FI.ICT-2011.1.8FINESCE

D7.4.2 Version 1.0

Consolidated Mid-term Analysis of Generic Enablers and Domain Specific Enablers Integration and Trial Impact

Contractual Date of Delivery to the CEC: *October 31, 2014*	
Actual Date of Delivery to the CEC:	
Author(s):	Padraic McKeever
Participant(s):	*RWTH Aachen*
Workpackage:	*WP7– FINESCE Project Management*
Estimated person months:	1.5
Security:	**PU**
Nature:	**R** = Report
Version:	1.0
Total number of pages:	**54**

Abstract:

This report is an intermediate consolidated report on the integration of GEs and DSEs into the FINESCE trial site infrastructures and on the results and impact achieved by the trials so far. It documents the current status of this ongoing activity. A final report will be prepared in February 2015.

Keyword list:

Generic Enabler, Domain Specific Enabler, FIWARE, Smart Energy, API

Executive Summary

This report is an intermediate consolidated report on the integration of GEs and DSEs into the FINESCE trial site infrastructures and on the results and impact achieved by the trials so far. It documents the current status of this ongoing activity.

The FINESCE project is performing field trials of the use of FI-WARE Generic Enablers (GE) in the smart energy sector and providing Domain Specific Enablers (DSE) and an Application Programming Interface (API) which will allow third-party clients (such as SMEs involved in Phase 3) 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:

- development of demand side response and demand-side management solutions for mixed-use buildings in a city district;
- efficient grid utilisation through demand-side management of prosumers;
- industrial demand side response interworking with a cross-border Virtual Power Plant (VPP);
- development of an energy marketplace to provide demand side response to varying energy production from Distributed Energy Resources (DER);
- controlling electrical vehicle charging to balance DER supply and improved utility communications.

FINESCE has performed an extensive evaluation of GEs to determine whether they can be used in the trials. This report gives an overview of the results of this GE evaluation activity.

Currently, FINESCE is at the end of M20 and is well advanced with integrating GEs and DSEs into the trial infrastructures. This report documents the current status of the GE and DSE evaluation and integration; for both the GE evaluation and integration, each individual vertical WP has produced detailed reports on their own activities and experiences.

The current status of GE integration is that twenty seven different instantiations of different GE implementations are already integrated into trial sites. Six trial sites have already integrated GEs.

FINESCE is developing twenty two DSEs, of which ten are already integrated. Open source code will be published for twelve of the DSEs, the rest will be published as open specifications.

Additionally, when FINESCE ends in September 2015, the trial sites will immediately continue to be available through the FINESCE API or to provide historical or aggregated data, also through the FINESCE API. Furthermore, the trial sites and the associated development work form the basis of further commercial developments by FINESCE project partners.

The FINESCE trial sites are largely already operational and are producing results. Experimentation will continue in the forthcoming months.

Authors

Partner	Name	e-mail
WP1		
E.ON	David Lillienberg	e-mail: david.lillienberg@eon.se
	Lars Norrman	e-mail: lars.norrman@eon.com
	Tommy Blom	e-mail: tommy.blom@eon.com
XLAB	Matej Artač	e-mail: matej.artac@xlab.si
WP2		
ACCIONA	José Luis Burón	joseluis.buron.martinez@acciona.com
ACCIONA	Ismael Rodriguez	ismael.rodriguez.cabado@acciona.com
Insero Software	Mikael Andersen	ma@inserosoftware.dk
Insero Software	Thomas Hune	tsh@inserosoftware.dk
Insero	Andy Drysdale	adr@insero.dk
Insero	Christina Ernst Jørgensen	chjo@insero.dk
WP3		
FIR	Julian Krenge	e-mail: Julian.Krenge@fir.rwth-aachen.de
	Marco Roscher	e-mail: Marco.Roscher@fir.rwth-aachen.de
RWTH	Adam Malik	e-mail: A.Malik@wzl.rwth-aachen.de
	Marija Stevic	e-mail: MStevic@eonerc.rwth-aachen.de
QSC	Ulrich Hacker	e-mail: ulrich.hacker@qsc.de
	Mirja Steels	e-mail: Mirja.Steels@qsc.de
	Adelheid Weinert	e-mail: adelheid.weinert@qsc.de
XLAB	Matej Artač	e-mail: matej.artac@xlab.si
	Marko Kuder	e-mail: marko.kuder@xlab.si
SOPTIM	Heiner Halbach	e-mail: heiner.halbach@soptim.de
WP4		
Engineering Ingegneria Informatica S.p.A.		
	Pasquale Andriani	e-mail: pasquale.andriani@eng.it
	Luigi Briguglio	e-mail: luigi.briguglio@eng.it
	Leandro Lombardo	e-mail: leandro.lombardo@eng.it
	Massimiliano Nigrelli	e-mail: massimiliano.nigrelli@eng.it
	Dario Pellegrino	e-mail: dario.pellegrino@eng.it
	Stefano De Panfilis	e-mail: stefano.depanfilis@eng.it
Synelixis Solutions Ltd.	ArtemisVoulkidis	e-mail: voulkidis@synelixis.com
	Theodore Zahariadis	e-mail: zahariad@synelixis.com
ISMB Istituto Superiore Mario Boella	Mikhail Simonov	e-mail: simonov@ismb.it
TeamWare Srl	Gianluca Zanetto	e-mail: gianluca.zanetto@teamware.it
Devolo AG	Markus Waechter	e-mail: markus.waechter@devolo.de
	AnilMengi	e-mail: anil.mengi@devolo.de
Grenoble Institute of Technology		
	José Sanchez Torres	e-mail:jose.sanchez@g2elab.grenoble-inp.fr
	Didier Boeda	e-mail: didier.boeda@g2elab.grenoble-inp.fr
	Raphaël Caire	e-mail: raphael.caire@g2elab.grenoble-inp.fr
WP5		
WIT	Miguel Ponce de Leon	e-mail: miguelpdl@tssg.org
	David Kirwan	e-mail: dkirwan@tssg.org
	Jason Whelan	e-mail: jwhelan@tssg.org
	Hisain ElShaafi	e-mail: helshaafi@tssg.org
ESB	John Howard	e-mail: john.howard2@esb.ie
Alcatel-Lucent Deutschland AG		
	Matthias Sund	e-mail: Matthias.Sund@alcatel-lucent.com
	Detlef Hartmann	e-mail: detlef.hartmann@alcatel-lucent.com
FUNITEC	Agustín Zaballos	e-mail: zaballos@salleurl.edu
	Ramon Martin de Pozuelo	e-mail: ramonmdp@salleurl.edu
	Joan Navarro	e-mail: jnavarro@salleurl.edu
	Alan Briones	e-mail: abriones@salleurl.edu
ALSTOMGrid	Amir Fazeli	e-mail: amir.fazeli@alstom.com
RWTH	Mohsen Ferdowsi	e-mail: MFerdowsi@eonerc.rwth-aachen.de
Orange Polska S.A.	RafałArtych	e-mail: rafal.artych@orange.com
WP7 Technical Management		
RWTH	Padraic McKeever	e-mail: pmckeever@eonerc.rwth-aachen.de
	Bettina Schäfer	e-mail: bschaefer@eonerc.rwth-aachen.de
	Antonello Monti	e-mail: amonti@eonerc.rwth-aachen.de

Table of Contents

1. Introduction

The purpose of this report is to consolidate the FINESCE WPs' individual reports on the Task X.3 – "Generic and Domain Specific Enabler Integration" (up to project month M20), to present an overview of the results already achieved and the continuing impact of the project after its end in September 2015 (M31).

At the time of writing this report, FINESCE is at the end of M20. In the Trial Design task (M1-M7) an initial assessment of the Generic Enablers (GEs) was made, and the architecture of the trial sites was defined, including definition of what functionality would be implemented using GEs and Domain Specific Enablers (DSEs). The Trial Preparation task (M4-M11) developed trial infrastructures.

This report is an output of the GE and DSE Integration task (M4-M24) where GEs and DSEs have been integrated into the trial infrastructures, which are ready for trial experiments to be performed (Trial Implementation task, M10-M26). The WPs already wrote individual Dx.3 reports providing an overview of, and status report on, the GE and DSE integration activities, covering the period up to M13 which were consolidated into the D7.4 report. This is the second report of three on this topic; a final report will be prepared in M24. In the period covered by this report, new partners have joined FINESCE and their GE integration work is included, as well as current information on the existing partners' work.

The current status of the integration of DSEs and other components into the FINESCE trial sites is described in Ch. 2. Because there are seven separate trial sites, each dedicated to a different area in the Smart Energy domain and each with its own physical and functional architecture, overviews of the individual trial sites' integration architectures and explanations of the purpose of the various components are presented in order to aid the reader's understanding of the status. The architectures presented in Ch. 2 complement the detailed descriptions of the FINESCE trial infrastructures that have already been published in the Dx.2 deliverables. Based on the presented integration architectures, the Domain Specific Enablers that have been, or will be, produced are then listed and their status described.

Ch. 3 gives an overview of the way in which FINESCE is using GEs and its experiences with them. There broad similarities across the FINESCE trial sites in the basic usage of GEs are shown, along with the way that each individual trial site also has its own particular uses of GEs. By its nature and by reason of limited space, a consolidated report is limited in the detail it provides: a more comprehensive description with more detailed comments on the GEs is contained in the underlying project-internal Wiki and Dx.3.2 and Dx.3 reports.

The GE selection and evaluation activities have been ongoing since the beginning of FINESCE and are still ongoing. The Dx.1 reports presented initial evaluation of the GEs' specifications with respect to the trial sites' functionality. Ch. 4 outlines how GEs have been selected and evaluated and gives an overview of the GE evaluation results and of GEs intended to be used in the trial sites. The GE evaluation also includes a comparison of selected GEs to commercial alternatives.

Outline information on trial results obtained so far are presented in overview in Ch. 5. Details are available in the underlying Dx.3 reports. The trial experimentation is continuing and more complete results will be presented in M31.

As FINESCE will end in less than a year's time, the question of what its lasting impact beyond that time will be is of especial interest to SMEs and web entrepreneurs who wish to develop their own applications using FINESCE API. Ch. 6 presents an overview of the impact and future exploitation of the trials.

Ch. 7 draws conclusions on the GE evaluation and integration, the DSE integration, the impact of FINESCE after September 2015 and the results already produced.

2. Integration of Domain Specific Enablers and other Components

In the FINESCE trial sites, FIWARE GEs have been integrated along with other software and hardware components, some of which have been defined by FINESCE as Domain Specific Enablers (DSEs). DSEs are open, re-usable software components. In order to be re-usable, specifications must be made available, thus allowing a developer to make his own implementation of the DSE. A reference DSE implementation must exist; in FINESCE the reference DSE implementations are in the FINESCE trial sites, or associated cloud infrastructure, hosting the respective DSEs

This chapter outlines how the FINESCE DSEs have been selected and gives an overview of their current status. An overview of the current functional architecture of the FINESCE trial sites is presented, indicating how the DSEs, GEs and other trial-site components are integrated.

2.1 DSE Selection Method

The generic functional architecture of the FINESCE trial system consists of a number of independent trial site systems which are externally accessible via a homogenised, service-oriented API which offers a single point of access to, and mediates towards, the individual APIs offered by the the trial sites.

Each trial site system consists of a set of interworking subsystems. Each subsystem, in turn, consists of lower-level subsystems, GEs, DSEs and other trial-site-specific components (which may have been developed in FINESCE or sourced outside the project). An overview of the structure of the FINESCE trial sites is given in the Chapters 2.3.1 to 2.3.5 below and also in deliverable D7.2[1].

Not all the components that are developed in FINESCE can be categorised as DSEs. In general, externally-developed components cannot be defined as DSEs due to licensing and ownership issues. Further reasons why a component developed in FINESCE may not be a DSE are because:

- its operation depends on the presence of the other trial-specific components with which it interacts and which may not be available as open source or may even be proprietary. In such a case, a higher-level “component” may be defined as a DSE, whose specifications but not whose code is published;
- it integrates sub-components which are developed by external organisations. Where such components are fundamental to the system, and cannot be separated; making software available on an open source basis would violate component licencing agreements.
- it is a component of the internal tool set of the project partner.

Another criterion for a component to be classified as a DSE is that its functionality is considered to be useful to developers in the Smart Energy domain.

2.2 Overview of FINESCE DSEs

FINESCE has analysed which components shall be made available as DSEs, with the result shown in Table 1. Currently a total of twenty two DSEs are planned.

There are two options for making DSEs available:
- Option 1: the specification of the DSE is made publicly available, royalty free.
- Option 2: the specification of the DSE is made publicly available, royalty free and the code is available Open Source.

[1] FINESCE D7.2 Version 1.0, “Consolidated Trial Description and Testbed Manual“

WP	DSE Name	DSE Description	T & C Options	Status
WP1	GE Integration Kit	Documented process including some examples for the established integrations within WP1	Option 1	Available Q4 2014
WP2 Horsens	Aggregator	Mediates various external device-related APIs to a trial-site internal format	Option 1	Integrated
	Configurator	Distributed configuration."Load Distributor" for the API Mediator, distributes the configuration of the entire system between all instantiated API Mediators	Option 1	Integrated
WP2 Madrid	Temporal Consistency	Check for mistakes in data from different sources based on historical data available	Option 1	Integrated
	Scene Configuration	Define customised scenes based on multiple parameters and get warning if a scene happens	Option 1	Integrated
WP3 Factory	Modbus Connector (ModConn)	Protocol Adapter from local shopfloor to Gateway Data Handling GE	Option 1	Available Q1 2015
	ODBC Event Sink (EvSi)	Local storage of shop floor events	Option 2	Available Q1 2015
	Production Planning and Control System Integrator (PPSI)	Connecting factory planning systems to the cloud	Option 1	Available Q1 2015
WP3 VPP	Generation Schedule Manager	Makes schedule from forecast data and actual generation data	Option 1	Integrated
	Production Schedule Manager	Makes schedule of possible variations of factory production schedules.	Option 1	Integrated
WP4	Contract Tailor Processor	Calculates a new contract for a specific customer based on an incentive plan (issued by a Retailer and approved by a Market Regulator Authority).	Option 2	Available Q1 2015
	Contract Processor	Stores all the contracts, their validation, and associated user profiles. . It supports operations like contract update/modification (by a Retailer) and contract approval (by end customers).	Option 2	Available Q1 2015
	Social Events Interface (Social2Orion)	REST-based client send Social events data loaded from an external provider to an instance of ORION Context Broker GE	Option 2	Integrated
	Weather Condition Interface (WeaFor2Orion)	Timer service that collects data from a weather forecasting service and sends them to to an instance of ORION Context Broker GE	Option 2	Integrated
	Contract Information (ContractInformation-2Orion)	REST service which allows clients (e.g. Retailers) to register data about cost of energy produced from the DERs, costs of transmission system and power plants, energy costs in an instance of ORION Context Broker GE.	Option 2.	Integrated
	Metering (Metering2Orion)	Gets metering data from AMM2Metering, translates it into NGSI10-compliant format and publishes it on an instance of ORION Context Broker GE	Option 2	Integrated
	Regulation Compliance Rule Engine	Checks and confirms the compliance of an incentive plan (issued by a Retailer) to established regulations and rules.	Option 2	Available Q1 2015
	Incentive Processor	Calculates a new incentive plan for a specific customer based on an issue resolution plan (issued by an Aggregator and approved by a DSO).	Option 2	Available Q1 2015
	AMR2AMI (Sensor2AMI)	Gathers Smart Meters information through DLMS/COSEM protocol and feeds it into an instance of IDAS (Backend) Device Management GE (which then sends it on to an instance of ORION Context Broker GE	Option 2	Available Q1 2015
	Issue Detector Processor	Detects issues related to weather forecast, social events, consumption/production aggregated data, power losses and voltage	Option 2	Available Q1 2015

WP	DSE Name	DSE Description	T & C Options	Status
		drops		
WP5	Charging Optimisation	Optimisation of Electric Vehicle Charging	Option 1	Available Q4 2014
Synelixis	FINESCE API	Homogenisation of the distinct API layers offered by each FINESCE trial	Option 2	Available Q1 2015

Table 1: FINESCE DSEs

Of the twenty two DSEs listed in Table 1 above, ten are already integrated in the architecture of their respective trial sites, four in WP2, two in WP3 and four in WP4. The other twelve DSEs will be integrated between Q4 2014 and Q1 2015. Open Source code will be published for twelve of the DSEs.

