



Grant Agreement No.: 723174

Call: H2020-ICT-2016-2017

Topic: ICT-38-2016 - MEXICO: Collaboration on ICT

Type of action: RIA



D3.1: SmartSDK Reference Models and Recipes

Revision: v.1.0

Work package	WP 3
Task	Task 3.3
Due date	28/02/2017
Submission date	28/02/2017
Deliverable lead	MARTEL
Version	1.0
Authors	Federico M. Facca (MARTEL), Tomas Aliaga (MARTEL), Jose Manuel Cantera Fonseca (TID)
Reviewers	Silvio Cretti (FBK), Daniele Pizzolli (FBK)

Abstract	This deliverable documents the initial activities reference architectures and data models for cloud native smart applications based on FIWARE Data Management and IoT Enablement Generic Enablers.
Keywords	Architecture Patterns, Data Models, FIWARE, Data Management, IoT Service Enablement

Document Revision History

Version	Date	Description of change	List of contributor(s)
V1.0	27/02/2017	Final version implementing reviewers comments	Federico M. Facca (MARTEL), Tomas Aliaga (MARTEL)
V0.9	24/02/2017	Version ready for review	Federico M. Facca (MARTEL), Tomas Aliaga (MARTEL)
V0.8	20/02/2017	Data Model section	Federico M. Facca (MARTEL), Jose Manuel Cantera Fonseca (TID)
V0.7	19/02/2017	Recipes	Tomas Aliaga (MARTEL)
V0.6	18/02/2017	Introduction and conclusion	Federico M. Facca (MARTEL)
V0.5	16/02/2017	Review of figures	Tomas Aliaga (MARTEL)
V0.4	15/02/2017	SmartSDK Architecture patterns	Tomas Aliaga (MARTEL)
V0.3	15/02/2017	Methodology and overview of FIWARE IoT and Data GEs	Federico M. Facca (MARTEL)
V0.2	01/02/2017	Cloud Architecture patterns	Tomas Aliaga (MARTEL)
V0.1	15/01/2017	Table of Content	Federico M. Facca (MARTEL)

Disclaimer

The information, documentation and figures available in this deliverable, is written by the SmartSDK (A FIWARE-based Software Development Kit for Smart Applications for the needs of Europe and Mexico) – project consortium under EC grant agreement 723174 and does not necessarily reflect the views of the European Commission. The European Commission is not liable for any use that may be made of the information contained herein.

Copyright notice

© 2016 - 2018 SmartSDK Consortium

Project co-funded by the European Commission in the H2020 Programme		
Nature of the deliverable:		R
Dissemination Level		
PU	Public, fully open, e.g. web	✓
CI	Classified, information as referred to in Commission Decision 2001/844/EC	
CO	Confidential to SmartSDK project and Commission Services	

* R: Document, report (excluding the periodic and final reports)

DEM: Demonstrator, pilot, prototype, plan designs

DEC: Websites, patents filing, press & media actions, videos, etc.

OTHER: Software, technical diagram, etc.

EXECUTIVE SUMMARY

SmartSDK provides a set of ready to use “recipes” to develop smart applications in the Smart City, Smart Healthcare, and Smart Security domains. Such recipes are based on: components (i.e. Generic Enablers and Specific Enablers), data models (i.e. NGSI formalisation of the data exchanged among components) and reference architectures (i.e. the combination of components and data models to support production grade requirements).

The work documented in this deliverable pave the way for the development of recipes and data models by providing guidelines for their development. More, it starts documenting ongoing developments for the definition of “shared” (i.e. not specific to a given application scenario) recipes and data models.

As regards the recipes, SmartSDK starts by analysing FIWARE Reference Architecture for Data/Context Management and IoT Service Enablement and how Cloud Architecture Patterns (such as High Availability and Scalability) can be applied to FIWARE Reference Architecture. Starting from the most relevant GE for the development of smart applications in SmartSDK, the deliverable documents how their deployment architecture can be modified to support high availability and scalability. Starting from the above deployment architecture and the analysis of best tools to implement it, defines concrete recipes for Docker Compose that allow for deploying and managing such architectures in few clicks. Finally, planned steps for the future development of recipes are presented.

Concerning the data models, SmartSDK analyse the current FIWARE data models in light of SmartSDK application scenarios: for each application scenario developed in the project, we identify with existing FIWARE data models that will be reused in SmartSDK. The re-usage of the data models will ensure the validation/revision of existing data models in the context of the EU-Mexico collaboration, and will allow SmartSDK to contribute to such data models with novel reference data sets. Starting from the reused data models, SmartSDK also defines a list of needed data models (which development is ongoing) to satisfy the application scenario requirements. Finally, planned steps for the future development of data models and related tools are presented.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	3
TABLE OF CONTENTS	4
LIST OF FIGURES	5
LIST OF TABLES	6
ABBREVIATIONS	7
1 INTRODUCTION	8
1.1 Design Principles	9
1.2 Structure of the deliverable.....	10
1.3 Audience	10
2 SMARTSDK ARCHITECTURE PATTERNS.....	11
2.1 Overview FIWARE Data/Context Management and IoT Services Enablement architecture patterns	11
2.2 Overview of cloud architecture patterns.....	13
2.2.1 Basic concepts related to distributed systems.....	13
2.2.2 Scalability Pattern.....	15
2.2.3 High Availability Pattern.....	15
2.2.4 Multi site pattern.....	15
2.2.5 Colocate Pattern.....	16
2.3 Methodology.....	16
2.4 SmartSDK Architecture Patterns	17
2.4.1 Context Broker.....	17
2.4.2 Comet Short Term History (STH)	19
2.4.3 Cygnus	19
2.4.4 IDAS	20
2.4.5 CKAN	21
2.5 Realization of Architecture Patterns recipes in SmartSDK	22
2.5.1 Foundation	23
2.5.2 Documenting and Developing Architecture Patterns recipes	25
2.5.3 SmartSDK Architecture Pattern recipes	26
2.5.4 Roadmap for SmartSDK Recipes	28
3 SMARTSDK DATA MODELS.....	31
3.1 Overview of existing FIWARE data models	31
3.2 Documenting and Developing Data Models.....	33
3.3 Novel SmartSDK Data Models	34
3.3.1 Smart City	34
3.3.2 Smart Security	35
3.3.3 Smart Health	35
3.4 Tools	35
3.5 Roadmap	37
4 CONCLUSIONS AND FUTURE WORK.....	38
REFERENCES.....	39

LIST OF FIGURES

Figure 1: SmartSDK's cookbook concept.....	8
Figure 2: SmartSDK Design Principles.	9
Figure 3: Data/Context Management and IoT Service Enablement Southbound-Northbound Architecture.....	12
Figure 4: Availability and Partition Tolerance.....	14
Figure 5: Consistency and Partition Tolerance.....	14
Figure 6: HA Deployment.	15
Figure 7: Multisite deployment.....	16
Figure 8: Queue-Centric Workflow.	16
Figure 9: High Available and Scalable Orion Context Broker Architecture Pattern.....	18
Figure 10: High Available and Scalable Comet STH Architecture Pattern.....	19
Figure 11: Cygnus Architecture Pattern.	19
Figure 12: High Available and Scalable Cygnus Architecture Pattern.....	20
Figure 13: IDAS Architecture Pattern.....	20
Figure 14: High Available and Scalable IDAS Architecture Pattern.	21
Figure 15: CKAN Architecture Pattern.	21
Figure 16: High Available and Scalable CKAN Architecture Pattern.	21
Figure 17: Docker Swarm.	24
Figure 18: Architecture of the MongoDB replicaset recipe.	26
Figure 19: Architecture of the MongoDB replicaset recipe.	27
Figure 20: Orion Context Broker in HA Docker Recipe.....	28
Figure 21: Data model portal on Schema.org.	36

LIST OF TABLES

Table 1: Recipes Roadmap Overview.29
Table 2: FIWARE data models by Scenario33
Table 3: Data models Roadmap Overview.37

ABBREVIATIONS

API	Application Programming Interface
DNS	Domain Name Server
GE	Generic Enabler
GEri	Generic Enabler reference implementation
HA	High Availability
HTTP	Hypertext Transfer Protocol
IoT	Internet of Things
IP	Internet Protocol
REST	REpresentational State Transfer
TCP	Transmission Control Protocol

1 INTRODUCTION

The creation of common principles is the base for any activity that aims at scaling beyond its initial promoters. In the case of Smart applications, i.e. applications applying intelligent decisions on top of acquired data, such common principles can be translated to common components (i.e. APIs), common architectures, and common data models. FIWARE is now reaching maturity and its adoption is spanning beyond Europe. The ability to leverage on the work done in Europe in other context, such as Mexico, is crucial to ensure that models and architectures defined are enough general to cover needs and regulations in the different countries worldwide. This process, also often referred to as “de-facto” standardisation, is crucial to facilitate the “replication” of platform and application deployments based on FIWARE.

While FIWARE so far mostly focused on the APIs definition and their basic interaction, it is now time to support wider adoption by promoting:

- ➔ Development of reference data models to be used in different smart scenarios, ensuring interoperability across different applications (tackling the same data) and facilitating the deployment of applications across different sites.
- ➔ Development of reference architectures that support the production grade deployment of applications. Such architectures need to deal with important requirements, such as High Availability and Scalability.

The work carried on in SmartSDK largely falls in this context. SmartSDK is the FIWARE’s “cookbook” for developing smart applications in the Smart City, Smart Healthcare, and Smart Security domains. SmartSDK is looking into applications developed so far within Europe and Mexico (using FIWARE or alternative Open Source technologies), analysing them and using them to create a cookbook for developing applications in the Smart City, Smart Healthcare, and Smart Security domains.

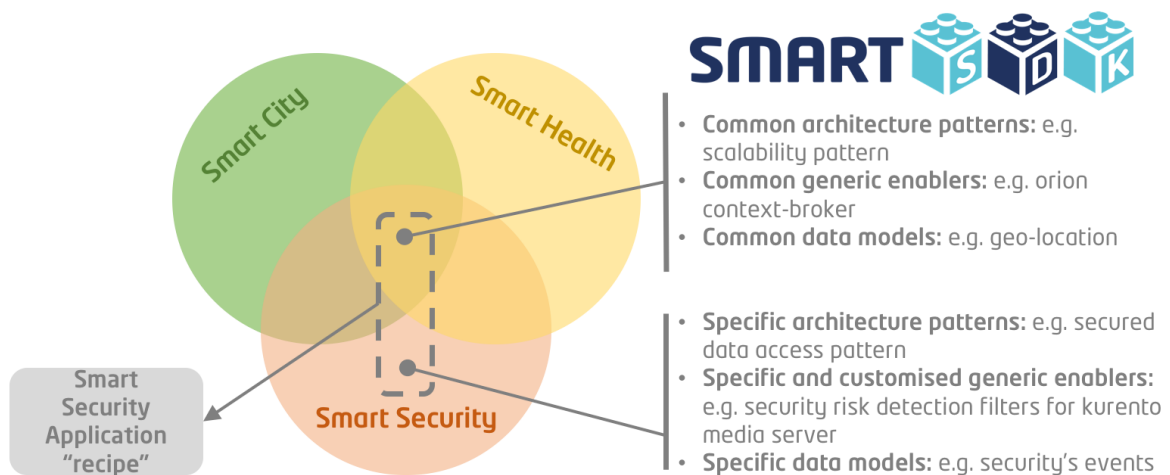


Figure 1: SmartSDK’s cookbook concept.

