

D4.1 Medolution Platform APIs and Specification V1

Medolution

Medical Care Evolution



ITEA3 – Project 14003

Document Properties

Edited by :	François Exertier, Bull
Authors	François Exertier (Bull), Mathis Gavillon (Bull), David Kuik (Norima), Mihai Mitrea (IMT), Anil Sinaci (SRDC), Béchir Taleb Ali (Prologue), Wolfgang Thronicke (Atos DE)
Date	22/11/2016
Visibility	Public
Status	Final

History of Changes

Release	Date	Author, Organization	Changes
0.1	03/2016	F. Exertier, Bull	Outline creation
0.2	07/2016	M. Gavillon, Bull	Bull contribution
0.3	02/09/2016	F. Exertier, Bull	Review and improve Bull contribution
1.0	21/10/2016	F. Exertier, Bull	Integration of contributions from Norima, IMT, SRDC, Atos DE, Prologue
1.1	23/10/2016	F. Exertier, Bull	Minor corrections
1.2	09/11/2016	F. Exertier, Bull	Minor updates from SRDC, Bull has added executive summary, introduction, elements on security, minor updates, completion of section 3.1 about platform positioning, section 8 Hosting Platforms, restructure section 4.5, completion of section 5, references
1.3	10/11/2016	F. Exertier, Bull	Minor updates, ready for review version
1.4	15/11/2016	D. Kuik, Norima	Review and adapt for grammar, formatting, and selective content areas.
1.5	18/11/2016	F. Exertier, Bull	Integrate review suggestions from D. Kuik
1.6	22/11/2016	F. Exertier, Bull	Final version, reviewed by Sopheon (H. Rutten), and Technolution (H. van den Brink)



Table of Contents

HISTORY OF CHANGES	2
1. EXECUTIVE SUMMARY.....	5
2. INTRODUCTION.....	6
3. MEDOLUTION PLATFORM DESCRIPTION	7
3.1. DATA AND SERVICES HUB PLATFORM.....	7
3.2. COMPONENTS CATALOGUE AND TOPOLOGIES.....	8
3.3. ORCHESTRATOR AND TARGET INFRASTRUCTURES	9
3.4. MEDOLUTION SECURITY AND PRIVACY	9
4. THE MEDOLUTION BIG DATA COMPONENTS CATALOGUE	10
4.1. CORE FOUNDATIONAL COMPONENTS	10
4.2. DATA COLLECTION AND STORAGE COMPONENTS.....	10
4.2.1. <i>Elasticsearch</i>	10
4.2.2. <i>Kafka</i>	11
4.2.2.1. <i>Kafka API</i>	12
4.2.2.2. <i>TOSCA component description</i>	12
4.2.2.3. <i>Kafka Relationships</i>	14
4.2.3. <i>Logstash</i>	14
4.2.4. <i>HBase</i>	15
4.2.4.1. <i>HBase Master</i>	15
4.2.4.2. <i>HBase Region Server</i>	16
4.2.5. <i>Hive</i>	17
4.2.6. <i>Drill</i>	17
4.2.7. <i>Flume</i>	18
4.3. DEVICE CONNECTOR COMPONENT	19
4.3.1. <i>Devices to be connected</i>	19
4.3.1.1. <i>Bluetooth enabled medical devices</i>	20
4.3.1.2. <i>Activity tracker wristband</i>	20
4.3.1.3. <i>Smart phone</i>	21
4.3.2. <i>Devices particularities with respect to data connexion</i>	21
4.3.3. <i>Medolution device connector when device data are exposed</i>	22
4.3.3.1. <i>Data to be collected</i>	22
4.3.3.2. <i>Device Connector Application Program Interface (API)</i>	24
4.3.3.2.1. <i>Binary payload</i>	24
4.3.3.2.2. <i>JSON payload</i>	25
4.3.4. <i>Medolution device connector when device data are not exposed</i>	26
4.3.5. <i>Conclusion</i>	26
4.4. DATA CONNECTOR COMPONENTS.....	27
4.4.1. <i>Sqoop</i>	27
4.4.1.1. <i>MapRSqoop2Server</i>	27
4.4.1.2. <i>MapRSqoop2Client</i>	28
4.5. DATA ANALYTICS COMPONENTS	28
4.5.1. <i>Rstudio</i>	28



4.5.2.	<i>Spark</i>	29
4.5.3.	<i>Pig</i>	30
4.5.4.	<i>Mahout</i>	31
4.6.	VISUALIZATION COMPONENTS.....	32
4.6.1.	<i>Kibana</i>	32
4.6.1.1.	Kibana Interface.....	32
4.6.1.2.	Kibana Configuration.....	32
4.7.	DATA DISTRIBUTION & RESOURCE MANAGEMENT COMPONENTS.....	33
4.7.1.	<i>MapR Filesystem</i>	33
4.7.1.1.	MapRCldb.....	33
4.7.1.2.	MapRFileserver.....	34
4.7.2.	<i>YARN</i>	35
4.7.2.1.	YARN Resource manager.....	35
4.7.2.2.	YARN Node manager.....	36
5.	MAIN MEDOLUTION BIG DATA TOPOLOGIES.....	39
5.1.	LAMBDA ARCHITECTURE.....	39
5.2.	CONCRETE ARCHITECTURE OF THE LVAD MEDICAL USE CASE.....	41
6.	SECURITY.....	44
6.1.	SECURITY PRINCIPLE AND ARCHITECTURE.....	44
6.2.	DATA ANONYMIZATION.....	44
6.2.1.	<i>Data Provider Interface</i>	45
6.2.2.	<i>Data Custodian Interface</i>	46
6.2.3.	<i>Data Access Interface</i>	46
6.2.4.	<i>Data Anonymization Concepts</i>	46
6.2.4.1.	Generalization.....	47
6.2.4.2.	Re-Identification Risk.....	47
6.2.4.3.	Suppression.....	48
6.2.4.4.	Quantifying Data Loss.....	49
6.3.	MANAGING ACCESS RIGHTS.....	50
6.4.	SECURED DATA TRANSMISSION.....	50
7.	MEDOLUTION PLATFORM APIS.....	51
8.	HOSTING PLATFORMS.....	52
9.	LIST OF FIGURES.....	53
10.	LIST OF TABLES.....	54
11.	REFERENCES.....	55

1. Executive Summary

This document is the first version of the Medolution Platform APIs and Specification. It describes the Medolution Core platform as of the first year of the project. The main role of the Medolution Core platform is to provide the execution support for Medolution applications in term of Big Data related services, software engineering support (design, deployment), and connection to devices and data warehouses. It makes the link between data sources (data warehouses, devices defined in WP3...), and the applications developed in WP5, that will be built using big data related services provided by the platform, and that will potentially run on the platform.

At this stage of the project, architecture and requirements were not entirely defined, and not all partners involved in the work package had the opportunity to contribute. This means this specification will be greatly upgraded in its second version, including the contributions of all partners.

Therefore, this first version of the document describes the platform principles, as initiated in the FPP document and based on the Bull Big Data Capabilities Framework (BDCF). BDCF provides capabilities to build Big Data Applications by composing Software components delivered in an associated catalogue, and then to deploy it on any Cloud but, in particular, on the WP4 hosting platform, which provides hardware features particularly suitable to Big Data processing.

This first version of the document includes the description of most of the components delivered with BDCF that are suitable for building Medolution applications dealing with Big Data processing. These are typically generic Big Data processing components and blueprints, including those related to Big Data Real Time stream processing. In addition, Atos DE has provided a blueprint representative of the LVAD (Left Ventricular Assist Device) Use Case, which should be implemented as a BDCF topology in the next phase of the project.

Among platform components, the device connector components have the role to collect data generated by devices and to route it to big data processing components. A first version of the specification on how such components should be handled by the platform has been provided by SRDC and IMT, based on Medolution requirements and on devices targeted in the project.

Finally, security is addressed. In particular, Norima has provided a first version of the specification of components implementing anonymization.

The APIs of the WP4 platform are provided in two categories:

- first, the APIs of the catalogue components which will be used by applications are described in section 4,
- and second, the platform API used to operate the platform are described in section 7.

2. Introduction

This document describes the WP4 platform and its APIs which satisfies the following objectives defined in the FPP:

- Software Platform
 - To connect Heterogeneous Healthcare Data Sources (Sensors, Data Warehouses, Internet sources...)
 - To provide Big Data related technical services (mediation, storage, analysis...)
 - To connect, compose and host applicative services (WP5) based on Big Data services
- Reactive Big Data Framework
 - low latency, deliver insights on-the-fly to healthcare staff
 - based on experimenting
 - in memory processing, stream processing, Lambda architectures
- Software Engineering
 - build and deploy easily “à la carte” Big Data Software Stacks
- Security
 - encryption, anonymization, tracking

As mentioned in Medolution deliverable D1.1 about the State of the Art Analysis, regarding the IoT Big Data platforms, the innovative aspects of this Core Platform will be to reduce the complexity of building Big Data software applicative stacks, however still allowing components choice and facilitating cloud deployment; the focus will also be set on integrating advanced real-time Big Data processing technologies. Taking into account access to Healthcare devices and Data Privacy aspects in this platform will also be part of the solution.

Section 3 describes the principle of the platform, i.e. a composition PaaS based on a catalogue of Big Data and Medolution related software components, providing scalability, reliability and security.

Section 4 describes the platform catalogue components, currently mainly generic big data processing related components (storage, collect, analysis, and visualization) delivered with the BDCF software, as well as a first specification of the Medolution Device Connector, which will enable connection of the IoT devices of WP3 with the WP4 platform.

Section 5 is dedicated to blueprints that will be delivered with the platform to satisfy Medolution applications architecture requirements. It currently provides a first real time streaming processing blueprint that is delivered as a topology in the BDCF catalogue, and a first design of a blueprint which should support the LVAD Medolution use case.

Section 6 about Security currently contains a first version of the specification of anonymization components; general security features as supported in BDCF are briefly described.

Finally, section 7 provides a short description of the API to operate the platform.

3. Medolution Platform Description

This section describes the principles of the Medolution Platform.

3.1. Data and services hub platform

The Medolution platform aims to support whole Big Data capacity and IoT connectors as a service. It is composed by three main layers:

- **Virtualization:** hardware resources are shared between all applications. An Infrastructure as a Service (IaaS) provides the ability to allocate these resources on demand.
- **Orchestration:** each application requires some resources (compute, network, disk space) which are allocated on any IaaS platform.
- **Provisioning:** manage software requirements, installation and configuration.

The Big Data Capability Framework (BDCF) aims to provide those three layers. It brings a packaged solution to create Big Data application clusters on demand, through an intuitive administrative user interface. Big Data applications initially targeted by BDCF are mainly distributed storage and processing (Hadoop) applications based on MapR or Hortonworks and Log Analysis applications based on Elastic Stack components. In addition, BDCF provides other useful components like a message broker (Kafka), a consensus system (Consul), a studio for data scientists (RStudio) and other technical components like Java and Python, allowing a detailed technical architecture design.

All components described above are part of a catalogue and can be connected all together to create distributed applications. Catalogue and application design is managed through a software named Alien4Cloud which provides an intuitive graphical user interface (GUI). It is a tool, part of BDCF, that aims to provide management for enterprise cloud and help enterprise to move their applications to a cloud, and based on project constraints, reach continuous delivery.

As moving to the cloud for an enterprise is a structural change, Alien4Cloud leverages the TOSCA¹ standard that is the most advanced and supported standard for the cloud. TOSCA stands for Topology and Orchestration Specification for Cloud Applications, it is an OASIS open standard that defines the interoperable description of services and applications hosted on the cloud and elsewhere; including their components, relationships, dependencies, requirements, and capabilities, thereby enabling portability and automated management across cloud providers regardless of underlying platform or infrastructure. These characteristics also facilitate the portable, continuous delivery of applications (DevOps) across their entire lifecycle.

The usage principle of BDCF is illustrated in the figure below, where the application designer defines the blueprints of the components (also known as topology) which compose her/his application software suite through the graphical TOSCA tool (Alien4Cloud), selecting her/his components in the BDCF catalogue, assembling and configuring them, and then benefiting from the automated orchestration of this topology.