Two additional DSEs (Issue Detector Processor and Sensor2AMI from WP4) have been added since the last report in D7.2 (March 2014). Also AMM2Metering is no longer a DSE as its functionality is the same as that offered by AMR2AMI (Sensor2AMI).

2.2.1 Availability of DSE Code as Open Source

As indicated in the two options listed for making DSEs available above, making the code available Open Source is optional. The source code will be available for some, but not all, of the FINESCE DSEs, as shown in Table 1. The decision to not make code available has been made for each DSE by the FINESCE project partner who owns it. The reason for not making code available may be one of the following:

- concerns about resourcing the provision of support to developers to use the code;
- components interface to proprietary components (not developed in FINESCE) whose interface is only available inside the trial site, so that the DSE code cannot be used except in the trial site;
- intention that code will be made available but it is too early to be explicit on it;
- difficulty in handling the code of the individual DSEs separately from other components it is linked to.

2.2.2 DSE Documentation

The DSE documentation is worked on internally in the project Wiki and, when it is ready to be published, the html can be exported from the Wiki and linked to the FINESCE project web page, http://www.finesce.eu/, where the DSE documentation and open source downloads will be publicly accessible. The layout of the DSE part of the FINESCE web site will follow that of the FIWARE Catalogue, so that the FINESCE DSEs can later easily be included in the FIWARE Catalogue.

2.3 WP Integration Architectures

The sub-chapters below give an overview of the FINESCE trial site functional architectures at the current stage of integration of DSEs, GEs and other trial-site components. The intention is to outline where the DSEs fit into the trial sites.

2.3.1 WP1 Integration Architecture

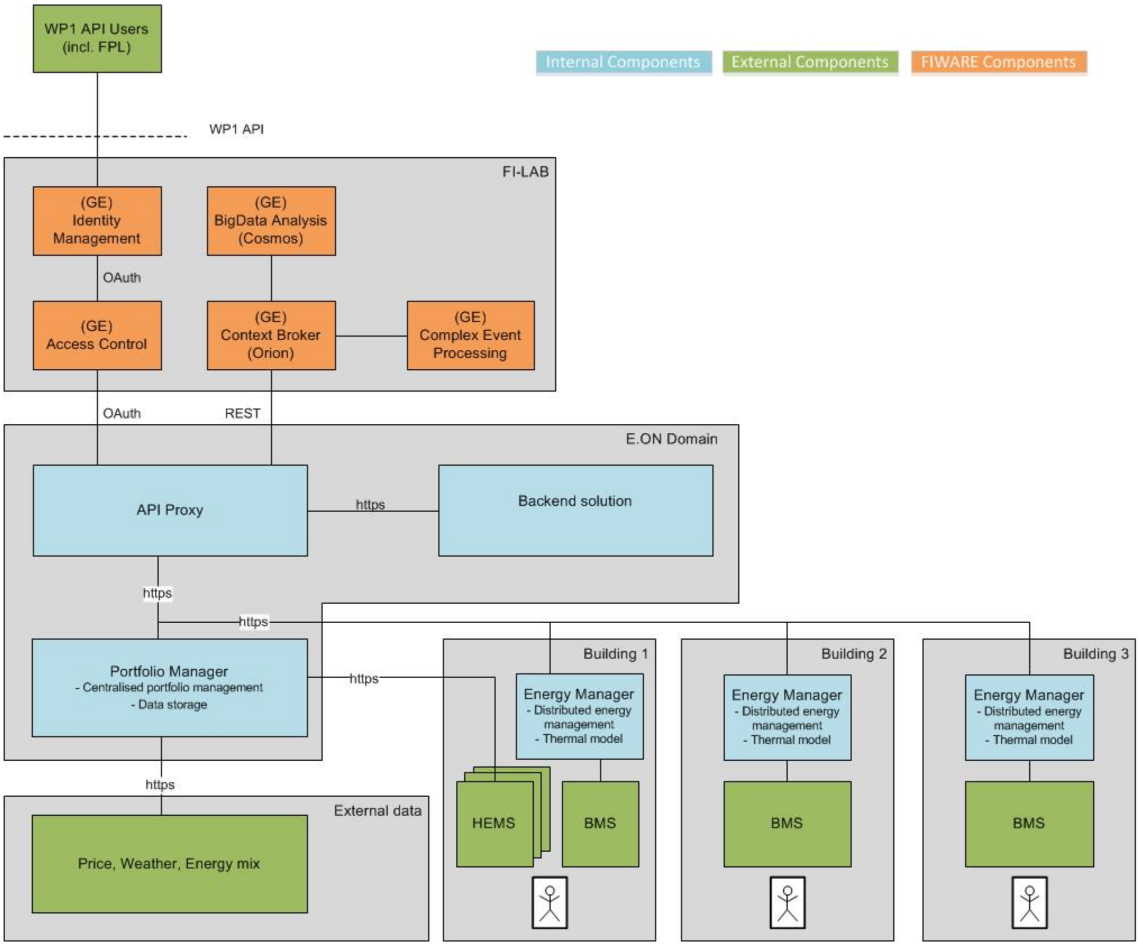

Figure 1: - WP1 Intermediate Integration Architecture

WP1's functional architecture, shown in Figure 1, comprises a distributed energy management function and a centralised energy portfolio management function.Figure 1 shows the following components which have been integrated into the WP1 trial site:

- *Backend solution,* handles WP1 API user access rights and roles towards E.ON IT systems;
- *API Proxy:* Provides the WP1 API, acting as a frontend integrator towards E.ON systems;
- *Portfolio Manager:* back-end server platform for centralised portfolio management
- *Energy Manager*, present in each building in WP1, performs distributed energy management
- *Building Management System (BMS)*, present in each building in WP1, computer based control system monitoring and steering the heating supply and ventilation;
- *Home Energy Management System (HEMS)*, present in all apartments in the first building in WP1 (Roth Fastigheter), computer based control system monitoring and steering the heat usage;
- *GE Integration Kit (DSE)* is a documented process for integrating GEs, with examples from WP1, to help others in such implementations; it is not a component in the architecture shown in Figure 1 above;
- *FINESCE Presentation Layer (FPL):* cloud-based visualisation app working towards WP1 and WP3 trial systems. It interworks with a graphical web app which FPL users run in their browsers. FPL will use *Identity Management (Keyrock)* and *Access Control* GEs for user authentication and *Wirecloud Mashup GE* for testing the graphical web app.

2.3.2 WP2 Integration Architecture

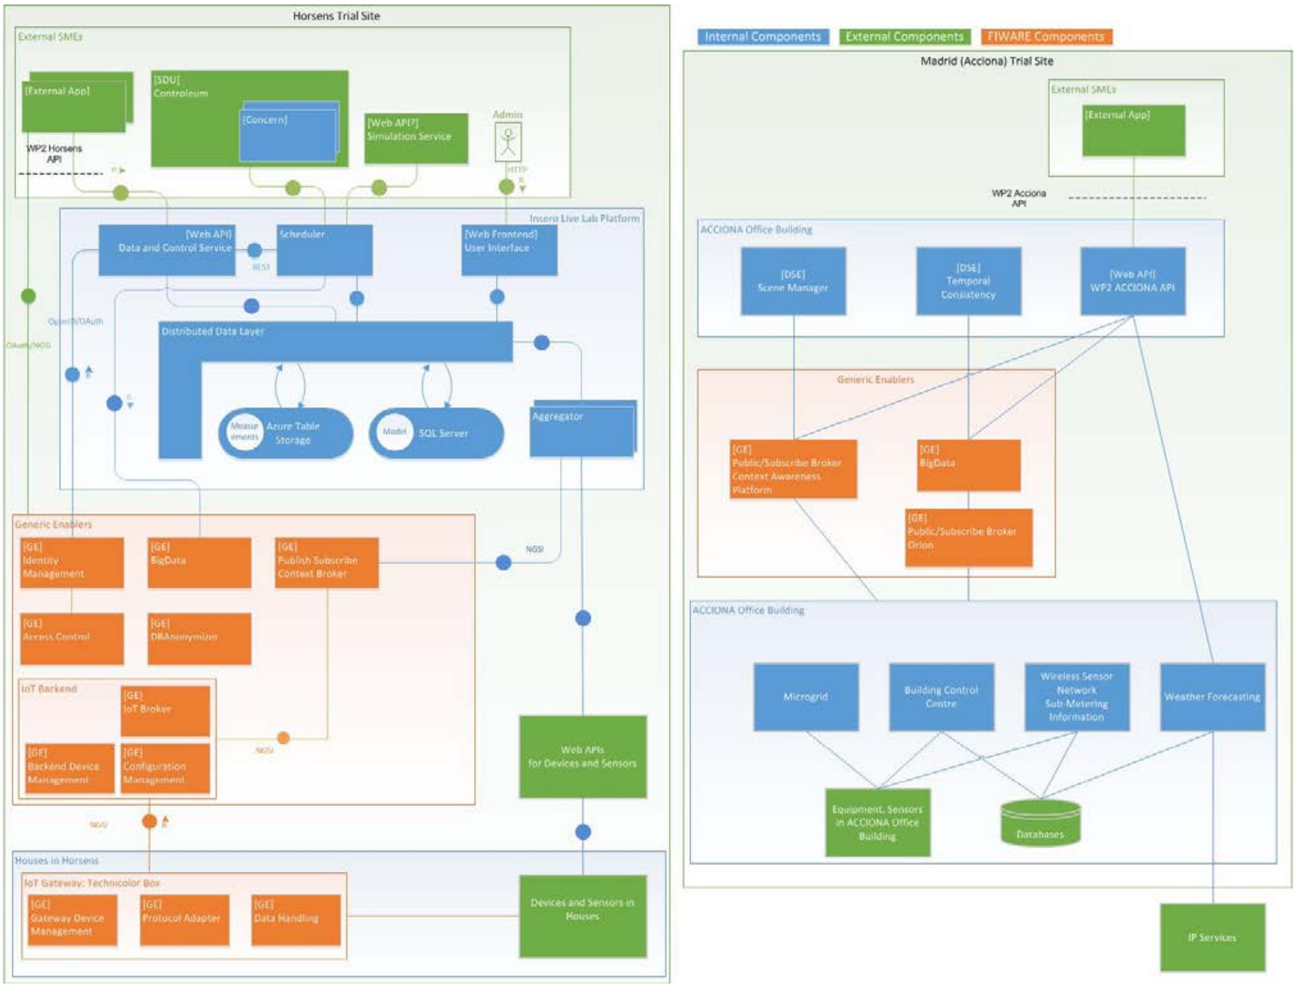

Figure 2 –WP2 Intermediate Integration Architecture

WP2 consists of two trial streams being implemented and executed independently but with the possibility to share data. The first trial stream consists of 20 single family houses in the Horsens area, Denmark. The houses will be equipped with energy producing and energy consuming components. In Madrid, Spain, a second trial stream consists of an office building equipped with a Building Management System, components for electricity production as well as electricity storage. Figure 2 depicts the current architecture of the WP2 Horsens and Madrid trial sites. It shows the following components which have been integrated into the Insero Live Lab Platform on the Horsens trial site:

- *User Interface (web service):* administrator's system management portal, used to manage everything that can be configured in the system and visualise some of the data that is being collected.
- *Data and Control Service* (formerly *Control Service and Historic Data Service*): exposes the API that all external services will use. Provides the historic data services to get raw measurements and aggregate measurements. Includes OData querying syntax for retrieving the sources of the measurements and the different types of measurements that exist. Provides API for the control services for the devices in the houses.
- *Aggregator (*formerly *API Mediator)* DSE*:* mediates the different APIs exposed from the web services that gather the data from the devices.
    - o it allows NGSI clients to retrieve information through the Publish Subscribe Broker;
    - o it stores all measurements from the houses into a NoSQL store for later retrieval and analysis;
    - o it stores external data from weather and energy services into a NoSQL store for later usage in the *Scheduler* and *Controleum*;

  - it performs monitoring on all equipment in the houses, so causes of errors can easily be identified and eliminated.
- *SQL Server* has a model of the physical configuration of equipment and the software deployment;
- *Scheduler* (formerly *Optimization Algorithm Service):* uses the measurements from the houses, simulation of heat loss for the individual houses and weather and electricity production/usage prognosis from various services in order to optimise the way energy is used in the smart grid as a whole. Control instructions for the houses are provided by *Controleum* and executed in this component.
- *Distributed Data Layer:* provides the communication bus facilitating the exchange of measurement data, queries and control messages between the components in the system. It features a distributed platform providing location-transparent communication between components.
- *Controleum* is a framework for multi-objective control problems where each objective represents a concern in the control domain, like maintaining the air quality within a comfort-band, or reducing electricity consumption by extending the comfort-band in case of a demand-response event. It is the responsibility of Controleum to find a Pareto optimal solution and to identify conflicts between objectives.
- *Azure Table Storage* is used for storing measurement.

Figure 2 shows the following components which have been integrated into the Madrid trial site:

- the four systems at the Acciona building which provide data from the building's equipment and sensors to FIWARE GEs;
- *Web API* module exposes, through appropriate security control, the the API that all external services will use;
- *Temporal Consistency* DSE pre-processes data stored in the Big Data GE by from any of the Madrid trial data sources: the Weather Forecasting module, the Building Control Centre, the Microgrid data concentrator, and/or the Smart Metering gateway. It detects inconsistencies and removes non-valid values. It uses HiveQL Client (Backend) to interface to Big Data GE.
- *Scene Manager* DSE allow configuration of a set of multiple parameters (scene), based on which different alerts can be triggered and offered to subscribed users. It works together with the Public/Subscribe Context Broker – Context Awareness Platform GE in order to perform the event configuration, receive alert notifications and manage the subscriptions to those events.

