

Article

# A Flexible IoT Stream Processing Architecture based on Microservices

Luca Bixio<sup>1</sup>, Giorgio Delzanno<sup>2</sup>, Stefano Rebora<sup>1</sup>, Matteo Rulli<sup>1</sup>

<sup>1</sup> Flairbit s.r.l, Italy; stefano.rebora@flairbit.io

<sup>2</sup> DIBRIS, University of Genova, Italy, giorgio.delzanno@unige.it

\* Correspondence: Giorgio Delzanno, email:giorgio.delzanno@unige.it

Version October 20, 2020 submitted to Journal Not Specified

**Abstract:** The Internet of Things (IoT) has created new and challenging opportunities for Data Analytics. IoT represents an infinitive source of massive and heterogeneous data, whose real-time processing is an increasingly important issue. Real-time Data Stream Processing is a natural answer for the majority of the goals of IoT platforms, but it has to deal with the highly variable and dynamic IoT environment. IoT applications usually consist of multiple technological layers connecting ‘things’ to a remote cloud core. These layers are generally grouped in two macro-levels: the edge-level (consisting of the devices at the boundary of the network near the devices that produce the data) and the core-level (consisting of the remote cloud components of the application). Real-time Data Stream Processing has to cope with a wide variety of technologies, devices and requirements that vary depending on the two IoT application levels. The aim of this work is to propose an adaptive microservices architecture for an IoT platform able to integrate real-time stream processing functionalities in a dynamic and flexible way, with the goal of covering the different real-time processing requirements that exist among the different levels of an IoT application. The proposal has been formulated for extending Senseioty, a proprietary IoT platform developed by FlairBit S.r.l., but it can easily be integrated in any other IoT platform. A preliminary prototype has been implemented as proof of concept of the feasibility and benefits of the proposed architecture.

**Keywords:** Cloud Computing; Service Oriented Computing; Internet of Things; Real-time Stream Processing; Query Languages

---

## 1. Introduction

Nowadays, with the rise of IoT, we have at our disposal a wide variety of smart devices able to constantly produce large volumes of data at an unprecedented speed. Sensors, smartphones and any other sort of IoT devices are able to measure an incredible range of parameters, such as temperature, position, motion, health indicators and so forth. More and more frequently, the value of these data highly depends on the moment when they are processed and the value diminishes very fast with time: processing them shortly after they are produced becomes a crucial aspect. Indeed, the aim of Real-time Stream Processing is to query continuous data streams in order to extract insights and detect particular conditions as quickly as possible, allowing a timely reaction. Possible examples are the alert generation of a medical device or the real-time monitoring of a production line. In Stream Processing, data are no more considered as static and persistent data stored in a database, but as continuous and potentially unbounded sequences of data elements (i.e. data streams) from which static queries (a.k.a. rules) continuously extract information. The systems that execute this processing phase in a very short time span (milliseconds or seconds) are defined real-time stream processing engines. The IoT world offers an infinite set of use cases where real-time stream processing functionalities can be applied, but IoT applications provide at the same time a heterogeneous environment with respect to requirements,

35 devices and technologies. For these reasons, integrating real-time processing engines in IoT platforms  
36 becomes a challenging operation that requires special attention.

37 An IoT platform provides tools, technologies and capabilities for simplifying the development,  
38 provisioning and management of IoT applications. Real-time stream processing engines are an  
39 increasingly popular and relevant technology, which the majority of the platforms are integrating  
40 in order to provide all the functionalities required by modern IoT applications. Indeed, Real-Time  
41 Stream Processing plays a crucial role in different and common IoT application scenarios, for instance:  
42 Anomaly and fraud detection; Remote monitoring; Predictive Maintenance; Real-time analytics  
43 (Sentiment analysis, Sports analytics, etc.). When integrating real-time stream processing engines in IoT  
44 platforms, the main difficulties arise from the high heterogeneity and dynamicity of the requirements  
45 and technologies of common IoT applications. At high level, a general IoT application consists of the  
46 following layers: The sensors/actuators layer, which includes the IoT devices; The edge layer, which  
47 includes all the devices near the sensors/actuators-level. These edge devices usually play the role of  
48 gateways, enabling the collection and the transmission of data; The core/cloud layer, which includes  
49 all the core functionalities and services of the application; The application/presentation layer, which  
50 includes all the client applications that have access to the core functionalities and services. Integrating  
51 real-time stream processing capabilities in IoT platforms imposes to face the following three main  
52 aspects:

- 53 • Twofold level of applicability. It is required often to apply Real-Time Stream Processing at two  
54 different levels: at edge level and at core/cloud level. Both approaches offer different benefits  
55 but the great difference between the devices and resources at edge level and core level imposes  
56 also quite different requirements that affect the choice of the stream processing engines.
- 57 • Technological pluralism. Due to the previous point, a natural consequence is to introduce  
58 different stream processing engines in the IoT platform because one stream processing technology  
59 rarely covers the edge level and the cloud level requirements. Having different stream processing  
60 engines means having different processing models and languages that must be handled for  
61 implementing stream processing rules.
- 62 • Rules' dynamicity. Usually, real-time IoT stream processing rules are based on a dynamic lifecycle.  
63 In the majority of IoT use cases, the functionalities implemented by real-time stream processing  
64 rules can be temporary functionalities (that are executed on demand and then removed after  
65 a while) or long-running functionalities never modified (e.g. a remote monitoring process).  
66 Moreover, it is often required to deploy rules directly on edge devices for reducing the response  
67 latency time or applying some pre-filtering operations, but when the workload increases, a  
68 scalable approach may be more preferable. For all these reasons, rules should have the possibility  
69 to be dynamically reallocated on different stream processing engines.

70 Considering these aspects, the goal of this work is to propose an adaptive solution for integrating  
71 real-time stream processing functionalities into an IoT platform, Senseioty by Flairbit [21], able to  
72 satisfy the different requirements imposed by the edge level and the cloud/level. Moreover, the  
73 proposal offers a dynamic mechanism for facilitating the dynamic management and relocation of  
74 stream processing rules, hiding at the same time the complexity introduced by the presence of different  
75 and heterogeneous stream processing engines. The innovative aspect of our solution, with respect to  
76 common IoT platforms, consists in limiting the expressive power for defining stream processing rules  
77 to a predefined set of templates, in favor of a much more flexible and dynamic deployment model. The  
78 proposal architecture has been designed following a microservices architectural pattern. Microservices  
79 are a natural and widely adopted solution for implementing software platforms. The majority of  
80 IoT platforms are based on a microservices approach, even when it is not explicitly mentioned. This  
81 happens because the microservices architectural style is a chameleonic style, which can be implemented  
82 in different ways. Indeed, several technologies exist for implementing microservices, but for our  
83 purposes we have selected a particular technology able to guarantee a significant level of flexibility  
84 and dynamicity.

## 85 *Plan of the paper*

86 In Section 2 we give an overview of the main features of the microservices architectural style and  
87 of a particular Java technology named OSGi, the main technology applied in Senseioty, the proprietary  
88 IoT platform developed by FlairBit. In Section 3 and 4 we present our proposal and a prototype  
89 implemented as a possible extension of Senseioty based on Siddhi and Apache Flink. In Section 5 we  
90 address some conclusions and future work.

