



# PROTEUS

Scalable online machine learning for predictive analytics and real-time  
interactive visualization

687691

---

## D5.4 Software implementation and integration with Apache Flink

---

**Lead Author: TREE**

**With contributions from: DFKI, BU**

**Reviewer: DFKI**

Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (PU)
Contractual delivery date:	M30 (30th May 2018)
Actual delivery date:	M30 (30th May 2018)
Version:	1.1
Total number of pages:	19
Keywords:	Big Data, data visualization, architecture

---

### *Abstract*

Advanced visualization of data analytics in real-time, user experience and usability are still open issues in the context of Big Data that also concerns the Arcelor Mittal dataset and use cases in PROTEUS. To meet the requirements of data analysis tasks for real-time data, visualization platforms need to be able to draw high volumes of data in real time while providing user interaction with a very low latency. This deliverable describes the demonstrator implemented and deployed during the project. As it was explained in deliverable D5.3, a 3-layer architecture has been designed for supporting incremental visual methods. It enables the analysts to explore both data-at-rest and data-in-motion with interaction times lower than onesecond.

This deliverable explains the implementation of the 3 layers of the proposed technical solution. Data Collector, and Incremental Analytics Engine layers will be implemented within the core of Apache Flink technology. The Visualization layer will be implemented as client-side library.

---

## Executive summary

Visualization of Big Data is still an open research question, especially when dealing with massive, unbounded streaming data at high frequencies that needs to be transformed, processed and merged with big volumes of batch data such as in the case of PROTEUS. Current visualization architectures add excessive complexity and costs to the development and maintenance of the system being developed, or they do not fully comply with the low-latency requirements presented in PROTEUS. High latency in a streaming visualization platform may cause delays that impede fluent interactions, hinder usability of the tools and make impossible to perform effective data analysis to make decisions on time.

In D5.3, we analysed the existing Big Data architectures that are commonly used for processing both streaming and batch data. Then, we introduced a 3-layer architecture that has been designed for PROTEUS.. The architecture, designed and implemented in WP5 solves the drawbacks of the Lambda and Kappa architectures, by focusing on the requirements of real time interactive visualization. The 3-layer architecture in PROTEUS is composed of three main components:

- The *Data Collector* is in charge of iteratively getting new data from data sources (both static and streaming) and serving them properly and efficiently to the incremental analytics engine and the visualization layer. It also collects and stores processed data and machine learning models, acting as a shared data layer for all the components of the architecture.
- The *Incremental Analytics Engine* processes data using classical analytics algorithms and also online machine learning algorithms. The gaps identified in T5.1 have been covered in this engine.. It outputs partial results which will be visualized by the third layer.
- The *Visualization Layer* enables views of the intermediate results at various iterations and allows the users to track and interact with those results in real time. It can combine the intermediate results with raw data, giving the user a *bigger picture perspective* for better understanding of the data.

To develop such architecture, we have developed the following contributions to the visual analytics ecosystem as part of WP5:

- Proteic.js: a general purpose data visualization library for the web, built for supporting high volumes of both streaming and static data.
- PROTEUS Visual Analytics Platform: a web application that provides a dashboard or user interface for performing incremental visual analytics on the PROTEUS data.

This task shows the demonstrator designed in task D5.4. This solution provides a realistic and ready-to-use software solution to achieve a remarkable industrial impact.

During demonstrator implementation we have had to keep in mind issues related to security, resiliency, expected response time, and expected impact on the rest of the system. Some conclusions about this have been showed in deliverable D.2.10 and will be updated in the future deliverables.

## Document Information

<b>IST Project Number</b>	687691	<b>Acronym</b>	PROTEUS
<b>Full Title</b>	Scalable online machine learning for predictive analytics and real-time interactive visualization		
<b>Project URL</b>	<a href="http://www.proteus-bigdata.com/">http://www.proteus-bigdata.com/</a>		
<b>EU Project Officer</b>	Martina EYDNER		

<b>Deliverable</b>	<b>Number</b>	D5.4	<b>Title</b>	Software implementation and integration with Apache Flink
<b>Work Package</b>	<b>Number</b>	WP5	<b>Title</b>	Real time interactive visualization

<b>Date of Delivery</b>	<b>Contractual</b>	M30	<b>Actual</b>	M30
<b>Status</b>	version 1.1		final <b>X</b>	
<b>Nature</b>	report <input type="checkbox"/> demonstrator <b>X</b> other <input type="checkbox"/>			
<b>Dissemination level</b>	public <b>X</b> restricted <input type="checkbox"/>			

