

Tailored IoT & BigData Sandboxes and Testbeds for Smart,
Autonomous and Personalized Services in the European
Finance and Insurance Services Ecosystem



D5.4 – Framework for Declarative and
Configurable Analytics - I

Lead Beneficiary	LXS / LeanXcale
Due Date	2020-11-30
Delivered Date	2020-11-30
Revision Number	3.0
Dissemination Level	Public (PU)
Type	Report (R)
Document Status	Release
Review Status	Internally Reviewed and Quality Assurance Reviewed
Document Acceptance	WP Leader Accepted and/or Coordinator Accepted
EC Project Officer	Pierre-Paul Sondag

HORIZON 2020 - ICT-11-2018



This project has received funding from the European Union's horizon 2020 research and innovation programme under grant agreement no 856632

Contributing Partners

Partner Acronym	Role ¹	Name Surname ²
LXS	Lead Beneficiary	Ricardo Jiménez-Peris
LXS	Contributor	Francisco Ballesteros Boyan Kolev, Javier López Moratalla, Patricio Martinez, Jacob Roldan, Rogelio Rodriguez, José María Zaragoza Alejandro Ramiro Luis Miguel Garcia
CTAG	Contributor	Andrea Becerra García
UBI	Internal Reviewer	Dimitris Miltiadou
NOVA	Internal Reviewer	Giovanni Orio
INNOV	Quality Assurance	Dimitris Drakoulis

Revision History

Version	Date	Partner(s)	Description
0.1	2020-10-23	LXS	ToC Version
0.2	2020-11-01	LXS	Input on Introductory section and Executive Summary
0.3	2020-11-23	LXS, CTAG	Input in sections 2,3,4
0.4	2020-11-24	LXS	Input in conclusions and preparation for internal review
1.0	2020-11-24	LXS	Submission for internal review
1.1	2020-11-25	NOVA	Internal review
1.2	2020-11-26	UBI	Internal review
1.3	2020-11-27	LXS	Addressing comments from internal review
2.0	2020-11-27	LXS	Submitted for internal QA
2.5	2020-11-30	INNOV	QA
3.0	2020-11-30	LXS	Document is finalized and published

¹ Lead Beneficiary, Contributor, Internal Reviewer, Quality Assurance

² Can be left void

Executive Summary

The goal of Task T5.3 “Declarative Real-Time Data Analytics” is to provide a framework that can be considered as the enabler for real-time (online) data analytics over the data management layer of INFINITECH, in a declarative way, hiding the complexities of the implementation on the framework itself. Typical scenarios in both finance and insurance sectors require the continuous ingestion of data coming from a variety of sources in high rates, while at the same time they require the ability to perform complex and time-demanding analytical operations over the operational dataset, as the data is being inserted. Due to the current technological obstacles of combining those two different types of workloads, complex architectures are usually proposed and integrated in the current enterprise applications that allow performing such analytical operations. However, these analytical operations are performed over a snapshot of a database that is outdated, and not on real-time data. Those architectures are from one hand hard to implement, and additionally very difficult to maintain due to their complexity.

The INFINITECH project aims to provide an innovative data management platform that overcomes the current technological barriers throughout the data lifecycle, from data collection and ingestion, to analytical processing. The outcomes of this task will be part of the *Advanced Analytical Processing* layer of the platform and are built on top of the fundamental pillars implemented in the work that has been carried out in the corresponding tasks of WP3, which provides the enablers for these advanced analytics. As a result, *the Declarative Real-Time Analytics* will support the declaration, configuration and execution of analytical queries over an operational datastore, and will be able to return the result of those analytics in real-time, or as we better call it, *online*. This is a key requirement for the acceleration of the parallelized algorithms and will additionally benefit the correlation of data *at-rest* with data *in-flight*, used in the streaming processing framework of INFINITECH. As the results can be retrieved online, they are acquired with minimum latency as opposed to traditional executions of analytic operations, and the results are always up-to-date with the operational data. We extend the vanilla data structure of our data schema with the addition of new analytical columns that aim to facilitate the query execution, thus facilitating the analytical algorithms to retrieve post-processing results by directly connecting to the operational data store, instead of using complex and hard to maintain alternative architectures. Therefore, we introduce the term *online aggregates* to refer to data analytics over real-time data, where the requirement is to retrieve results with the minimum latency possible, while ensuring data consistency on the same time. The execution of those *online aggregates* has been designed to be declarative, as they can be defined and configured via SQL DDL statements and their execution relies on the standard SQL syntax.