The “cookbook” is based on: a set of architecture patterns (i.e. the basic cooking processes), a set of Generic Enablers (i.e. the basic ingredients) and a set of data models (i.e. the spices and flavours binding the ingredients through the cooking process, depicted in Figure 1).

While application scenario specific activities are carried out in the Application work package (WP2), this deliverable provides general reference architectures and data models for the development of smart applications at scale, based on a set of design principles derived from FIWARE best practises. This deliverable provides guidelines for the development of new recipes and new data models. Future activities are outlined for both recipes and data models.

1.1 Design Principles

SmartSDK follows the design principles of FIWARE and brings it to its full potential, by moving from single enabler architectures to those which include multiple enablers. To this aim, SmartSDK focuses the fundamental glue among enablers: data models and architectures.



Figure 2: SmartSDK Design Principles.

- ➔ **Restful APIs.** Restful APIs are programmable interfaces that exploit REST¹ (REpresentational State Transfer) architectural style to produce and consume data in a lightweight communication infrastructure. REST, which typically runs over HTTP (Hypertext Transfer Protocol), has several architectural characteristics: 1) Decouples consumers from producers; 2) Stateless existence; 3) Able to leverage a cache; 4) Leverages a layered system; 5) Leverages a uniform interface. All FIWARE Enablers offer Restful APIs, and SmartSDK, in contributing to FIWARE, will follow this design principle.
- ➔ **Reusability and Openness.** Software development activities often face problems already met by other developers before. Thus, the ability to capitalize on these developments is crucial to allow developers to focus on their core goals rather than having to build the foundation of their software before focusing on the core specific features of their applications. Reusability is at the core of FIWARE design: Generic Enablers are reusable software components, that developers may compose as Lego® bricks to build their application. Beyond that, through the evolution process from a European project to an Open Source community with its own governance, FIWARE embraced the Open Source principles to support the wider possible diffusion and impact of its achievements. Beyond the initial Open specifications, as of today all official Generic Enabler reference implementations are available as Open Source (mostly adopting the Apache License v2) and Open Data support has been included in FIWARE (through CKAN GE). SmartSDK will follow the same path and contribute to the FIWARE community the newly and enhanced developed Enablers. Beyond that it will contribute added value on top of them through: open and reusable reference architecture implementation for Smart applications based on FIWARE, open and reusable reference data models for Smart applications encoded using NGSIv2.
- ➔ **Cloudification and Microservices.** Nowadays new services are embracing the cloud paradigm. In this paradigm, services can be self-provisioned by users or developers and their configuration is automatically managed to guarantee their scalability and fault tolerance. FIWARE made the development of Enablers to support cloud hosting one of its core activities. Indeed, FIWARE Lab itself is based on the FIWARE Cloud Hosting Chapter and, from FIWARE Release 4, Generic Enablers are available as Dockerized services. Even though the term “microservices” has been around from some years, with the advent of cloud-native applications it gained momentum. The idea of decomposing the traditional monolithic business systems into small, independently deployable services fits well with the objectives of a cloud architecture of promoting scalability, automated deployment, decentralized control, fault tolerance and resilience. The SmartSDK reference architecture will be inspired by the microservices architectural style and will take advantage of the nature of the FIWARE Generic

¹ https://en.wikipedia.org/wiki/Representational_state_transfer

Enablers that can be considered as atomic units of services that still deliver a business value. SmartSDK will complete this picture by providing deployable bundles of Generic Enablers that implement cloud native patterns for the development of Smart applications.

- ➔ **Market and community relevance.** Following the launch of third phase of the FI-PPP programme, FIWARE roadmap became more and more driven by market needs and by the FIWARE adopters' feedbacks. Feedbacks from the FIWARE Accelerator programme have been a key instrument to improve the quality and market readiness of FIWARE and influenced the development of features and functionalities of FIWARE platform. Today FIWARE, with the launch of the FIWARE Foundation, becomes as a matter of fact an Open Source community where different companies and single developers can contribute following the principles regulating the community. SmartSDK will plug into the FIWARE community and go through their validation process to submit to the mainstream FIWARE community its contributions.

1.2 Structure of the deliverable

The deliverable is structured as follow:

- ➔ Section 2 overviews the SmartSDK Architecture Patterns, i.e. a set of basic patterns that allow to deploy production grade Smart applications, by combining FIWARE Reference Architecture and Cloud Architecture Patterns. The section discusses how cloud architecture patterns can be applied to the different core components of FIWARE Data Management and IoT service enablement chapters. It also provides an example of a “recipe” implementation based on Docker.
- ➔ Section 3 presents plans for the adoption and development of FIWARE data models in SmartSDK. The section starts with a presentation of existing FIWARE data models and discuss which ones will be adopted (and hence validated) in SmartSDK application scenarios. The section overviews guidelines for the development new data models and discusses novel required models by the SmartSDK applications scenarios in relation to existing standards.
- ➔ Section 4 summarizes the plan for future work in relation to architecture patterns and data models in the context of SmartSDK.

1.3 Audience

This deliverable is mainly intended for:

- ➔ Developers and Operators interested into deploy FIWARE Smart applications in a production context.
- ➔ Developers and Knowledge modellers interested into adopting FIWARE data models or contributing to the initiative.

2 SMARTSDK ARCHITECTURE PATTERNS

When a software platform reaches maturity, developers look for a reference architecture to adopt it. A reference architecture provides blueprints on how to plumb the different services composing a given platform. Still, the same reference architecture can be implemented in different ways depending on the purpose of the deployed platform. For example, for a test environment you may not need the platform to scale up to thousands of requests per second, neither to be resilient to hardware failures. Concretely, this means that you can easily, for your test environment, deploy your whole architecture in a single (virtual) server.

Later, when you move the platform into production, you face additional challenges: services should be resilient to hardware failure, services should be able to scale with the number of requests, and so on. A single server is not more suitable to host your platform: you will not be able to serve requests beyond its maximal capacity and you can be sure that sooner or later it will have a failure. i.e., to provide a production level platform, you need to distribute the services that compose architecture of your platform.

Distributing an architecture introduces some complexity related to the consistent management of its services. The introduction of cloud computing resulted in the simplification of deployment of distributed architectures and in the definition of the so called “cloud architecture patterns” that help developers improve the architecture of their platform to benefit from cloud computing capabilities.

FIWARE is now becoming a mature platform and developers adopting it are starting to face the problem of deploying its services into production over cloud infrastructures. While Generic Enablers reference implementations (GERis) can be easily provisioned over the Cloud (thanks to their availability as virtual machine images and Docker images), composing them do provide a resilient and scalable platform is not documented in detail.

One of the goals of SmartSDK is simplifying life of developers by providing them easy to reuse architecture patterns, that not only allow them to deploy complex architectures made of different IoT Management and Data/Context Management enablers, but to do so in such a way that the deployment is dealing with the different production grade requirements.

This section first provides an overview of basic FIWARE IoT Management and Data/Context Management reference architecture (i.e. without considering production requirements). Then it introduces a set of Cloud architecture patterns and discuss how such patterns can be used to extend FIWARE reference architecture to support production requirements over cloud infrastructures. Finally, it discusses how SmartSDK implements such patterns in reference recipes.

2.1 Overview FIWARE Data/Context Management and IoT Services Enablement architecture patterns

The Data/Context Management and IoT Services Enablement chapters contain the most relevant components to build so called Smart services. Thus, the work of SmartSDK focuses on them. The components of the two chapters, contrary to other ones, are also well interchained due to the adoption of a common interface among them: FIWARE NGSIv2 [1].

The components of the two chapters are meant to be combined into a southbound (i.e. data/context producers, including IoT Devices) - northbound architecture (i.e. data processors of NGSI harmonised data). The central element in this picture, is the ContextBroker that allows data processors to access in different ways to data produced by sensors and other data producers.

SmartSDK's Figure 3 presents the different components of the Data/Context Management chapter and of the IoT Service Enablement chapter as a southbound- northbound architecture. Of course, developers, depending on their needs, can select to adopt a given subset of components in their application.

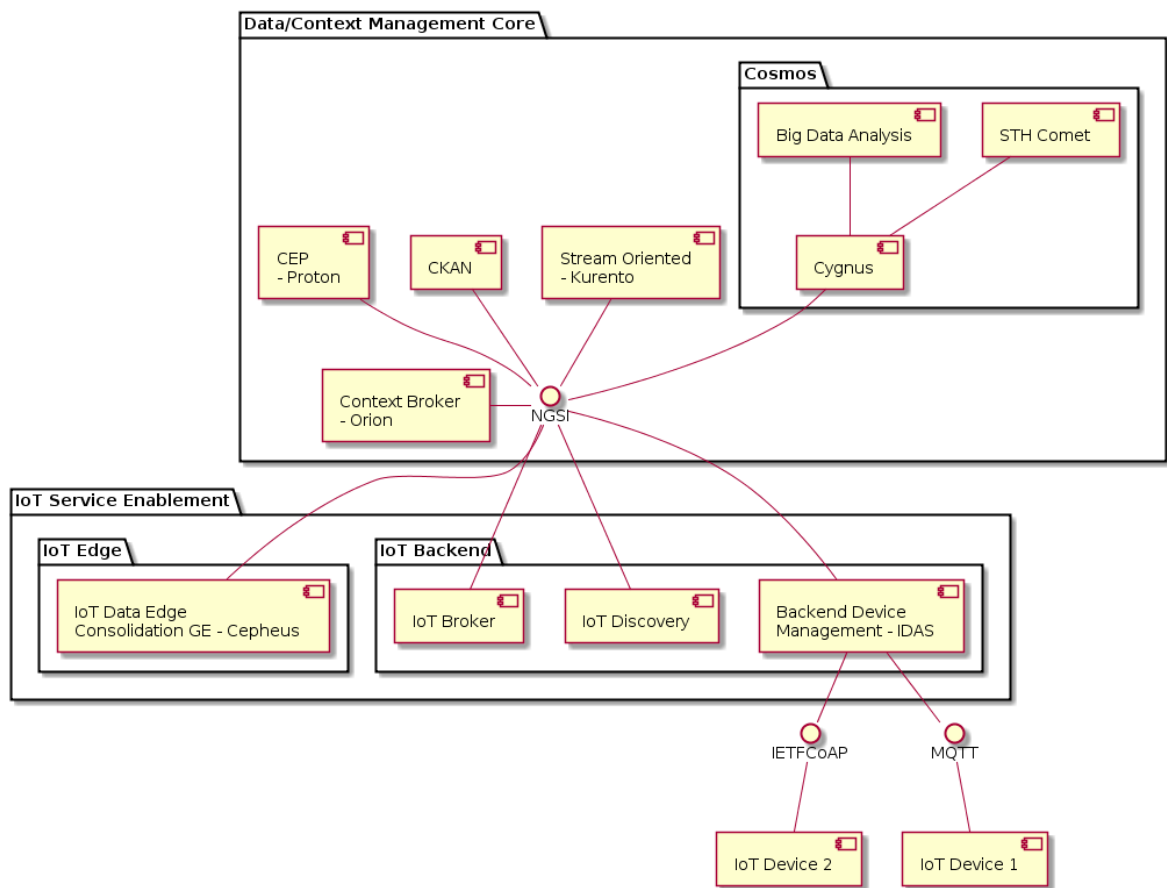


Figure 3: Data/Context Management and IoT Service Enablement Southbound-Northbound Architecture.