<b>Authors (Partner)</b>	Javier de Matías Bejarano (TREE)			
<b>Responsible Author</b>	<b>Name</b>	Javier de Matías Bejarano	<b>E-mail</b>	javier.dematias@treeologic.com
	<b>Partner</b>	TREE	<b>Phone</b>	

<b>Abstract (for dissemination)</b>	<p>Advanced visualization of data analytics in real-time, user experience and usability are still open issues in the context of Big Data that also concerns the Arcelor Mittal dataset and use cases in PROTEUS. To meet the requirements of data analysis tasks for real-time data, visualization platforms need to be able to draw high volumes of data in real time while providing user interaction with a very low latency. This deliverable describes the demonstrator implemented and deployed during the project. As it was explained in deliverable D5.3 a 3-layer architecture has been designed for supporting incremental visual methods. It allows the analysts to explore both data-at-rest and data-in-motion with interaction times lower than 1 second.</p> <p>This deliverable explains the implementation of the 3 layers of the proposed technical solution. Data Collector, and Incremental Analytics Engine layers will be implemented within the core of Apache Flink technology. The Visualization layer will be implemented as client-side library.</p>
<b>Keywords</b>	Big Data, data visualization, architecture

<b>Version Log</b>			
<b>Issue Date</b>	<b>Rev. No.</b>	<b>Author</b>	<b>Change</b>
2018-04-18	1.0	Javier de Matías Bejarano	First draft
2018-05-24	1.1	Javier de Matías Bejarano	Changes proposed by Alireza (DFKI)

## Table of Contents

Executive summary.....	3
Document Information.....	4
Table of Contents.....	5
List of figures and/or list of tables.....	6
Abbreviations.....	7
1 The 3-layer architecture.....	8
1.1 Data collector.....	9
1.1.1 Message transport system.....	9
1.1.2 NoSQL database.....	9
1.2 Incremental analytics engine.....	10
1.3 Visualization layer.....	12
1.3.1 Visual Analytics Platform: back end.....	13
1.3.2 Visual Analytics Platform: PROTEUS Dashboard.....	14
1.3.3 Datasources API.....	14
1.3.4 Proteic.js.....	14
2 Demonstrator source code.....	16
3 Demonstrator deployment.....	17
4 Conclusions.....	18
References.....	19

---

## List of figures and/or list of tables

Figure 1: The 3-layers incremental visualisation architecture in PROTEUS.....	8
Figure 2: Architecture of the Incremental Hybrid Engine in PROTEUS.....	11
Figure 3: User interface of the Visual Analytics Platform.....	13

## Abbreviations

**API:** Application Programming Interface

**HDFS:** Hadoop Distributed File System

**HTTP:** HyperText Transfer Protocol

**JSON:** JavaScript Object Notation

**PEACH:** PROTEUS Elastic Cache

**POJO:** Plain Old Java Object

**SOLMA:** Scalable Online Machine Learning Algorithm

**SQL:** Structured Query Language

**SPA:** Single Page Application

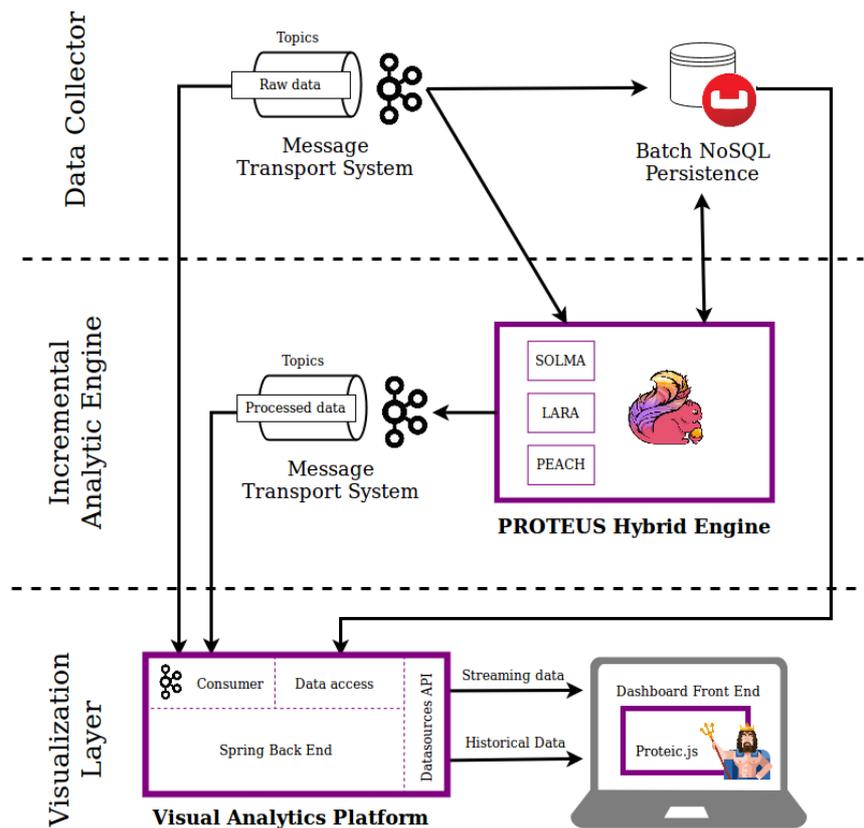
**STOMP:** Simple Text Orientated Messaging Protocol