This deliverable introduces the *online aggregates* and describes the general concepts and design principles behind their implementation. It provides examples explaining the problem statement and the motivation behind them. At this phase of the project, we provide the initial design of this framework, which is based on the semantic multi-version concurrency control implementation of the central data repository of the project. The design of this prototype along with an initial implementation has been already delivered, integrating the core mechanism of this component with the query engine of central data repository, whose improvements and features have been already described in the corresponding deliverables of T3.1 and T3.2. This integration allows the support of the definition of the *online aggregates* in a declarative way, using standard SQL syntax and delegating their execution to the query engine of the data management layer. As a result, significant effort has been put into the design of the SQL extensions and the documentation on how to use the real-time analytics. This documentation has been provided in deliverable, to allow the developers of the analytical algorithms of the pilot cases to utilise them. As the integration with the query engine is currently in progress, more advanced features and combination of aggregate operations will be supported, and will be described in the second version of the deliverable. In that second version, more technical information regarding the implementation of the core engine that allows for the

online aggregates will be given, along with benchmarking results, comparing our implementation with similar approaches used by other data base vendors. The third and final version of this document will provide the complete documentation of this framework, along with its validation using experimentation and use cases demonstrating its use.

Table of Contents

Contributing Partners	2
Revision History	2
Executive Summary	3
Table of Contents	5
List of Tables	5
Abbreviations	6
1. Introduction	7
1.1. Objective of the Deliverable	8
1.2. Insights from other Tasks and Deliverables	8
1.3. Structure	8
2. Motivation and Design Principles of Online Aggregates.....	9
2.1 A BigData application scenario as a reference example	9
2.2 Introducing Aggregate Tables	10
2.3 Problems Using Aggregate Tables with Traditional Approaches	10
2.4 Our solution: Online Aggregates.....	12
3. Using Declarative SQL to enable Real-Time Data Analytics	14
4. Real-Time Data Analytics in practice.....	17
4.1 Setting up the application.....	17
4.2 Implementing the application using online aggregates.....	18
5. Conclusions and next steps.....	23
6. References	24

List of Tables

Table 1: Example of sample rows in a table containing transactions.....	9
Table 2: Aggregate table containing summary of transactions	10
Table 3: Table containing information from a visitor's cookies.....	18

Abbreviations

2PL	Two-phase Commit protocol
BI	Business Intelligence
CSV	Comma-separated values
CTS	Commit Timestamp
DDL	Data Definition Language
ETL	Extract, Transform, Load
HTAP	Hybrid Transactional and Analytical Processing
IoT	Internet of Things
JDBC	Java DataBase Connection
NoSQL	No/Not only SQL
OLAP	Online Analytical Processing
OLTP	Online Transactional Processing
SQL	Structured Query Language
WP	Work Package

1. Introduction

The term real-time analytics refers to the ability of the framework or the architecture to provide analytical processing capabilities over the live-data, and not over a snapshot of the dataset that has been taken either by a periodically executed process (i.e. using ETLs) or via the data migration from the operational datastore (where the real data lives) to a data lake or an analytical data warehouse via micro batching using intermediate data queues.

Typically, real-time analytics are very hard to achieve as a major requirement is the ability to ingest data at very high rates, while at the same time, compute large aggregate queries over the real-time data.

This scenario is often a common use case in BigData applications where data is being ingested at high rates and the analytical algorithms need to compute KPIs or other metrics over the real or live data, as it has been ingested. Such use cases can be found in the area of performance monitoring, in cases where there is the need to calculate aggregated values of IoT sensors, e-Advertisements, online detection of opportunities of risk assessment, smart grids, industry 4.0, etc.