¹ <https://www.oasis-open.org/committees/tosca>

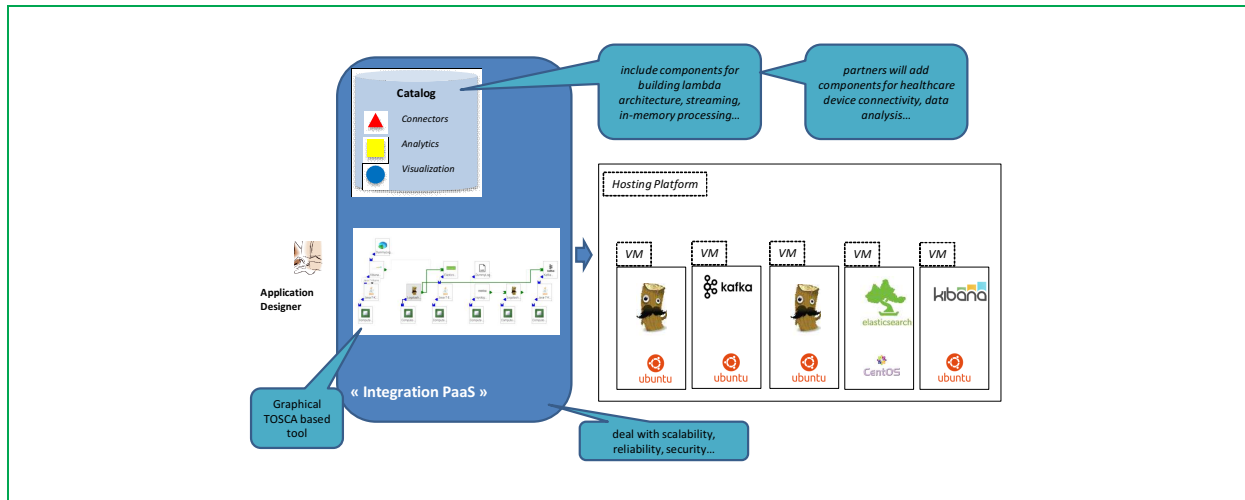


Figure 1: BDCF Platform

The Orchestration layer also brings reliability and scalability, which provides the elasticity to use the right amount of resources.

- **Reliability:** algorithms monitor the health of software components and virtual machine state. If one of those is falling down, the orchestration tool is able to instantiate new resources, to reconfigure and/or restart the service.
- **Scalability:** application designer can scale out some resources at deployment step or dynamically allocate more or fewer resources post deployment according to its needs.

This framework fits in the Medolution architecture as it enables the setting up and deployment of Big Data Applications that will:

- rely on big data analysis components to implement the smart Medolution applications of WP5 and the Medolution Use Cases. As the Medolution requirements are elaborated during the project, additional components and topologies (blueprints representing components assemblies) will be added to the platform catalogue.
- rely on connector components to interconnect Medolution identified healthcare devices (WP3) and Medolution healthcare data warehouses to the data processing services. Such connector components will also be introduced during the project.

3.2. Components catalogue and topologies

Big Data Capability Framework (BDCF) provides a catalogue of tools which allow Big Data actors to store, analyse, explore, and visualize data coming from some different kinds of sources. Those can be IoT devices, applications logs, etc...

In Medolution platform, these tools are described following the TOSCA standard. Each of them is called a component. Alien4Cloud allows connecting these components to each other to create a distributed application. This arrangement is called a topology.

Once it has been designed, the application can be deployed, which will install, setup and start software components.

Each component which follows the **TOSCA** standard must define at least some elements:



- **Properties:** fields which allow end-users to customize this component and could be used at installation, configuration or start step.
- **Attributes:** properties which will be generated at runtime.
- **Requirements:** define a dependency from this node to other ones.
- **Capabilities:** simple description of what kind of operation the component will do.
- **Artefacts:** attached items which will be uploaded and used during component instantiation.

Creating a BDCF component from existing software is a matter of defining the associated TOSCA elements and developing the scripts that will allow configuring, installing, starting, and stopping the component. A BDCF component developer guide [1] has been written and provided to Medolution partners (available on the project SharePoint).

Components of BDCF to be used in Medolution, as well as those to be developed and included in the catalogue for Medolution, are described in section 4.

3.3. Orchestrator and target Infrastructures

Big Data Capability Framework relies on an orchestrator to create applications over a Cloud infrastructure. The aim is to coordinate Cloud resources to be able to deploy a defined distributed application, based on its TOSCA description. The orchestrator will check available resources, will hold it and then use it. Some post deployment step can be defined to manage resiliency, scaling and monitoring. The current orchestrator supports some Cloud infrastructure like Amazon Web Services, VMWare or Azure. Medolution infrastructure is currently based on Openstack, an Open source Infrastructure as a Service (IaaS) solution.

3.4. Medolution security and privacy

The main goal in the Medolution project about security in WP4 is to ensure Data Privacy in the context of Healthcare Big Data processing. This involves dealing with data privacy at the level the Big Data storage capabilities, and at the level of Data collection capabilities, i.e. HealthCare Data Warehouse connectors (like EHR connectors), Device Connectors.

In this first stage of the project only two aspects have been addressed:

- The Big Data Platform security by itself, i.e. securing data access and data communication with the platform itself
- Data Anonymization capabilities, where a first level of specification regarding Data Anonymization components has been provided.

Both are described in section 6.

4. The Medolution Big Data components catalogue

This section describes the components that will be made available in the Platform catalogue to enable Medolution application providers to compose their application.

This includes:

- the TOSCA description of each component and corresponding relationships (how they connect to other components),
- their APIs,
- their high availability and scalability capabilities
- their security features

4.1. Core Foundational Components

Some of the components described below may use one or both the technical components below

- The Java component is a technical component allowing other software components to choose the required Java version.
- Consul is a technical component allowing other software components to discover each other in a flexible, highly available and fault tolerant way. The same Consul component is used for two different purposes; it can run as a Consul Server or as a Consul Agent. Consul Servers are responsible for storing and replicating data between themselves, while Consul Agents are responsible for registering functional services, store services distributed configuration, performing health checks on nodes hosting those services and serving service discovery requests through their DNS interface.
- The Compute Node component represents a Virtual Machine to be allocated and on which the components will run.

4.2. Data collection and storage components

This section describes a set of BDCF components which are concerned with collecting data from various sources and storing them for future processing. Subsequent sections will describe components for processing data.

4.2.1. Elasticsearch

Elasticsearch is a search server based on Apache Lucene. It provides a distributed, multitenant-capable full-text search engine with a RESTful Web interface and schema-free JSON documents.

In the Elastic Stack architecture, Elasticsearch is the database that stores the logs sent by Logstash, and provides to Kibana the needed data for visualization.

The infrastructure administrator needs to understand the following Elasticsearch internal concepts: For Elasticsearch underlying concept, management and configuration refer to the official documentation: <https://www.elastic.co/guide/en/elasticsearch/reference/current/index.html>

For more information on Apache Lucene, refer to <http://lucene.apache.org/>.

Elasticsearch API

The Elasticsearch tool is accessible via a REST API which offers a way to read and write data on the datastore.

Elasticsearch Relationships

The only prerequisite of the Elasticsearch component is to be attached to a Java component that is configured to use Java 7 or greater. This Java component must be hosted on a Compute component.

The Elasticsearch component must be connected to a Consul agent hosted on the same Compute node, to perform cluster discovery. This connection to a Consul agent is prerequisite when:

- the Elasticsearch cluster has several nodes, or
- the Elasticsearch search_endpoint capability is used.

Elasticsearch has the capability to be used as a search_endpoint by the Logstash and Kibana components. This means that you can connect the components that use Elasticsearch by setting their search_endpoint prerequisite.

In the Elastic Stack chain, Logstash and Kibana use Elasticsearch as target database. The data is stored in indexes dedicated to those components.

Elasticsearch includes indexes containing the logs provided by Logstash. By default Logstash creates one different index per day, named Logstash-YYYY-MM-dd. It is possible to override this name in the Logstash configuration file by modifying the index option.

Kibana stores its configuration into Elasticsearch under the kibana index. Do not modify it manually or remove it if you want Kibana to work properly.

4.2.2. Kafka

Kafka is a distributed, partitioned and replicated publish-subscribe messaging system.

In the Elastic Stack architecture, Kafka offers the following features:

- Store into a Highly Available, fault tolerant system, the logs ingested by the Elastic Stack before their indexing by Logstash.
- Load-Balance data across the Kafka cluster and between its subscribers (Logstash Indexer in our case) by natively partitioning data across the cluster.
- Decouple the processing part of the Elastic Stack architecture from the ingestion part by implementing a retention queue.

The Kafka community uses the following terminology:

- Kafka maintains feeds of messages in categories named topics. Topics are split into partitions, which brings data reliability between all nodes of the Kafka cluster;
- Processes that publish messages to a Kafka topic are named producers
- Processes that subscribe to topics and process the feed of published messages are named consumers.
- Kafka runs as a cluster of one or more servers. Each of these servers is named a broker.

4.2.2.1. Kafka API

Kafka provides some Java API which allows a user to consume, publish or stream the content of a Kafka broker.

4.2.2.2. TOSCA component description

Kafka

In the Elastic Stack architecture, Kafka topics are connected with a Logstash instance named the shipper, which is responsible to ingest logs from different inputs (for instance Rsyslog or lumberjack) and which acts as the producer by publishing those logs into Kafka topics. On the other hand, Kafka topics are connected with another Logstash instance named the indexer, which consumes published logs in order to process them.

A Kafka node is hosted on a Java node, which is hosted itself on a Compute node. A Kafka node requires to be related to a Consul agent hosted on its Compute node. On top of this stack you can deploy as many different topics as you need. Each topic has its own configuration.

The minimum and recommended version of Java is JRE7.

Properties

kf_heap_size: This property allows setting the heap memory size that is allocated to Kafka java process, It allocates the same value to both initial and maximum values (ie -Xms and -Xmx java options).

default: "1G"

zk_heap_size: This property allows setting the heap memory size that is allocated to Zookeeper java process (ZooKeeper is a coordination service for distributed applications, the corresponding BDCF component, built on the MAPR Zookeeper element, is not detailed in this documentation). It allocates the same value to both initial and maximum values (ie -Xms and -Xmx java options).

default : "500M"

repository: This property gives the opportunity to specify an alternative download repository for this component artefacts. It is your responsibility to provide an accessible download url and to store required artefacts on it. You should specify only the base repository url. Artifacts names will be appended to it, so this property could be shared among several components using the inputs feature.

default : ""

Requirements

java: Kafka should be hosted on a Java component. Java 7 or greater is required.

host: Kafka component has to be hosted on a Compute.

consul: Kafka component has to be connected to a local (hosted on the same Compute) Consul Agent. This is required to perform cluster discovery.

filesystem_endpoint: *Kafka may be connected to a filesystem in order to store its runtime data on it. A typical use case would be to link this filesystem to a block storage in order to achieve data resilience and recovery.*

Kafka Topic

A Kafka Topic should be hosted on a Kafka component and may be configured through the following properties.

Properties

topic_name: *The topic name (value should match the following pattern: [-_A-Za-z0-9]+)*

default: ""

partitions: *Number of partitions for this topic*

default : 1

replicas: *Number of replicas for this topic. Should be at most the number of hosting Kafka Component instances.*

default : 1

min_in_sync_replicas: *When a producer sets request_required_acks to in_syncs, min_insync_replicas specifies the minimum number of replicas that must acknowledge a write for the write to be considered successful. If this minimum cannot be met, then the producer will raise an exception (either NotEnoughReplicas or NotEnoughReplicasAfterAppend). When used together, min_insync_replicas and request_required_acks allow you to enforce greater durability guarantees. A typical scenario would be to create a topic with a replication factor of 3, set min_insync_replicas to 2, and produce with request_required_acks of in_syncs. This will ensure that the producer raises an exception if a majority of replicas do not receive a write.*

default : 1

retention_minutes: *The number of minutes to keep a log file before deleting it.*

default: 10080 (7 days)

segment_minutes: *The number of minutes after which Kafka will force the log to roll even if the segment file isn't full to ensure that retention can delete or compact old data.*

default: 10080 (7 days)

segment_bytes: *Segment file size for the log.*

default: 1073741824 (1GB)

Requirements

kafka_host: *Kafka topics are hosted on Kafka components.*

4.2.2.3. Kafka Relationships

Any Kafka node is related to a Consul agent hosted on the same Compute node. This relationship is obtained by binding the consul prerequisite of the Kafka node to the agent capability of the Consul node.

As explained above, in the Elastic Stack architecture Kafka topics are connected with a Logstash that publishes messages and another Logstash that consumes those messages.