# 1 The 3-layer architecture

To confront the challenges of developing low-latency incremental visualizations for the PROTEUS industrial use case we have implemented a visualization architecture. It is mainly inspired in the Lambda architecture and it is designed to support incremental visual methods. This incremental visual approach consists of three main layers: *Data Collector*, *Incremental Analytics Engine* and *Visualization Layer*.

- The *Data Collector* is in charge of iteratively getting new data from data sources (both static and streaming) and serving them properly and efficiently to the incremental analytics engine and the visualization layer. It also collects and store processed data and machine learning models, acting as a shared data layer for all the components of the architecture.
- The *Incremental Analytics Engine* processes data using classical analytics algorithms and also online machine learning algorithms. The gaps indentified in T5.1 have been covered in this engine.. It outputs partial results which will be visualized by the third layer.
- The *Visualization Layer* enables views of the intermediate results at various iterations and allows the users to track and interact with those results in real time. It can combine the intermediate results with raw data to enhance insight into the data.

The following figure 1 identifies how these layers have been implemented in the PROTEUS implementation. Components coloured in purple have been developed within the project:



**Figure 1: The 3-layers incremental visualisation architecture in PROTEUS**

As shown in the previous figure 1, the Data Collector Layer collects and stores raw data entering the system and also processed data so it can be shared between the other two layers. This shared data layer allows the PROTEUS Hybrid Engine to perform operations with both streaming and batch data, while connection with the Visualization Layer enables visualizations with both kinds of data and mixing raw and processed data.

This architecture has been designed to support the requirements of real-time interactive visualization of incremental results in PROTEUS with a very low latency. This has been achieved following the guidelines

specified in the deliverable D5.2, which describes the technologies, techniques, data formats and protocols used to develop the visualization layer.

The following sections provide further details about each of the elements in the architecture and its implementation to support real-time interactive visualisations in PROTEUS. The individual components used and developed for implementation of the visualization layer has been described more thoroughly, since they play the most prominent role to meet the visualization requirements. Information about the deployment of the architecture in a production environment and how they are integrated for the prototypes can be found in the deliverable D2.9.

## 1.1 Data collector

This layer mainly deals with the challenges associated with the velocity and volume properties of the data. It involves the connection to sources of raw data, so it can be collected and distributed to other components of the architecture. This is achieved by a message transport technology in charge of writing data to distributed topics in the cluster where other components will read and treat them accordingly. The data collector is also in charge of gathering data emitted from the Visual Analytics Engine, such as intermediate results and machine learning models, so it can be aggregated with raw data in the visualization layer. This enables better insight for operators in the industrial setting, who may find useful correlations in the manufacturing process as they can visualize raw data from the coils being produced with the analytics results about the whole dataset of past coils.

All data collected in the PROTEUS is also integrated by this layer into a master dataset stored in a NoSQL database to enable querying from the visualization layer. This database also stores the batch datasets used to generate the machine learning models at the incremental analytics engine.

As the entry point of all data in the system, technologies chosen for this layer must be as fast and efficient as possible to avoid bottlenecks that may halt the rest of the pipeline or cause data loss. In this case, we have implemented this layer using the technologies described in the following sections.

### 1.1.1 Message transport system

The message queue has been implemented using Apache Kafka but any other message queue system can perform this role, e.g., Apache Flume or RabbitMQ. We chose Kafka over others because it is an open source technology that is becoming an industry standard for message processing [2]. Also, Kafka is a general-purpose solution, while Flume is more oriented towards the Hadoop ecosystem and works with a push model in which data sinks must keep up the pace of events, causing bottlenecks at event peaks [3]. RabbitMQ is also a proven and well-known technology, but it does not reach the performance of Kafka [4].