Core components of the Data/Context Management² chapter include:

- ➔ Context Broker - Orion, the central element of the Data Context Management architecture, provides a publish/subscribe service for context data.
- ➔ Big Data Analysis - Cosmos, the solution for the analysis of NGSI data sets, made of different tools including the short time historical data storage (STH Comet) and the integration with different datastores (relational databases, big data filesystems, etc.) thanks to the adaptation capabilities provided by Cygnus.
- ➔ Stream Oriented - Kurento, an NGSI integrated multimedia processing server.
- ➔ CKAN - a repository for Open Data sets.
- ➔ Complex Event Processing (CEP) - Proton, an event processing tool that identify patterns over NGSI data and generate response over identified patterns.

The IoT Services Enablement³ chapter include the following core components:

- ➔ Backend Device Management - IDAS, a set of IoT Agents which allow to register IoT devices and to transform collected data into NGSI compliant format.
- ➔ IoT Data Edge Consolidation GE - Cepheus, a solution for edge processing and aggregation of sensor data NGSI compliant.
- ➔ IoT Broker, an NGSI middleware that support the aggregation of data from multiple sensors.

² <https://catalogue.fiware.org/chapter/datacontext-management>

³ <https://catalogue.fiware.org/chapter/internet-things-services-enablement>

- ➔ IoT Discovery, an NGSI middleware that support the discovery of context producers by context consumers.

It is worth mentioning that some of these services are meant to run only as Global Instance of FIWARE Lab. This is the case for example of Big Data Analysis Cosmos. As such, it is outside of the scope of SmartSDK to design High Available (HA) patterns for Cosmos (that has a matter of fact is HA and scalable in his essence being based on Apache Hadoop⁴). Rather relevant, instead is working on simplifying the connectivity of other GEs to FIWARE Global Instances as in the case of Cosmos.

2.2 Overview of cloud architecture patterns

Throughout the history of software development projects, specific software engineering practices have been proved successful to solve certain recurring problems, particularly, problems related to hardware limitations, business risks or performance requirements. When developers started noticing they were applying the same kind of solutions for certain types of problems, these solutions were abstracted from the project details and taken as “Patterns”.

Patterns are therefore well-defined, well-known and maybe most importantly, reusable solutions for common software development problems. They can emerge from the lower level of software implementation details, as in the case of Design Patterns, or from the software deployment structure and the relationship of its main components. We will refer mostly, though not exclusively, to the latest form, commonly known as Architecture Patterns.

The following is a simple summary of the most relevant patterns in the context of cloud applications. Its purpose is to give a better understanding of the patterns by exemplifying some of them. Interested readers of this subject can refer to literature such as Design Patterns [3] or Cloud Architecture Patterns [4].

2.2.1 Basic concepts related to distributed systems

As mentioned in the introduction of Section *SmartSDK Architecture patterns*, moving from a centralised to distributed architecture poses some challenges. Most of the challenges relates to the fact that often a service is not *stateless* and distributing the service on multiple nodes implies also that the status must be distributed across the nodes.

2.2.1.1 Stateless versus Stateful services

A service is *stateless* if at any point in time the value of its output(s) depends only on the value of the input(s). A service is *stateful* if at any point in time the value of the output(s) depends on the value of the input(s) and of an internal state.

In cloud computing platforms, stateless services are ideal for scaling horizontally⁵ and for high-availability, since they can be deployed n-times without any interdependency; whereas stateful services are more suitable for scaling vertically⁶. High availability and horizontal scaling, in the case of stateful services, require ways to recover or synchronize the status of the system.

2.2.1.2 CAP Theorem

In the case of distributed stateful services, solutions to synchronize the status of the service across all the nodes need to be applied to guarantee that the distributed service acts coherently as a single service. In the context of a reliable and resilient distributed system, we can identify the

⁴ Readers interested to learn about deploying Hadoop in High-Availability can refer to [2].

⁵ Horizontal scalability is the ability to increase the capacity of a service by running more instances of the same service.

⁶ Vertical scalability is the ability to increase the capacity of a service by increasing the number of hardware resources (e.g. CPU, RAM) allocated to the service.

following desirable properties:

- ➔ Consistency: All nodes composing the service have the same status at the same time.
- ➔ Availability: The service can provide a reply at any given time.
- ➔ Partition tolerance: The service continues to work even though some of the nodes composing it are not reachable temporarily or definitively.

Unfortunately, Brew's conjecture [5], better known as CAP Theorem, states that a distributed system can have simultaneously only two out of three of the above properties. While of course the theorem is valid under certain definition of Consistency, Availability, and Partition Tolerance, it is generally valid. As a matter of fact, distributed systems are built on networks and networks aren't completely reliable. This implies that distributed systems must tolerate partition. According to the CAP theorem, this means we are left with two options:

- ➔ **Availability and Partition Tolerance.** Most of the data replication solutions on the market today adopt Availability and Partition Tolerance as properties and implement mechanism to ensure that the system reconciles in a short time, minimizing the time window where an observer would find the status of the system inconsistent. These systems are generally called *eventually consistent*.




Figure 4: Availability and Partition Tolerance.

- ➔ **Consistency and Partition Tolerance.** Solutions based on Consistency and Partition Tolerance, basically implements atomic reads and write. The system will return a response only when the global status is consistent (and hence if the system is not partitioned). If for any reason a node is not reachable, the system will return no response (or an error). I.e, the system is *eventually available*.




Figure 5: Consistency and Partition Tolerance.

2.2.2 Scalability Pattern

The main principle of the Scalability Pattern is that applications should be prepared to easily adapt to a sudden change in its demand of the resources. Scalability is therefore a measure how well an application can adapt to, for example, attend more users.

If an application is not truly scalable, having more hardware resources at disposal will not necessarily make it run faster and better, allowing for example more users. Ideally, applications should be developed in such a way that its scalability is ultimately bounded only by hardware resources. This means, you can attend more users by simply buying more hardware without the need of rewriting your application. This is because software development tends to be more expensive than hardware acquisition, particularly in the cloud.

2.2.3 High Availability Pattern

Imagine a Bank was running their Home Banking service on a single server and suddenly, something bad happened to that server. For example, the computer overheats and turns down. If the application was not prepared to react to this scenario, Bank users would instantly lose access to the service.

The main idea behind this pattern is that applications should be resilient to hardware failures and have either an extra instance of the application ready to cover the demands of the failed instance or enough resources in the running instances to distribute among them the load the failed instance was attending. Figure 6 shows an example of a deployment of an App which must attend 10M users. Since each server can serve up to 5M users, the deployment on 3 servers can be considered in HA with tolerance to 1 server failure.

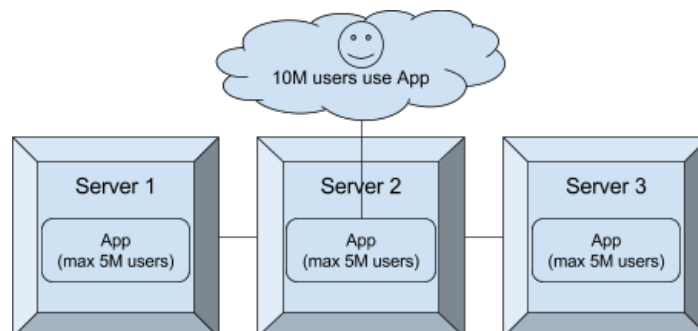


Figure 6: HA Deployment.

2.2.4 Multi site pattern

Imagine that Youtube had its services running only in California. If this was the case, users in the other side of the world would experience a terribly slow service because of the distance data would need to travel in the network.

In order to give all users a rather similar quality of service and response times, the pattern (shown in Figure 7) suggests that software should actually be deployed in servers located in different points of the globe. Of course, a synchronization mechanism among sites should be adopted so that the content is shared and seen equally everywhere.




Figure 7: Multisite deployment.

2.2.5 Colocate Pattern

This pattern complements the Multi-Site, and suggests that services that usually work together depending on each other should be deployed to the same server so as to reduce the delays produced by communication. Packet transmission times in network connecting different servers or even worse different data centers can introduce several performance losses.

Queue-Centric Workflow Pattern

The idea of this pattern, illustrated in Figure Figure 8, is that whenever there are two main components of a software that need to communicate with each other, and both have different performance requirements, it is good practice to place a queue between them to handle asynchronous messages and let each part run at its pace.




Figure 8: Queue-Centric Workflow.

2.3 Methodology

So far, we have listed some FIWARE Architecture Patterns, i.e. how FIWARE services can be combined to define a FIWARE reference platform for IoT and Data Management, and Cloud Architecture Patterns, i.e. how services can benefit from cloud architecture to provide higher scalability and resiliency.

The next step is to combine the two patterns so as to apply Cloud Architecture Patterns to FIWARE Architecture Patterns and provide solutions for deploying FIWARE IoT and Data Management in a scalable and resilient way.

Each FIWARE GERi has its own architecture, so this means that there is not one single solution to make it scalable and resilient. Each FIWARE GERi needs to be analysed in order to understand how to apply Cloud Architecture Patterns to its architecture.

The methodology we follow in the definition of the “SmartSDK Architecture Patterns” (i.e. the patterns that combine FIWARE and Cloud Architecture Patterns) is based on a simple principle:

composability. This principle simply means that the definition of patterns is guided by the idea that we break down a larger pattern in smaller ones, so that recurring patterns can be easily reused. For example, MongoDB is used in different FIWARE GERis as backend solution, so we can define a specific pattern for MongoDB that then can be reused by other patterns, for example the ones that deals with High Available Context Broker.

The methodology works as follows:

1. Analyse the GErI architecture and identify its components;
2. For each component identified, understand if it is stateful or stateless;
3. For each component identified, apply Cloud Architecture Patterns according to its type (stateful vs stateless) and supported solutions for distributing its state across multiple nodes;
4. Document the architecture pattern of the different components composing the GErI;
5. Compose the overall architecture pattern for the GErI from the sub-patterns;
6. Translate the defined patterns into deployable recipes.