2.3.3 WP3 Integration Architecture

WP3's trial sites comprise a Virtual Power Plant (VPP) comprising over 10 Distributed Energy Resource sites in Belgium and Germany and a Smart Factory in Aachen. Figure 3 shows how the VPP, the Smart Factory, a simulation of the VPP and the FPL (described in Ch. 2.3.1 on WP1 above) are connected via a FIWARE cloud-based infrastructure.

The FINESCE Presentation Layer (FPL) is a component developed by the new partner XLAB which implements a common data representation of several applications (in our case WP1 and WP3 trial sites), performs data aggregation and visualisation.

The WP3 trial site comprises the following components:

- *VPP Power Plant* sites (windmills, PV plants and biogas plants): here, proprietary QSC *Gateways* collect energy data from the DERs' meters and send it in encrypted form to a central proprietary QSC *Gateway Server* which decrypts and stores the data locally and forwards it to the *Complex Event Processing* GE on the FI Testbed.
- The DSE *Generation Schedule Manager* provides information about the VPP's energy generation. The DSE *Production Schedule Manager* processes the factory's production steps (including the associated power requirement) in to a production plan. The *Complex Event Processing* GE takes the output of these two DSEs to plan how to balance the energy production and consumption.

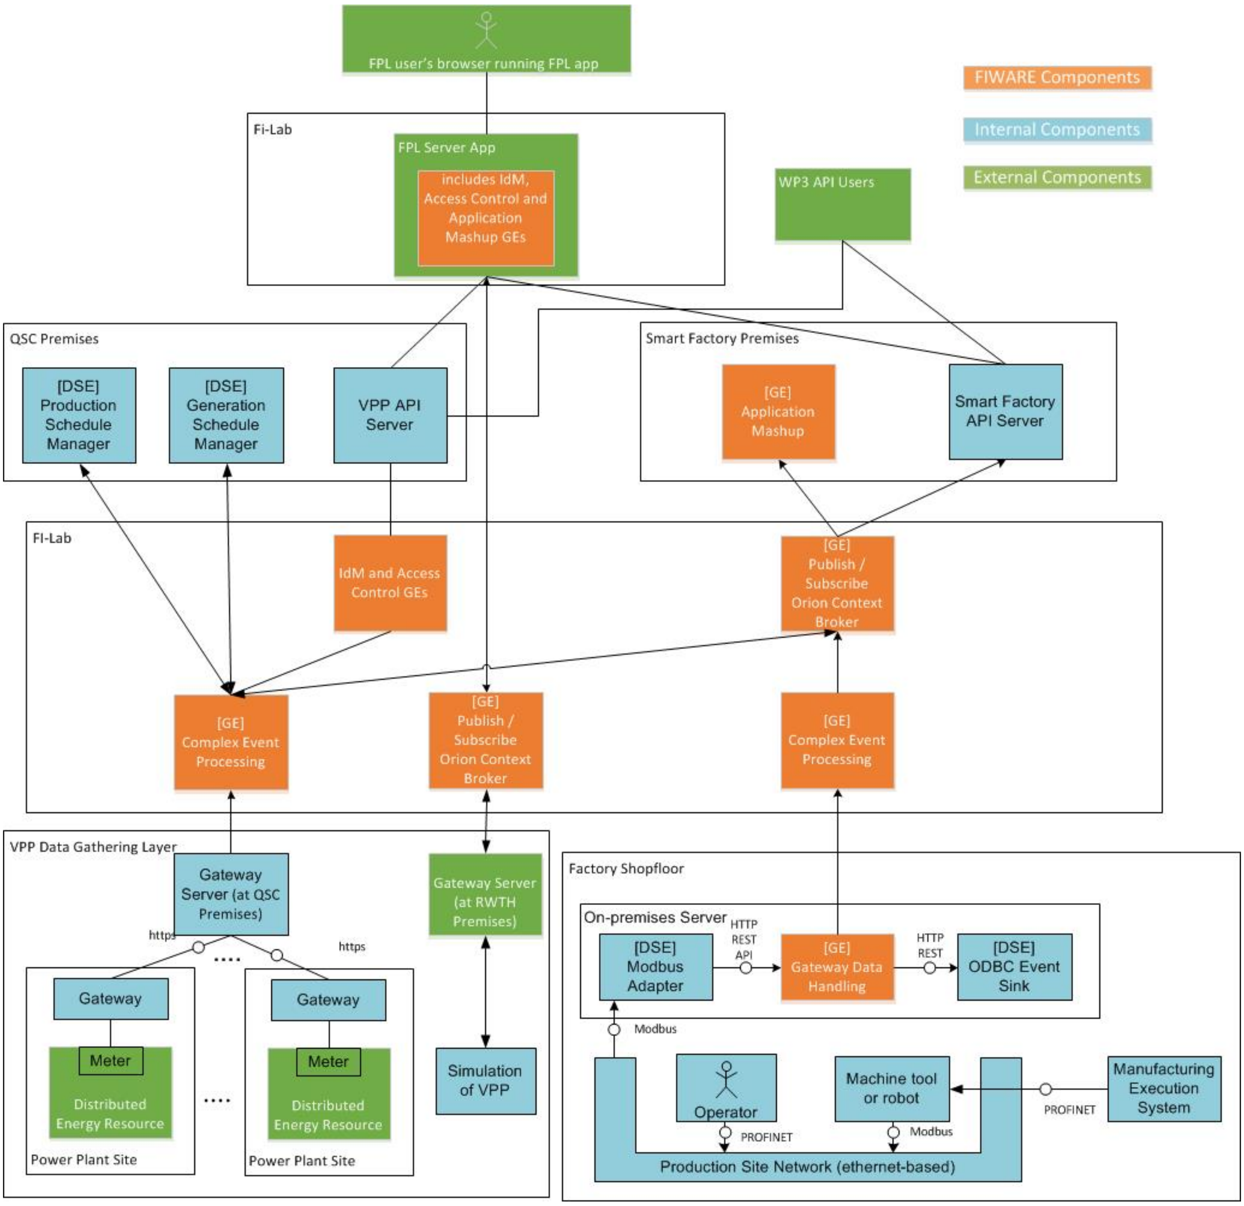

Figure 3: - WP3 Intermediate Integration Architecture

- The DSE *Production Planning and Control System Integrator* (not yet integrated) will allow feedback of information from the cloud to the factory planning and control systems.
- The *Modbus Connector* DSE supports the connection of the Factory shopfloor infrastructure to the *Gateway Data Handling* GE. The *ODBC Event Sink* DSE provides local storage of data, which is typically mandated by manufacturers. Both these DSEs and the GE are instantiated locally at the factory.
- In the Factory, the *Application Mashup* GE will allow integration of the factory-related events into existing factory production management systems (not yet performed).

2.3.4 WP4 Integration Architecture

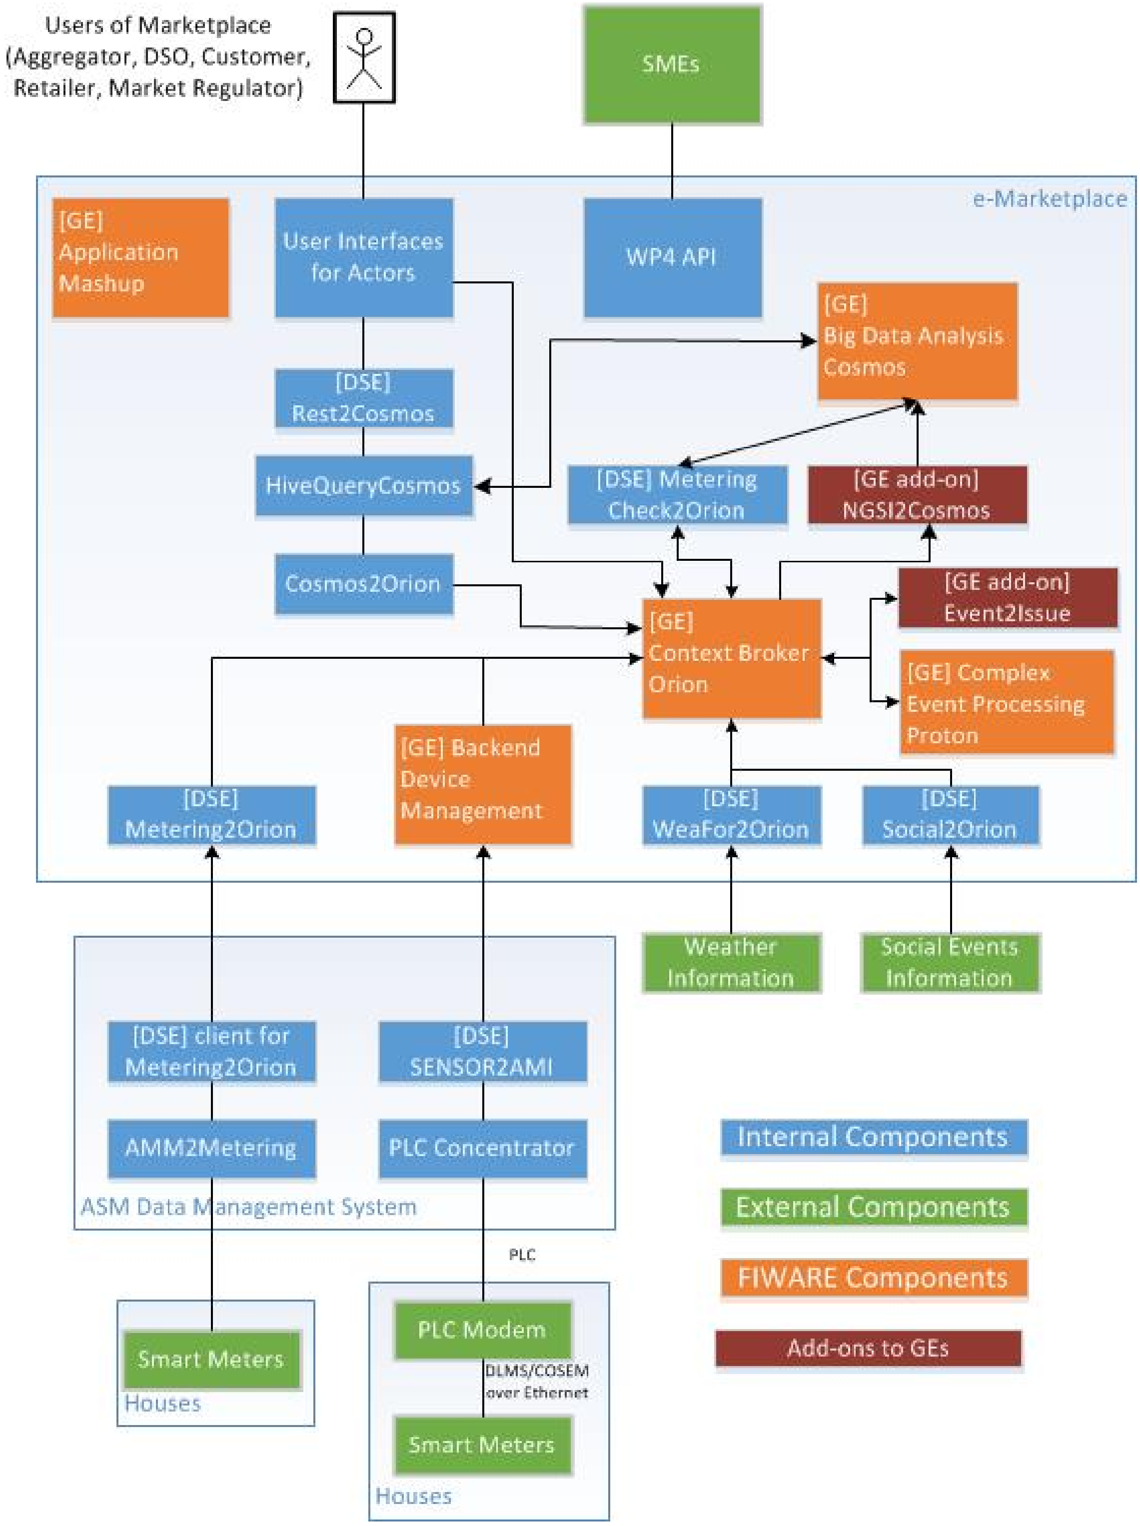

Figure 4–WP4 Intermediate Integration Architecture

Figure 4 depicts the current architecture of the WP4 Terni trial site. It shows the following components which have been integrated into the Terni trial site:

- *AMM2Metering*, which retrieves “raw” consumption and production data from the smart meters installed at the trial site and passes them over IP to *Metering2Orion*;
- *Metering2Orion* DSE, which translates metering data coming from *AMM2Metering* into an NGSI10-compliant format (ORION context events) and finally publishes them on the ORION Context Broker;

- *WeaFor2Orion* DSE, which collects data from a weather forecasting service every five minutes and sends them to ORION Context Broker.
- *Social2Orion* DSE, a REST-based client that exposes a @POST method via which an external provider can send data on social events (such as concerts, football matches, etc.) that can affect consumption/production in the trial site area.
- *NGSI2Cosmos* is a special data injector connecting ORION to COSMOS. It subscribes to the data to be persisted) and when their values change, it automatically appends the new value in a COSMOS file;
- *Cosmos2Orion* is a Timer service which retrieves aggregated information on total consumption and production for the trial site area and sends it to ORION;
- *HiveQueryCosmos*, which analyses and retrieves the data from COSMOS GE via HIVE, It establishes the connection to the HIVE Data Warehouse, executes the Hive Query in HQL language, retrieves aggregated data and sends them to the data model.
- *Rest2Cosmos* is a REST-based client that exposes methods (GET) to retrieve the aggregated data from COSMOS GE via the HiveQueryCosmos module.
- *AMR2AMI*, is a metering capture system which is currently being implemented. At the trial site, the Smart Meters, which are of a different type to those communicating with *AMM2Metering*, communicate using the DLMS/COSEM protocol over Ethernet to a PLC modem. A PLC concentrator at the substation terminates the PLC and communicates over Ethernet to the *SENSOR2AMI* DSE, which comprises
    - IAM-Reader, which collects real time metering data, converting it to the DLMS/COSEM protocol if needed;
    - IAM-Server Relay, a middleware server which receives DLMS/COSEM metering objects from IAM-Readers and posts them to IAM2IDAS;
    - IAM2IDAS, which converts the DLMS objects to the format required the *(Backend) Device Management (IDAS)* GE;
- *ContractInformation2Orion* DSE is a REST-based client that exposes methods (POST) whereby an external provider can send data on the cost of energy, the cost of energy produced by DERs and the cost of system transmission power plants;
- *Issue Detector Processor* DSE (not shown in Figure 4), which is composed of the following sub-modules:
    - *Event2Issue* listens to updates in the CEP singleton entity, which is sent by PROTON to ORION as NGSI10 notifyContext after the evaluation of pre-defined patterns of behaviour, processes the information and generates (or updates) the corresponding Issue in the Orion Context Broker (using NGSI10 updateContext request).
    - *Cosmos2SCILAB* DSE is a Timer service that retrieves weather forecast and historical consumption/production data from COSMOS GE (via *Rest2Cosmos*) and then stores it in a configurable directory accessible to the SCILAB simulator software;
    - *SCILAB2Orion* DSE(not shown in Figure 4) is a Timer service that retrieves “power losses”, “voltage drops” and consumption/prediction data from SCILAB and sends them to ORION;
    - *INP SCILAB* DSE (not shown in Figure 4) accepts weather forecast and historical consumption/production data retrieved by *Cosmos2SCILAB* via *Rest2Cosmos* from *Big Data Analysis (Cosmos*) GE. It performs a simulation and returns to *SCILAB2Orion* information on power losses and voltage drops in the trial site grid) and short time (next twenty-four hours) predictions on consumption/production. *INP SCILAB* is developed on SCILAB, an open licensed software that is compatible with the Linux version installed on the FI-Lab VM where it has been deployed.
- e-Marketplace offers an external interface over an API, shown in Figure 4 as *WP4 API*.