The streams of data received at the message queue are written to the following distributed topics according to the coil production data flow (for further details see the deliverable D2.8, section 1.2, pp. 11):

- **Real-time:** manages 1-dimensional and 2-dimensional time-series data produced and available in real time during the coil production.
- **HSM:** manages all HSM data, produced as aggregated information at the end of each coil production and available only once the coil production has finalised.
- **Flatness:** manages the flatness data variables, produced as measures of the flatness of the resulting coil, and available only after a certain delay after the coil production has finalised.

The Incremental Analytics Engine also writes the incremental result of each algorithm to different topics in the queue system. All the topics will later be read by two consumers developed purposefully for the kind of data handled in the project. The first one treats the unbounded stream of data for the incremental analytics engine, while the second one writes the stream to a distributed NoSQL database where it is merged with all the historical data.

### 1.1.2 NoSQL database

PROTEUS provides innovative solutions based on incremental visual methods that enable end-users to explore data in real time efficiently to make well-informed decisions. This is based on a hybrid approach, i.e. not only for data-in-motion but also for both data-at-rest. For the historical or batch data, we have chosen

Couchbase Server as the technology for the storage of historical records of raw data and results from the incremental analytics engine in a reliable way. This database also stores batch data from the HSM dataset. The Data Collector is in charge of iteratively handling new data from data sources (both static and streaming), managing updates and queries towards serving them properly and efficiently for incremental analytics operations.

## 1.2 Incremental analytics engine

The next step in the data pipeline is to process the data so it can be cleaned and modelled based on the visualization layer requirements. It is also in charge of performing analytic operations that will be visualized by users, providing them with a deeper understanding of the data. In order to deliver the results in real time while enabling low-latency user interaction, this engine must use incremental approaches to the processing that generate partial results as new data enters the platform. This kind of processing is achieved by applying online scalable algorithms over distributed stream processing technologies that follow hybrid architectures.

Instead of using two parallel processing engines like traditional Lambda Architectures, one for processing batch data and another one for streaming data, we have developed a hybrid processing engine based on Apache Flink to process both kinds of datasets. The PROTEUS Hybrid Engine is a fork of Apache Flink that features an extended support to operations that mix streaming and batch processing. This is achieved through side inputs for stream operators, i.e., an operator processes records coming from the stream as well as preloaded historical data retrieved from the NoSQL database. As a result, side inputs enable hybrid processing within a stream operator.

Results generated at this layer are directed again to the Data Collector layer, implemented as a Kafka cluster, where results from each algorithm will be written in different topics. The visualization layer will read from these topics incrementally to generate real-time visualizations assuring low-latency user interaction with partial results. Besides, these results will also be stored in the NoSQL database as batch data in order to support later visualization of historical results.