Of course, it may not be possible to apply Cloud Architecture Patterns to all GErI. Some GErI may have architectures that are not suitable because, for example, do not provide ways to expose part of the status and share it across multiple nodes. It is outside the scope of SmartSDK to modify the code of such GERis to enable the application of Cloud Architecture Patterns to them. Nevertheless, the analysis will allow to provide feedbacks to GErI owners regarding the limitations of their current architecture.

2.4 SmartSDK Architecture Patterns

Many of the ideas explained so far regarding Architecture Patterns can be applied to different FIWARE Generic Enablers. In this section, we will highlight how some Generic Enablers would benefit the application of the patterns and which considerations must be considered for the applicability of those. In particular, we have chosen the most used generic enablers in the SmartSDK applications scenarios, which are also among the most popular GEs in the FIWARE community.

2.4.1 Context Broker

The Context Broker lies at the core of the Data Management chapter of the FIWARE Architecture and its reference implementation is called Orion. Orion is basically composed of two complementing parts, namely a front-end and a backend.

The frontend is the Orion service implementing the NGSI REST API providing users features for queries and subscriptions/notifications for context. This frontend is stateless, in the sense that there is no state being persisted on the service instance and if an instance dies, no data is lost.

On the other hand, the backend is the stateful part of Orion that gives persistence to its data, up to a certain point. It is usually implemented as a MongoDB database running in a different process and known to Orion as the “db”.

2.4.1.1 Scalability

A simple Orion configuration with one service in the frontend and one in the backend will be able to attend a certain maximum number of users. This limit, known as “capacity” is usually dictated by both hardware and software characteristics of the deployment. Moreover, such scenario is not well-suited for failure resiliency as explained in previous sections, in the sense that if one of the parts fails, the whole service falls apart.

To increase the capacity of the service, or scale it down (unused resources in the cloud usually mean wasted money), the administrators can scale both frontend and backend separately. However, there are

important considerations to keep in mind.

The frontend, as a stateless service, can be scaled up or down without worrying about data losses. The only thing to keep in mind is that those services going down should not have been directly exposed to the public using individual details (e.g specific IPs). This is to avoid users attached to the fallen service not being able to reach the service anymore. To avoid this, a common practice, as shown in Figure 9, is to put a load balancer in-front of all the frontend services to act as an entry-point whose connection details remain the same. Another DNS-based solution, would be dynamically associating the IPs of the frontend to a specific Orion service name clients will interact with.




Figure 9: High Available and Scalable Orion Context Broker Architecture Pattern.

Nevertheless, scaling the backend requires a more sophisticated approach, because data should be the same across all instances to guarantee users always see the same set of data regardless of the service responding their requests. For example, when data is added to one of the instances, this new information should be replicated to the rest of the instances as soon as possible.

In general, changes in the data introduce inconsistency into the system that needs to be sorted out. To solve this, distributed algorithms based on the quorum principle exist so that instances can reach consensus and keep the same data across the dataset. These algorithms, however have certain requirements that should be met. For example, in MongoDB version 3.0 there is a maximum of 50 members out of which a maximum of 7 can be voters. Also, there must be an odd number of voting members in order to have a tie-breaker in the case of tied elections. Further details can be read in the Replication Manual of MongoDB⁷.

2.4.1.2 High Availability

An administrator of the Orion service, knowing the capacity of the current service and having the ability to scale up both frontend and backend, will be able to run the Orion service in High-Availability. There are, though, two more things required to do so.

Scaling up the frontend services would be useless without the ability to actually distribute the incoming requests among them. This is the main task of the Load Balancer, an essential element in the HA architecture.

Also, in the case of failures, if the services keep on failing one after another at some point there will be no more redundancy of services and Orion will be down. Consequently, a reconciliation mechanism instrumented to recover the failed instances will be necessary. This could be either manual (i.e, an administrator re-launching a failed instance) or a software-based automated solution.

2.4.1.3 Multisite

This pattern becomes relevant only for huge deployments of Orion across many geographical areas.

⁷ <https://docs.mongodb.com/manual/replication/>

The benefits for simpler or smaller deployments of Orion will not outweigh the challenges of this kind of deployment.

The multisite deployment of the frontend might not be technically as challenging as the backend, but still both would require significant efforts. The frontend would require configuring load balancers to send requests to different data-centers based on physical proximity of the users to the Orion instances. The backend, on the other hand, would require a robust replication mechanism in order to keep data consistency while performing within reasonable time limits, considering the distances data will have to travel.

2.4.2 Comet Short Term History (STH)

Comet is a component of the FIWARE ecosystem in charge of storing and retrieving historical raw and aggregated time series information about some context data. For example, entity attribute values registered in an Orion Context Broker instance.

Comet interacts with an NGSI data generator such as Orion Context Broker via notification mechanisms and requires a database, typically a MongoDB, to store raw and aggregated data from these notifications.

In a scenario consisting of heavy loads of notifications, having a single instance of Comet to attend and process all the data could easily become a bottleneck for the overall architecture. It is therefore desired to have a receptive frontend with a distributed load of notifications as show in Figure 8 below.




Figure 10: High Available and Scalable Comet STH Architecture Pattern.

Clearly, due to the architectural similarities with the aforementioned Context Broker, the same mentioned patterns could be applied also to Comet STH.

Regarding the database component shown in the figure, it could be either a single database instance or a single logical database composed of a physical replicaset of instances, as explained in section 2.5.3.1.

2.4.3 Cygnus

Cygnus is a connector in charge of persisting certain sources of data in certain configured third-party storages, creating a historical view of such data. In particular, we will be working with the “cygnus-ngsi” version of Cygnus. The simplest deployment of Cygnus is depicted in the image below.




Figure 11: Cygnus Architecture Pattern.

As it can be seen from the Figure 11, the main component of the architecture is the Agent, whose source consumes notifications from an NGSI origin such as Orion and dispatches those events to

channels, each of which is connected to some sort of sink. Cygnus recipes should simplify for users the process of their agents deployment with different types of sinks.

As explained in its official documentation⁸, Cygnus provides users the possibility to configure its internal architecture applying patterns such as the Queue-centric workflow to improve its performance for some specific scenarios. However, in order to improve the reliability of the Cygnus deployment applying the High Availability pattern, it is required to deploy multiple agents at the same time. There could be a recipe whose purpose would be exactly this one, allowing users to easily deploy multiple agents, as show in Figure 12 below.




Figure 12: High Available and Scalable Cygnus Architecture Pattern.

2.4.4 IDAS

IDAS is the reference implementation of the Backend Device Management Generic Enabler⁹. Its main usage is the Typical IoT use-case Scenario I: Common Simple Scenario of the overall FIWARE IoT architecture¹⁰.

As of now, IDAS can be seen as a group of IoT Agents supporting different protocols on top of different transport options. Some of the supported protocols for the different available Agents are UL2.0, JSON, OMA-LWM2M and SIGFOX. It could be said though, that an average Agent deployment architecture typically looks like Figure 13.




Figure 13: IDAS Architecture Pattern.

To improve its resilience, multiple instances of the internal services (IoTAgent and MongoDB) could be deployed in the way presented in Figure 14.

⁸ <https://fiware-cygnus.readthedocs.io/en/latest/architecture/index.html>

⁹ <https://forge.fiware.org/plugins/mediawiki/wiki/fiware/index.php/FIWARE.ArchitectureDescription.IoT.Backend.DeviceManagement>

¹⁰ [https://forge.fiware.org/plugins/mediawiki/wiki/fiware/index.php/Internet_of_Things_\(IoT\)_Services_Enablement_Architecture#Typical_IoT_use-case_Scenarios_.28I.29:_Common_Simple_scenario](https://forge.fiware.org/plugins/mediawiki/wiki/fiware/index.php/Internet_of_Things_(IoT)_Services_Enablement_Architecture#Typical_IoT_use-case_Scenarios_.28I.29:_Common_Simple_scenario)

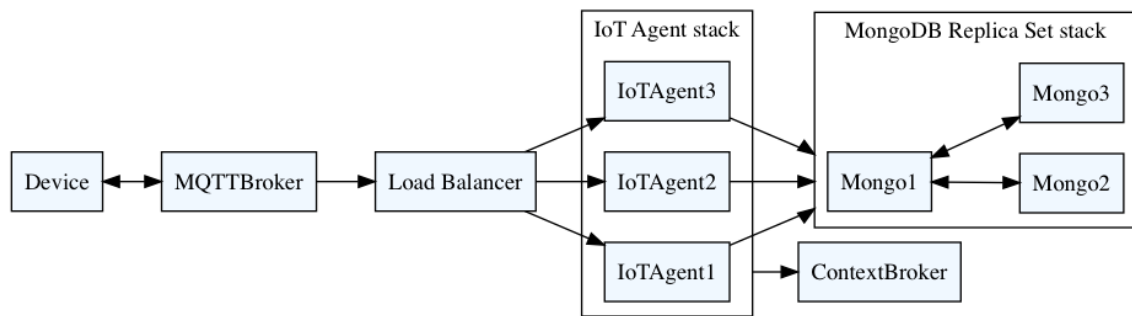


Figure 14: High Available and Scalable IDAS Architecture Pattern.

2.4.5 CKAN

A basic deployment of the CKAN Generic Enabler consists of the deployment of group of different services, some of which are stateless whereas others are stateful. The diagram in Figure 15 shows the architecture of a simple deployment of these services.




Figure 15: CKAN Architecture Pattern.

The scalability of the CKAN GE could be approached first with a multiple-instance deployment of the stateless services, namely CKAN and Datapusher. A load balancer in front of each stateless service would distribute requests among the multiple replicas of the service. This would make the architecture look something like in Figure 16 below.

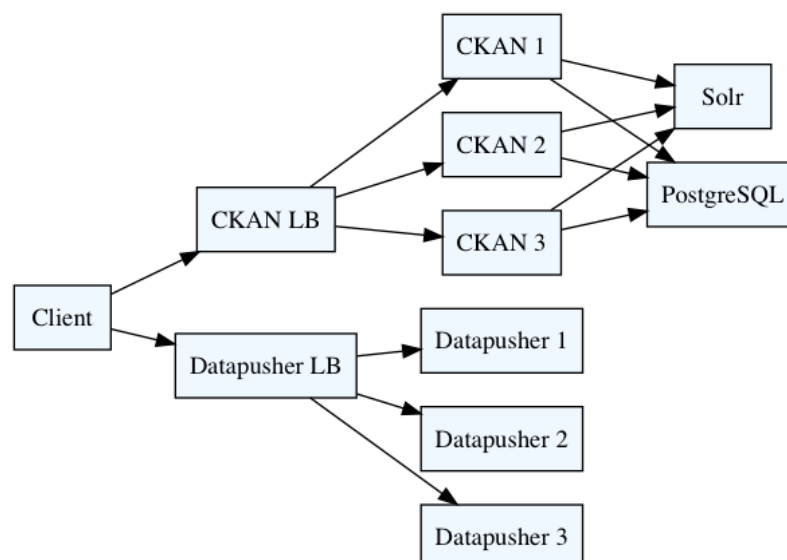


Figure 16: High Available and Scalable CKAN Architecture Pattern.