## 91 **2. The Microservices Architectural Style and Java OSGi**

92 The microservices architectural style, see e.g. [25], is born to address the problems of the  
93 traditional monolithic approach. When you start to design and build a new application, the easiest  
94 and most natural approach is to imagine the application as unit composed by several components.  
95 The application is logically partitioned in modules and each one represents a functionality, however  
96 it is packaged and deployed as a unit. This monolithic approach is very simple and comes naturally.  
97 Indeed, all the IDE's are necessarily designed to build a single application and the deployment of a  
98 single unit is easy and fast. Also scaling the application is trivial because it requires only running  
99 multiple instances of the single unit. This approach initially works and apparently quite well: the  
100 question is what happens when the application starts to grow. When the number of functionalities  
101 increases and the application becomes bigger and bigger, the monolithic approach shows its natural  
102 limit with respect to human capacities. In a short while, the dimensions of the application are such that  
103 a single developer is unable to fully understand it and this leads to serious problems. For example,  
104 implementing a new functionality becomes harder and time consuming and fixing bug even worse.  
105 The whole code is inevitably too complex; therefore adopting new frameworks and technologies is  
106 discouraged. In addition, the deployment and the start-up time are obviously negative affected by  
107 the huge size of the application. The main consequence is the slowdown of the entire development  
108 phase and any attempts of continuous integration and other agile practises fails. Moreover, all the  
109 components run in the same process or environment causing serious problems of reliability: a failure  
110 or a bug in a single component can compromise the entire application. In few words, the overall  
111 complexity of a huge monolith overwhelms the developers. The microservices architectural style  
112 was created specifically to address this kind of problems and to tackle the complexity. The book [24]  
113 describes a three-dimensional scale model known as scale-cube: Horizontal Scaling (running multiple  
114 instances behind a load-balancer); Functional Scaling (decomposing a monolithic application into a  
115 set of services, each one implementing a specific set of functionalities); Scaling of Data Partitioning  
116 (data are partitioned among the several instances and each copy of the application). These concepts  
117 are strongly connected to the idea of Single Responsibility Principle (SRP) [31]. The functionalities are  
118 exposed through an interface, often a REST API, and can be consumed by other services increasing  
119 the composability. The communication between microservices can be indifferently implemented by  
120 synchronous or asynchronous communication protocol and each microservice can be implemented  
121 with a different and ad-hoc technology. Moreover, each microservice has its own database rather  
122 than sharing a single database schema with other services. This makes a microservice an actual  
123 independently deployable and loosely coupled component. In this setting communication is provided  
124 via an API Gateway [28].

125 The API Gateway is similar to the Facade pattern from object-oriented design: it is a software  
126 component able to hide and encapsulate the internal system details and architecture, providing a  
127 tailored API to the client. It is responsible for handling the client's requests and consequently invoking  
128 different microservices using different communication protocol, finally aggregating the results.  
129 Microservices architectures often provide a service discovery mechanism typically implemented  
130 via a shared registry which is basically a database that contains the network locations of the associated  
131 service instances. Two important requirements for a service registry are to be highly available and up to  
132 date, thus it often consists in a cluster of servers that use a replication protocol to maintain consistency.  
133 One of the main principles at the heart of the microservices architecture is the decentralization of

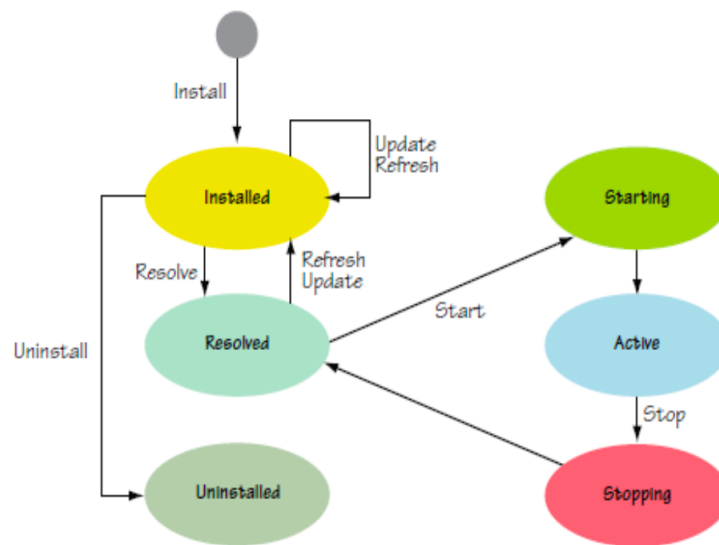


Figure 1. Bundle life cycle.

134 data management: each microservice encapsulates its own database and data are accessible only  
 135 by its API. This approach makes microservices loosely coupled, independently deployable and  
 136 able to evolve independently from each other. In addition, each microservice can adopt different  
 137 database technologies depending on its specific requirements, for example for some use cases a NoSQL  
 138 database may be more appropriate than a traditional SQL database or vice versa. Therefore, the  
 139 resulting architecture often uses a mixture of SQL and NoSQL databases, leading to the so-called  
 140 polyglot persistence architecture. In this setting, data consistency is often achieved via an event-driven  
 141 architecture [29]. A message broker is introduced into the system and each microservice publishes  
 142 an event whenever a business entity is modified. Other microservices subscribe to these events,  
 143 update their entities and may publish other events in their turn. The event-driven architecture is  
 144 also a solution for the problem of queries that have to retrieve and aggregate data from multiple  
 145 microservices. Indeed, some microservices can subscribe to event channels and maintain materialized  
 146 views that pre-join data owned by multiple microservices. Each time a microservice publishes a new  
 147 event, the view is updated. The last key aspect of the microservices architecture is how a microservices  
 148 application is actually deployed. Three main different deployment patterns exist [30]: Multiple Service  
 149 Instances per Host Pattern; Service Instance per Host Pattern sub-divided in (Service Instance per  
 150 Virtual Machine Pattern/Container Pattern); Serverless Deployment Pattern (e.g. AWS Lambda;  
 151 Google Cloud Functions; Azure Functions).

152 Java OSGi OSGi [14] consists of a set of specifications established by the OSGi Alliance. The  
 153 OSGi architecture [15] appears as a layered model. The bundles are the modules implemented by  
 154 the developers. A bundle is basically a standard JAR file enriched by some metadata contained in  
 155 a manifest [27]. The manifest and its metadata make possible to extend the standard Java access  
 156 modifiers (public, private, protected, and package private). A bundle can explicitly declare on which  
 157 external packages it depends and which contained packages are externally visible, meaning that the  
 158 public classes inside a bundle JAR file are not necessarily externally accessible. The module, life cycle  
 159 and services layer constitute the core of the OSGi framework: The module layer defines the concept of  
 160 bundle and how a bundle can import and export code; The life cycle layer provides the API for the  
 161 execution-time module management; The service layer provides a publish-find-bind model for plain  
 162 old Java objects implementing services able to connect dynamically the bundles. Finally, the security  
 163 layer is an optional layer, which provides the infrastructure to deploy and manage applications that  
 164 must run in fine-grained controlled environments, and the execution environment defines the methods  
 165 and classes that are available in a specific platform. A Bundle object logically represents a bundle into

166 OSGi framework and it defines the API to manage the bundle's lifecycle. The BundleContext represents  
167 the execution context associated to the bundle. It basically offers some methods for the deployment and  
168 lifecycle management of a bundle and other methods for enabling the bundle interaction via services.  
169 It is interesting to notice that the BundleContext interface has methods to register BundleListener and  
170 FrameworkListener objects for receiving event notifications. These methods allow to monitor and to  
171 react to execution-time changes into the framework and to take advantage of the flexible dynamism  
172 of OSGi bundles. Finally, the BundleActivator offers a hook into the lifecycle layer and the ability to  
173 customize the code that must be executed when a bundle is started or stopped. The class implementing  
174 the BundleActivator inside a bundle is specified adding the Bundle-Activator header to the bundle  
175 manifest.

176 As shown in Figure 1, firstly, a bundle must be installed into OSGi framework. Installing a bundle  
177 into the framework is a persistent operation that consists in providing a location of the bundle JAR  
178 file to be installed (typically a URL) and then saving a copy of the JAR file in a private area of the  
179 framework called bundle cache. Then, the transition from installed to resolved state is the transition  
180 that represents the automated dependency resolution. This transition can happen implicitly when  
181 the bundle is started or when another bundle tries to load a class from it, but it can also be explicitly  
182 triggered using specific methods of lifecycle APIs. A bundle can be started after being installed into  
183 the framework. The bundle is started through the Bundle interface and the operations executed during  
184 this phase (e.g. operations of initialization) are defined by an implementation of the BundleActivator.  
185 The transition from the starting to the active state is always implicit. A bundle is in the starting state  
186 while its BundleActivator's start() method executes. If the execution of the start() method terminates  
187 successfully, the bundle's state transitions to active, otherwise it transitions back to resolved. Similarly,  
188 an active bundle can be stopped and an installed bundle can be uninstalled. When uninstalling an  
189 active bundle, the framework automatically stops the bundle first. The bundle's state goes to resolved  
190 and then to installed state before uninstalling the bundle.