As data need to be stored in a persistent medium, data store solutions coming both from the SQL or NoSQL ecosystems seem to be inadequate to perform such operations.

In particular, SQL databases find this kind of workloads troublesome, as they are not efficient at ingesting data at very high rates, and in the meantime, aggregate analytical queries are very expensive, as they require to traverse very large amounts of data, and in many cases, the whole data table, very frequently. This, in combination with the fact that traditional approaches rely on the 2PL (two-phase locking) mechanism for ensuring transactional consistency, makes it impossible for them to execute analytical query operations on the operational and real-time data.

On the other hand, NoSQL databases can handle the data ingestion at high rates, thus being effective on serving this type of workload. However, they cannot be used to execute analytical queries, and in many cases they rely on external analytical frameworks to provide this type of functionality to the data user. They cannot be used in scenarios where operational guarantees in terms of transactional semantics are required, and they lack support of transactions.

As a result, modern approaches in the industry rely on solutions that need to combine different data management technologies to solve these use cases. This yields complex architectures that are very hard to be implemented successfully, and in any case, they are also very expensive to maintain.

The INFINITECH platform differentiates from those approaches by providing a framework for online data analytics, in a declarative manner, which is introduced in this report. We call this differential feature as *online aggregates*, a new technology that is being developed in the scope of this task, which is based on the brand new semantic multi-version concurrency control mechanism and enables the computation of aggregates (i.e. max, min, count, sum and average) in an incremental and real-time manner, using additional data structures in the storage engine. This creates new analytical columns that can be used by analytical algorithms. The latter can be facilitated by using those columns in a seamless way while directly connecting to the data management layer of the platform.

The specific benefit the *online aggregates* offer is to enable the update of the relevant aggregate tables, as data is being ingested, so aggregates are always pre-computed and the values can be retrieved online without the need to traverse the whole data table, an operation that has significant implications on the latency. Thus, a typical large and expensive analytical aggregate query, becomes an almost costless query that reads one or few rows of those introduced aggregate tables that contain the analytical columns. With our approach, data ingestion becomes slightly more expensive in terms of latency, however, the implementation of the Online Transactional Processing (OLTP) engine of the INFINITECH data management layer, as described in D3.1 (“Hybrid Transactional/Analytics Processing for Finance and Insurance Applications –I”), enables the sufficient

execution and can handle very high throughputs. The slightly more expensive data ingestion removes the cost of computing the aggregates in real-time.

In this report, we introduce the concept of online aggregates, showcasing how both a traditional SQL database and a NoSQL datastore fail to solve the problem, and we explain how this new technology developed in the scope of the INFINITECH project, succeeds at providing these advanced analytical capabilities. We also demonstrate how to use this feature by providing documentation with examples.

1.1. Objective of the Deliverable

The objective of this deliverable is to report the work that has been done in the context of task T5.3, at this phase of the project (M14). Two additional versions of this document will be released, extending and modifying when necessary the content of this deliverable. The work that has been currently delivered during the first phase of the project (M06-M14) was focused on the implementation of the brand new semantic multi-version concurrency control mechanism of the storage engine and its integration with the data management layer, in order to facilitate the provision of the *online aggregates*, that formulates the real-time analytics. In this deliverable, we report the rationale and motivation along with our approach and the design principles of our implementation. Additionally, a documentation on how to use this mechanism in a declarative manner has been added.