4.2.3. Logstash

Logstash is a tool for receiving, processing and outputting logs. All kinds of logs are concerned: system logs, webserver logs, error logs, application logs, etc. Logstash provides a powerful pipeline for storing, querying, and analysing logs. It includes an arsenal of built-in inputs, filters, codecs, and outputs.

The Logstash event processing pipeline includes three stages:

Inputs → Filters → Outputs

Inputs generate events, Filters modify them, and Outputs ship them elsewhere.

- Inputs are used to get data into Logstash (see <https://www.elastic.co/guide/en/logstash/2.3/input-plugins.html>). Some of the most commonly-used inputs are:
 - **File:** Reads from a file on the filesystem.
 - **Syslog:** Listens on the port 514 for syslog messages and parses according to the RFC3164 format.
 - **Lumberjack:** Processes events sent in the lumberjack protocol.
 - **Elasticsearch:** Reads query results from an Elasticsearch cluster.
- Filters are intermediary processing devices in the Logstash pipeline (see <https://www.elastic.co/guide/en/logstash/2.3/filter-plugins.html>). You can combine filters with conditionals to perform an action on an event if it meets certain criteria. Some useful filters include:
 - **Grok:** Parse and structure arbitrary text. Grok is currently the best way in Logstash to parse unstructured log data into something structured and queryable.
 - **Mutate:** Perform general transformations on event fields. You can rename, remove, replace, and modify fields in your events.
 - **Drop:** Drop an event completely, for example, debug events.
 - **Clone:** Make a copy of an event, possibly adding or removing fields.
 - **Geoip:** Add information about geographical location of IP addresses.
- Outputs are the final phase of the Logstash pipeline (see <https://www.elastic.co/guide/en/logstash/2.3/output-plugins.html>). An event can pass through multiple outputs, but once all output processing is complete, the event has finished its execution. Some commonly used outputs include:
 - **Elasticsearch:** Send event data to Elasticsearch. If you plan to save your data in an efficient, convenient, and easily queryable format Elasticsearch is the way to go.
 - **File:** Write event data to a file on disk.

A Logstash node should be hosted on a Java node, which is hosted itself on a compute node.

The TOSCA component description exposes the following properties and artefacts:

Properties

version: *Version number*

repository: *Allows you to specify an alternative download repository for a Logstash binary.*

auto_reload: *Boolean property set to false by default. If set to true, Logstash monitors configuration changes and reloads configuration whenever it is changed.*

reload_interval: *If auto-reload is true, this property specifies how frequently to poll the configuration location for changes, in seconds.*

stdout: *Boolean property set to false by default. When set to true, the simple stdout output is added to the output configuration of Logstash.*

heap_size: *Allows setting the heap memory size for Logstash java process. Default value is 500M. It allocates the same value to both initial and maximum values (ie -Xms and -Xmx java options).*

log_level: *Define Logstash log level. By default it is quiet and very few logs are generated by Logstash. All logs are redirected to a file except if you set 'stdout' to 'true'.*

Artefacts

input_conf: *input configuration*

output_conf: *output configuration*

filter_conf: *filter configuration*

4.2.4. HBase

HBase is a non-relational database based on Hadoop distributed filesystem which allows storing billions of rows of millions of columns. This is a distributed and scalable database system.

HBase is made up of two distinct components: HBase Master and HBase Region server.

4.2.4.1. HBase Master

HBase Master is the database system endpoint. It allows user to ask for accessing or pushing data inside its database and manages HBase Region Server instances.

HBase API

HBase provides some Java classes which allow to interact with HBase master and to create table, upload, read data... A Web interface is also available for end user.

TOSCA component description

HBase Master implementation in BDCF is part of MapR distribution. The component name in the BDCF catalogue is MapRHBaseMaster.

Properties

MapRHBaseMaster does not have any properties.

Requirements

MapRHBaseMaster requires to be hosted on MaprWarden².

Capabilities

hbase_master_endpoint: *allows connecting any Alien4Cloud component to this one. It can be used for MapRHUE³ requirement.*

Artefacts

tools: *MapR utility scripts.*

utils_scripts: *Common utility scripts for whole BDCF components.*

4.2.4.2. HBase Region Server

HBase Region server is the low level database management component. It stores data and answer user requests. In general, several instances are instantiated.

TOSCA component description

HBase Region Server implementation in BDCF is part of MapR distribution. The component name in A4C catalogue is MapRHBaseRegionServer.

Properties

MapRHBaseRegionServer does not have any properties.

Requirements

MapRHBaseRegionServer requires to be hosted on MaprWarden.

Capabilities

hbase_region_server_endpoint: *allows connecting any Alien4Cloud component to this one.*

Artefacts

tools: *MapR utility scripts.*

² Warden is a light Java application that runs on all the nodes in a MAPR Hadoop cluster and coordinates cluster services. The corresponding BDCF component is not detailed in this document.

³ HUE is the User Interface to interact with Apache Hadoop and its ecosystem components, such as Hive, Pig, and Oozie. The corresponding BDCF component is not detailed in this document.

utils_scripts: Common utility scripts for whole BDCF components.

4.2.5. Hive

Hive is a data warehouse system for Hadoop that facilitates easy data summarization, ad-hoc queries, and the analysis of large datasets stored in Hadoop-compatible file systems, such as the MapR Data Platform (MDP). Hive provides a mechanism to project structure onto this data and query the data using a SQL-like language called HiveQL. At the same time this language also allows traditional map/reduce programmers to plug in their custom mappers and reducers when it is inconvenient or inefficient to express this logic in HiveQL.

Hive API

Hive provides Java API to interact with HiveServer2 and Hive Metastore.

TOSCA component description

Properties

MapRHive does not have any properties.

Requirements

MapRHive requires to be hosted on MapRWarden. If you have multiple MapRHive components, you need to resolve these needs:

hive_registerTo_haproxy: can be connected to a HAProxyTCPHive component to manage load-balancing.

hive_connectsTo_mysql: can be connected to a MySQLDatabase component to externalize the Hive Metastore.

Capabilities

hive_endpoint: allows connecting any Alien4Cloud component to this one. It can be used for MapRHUE requirement.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

hive_files: Required MapRHive template configuration files.

4.2.6. Drill

Apache Drill is an open source, low-latency query engine for big data that delivers secure and interactive SQL analytics at petabyte scale. With the ability to discover schemas on-the-fly, Drill is a pioneer in delivering self-service data exploration capabilities on data stored in multiple formats

in files or NoSQL databases. Drill is fully ANSI SQL compliant and integrates seamlessly with visualization tools.

Drill API

Drill provides a REST API to connect and interact with the service running queries, performing storage plugin tasks, such as creating a storage plugin, obtaining profiles of queries, and getting current memory metrics. A Web UI is also available.

TOSCA component description

Properties

MapRDrill does not have any properties.

Requirements

MapRDrill requires to be hosted on MapRWarden.

Capabilities

drill_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.2.7. Flume

Apache Flume provides a way to easily and efficiently move data from a source to a sink. It is based on a simple and flexible architecture which uses streaming data flows.

Flume API

Flume offers a way to interact with its service using a Java API.

TOSCA component description

Flume implementation in BDCF is part of MapR distribution. The component name in A4C catalogue is MapRFlume.

The output configuration is bound to MapR filesystem, although no relationship is necessary for that purpose.

When MapRFlume is not connected to Kafka, it is configured to open a port with Netcat on port 44444 and store data in MapR filesystem.

Properties

component_version: Version of the Flume component integrated

default: 1.6.0

hdfs_path: Flume is configured to use HDFS as sink. An alternative path in MapRFS can be provided.

default: maprfs:///user/mapr/flume/kafka_input/%y-%m-%d-%H

Requirements

MapRFlume requires to be hosted on MapRWarden

kafka_input: Input must be set to a Kafka Topic

Capabilities

flume_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

flume_files: Required Flume template configuration files. An example of this file is given in BDCF user documentation.

4.3. Device connector component

Another way to collect data is to get it directly from devices. A device connector component is a Medolution component available in the platform catalogue, that can be easily incorporated in an application topology in order specify the input data flow for such an application. Its purpose is to collect a data flow from outside the platform, and to route it toward any Big Data component to be used for the application needs, i.e. an analytics processing component, a data storage component, etc.

In this section, different kinds of devices to be connected in the Medolution project are described, with standard related information (section 4.3.1). The issue with the devices whose data can only be accessed from proprietary application is explained in section 4.3.2, which results in two categories of devices to be considered. For the first category, devices whose data is directly accessible (section 4.3.3), the data to be collected with their respective standard and formatting information is described in section 4.3.3.1, and the Device Connector Component external API, i.e. the API to be used from outside the Medolution platform to inject data, is described in section 4.3.3.2. In a next version of this document, the Device Connector component description will be completed with its TOSCA description (properties, requirements) and relationships with other components of the Medolution Big Data platform.

4.3.1. Devices to be connected

There are several devices to be integrated into the Medolution platform to collect information about the patients and deduce meaningful results for the physicians through health data analytics.

Although Medolution prefers to use international and national standards for integration, most of the devices used in the healthcare sector are not conformant to the standards. It can be said that each manufacturer creates its own data models, serialization formats and/or protocols and this makes the integration impossible without their inclusion in the development process. For example, considering the LVAD use-cases, there are different LVAD manufacturers which are currently implanted into several patients. Most of them are even not providing any remote monitoring capability, though the one providing such a capability is not doing this through an open perspective and this makes the integration impossible.

Apart from the proprietary devices, Medolution promises to collect information from many other devices which conform to international protocols. This section covers the standards based connector components because the proprietary ones will be integrated through custom implementations for each device used for the respective use-cases.

The standards based interaction with the Medolution devices is going to be one-way; data will be collected from the devices into the Medolution platform. LVAD use-cases include devices (i.e. LVAD pumps) which require two-way communication; some parameters of the pumps are needed to be adjusted. These kinds of interactions will be solved through custom implementations during the project lifetime.

4.3.1.1. Bluetooth enabled medical devices

Personal health devices that use Bluetooth as a communication medium, based on ISO/IEEE 11073⁴ standards, will be used as data sources to Medolution Big Data platform. These devices can be blood pressure monitors, glucometers and the like Continua Health Alliance compliant medical devices covering the Medolution use-cases defined in *D1.2 – As-Is Landscape of Clinical Care at Pilot Sites and To-Be Pilot Application Scenarios*.

ISO/IEEE 11073

ISO/IEEE 11073 Health Informatics - Medical / health device communication standards are a family of ISO, IEEE, and CEN joint standards addressing the interoperability of medical devices. Continua Health Alliance(CHA)⁵ conformant products make use of the ISO/IEEE 11073 Personal Health Data (PHD) standards that specifically address the interoperability of personal health devices. The standard has the concept of agents and managers that defines communicating parties. Agents are typically small, battery-powered medical devices whereas managers are typically computers or smart phones. Agents are restricted to communicate with a single manager at a time yet managers can communicate with multiple agents. An example would be a smart phone as a manager, connected to multiple medical devices of a patient, e.g. weighing scale, thermometer, and insulin pump. Communication is bi-directional; weighing scale can send measurements to phone as well as phone can control insulin pump. The standard only defines messages that travel between agents and managers but not the transport protocol. Bluetooth, Zigbee or USB can be used to transfer messages.

4.3.1.2. Activity tracker wristband

We plan to use an activity tracking wristband as an additional data source to be able to collect more useful information from patients. Most of the commercial activity tracker wristbands has their own restricted APIs and does not allow direct data access with Bluetooth or other communication

⁴ <http://www.11073.org/>

⁵ <http://www.continuaalliance.org/>

protocols. As a result of investigating other open options, Angel Sensor⁶ stands out as it offers fully open access to its data with Bluetooth Low Energy communication.

Angel Sensor provides Heart Rate, Health Thermometer measurements as standard Bluetooth Low Energy (BLE) services as well as step count, blood oxygen saturation, fall detection, accelerometer, gyroscope measurements as custom BLE services.

4.3.1.3. Smart phone

In consequence of having Bluetooth enabled health and activity data producing devices to collect data from, it is planned to use a smart phone as a gateway between these data sources and the Medolution big data platform (BDCF). Both Android and iOS mobile devices can be used as a gateway since the requirements are only Bluetooth and an Internet connection. The devices are going to send their measurements to the smart phone and an application (possible a dedicated Medolution app to collect information and show some statistics to the user) is going to be responsible to pass this information conveniently to the Medolution BDCF servers.