191 The OSGi environment is dynamic and flexible and it allows to update a bundle with a newer  
192 version even at execution-time. This kind of operation is quite simple for self-contained bundles  
193 but things get complicated when other bundles depend on the bundle being updated. The same  
194 problem exists when uninstalling a bundle, both the updating and uninstalling operations can cause  
195 a cascading disruption of all the other bundles depending on it. This happens because, in case of  
196 updating, dependent bundles have potentially loaded classes from the old version of the bundle,  
197 causing a mixture of loaded old classes and new ones. The same inconsistent situation occurs when a  
198 dependent bundle cannot load classes from a bundle that has been uninstalled. The solution for this  
199 scenario is to execute the updating and uninstalling operation as a two-step operation: the first step  
200 prepares the operation; the second one performs a refreshing. The refreshing allows to recalculate the  
201 dependencies of all the involved bundles, providing a control of the moment when the changeover  
202 to the new bundle version or removal of a bundle is triggered for updates and uninstalls. Therefore,  
203 each time an update is executed, in the first step the new version of the bundle is introduced and two  
204 versions of the bundle coexist at the same time. Similarly, for uninstalling operations, the bundle is  
205 removed from the installed list of bundles, but it is not removed from memory. In both cases, the  
206 dependent bundles continue to load classes from the older or removed bundle. Finally, a refreshing  
207 step is triggered and all the dependencies are computed and resolved again. In conclusion, the lifecycle  
208 layer provides powerful functionalities for handling, monitoring and reacting to the dynamic lifecycle  
209 of bundles. The next section presents the last but not the least layer of the OSGi framework: the service  
210 layer.

211 Java OSGi and Microservices OSGi allows the combination of microservices and nanoservices.  
212 Leveraging the OSGi service layer, it is possible to implement microservices internally composed by  
213 tiny nanoservices. The final resulting architecture will be composed by a set of microservices, each  
214 one running on its own OSGi runtime and communicating remotely with the other microservices.  
215 Internally, a single microservice may be implemented as a combination of multiple nanoservices that



216 communicate locally as a simple method invocations. Secondly, OSGi offers an in-built dynamic  
217 nature. Developing microservices using OSGi means having a rich and robust set of functionalities  
218 specifically implemented for handling services with a dynamic lifecycle. Even more, it makes the  
219 microservices able to be aware of their dynamic lifecycle and react consequently to the dynamic  
220 changes. The OSGi runtime and its service layer were built upon this fluidity; therefore, the resulting  
221 microservices are intrinsically dynamicity-aware microservices. Last but not the list, OSGi makes  
222 the microservices architecture a more flexible architecture with respect to service decomposition.  
223 One drawback of microservices is the difficulty of performing changes or refactoring operations that  
224 span multiple microservices. When designing a microservices architecture, understanding exactly  
225 how all the functionalities should be decomposed into multiple small microservices is an extremely  
226 difficult task, which requires defining explicit boundaries between services and establishing once for  
227 all the communication protocols that will be adopted. If in future, the chosen service decomposition  
228 strategy turns to be no more the best choice or only a modification involving the movement of one  
229 or multiple functionalities among different microservices is required, performing this change may  
230 become extremely difficult because of the presence of already defined microservices boundaries and  
231 communication protocols. On the contrary, the OSGi Remote Services offers a flexible approach for  
232 defining the microservices boundaries. Indeed, a set of functionalities implemented by an OSGi service  
233 can be easily moved from a local runtime to a remote one without any impact on other services.  
234 Therefore, an already defined microservices decomposition strategy can be modified by reallocating  
235 the functionalities offered by services at any time. For example, one of two OSGi services previously  
236 designed for being on the same runtime (i.e. within the same microservice boundary) can be moved  
237 on another remote OSGi runtime without any difficult changes. The interaction between a distribution  
238 provider and a distribution consumer in OSGi takes place always as the two entities were on the  
239 same and local runtime: the distribution manager provided by the Remote Services specification  
240 transparently handles the remote communication. Moreover, this remote communication is completely  
241 independent from the communication protocols; therefore, any previous choice is not binding at all. In  
242 conclusion, OSGi enriches the microservices architecture with new and powerful dynamic properties  
243 and a flexible model able to support elastic and protocol-independent service boundaries. Moreover, it  
244 provides a level of service granularity highly variable allowing the combination of microservices and  
245 nanoservices. The only but very relevant drawback of OSGi with respect to microservices architectural  
246 pattern is the complete cancellation of technological freedom that characterizes microservices. OSGi  
247 is a technology exclusively designed for Java and implementing a microservices architecture based  
248 on OSGi necessary requires to adopt Java for developing the microservices. This does not mean  
249 that a microservice implemented using OSGi cannot be integrated with other services implemented  
250 with different technology; an OSGi remote service, for example, can be exposed externally also for  
251 not-OSGi service consumers, loosing however all the OSGi service layer benefits. It actually means  
252 that if Java and OSGi are not widely adopted for implementing the majority of the microservices of  
253 the architecture, the OSGi additional features lose their effectiveness. OSGi represents also a very  
254 powerful and dynamic service-oriented platform due to the several features offered by its service layer  
255 [17].

256 Finally, the last relevant and powerful feature of the OSGi service layer is the flexibility offered  
257 by the Remote Services Specification [20]. The OSGi framework provides a local service registry for  
258 bundles to communicate through service objects, where a service is an object that one bundle registers  
259 and another bundle gets. However, the Remote Services Specification extends this behaviour in a very  
260 powerful and flexible manner, allowing the OSGi services to be exported remotely and independently  
261 from the communication protocols. The client-side distribution provider is able to discover remote  
262 endpoints and create proxies to these services, which it injects into the local OSGi service registry. The  
263 implementation of the discovery phase depends on the chosen distribution provider implementation  
264 (e.g. The Apache CXF Distributed OSGi [3] implementation provides discovery based on Apache  
265 Hadoop Zookeeper [9]). Another additional and powerful feature of OSGi Remote Services is the ability

266 to be independent from the underlying communication protocol adopted for the service exportation.  
267 A distribution provider may choose any number of ways to make the service available remotely. It can  
268 use various protocols (SOAP, REST, RMI, etc.), adopting a range of different security or authentication  
269 mechanisms and many different transport technologies (HTTP, JMS, P2P, etc.). The Remote Services  
270 specification offers a layer of indirection between the service provider and the distribution provider,  
271 leveraging the concepts of intents and configurations. They basically allow the service provider  
272 to specify just enough information to ensure that the service behaves as expected, then the task of  
273 the distribution provider is to optimize the communications for the environment in which they are  
274 deployed.