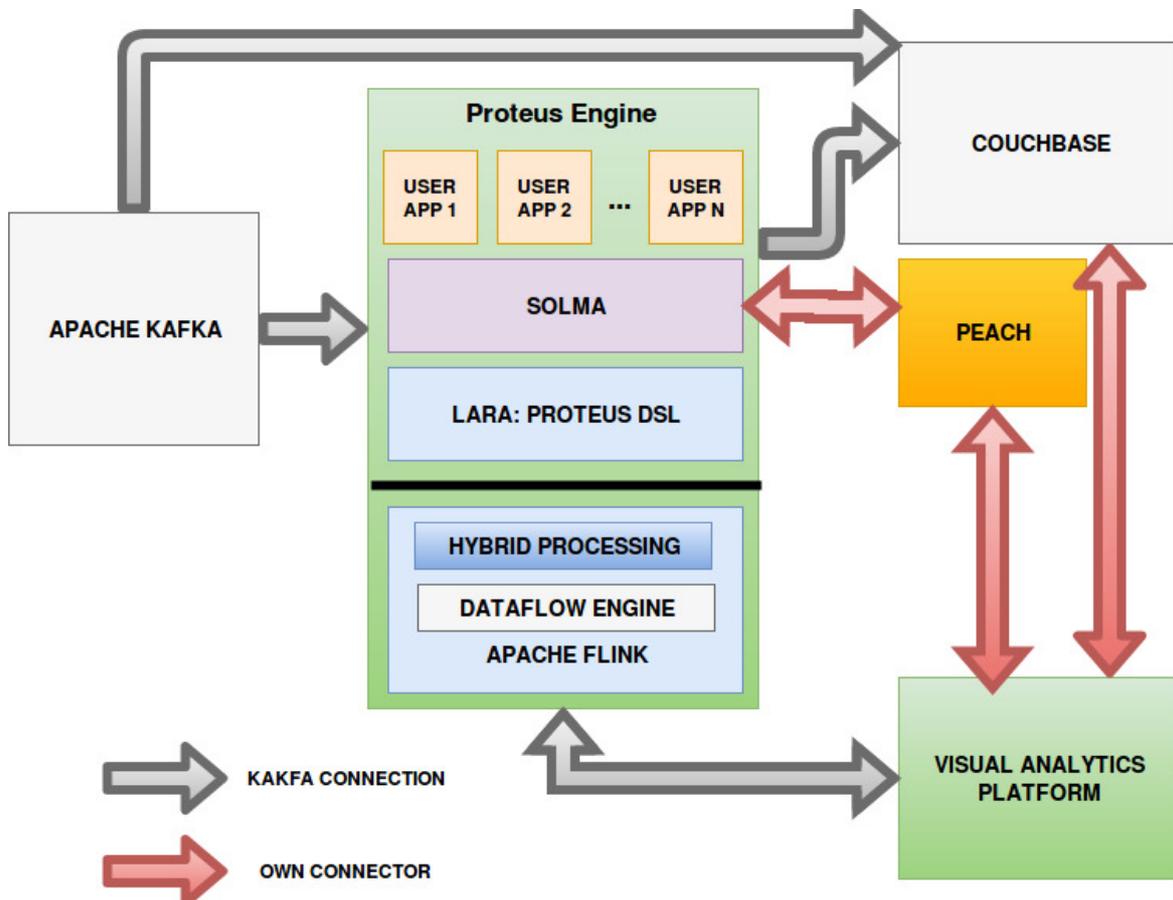


Figure 2: Architecture of the Incremental Analytics Engine in PROTEUS

The architecture of this layer, depicted in Figure 2, has been implemented using the following technologies developed in PROTEUS as the core contributions (see the WP3 and WP4 deliverables for the details):

- **PROTEUS Hybrid Engine:** a fork of Apache Flink. It features an extended support to operations that mix streaming and batch processing. This is achieved through side inputs for stream operators, i.e., an operator processes records coming from data stream as well as preloaded historical data. As a result, side inputs enable hybrid processing within a stream operator.
- **LARA:** LARA is an engine-independent declarative language that enables programmers to author complete end-to-end data analysis programs that are automatically parallelized. The main observation behind LARA is that data analytics first perform some pre-processing, which naturally fits the dataflow paradigm, and then they perform some machine learning tasks, which are expressed as linear algebra operations with iterations. Therefore, LARA aims to provide a common environment that enables to define pipelines that mix relational and linear algebra
- **Distributed Cache: This component store temporarily some data (e.g. machine learning algorithm models) which is needed to be cached during machine learning algorithm performance. We trust in two solutions to implement this component. Firstly PEACH was developed to execute this mission (deliverable D3.5).** Peach (PROTEUS Elastic Cache) is a distributed key-value cache that can be used both inside and outside of the PROTEUS scope. The cache aims to provide low latency responses on a distributed elastic deployment with fault-tolerance capabilities. As a generic design, the cache could be integrated within Apache Flink to speedup computing processes. In addition to PEACH, Flink Paramer Server [9] also was used in the project to perform this important mission. It is an open source solution which has helped PROTEUS team in the development of some of machine learning algorithms.
- **SOLMA:** scalable online machine learning library includes classification, clustering, regression, ensemble algorithms, graph oriented algorithms, linear algebra operations, and anomaly and novelty detection, all of them implemented on top of Apache Flink using the hybrid processing capabilities.

Deliverable D2.9 in section 2.2.3 describes each of these contributions and their deployment in production with more detail. Further information about each contribution can be found in the following deliverables:

- D4.3: description of the novel online learning algorithms implemented in SOLMA.
- D4.4: description of additional online learning algorithms implemented in SOLMA
- D3.7: design and implementation of the declarative language LARA.
- D3.2: design and implementation of the hybrid engine.
- D3.5: design and implementation of the state management system PEACH.

### 1.3 Visualization layer

Big data visualizations require a set of features that are uncommon in traditional visualization tools. For instance, they should be able to draw high volumes of both processed and raw data in real time. They provide querying capabilities so users can get an overview of the dataset while being able to explore the data in detail. Interface and presentation technologies provide users with capabilities to visualize and sense data from mobile devices to fully immersive systems. Interaction technologies let users change their visualization needs and adapt it according to their context.

As mentioned in the previous section, the Data Collector Layer aggregates all data in PROTEUS including raw data (streaming or batch) and results from the Incremental Analytics Engine Layer. This means that the visualization platform is able to access both raw data and processed data in real time and it can generate rich visualizations combining both data sources thanks to the low latency of the architecture and technologies used. Connection to the NoSQL database enables querying of historical and batch data, allowing for comparison of ongoing streaming visualization with previous results and fast decision making for operators in the industrial scenario.