Google Fit and Apple HealthKit

Using smart phones as a gateway enables us to gather more information from patients. Based on the gateway device in use, Google Fit⁷ or Apple HealthKit⁸ is planned to be utilized in order to collect more data from patients. Both Google Fit and Apple HealthKit is used as a health and activity data repository by other health and fitness applications, furthermore Google Fit collects and derives its own data using phone sensors and GPS location.

Google Fit provides various data types including heart rate, step count, nutrition and allows developers to create custom data types⁹. Apple HealthKit also provides similar fitness related data types along with health related types including blood glucose, oxygen saturation, body temperature and the like¹⁰.

Mobile sensing library

Smart phones have various sensors and components that contain valuable information about patients. To reach out this data, we are planning to use a mobile sensing library which will be configured to collect desired data from mobile phone components periodically and store them in a preferred format so that the mobile application can use this stored data to submit to the Medolution Big Data platform.

4.3.2. Devices particularities with respect to data connexion

The section above brings to light that from the functional point of view, the MEDOLUTION DEVICE CONNECTOR shall output a unitary data format, able to accommodate the large variety of heterogeneous data formats inner to the devices required by the MEDOLUTION use cases.

⁶ <http://www.angelsensor.com>

⁷ <https://www.google.com/fit/>

⁸ <https://developer.apple.com/healthkit/>

⁹ <https://developers.google.com/fit/android/data-types>

¹⁰ https://developer.apple.com/reference/healthkit/1627060-healthkit_constants

Thinking beyond the strict MEDOLUTION framework, such a unitary data format should also be interoperable.

Actually, from the practical point of view, two different situations are encountered in practice. The first situation corresponds to the case in which the device manufacturer exposes the data generated by the application as well as the proprietary data format as direct output of the device. In such a situation, the MEDOLUTION DEVICE CONNECTOR ensures a syntax management function. Cf. Figure 2. From a conceptual point of view, the MEDOLUTION DEVICE CONNECTOR includes data formats and API.




Figure 2: MEDOLUTION DEVICE CONNECTOR when the device proprietary data format is exposed.

The second situation corresponds to the case in which the device manufacturer does not expose to the user the data generated by the device; instead, it encapsulates them into a proprietary wrapper and subsequently forward them to a proprietary application which is supposed to process and to present them to the user. In such a situation, the MEDOLUTION DEVICE CONNECTOR becomes more complex, requiring an additional DEVICE PROXY whose functionality is to virtually de-capsulate the data, cf. Figure 3 .

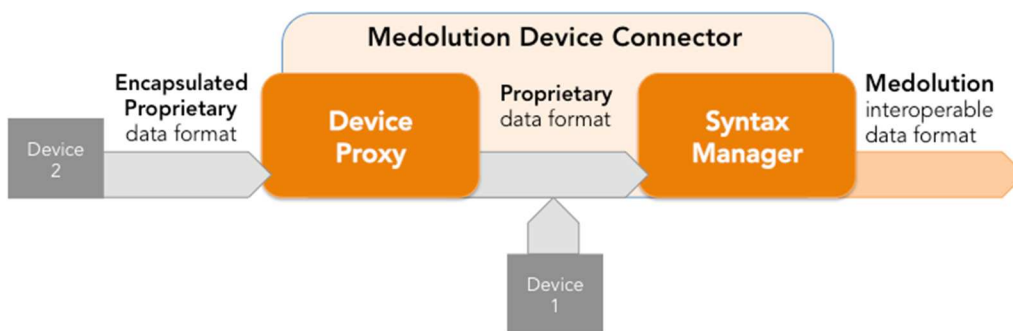


Figure 3: MEDOLUTION DEVICE CONNECTOR when the device proprietary data format is encapsulated (Device 2) vs. exposed (Device 1).

Of course, from the data format and interfaces point of view, the Device Proxy and the Syntax Manager should follow the same philosophy.

4.3.3. Medolution device connector when device data are exposed

4.3.3.1. Data to be collected

With the use of the devices and sensors, medical and personal data will be transferred to the big data processing environment of Medolution. Types of such data form an abstraction layer can be listed as follows:

- Personal health data from medical devices (e.g. pulse, blood pressure, blood glucose)
- Sensor data from implanted devices (e.g. LVAD)
- Fitness and health data from activity tracker wristband (e.g. heart rate, skin temperature, step count)
- Health and Fitness records from Google Fit or Apple Healthkit.
- Various information from mobile sensing library (e.g. GPS location, Wifi scan results)

Custom devices (which do not conform to the international standards, i.e. LVAD) serve data in proprietary formats, hence they will be processed through custom implementations during the project lifetime based on the specific data handling requirements.

For the ISO/IEEE 11073 based device integration and fitness/health sensor (wristlet) integration, smart phones will play the intermediary role to collect data from these personal devices and send data to the Medolution Big Data platform through the specific APIs. This deliverable focuses on the second part: transferring data to the platform through the APIs. There might be other sources of data which can directly transfer data to the platform APIs without any Bluetooth intermediary such as a smart phone.

The API structure is going to follow the REST principles and the data model of the payloads is designed to be based on Open mHealth.

Google Fit and Apple HealthKit

Open mHealth is a non-profit organization aiming integration of digital health data. They provide common data schemas for mobile health data types.

Data point schema is planned to be used as the common format, as illustrated in Figure 4.

```

{
  "body": {
    "heart_rate": {
      "value": 111,
      "effective_time_frame": {
        "date_time": "2016-03-15T07:30:11.099Z+02:00"
      },
      "unit": "beats/min"
    }
  },
  "header": {
    "id": "FBB8640D-8A2F-4584-9AB5-B857044E272C",
    "creation_date_time": "2016-03-15T13:09:08.802ZZ",
    "schema_id": {
      "version": "1.0",
      "namespace": "omh",
      "name": "heart-rate"
    },
    "patient_id": "A8BBACEE-CFFC-42A9-A147-92E90474720A"
  }
}

```

Figure 4: An example of Open mHealth Data Point Schema

Basically, “header” contains metadata about the resource whereas “body” contains actual measurement or value. Another “body” example is blood glucose schema as shown in Figure 5.

```

{
  "blood_glucose": {
    "unit": "mg/dL",
    "value": 120
  },
  "effective_time_frame": {
    "time_interval": {
      "start_date_time": "2013-02-05T07:25:00Z",
      "end_date_time": "2013-06-05T07:25:00Z"
    }
  },
  "descriptive_statistic": "minimum"
}

```

Figure 5: "Body" if a blood glucose measurement through Open mHealth

4.3.3.2. Device Connector Application Program Interface (API)

The associated API for device data collection is going to act as an "Upload Service" which will be a Restful API and used by the Bluetooth intermediaries (smart phones) to transfer the collected obtained data from the devices to the Medolution Big Data platform. The API will accept binary and JSON payloads for POST operation.

4.3.3.2.1. Binary payload

Binary payload endpoint will accept SQLite database files containing two tables: Data and FileInfo. The table structure is presented in the ER diagram in Figure 6. The tables and the field can be described as follows:

- Data
 - id: Holds database entry id.
 - name: Holds the source of the value.
 - timestamp: Holds the timestamp of the value.
 - value: Holds the Open mHealth Data Point schema formatted value of the resource.
- FileInfo:
 - id: Holds database entry id.
 - dbname: Holds database name.
 - device: Holds unique device id which will be used to identify patients.
 - installation: Holds unique installation id of the mobile application.
 - uuid: Holds unique database file id.
 - created: Holds the creation time of the database file.

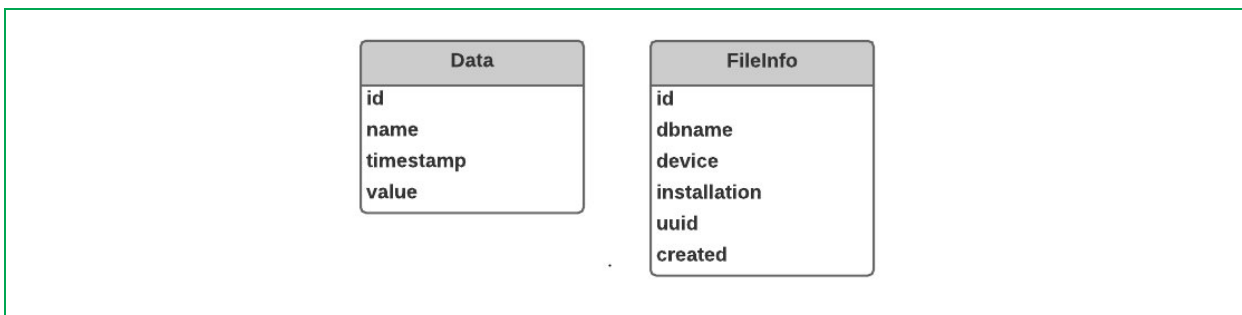


Figure 6: Simple ER diagram of the binary payload of the API

The HTTP Post command should be issued as follows:

- POST [base]/data

- The [base] indicates base URL for REST service e.g. www.medolution.org/service/rest/upload

The content will be provided using multipart form data format, with the parameter name “uploadedfile”. The service will parse the given database file, and forwards its entries to relevant Big Data Platform entry points.

Table 1: Binary payload of the API

Binary Payload	
Allowed HTTP methods	POST
Path including path parameters	data
Query parameters	Not applicable
POST data	uploadedfile: {Database file content in binary}
Success POST response	200 ok empty
Error responses	400 Bad Request - if file is not in desired structure. 415 Unsupported Media Type - if content-type is not multipart form data 500 Internal Server Error in case of internal exception

4.3.3.2.2. JSON payload

JSON payload endpoints will accept Open mHealth formatted JSON lists. HTTP POST command will be used as follows:

- POST [base]/[source]/[type]
 - The [base] indicates base URL for REST service e.g. www.medolution.org/service/rest/upload
 - The [source] indicates source of the data e.g. healthkit, ios.
 - The [type] indicates type of the data e.g. location, wifi.

The content will be provided with application/json format. The service will forward each element in the given JSON List to relevant Big Data Platform entry points.

Table 2: JSON payload of the API

JSON Payload	
Allowed HTTP methods	POST
Path including path parameters	/{source}/{type}
Query parameters	Not applicable
POST data	JSON list string

Success POST response	200 ok empty
Error responses	400 Bad Request - if string is not in desired structure. 404 Not Found - if {source} or {type} is invalid 500 Internal Server Error in case of internal exception

4.3.4. Medolution device connector when device data are not exposed

In such a situation (i.e. case 2 in Figure 3), the data generated by the device are encapsulated into a proprietary wrapper and subsequently forward them to a proprietary application which is supposed to process and to present them to the user

As explained above, in such a situation, a Device Proxy should be considered in conjunction with the Syntax Manager.

The DEVICE PROXY has three sub-modules: Virtual Access to the Application, the Virtual Control of the Application and the Semantic Data interpreter, cf. Figure 7.




Figure 7: The Device Proxy overview

The interfaces towards and from the device proxy will follow the same philosophy as the Semantic Manager: they are based on REST API and JSON formats.

At the time of writing of this document, the Device Proxy was instantiated for a medical card reader (the hardware device ensuring the authentication of the medical professional and patients in France).

4.3.5. Conclusion

The MEDOLUTION Device Connector can be illustrated as in Figure 8 below.

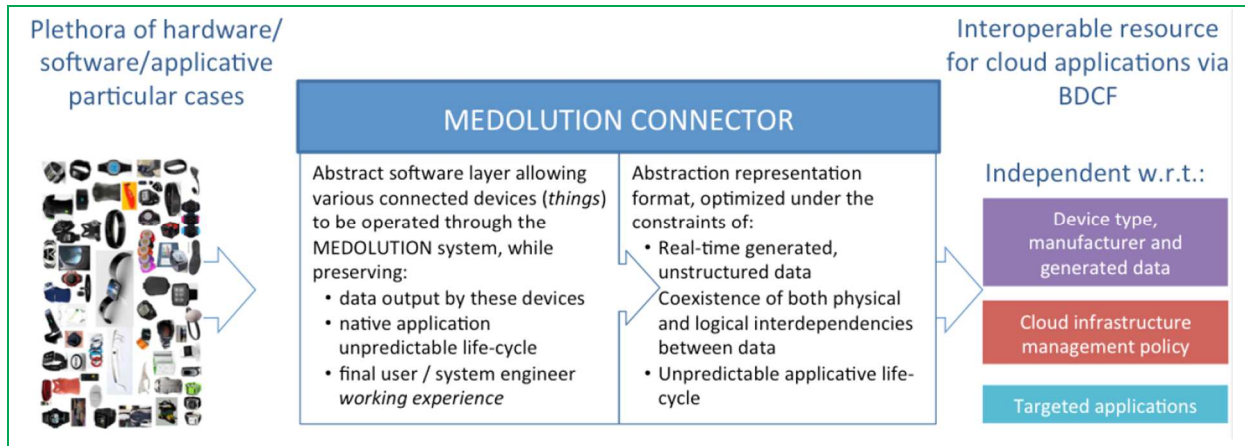


Figure 8: Global view on the device Connector component and on its possible solution.