### 275 **3. A Microservices Architecture for Adaptive Real-time IoT Stream Processing**

276 Senseioty [21] is an IoT platform designed to accelerate the development of end-to-end solutions  
277 and verticals, revolving around the concept of insights-engineering, the seamless integration between  
278 data ingestion and distribution, data analytics and on-line data analysis. Senseioty is developed in  
279 Java as a set of highly cohesive OSGi microservices. Each Senseioty microservice can either be used  
280 together with Amazon AWS or Microsoft Azure managed services or deployed on private cloud or  
281 on-premises to accelerate and deliver full-fledged end-to-end IoT solutions for the customer. Senseioty  
282 features also an SDK to implement rapid verticalizations on top of its rich set of JSON RESTful APIs  
283 and analytics services. Senseioty automates the integration of IoT operational data with analytics  
284 workflows and provides a common programming model and semantics to ensure data quality, simplify  
285 data distribution and storage and enforce data access policies and data privacy. Senseioty is natively  
286 integrated with both Microsoft Azure IoT and Amazon AWS IoT and it can also operate on private and  
287 hybrid cloud to provide the maximum flexibility in terms of cloud deployment models Senseioty offers  
288 a wide variety of interesting and flexible functionalities that should give an idea of the flexibility and  
289 interoperability offered by a microservices architecture in the IoT context: Single-sign-on services for  
290 user and devices along with user management; Access policies microservice to protect resources and  
291 devices against unauthorized access and to guarantee data privacy; Flexible and unified programming  
292 interface to manage and provision connected devices.; Persistence of time series in Apache Cassandra  
293 clusters; Powerful and flexible way to communicate different microservices together and to implement  
294 remote services discovery based on the OSGi Remote Service specification; Senseioty microservices  
295 can be deployed at the three different layers of the hybrid-cloud stack (cloud layer, on-premises layer  
296 and edge layer); Deep-learning workflows based on neural networks and to push them on connected  
297 devices, in order to run analytics workflow on the edge. Senseioty integrates Apache Spark, a powerful  
298 Distributed Data Stream Processors engine to analyse data stream in real-time and provide on-line data  
299 analytics on the cloud, leveraging both neural network and statistical learning techniques to analyse  
300 data. Finally, Senseioty provides a rich set of IoT connectors to integrate standard and custom IoT  
301 protocols and devices.

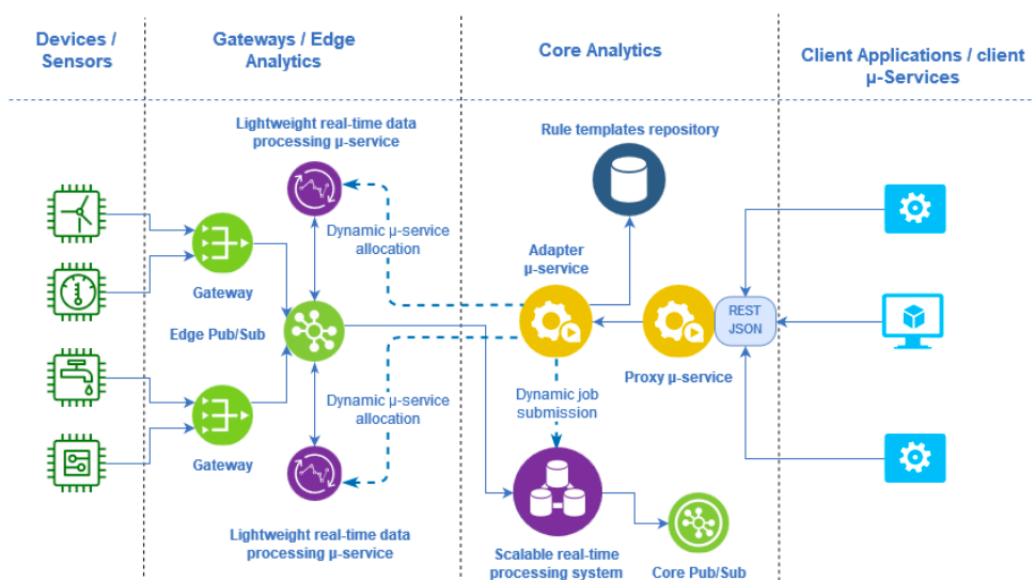
302 FlairBit extensively adopts in its platform Apache Karaf [6], a powerful and enterprise ready  
303 applications runtime built on top two famous OSGi implementation (Apache Felix and Eclipse Equinox)  
304 that offers some additional and useful functionalities, such as the concept of feature.

305 One problem exposed by FlairBit, which is usually a problem common to the majority of IoT  
306 platform, is to have two different levels of data stream processing: The edge level; The core/cloud level.  
307 The two levels offer different benefits but impose quite different requirements. The term “edge” in IoT  
308 platforms generally means the location at the boundary of the network near the devices that produce  
309 the data. Edge devices are usually quite simple devices that play the role of gateways, enabling  
310 the collection and the transmission of data. However, modern edge devices can also offer enough  
311 computational resources to enable more complex functionalities, such as pre-processing, monitoring or  
312 pre-filtering. Moving the stream processing elaboration directly on the edge of the IoT platform takes  
313 the name of Edge analytics and the consequent benefits are quite notable: Lowest possible latency,  
314 having a stream processing unit deployed directly on an edge devices makes possible to respond

315 quickly to events produced locally, avoiding to send data to the remote cloud/core of the platform over  
 316 the network; Improved reliability, moving the stream processing rules on the edge allows the edge  
 317 devices to operate even when they lose the connection with the core platform; Reduced operational  
 318 costs, pre-processing and pre-filtering data directly on the edge makes possible to save bandwidth,  
 319 cloud storage and computational resources consequently lowering operational costs.

320 On the other hand, edge analytics imposes some stringent requirements in term of computational  
 321 power. The modern edge devices are becoming more and more powerful, but the computational  
 322 resources offered by this kind of devices are limited. Therefore, the technologies installed on the edge  
 323 must be lightweight and it is likely that they are quite different technologies from those applied on the  
 324 core platform. Indeed, the stream processing units on the edge usually deals with simple filtering rules  
 325 and streams of data restricted to the local sensors or devices, without the need of scaling the stream  
 326 processing job across multiple machines.

327 On the core platform, the context is completely different. In this scenario, the stream processing  
 328 engines must be able to deal with different workloads and the computational resources abound. They  
 329 must be able to scale the computation across a cluster of machines in order to handle large volume of  
 330 data and more intensive tasks, for example joining and aggregating different events from different  
 331 streams of data. Therefore, the need of scaling capabilities overcomes the limit of the computational  
 332 resources. The main goals of the proposed extension of the Senseioty architecture are as follows: 1)  
 333 Providing adaptivity, meaning that the stream processing units can be indifferently allocated on the  
 334 edge or on the core and moved around. This makes possible to cover the two different levels of data  
 335 stream processing, the edge level and the core/cloud level, and exploiting all their different benefits. 2)  
 336 Providing flexibility, allowing a punctual and on-demand deployment of the stream processing units.  
 337 The user or the client application/service defines when and where allocating, starting, stopping and  
 338 deallocating the stream processing rules. 3) Providing a set of portable and composable rules that can  
 339 be defined in a standard way and then automatically deployed on different stream processing engines  
 340 without depending on their own languages and models. The rules can be combined together, in order  
 341 to apply a sort of stream processing pipeline. The rules are not only dynamically manageable, but  
 composable and engine-independent. The reference structure of the resulting architecture is shown in



342 **Figure 2.** Reference Architecture

343 Fig. 2. There are two main components in our architecture:

- 344 • The proxy  $\mu$ -service: the entry point of the architecture, offering a RESTful API for installing,  
 345 uninstalling, starting, stopping and moving stream processing rules on demand.



- The adapter  $\mu$ -service. It is responsible for physically executing the functionalities offered by the proxy  $\mu$ -service, interacting with the different stream processing engines available on the edge and on the core of the architecture.