The replication of the stateful services “Solr” and “PostgreSQL” would require a more careful

configuration. Some hints about how to run a distributed solr are mentioned in the official Solr's Docker image page¹¹ whereas the replication of the "PostgreSQL" database is covered in the replication wiki page¹².

The applicability of different Architecture Patterns to other Generic Enablers not listed above is still under study. If time allows within the scope of the project, we might include recipes for some other useful Generic Enablers. In this case, the priority will be given to those Generic Enablers not mentioned in this document but used in the development of the SmartSDK application examples of the Smart Healthcare, Smart City and Smart Security scenarios.

2.5 Realization of Architecture Patterns recipes in SmartSDK

As shown in previous section, Architecture Patterns are powerful tools that can be applied to make software applications more robust. In line with its goal of simplifying and improving the development of FIWARE Applications, SmartSDK aims to include these patterns in its toolbox in the form of recipes.

FIWARE Applications are typically developed with the usage and composition of different FIWARE Generic Enablers (GE). A simple look at the documentation of complex GEs such as Cygnus¹³ or even at the composition of different GE such as the cases mentioned in the page for Data Management Architecture shows that there is room for applying software architecture patterns.

GEs are relatively complex pieces of software running as services, developed in either one or, most likely, multiple programming languages. As such, these services require a fair amount of configuration and pieces from other systems to work. Furthermore, many times services need to be re-deployed to different places, and hence must be able to run in different platforms. All these challenges are not specific to FIWARE but common in the software industry. To overcome these difficulties, services are nowadays being containerized.

There are many technologies supporting containerization nowadays, the most prominent one in the market being Docker. On top of the containerization technology, orchestration systems are built in order to control and manage the lifespan of containers. Orchestration typically involves scheduling, launching and monitoring containers running in different nodes of a cluster.

There are many orchestration technologies in the market nowadays, each with its pros and cons, and the most popular are Kubernetes, Swarm and Mesos/Marathon. Among the three, Swarms stands out as the simplest and easiest to use. Moreover, it is developed by Docker, so it integrates quite well with the tooling of the Docker ecosystem, making it the easiest to adopt for people already working with Docker containers. It is true though that both Mesos/Marathon and Kubernetes are both older and more robust systems than Swarm in terms of offered features for containers orchestration and scaling capabilities. Nevertheless, it is also well known that this power comes with the cost of bigger complexity and harder learning curves.

Docker has been the chosen container solution in the FIWARE community. In line with this decision, and the previously mentioned advantages, SmartSDK decided to adopt all the technologies of the Docker ecosystem, including Swarm. SmartSDK acknowledges some features are still missing in Swarm, as it will be pointed out in later sections of this document. However, considering Docker's fast growth^{14,15}, we see the distance with its competitors is being shortened.

The following subsection will help the reader understand how the patterns will come to life with the underlying usage of Docker-based technologies.

¹¹ https://hub.docker.com/_/solr/

¹² https://wiki.postgresql.org/wiki/Replication,_Clustering,_and_Connection_Pooling

¹³ <http://fiware-cygnus.readthedocs.io/en/latest/architecture/index.html>

¹⁴ <https://www.docker.com/docker-news-and-press/docker-automates-and-democratizes-container-orchestration>

¹⁵ <https://dzone.com/articles/docker-swarm-lessons-from-swarm3k>

2.5.1 Foundation

2.5.1.1 Docker Compose

Deploying containers is easier than manually handling custom deployed services on either virtual machines or bare metal. However, when applications grow in complexity and they start being divided into multiple services, handling multiple containers can become a very challenging task.

To simplify the orchestration of containers, Docker introduces a native tool called Compose¹⁶. Among the main characteristics that make Compose attractive to the audience of SmartSDK users, we could highlight the following three.

- ➔ **Ease of use:** With Docker Compose, users can launch multiple services properly configured with the ease of a single command line.
- ➔ **Declarative approach:** Compose adopts a declarative approach for the configuration, which can be read in a simple script file. This is much easier than having to dig and mine multiple files of code to customize certain parameters of the application.
- ➔ **Can be easily reused:** This brings the power of writing once and reusing multiple time, allowing users to attend new use-cases by combining pre-scripted recipes.

Below, a simplified example of a Docker Compose file:

```
version: '2'

services:

  mongo:
    image: mongo:${MONGO_VERSION}
    command: --nojournal

  orion:
    image: fiware/orion:${ORION_VERSION}
    links:
      - mongo
    ports:
      - "1026:1026"
```

With such a file, a user can simple execute a command like “docker-compose up” and have all the stack of services up and running, showing how easy-to-use compose is.

Also, the declarative approach of the file simplifies both understanding and customization. For example, any reader could tell how many services are being deployed and which port is being used by the Orion service.

Finally, if a user wanted to use Docker Compose for their microservices-based application using Orion, they could reuse the same file removing the need to redefine Orion’s configuration in their custom file.

Compose has been adopted by many FIWARE Generic Enablers for simple scenarios such as the Tour Guide¹⁷ and will be adopted by SmartSDK for more complex recipes. However, SmartSDK requires a special set of features available only from Docker Compose version 3, as explained in the following section.

¹⁶ <https://www.docker.com/products/docker-compose>

¹⁷ <http://fiwaretourguide.readthedocs.io/en/latest/>

2.5.1.2 Docker Swarm

Swarm is Docker's native tool for defining, configuring and managing clusters of Docker hosts. It started as a standalone application but since the release of Docker 1.12 a new Swarm, known as swarmkit or swarm mode, was integrated into the Docker engine. This last one, swarm mode, is the Swarm¹⁸ this document will be referring to.

However, Compose was not ready to fully integrate with the new Swarm until the release of Docker version 1.13. Docker 1.13 is the first Docker version supporting Docker Compose files version 3, giving users the features of Docker Compose mentioned in the previous section with the power of the new swarm features explained further below. Therefore, it was decided not to support Docker versions older than 1.13 in SmartSDK. Due to this commitment to follow the state of the art of the Docker releases, as the project evolves, we may push the minimum required version to a higher value, of course, always respecting the official available releases so that users do not need to "customize" installations to beta versions.

Before we mention the new Swarm features SmartSDK is taking advantage of, it is worth explaining what a Swarm is in the first place. As it can be seen from Figure 17, a Docker Swarm is basically a collection of nodes (usually called a cluster). Nodes are machines running Docker which could be physically distributed in datacenters of different providers, such as AWS or FIWARE lab, but must be interconnected through the same network. On top of the cluster, users can therefore deploy their stacks (i.e., collection of services) and, according to the deployed stack and related constraints, Swarm will orchestrate the services part of the stack and the nodes part of the cluster, providing the needed synchronization and communication between them.

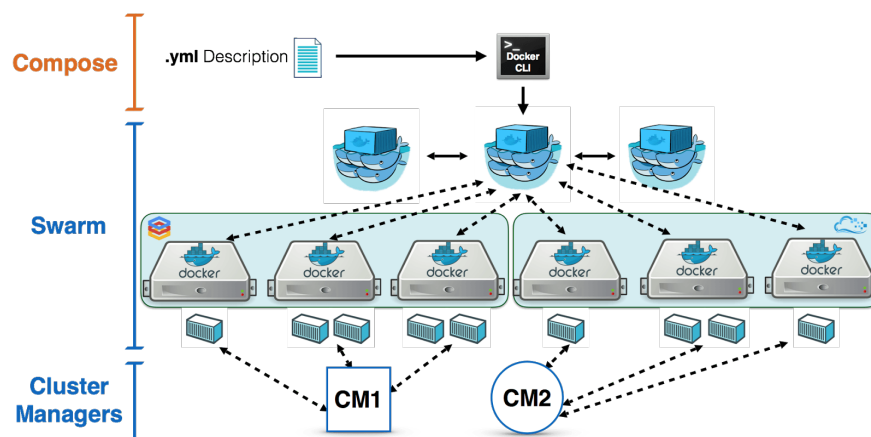


Figure 17: Docker Swarm¹⁹.

The first concept introduced by the new swarm mode is the concept of a service. The idea is that users should declare services instead of containers. A service could consist of either one or multiple containers (called tasks) and includes definitions such as which networks the containers of that service will connect to, how many replicas of the service there will be, etc.

The new mode also introduces an internal implementation of service reconciliation. This means, if the user declares that he needs 3 tasks of the Orion service, when one task fails, Docker will automatically reschedule a new task to conciliate the required status of 3 tasks.

With an internal DNS in the swarm nodes, there is no longer need to reference the deployed services containers by their IPs. Being able to reference them by name simplifies the operation of the services in the swarm, particularly considering the dynamic nature of the containers and the fact that, due to redeployments, those IPs might be subject to change.

¹⁸ <https://docs.docker.com/engine/swarm/swarm-tutorial/>

¹⁹ Image taken from <https://blog.docker.com/2015/11/deploy-manage-cluster-docker-swarm/>

When having multiple instances of a service, a virtual IP is created for the service. Swarm managers use an internal load balancer to distribute requests on that virtual IP among the IP addresses assigned to the single task instances part of a service.

Services can be deployed in different modes, but there are mainly two options: global or replicated. A global service means Docker will deploy a task for that service on each node of the cluster. The replicated mode instead will let the user define the number of replicas (tasks) for a given service and Docker will spread those replicas among the available swarm nodes. Below an extract of a Docker Compose file used in the recipe for the Context Broker in HA.

```
services:
  orion:
    image: fiware/orion:1.3.0
    ...
    deploy:
      replicas: 3
    networks:
      - mongo-replica_replica
```

Moreover, Docker provides a mechanism to adjust the deployment strategy when particular requirements should be met. This is achieved with a combination of the use of the “placement” feature of the Docker Compose file with the labelling mechanisms for swarm nodes. The following is an extract example of how a service is declared to be executed only in manager nodes running Ubuntu 14.04.

```
...
placement:
  constraints:
    - node.role == manager
    - engine.labels.operatingsystem == ubuntu 14.04
```

This “constraints” feature could be helpful for example in the applicability of a multi-site pattern in the deployment of a service, where the location of service replicas should be constrained to a specific set of nodes of a multi-site cluster.

Nevertheless, it is worth mentioning that though very powerful, Docker technologies are not a silver bullet to attend all the requirements of SmartSDK. There are still some challenges and open issues to be tackled. This is partly due to the fact that Compose and Swarm are relatively new technologies and, even though they have a really fast pace of development, there are still much work in progress, as it can be seen in their Github Issues page²⁰. We will discuss some examples of these requirements, when needed, in the following sections.