4.4. Data connector components

Data to be analysed could also be extracted from data warehouses like legacy databases. In this case it will be useful to benefit from components which encapsulate the reusable data extraction code, and that will be easy to connect to Big Data storage or analytics processing components. At the time being, only one data connector component exists to extract data from relational databases, it is planned in the project to develop additional connectors to access data from medical databases (e.g. HL7).

4.4.1. Sqoop

Sqoop is a tool for efficiently transferring data from relational database management system (RDBMS) to Hadoop. Sqoop is designed as client-server. Sqoop implementation in BDCF is part of MapR distribution. Both client and server Alien4Cloud components are respectively named MapRSqoop2Client and MapRSqoop2Server.

4.4.1.1. MapRSqoop2Server

API

Server side is reachable via a REST API. Full documentation is available here: <https://sqoop.apache.org/docs/1.99.3/RESTAPI.html>.

TOSCA component description

Properties

MapRSqoop2Server does not have any properties.

Requirements

MapRSqoop2Server requires to be hosted on MaprWarden.

Capabilities

sqoop2_server_endpoint: allows connecting any Alien4Cloud component to this one. It can be used for MapRHUE requirement.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.4.1.2. MapRSqoop2Client

API

Client side provide a Java API which facilitates Sqoop integration in Java applications.

TOSCA component description

Properties

MapRSqoop2Client does not have any properties.

Requirements

MapRSqoop2Client requires to be hosted on MaprWarden.

Capabilities

sqoop2_client_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.5. Data analytics components

The following subsections describe the BDCF components which allow application developers to analyse the data which has been gathered and stored using the components from the previous sections.

4.5.1. Rstudio

RStudio Server is a web based integrated development environment (IDE) for R. It includes a console, syntax-highlighting editor that supports direct code execution, as well as tools for plotting, history, debugging, and also workspace management.

Rstudio API

Rstudio has been developed to provide an integrated development environment as Web UI to allow end user to easily design an R application.

TOSCA component description

Properties

proxy_to_use: Setup a proxy configuration in Renviron.site in order to allow downloading remote packages. If a value is set for this property then it will be used as http and https proxy (it should honour the unix http_proxy env var format). If not set or set to an empty string then the default environment proxy settings on the compute will be used (http_proxy, https_proxy and no_proxy). This is the default. If set to 'None' then proxies are not configured at all.

default : ""

cran_mirror_to_use: Mirror for R packages downloads to use. This allows setup and using your own mirror.

default : "http://cran.r-project.org"

user_to_create: Name of the unix account to be created and used to connect to RStudio

default : "rstudio"

password_to_create: Password of the unix account to be created and used to connect to RStudio

default : "rstudio"

repository: This property give the opportunity to specify an alternative download repository for this component artifacts. It is your responsibility to provide an accessible download url and to store required artifacts on it. You should specify only the base repository url. Artifacts names will be appended to it, so this property could be shared among several components using the inputs feature.

default : ""

Requirements

host: RStudio server component has to be hosted on a Compute.

Attributes

url: This attribute contains the URL used to access the RStudio UI. This attribute is only valid if the Compute on which RStudio is hosted is connected to a public network.

4.5.2. Spark

Spark is a fast and general-purpose cluster computing system which aims to analyse a large amount of data using different development languages like Java, Scala, Python and R. It provides some frameworks for machine learning, streaming and graph processing. Command line interface is the only way to execute Spark Jobs.

Spark API

BDCF offers three ways to launch Spark jobs:

- Standalone: jobs are started on compute Spark is installed. It is an easy way to test scripts at development step. The command is: `spark-submit [spark-options] <path/to/spark/job>`
- YARN client mode: jobs placements are managed by YARN resource manager (see section 4.7.2). Unlike YARN cluster mode, Spark driver is executed on the compute Spark is installed which gives an easy access to outputs. It uses all cluster resources, which allows launching more consumer jobs. The command is: `spark-submit --master yarn-client [spark-options] <path/to/job/script>`
- YARN cluster mode: all parts of Spark jobs are managed by YARN. Even Spark driver is launched on a YARN container. It is the recommended way to execute Spark jobs in production environment. The command is: `spark-submit --master yarn-cluster [spark-options] <path/to/job/script>`

TOSCA component description

Spark implementation in BDCF is part of MapR distribution. The component name in BDCF catalogue is MapRSpark.

Properties

MapRSpark does not have any properties.

Requirements

MapRSpark requires to be hosted on MapRWarden.

Capabilities

spark_endpoint: *allows connecting any Alien4Cloud component to this one.*

Artefacts

tools: *MapR utility scripts.*

utils_scripts: *Common utility scripts for whole BDCF components.*

4.5.3. Pig

Apache Pig is a high-level development framework which allows creating MapReduce jobs to analyse large amounts of data. It provides parallelization mechanisms to improve performances. Jobs are written using Pig Latin, a dedicated development language to facilitate high level comprehension.

Pig API

The only way to launch pig jobs is to use its command line interface: `pig <path/to/pig/job>`

TOSCA component description

Pig implementation in BDCF is part of MapR distribution. The component name in BDCF catalogue is MapRPig.

Properties

MapRPig does not have any properties.

Requirements

MapRPig requires to be hosted on MapRWarden.

Capabilities

pig_endpoint: allows connecting any Alien4Cloud component to this one. It can be used for MapRHUE requirement.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.5.4. Mahout

The aim of Apache Mahout project is to provide implementations of machine learning algorithms.

Mahout API

Mahout algorithms are designed to be implemented in Java. Please refer to the official documentation (<https://mahout.apache.org/>) for more information.

TOSCA component description

Mahout implementation in BDCF is part of MapR distribution. The component name in BDCF catalogue is MapRMahout.

Properties

MapRMahout does not have any properties.

Requirements

MapRMahout requires to be hosted on MapRWarden.

Capabilities

mahout_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.6. Visualization components

This section describes the components used to visualize the results of data analysis performed by components previously described.

4.6.1. Kibana

Kibana 4 is an analytics and visualization platform that uses Elasticsearch to give you a better understanding of your data. It is a part of the Elastic Stack, which includes Elasticsearch, Logstash, Kibana.

A complete documentation is available at <https://www.elastic.co/guide/en/kibana/current/index.html>.

4.6.1.1. Kibana Interface

The Kibana interface includes four main sections:

- Discover
- Visualize
- Dashboard
- Settings

These sections are reached by tabs on the top of the main interface.

- **Kibana Discover:** The Discover page displays all of the Elastic Stack most recently received logs. Here, you can filter through and find specific log messages based on search queries then narrow the search results to a specific time range with the time filter.
- **Kibana Visualize:** The Visualize page is where you will create, modify, and view your own custom visualizations.
- **Kibana Dashboard:** The Dashboard page is where you can create, modify, and view your own custom dashboards. A dashboard allows you to combine multiple visualizations onto a single page. Dashboards are useful to get an overview of your logs, and to make correlations among various visualizations and logs. Dashboards can be refreshed automatically with a configurable period. They can be filtered by entering a search query, changing the time filter, or by clicking on the elements on the visualization.
- **Kibana Settings:** The Settings page lets you change various parameters like default values or index patterns.

4.6.1.2. Kibana Configuration

Basic configuration

A Kibana node must be hosted on a Java node, which is hosted on a Compute node. Kibana Dashboard nodes can be attached to the Kibana node using dashboard_host prerequisite.

A Kibana node requires to be related to an Elasticsearch node. Use the search_endpoint prerequisite to establish the relation with an already created Elasticsearch node.

A Kibana node is related to a Consul agent hosted on its Compute node. This is required to perform Elasticsearch cluster discovery. Use consul prerequisite to connect Kibana to its Consul Agent.

Kibana Properties

- repository: This property allows you to specify an alternative download repository for the Kibana binary.
 - Default : ""
- es_heap_size: This property allows setting the heap memory that will be allocated to the java process of the Elasticsearch client node associated to Kibana. It allocates the same value to both initial and maximum values (ie -Xms and -Xmx java options).
 - Default : "1G"

Kibana Capabilities

The Kibana component can be used as a dashboard_host by Dashboard components. The role of a Dashboard component is to carry a dashboard configuration (it has an artifact named dashboard_file). This configuration is described in a .JSON file. Several Dashboard components can be connected to a Kibana component by using their dashboard_host prerequisite.

Kibana URL

Kibana URL is an address generated at the end of the topology deployment.

4.7. Data distribution & Resource management components

Performing Big Data analysis requires complex distributed computing configuration, managing data distribution and processing distribution. The components described in this section allow managing this distribution.

4.7.1. MapR Filesystem

MapR distribution has designed its own distributed filesystem dedicated to Big Data applications. It is an improved Hadoop filesystem with better performances, reliability, efficiency, maintainability and ease of use.

MapR filesystem management in BDCF is split into two different components: MapRCldb and MapRFileserver.

4.7.1.1. MapRCldb

The Container Location Database (CLDB) service is the equivalent of Hadoop Namenode. It tracks information about every container in MapR filesystem (location, size, volume, policies, quotas, usages, etc...). It also tracks file servers and node activities in the cluster.

Running the CLDB service on multiple nodes distributes lookup operations across those nodes for load balancing and also provides high availability.

MapR distribution is based on a licensing model. When a cluster is deployed, a MapR Base Edition license is registered which doesn't provide all functionalities (e.g. the NFS gateway can't be instantiated). The MapR Community Edition provides all functionalities with unlimited access for free. The MapR Enterprise Edition is dedicated to production environments, with high availability capabilities. For a full description of the MapR software editions, follow this link: <https://www.mapr.com/products/mapr-distribution-editions>.

TOSCA component description

Properties

replication: You can set the replication level of your cluster. Replication is the number of times your data are copied in your cluster to provide data safety.

Default value: 2.

Requirements

MapRCLDB requires to be hosted on MaprWarden.

MapRCLDB requires MapRFileServer to be installed on the same node.

Capabilities

cldb_endpoint: allows connecting any Alien4Cloud component to this one.

Custom command

`apply_license (license_url)`

license_url: The URL of the new MapR license file to apply.

Artefacts

licenseFile: The MapR license file to apply on the cluster at deployment step.

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.7.1.2. MapRFileserver

The MapRFileServer component is the access point to MapR filesystem. When an application want to access to a file into the filesystem, it first communicates with CLDB and then with MapR fileserver, which is the equivalent of Hadoop DataNode.

TOSCA component description

Properties

storage_pool_size: Allow to customize the number of disks in a storage pool (see MapR doc for more details)

Default: 3 (MapR default value)

Requirements

MapRFileServer requires to be hosted on MapRWarden.

volume: it requires to be connected to a block storage (30GB minimum) in order to format it as a MapR filesystem storage.

host: it requires to be hosted on a MapR Warden component.

Capabilities

fs_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

4.7.2. YARN

YARN is the main Big Data application resource manager. It embeds last MapReduce algorithms version (MR v2.0) and allows launching other types of applications. YARN is made up of two different components: YARN Resource manager and YARN Node manager.

YARN API

Both Resource and node managers are available through a Web UI and an RPC API. This allows managing and monitoring job executions.

4.7.2.1. YARN Resource manager

The MapRYarnRM component is the YARN Resource Manager of Hadoop. Using YARN, Big data applications run inside containers which are instantiated on YARN Node Manager. The Resource Manager is big data YARN application endpoint. It manages all resources available on the cluster and schedules deployment of applications on containers.

TOSCA component description

Properties

yarn.scheduler.minimum-allocation-vcores: Defines the minimum number of vcores allocated to a container.

default: 1

yarn.scheduler.maximum-allocation-vcores: Defines the maximum number of vcores allocated to a container.

default: 32

yarn.scheduler.minimum-allocation-mb: Defines the minimum memory allocation available for a container in Mb.

default: 1024

yarn.scheduler.maximum-allocation-mb: Defines the maximum memory allocation available for a container in Mb

default: 8192

Requirements

MapRYarnRM requires to be hosted on a MapRWarden.