Note that Figure 4 does not show WP4’s usage of Object Storage or other Cloud GEs used to deploy e-Marketplace in FI-Lab.

2.3.5 WP5 Integration Architecture

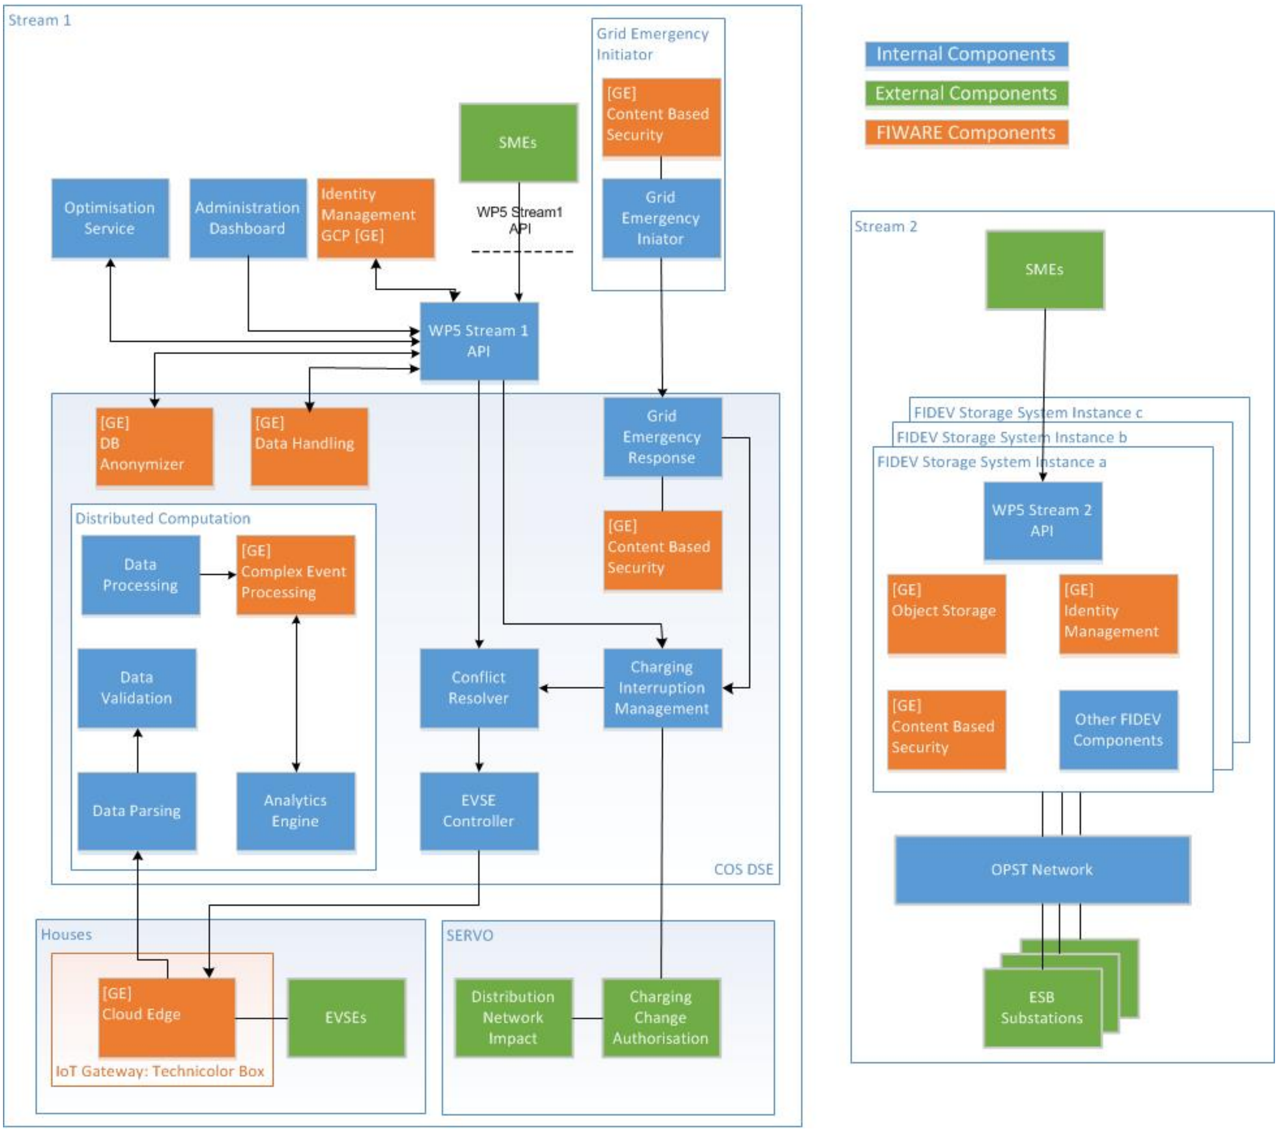

Figure 5–WP5 Intermediate Integration Architecture

The Stream 1 trial comprises the following components:

- the *Electrical Vehicle Supply Equipment (EVSE)* to charge electrical vehicles parked at the houses;
- *Cloud Edge* GE at the houses supporting *COS – EVSE communication;*
- the *Charging Optimisation System (COS)* DSE controls EVSE charging;
- *SERVO,* an external DSO system which authorises *EVSE* charging based on its knowledge of the effect that a given EVSE's charging has on the LV and MV grid conditions; *SERVO* exposes an OpenADR VTN interface to the COS; *SERVO* is not currently fully operational, so COS is currently working on the assumption that *SERVO* always authorises charging requests;
- The *Grid Emergency Initiator* allows an grid emergency state to be defined and communicated encrypted to *COS;*
- *Optimisation Service:* algorithm which generates an EV charging schedule using EVSE state information retrieved from the COS and sends it to the COS for implementation during the next optimisation cycle (15 minutes);
- *WP5 Stream 1 API* provides an API for both internal WP usage and for external usage

The Stream 2 trial comprises the following components:
an *Optical Packet Switch and Transport Network* connecting several MV substations via optical switches and optical fibre

FIDEV Storage System is a distributed, cloud-based data storage system. Currently it is planned to use EVSE data which will be fed from *COS* to the OPST network, rather than DSO data from the substations, as originally intended. The provision of an API by FIDEV, as part of FINESCE API is under consideration.

3. Integration of Generic Enablers

3.1 Integration Status and Experiences

FINESCE is organised with seven independent trial sites. Although independent, there are broad similarities in their use of GEs. The predominant pattern 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.

GEs from the 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. This typical GE usage pattern is shown in Figure 6 for the different trial sites; the letters (A-E) used for the GE groups in Figure 6 is referred to in the per-WP descriptions in the sub-chapters below to show which GE groups are used in particular trial sites; if a GE group is not mentioned below then it is not used by the given trial site.

Please note that Figure 6 is meant as a simplified overview only. It shows the broad pattern of GE usage, but it does not show all GEs used (some WPs use GEs from FIWARE chapters which are not indicated). The purpose for which individual GEs are used by the WPs is not described below: for this, and for detailed comments on the GEs, please refer to the underlying WPs' Dx.3.2 reports.

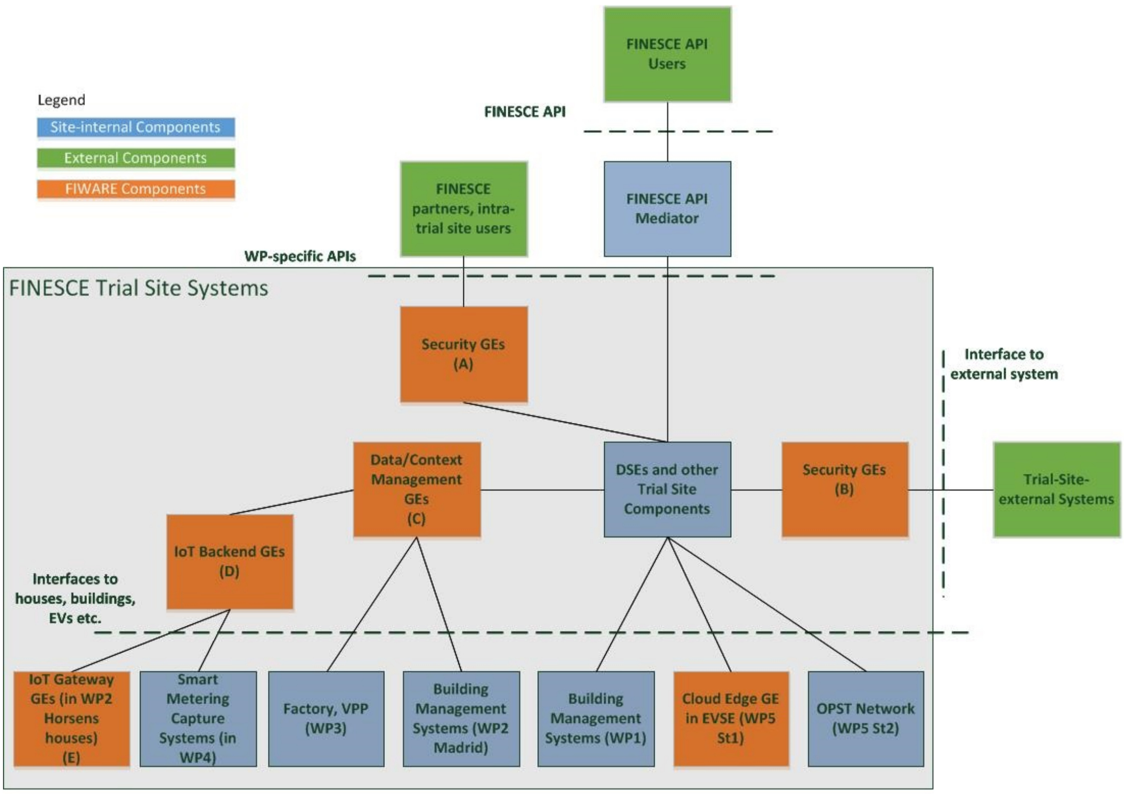

Figure 6 – Broad Pattern of GE Usage Across FINESCE Trial Sites

3.1.1 WP1

WP1's GEs are integrated via the API Proxy, as depicted Figure 1, where the *Context Broker* GE has a very central position.

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Security (A)
Use of *Identity Management (One)* and *Access Control (Thales)* to bridge between the users of the FINESCE WP1 API and E.ON's internal security system is under study.

Data/Context Management (C)
Integration of the *Big Data Analysis (Cosmos)* and *Publish/Subscribe Context Broker (Orion)* GEs is being completed and the experience has been positive. A new module "Cygnus" is used to notify changes from Orion to Cosmos.

IoT Backend (D)
A backend comprising the *Backend IoT Broker (NEC)* to enable handling larger numbers of Energy Managers and *Backend Configuration Manager (Orion)* for version management of software in Energy Managers and *Backend Device Manager (IDAS)* is under study. WP1's IoT Backend would be fed from the *Publish/Subscribe Context Broker (Orion)*, i.e. not directly from the buildings.

3.1.2 WP2

Integration of the GEs into the trial is ongoing along with development of DSEs and other software components. The GE evaluation of the GEs is based on experiences from the integration process and the communication with - and support from - the individual GE developers.

Nine GEis are integrated, and they are showing satisfactory results in terms of stability. The main challenge during the integration has been compatibility between GEis intended to work together in the design. In general, the stability of the GEis has improved since the last report and is no longer considered an obstacle for the integration.

3.1.2.1 WP2 Horsens

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Security (A)
The trial site is experimenting with using *Access Control (Thales)* to manage authorisation policy and decisions and has successfully integrated *Identity Management (GCP)* to authenticate trial site users (external apps and owners of houses involved in the trial).

Usage of the *DB Anonymizer* GE to evaluate data anonymisation policies is under question due to its removal from the FIWARE Catalogue and its use of MySQL, which is not used in the trial.

Data/Context Management (C)
Problems with trying to setup a dedicated instance of the *Orion Context Broker* GE on the FI-Lab testbed were experienced but overcome with help from the GE developer and the GE in integrated. However, the GE does not support access control functionality, so that it is only being used internally in the site, not towards the public API.

IoT Gateway (E)
A number of GEs are located in the Technicolor hardware box in the houses participating in the trial. As delivered, this box hosts the *Cloud Edge* GEi (which can locally execute applications) but this is not used by WP2. Rather, WP2 is considering hosting *Protocol Adapter (ZPA)*, *Gateway Data Handling (EspR4FastData)* and *Gateway Device Management (Fraunhofer)* on the box. The latter GE supports the use of an ETSI M2M format between the IoT Gateway and IoT Backend: as WP2 doesn't have devices which use this format, this GE is not a must for WP2. WP2 is currently addressing compatibility problems with using NGSI between the *ZPA* and *EspR4FastData* GEs and towards the IoT Backend. The withdrawal of Technicolor from FIWARE is causing concern about continuing commercial availability of the box.

IoT Backend (D)

A backend comprising the *Backend IoT Broker (NEC)* and *Backend Configuration Manager (Orion)* is integrated.

3.1.2.2 WP2 Madrid

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Security (A)
It is planned to use *Identity Management (GCP)* in a similar fashion to the Horsens trial.

Data/Context Management (C)
Integration of the *Big Data Analysis (Cosmos)* and *Publish/Subscribe Context Broker (Context Awareness Platform)* GEs is being completed and the experience has been positive. A new module "Cygnus" has replaced the unsatisfactory "ngsi2Cosmos" for converting XML data streams input to the *Publish/Subscribe Context Broker (Orion)*, which has been integrated.

3.1.3 WP3

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Security (A)
VPP plans to use *Access Control (Thales)* and *Identity Management (GCP)* in a similar fashion to the Horsens trial.