The proxy  $\mu$ -service represents the entry point of the architecture. It offers a RESTful JSON interface, a standard choice in microservices architecture (and it is usually adopted in Senseioty) in order to offer a solution as much compatible and reusable as possible. The REST API offers the following functionalities:

URL Method	Request Body	Response Body
/api/install POST	JSON installation object	JSON jobinfo object
/api/uninstall POST	JSON jobinfo object	JSON jobinfo object
/api/start POST	JSON jobinfo object	JSON jobinfo object
/api/stop POST	JSON jobinfo object	JSON jobinfo object
/api/move POST	JSON relocation object	JSON jobinfo object

The method *install* installs the rule on the required resource and engine defined by the json installation object. The method *uninstall* uninstalls the rule identified by the jobinfo request object. The method *start* runs the rule identified by the jobinfo request object. The method *stop* stops the execution of the rule identified by the jobinfo request object. The method *move* moves the rule identified by the jobinfo request object to the target runtime defined by the relocation object. Interaction with the proxy  $\mu$ -service is carried out through the JSON objects of the following form:

```

360 // JSON INSTALLATION OBJECT
361 {
362     "headers":{
363         "runtime":<ENGINE>,
364         "targetResource":<URL>,
365         "jobType":<JOB_TYPE>
366     },
367     "jobConfig":{
368         "connectors":{
369             "inputEndpoint":<STRING>,
370             "outputEndpoint":<STRING>
371         },
372         "jobProps":{
373             "condition":< ">" | ">=" | "=" | "<" | "<=" >,
374             "threshold":< INT | FLOAT | DOUBLE | STRING >,
375             "fieldName":<STRING>,
376             "fieldJsonPath":<JSON_PATH>
377         }
378     }
379 }
380 // JSON JOBINFO OBJECT
381 {
382     "runtime":<ENGINE>,
383     "jobId":<STRING>,
384     "jobType":<JOB_TYPE>,
385     "jobStatus":<INSTALLED|RUNNING|STOPPED|UNINSTALLED>,
386     "configFileName":<STRING>
387 }
388 // JSON RELOCATION OBJECT
389 {
390     "target_runtime":<ENGINE>,

```

```

391     "targetResource":<URL>,
392     "jobInfo":<JSON_JOBINFO_OBJECT>
393 }

```

394 The JSON installation object is the object that the client must provide to the proxy in order to describe
395 the stream processing rule to be allocated. The headers field indicates the runtime engine that will
396 execute the rule (the <ENGINE> value depends on the engines supported by the implementation), the
397 target resource which is the machine on which allocating the rule (the value can be an URL or a simple
398 ID, depending on the architecture implementation) and the job type, which indicates the kind of rule
399 that the jobProps field contains. The implementation of the architecture supports a set of predefined
400 rule templates identified by a unique name that must be inserted in the jobType field (e.g. single-filter,
401 sum-aggregation, avg-aggregation, single-join etc.). Ideally, we would like to have a solution able to
402 support any kind of rule expressible with a standard query stream language (e.g. the Stanford CQL
403 [23]), but in practice this is not achievable because each stream processing engine has its own model
404 and language with its own level of expressiveness. Therefore, it is extremely complicated to implement
405 a compiler able to validate an arbitrary query and to compile and translate it to the model or language
406 of the underlying stream processing engine. Considering this scenario, we provide an architecture
407 able to support a set of predefined rule templates. A possible subset that should be compatible with
408 the majority of stream processing engines includes (using a SQL like syntax):

- 409 • Filtering query (e.g. SELECT \* FROM inputEvents WHERE field > threshold)
- 410 • Aggregation query over a window (e.g. SELECT SUM(field) FROM inputEvents[5 s])
- 411 • Joining query between two streams over windows (e.g. SELECT field1 field2 FROM stream1[1m]
- 412 JOIN stream2[1m] ON stream1.field3 = stream2.field)

413 This is of course only a possible subset, which must be verified and extended considering the engines
414 selected for the implementation.

415 The connectors field specifies the information needed for reading and writing the events consumed
416 by the rule from/to a pub-sub broker. Again, the format of these fields depends on the pub-sub broker
417 adopted in the implementation, but it in general the required parameters are a simple URL or a queue
418 or topic name. It is important to notice that the presence of two pub-sub brokers (one on the edge
419 and one on the core) makes possible to combine and compose the rules in order to obtain stream
420 processing pipelines. The jobProps field contains the parameters needed for allocating the stream
421 processing rules. The format of this field depends on the rule template specified in the jobType field.
422 The JSON jobInfo object is the object that contains all the necessary information that must be provided
423 in order to perform all the other operations (starting, stopping, uninstalling or moving the rule) and it
424 is created by the adapter  $\mu$ -service and returned to the client by the proxy  $\mu$ -service. It contains some
425 information specified by the JSON installation object, with the addition of a jobId (a unique identifier
426 for the installed rule instance), a jobStatus (it indicates the current execution status of the rule) and a
427 configFileName (the name of the configuration file that represents the materialization of the jobConfigs
428 field specified in the JSON installation object). The role of the configuration file will be clarified shortly
429 when describing the adapter  $\mu$ -service. The JSON relocation object is the object required for moving a
430 rule from the current runtime to a target runtime. It contains the jobinfo object describing the selected
431 rule and information regarding the target runtime (the engine and the resource URL or ID identifying
432 the target machine).

433 The adapter  $\mu$ -service is responsible for actually executing the functionalities offered by the proxy
434  $\mu$ -service. It offers the following procedures: A procedure for installing a new rule; A procedure for
435 starting/stopping/uninstalling an existing rule; A procedure for moving an existing rule from its
436 current runtime to another one. During the installation procedure, the adapter  $\mu$ -service translates the
437 information received from the proxy  $\mu$ -service into executable rules via a sort of parametrization as
438 shown below: The adapter  $\mu$ -service has access to a repository from where it can download the rule
439 template corresponding to the jobType and runtime fields expressed in the JSON installation. The

440 rule template is any sort of predefined executable file (for our purposes will be a JAR archive) that  
441 can be modified injecting a configuration file containing the rule parameters specified by the JSON  
442 installation object. Therefore, in case of rule installation, the adapter  $\mu$ -service downloads the relative  
443 rule template, creates and injects the configuration file and then install the rule on the target runtime.  
444 If the target runtime is a distributed stream processing engine, the executable template is actually  
445 an executable job that is submitted to the cluster manger. If the target runtime is a lightweight and  
446 non-distributed stream processing engine for the edge, the rule template is actually an independent  
447  $\mu$ -service that is installed on the target machine on the edge. Finally, the adapter  $\mu$ -service creates  
448 the JSON jobinfo object with the necessary information that will be returned to the client. In case of  
449 starting, stopping and uninstalling operations, the adapter  $\mu$ -service acts always depending on the  
450 runtime engine associated to the rule, as shown below.

451 In case of distributed stream processing engine, it communicates with the cluster manager for  
452 executing the required operation. On the other hand, in case of lightweight stream processing  $\mu$ -service  
453 on the edge, the implementation must provide a mechanism to interact dynamically with target  
454 runtime. It is intuitive to understand that a technology like OSGi and its bundle lifecycle naturally  
455 fits this scenario. OSGi is the main technology adopted in the prototype that will be described in  
456 the next section, but this architecture description section is intentionally lacking of technical and  
457 implementation details in order to be as much general as possible. The idea is to offer a guideline  
458 proposal that must be refined with respect to technologies selected for the implantation, which may be  
459 completely different from those selected for our prototype. Indeed, one benefits of a microservices  
460 architecture is the technological freedom.

461 Finally, for moving an existing rule across different runtimes the adapter  $\mu$ -service acts as follows.  
462 First, it checks if the rule to be moved is running and eventually it stops its execution. Secondly, it  
463 uninstalls the current rule, it downloads the new template for the new target runtime and it injects  
464 the previous configuration file. Finally, it installs the new rule on the new target runtime, starting the  
465 execution of the new rule if it was previously running. In our architecture we introduce two pub/sub  
466 brokers. Having two event dispatcher systems in the architecture, one for edge level and another one  
467 for the cloud/core level, makes possible to implement composable stream processing rules. Indeed, the  
468 connectors field in the JSON installation object allows to specify the queue or topic names from where  
469 reading events and where writing the output events. This means that any rule can be concatenated  
470 with other rules in order to implement a stream processing pipeline. For example, two filtering rules  
471 can be combined on the same edge-device leveraging the edge pub/sub broker in order to create a  
472 two-step filter. Moreover, a pre-filtering rule on the edge can be applied on top of an aggregation rule  
473 executed at cloud/core level in order to reduce the amount of data sent over the network.