This layer is designed as a client-server web application that is connected to topics in the message queue of the Data Collector layer, and also to the historical data in the NoSQL database. The client or front end of this application is a dashboard where operators can monitor the coil production process and query historical data,

while the server or back end is in charge of providing authentication, authorisation, and access control to the platform and to provide services for querying and connecting the dashboard to the Data Collector layer. Performance of the protocols and technologies chosen for this layer are crucial, since it lives outside of the cluster infrastructure and the front end will be run as a JavaScript application on commodity hardware like an ordinary laptop. WP2 will define KPIs to measure such performance.

To meet these criteria, we have implemented such architecture by developing the Visual Analytics Platform that integrates the following contributions in PROTEUS:

- The *back end* of the visual analytics platform is built as a Java EE application on top of the Spring Framework and can be deployed in an Apache Tomcat Server.
- *PROTEUS Dashboard* is the user interface for the platform. It is a cross-platform application that can run either as a web application or a desktop application by means of the Electron framework. The application enables the user to access a dashboard through a login screen where he/she can create and customize visualizations of high-dimensional and streaming datasets.
- The *Datasources API* is part of the back end in the Visual Analytics Platform. It allows bidirectional communication of the visualization at the client side of the platform with the server side through various protocols, namely WebSocket, HTTP or HTTP/2.
- *Proteic.js* is the data visualization library for the web created for the interactive visualization of the PROTEUS data and the results of the incremental analytics engine. It is integrated in the PROTEUS Dashboard (Figure 3) to display data retrieved from the back-end of the visual analytics platform and has been designed to support high volumes of data with optimal performance.

In the following sections, we describe the previous contributions in more detail. Information about the deployment in production of these technologies can be found in deliverable D2.9, and further details about the technologies in the Visual Analytics Platform, and the guidelines followed for its design are included in deliverable D5.2.