Data/Context Management (C)
VPP has integrated *Complex Event Processing* GE to collect events from the VPP and Smart Factory and distribute processed events. The Smart Factory has integrated *Complex Event Processing* GE and *Publish/Subscribe Context Broker (Orion)* for event collection and distribution.

IoT Gateway (E)
Smart Factory has integrated *Gateway Data Handling (EspR4FastData)* GE locally on a factory server to act as a gateway between the factory and the cloud.

3.1.4 WP4

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

WP4 has stopped considering GEs that are not included in the FIWARE Catalogue (*Cloud Proxy, Data Handling GE, DB Anonimyzer GE, Content Based Security*).

Cloud
The FI-Lab implementation of the IaaS Data Center Resource Management and Self Service Interfaces GEs is being used. These GEs were integrated already at M13 and have worked well.

The public FI-Lab instance of *Object Storage* GE is included in the design and integrated. It is being used to store and share files among Wirecloud users and works well.

Apps
Application Mashup (Wirecloud) works as expected and is fast and stable. Support has been excellent. Its integration continues as more "panels" are being included in the UI for the different stakeholders included in the process.

Security (A)
Identity Management (Keyrock) is included in the design and integrated. It is used to authorise users and give them a single sign-on to FI-Lab and Wirecloud.

Data/Context Management (C)

Big Data (COSMOS), *Context Broker (ORION)* and *Complex Event Processing (PROTON)* were already integrated at M13, *CEP* and *Orion* work as expected, without issues. Also, *Big Data* works as expected, although the Shark/Spark component is sometimes overloaded thus resulting in poor performance (and sometimes in system crashes).

IoT Backend (D)
There are now two ways that the Smart Meters are connected to the Terni trial infrastructure: in the original trial the meters are connected via the AMM2Metering and Metering2Orion DSEs to the Orion Context Broker; now, the new partners from the Open Call winning consortium (FIPONET and IAM) are developing another metering capturing system (AMM2AMI) which collects data from a different smart meter type than the original trial and sends the data via the *Backend Device Management (IDAS)* GE to the *Orion Context Broker*. The integration of *IDAS* is currently ongoing. The experience has been that the time taken to familiarise with the GE was greater than expected, that support from *IDAS* developers was needed but not always promptly available.

3.1.5 WP5

3.1.5.1 WP5 Stream 1

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Security (A)
Identity Management (GCP) and *Data Handling* are included in the design and integrated. *IdM GCP* is used to authenticate API users and *Data Handling* to enforce privacy of EVSE data. *DB Anonymizer* is included in the design.

Security (B)
Content-based Security is included in the design for encryption of data between COS and external TSO systems.

Data/Context Management (C)
Complex Event Processing (IBM) is included in the design to support provision of historical EVSE data.

3.1.5.2 WP5 Stream 2

Note: the letters in brackets e.g. "Security (A)" refer to the GE Groups illustrated in Figure 6 above.

Object Storage GE will be used for storage of data in FIWARE cloud

Security (A)
Identity Management (Keyrock) will be used for user access and management.

Security (B)
Content-based Security will be used for encryption between FIDEVs.

3.2 Where the GEs Are Instantiated

There are, in principle, four different ways to instantiate the GEis in a trial site: FI-PPP Testbed (FI-WARE-internal but available to Use Case projects with FI-WARE support), FI-Lab (for public use, a shared resource for everyone, not scalable), XIFI (allows Use Case project to have more individual use with its own GE instance, is scalable) or own local instantiation.

The way each test site intends to instantiate its GEis is shown in Table 2.

WP	**Instantiation**
WP1	FI-Lab, considering using XIFI
WP2 Horsens	FI-Lab, except for • GEs on hardware located at houses in trial. • Identity Management GCP, where the multi-tenant instantiation hosted by the GE owner which must be used
WP2 Madrid	• FI Testbed for Publish/Subscribe Broker, Data Handling, Orion Context Broker and IdentityManagement • Telefonica I+D cluster for BigData
WP3 Factory	• FI-Testbed for Object Storage, Big Data Analysis • FI-Lab for Gateway Data Handling, CEP, Publish/Subscribe Broker Orion • Gateway Data Handling local.
WP3 VPP	FI-Lab cloud.
WP4	FI-PPP Testbed (until M18). From M19, migration to FI-Lab / Spanish XIFI node completed.
WP5 Stream 1	Local, except for Identity Management GCP (see WP2 Horsens above).
WP5 Stream 2	Local

Table 2: Instantiation of GE Implementations by WPs

4. Selection and Evaluation of Generic Enablers

This chapter discusses firstly how GEs have been selected and evaluated in FINESCE and then gives an overview of the evaluation results, including a number of comparisons of GEs with commercial competitors. The evaluation results comprise a lot of detailed comments and observations accumulated during the process of studying and integrating the GEs; the value of this information lies in its rich detail. The detailed GE evaluation results are contained in the WPs' Dx.3 and Dx.3.2 reports. The comparisons of the GEs to alternative commercial / Open Source products are collected in ANNEX 2.

A word on terminology: the term "selection" is used here to refer to the complete process of choosing to use a GE; selection is a continuous iterative per-trial-site process which continues until the GE has been successfully integrated in the trial site or until it has been decided not to use the GE. The term "evaluation" refers to the process of studying, working with, testing and integrating the GE. Evaluation is also a per-trial-site, continuous and iterative process. Its output is a documented GE evaluation. This evaluation is used to make GE selection decisions.

4.1 GE Selection and Evaluation Methodology

The FINESCE DoW outlined each WP's original scope and trial site architecture, including a proposed usage of GEs. The Trial Design task further developed the site architectures and included a study and selection of the GEs. Each WP has each developed plans for the evaluation of GEs, initial selection of the GEs upon which trial design is based and integration of the chosen GEs into their trial site infrastructures, as documented in the Dx.1 deliverables. This process has continued in the subsequent Trial Preparation and the current GE and DSE Integration tasks. The basic organisation of FINESCE into independent trial sites carried over into the GE evaluation process, with each WP being responsible for selecting and evaluating its own used GEs. The project Wiki has been used to share and make visible the evaluation results. The selection of GEs has, therefore, been a continuing, iterative process since before FINESCE even started 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 involves 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 is included in the trial site architecture then it will subsequently undergo a process of integration into the trial site, during which their functionality continues to be evaluated and more deeply understood as they are 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, as described in Ch. 3.1 above.
- the GE's terms and conditions and availability from FIWARE (in either the FI-PPP Testbed, FI-Lab or as a downloadable product, as per the trial site's specific needs);
- whether the GE's documentation is of sufficient quality to allow the GE to be studied and, later, to be integrated into the trial site;
- whether there is sufficiently good support of the GE by its developers.

The GE Evaluation began during the trial design activity where WPs studied the available GEs in order to see how the GEs could be used to implement their trial functionality. GEs for which the initial result was that the GE was of potential interest to the WP have been further evaluated over a longer period; unless a decision has been made by the trial site not to select the GE, this process of evaluation is still ongoing and will continue during the integration of the GEs into the trial infrastructures.

The FINESCE WPs are each building separate trial site infrastructures. Certain WPs provide more than one trial site: in WP2, the Horsens and Madrid trials are separate, as are Streams I and II in WP5; in WP3, the Factory and VPP trials are separate but linked by a cloud-based infrastructure. This is the reason why the two parts of both WP2, WP3 and WP5 are presented separately in this report.

The GE evaluation activity has, except for cross-WP discussions and information exchange, been performed separately by each WP.

The level to which the GE evaluation has been performed varies from GE to GE depending on the GE's perceived usefulness and the length of time it has been worked with. For example, if a WP looked at a GE and decided that its functionality was not relevant to the trial infrastructure and, hence, that it was not going to be used in the trial, then the evaluation result just documents this and no detailed evaluation is made.

The more usual case is that the GEs have been evaluated over a period of time and that a blog post has been made in the FINESCE Wiki detailing and recording the progress. Additionally, the WPs document their GE Evaluations in their Dx.3 and Dx.3.2 reports and by means of the formal GE Evaluation. These evaluation methods are described in the chapters below:

4.1.1 The GE Evaluation Wiki

FINESCE is using a project-internal Redmine Wiki to document the ongoing GE Evaluation work. The Wiki is being updated on an ongoing basis. It gives an overview (per-trial site) of the GEs that have been looked at. The purposes of documenting the evaluation are firstly to document the development of the WPs' work with the GEs and secondly to allow feedback to be given to FI-WARE on the GE experimentation in addition to the normal ongoing contact and discussions between FINESCE project members and GE developers about ongoing issues.

The GE Evaluation Wiki is organised in a tree structure with three levels.

The (level 1) entry page of GE Evaluation Wiki consists of a list of the GE implementations (GEi) that at least one WP is experimenting with, has already included in its trial site design or has already included in its trial site (marked by the codes "E", "U" and "D" which are explained before Table 1 below).

The user can click on the title of a GEi to reach the evaluation entry page for that GEi (level 2). Here there is information per-WP on the GEi's usefulness and status. The user can access an individual WP's evaluation of the given GEi (level 3) by clicking on the corresponding link.

At level 3 there are individual GEi evaluations on a per-WP and per-GEi basis. The WPs with two sites can, if both sites have evaluated the GEi, enter separate evaluations here. The evaluation typically gives a log of the evaluation, comments on the GEi and status information. The excel sheet detailing the formal GE evaluation is also linked here.

4.1.2 Evaluations in Dx.3 and Dx.3.2 Reports

The Dx.3.2 reports are five individual mid-term reports on the GE and DSE Integration which the present report is consolidating. The Dx.3 reports are the five corresponding preliminary reports. Both the Dx.3s and Dx.3.2s contain detailed evaluation comments on the GEs. A short overview of the evaluation results is given in Ch.3.1 above.

4.1.3 Formal GE Evaluation

In addition, a formal evaluation of the GEs will be 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, as shown in Table 8 in ANNEX 1: Functional suitability, Performance Efficiency, Compatibility, Usability, Reliability, Security, Maintainability and Portability. Under each of these headings, a set of sub-criteria has been developed under which each GE receives a score (1-5). Each criterion is allocated a Rank, indicating its importance; the ranks are (M)andatory, (D)esirable, or (O)ptional, to which Weighting factors (WF) are assigned as 5 for mandatory, 3 for desirable, and 1 for optional. The weighted scores (i.e. score multiplied by weighting factor) are summed to give a numerical evaluation result. This is also expressed as a percentage of the maximum possible score.

This formal evaluation is only possible for those GEs which have been worked with over a considerable time period and tested thoroughly; in effect, this means that the formal evaluations will mainly be done for the GEs which are being considered for integration into the trial infrastructures. The formal evaluations are presented in Ch. 4.3.

4.2 Overview of GE Evaluation Results and GE Usage Plans

This chapter gives an overview of the WP's GE evaluation results and plans for using GEs. The total number of GE instantiations (GEis) that have been evaluated is 175. This includes many GEis which are not in the current FIWARE Catalogue anymore, or never were included in it. Those GEis which were experimented with by at least one WP are shown in Table 3 below.

In the first column the name of the GEi is given.

The second column gives the GEi name and the GEi owner: for several GEs, there is more than one implementation (GEi). In some cases, the GEi assessment was made before the GEi was entered in the FI-WARE Catalogue and, subsequently, no GEi was entered in the Catalogue: these GEis are marked with "None" in the second column.

The other columns give the current evaluation result for each WP (and for the individual trial sites WPs 2, 3, 5). The meaning of the entries is as follows:

- a blank means that the WP did not evaluate the GEi;
- "N" means that the GEi was evaluated but it was neither experimented with nor taken into consideration in the trial design;
- the codes "E", "U" and "D" are the same as those used in the FI-PPP GE Cockpit and mean that the GE is under evaluation or has already been included in the trial infrastructure:
    - WP has already integrated the GEi into its trial infrastructure ("D")
    - WP has already taken the GEi into consideration in its trial infrastructure design ("U")
    - WP plans to experiment with it and consider it based on results ("E")

The progression between these codes is from "E" to "U" and finally to "D". If a GEi does not progress from "E" or "U" to "D", but is discarded, then this is indicated by "EX" or "UX", respectively.

In Table 3, the number of GEis per-trial in each of the states "E", "U" and "D" is shown, giving an overview of the GEs that are under evaluation, included in the trial design or already integrated into the trial infrastructures. All trial sites (except WP5 Stream 2) have already integrated GEis.

Table 3 shows that twenty seven different instantiations of different GE implementations have already been integrated into trial sites, that an additional eleven such different instantiations are included in the trial designs but not yet integrated and that an additional twenty four such instantiations are still being considered for inclusion in the trial design.

Many of the GEis have been evaluated separately by more than one WP. Just considering the GEis in statuses "E", "U" or "D"; there is a total of eighty instances of twenty nine different GEis in these statuses, i.e. twenty nine different GEis have reached at least status "U" and, counting that some of them are being used by several trial sites, there are eighty such instances.

FI-WARE GEs	GEi Name	WP1	WP2 Horsens	WP2 Madrid	WP3 Factory	WP3 VPP	WP3 FPL	WP4	WP5 Str.1	WP5 Str. 2
Cloud Chapter										
IaaS Data Center Resource Management	-	N						D	N	EX
Self Service Interfaces	Self Service Interfaces - Cloud Portal	N						D	N	
Object Storage	-	N			E			D	N	U
Data/Context Management Chapter										
Publish/Subscribe Broker	Orion Context Broker	D	D	D	D		D	D	N	
Complex Event Processing (CEP)	IBM PROactive Technology ONline (PROTON)	E			D	D		D	U	N
BigData Analysis	COSMOS	D	E	D	E		E	D	N	
Publish/Subscribe Broker	Context Awareness Platform	N		D	N				N	N
Apps Chapter										
Application Mashup	WireCloud	N			U		D	D	N	
Store	-						E		N	
IoT Chapter										
(Backend) Device Management	IDAS DCA - TID	E	E					U	N	
(Backend) IoT Broker	IoT Broker – NEC	E	D					EX	N	
(Backend) Configuration Management	Orion Context Broker	E	D	E				UX		
(Backend) Configuration Management	IoT Discovery - UNIS	N						UX	N	
(Gateway) Data Handling	EspR4FastData	N	D		D			UX	N	
(Gateway) Device Management	OPENMTC	N	E					EX		
(Gateway) Protocol Adapter	ZPA	N	D		N			UX	N	
(Backend) Template Handler	Template Handler							EX		
(Gateway) Device Management	Ericsson IoT Gateway	N			E			EX	N	
(Gateway) Data Handling	SOL-CEP	N						UX		
Security Chapter										
Identity Management	KeyRock					E	U	D	E	U
AccessControl	-	E	E			E	U	E	N	N
Privacy	- /	N						E	EX	N
Identity Management	GCP	N	D	U					D	N
Identity Management	One-IDM	E							N	N
Data Handling	PPL	N		E				UX	D	N
DB Anonimyzer	DBA	N	E					UX	U	
Content-based Security	CBS	N						EX	E	U
I2ND Chapter										
Cloud Proxy a.k.a. Cloud Edge	CloudProxy	N	EX		N			EX	U	EX
(I2ND) Network Information and Control (NetIC)	Altoclient								N	EX
Total Evaluated	66 different GEis evaluated	29	13	7	13	3	6	24	59	21
Under consideration "E"	24	6	5	2	3	2	2	2	2	0
Considered but discarded “EX”	11	0	1	0	0	0	0	6	1	3
In design "U"	11	0	0	1	1	0	2	1	3	3
Was in design but then discarded, "UX"	7	0	0	0	0	0	0	7	0	0
Integrated "D"	27	2	6	3	3	1	2	8	2	0
Total in status “E”, ”EX”, “U”, “UX” or “D	80 instances of GEis / 29 individual GEis”	8	12	6	7	3	6	24	8	6

Table 3: Overview of GE Evaluation

FI-WARE GE Chapter	Number of "E" or “EX”	Number of "U" or “UX”	Number of "D"
Cloud	2	1	3
Data/Context Management	4	1	13
Apps	1	1	2
IoT	11	6	5
Security	13	8	4
I2ND	4	1	0
	35	18	27

Table 4:- GEis Integration Status per FIWARE Chapter

Table 4 shows the FIWARE chapters which have been of most interest to the FINESCE WPs. Most of the integrated GEis come from the Data/Context Management, IoT and Security chapters. A lot of GEis from the IoT and Security chapters have been considered for inclusion in the trial designs but have not progressed beyond status “E”, indicating that these GEis were not considered useful. The Security and IoT chapters also contain the most GEis that have been included in the trial design but not yet integrated (“U”), which is a further indication of the importance of these chapters for FINESCE.