474 The last aspect to consider is the client application layer. As already explained, the proxy  
475  $\mu$ -service offers a simple RESTful JSON API accessible from any kind of client. For this reason, the  
476 functionalities offered by the API can be employed by other  $\mu$ -services in the context of a larger  
477 platform (e.g. Senseioty) and combined with other additional functionalities (e.g. the authentication  
478 and authorization  $\mu$ -services offered by Senseioty). Moreover, it is possible to provide a web interface  
479 that lets a user to interact directly with the proxy  $\mu$ -service, defining and managing the rules. The  
480 API offers all the functionalities needed for implementing an adaptive monitoring rule relocation  
481 procedure. The only prerequisites are: Having access to a stream of events logging statistics about  
482 the performance and workload of the edge-devices; Having the possibility to store and update the  
483 information mapping the rules (i.e. the JSON jobInfo objects) to the IDs of the edge devices that are  
484 executing the rules.

#### 485 4. Prototype Implementation

486 In order to implement a prototype of the proposed architecture we considered a lightweight stream  
487 processing engine for edge analytics called Siddhi [22]. Siddhi is an effecting streaming processing  
488 engine that provides an SQL-like stream language with a rich expressive power. It allows any sort of

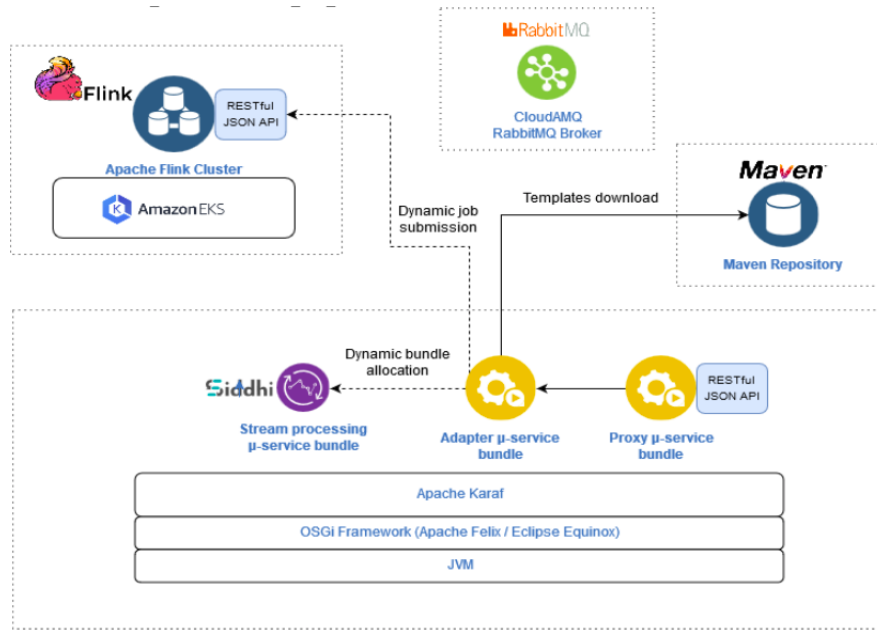


Figure 3. Prototype Structure

489 stateful and stateless operation, timing and counting windows, aggregation and join functions. It also  
 490 supports different event formats (JSON, XML, etc.) for specifying event patterns for complex event  
 491 processing (CEP). It provides a rich set of external event source integration, such as Kafka, MQTT,  
 492 RabbitMQ and other brokers and provides a lightweight runtime compatible, e.g., with Android  
 493 devices. The Siddhi libraries were transformed and wrapped into well-defined OSGi bundles. The  
 494 Senseioty SDK provides some Java project templates explicitly configured for applying OSGi specific  
 495 tools (e.g. Bnd tools [10]) able to create a JAR with OSGi meta data (i.e. a bundle) based on instructions  
 496 and the information in the class files. A feature for the Karaf runtime, collecting all the bundles  
 497 needed by Siddhi as dependencies, was created. The Senseioty SDK offers some functionalities able to  
 498 discover all the dependencies and transitive dependencies required by a bundle and then to materialize  
 499 them in the form of a Karaf feature. Therefore, the provisioning phase of a Siddhi application on a  
 500 Karaf runtime (by provisioning application, it means install all modules, configuration, and transitive  
 501 applications) requires now only a simple and automatic feature installation. Although this step  
 502 required a lot of technical passages, a detailed description is beyond the scope of the paper. Once  
 503 obtained a fully OSGi-compliant stream processing engine for edge analytics purposes, the second  
 504 step consisted in exploring and selecting another engine able to scale across a cluster of machines for  
 505 core/cloud analytics purposes. During this phase, an analysis and some implementation of spike  
 506 test programs were performed for the following stream processing technologies: Ignite [5]; Samza [7];  
 507 Flink [4]; Storm [8]; Streams [11]. Apache Flink turned out to be the most flexible available solution.  
 508 It has a rich and complete API that follows a declarative model very similar to the Spark Streaming  
 509 one and it has also a powerful additional library for complex event processing for specifying patterns  
 510 of events. Moreover, it has a rich set of out-of-the-box external source connectors, a flexible resource  
 511 allocation model based on slots independent from the number of CPU cores, and it is very easy to  
 512 deploy a Flink cluster on Kubernetes. Based on all the above considerations, the structure of the  
 513 implemented prototype is shown in Fig. 3. The proxy and the adapter  $\mu$ -services are implemented  
 514 as OSGi-bundles deployed on a Karaf runtime. The proxy  $\mu$ -service offers a RESTful JSON API with  
 515 the following functionalities. First, it provides an installation functionality for installing a filtering  
 516 rule for events in JSON format. The rule can be indifferently instantiated as an independent Siddhi  
 517  $\mu$ -service (implemented in the form of an OSGi bundle) or deployed as a distributed job on a Flink  
 518 cluster. It also provides a starting, stopping and uninstalling functionalities for removing or handling

519 the rule execution, and a moving functionality for relocating a rule from a Siddhi runtime to a Flink  
 520 runtime or vice versa. The RESTful API was implemented using Apache CXF [2], an open-source and  
 521 fully featured Web services framework. In this preliminary implementation, the runtime supported  
 522 are Siddhi and Apache Flink and only one rule type is available: a threshold filter for events in JSON  
 523 format. The client can specify a filtering rule defining the following jobProps in the JSON installation  
 524 object:

```
525 // JSON INSTALLATION OBJECT
526
527 { ...
528     "jobProps":{
529         "condition":< ">" | ">=" | "=" | "<" | "<=" >,
530         "threshold":< INT | FLOAT | DOUBLE | STRING >,
531         "fieldName":<STRING>,
532         "fieldJsonPath":<JSON_PATH>
533     }
534 }
535 }
```

536 The prototype supports one rule type: a threshold filter for events in JSON format. The parameters  
 537 specified by this rule type will be injected into two different rule templates that are implemented  
 538 using the model and libraries provided by Siddhi and Flink. In practice, the rule parameters can be  
 539 instantiated in two different rule templates:

```
540 // PROTOTYPE INSTALLATION ADAPTER SERVICE PROCEDURE
541 private bundleContext;
542
543 Install_rule (JsonInstallationObj req) {
544     // Get and configure the right template
545     mavenUrl = get_template_maven_url (req.headers.runtime, req.headers.jobType)
546     ruleTemplate = download_from_maven_repo(mavenUrl)
547     configurationFile = create_config_file(req.headers, req.jobConfig)
548     configFileName = save(configurationFile)
549     deployableRule = inject_config_file(ruleTemplate, configurationFile)
550
551     // Install the rule as an independent Siddhi service
552     if (req.headers.runtime == SIDDHI)
553         job_id = install_OSGi_bundle ( bundleContext, deployableRule)
554
555     // Submit the rule to the remote Flink Cluster
556     if (req.headers.runtime == FLINK)
557         job_id = submit_to_cluster_manager (deployableRule)
558
559     jobInfo = new JobInfo(runtime, jobId, jobType,
560         status.INSTALLED, configurationFileName )
561     return jobInfo
562 }
563 }
```