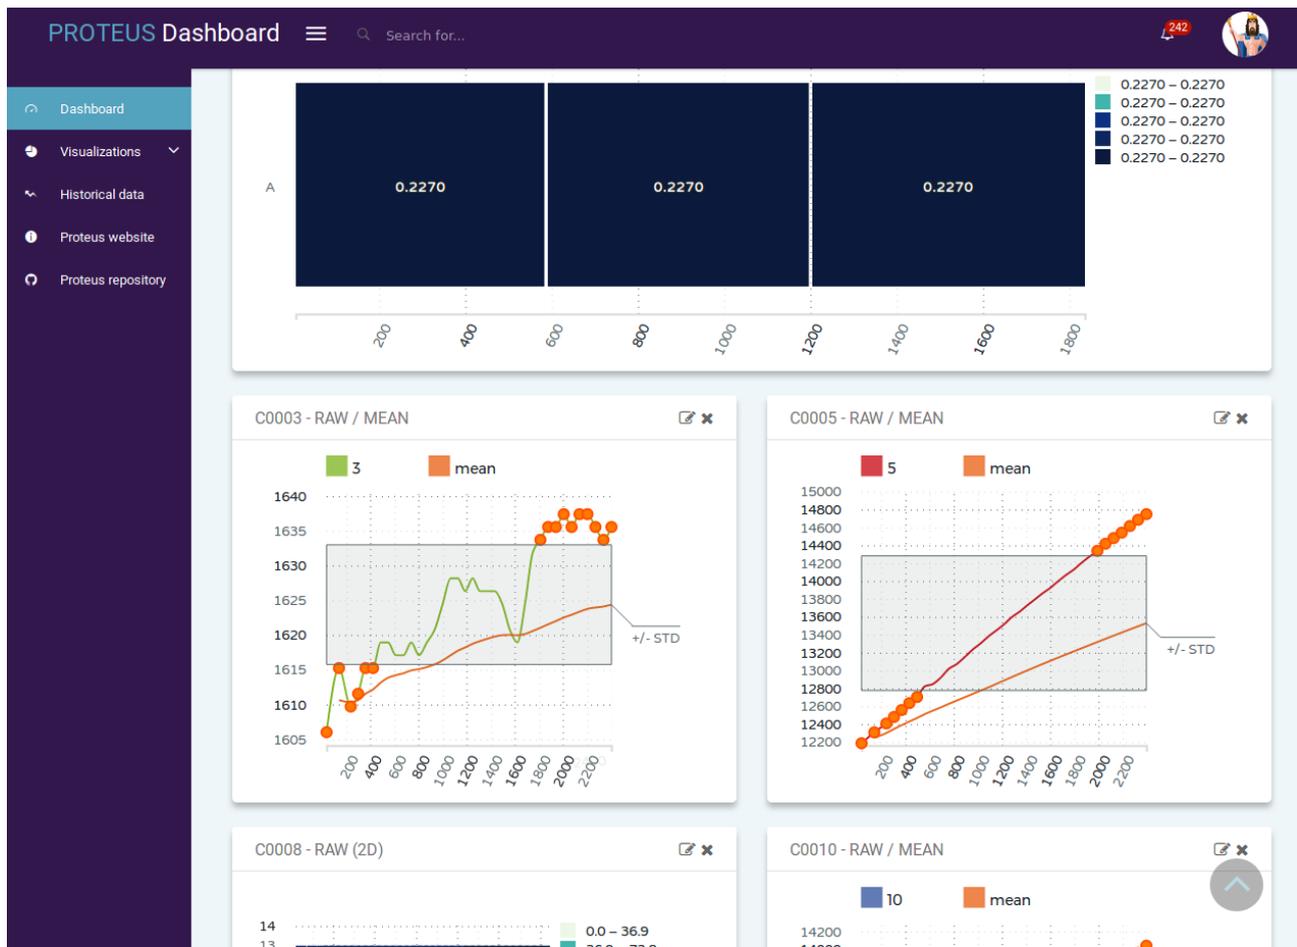


Figure 3: User interface of the Visual Analytics Platform

### 1.3.1 Visual Analytics Platform: back end

The backend for the Visual Analytics Platform is implemented as a Java EE application, built on top of the Spring Framework and deployed in an Apache Tomcat Server. It provides security features via Spring Security, querying of the historical database using Spring Data and N1QL code in order to retrieve data from CouchBase server, and a Kafka consumer to collect the messages from the processing engine.

#### 1.3.1.1 Security

The Visual Analytics Platform incorporates the Spring Security framework to preserve confidentiality of the PROTEUS dataset. This framework provides customizable authentication, authorization, access control, and protection against attacks like session fixation, click-jacking, cross site request forgery, etc.

#### 1.3.1.2 Database connection

The visual analytics platform is able to query historical data stored in the CouchBase server using N1QL (declarative query language that extends SQL for JSON documents), but it could be connected to other kinds of databases for other use cases, both relational and non-relational. To enable these queries, the backend employs Spring Data with the Couchbase connector [6], which provides a common programming model based on POJO for the data access layer to abstract the programming model from the underlying data store [7].

#### 1.3.1.3 Kafka Consumer

All data in the cluster infrastructure is written into Kafka topics, including raw data (real-time, HSM, flatness) and process data (results from the algorithms in the Incremental Analytics Engine). These topics are later read by a Kafka consumer in the visual analytics platform, and its data can be accessed from the

PROTEUS Dashboard using the Datasources API through various endpoints for different communication protocols, for instance STOMP and WebSocket.

### 1.1.1 Visual Analytics Platform: PROTEUS Dashboard

PROTEUS Dashboard is the user interface for the visual analytics platform. It is a cross-platform application that can run either as a web application or a desktop application through the Electron framework. The application enables the user to access a dashboard through a login screen where he/she can create and customize Proteic visualizations. These can then be connected to a set of data sources in the back-end of the visual analytics platform. The dashboard is based on the ng2-admin framework [8], which is built with the following cutting edge open source technologies:

- Angular: open-source web application framework based on TypeScript for developing front-end web applications, usually in the form of Single Page Applications (SPA). This is used in PROTEUS for building the web interface of the Visual Analytics Platform. The interface is designed as a set of reusable modules that enables creating, editing and deleting visualizations contained in cards.
- Bootstrap 4: open-source CSS framework for designing websites and web applications. It brings responsive and reusable front-end components that ease the design process of the platform.
- Electron: open-source framework for the development of desktop applications using front and back end components originally developed for web applications. This allows the Visual Analytics Platform to be used as a standalone desktop application.

### 1.1.2 Datasources API

The connection between client and server in the visualization layer is achieved by means of *Datasources APIs*, which enable bidirectional communication of the visualization at the client side of the platform with the server side. This API can be used from the client-side to list all the available data sources, either batch or streaming, and receive data through various protocols, for example, WebSocket, HTTP or HTTP/2. Thanks to WebSocket connections, we aim to achieve low-latency visualization of the streaming datasets. They provide asynchronous and bidirectional communication between client and server, allowing data to be drawn incrementally as messages arrive to the visualization [5]. Performance of this technique will be evaluated as part of WP2.

