimization, Simulations, and Other.

WP1 has a gained a positive experience with FIWARE, mainly linked to data processing and security chapters.

The developed trial infrastructure has been to proven to be a very flexible system with regard to handling different use cases and business models. One of the infrastructure's many strengths is the ability to deliver benefits both on a local level, optimization in the building, and at the same time on a global level, system optimization.

Concerning the CO2 model simulations, the analysis showed that the model trained with historic production from two years can be used to evaluate a CO2 signal. This signal is able to be used for DSM concepts in district heating grids to give incentives to move consumption into times of lower CO2 emissions. In general the three-seasonal model according to VDI 4655 should be preferred due to a lower SSE meaning a better performance. However, further training data is desired to reduce the SSE.

Lastly, all of the desired outcomes for WP1 (mentioned in the Introduction) have been met and documented in different deliverables.

## 2.8 References

[1] International Energy Agency. Technology roadmap energy-efficient buildings: heating and cooling equipment. Tech. Rep 2011.

## 2.9 List of Abbreviations

| | |
|---|---|
| B2B | Business to Business |
| BMS | Building Management System |
| CAPEX | CAPital EXpenditure |
| CENELEC | European Committee for Electro technical Standardization |
| CEP | Complex Event Processing |
| COTS | Commercial off-the-shelf |
| CPMS | Charge Point Management System |
| CSA | Cloud Security Alliance |
| DER | Distributed Energy Resources |
| DMS | Distribution Management System |
| DMTF | Distributed Management Taskforce |
| DSE | Domain Specific Enabler |
| EAC | Exploitation Activities Coordinator |
| EMS | Energy Management System |
| ERP | Enterprise Resource Planning |
| ESB | Electricity Supply Board |
| ESCO | Energy Service Companies |
| ESO | European Standardisation Organisations |

| | |
|---|---|
| ETP | European Technology Platform |
| ETSI | European Telecommunications Standards Institute |
| GE | Generic Enabler |
| HEMS | Home Energy Management System |
| HV | High Voltage |
| I2ND | Interfaces to the Network and Devices |
| ICT | Information and Communication Technology |
| IEC | International Electro-technical Commission |
| IoT | Internet of Things |
| KPI | Key Performance Indicator |
| LV | Low Voltage |
| M2M | Machine to Machine |
| MPLS | Multiprotocol Label Switching |
| MV | Medium Voltage |
| NIST | National Institute of Standards and Technology |
| O&M | Operations and maintenance |
| OPEX | OPerational EXpenditure |
| PM | Project Manager |
| PMT | Project Management Team |
| PPP | Public Private Partnership |
| QEG | Quality Evaluation Group |
| S3C | Service Capacity; Capability; Connectivity |
| SCADA | Supervisory Control and Data Acquisition |
| SDH | Synchronous Digital Hierarchy |
| SDN | Software defined Networks |
| SDOs | Standards Development Organisations |
| SET | Strategic Energy Technology |
| SET | Strategic Energy Technology |
| SG-CG | Smart Grid Coordination Group |
| SGSG | Smart Grid Stakeholders Group |
| SME | Small & Medium Enterprise |
| SoA | State of the Art |
| SON | Self Organizing Network |
| SS | Secondary Substation |
| TL | Task Leader |
| TM | Technical Manager |
| VPP | Virtual Power Plant |
| WP | Work Package |
| WPL | Work Package Leader |