MapRYarnRM component is required on MapR cluster to use ecosystem tools.

Capabilities

rm_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

yarn_files: Required YARN template files.

4.7.2.2. YARN Node manager

The MapRYarnNM component is the YARN Node Manager of Hadoop. It is managed by YARN Resource Manager and allows deploying containers which will execute big data applications.

TOSCA component description

Properties

yarn.nodemanager.resource.cpu-vcores: Defines the number of CPUs available to process YARN containers on this node.

Default: Variable. This value is calculated by Warden.

yarn.nodemanager.resource.memory-mb: Defines the memory available to process Yarn containers on the node in Mb.

Default: Variable. This value is calculated by Warden.

yarn.nodemanager.resource.io-spindles: Defines the number of disks available to process YARN containers.

Default: Variable. This value is calculated by Warden.

mapreduce.map.cpu.vcores: Defines the number of CPUs available to process map operation.

Default: 1

mapreduce.map.memory.mb: Defines the container size for map tasks in MB.

Default: 1024

mapreduce.map.java.opts: Java options for map tasks.

Default: -Xmx900m

mapreduce.map.disk: Defines the number of disks a map task requires.

Default: 0.5

mapreduce.reduce.cpu.vcores: Defines the number of CPUs available to process reduce operation.

Default: 1

mapreduce.reduce.memory.mb: Defines the container size for reduce tasks in Mb.

Default: 3072

mapreduce.reduce.java.opts: Java options for reduce tasks.

Default: -Xmx2560m

mapreduce.reduce.disk: Defines the number of disks that a reduce task requires.

Default: 1.33

Requirements

MapRYarnNM requires to be hosted on MapRWarden.

A MapRYarnRM must be set on the cluster.

Capabilities

nm_endpoint: allows connecting any Alien4Cloud component to this one.

Artefacts

tools: MapR utility scripts.

utils_scripts: Common utility scripts for whole BDCF components.

yarn_files: Required YARN template files.

5. Main Medolution Big Data Topologies

This section describes typical Medolution components assemblies to satisfy Medolution application requirements. Big Data Capability Framework (BDCF) provides many Big Data components assemblies that cover some general use cases and can be re-used and/or enhanced inside the Medolution project. The figure below illustrates an assembly, also called topology, for an ElasticSearch, Logstash, and Kibana assembly, as it can be viewed in the BDCF GUI.

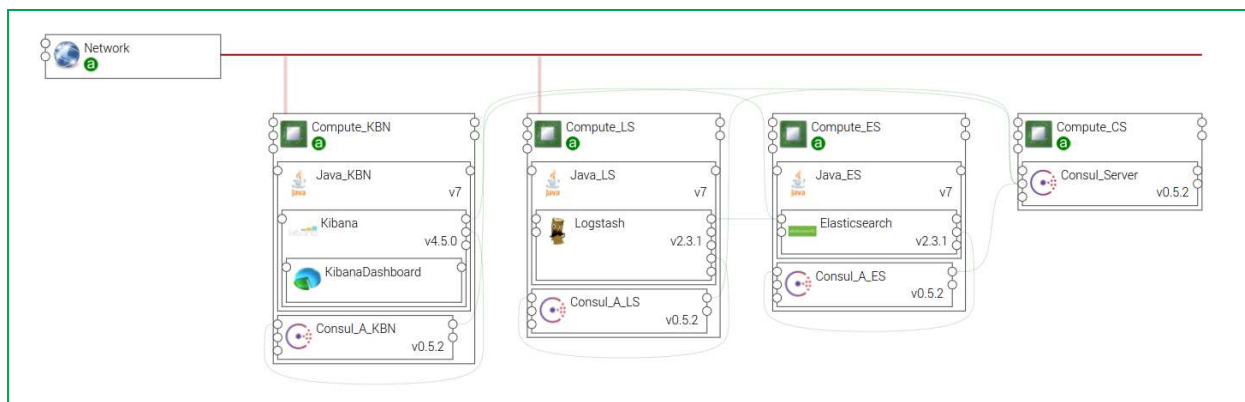


Figure 9: Smart Log Analytics assembly example

The complete list of available topologies is part of the BDCF user guide [2]; it contains topologies for Hadoop clusters configuration with MAPR or HortonWorks, topologies for MongoDB, etc... In the following subsections are described the Lambda Architecture topology and the LVAD Use Case initial topology, which typically address identified Medolution requirements at the time being.

5.1. Lambda Architecture

The Lambda Architecture is generic, low latency, scalable and fault-tolerant data processing architecture. It aims to analyze large set of data both in a batch mode and in a real-time mode. It is composed of the following layers:

- The ingestion part is where data is coming usually in a continuous flow, and whose goal is to transfer it to both batch and speed layers. Scalability and retention capabilities of this layer are very important to cope with large amount of incoming data and to ensure that no data will be lost in case of failure.
- The batch layer is processing a huge set of data and performs very precise calculation on the global data set. This layer ensures persistency of the global data set and is usually set up on a distributed infrastructure.
- The speed layer is processing data which is directly coming from sources and performs live analysis. The idea is to minimize the latency to provide real time views on most recent data. Results are stored in the serving layer allowing data scientist to visualize and query these results.
- The serving layer is usually the place where results of the batch analysis and indexed to provide a fast and easy access. In the example below, it is also used by the speed layer.

- Visualization layer is the place where results are visualized. BDCF provides a ready to use topology built from Alien4Cloud catalog components.

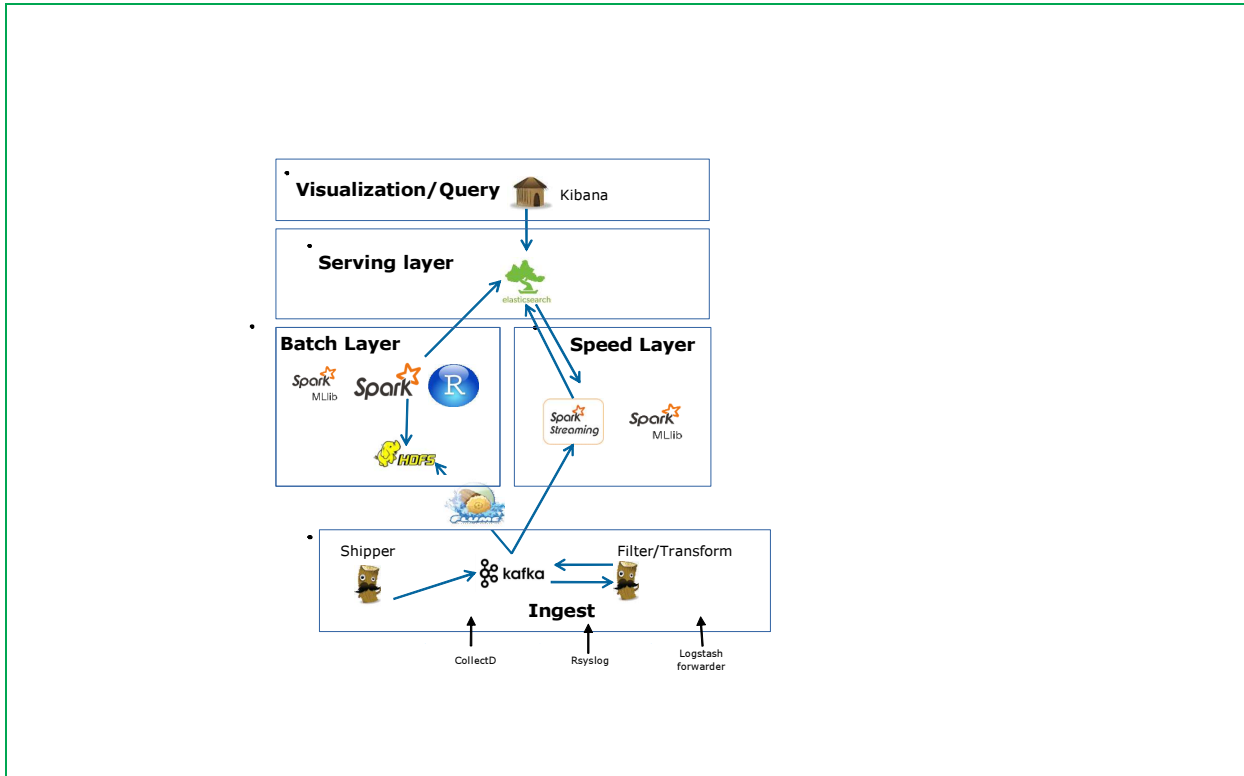


Figure 10: Lambda architecture

Figure 10 above illustrates this topology, the TOSCA topology diagram as visible in the BDCF GUI is shown below in Figure 11 (although not very readable this document, it can be zoomed in the GUI).

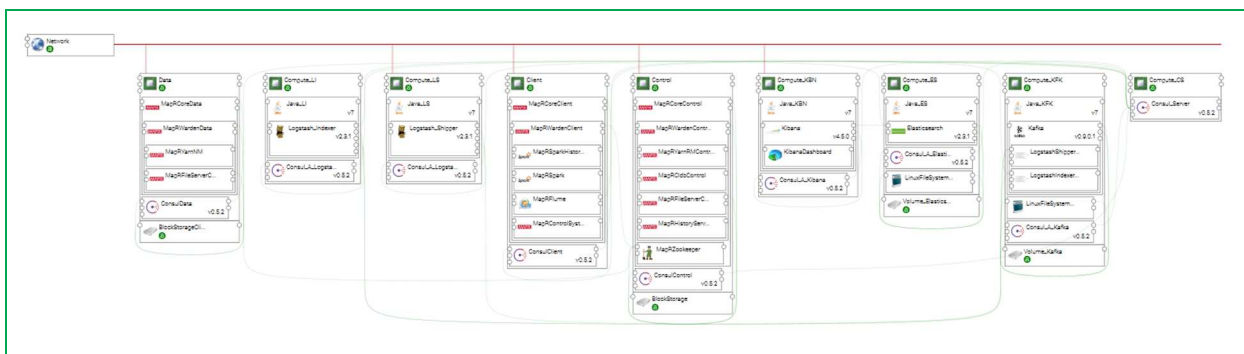


Figure 11: BDCF Lambda architecture representation

These components have been selected to match the requirements of each lambda architecture layer:

- Batch layer:** The master data management & batch computation would be done on a MapR cluster. This cluster includes MapR Spark as computation framework.
- Speed layer:** The streaming processing would be done on the same MapR cluster than batch layer.
- Serving layer:** Elasticsearch is used to index results from batch and speed layer. These indexed results can be queried by Kibana.

The data ingestion in the processing architecture is made using a model similar to ELK-broker. After the data is consumed by Kafka coming from LogstashShipper, it can follow 3 streams:

- The data could follow the classic ELK-broker flow: LogstashIndexer, ElasticSearch and Kibana.
- The data could be acquired by a Spark streaming program running on MapR cluster.
- The data could be stored on MapRFS using MapRFlume for batch processing.

5.2. Concrete architecture of the LVAD medical use case

The architecture below is a first draft that has been built to support the Medolution LVAD medical use case. This is typically a topology that will be supported in the Medolution platform, and corresponds to the kind of architecture that is straightforward to build with the BDCF integration PaaS. When the technical choices will be definitive, the corresponding components will be integrated in the Medolution platform catalogue (note that some of them are already there, like Kafka and R).