564 In the form of an OSGi bundle (i.e. a  $\mu$ -service ) encapsulating a Siddhi runtime executing the filtering  
 565 rule; In the form of a Flink job, which can be submitted to a Flink cluster. In this preliminary version of  
 566 the prototype, the Siddhi bundles are installed and executed on the same OSGi runtime of the proxy



567 and adapter  $\mu$ -service. The remote installation on an edge-device can be easily integrated in future.  
568 Instead, the Flink runtime is installed on a remote Kubernetes cluster on the Amazon EKS service.  
569 The two rule templates previously cited are implemented in the form of a JAR file. Both templates  
570 are stored as Maven artifact into a Maven repository. Maven [11] is a tool used for building and  
571 managing any Java-based project and a Maven repository is basically a local or remote directory where  
572 Maven artifacts are stored. A Maven artifact is something that is either produced or used by a project  
573 (e.g. JARs, source, binary distributions, WARs etc.). In this case, both templates are implemented as  
574 JAR files. In order to download a Maven archetype from a Maven repository, an OSGi bundle (i.e.  
575 the adapter  $\mu$ -service) needs only to specify a Maven URL identifying the artifact. Then the URL  
576 resolution and the JAR download is handled by Pax URL [12], a set of URL handlers targeting the  
577 OSGi URL Handler Service. This mechanism is applied by the adapter  $\mu$ -service for downloading  
578 the rule template for installing the rule on the required stream processing engine. The template to be  
579 download (and its relative Maven URL) depends on the jobType and runtime fields specified in the  
580 JSON installation object. Therefore, the adapter  $\mu$ -service must have some predefined information  
581 that bind a Maven URL to a specific jobType and runtime. In this preliminary implementation, the  
582 above mentioned information are stored in memory into a simple hashTable, but for real purposes a  
583 simple database is required. The adapter  $\mu$ -service provides the implementation of the procedures  
584 for installing, starting, stopping, uninstalling and moving the rules and it is responsible for injecting  
585 the rule parameters into the two different templates previously cited. When the adapter  $\mu$ -service has  
586 to install a new rule, considering the jobType (in this case there is only one jobType: a filter) and the  
587 runtime (Siddhi or Flink) specified by the JSON installation object, it downloads the corresponding  
588 JAR file template from the Maven repository. Once obtained, the adapter  $\mu$ -service translates the  
589 jobProps in a configuration file that is injected into the JAR template file. At this point, depending  
590 on the runtime chosen, the template rule is installed in two different ways. In case of a Flink job, the  
591 JAR template is sent to the Flink cluster manager using a REST API offered directly by Flink. On the  
592 other hand, in case of an OSGi bundle implementing the Siddhi filtering application, the bundle is  
593 installed on the Karaf runtime using the OSGi methods offered by the lifecycle layer. In this preliminary  
594 prototype, for the sake of simplicity, the OSGi bundle is installed on the same runtime of the proxy  
595 and adapter  $\mu$ -service, but actually it should be installed on a remote runtime (i.e. a gateway device)  
596 on the edge of the IoT platform. Once the required rule is correctly installed on the target runtime, the  
597 adapter  $\mu$ -service creates a JSON jobInfo object collecting all the relevant information about the just  
598 installed rule. In particular, it keeps trace of a jobId (corresponding to a bundle id for a Siddhi rule  
599 and to a jobId for Flink rule) and a configFileName (corresponding to a unique name of the generated  
600 configuration file, useful for reusing the file when moving the rule for one runtime to another). For  
601 all the other operations (starting, stopping, uninstalling and moving), the adapter  $\mu$ -service uses the  
602 information provided by the jobInfo object and the methods offered by the OSGi lifecycle layer or the  
603 Flink REST API. The Siddhi OSGi bundles are installed on the same runtime of the proxy and adapter  
604  $\mu$ -service, but actually they should be installed on a remote runtime (i.e. a gateway device) on the  
605 edge of the IoT platform. This behaviour has been successfully implemented in Senseioty by FlairBit,  
606 which has extended the OSGi functionalities for communicating with remote runtime and it can be  
607 easily integrated in this prototype implementation in future. In practice, a remote OSGi runtime is  
608 connected to the core platform through two communication channels. A bidirectional channel used for  
609 communicating configurations options and statements. In this scenario, the adapter  $\mu$ -service would  
610 use this channel to notify the target runtime about downloading the required bundle rule: it requires  
611 only a symbolic ID or URL to identify the target runtime. Possible communication protocols adopted  
612 for this channel are MQTT or TCP. A one-directional channel used by the remote OSGi runtime for  
613 download a remote resource, in this case the bundle rule notified by the adapter  $\mu$ -service. A possible  
614 example of communication protocol adopted for this channel is FTP. This communication mechanism  
615 can be used by the adapter  $\mu$ -service for executing all the required interactions with a remote OSGi  
616 runtime (installing, starting, stopping and uninstalling a Siddhi bundle). Another relevant feature

617 implemented by this prototype is the rule composability, meaning that multiple filtering rules can  
618 be concatenated in order to obtain a multiple-step filtering pipeline. Indeed, the currently supported  
619 filtering rules are easily composable because they read and write events from a RabbitMQ broker.  
620 RabbitMQ [19] is an open source message broker supporting multiple messaging protocols and it  
621 was chosen for this prototype implementation because both Siddhi and Flink provide out-of-the box  
622 connectors for consuming and writing event from a RabbitMQ broker. More specifically, RabbitMQ is  
623 adopted in this prototype for handling streams of events in JSON format using the AMQP protocol [1].  
624 The role of an AMQP messaging broker is to receive events from a publisher (event producer) and  
625 to route them to a consumer (an application that processes the event). The AMQP messaging broker  
626 model relies on two main components:

- 627 • *Exchanges*, which are components of the broker responsible for distributing message copies to  
628 queues using rules called bindings. There are different exchange types, depending on the binding  
629 rules that they apply. This prototype uses only exchanges of type direct, which delivers messages  
630 to queues based on a message routing key included when publishing an event.
- 631 • *Queues*, which are the component that collect the messages coming from exchanges. A consumer  
632 reads the events from a queue in order to process the messages.

633 Therefore, when specifying the `jobConfigs` field in the JSON installation object, a client must provide in  
634 the `connectors` field the information needed for reading and writing events from/to an AMPQ queue.  
635 More specifically, is necessary to specify the parameters in the `connectors` field: For specifying the  
636 input source for the event, the following information are needed:

- 637 • *inputEndPoint*: the URL for connecting to the RabbitMQ broker (it might be different from  
638 *outputEndPoint*).
- 639 • *inputExchange*: the name of the exchange from which the input queue will read the messages. If  
640 the exchange does not already exist, it is created automatically.
- 641 • *inputQueue*: the name of the queue that will be bind to the *inputExchange*. If the queue does not  
642 already exist, it is created automatically.
- 643 • *inputRoutingKey*: the routing key that is used for binding the *inputExchange* to the *InputQueue*.

644 On the other hand, for specifying the output source of events, these information are required:

- 645 • *outputEndPoint*: the URL for connecting to the RabbitMQ broker.
- 646 • *outputExchange*: the name of the exchange where to publishing the events. If the exchange does  
647 not already exist, it is created automatically.
- 648 • *outputRoutingKey*: the routing key that is included to the event when publishing it.