2.5.2 Documenting and Developing Architecture Patterns recipes

SmartSDK recipes will be developed mainly in the form of Docker Compose files with supplementary script files when needed. All these files will be kept in a GitHub repository along with the corresponding documentation for developers and users. The repository is already available and can be explored at SmartSDK GitHub [6]. Moreover, the official documentation is available online [7].

²⁰ <https://github.com/docker/docker/issues>

Recipes are organized in different folders, mainly reflecting their relation to different components of the FIWARE Architecture. For example, at the root of the folders structure there is currently one folder for the Data Management chapter, where recipes for Generic Enablers such as Orion are placed. Likewise, there is a different folder called IoT Management intended to host the recipes for the IoT stack generic enablers.

Each folder must contain the following elements:

- ➔ An introduction explaining the purpose of the recipe, explained in a Markdown²¹ file. For example, which GE or combination of GEs will be used by the recipe.
- ➔ Links to proper pieces of external documentation. For example, instead of explaining details of the Generic Enablers, links to the corresponding FIWARE documentation sites should be provided. It is not a good practice to repeat documentation because it is error prone and makes content more likely to become outdated.
- ➔ A high-level overview diagram explaining the main ideas of the recipe, if not already covered by the previous point. For example, Figure 18 shows the high-level overview for the MongoDB replicaset recipe.

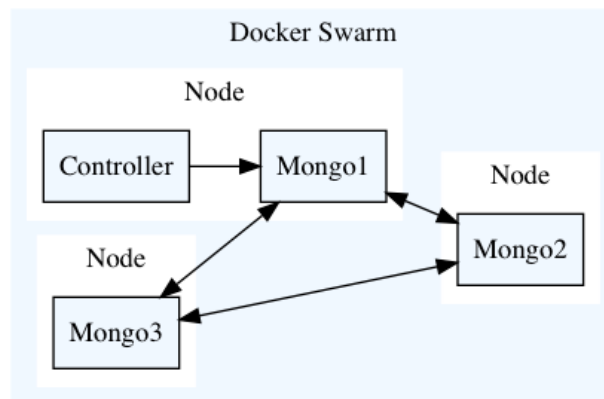


Figure 18: Architecture of the MongoDB replicaset recipe.

- ➔ Instructions on how to use the recipe, and eventually customize it. Since each user might desire different kind of customizations, it may not always be possible to cover all customization options. However, something like a walkthrough showing why default values work as they do will let users understand how to implement the changes according to their needs.
- ➔ The recipe itself and all complementary files required for proper use.

Furthermore, the repository will contain required instructions to create and share either new recipes or enhancements of the existing one. This way, developers can improve the SmartSDK repository by sharing their expertise and tools. Guidelines will also be developed to keep an order in the toolbox. For example, stating which documentation tools to use, standard folders structure and file formats, etc. All of this will be developed under a “How to Contribute” section of the documentation.

2.5.3 SmartSDK Architecture Pattern recipes

At least during the first phases of the project, SmartSDK recipes will be aimed at mainly two of the FIWARE Architecture Chapters, namely, the Data Management and Internet of Things Services Enablement. A third category will group miscellaneous recipes that may or maybe not be strictly related to FIWARE but that are of great utility for the development of the rest of the recipes. This section will briefly explain the first recipes that were developed as well as outline some of the foresaw

²¹ <https://daringfireball.net/projects/markdown/>

applicable recipes for different generic enablers of the mentioned chapters.

It is worth noting that this is not an exhaustive list and new recipes might be developed throughout the lifespan of the project. Moreover, recipes should be designed in a lean and modular way so that they can be combined to create new ones. For example, reusing the recipe for a MongoDB replicaset, users can create a recipe to use the Orion Context Broker storing its data in that replicaset. Thus, it becomes unfeasible to list in this document all possible combinations of recipes.

2.5.3.1 Orion Context Broker in HA with MongoDB replicaset

The first goal was to test a deployment of the Orion Context Broker applying the High Availability Pattern. As explained in a previous section, this task was divided into two recipes, one for the MongoDB replicaset and another with a deployment of Orion frontend.

MongoDB Replicaset

The first developed recipe was to setup of a MongoDB replicaset. The task involved deploying a service of MongoDB instances in the Docker Swarm. In the spirit of the HA pattern, the idea was to have each replicaset member automatically deployed to a different machine (node) of the swarm. This way, if a node failed, it would not take away all the replicas at once. This was achieved using the “global” deploy mode of Docker Compose, which instructs Docker to deploy one container (task) of that service in each available node. The idea is shown in *Figure 19*.

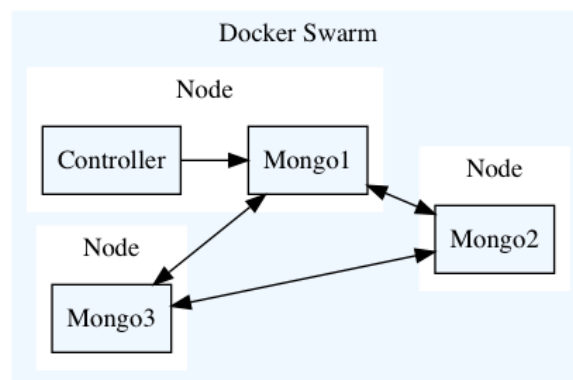


Figure 19: Architecture of the MongoDB replicaset recipe.

One of the main challenges was automating the configuration of the replicaset. Configuring a replicaset is not so complicated when you have the IP addresses of each member of the replica. But, the problem in a Docker Swarm is that by default you don’t know a priori which IPs each container will be assigned in the overlay network. Moreover, at the time being, Docker does not provide a safe introspective API²² which containers could request for these pieces of information.

To circumvent this issue, a side-car controller script was developed to retrieve these required pieces of information once everything is deployed and connect to a mongo container to configure the replicaset. For the time being, it is using the Docker API of the host daemon, a very popular workaround in the Docker community, though not ideal from a security point of view. Nonetheless, the script can be easily updated once a safer mechanism is provided by Docker. The nitty-gritty of the replicaset configuration remains basically the same. For further details please refer to the implementation files²³.

Orion in HA

The following step was to create a new recipe for the frontend of Orion Context Broker in HA mode that would reuse the MongoDB replica recipe. The scenario would then look like Figure 20.

²² <https://github.com/docker/docker/issues/8427>

²³ <https://martel-innovate.github.io/smartsdk-recipes/utis/mongodb/replica/readme/>

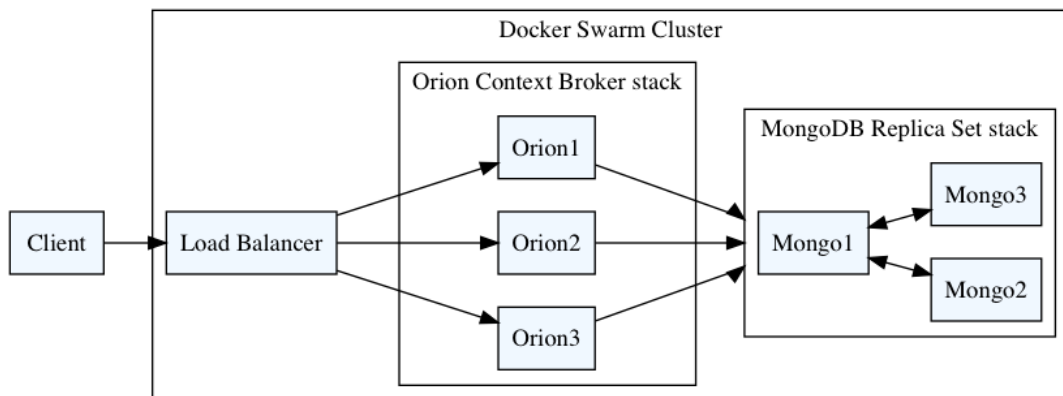


Figure 20: Orion Context Broker in HA Docker Recipe.

Client requests for Orion will arrive at Docker’s internal Load Balancer (LB). This component will make sure that requests are equally distributed to all the tasks (containers) of the Orion service. It is not that the same request arrives at all containers at once, rather, N request in a row will be dispatched to as many different containers as possible.

Thanks to Docker’s embedded DNS server, the Orion service needn’t attach to any specific mongod instance. It was linked to the “MongoDB” service, who has a virtual IP, whose traffic is then automatically directed to a final task of that service.

Users could then manually scale up the service running a simple command like “*docker service scale orion=5*” to scale from 3 (in the example above) to 5 instances of the Orion service. Ideally, the scaling action would be automated to become a function of a specific group of metrics. For example, a new replica should be deployed whenever the average response time falls below a certain threshold. However, the support for this kind of orchestration is something still not fully developed in Docker and hence the recipes will have to use operator’s defined number of replicas.

2.5.4 Roadmap for SmartSDK Recipes

This section is a high-level overview of the next recipes to be developed in SmartSDK. It is not intended to give full details of the activities nor the precise schedules, because like all roadmaps, it might be subject to minor adjustments down the road.

As a reminder, the SmartSDK project is being tracked in JIRA²⁴ and all the tasks related to the development of the recipes are identified by the JIRA component “Generic Architecture Patterns”.

Below we will now present a brief outline of the recipes planned to be developed, which can be grouped into the following three groups.

2.5.4.1 Data Management

In addition to the Orion Context Broker Recipes, for the Data Management chapter we plan to develop a couple of recipes for Cygnus and STH Comet generic enablers.

For STH, we would like to let users deploy multiple instances of the hapi service in case high levels of input from Orion notifications as well as output to User requests are needed.

As regards Cygnus, the plan is to first develop a docker-based recipe to deploy a simple cygnus scenario (one agent) with an easy-to-customize configuration to support different sinks. Then, a second recipe will be aimed to apply the High Available pattern to the the cygnus agent.

²⁴ <https://jira.fiware.org/projects/SMART>

2.5.4.2 IoT Management

For IDAS, we plan to develop a recipe to allow the deployment of Agents in HA so as to overcome failures in IoT agents. Of course, since Agents are deployed on the Edge, certain geographical constraints need to be accounted for the distribution of the replicas.

Regarding the IoT Broker, the same recipes developed for Orion Context Broker will be recreated for this smaller Context Broker suited for IoT scenarios.

2.5.4.3 Utils

This category of recipes deals with complementary recipes which are helpful to support the deployment of Data Management and IoT Management Generic Enablers deployments following Cloud Architecture Patterns.

For example, to support Orion Context Broker and IoT Agents deployment in High Availability, also MongoDB needs to be deployed in High Availability. Accordingly, the first recipe in this category tackle the support for MongoDB replicaset management.

Since these recipes are complementary to the main ones, new unplanned recipes might evolve during the development of the project. At the time being, we expect to have at least two more:

- ➔ A recipe to support the registration of NGSI data sources in the Open Data portal of FIWARE Lab.
- ➔ A recipe to support the protection of services through authorization and authentication patterns. When the authorization of a service is delegated to another service, for example using oauth, the procedures for retrieving tokens and exchanging logging information can be simplified for the developers of applications.