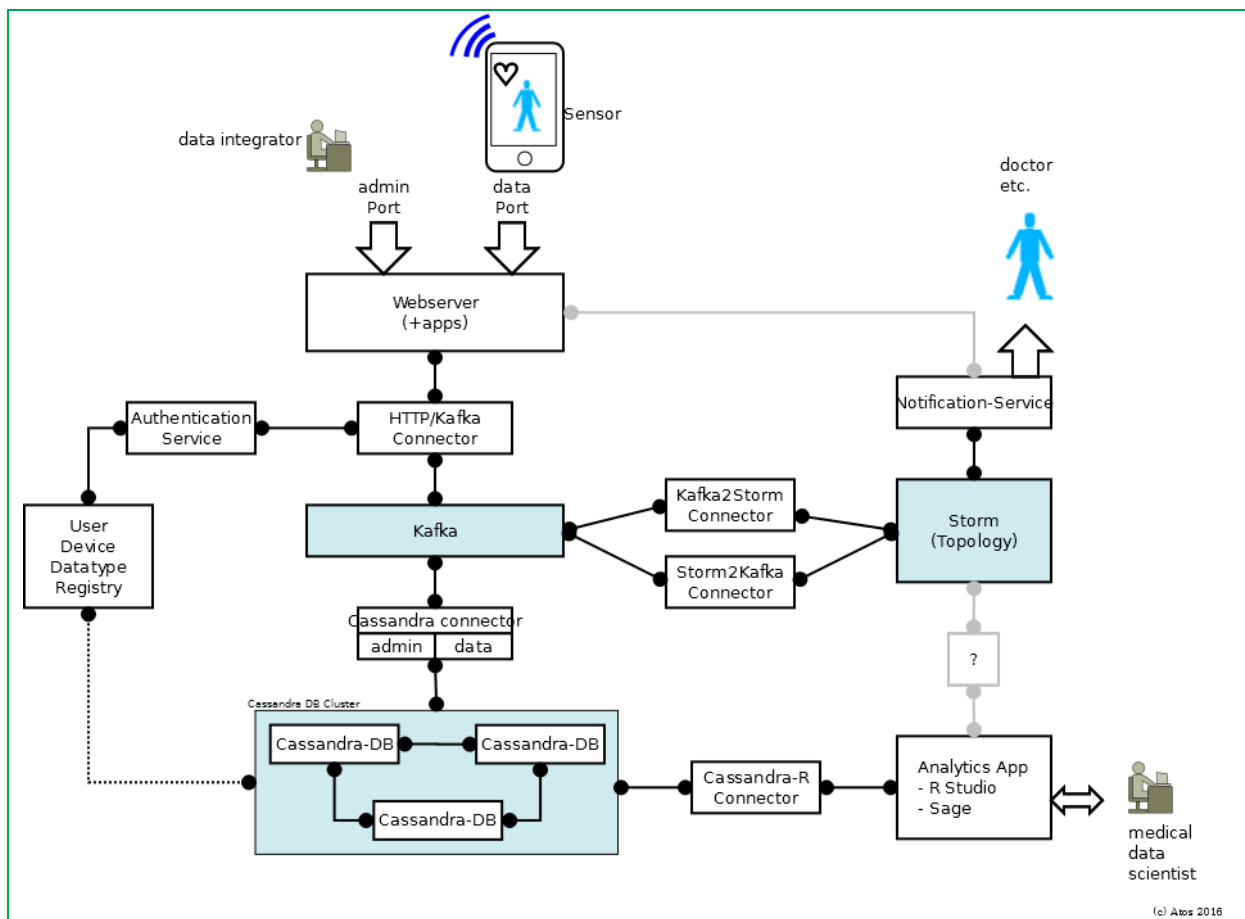


Figure 12: LVAD Medolution Cloud Scenario

In the LVAD use-case data from the patients LVAD devices and potential other sources is sent to a data-analytics and store component which will be realized as virtualized services. The use-case supports two kinds of processing: Stream processing of the actual received data packages and batch-oriented analysis of historical data. The concrete architecture is a variant of the lambda



architecture for data analytics but it has been extended to support the requirements of the Medolution platform. As depicted in Figure 12 the interface to the non-cloud components of Medolution (which is typically a local IoT platform or gateway) is a standard REST-based web interface where input data is received.

The processing in the cloud is designed as routing of messages to the respective processing elements. Therefore the data is validated and checked for the authorization. If it is identified as valid data from a valid source it is handed over to a Kafka component which serves as high-performance message queue. The architecture has been layered out so that arbitrary processing components and storage can be attached to Kafka by registering for target types of data. All data will be stored in a permanent database. In this scenario Cassandra is the database of choice, because it can be configured to create dependable and secured storage. For redundancy a minimum of three Cassandra instances is deployed to build a database cluster.

For near real-time analysis of data Storm is used to set up a processing network (topology). In this particular use-case there will be recognition nodes in Storms hub and spoke sub-architecture to detect patterns in the data-stream from LVAD devices which allow to detect or even to forecast critical conditions of the device. Such events will be sent to the notification service which will relay the message to a doctor or clinical personal.

In the more batch-oriented part standard tools like R studio, Sage or Jupyter can be used for processing. All three of them support a number of algorithms suitable for a data scientist to extract new findings from the set of received data.

The architecture is augmented by the required connector components which create the links between the elements. In these components also the semantic of the processing is captured because they contain the definitions which types of data are to be send where in the architecture. Additionally components are necessary to provide security and authorization mechanism for the connected devices. Therefore a registry of allowed communication endpoints (devices or gateways) and the associated data-types has to be maintained. With this a certificate or token based identification can be implemented.

The described architecture can be deployed in the cloud. In order to isolate the component and to be able to tune the performance and allow scalability the following deployment configuration is envisaged:

- One instance for Kafka and Kafka adapters
- For each Cassandra instance a separate instance (for dependability reasons the instances should be allocated on different hardware components in the cloud)
- One instance for the storm topology (This is the start configuration. If the Storm topology grows more instances are necessary to ensure the performance).
- One instance for the batch-oriented tools.

This basic configuration will be refined during the progressing of the project. It is useful to create a blueprint from this architecture so for every hospital an individual cloud environment can be created.

Moreover, with such a blueprint integrated as a BDCF topology, the scalability can also be taken in charge automatically (i.e. assigning a “scalable instance” to a component for its deployment, and then it is possible increase at runtime the number of instances running this component). It has also to be noted that with the Medolution platform, deployment target may be any of public cloud,



private cloud, on premises... and that most of the time for Medolution healthcare apps, private clouds, on premises (like the hosting platform of WP4) will be used.

6. Security

6.1. Security Principle and Architecture

This section addresses how security aspects are covered within an application composed of Medolution components. In this first version of the document, the global security architecture and principles is not yet finalized. In particular the device part will need to be addressed. However some aspects which will be required on this architecture can already be defined:

- **Authentication:** Medolution platform will have to define a protocol to authenticate its users to ensure unauthorized people cannot access to the platform;
- **Authorization:** to define which resources can be accessed by an authenticated user (it requires the authentication process to be setup);
- **Encryption:** to apply algorithms at least on data which is sent on the network, to ensure it cannot be decoded by unauthorized user;
- **Auditing:** log on the system all activities related to the three previous items;
- **Anonymization:** mechanisms to anonymize data before it is made available to Medolution users.

All these aspects will be discussed later on in the project and will be defined on the next version of this document. The progress on some of them is described in the following sections.

6.2. Data Anonymization

The Data Anonymization Engine is a Medolution Core Platform Service component that provides a common mechanism for Data Providers to safely publish data to Medolution with the knowledge that any access to this data will be governed by their Data Custody rules around re-identification risk and limited to the Smart Apps & Researchers who have been explicitly approved. Conceptually, it is illustrated as follows:

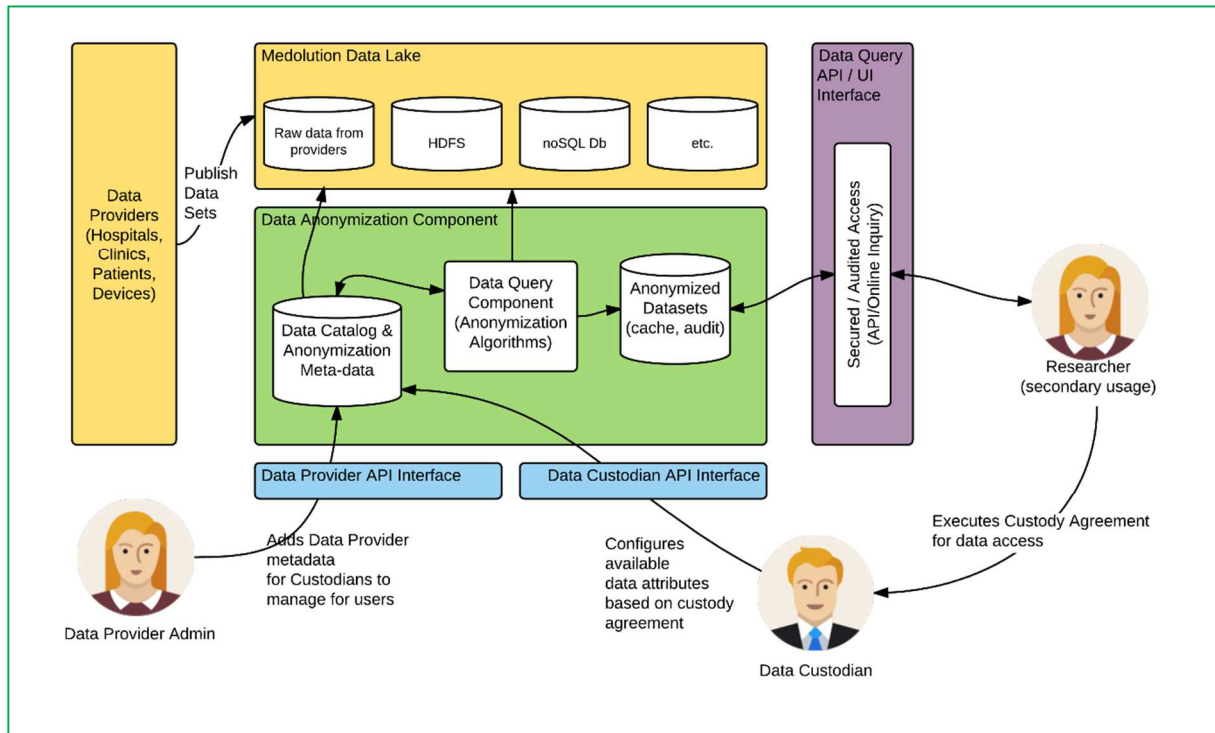


Figure 13: Data Anonymization

The Anonymization Engine will be divided into three sub-components and the API capabilities will be expressed via these sub-components:

- **Data Provider Interface** – mechanism for a Data Provider to describe the data attributes and its pseudo-identifiers which they will publish to the Medolution data lake. The actual data publishing will require discussion with partners because of the many BDCF technical options available.
- **Data Custodian Interface** – mechanism for a Data Provider to setup the custody terms around usage as well as configuration settings such as re-identification risk thresholds, suppression limits, etc.
- **Data Access/Query Interface** – mechanism for a Smart App or Researcher to query the Medolution data lake in accordance with the Data Custody terms.

The following sections describe these three sub-components in more detail.

6.2.1. Data Provider Interface

The Data Provider sub-component will provide an API mechanism to enable Providers to describe their data sets and the attributes which must be anonymized within them. Specifically, the API will accept a JSON document which describes the data attributes and the quasi-identifiers (also called pseudo-identifiers) within the data set so that the Anonymization Engine can apply the appropriate algorithms to reduce the re-identification risk. A quasi-identifier is an attribute which can be combined with other data elements to enable an antagonist to identify the subject of the data record (e.g. visit data, gender, age, diagnosis, etc). Each quasi-identifier by itself doesn't identify

the subject but when used in combination with other data, especially external data, can be used to re-identify a subject.

The JSON document will optionally include configuration attributes for re-identification risk thresholds as well as other constraints such as suppression maximums, etc.

6.2.2. Data Custodian Interface

The Data Custodian sub-component will provide an API mechanism for Data Providers to setup their Custody constraints to manage access to their data for different Smart Apps & Researchers. More specifically, this element of the API will allow Smart Apps & Researchers to:

- Discover the Data Providers and attributes which are available from the Engine and the terms which they must adhere to.
- Request access to the Data Providers and attributes.
- Acknowledge the terms of use and agree to data custody terms.

The API will also allow Data Custodians to:

- Publish the terms of their data custody agreement, including restrictions (possibly legal terms) about combining data with external data sets for the purposes of re-identifying subjects.
- Grant or Deny access to the Data attributes based on these requests.
- Record the acceptance of data custody terms.

6.2.3. Data Access Interface

The Data Access sub-component of the Anonymization Engine provides the main functionality of the Engine. It will provide an API mechanism for a Smart App or Researcher to submit a SQL-like query which includes the attributes which they want to retrieve. The attributes requested in the query will be looked up in the Data Provider catalogue to identify the quasi-identifiers and build up the anonymization profile of the query. The Data Access component will execute this query against the Medolution data lake and pass the results through the anonymization algorithm to ensure the pseudo-identifiers have been generalized or suppressed according to the anonymization profile. The query interface will also allow the specification of data loss precision and the maximum suppression level. These two constraints, when combined with re-identification risk (prescribed by the data provider above), guide the algorithms as described below.