GEi Status for WP ↓	WP1	WP2	WP3	WP4	WP5
Total under consideration "E"	6	7	7	2	2
Total in design "U"	0	1	3	1	6
Total in trial "D"	2	9	6	8	2

Table 5 – GEis Integration Status per WP

Table 5 gives a per-WP view of the the current status of GEi usage in FINESCE. It shows that the WPs are in the middle of the GEi integration activity: a good number of GEis have already been integrated but the integration is still a work in progress.

4.3 Results of the GE Evaluation

The process of evaluating the GEs, culminating with their integration into the trial sites involves an intensive learning about the GE and its behaviour in the field. Support from FI-WARE developers is essential to the success of the FINESCE trials. The feedback from the WPs on the interaction with the GE developers is generally very positive.

The GE evaluation results so far are documented in the individual Dx.3 and Dx.3.2 reports of the WPs and in the GE Evaluation Wiki (which is project internal). The WP’s Dx.3 and Dx.3.2 reports and Wiki entries log their experiences with the GEs. The overall experience with the GEs has been positive, although using the GEs has been a learning experience and has meant working with what are, in effect, products that are still in the field trial stage.

The formal GE evaluation method described in Ch. 4.1 using the criteria listed in ANNEX 1 has already been applied to a number of GEs. The graphs in Table 6 below show the results on a scale of zero to one for each of the main evaluation criteria. A score of one corresponds to the GE receiving full marks.

The diagrams show the results under the main evaluation headings only. The detailed evaluations for each sub-criterion are not shown here. In some cases a value of zero is given for the evaluation. This should not be interpreted as a bad score: it either means that the evaluation could not yet be performed or that the criterion was not considered to be applicable. Hence, these zero values can be ignored in the concerned diagrams as being not applicable.

Generally, WPs have applied this method to the main GEs they have worked with, as indicated by the labelling of the GE as “integrated” in the graphs. It is therefore reasonable that such GEs should receive good scores for categories such as “Functional suitability”.

In some cases GEs have been formally evaluated which have only reached the stage of “Experimentation” or “Design” and have then not been selected for integration, e.g. Complex Event Processing by WP1 or Backend Device Management by WP2 Horsens: for such GEs, the low scores in some categories reflect the concerned WPs’ decisions not to use the GE.

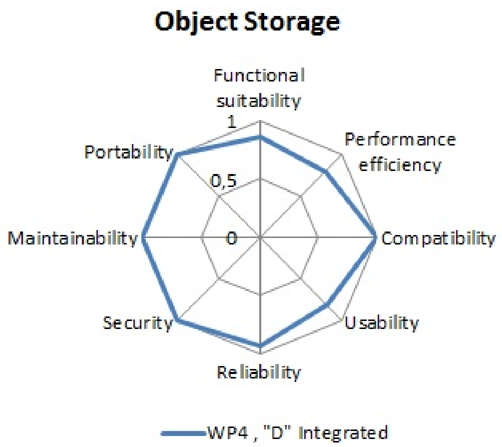

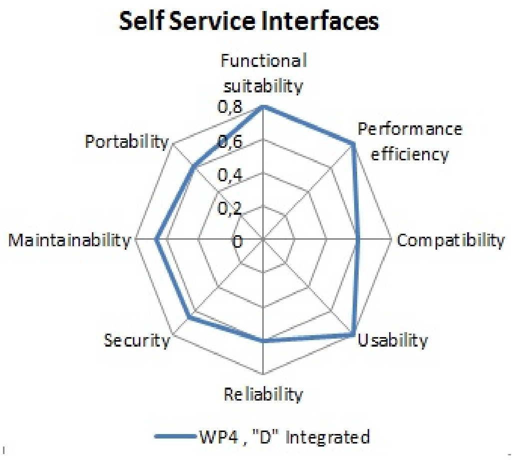

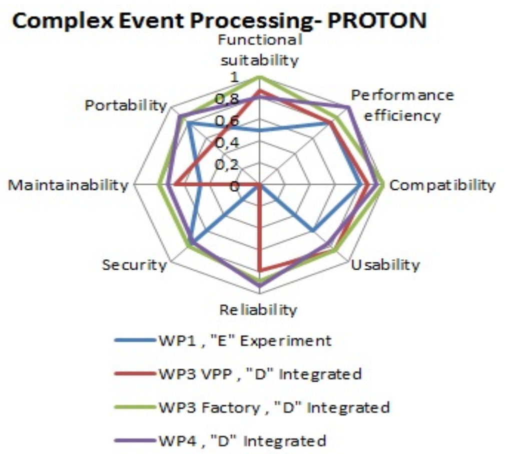

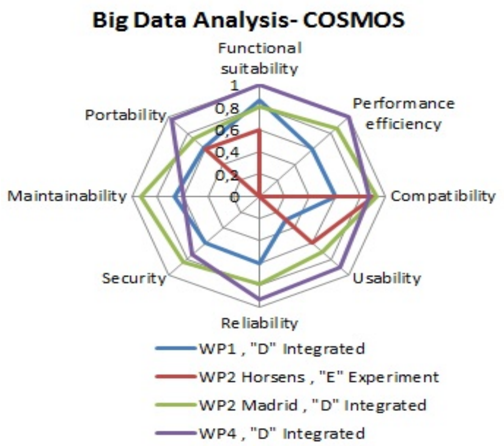

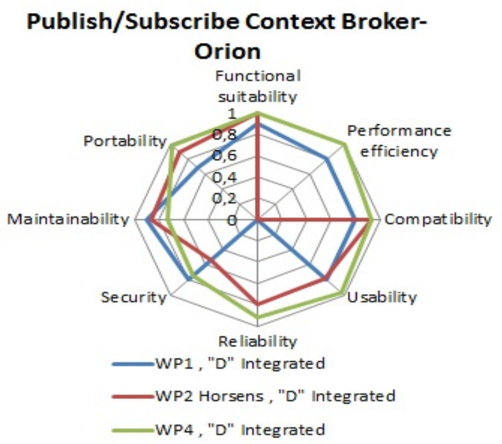

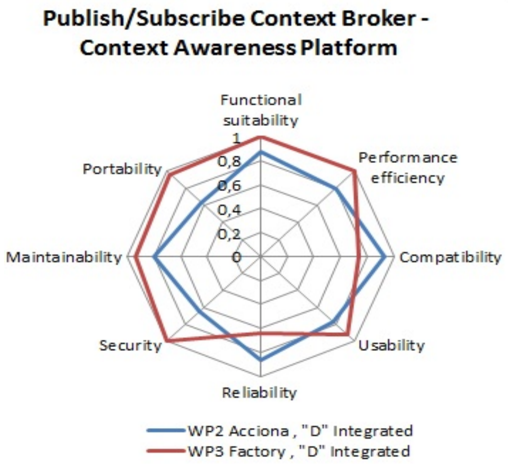

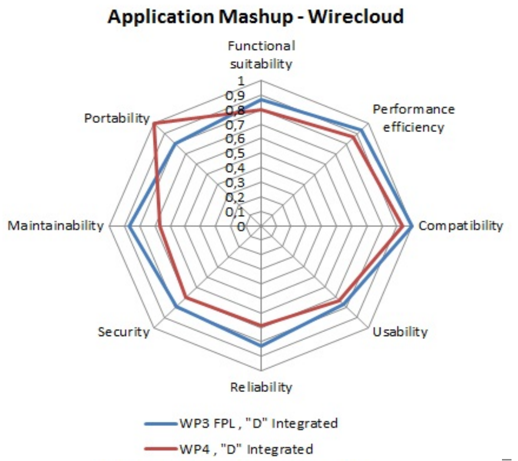
Application Mashup - Wirecloud
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP3 FPL , "D" Integrated
WP4 , "D" Integrated

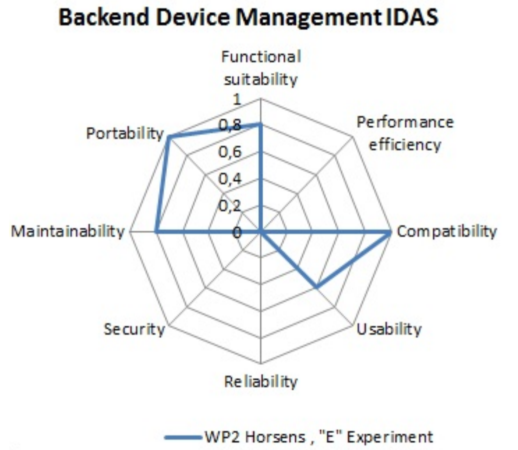
Backend Device Management IDAS
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "E" Experiment

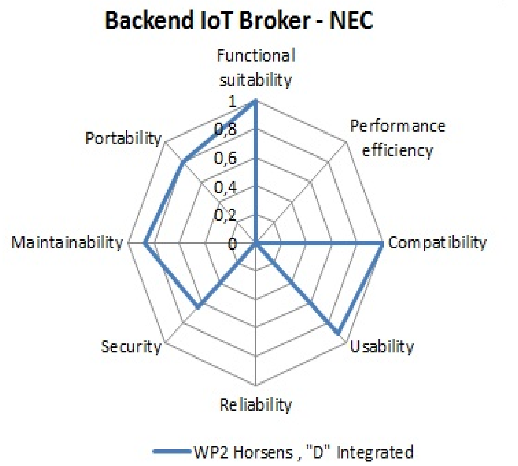
Backend IoT Broker - NEC
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "D" Integrated

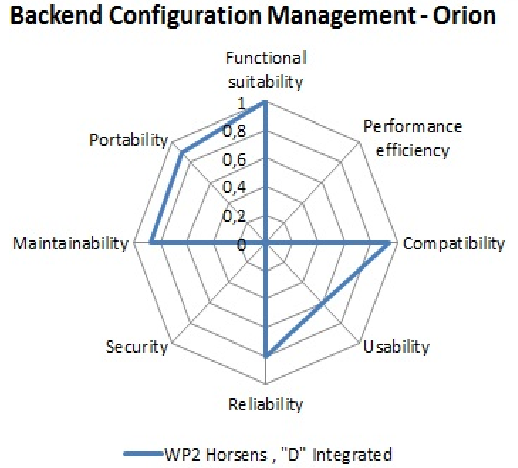
Backend Configuration Management - Orion
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "D" Integrated

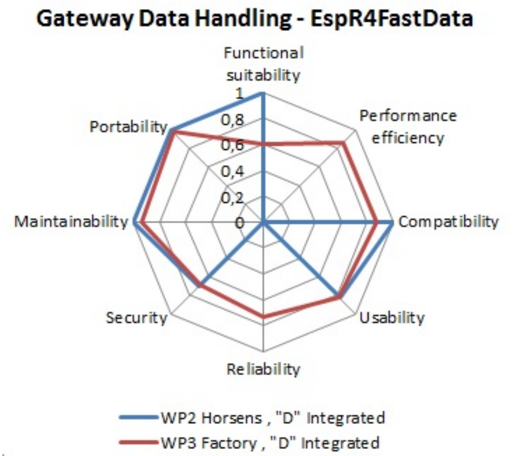
Gateway Data Handling - EspR4FastData
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "D" Integrated
WP3 Factory , "D" Integrated

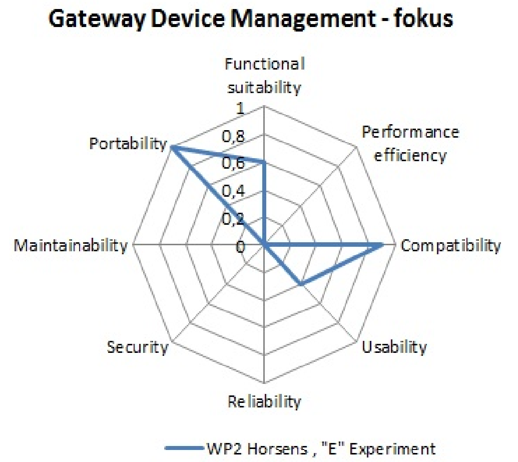
Gateway Device Management - fokus
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "E" Experiment

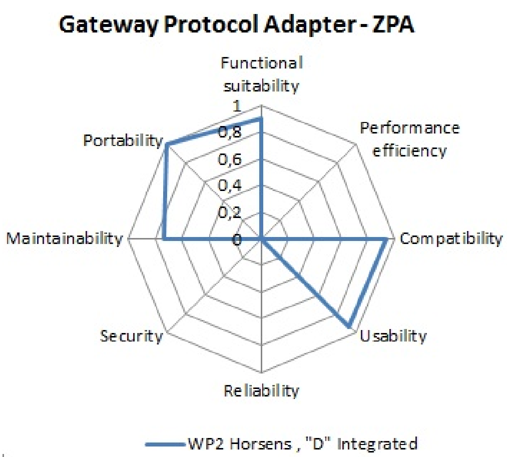
Gateway Protocol Adapter - ZPA
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP2 Horsens , "D" Integrated

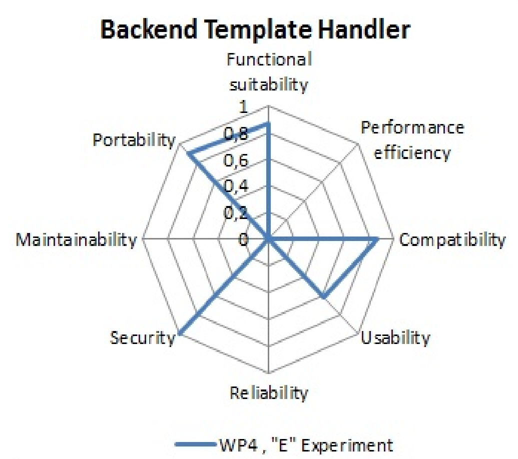
Backend Template Handler
Functional suitability
Performance efficiency
Compatibility
Usability
Reliability
Security
Maintainability
Portability
WP4 , "E" Experiment