### 1.1.3 Proteic.js

Proteic is the data visualization library for the web created for the interactive visualization of the PROTEUS data and the results of the incremental analytics engine. It is integrated in the PROTEUS Dashboard to display data retrieved from the back-end of the visual analytics platform. The library features a convenient declarative API for the creation of multiple kinds of visualizations suited for highvolume streaming and static data. Its API enables the retrieval of data over the various communication protocols available in the Datasources API of the back-end component.

This component is built and distributed as a standalone open source library that can be used in any data visualization scenario for the web. It is implemented in the TypeScript language, and follows the latest standards and good practices in web development to achieve high performance, compatibility and ease of development. The design process of the library accounts for usability and accessibility of the various visualizations. It also follows design guidelines like responsive web design and offering alternative colour palettes for colour blind-users.

The Proteic library brings the following features to the visual analytics platform:

- Time series visualizations using dynamic and interactive line charts, stream graphs, and area charts. These charts include a set of tools for time series analysis such as dynamic thresholds, outlier alerts based on thresholds, dynamic error or confidence bands and pause / resume capabilities.
- User interaction features like brushing, zoom, highlighting, pause / resume; following the Shneiderman's mantra: *Overview first, zoom and filter, and details on demand* [1].

- Usability and accessibility features, such as interactive legends, annotations, accessible colour palettes, customizable themes, error warnings and a throbber that is shown while the visualization is waiting for data to arrive from data sources.
- High-dimensional geometry and multivariate data visualization by means of parallel coordinates visualization, including features for visualization of missing data.
- Innovative real-time visualizations, namely real-time heatmaps and swimlane charts.

## 2 Demonstrator source code

All the code developed during the project to implement this demonstrator has been created under open source licenses and is available in public GitHub repositories. Next list gives a brief description of all these repositories:

- Proteic.js (<https://github.com/proteus-h2020/proteic>)
- Proteus-web (<https://github.com/proteus-h2020/proteus-web>): It is the dashboard implementation. It is a web application with a backend (Spring), a front-end (AngularJS, Proteic.js) and WebSockets to communicate between them.
- Proteus-engine (<https://github.com/proteus-h2020/proteus-engine>): Fork from official Apache Flink repository with the hybrid engine implementation
- Proteus-SOLMA (<https://github.com/proteus-h2020/proteus-solma>): Online Machine Learning algorithms implemented on top data stream API of Flink.
- Proteus-job (<https://github.com/proteus-h2020/proteus-job>): Arcelor use case functionalities implemented with an Apache Flink streaming job.
- PEACH (<https://github.com/proteus-h2020/peach>): Distributed cache implementation.
- LARA (<https://github.com/proteus-h2020/streaming-lara>): LARA is an engine-independent declarative language that allows programmers to author complete end-to-end data analysis programs that are automatically parallelized

### **3 Demonstrator deployment**

Demonstrator deployment was already detailed in deliverable D2.10. Section 2.4 of that document describes it step by step.

## 4 Conclusions

Visualization of very large data sets is still an open research question, especially when dealing with streaming data at high frequencies that needs to be transformed, processed and merged with big volumes of batch data like it is the case in PROTEUS. These operations on streaming data cause delays that impede fluent interactions, and hinder usability of the tools and making impossible to perform effective data analysis. Existing big data architectures, like the Lambda Architecture and Kappa Architecture which add excessive complexity in the first case (Lambda), or do not fully meet the requirements of the PROTEUS industrial use case in the second case (Kappa), since it requires processing of both data-at-motion and data-at-rest.

This deliverable describes the demonstrator which implements the proposed 3-layer architecture designed (deliverable D.5.3) for PROTEUS as part of WP5.

## References

- [1] B. Shneiderman, “The eyes have it: a task by data type taxonomy for information visualizations,” Proc. 1996 IEEE Symp. Vis. Lang., pp. 336–343, 1996.
- [2] “The Continued Rise of Apache Kafka – Charting Stacks.” 2017. [Online]. Available: .
- [3] “Flume or Kafka for Real-Time Event Processing | LinkedIn,” 2015. [Online]. Available: .
- [4] “Understanding When to use RabbitMQ or Apache Kafka.” [Online]. Available: <https://content.pivotal.io/blog/understanding-when-to-use-rabbitmq-or-apache-kafka>. [Accessed: 24-Nov-2017].
- [5] C. Marion and J. Jomier, “Real-time collaborative scientific WebGL visualization with WebSocket,” in Proceedings of the 17th International Conference on 3D Web Technology - Web3D '12, 2012, p. 47.
- [6] “Spring Data Couchbase.” [Online]. Available:
- [7] C. Walls, Spring in Action. 2014.
- [8] “ngx-admin” [Online]. Available:
- [9] “Flink Parameter Server” [Online]. Available: <https://github.com/gaborhermann/flink-parameter-server>