SmartSDK Release	Expected Date	Group	SmartSDK Architecture Pattern Recipes ²⁵
1	01/2017	Data Management	Context Broker
		Utils	MongoDB replicaset
2	04/2017	Data Management	Comet STH
			Configurable Simple Cygnus
3	07/2017	Data Management	Cygnus
4	10/2017	Data Management	CKAN
		Utils	Open Data integration
5	01/2018	IoT Management	IoT Broker
			Multiple IDAS IoT Agents
6	04/2018	Utils	Simplified Auth for Apps
7	07/2018	-	No new recipes, rather fixes and upgrades to existing ones.

Table 1: Recipes Roadmap Overview.

There are 7 planned releases throughout the lifespan of the project. Table 1 summarizes the estimated

²⁵ The different recipes, where meaningful, will cover all the different Cloud Pattern identified as relevant: HA, Scalability and Multi-site

distribution of the planned recipes among the different releases.

2.5.4.4 Testing recipes and their scalability

To understand the behaviour of the proposed reference architectures and their implementation into recipes, it is important to test the deployed recipes with proper workloads. To this aim, we developed a tool: the NGSI Load Tester²⁶. The tool will be adopted in the next future to test the recipes and measure their behaviour with the increase of load and in case of service failures.

²⁶ <https://github.com/martel-innovate/ngsi-load-tester>

3 SMARTSDK DATA MODELS

The FIWARE Community is promoting the development of re-usable and harmonized data models under the umbrella of the FIWARE data models²⁷ initiative. These data models are reusing and extending the work performed under the GSMA IoT Big Data initiative. The work conducted under FIWARE is evolving on a daily basis by taking into consideration requirements from adopters of data models. All FIWARE data models, coherently with the centric role of NGSIv2 API, are expressed to be used with such API.

SmartSDK contributes to this activity in different ways:

- ➔ Developing Applications that reuse existing FIWARE data models;
- ➔ Providing adapters to convert existing data into FIWARE data models;
- ➔ Populating the Global FIWARE Context Broker with data compliant with existing FIWARE data models;
- ➔ Proposing extensions to existing FIWARE data models according to the SmartSDK applications' requirements;
- ➔ Developing novel FIWARE data models according to the SmartSDK applications' requirements;
- ➔ Providing tools that support the FIWARE data models activity;

At this point in time, the work on SmartSDK applications is still preliminar, as such data models are not yet completed and ready to be included in the deliverable. Thus, this section mainly focuses on: overviewing existing FIWARE data models and how they are adopted in SmartSDK's applications; providing guidelines for the development of novel data models; introducing planned data models and tools to support FIWARE data models activities.

3.1 Overview of existing FIWARE data models

As early mentioned, the FIWARE data models initiative is steering the development of harmonised data models to provide a set of common models within the FIWARE community to enable data portability at the application layer.

Adoption of common models, beside the adoption of common APIs (NGSIv2 in the case of FIWARE), is a fundamental step to allow for a scalable data ecosystem that supports the interoperability and re-usage of data and applications working on top of data.

In line with this principle FIWARE has harmonised so far, the following set of data models:

- ➔ **Environment.** A model to enable the monitoring of air quality and other environmental conditions for a healthier living. In particular, covered entity types include:
 - AirQualityObserved: an observation of air quality conditions at a certain place and time.
 - WaterQualityObserved: capture all the parameters involved in Water Quality scenarios.
 - NoiseLevelObserved: represents an observation of those parameters that estimate noise pressure levels at a certain place and time.
- ➔ **Civic Issue tracking.** This set provides entity types for civic issue tracking interoperable with the de-facto standard Open311. In particular, covered entity types include:

²⁷ schema.fiware.org

- ServiceType. A type of service a citizen can request.
- ServiceRequest. A specific service request (of a service type) made by a citizen.
- ➔ **Street Lighting.** It models street lights and all their controlling equipment towards energy-efficient and effective urban illuminance. The covered entity types include:
 - Streetlight: a particular instance of a streetlight. A streetlight is composed by a lantern and a lamp. Such elements are mounted on a column (pole), wall or other structure.
 - StreetlightGroup: a group of streetlights being part of the same circuit and controlled together by an automated system.
 - StreetlightModel: a model of streetlight composed by a specific supporting structure model, a lantern model and a lamp model. A streetlight instance will be based on a certain streetlight model.
 - StreetlightControlCabinet: an automated equipment, usually on street, typically used to control a group or groups of streetlights, i.e. one or more circuits.
- ➔ **Device.** This set of entity types describes IoT devices (sensors, actuators, wearables, etc.) with their characteristics and dynamic status. The covered entity types include:
 - Device: an electronic apparatus designed to accomplish a particular task.
 - DeviceModel: the static properties common to multiple instances of a Device.
- ➔ **Transportation.** Transportation data models for smart mobility and efficient management of municipal services. The covered entity types include:
 - TrafficFlowObserved: a recorded observation of traffic flow.
 - Road: a geographic and contextual description of a Road.
 - RoadSegment: a geographic and contextual description of a road segment.
 - Vehicle: a specific vehicle instance.
 - VehicleModel: a model of vehicle, capturing its static properties such as dimensions, materials or features.
- ➔ **Indicators.** It models key performance indicators intended to measure the success of an organization or of a certain activity in which the organisation is engaged.
- ➔ **Waste Management.** This model enable efficient, recycling friendly, municipal or industrial waste management using containers, litters, etc. The covered entity types include:
 - WasteContainerIsle: the isle that holds one or more containers.
 - WasteContainerModel: a model of waste container, capturing its static properties such as dimensions, materials or features.
 - WasteContainer: a certain instance of waste container placed at a particular isle or place.
- ➔ **Parking.** This model provides real time and static parking data (on street and off street) interoperable with the EU standard DATEX II. This model includes the following entity types:
 - OffStreetParking: an offstreet parking site with explicit entries and exits.
 - ParkingAccess: an access point to an off-street parking site.
 - OnStreetParking: an on street, free entry (but might be metered) parking zone which contains at least one or more adjacent parking spots.
 - ParkingGroup: a group of parking spots.

- **ParkingSpot:** an individual, usually monitored, parking spot.
- ➔ **Weather.** Weather observed, weather forecasted or warnings about potential extreme weather conditions.

The FIWARE data models have been adopted in different applications scenarios enabling their validation, for example in different cities.

SmartSDK is developing demonstrators in three domains:

- ➔ **Smart Cities:** the planned demonstrator will provide an application that helps citizen to move around their city in environment-friendly and healthy way.
- ➔ **Smart Security:** the planned demonstrator monitor parking spots and buildings for security threats.
- ➔ **Smart Health:** the planned demonstrator collect and analyse patient behaviour data through mobile devices.

Clearly, several of the data models already part of FIWARE data models initiative are key to SmartSDK. Table 2: FIWARE data models by Scenario summarises a mapping between scenarios and FIWARE data models.

FIWARE Data Model	Scenario	Comments
Transportation	Smart City	The model will be used for tracking public and private vehicles to support computation of eco-friendly and health-friendly paths.
Weather	Smart City	The model will be used for the weather forecast to support computation of eco-friendly and health-friendly paths.
Environment	Smart City	The model will be used for air quality data to support computation of eco-friendly and health-friendly paths.
Device	Smart Security	The model will be used to define and track the security devices deployed in the building and parkings to monitor security.
Transportation	Smart Security	The model will be used to track vehicles detected by camera and sensors.
Parking	Smart Security	The model will be used to support the tracking of vehicles in parking lots.

Table 2: FIWARE data models by Scenario

On the other side, as it is clear for example for what regards the Smart Health scenario, more models are required beyond the ones already part of FIWARE data models catalogue.

3.2 Documenting and Developing Data Models

The existing FIWARE data models are documented using markdown format as detailed in the data model template²⁸, and they are formalised using JSON Schema draft version 4 [8]. The data model documentation must include:

- ➔ A general description of the data model.

²⁸ The data model template is available at: https://github.com/Fiware/dataModels/blob/master/datamodel_template.md

- ➔ A description of the attributes of the data model that are formalised in the JSON Schema.
- ➔ An example of data model instance as JSON code fragment.
- ➔ A description of reference endpoints that store data following the specified data model.
- ➔ A list of activities yet to be completed on the data model.

The data model template is complemented by a set of basic data model references that are often reused in the different data models. These include:

- ➔ Common general properties (e.g. name, data source, ...) inherited from the GSMA initiative [9].
- ➔ Common properties to define location (e.g. address, position, ...) based on GeoJSON.

Currently there is an ongoing discussion on how to better support the navigation of data model documentation. The plan is to provide FIWARE data models descriptions through a portal that follows the principles of the schema.org²⁹ initiative.

Besides providing a template for the data model development, FIWARE defined a set of guidelines and related design principles that should be followed to develop a FIWARE Data Model³⁰:

- ➔ **Syntax guidelines.** They define naming conventions for entity types or attributes. For instance, the camelcase³¹ naming pattern is recommended.
- ➔ **Reusing guidelines.** They promote reusability encouraging the authors of new data models to check existing work (eg. Openstreetmap or schema.org) and base on it when feasible.
- ➔ **Modelling guidelines.** They provide modelling design patterns to be followed when representing real world entities and attributes. They define the recommended practices on modelling dates, location, civic addresses, etc. or how to express extra attribute metadata, for instance, timestamps.
- ➔ **Representation guidelines.** They recommend, particularly, how to model and represent quantitative values. Percentage-based magnitudes, such as relative humidity must be represented in parts per one. An important aspect of a quantitative value is the unit of measurement. To this aim, the guidelines advice on the usage of the International System of Units [10] by default.

To contribute to the existing Data Model repository on Github³², developers are recommended to fork the repository in Github, add the new models or revise the existing ones and provide your contributions as a pull request to the original Github repository.

3.3 Novel SmartSDK Data Models

In addition to the models identified in Section 3.1, SmartSDK scenarios require additional data models that are not yet covered by FIWARE data models initiative. The development of the data models for the above scenarios is still ongoing, a first released will be made available in D2.1 at M9.

In this deliverable, for each scenario we provide a short discussion on the ongoing developments, and table summarizing the novel models.

²⁹ <https://schema.org>

³⁰ The on going work on recommendations is described at <http://fiware-datamodels.readthedocs.io/en/latest/guidelines/index.html>

³¹ https://en.wikipedia.org/wiki/Camel_case

³² <https://github.com/Fiware/dataModels>

3.3.1 Smart City

The Smart City scenario, besides reusing the harmonized transportation models developed by FIWARE, requires models for the *Transportation Schedule, User's Profile and Diseases*³³. The Transportation Schedule data model is required to support the route planning based on public transports routes. The model will be based on the standard de facto General Transit Feed Specification (GTFS) model. The User's Profile will model the user context needed to provide a customised route planning. The model will map to schema.org general attributes to describe people. The Disease model will be used to customise the route planning according to the User's health conditions. The model will leverage on the health-lifesci schema³⁴.