Note: The precise mechanism of querying the Medolution Big Data repository is variable based on the technology which is selected in the BDCF and will be elaborated further during the project. All queries to the Data Access sub-component will be stored for auditing purposes.

6.2.4. Data Anonymization Concepts

This section explains the concepts used in the APIs described above. Re-identification risk, data loss precision and maximum suppression level are parameters of these APIs and are tightly related to the de-identification algorithms, which are built on generalization and suppression techniques.

6.2.4.1. Generalization

The de-identification algorithms perform varying levels of generalization on a particular quasi-identifier. The set of algorithms assigned to a quasi-identifier is an ordered set; where the order is based on the level of generalization performed by the algorithm.

When discussing generalization, we also need to understand the concept of an equivalency class. All records within a dataset belong to the same equivalency class if they have the same values for the set of quasi-identifiers. When quasi-identifiers are generalized, the number of equivalency classes will decrease and the size of the equivalency classes will increase as the level of generalization increases. The converse is also true, as the level of generalization decreases, then number of equivalency classes will increase and the size of the equivalency classes will decrease. Therefore, as the data becomes more generalized, the probability of re-identification is lessened due to larger equivalency class sizes (at the expense of data loss).

Assume the following:

- Visit date is a quasi-identifier
- For the 2015 calendar year there were 118 visits on each date except for the following:
 - 20150724 – 3 visits
 - 20150725 – 18 visits
 - 20150726 – 333 visits
- This gives us a total of $(362 * 118) + 3 + 18 + 333 = 43\ 070$ visits in total
- The set of de-identification algorithms for visit date are
 - *a0* – output date in *yyyymmdd* format
 - *a1* – output date in *yyyymm* format
 - *a2* – output date in *yyyy* format
- The set of algorithms $\{a0, a1, a2\}$ are ordered based on the level of generalization they perform (from most specific to least specific).

Applying the algorithms to de-identify a dataset we can see that the number of equivalency classes generated by *a0* will be greater than the number of equivalency classes generate by applying *a1*. Similarly, the number of equivalency classes generated by *a1* will be more than the number generated by *a2*.

6.2.4.2. Re-Identification Risk

The stakeholder will also determine an acceptable level of risk that an individual could be re-identified from the de-identified dataset. For the purpose of this document we will discuss the maximum risk of re-identification. A commonly used de-identification criterion is *k*-anonymity; which stipulates that each record in a dataset is similar to a least *k*-1 records. Therefore, we can see that maximum risk is equivalent to $1/k$. When testing for maximum risk we need to look at the equivalency classes whose size is less than *k* as these represent the classes with the highest probability of re-identification.

Building on the example in the previous section, assume the following:

- The stake holder has specified the maximum risk of re-identification threshold is .01 (k=100)

If we were to apply de-identification algorithm *a0* (yyyymmdd), the size of the equivalency class for 20150724 is 3 and 20150725 is 18. This results in 1/3 and 1/18 chance of re-identification respectively, which are both higher than our risk threshold of .01; therefore, we could not release the data using *a0*. However, if we were to apply *a1* (yyyymm) we would end up with 12 equivalency classes, and based on our stated distribution the smallest equivalency class would be for 201502. The size of 201502 would be 28 days * 118 visits / day = 3304. The 1/3304 chance of re-identification is well within our stated threshold of .01. Based on this we could safely release the data using *a1* as our algorithm.

6.2.4.3. Suppression

The third aspect to the approach is the concept of suppression. This deals with what needs to be done to de-identified data to ensure the risk threshold is met. If we have an equivalency class with a re-identification probability greater than our risk threshold we will need to perform suppression on the data so that our threshold can be met.

Suppression, like generalization, introduces data loss. When we are evaluating de-identification algorithms to find the optimal set, the maximum amount of suppression that is acceptable must be specified. For example, if a set of de-identification algorithms require a suppression 80 percent of our records to achieve a .01 maximum risk to identification, it would clearly not be acceptable.

Introducing the concept of maximum amount of suppression to our running example, assume the following:

- The stakeholder has specified the maximum amount of suppression is .10 (10 percent).

If we were to apply algorithm *a0* (yyyymmdd) we would have to suppress the records for the 20150724 (3 records) and the 20150725 (18 records) equivalency classes as they exceed the maximum risk of re-identification threshold of .01. This gives us a suppression ratio of $(3+18) / 43070 = .00048$. This is well within the specified maximum amount of suppression (.10). Therefore, *a0* would indeed be an acceptable de-identification algorithm as it results in an acceptable amount of suppression to achieve our risk threshold.

Without the concept of suppression, we would have rejected *a0* and we would have selected *a1*.

Now that we have identified the records to suppress, we need to specify the techniques to use when performing the suppression. We could do some of the following:

- Case-wise deletion.
Remove all 21 records from the release data.
- Quasi-identifier omission.
Omit the quasi-identifier (visit date) for the 21 records, but release all other data associated with the visits. (Ignore the fact that we can infer the visit date for the 21 records as they are the only 2 dates not represented in our released file; this is a contrived example to add context to our abstract discussion).



- Quasi-identifier generalization.
Generalize the quasi-identifier for the 21 records. This would result in the visit date being generalized to yyyyymm or 201507.
- Local cell suppression.
We could suppress some, but not all, of the quasi-identifiers. When we remove a quasi-identifier the resultant equivalency class is less detailed than the original and therefore would be a larger set (if the entire data set was taken into account). We would keep omitting quasi-identifiers until the size of resultant equivalency class satisfies our maximum risk threshold. If had to remove all of the quasi-identifiers to satisfy our maximum risk threshold we effectively have “Quasi-identifier omission”

As we can see, not only do we need to know the risk tolerance when determining which de-identification algorithms should be used, we also need to know the maximum amount suppression that will be allowed.

6.2.4.4. Quantifying Data Loss

Generalization performed by applying a de-identification algorithm will cause data loss due to the amount of precision lost during the operation. A de-identification algorithm has a precision loss factor that indicates the degree of generalization performed. For example, generalizing a date to yyyyymmddd will result in 365 discrete values for a given year, generalizing a date to yyyyymm results in 12 discrete values for a given year, and generalizing to yyyy will result in a single value for a given year. As we can see the different algorithms have different precision loss factors. Also note that the precision lost between different algorithms is not constant, the precision lost between yyyyymmddd and yyyyymm is different than the precision lost between yyyyymm and yyyy.

Suppression also introduces data loss. To aid in quantifying data loss encountered during suppression we need to know the clinical significance of the quasi-identifiers.

During suppression, when we are performing local cell suppression, we would need to know which quasi-identifier is the least clinical significant. This would allow us to omit the quasi-identifiers in clinical significance order, omitted the least significant first.

The clinical significance of a quasi-identifier, along with the precision loss factor of the de-identification algorithms, will allow us to quantify data loss created by generalization and suppression for a given set of de-identification algorithms. This gives us the capability to compare the amount of data loss generated by different sets of de-identification algorithms so that we can choose the one with the least of amount of data loss.

As precision loss increases, the clinical significance will decrease. For example, a visit date in yyyyymmddd format has more clinical significance than a visit date in yyyyymm format. We have gone from 365 discrete values to 12, a factor of 30.4. However, this does not imply that the clinical significance has decreased by a factor of 30.4.

The magnitude of the change in clinical significance is independent of the precision loss factor.

6.3. Managing Access Rights

The global management of access rights in Medolution will be discussed later on in the project. However, some elements regarding the Big Data platform are already available.

Big Data distributions like MapR, Hortonworks or Cloudera come with their own security process or give a way to connect to dedicated components, which helps end-user to securely interact with the entire ecosystem. The 2.0 version of Big Data Capability Framework integrates such security features on MapR distribution only, allowing enabling the MapR wire-level secure mode at application deployment. This is documented inside the BDCF user guide [2].

MapR core system proposes an access rights management based on token principle. A user authenticates giving its credentials (user/password, Kerberos ticket...) and then the system returns a ticket which gives him rights to access to MapR services. Authorization is one part of MapR wire-level secure mode, as encryption is another.

6.4. Secured Data Transmission

Global requirements and specification about secured data transmission in the project have not yet been discussed at this phase of the project and will be handled later.

As part of BDCF, such a feature is also already available on the Big Data platform, through the MapR wire-level secure mode, together with the access rights management.

When this mode is enabled, all MapR core services encrypt their communications, MapR file system in particular. That will prevent potential attacks on decoding for example healthcare related data.

7. Medolution Platform APIs

Medolution Platform API is a RESTful interface that gives the ability to plug with existing enterprise systems (Development and Operations) and other stakeholders. This layer implements the OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) standard. In fact, it makes possible to build and manage a full application lifecycle on hybrid infrastructure in the same way as using the GUI. According to the Medolution project, Medolution Platform API could be used for the management of some preconfigured templates customized for Medolution use cases. It also allows the management and configuration of the interface between the deployed applications and other components such as IoT platforms, IoT gateways, portal platforms.

Using this RESTful interface, management features can be summarised as follows:

- Listing applications catalogue: allow Medolution user to list existing templates for on demand usage.
- Building a new application: Medolution API implements the TOSCA standard, which is used to build and orchestrate PaaS applications in the cloud.
- Deployment of applications: creation of the application that will be hosted on specific cloud platforms.
- Versioning and application lifecycle management
- Configuration of application APIs and access management.

Other features could be used, but may not be used in the scope of Medolution use cases, such as the Identity and Access Management.

8. Hosting Platforms

At this stage of the project, a first delivery of the BDCF framework has been made available on the hosting platform delivered by Bull in task 4.4 of the work package.

This hosting platform is described in the document describing the merged D4.2/D4.3 work package deliverable [3], it is composed of

- HW delivering high memory, compute and storage capabilities: a cluster architecture hosted in Bull French facilities providing a small datacentre for Cloud/Big Data, including a mix of standard x86 servers, large In-Memory Servers (Bullion 4/8TB up to 16 sockets), as well as different types of storage systems.
- An OpenStack IaaS, Liberty version, to virtualize this HW
- The BDCF framework to provide a PaaS access to this platform, i.e. creating VMs, deploying applications built from software components of the BDCF catalogue.

Within this platform, partners can bring their data and then install, test and validate their developments throughout the project whether they are technology providers or applications developers. The platform is located in a DMZ zone so as to allow all projects partners to have access to the platform resources and to their own private dedicated network zone and space.

Each partner may ask an access to this platform, the procedure is described in the Medolution SharePoint wiki, he will receive VPN and login information to access the BDCF portal (Alien4Cloud).

At this date, IMT, UPEC, Maidis, Prologue, Norima and Sopheon have received access to the platform.

That way, any partner has the ability to test the platform by designing and deploying applications built from components already available in the catalogue. Of course partners have also the ability to integrate new components in the catalogue, a document "Components Developer Guide" has been provided for this purpose. Currently version 2.0 of the BDCF framework is available on the platform. "BDCF User Guide" [2] and "Component Developer Guide" [1] are available on the Medolution SharePoint.

9. List of Figures

Figure 1: BDCF Platform	8
Figure 2: MEDOLUTION DEVICE CONNECTOR when the device proprietary data format is exposed.	22
Figure 3: MEDOLUTION DEVICE CONNECTOR when the device proprietary data format is encapsulated (Device 2) vs. exposed (Device 1).	22
Figure 4: An example of Open mHealth Data Point Schema	23
Figure 5: "Body" if a blood glucose measurement through Open mHealth	24
Figure 6: Simple ER diagram of the binary payload of the API	24
Figure 7: The Device Proxy overview	26
Figure 8: Global view on the device Connector component and on its possible solution.	27
Figure 9: Smart Log Analytics assembly example	39
Figure 10: Lambda architecture	40
<i>Figure 11: BDCF Lambda architecture representation</i>	40
Figure 12: LVAD Medolution Cloud Scenario	41
Figure 13: Data Anonymization	45



10. List of Tables

Table 1: Binary payload of the API	25
Table 2: JSON payload of the API.....	25



11. References

- [1] Atos Bull, *BDCF Component Developer Guide*, 2016.
- [2] Atos Bull, *BDCF 2.0 User Guide*, 2016.
- [3] Atos Bull, *Medolution D4.2 / D4.3, Medolution Core Platform and Hosting Platform*, 2016.