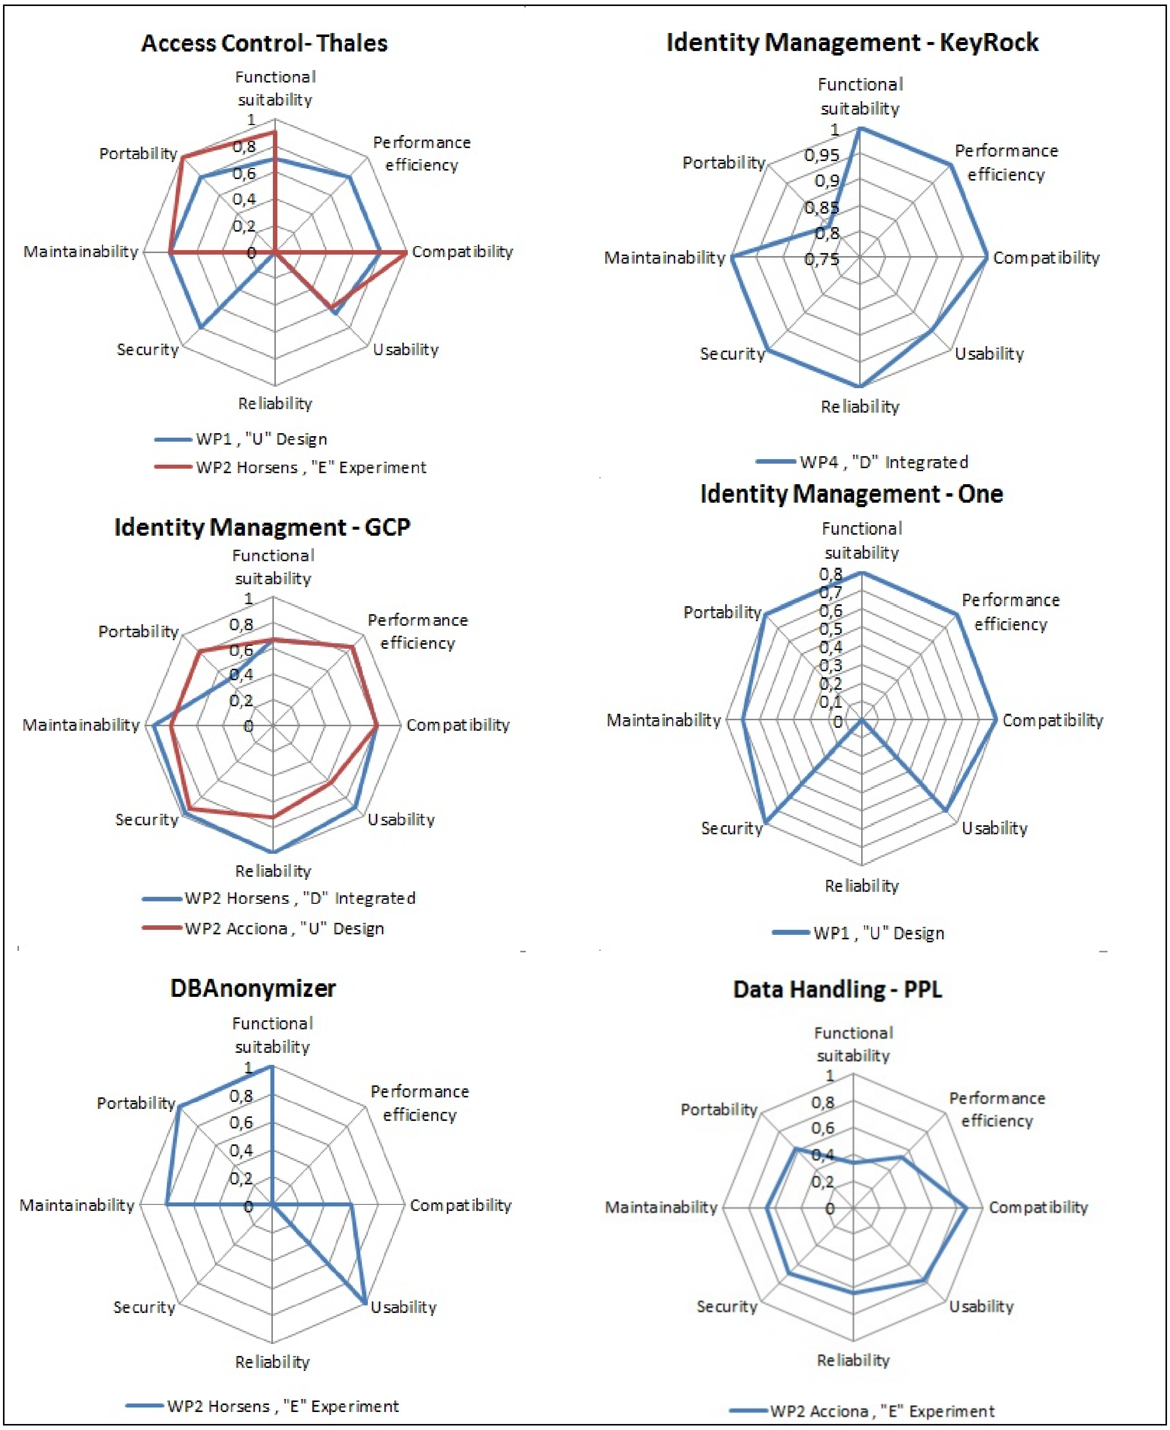

Table 6: - Diagrams showing Results of Formal Evaluation of GEs

4.4 Comparison of GEs With Other Commercial Alternatives

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: as outlined in Ch. 3.1 above, several GE chapters are particularly useful in the Smart Energy domain. 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).

A comparison of some of the most important GEs for FINESCE with commercial competition has been performed by FINESCE WPs. In general, the WPs have done the comparison for GEs they have worked extensively with and whose commercial competitor(s) they are familiar with. The details of the comparisons are given in ANNEX 2.

WP1 has compared *Big Data* GE with Hortonworks using the evaluation criteria of ANNEX 1. Hortonworks scored better for Usability (better quality of documentation, learnability).

WP2 compared *Identity Management GCP* GE with ASP.NET, using the evaluation criteria of ANNEX 1. The GCP is hosted by Deutsche Telekom, limiting its functional suitability and performance efficiency compared to the competitor, which is hosted within the application itself. GCP is considered easy to learn.

WP2 also compared *Big Data* GE with Google's BigQuery using the evaluation criteria of ANNEX 1. *Big Data* GE scores similarly to BigQuery for all categories except Usability, where it scores significantly lower.

WP3 discusses the *Public/Subscribe Context Broker- Orion* GE in detail, pointing out several positive features of the GE, including its use of open standards and its effective web interface. It is not compared to a competitor because there is no known competitor offering equivalent functionality.

WP3 also compares the GEs *Complex Event Processing* and Gateway *Data Handling* with several commercial competitors. The conclusion is that the GEs are as good as the competition in terms of functionality, configurability, scalability and interoperability.

WP4 has analysed the advantages of a number of GEs. The key factor for using GEs is the availability of support from the GE owner.

WP5 compared, using the evaluation criteria of ANNEX 1, *Data Handling* GE with OpenPDS, *DB Anonymizer* GE with ARX Data Anonymization Tool, *IdM GCP* GE (OAuth implementation) with the Spring Security OAuth extension and *Content Based Security* GE with an open source alternative FileSender.

5. Preliminary Trial Results

A separate report will be prepared on the trial results and published in October 2015. This chapter outlines the preliminary trial results that have already been achieved and outlines the next results expected. More details on the preliminary results can be found in the WP's individual Dx.3.2 deliverables.

5.1 WP1 Preliminary Trial Results

WP1 presented different energy optimisation results and an interesting commercial opportunity linked to optimisation of loads connected to the district cooling grid in Malmö during the Review meeting in Malmö, May 2014.

E.ON is currently exploring the opportunity to process customer consumption data (district heating and districting cooling) for 500+ buildings (i.e. data for all customers within E.ON's district heating and district cooling systems throughout Sweden) with the GEs *BigData Analysis* and *Orion Context Broker* within the WP1 infrastructure. The data is made available to the WP1 trial infrastructure through the E.ON Backend solution, see Figure 1. As the consumption data is very sensitive it may not be available via the FINESCE API. The ambition is to start processing the data later on in 2014. Anonymising the data, or at least customers IDs, would be required and hosting of the GEs within E.ON infrastructure might be necessary.

5.2 WP2 Horsens Preliminary Trial Results

In the Horsens trial, the software platform for collecting data from the equipment in the houses in Stenderup is working as planned. There are two initial applications for the data collected:

- A monitoring application to support and maintain the equipment in the houses. The user interface gives an overview of the status of the houses' equipment and indoor comfort.
- The data from the individual houses is shared with the residents in the house to inform them about the energy consumption. This is a key element in the effort to change the residents' perception of energy and their pattern of energy consumption. The data is presented to the residents using a web tool called eButler from the company Saseco – a local SME. Data is received by eButler using a preliminary version of the FINESCE API.

Controlling the consumption patterns in a smart way is a key objective of the trial. Until the start of September 2014, the focus has been on understanding the residents' energy consumption patterns and their attitude towards energy consumption. Especially the introduction of solar cells (own production) has resulted in a significant focus on consuming electricity when the sun is shining. During September 2014, the first experiments controlling the consumption (initially the charging of the electric vehicles) took place.

5.3 WP2 Madrid Preliminary Trial Results

In the Madrid trial good progress has been achieved in integrating the three GEs and the two DSEs shown in Figure 2. Integration with the Weather Forecasting component is nearly finished. The next milestone will be the completion of the integration of the GEs with Wireless Sensor Network component and the deployment of additional metering infrastructure in the trial, allowing more fine-grained measurements of electricity consumption, which can be shown to the trial building end users in order to increase awareness about their energy consumption patterns at their working place. The last step will be to complete the integration of the Microgrid and Building Control Centre components.

5.4 WP3 Preliminary Trial Results

FPL is loading data from the WP1 trial site and aggregating the data over time and by category. The data concerned is heat and electricity power consumption history, measured and forecast outside temperature, and the price of energy in the region. A web application for monitoring the power consumption has been created.

The RWTH simulation of the VPP is working stand-alone with a simulated VPP. The next step is to connect it to the real VPP monitoring and control platform via the VPP cloud infrastructure.

5.5 WP4 Preliminary Trial Results

In the Terni trial site, Smart Meters have been deployed, a technological stack based on open standards and publicly available as FIWARE GE implemented and the DSO ASM Terni can benefit from a DSO dashboard (delivered by using the WIRECLOUD GE) giving:

- Near real-time load profile (consumption and production);
- Minimise voltage drops and power losses thanks to energy marketplace;
- Attempt to correlate between weather and metering data in order to forecast consumption and production and so link it to the energy marketplace;

As regards the implementation of a market mechanism based on Demand Response aiming at minimizing power losses due to the reverse power flow effect, a dashboard customised for each different marketplace actor has been implemented, allowing the Aggregator to create Issue Resolution Plan(s), the DSO to monitor the network status and to approve or reject the Aggregator's proposal(s) and the Retailer to develop plan(s) including a list of incentives/disincentives for its own customers based on the content of the Issue Resolution plan(s).

5.6 WP5 Preliminary Trial Results

<u>COS Response Time</u>

An important output of the EV Integration Trial is a realistic estimate of the response time of the Charging Optimisation System. The faster the response time the more valuable the Charging Optimisation System would be in providing grid stability. If a system response time of under a few seconds could be achieved it could be used to support grid frequency stability, potentially the most economically valuable area of control, providing an alternative to spinning reserves in conventional power stations. The system response times can be broken down into a number of latency elements: the IT latency; 4G network ingress time, 4G network flight time, and EVSE latency.

The IT latency is the time it takes the Charging Optimisation System's software and hardware to process a charging interruption request and to issue the interrupt command to the EVSEs. This value, currently estimated at 500 ms, will be measured during the trial. One uncertainty regarding this value is the level of security required and the impact of security on latency.

The 4G network ingress time, is the time it take for the ingress port of a 4G network to absorb the large volume of interrupt commands. It is estimated at 150 ms, based on calculations and simulations.

Network flight time: initial results gives an average one way latency for sending a command to an EVSE of about 140 ms. This figure will be further refined during the course of the Trial.

EVSE latency: the response of the EVSEs to an interrupt command is estimated to take 300 ms.

Therefore the current estimate of the overall system latency is just over one second which is very satisfactory.

<u>EVSE Interworking with EVs</u>

A significant issue with interworking between the EVSE and electric vehicle was found during testing: when a charging interrupt command was sent to the EVSE, this operates a circuit breaker in the EVSE which interrupts the power. Following the interruption it was discovered that the EVs were entering a "sleep mode" from which they could not be awoken by restoring power. Instead the EVSE power cable had to be manually removed and replaced. Clearly this is not a procedure that would work in the field. The source of the issue was reviewed by the ESB and the supplier of the EVSEs and the initial conclusion is that the current international standards for EVSE-EV interaction using the PWM protocol predicated this behaviour, so that the EVSE and EVs were behaving normally and to standard, indicating a gap in the standard.

Network Simulation

A simulation of the SERVO system using a data-driven approach based on Artificial Neural Networks to forecast the grid operating conditions following implementation of schedules proposed by the COS has shown that such a SERVO system can enable DSOs to effectively assess the impacts of different EV charging schedules (proposed by the COS) on their distribution systems and determine whether these schedules can be directly approved or need to be modified.

API User Tests

The part of the COS DSE API available to SMEs has been tested internally in WP5.

6. Continuing Impact of Trials after FINESCE Project End

FINESCE is currently in the middle of its main work: the GE/DSE Integration and Trial Implementation tasks which runs from project month M4 to M24 and M10 to M26 (April 2015), respectively. Additionally the FINESCE Open Call has been run, resulting in new project partners joining, and FINESCE will support use of its trial site infrastructures by means of the FINESCE API. However, as the FINESCE project will end in September 2015, the question arises of what parts of the FINESCE trial site infrastructures will continue to be provided beyond that date and what the plans of the FINESCE partners are to further exploit their FINESCE results. Answers to these questions are given below in Table 7. More details of the individual WPs plans the the individual partners' plans can be found in the Dx.3.2 reports.