3.3.2 Smart Security

The Smart City scenario, besides reusing the harmonized for parking, devices and transportation developed in FIWARE, requires models to support the description of **Building, Video, Person and Objects**. The Building data model will be used to describe surveilled buildings. It will be based on the Building model proposed by GSMA IoT Big Data Harmonised Data Model initiative. The Video data model will be need to store metadata about recorded surveillance video and to annotated them with in respect detected security events. The model will be based on relevant video metadata models. The Person model will be used to describe identified person in a video streaming. The Object model will be used to describe identified objects in a video streaming.

3.3.3 Smart Health

The Smart Health scenario, focusing on acquisition of health parameters from mobile devices, covers topics not yet covered in the FIWARE Data Model initiative, neither by previous FIWARE projects in the healthcare area (FI-STAR). The work conducted in the scenario builds on top of the Open Mobile Health³⁵ standard. Entities modelled include: the **User**, the **Patient**, the **Physical Test**, the **Questionnaire**, and the **Control Test**. The User entity collects contextual information on the application user (e.g. a doctor); the Patient entity collects demographic data to identify patients; the Physical Test collect information from sensor in relation to Patients; the Questionnaire entity contains information collected by doctors after a Physical Test; and the Control Test collect several patient measurements to be related with the Physical Test.

3.4 Tools

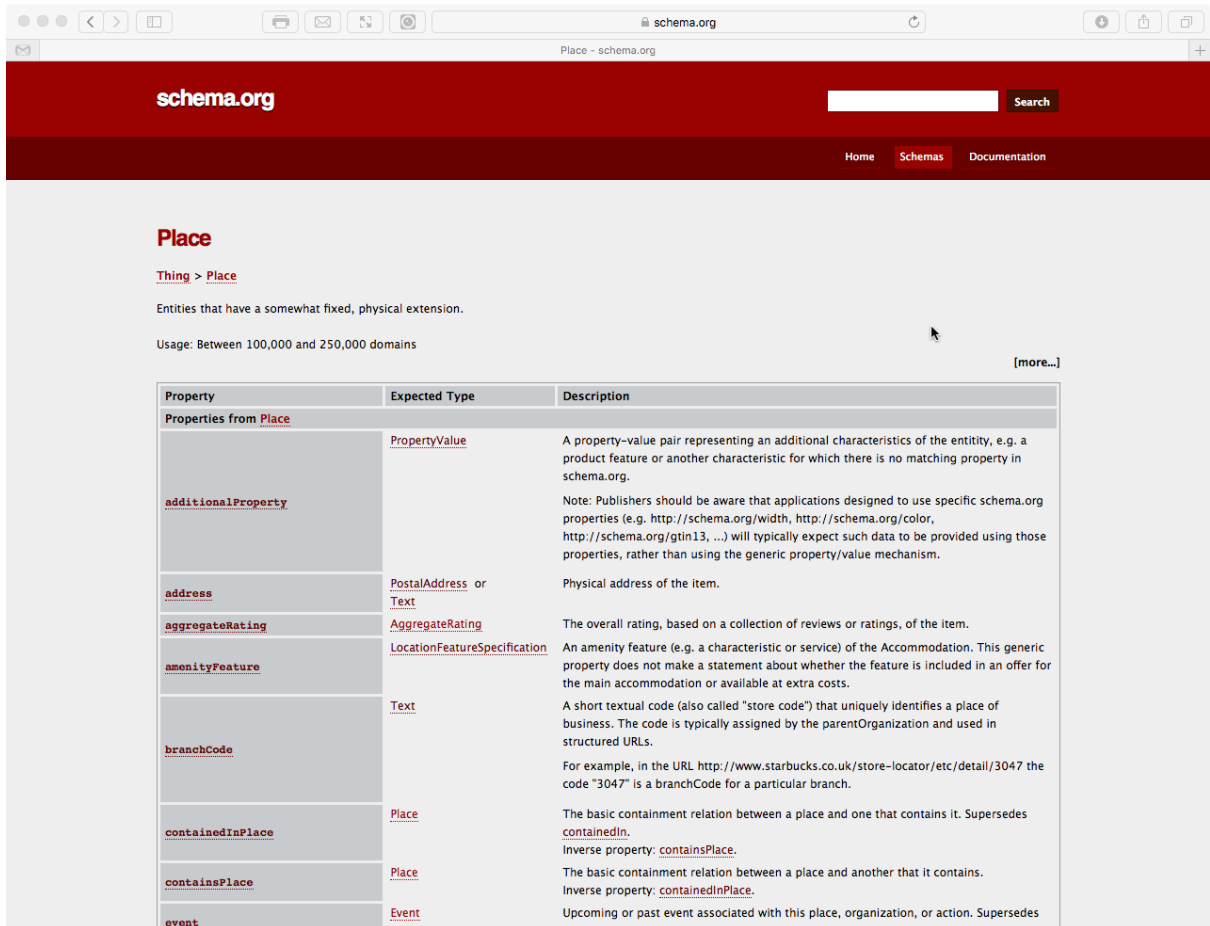
SmartSDK to support the FIWARE data models related activities, plans to provide two tools dealing with the presentation and documentation of FIWARE data models.

- ➔ A command line script that support the generation of a UML-like graphical representation of FIWARE data models schema as part of their documentation to facilitate the understanding of the developers of the models and their relations.
- ➔ A schema.org like portal to support the navigation of FIWARE data models (see Figure 21). This activity will require enriching the current JSON Schemas with proper annotations to support their serialization into RDF format. This activity also requires alignment with the ongoing work on JSON-LD support for NGSI.

³³ <https://developers.google.com/transit/gtfs/>

³⁴ <https://health-lifesci.schema.org>

³⁵ <http://www.openmhealth.org>



Place

[Thing](#) > [Place](#)

Entities that have a somewhat fixed, physical extension.

Usage: Between 100,000 and 250,000 domains

[more...]

Property	Expected Type	Description
Properties from Place		
additionalProperty	PropertyValue	A property-value pair representing an additional characteristics of the entity, e.g. a product feature or another characteristic for which there is no matching property in schema.org. Note: Publishers should be aware that applications designed to use specific schema.org properties (e.g. http://schema.org/width , http://schema.org/color , http://schema.org/gtin13 , ...) will typically expect such data to be provided using those properties, rather than using the generic property/value mechanism.
address	PostalAddress or Text	Physical address of the item.
aggregateRating	AggregateRating	The overall rating, based on a collection of reviews or ratings, of the item.
amenityFeature	LocationFeatureSpecification	An amenity feature (e.g. a characteristic or service) of the Accommodation. This generic property does not make a statement about whether the feature is included in an offer for the main accommodation or available at extra costs.
branchCode	Text	A short textual code (also called "store code") that uniquely identifies a place of business. The code is typically assigned by the parentOrganization and used in structured URLs. For example, in the URL http://www.starbucks.co.uk/store-locator/etc/detail/3047 the code "3047" is a branchCode for a particular branch.
containedInPlace	Place	The basic containment relation between a place and one that contains it. Supersedes containedIn .
containsPlace	Place	The basic containment relation between a place and another that it contains. Inverse property: containedInPlace .
event	Event	Upcoming or past event associated with this place, organization, or action. Supersedes startDate .

Figure 21: Data model portal on Schema.org.

3.5 Roadmap

As mentioned in Section 3.4, several activities related development of novel FIWARE Data Model within the SmartSDK application scenarios are ongoing. Table 3 summarizes their release plan according SmartSDK releases. Additional data models may be included in the roadmap (or dropped) with the evolution of the works in SmartSDK applications (WP2).

SmartSDK Release	Expected Date	Data Model / Tool
1	01/2017	Patient
		Physical Test
		Questionnaire
		Control Test
2	04/2017	Building
		Alarm
		User Profile
3	07/2017	Video
		Person
		Objects
4	10/2017	Transportation Schedule
		Serialisation of data models into UML Graphs
5	01/2018	Review of data models
6	04/2018	Data models portal

Table 3: Data models Roadmap Overview.

4 CONCLUSIONS AND FUTURE WORK

This deliverable outlines one of the main contributions of SmartSDK to the FIWARE Community and to the reinforcement of the collaboration between Europe and Mexico around FIWARE. The ability to replicate, in an easy way, the EU experience in Mexico and vice-versa relies on defining reference implementations for a given application scenario. A reference implementation, in the case of a Smart application built using FIWARE, relies on: reference components, reference data models and reference architectures to glue everything together.

SmartSDK, with the aim of facilitating FIWARE adoption beyond simple proof of concepts, invests in the development of reference architectures that, benefitting from Cloud Architecture Patterns, provide platform deployments that deals with production grade requirements such as fault tolerance and ability to scale.

While several activities have yet to be completed in this context within SmartSDK, this deliverable deals with essential work needed to further develop recipes and data models:

- ➔ Provides guidelines for their development;
- ➔ Analyses the current available assets;
- ➔ Provides a roadmap for the next activities.

This last item is particularly relevant to give FIWARE developers an overview of expected contributions from SmartSDK, and lies out path for potential collaborations within the FIWARE community, beyond the SmartSDK consortium.

REFERENCES

- [1] FIWARE NGSIv2: <http://telefonicaid.github.io/fiware-orion/api/v2/stable/>
- [2] Hortonworks Data Platform. (Dec 21, 2015). Hadoop High Availability. https://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.3.4/bk_hadoop-ha/bk_hadoop-ha-20151221.pdf
- [3] Gamma, E., Helm, R., Johnson, R. E., & Vlissides, J. (2016). *Design patterns: elements of reusable object-oriented software*. Boston, MA: Addison-Wesley.
- [4] Wilder, B. (2012). *Cloud architecture patterns*. Beijing: O'Reilly
- [5] Eric Brewer, “CAP twelve years later: How the "rules" have changed”, Computer, Volume 45, Issue 2 (2012), pg. 23–29.
- [6] SmartSDK Github Repository: <https://github.com/martel-innovate/smartsdk-recipes>
- [7] SmartSDK Documentation: <https://martel-innovate.github.io/smartsdk-recipes/>
- [8] JSON Schema draft version 4: <https://tools.ietf.org/html/draft-zyp-json-schema-04>
- [9] GeoJSON format: <https://tools.ietf.org/html/rfc7946>
- [10] Taylor, B. N. (n.d.). The International System of Units (SI) (National Institute of Standards and Technology (U.S.)). Retrieved from <http://physics.nist.gov/Pubs/SP330/sp330.pdf>
- [11] IoT Big Data Harmonised Data Model(Publication). (2016, October 25). Retrieved <http://www.gsma.com/connectedliving/wp-content/uploads/2016/11/CLP.26-v1.0.pdf>
- [12] SmartSDK Consortium. Description of Action. July 2016.