649 Leveraging these features, the prototype allows to create stream processing pipelines of arbitrary  
650 complex. For example, multiple Siddhi filters can be concatenated with other filters executed on  
651 Flink. In practice, there is the need of two message brokers: one for the edge level and one for the  
652 cloud/core level. This aspect makes possible to concatenate multiple edge rules without the need  
653 of sending events to a remote broker in the core of the platform, avoiding to introduce unnecessary  
654 latency. RabbitMQ may be a reasonable choice for the cloud/core level scenario, but for the edge level,  
655 the choice must be carefully evaluated for not overloading the edge/gateway devices. For FlairBit and  
656 Senseioty purposes, considering that the edge/gateway devices are provided with an OSGi runtime, it  
657 may be a reasonable choice to take advantage of the OSGi Event Admin Service [16]: an inter-bundle  
658 communication mechanism based on an event publish and subscribe model. This sort of OSGi message  
659 broker can be easily paired with the Remote Service functionalities in order to connect multiple OSGi  
660 runtimes. This solution makes possible to obtain a message broker at edge level, without the need of  
661 adding an external and additional technology. The drawback is that we have to develop a customized  
662 connector implementation for each stream processing engine, in order to consume events from the  
663 OSGi broker

## 664 5. Related Work and Conclusions

665 In this paper we have proposed an adaptive solution for satisfying the dynamic and heterogeneous  
666 requirements that IoT platforms are inevitably facing. During the research and development path that  
667 led to our proposal, we investigated all the features of the microservices architectural pattern, with the  
668 aim of deeply understanding the level of flexibility and dynamicity that this approach is able to offer.  
669 OSGi turned out to be the perfect booster for those dynamic and flexible features that we were looking  
670 for. Then, on the basis of the industrial experience of FlairBit, we formulated a proposal architecture  
671 accompanied by a preliminary prototype implementation. Our solution meets the need to introduce  
672 different real-time stream processing technologies in IoT platforms, in order to offer streaming analytic  
673 functionalities on the different architectural levels of IoT applications. The innovative aspect resides in  
674 a limitation of the expressiveness power for defining stream processing rules, in favour of a much more  
675 flexible and dynamic deployment model. Streaming rules are restricted to a predefined and manageable  
676 set of templates, which allows to handle rules as resources dynamically allocable, composable and  
677 engine independent. These resources can be indifferently deployed at edge-level or core-level and  
678 moved around at any time.

679 Comparing our proposal with similar real-time streaming functionalities offered by the IoT  
680 platforms of Amazon, Azure and Google, the dynamic features of our solution can be potentially  
681 promising and innovative. Amazon offers AWS IoT Greengrass [16] as a solution for moving analytical  
682 functionalities directly on edge devices. It is basically a software that once installed on an edge  
683 device enables the device to run AWS Lambda functions locally. AWS Lambda enables to run code  
684 without provisioning or managing servers. They offer a great level of expressivity with respect to  
685 our proposal because they support function implemented with all the most common programming  
686 languages. However, AWS IoT Greengrass does not provide any functionality for dynamically moving  
687 the Lambda computation back and forth between the edge-level and cloud-level and it is bound  
688 to the Lambda execution model. It does not offer any integration with external stream processing  
689 engines, which on the other hand can be integrated in our solution as pluggable components as long  
690 as template implementations of the supported rule types are provided. Microsoft Azure offers similar  
691 functionalities with Azure Stream Analytics on IoT Edge [13]. It empowers developers to deploy  
692 near-real-time analytical intelligence, developed using Azure Stream Analytics, to IoT devices. The  
693 principle is the same of AWS IoT Greengrass: installing the Azure IoT Edge software we enable the  
694 edge devices to locally execute Azure Stream Analytics rules. Azure Stream Analytics is a real-time  
695 analytics and complex event-processing engine where streaming rules and jobs are defined using a  
696 simple SQL-based query language. Again, the power of expressiveness is much wider with respect to  
697 our proposal, but the resulting solution is inevitably bound to the only Azure Stream Analytics engine  
698 and no mechanisms for the dynamic relocation of rules between edge and cloud are provided. Finally,  
699 Google Cloud IoT [18] integrates the Apache Beam SDK [21], which provides a rich set of windowing  
700 and session analysis primitives. It offers a unified development model for defining and executing  
701 data processing pipelines across different stream processing engines, including Apache Flink, Apache  
702 Samza, Apache Spark and other engines. However, Apache Beam supports only scalable engines  
703 suitable for the core-cloud level and it is not designed for supporting edge analytics.

704 Concerning future research directions, one possibility is to improve the architecture by  
705 investigating possible solutions for simplifying the rules' definition. Our API requires to define  
706 a JSON object containing the rules' parameters, but for example a web interface or an SDK similar  
707 to Apache Beam may offer a higher level approach. In this case, it is required to identify the right  
708 trade-off between the level of expressiveness offered by a possible unified model or language and  
709 the limits imposed by the presence of predefined rule types and templates. Another interesting point  
710 consists in integrating a monitoring  $\mu$ -service in the application. This monitoring functionality, which  
711 is presented as possible application scenario of our proposal, can be formalized in more details in order  
712 to become an integral part of our solution. Providing an out-of-the-box monitoring behaviour can be a  
713 powerful additional feature useful in many IoT use cases.

714

- 715 21. Senseioty. <http://senseioty.com/>.
- 716 25. Fowler, M.; Lewis, J. Microservices - A definition of this new architectural term. <https://martinfowler.com/articles/microservices.html>, 2014.
- 717
- 718 24. Abbott, M.; Fisher, M. *The Art of Scalability: Scalable Web Architecture, Processes, and Organizations for the*
- 719 *Modern Enterprise*; Addison-Wesley Professional, 2015.
- 720 31. Robert, M. *Agile Software Development: Principles Patterns And Practices*; Pearson, 2003.
- 721 28. Richardson, C. Building Microservices: Using an API Gateway. <https://www.nginx.com/blog/introduction-to-microservices/>, 2015.
- 722
- 723 29. Richardson, C. Event-Driven Data Management for Microservices. <https://www.nginx.com/blog/event-driven-data-management-microservices/>, 2015.
- 724
- 725 30. Richardson, C. Choosing a Microservices Deployment Strategy. <https://www.nginx.com/blog/deploying-microservices>, 2016.
- 726
- 727 14. OSGi Alliance. <https://www.osgi.org>.
- 728 15. OSGi Architecture. <https://www.osgi.org/developer/architecture/>.
- 729 27. Hall, R.; Pauls, K.; McCulloch, S.; Savage, D. *OSGi in Action, Creating Modular Applications in Java*; Manning, 2011.
- 730
- 731 17. OSGi Service Layer. <https://osgi.org/specification/osgi.core/7.0.0/framework.service.html>.
- 732 20. Remote Services Specification. <https://osgi.org/specification/osgi.cmpn/7.0.0/service.remoteservices.html>.
- 733
- 734 3. Apache CXF Distributed OSGi. <https://cxf.apache.org/distributed-osgi.html>.
- 735 9. Apache Zookeeper. <https://zookeeper.apache.org/>.
- 736 6. Apache Karaf. <https://karaf.apache.org/>.
- 737 23. A. Arvind, B.S.; Jennifer, W. The CQL Continuous Query Language: Semantic Foundations and Query Execution. <http://ilpubs.stanford.edu:8090/758/1/2003-67.pdf>, 2003.
- 738
- 739 22. Siddhi Streaming and Complex Event Processing System. <https://siddhi.io/>.
- 740 10. Bnd tools. <https://bnd.bndtools.org/>.
- 741 5. Apache Ignite. <https://ignite.apache.org/>.
- 742 7. Apache Samza. <http://samza.apache.org/>.
- 743 4. Apache Flink. <https://flink.apache.org/>.
- 744 8. Apache Storm. <https://storm.apache.org/>.
- 745 11. Kafka Streams. <https://kafka.apache.org/documentation/streams/>.
- 746 2. Apache CXF. <http://cxf.apache.org/>.
- 747 12. Kubernetes. <https://kubernetes.io/>.
- 748 19. RabbitMQ. <https://www.rabbitmq.com/>.
- 749 1. AMQP Protocol. <https://www.amqp.org/>.
- 750 16. OSGi Event Admin Service. <https://osgi.org/specification/osgi.cmpn/7.0.0/service.event.html>.
- 751 13. Maven. <https://maven.apache.org/>.
- 752 18. Pax URL. <https://ops4j1.jira.com/wiki/spaces/paxurl/overview>.

753 © 2020 by the authors. Submitted to *Journal Not Specified* for possible open access  
754 publication under the terms and conditions of the Creative Commons Attribution (CC BY) license  
755 (<http://creativecommons.org/licenses/by/4.0/>).