	Can SMEs use the WP's API after Sept 2015? Data available?	Product development based on FINESCE work or further development of trial site.
WP1	Platform will continue to be available. WP1 API will be up and running some few months after September, 2015: it depends partly on the level of interest experienced during Q1 through Q3 2015. Historical data will be available.	E.ON platform will continue to be developed after FINESCE as promising commercial aspects have been found. E.ON plans further product development based on the FINESCE trial site and results.
WP2 Horsens	Historical data will be available through general FINESCE repository	Platform will continue to be used and developed.
WP2 Madrid	Yes	Yes. trial subsystems, especially the BCC, are expected to undergo further development beyond the FINESCE project scope
WP3	Data will continue to be available (anonymised energy generation data from VPP, data from Smart Factory: visualisation service from XLAB)	Further product development: yes. Factory site will remain on-line.
WP4	Yes, access to the trial site infrastructure with integrated GEs as well as access to the data gathered during the FINESCE project will be allowed via the FINESCE API (investigations of safety, security and data protection issues will be needed case by case and terms for use of GEs may possibly need to be agreed with the trial site owners).	Yes, both for marketplace product development by means of FIWARE GEs and for Terni trial by implementing real demand response after the end of FINESCE project.
WP5	Yes, the trial site will be maintained for 6 months after September 2015.	Trial site maintained for 6 months at least, subject to agreement with SMEs. Commercial exploitation under consideration.

Table 7:- Overview of WPs' Plans for after September 2015

7. Conclusion

The trial site architectures have been defined, the basic trial infrastructures developed and the integration of GEs and DSEs into the trial infrastructures has made excellent progress and is on track to be completed in early 2015 according to plan.

FINESCE has put substantial effort into GE selection and evaluation. The evaluation results have been documented in project deliverables, the GE Evaluation Wiki and through working with the GE developers and give rich, detailed information on FINESCE's experiences with the GEis.

Twenty seven GE instantiations have already been integrated into trials, which means that these GEis have been given a rigorous field testing in FINESCE. The GEis which have been integrated in more than one trial in FINESCE include the *Orion Publish/Subscribe Broker, Proton Complex Event Processing, Cosmos Big Data Analysis, Wirecloud Application Mashup* and *GCP Identity Management* GEs: these GEs have been the most useful in FINESCE. Six of the seven trial sites have already integrated GEs. Three of the new partners who have joined FINESCE have integrated GEis or plan to do so. The GEi integration activity is still ongoing.

Each individual trial site has its own particular uses of GEs. However, the GEs used in FINESCE reflect this the FINESCE trials are broadly concerned with securely gathering and processing data from equipment and sensors in buildings, grids and cars. The GEs that have been of most interest to FINESCE have come from the Data/Context Management, IoT and Security chapters. GEs from these chapters are being given an intensive field trial in FINESCE.

FINESCE plans to produce 22 DSEs, of which ten are already integrated. For 12 of these, the source code will be available as open source. All DSEs will be published by end of Q1 2015.

The trial sites of WP1, WP2 Madrid, WP3 Factory, WP4 and WP5 will continue after the project ends in September 2015. In the case of WP5, the period of continuance is limited to some months.

The following trial sites plan to continue to support the FINESCE API after the project ends in September 2015: WP1 (for some months), WP2 Madrid, WP3 Factory, WP4, WP5 (for six months).

The following trial sites plan to provide historical trial data after the project ends in September 2015: WP1, WP2 Horsens, WP3 VPP, WP3 FPL, WP4, WP5.

The further commercial development of products based on the FINESCE results is planned by FINESCE partners according to details given in the Dx.3.2 reports.

The trials are already in the implementation phase, where trial results are being gathered. Preliminary results have been summarised in this report, with details available in the Dx.3.2 reports. Comprehensive trial results will be presented in a subsequent dedicated deliverable.

8. List of Abbreviations

API	Application Programming Interface
DER	Distributed Energy Resources
DSE	Domain Specific Enabler
ERP	Enterprise Resource Planning
EVSE	Electrical Vehicle Supply Equipment
FI	Future Internet
GE	Generic Enabler
I2ND	Interfaces to the Network and Devices
IaaS	Infrastructure as a Service
ICT	Information and Communication Technology
IoT	Internet of Things
KPI	Key Performance Indicator
NGSI	Next Generation Services Interface
ODBC	Open Database Connectivity
OPST	Optical Packet Switch and Transport
SME	Small and Medium-Sized Enterprises
VPP	Virtual Power Plant
VTN	Virtual Tenant Network
WP	Work Package

ANNEX 1 Formal GE Evaluation Criteria

The criteria used for the formal GE Evaluation, based on the ISO/IEC 25010:2011 specification are given in Table 8 below. Each criterion is given a weighting factor ("WF") according to its perceived importance. The GE is evaluated by entering a score (5 = Outstanding, 4 = Good, 3 = Satisfactory, 2 = Poor, 1 = Unsatisfactory) per criterion, which is multiplied by WF to give the weighted score per criterion.

Category/Criteria	Category/Criteria Explanation	Rank	WF
Functional suitability	degree to which enabler provides functions that meet stated and implied needs when used under specified conditions.		
Functional completeness	degree to which the set of functions covers all the specified tasks and user objectives.	M	5
Functional correctness	degree to which enabler provides the correct results with the needed degree of precision.	M	5
Terms and Conditions	degree to which Terms and Conditions of usage of enabler fulfil Phase 2 and Phase 3 needs. E.g. is product available under Open Source conditions?	M	5
Performance efficiency	performance relative to the amount of resources used under stated conditions. Resources can include other software products, the software and hardware configuration of the system.		
Time behaviour	degree to which the response and processing times and throughput rates of enabler, when performing its functions, meet requirements	M	5
Resource utilisation	degree to which the amounts and types of resources used by enabler, when performing its functions, meet requirements.	M	5
Capacity	degree to which the maximum limits of enabler meet requirements. Parameters can include the number of items that can be stored, the number of concurrent users, the communication bandwidth, throughput of transactions, and size of database, scalability.	M	5
Compatibility	degree to which enabler can exchange information with other enablers, systems or components, and/or perform its required functions, while sharing the same hardware or software environment.		
Co-existence	degree to which enabler can perform its required functions efficiently while sharing a common environment and resources with other products, without detrimental impact on any other product.	D	3
Interoperability	degree to which enabler can exchange information with other enablers, systems, or componentsand use the information that has been exchanged	D	3
Standards compliance	degree of alignment with existing de-facto standards like REST, XML, JSON etc.	D	3
Usability	degree to which a enabler can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use.		
Appropriateness recognisability	degree to which users can recognise whether enabler is appropriate for their needs from initial impressions of the enabler and/or any demonstrations, tutorials, documentation.	D	3
Learnability	degree to which users can learn to use enabler with effectiveness, efficiency, freedom from risk and satisfaction in a specified context of use. Includes the availability of, and quality of, product documentation and training material.	D	3
Quality of Documentation	quality of the documentation.	M	5
Operability	degree to which enabler has attributes that make it easy to operate and control.	D	3
Support for Implementation	degree and quality of support provided to users implementing with the enabler.	M	5
Reliability	degree to which enabler performs specified functions under specified conditions for a specified period of time.		

Category/Criteria	Category/Criteria Explanation	Rank	WF
Maturity	degree to which enabler meets needs for reliability under normal operation.	M	5
Availability	degree to which enabler is operational and accessible when required for use.	M	5
Fault tolerance	degree to which enabler operates as intended despite the presence of hardware or software faults	D	3
Recoverability	degree to which, in the event of an interruption or a failure, enabler can recover the data directly affected and re-establish the desired state of the system.	D	3
Security	degree to which enabler protects information and data so that persons or other products or systems have the degree of data access appropriate to their types and levels of authorisation.		
Confidentiality	degree to which enabler ensures that data are accessible only to those authorised to have access.	M	5
Integrity	degree to which enabler prevents unauthorised access to, or modification of, computer programs or data	M	5
Non-repudiation	degree to which actions or events can be proven to have taken place, so that the events or actions cannot be repudiated later.	D	3
Accountability	degree to which the actions of an entity can be traced uniquely to the entity.	D	3
Authenticity	degree to which the identity of a subject or resource can be proved to be the one claimed.	M	5
Maintainability	degree of effectiveness and efficiency with which enabler can be modified by the intended maintainers. Modifications can include corrections, improvements or adaptation of the software to changes in environment, and in requirements and functional specific		
Modularity	degree to which enabler is composed of discrete components such that a change to one component has minimal impact on other components.	O	1
Reusability	degree to which enabler can be used in more than one system, or in building other systems or enablers.	O	1
Testability	degree of effectiveness and efficiency with which test criteria can be established for enabler and tests can be performed to determine whether those criteria have been met.	O	1
Portability	degree of effectiveness and efficiency with which enabler can be transferred from one hardware, software or other operational or usage environment to another.		
Adaptability	degree to which enabler can effectively and efficiently be adapted for different or evolving hardware, software or other operational or usage environments. Adaptability includes the scalability of internal capacity.	O	1
Installability	degree of effectiveness and efficiency with which enabler can be successfully installed and/or uninstalled in a specified environment. E.g. can the product be downloaded and installed locally?	M	5
Replaceability	degree to which enabler can replace another version of the enabler for the same purpose in the same environment. Replaceability of a new version of a software product is important to the user when upgrading.	M	5

Table 8: Formal GE Evaluation Criteria

ANNEX 2 Details of Comparison of GEs to Competitors

Details of the comparisons performed by individual WPs are given below.

2.1 WP1

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%	

2.2 WP2

2.2.1 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.

2.2.2 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.

2.3 WP3

2.3.1 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[2], Riak[3], Scalaris[4] 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 an 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.

[2]Memcached: http://memcached.org/
[3]Riak: http://basho.com/riak/
[4]Scalaris: https://code.google.com/p/scalaris/

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 the conditions for triggering the push are met. This is a functionality that could be provided with solutions such as the Java Message Service (JMS)[5] or one of the Advanced Message Queuing Protocol (AMQP)[6] (e.g., RabbitMQ[7]).
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.
Even if this feature set appears to be limiting compared to the more general solutions, the GE is very powerful when implementing an IoT application. It provides features that would need to be custom implemented if using the the competitors. Since it is a GE, the administrators can easily deploy and use it, making it readily available to accept entities and context updates.
The use of open standards, namely the FIWARE NGSI-9 and NGSI-10 has further advantages. Rather than using a binary protocol which 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 to 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.
Finally, instances of the Publish/Subscribe Context Broker - Orion may be used as an effective web interface, where the data consumer may choose from a passive receiving of the data or active polling. 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 provides different entity attribute availability. This is useful when, for instance, an openly available access point should provide only some of the attributes from the full set, while the latter is available only from a protected network.

2.3.2 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[8]
- Cornell University[9]
- UC Berkeley[10]

[5]Java Message Service: http://www.oracle.com/technetwork/java/index-jsp-142945.html
[6]Advanced Message Queuing Protoco: http://www.oracle.com/technetwork/java/index-jsp-142945.html
[7]RabbitMQ: http://www.rabbitmq.com/
[8]Stanford Stream Data Manager: http://infolab.stanford.edu/stream/
[9]Cayuga: http://www.cs.cornell.edu/bigreddata/cayuga/

- Karlsruhe Institute of Technology[11]

They all offer prototypes or ready to use software packages for the described tasks of analyzing incoming events in real-time.
On the commercial market, also different vendors offer their solutions, e.g.:

- Microsoft StreamInsight[12]
- JBoss Drools[13]
- SAP Event Stream Processor[14]
- IBM Active Middleware Technology[15]
- Software AG Apama Analytics & Decisions Platform[16]
- Esper[17]

The last software also builds the foundation for the GE Gateway Data Handling. Which shows another 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 a cloud based architecture, whereas the Gateway Data Handling is meant for on-premise deployment and a pre-filtering of data. However, the FIWARE documentation gives not a clear point when to switch from one solutions 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 etc. Hence, the distinct advantage of both GEs is their embeddedness into FIWARE and the surrounding software tools such as the Publish/ Subscribe Context Broker.

[10]TelegraphCQ: http://telegraph.cs.berkeley.edu/

[11]SpoVNet: http://www.spovnet.de/

[12] Microsoft StreamInsight: http://technet.microsoft.com/de-de/library/ee362541(v=sql.111).aspx

[13]JBoss Drools: http://www.drools.org/

[14] SAP Event Stream Processor: http://scn.sap.com/community/developer-center/esp

[15] IBM AMT: http://www.research.ibm.com/haifa/dept/services/papers/amt_fact_sheet.pdf

[16] Software AG Apama: http://www.softwareag.com/corporate/products/bigdata/apama_analytics/overview/

[17]Esper: http://esper.codehaus.org/

2.4 WP4

We have performed a comparison between FIWARE GEs and products available on the market offering the same functionalities .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”.

8.1.1 FIWARE GEs alternatives on the market

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 dist, HortonWorks’s Hadoop dist, MapR’s Hadoop dist, EMC-spinoff PIVOTAL, IBM InfoSphereBigInsights
WIRECLOUD Application Mashup GE	Mashup [18] platforms such as: iMashup, iGoogle, Apache Shindig, Apache Rave, Apache Wookie& Cordova
ORION Context Broker GE	Message Brokers - Distributed publish-subscribe Messaging Systemsuch as: Redis, RabbitMQ, Apache Kafka Apache ActiveMQ, and Kestrel[19],
PROTON Complex Event Processing GE	Event Processing Software [20]such as: Oracle Event Processing, TibcoStreambase, Esper, Drools and IBM Infosphere
OBJECT STORAGE GE	Cloud storage products/services based on CDMI such as those listed in[21]
KEYROCK Identity Management GE	Different implementation of the OAuth2 standard[22]
IDAS (Backend) Device Management GE	IoT Device Management products such as: Oracle’s Internet of Things platform, Axeda Ready M2M, Device Cloud by Etherion and Wind River® Intelligent Device Platform XT

Table 9 - FIWARE GEs vs. alternative products/services

[18]Mashup - http://en.wikipedia.org/wiki/Mashup_(web_application_hybrid)

[19] 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

RabbitMQvs Kafka - http://www.quora.com/RabbitMQ/RabbitMQ-vs-Kafka-which-one-for-durable-messaging-with-good-query-features

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

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

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

8.1.2 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).

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)	X	X	X	X	X	X	X	X
no hardware to be purchased	X	X						
free usage (no costs per hour / no licensing costs)	X	X	X	X	X	X	X	X
cloud based (available as SaaS) + no hardware to be purchased + no installation activities to be carried forward (GE Global instance)		X		X		X		X
openness: no vendor lock in		X	X	X	X	X	X	X
no need of extreme hardware configuration		X				X		X
comes with native integration with other GEs, thus being part of an ecosystem		X	X	X	X		X	X
cloud based: installable from image available on a catalogue (coming with a default configuration) on a cloud-based VM (GE private instance)			X	X	X			
natively integrated in the FIWARE cloud infrastructure						X	X	

Table 10 - "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.

2.5 WP5

2.5.1 Data Handling GE

A comparison was undertaken between the Data Handling GE and the OpenPDS [deMontjoye] [23].

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 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.

[23] 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

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

2.5.2 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[24]. 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
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.

2.5.3 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.

[24] 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

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.

2.5.4 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.