1.2. Insights from other Tasks and Deliverables

The work that has been carried out in the scope of T5.3 (“Declarative Real-Time Data Analytics”) has relied on the outcomes of the tasks of WP2 that define the overall user stories and requirements of the use cases of INFINITECH, and has been aligned with the INFINITECH RA that was previously introduced in D2.13 (“INFINITECH Reference Architecture – I”), under the scope of T2.7 (“Reference Architecture for BigData, AI and IoT in Financial Services Industry”). Apart from this, it relies on the outcomes of T3.1 (“Framework for Seamless Data Management and HTAP”), which implements the Hybrid Transactional and Analytical Processing (HTAP) framework that allows on the one hand for data ingestion in very high rates, and on the other hand, the combination of operational with analytical operations. The output of this task will be also beneficial to T3.3 (“Integrated Querying of Streaming Data and Data at Rest”) that implements the unified data query processing framework, which in turn allows the correlation streaming with batch processing. In particular, the provision of real-time analytics via the *online aggregates* can be used in the query processing, as it allows for the execution of analytical operations in real-time, with the minimum latency possible. Finally, this task also gives input to T5.2 (“Incremental and Parallel Data Analytics”), as the latter depends on the real-time analytics for the effective parallelization of its algorithms.

1.3. Structure

This document is structured as follows: Section 1 introduces the document and section 2 provides the rationale and motivation of the Declarative Real-Time Analytical framework, including the problem stating and introducing the notion of *online aggregates* that this framework is based upon. Section 3 provides the documentation of this framework, while section 4 includes a hands-on demonstrator on how to use *online aggregates* in an application, providing code examples as guidelines for the application developers and data scientists. Finally, section 5 concludes the document.

2. Motivation and Design Principles of Online Aggregates

This section describes the motivation and the design principles of the *online aggregates*. We describe a use case that is common in BigData applications in the finance sector, and we describe the fundamental pillar of our design: the *aggregate table*. After investigating why our approach is not feasible in traditional solutions relying either on relational SQL databases or NoSQL datastores, we explain the basic concepts of the *online aggregates* that make use of our semantic concurrency control mechanism and how our design solves the problems arisen with other approaches and respect the *Snapshot Isolation* paradigm that our transaction processing mechanism follows.

2.1 A BigData application scenario as a reference example

BigData applications often require the support of data ingestion at high rates, where data can be collected from a variety of sources. Such sources can vary from data produced by IoT sensors installed in the soil or in a vehicle and transmitted to an insurance organization, to the finance transactions of customers of a finance institution, records containing meta-information of a phone call collected by a telecom provider or data containing application metrics that is being acquired by a cloud provider in order to detect anomalies in the application behaviour for application performance monitoring.

For a finance institution, it might be crucial to monitor the maximum and average amount of money being used in a finance transaction in order to detect anomalies that might indicate fraud detection or money laundry operation from a client. Detecting those anomalies requires tracking the regular customer behaviour. Typically, this regular behaviour is modelled by several metrics: the maximum amount of money transferred in a single transaction, the average amount, the number of money transfers every given period of time and the overall amount being transferred. Additionally, there is a need to monitor and calculate these metrics at different aggregation levels: per customer, per region/area, per month etc.

In this example, we will describe how to compute this aggregation hierarchy in real-time, in an equivalent manner in terms of data consistency and in a cost effective way. Let's assume that there is a data table where each transaction of a client is being recorded. We will focus on two columns of that table, the name that identifies the client and the amount of money being transferred in a finance transaction. For instance, some sample rows of this table are illustrated in Table 1 below:

Table 1: Example of sample rows in a table containing transactions

Name	Value
Aleka	500
Pedro	250
José Maria	125
Aleka	200
Ricardo	480
Aleka	150
Ricardo	220

Large finance organizations like national banks or other similar institutions with millions of customers having different accounts have to keep track of all transactions that are being currently processed in their organizations. Depending on the size of the institutions, this might end up in the

ingestion of hundreds of thousands of records per second. As the bank might have millions of customers to keep track of, and for each one of them, it is required to keep track of the aforementioned different types of aggregations, the execution of each of those operations for each of its customers would be required, in a per-minute basis (or even less in case of money laundry and fraud detection cases). As each of those operations require the traverse of the majority of the data table, this is very resource demanding from the database management system perspective.

2.2 Introducing Aggregate Tables

In order to reduce the cost of each particular aggregate operation and the overall cost of the analytical queries, one solution could be to compute those aggregations incrementally in a separate structure, whose values can be retrieved in a more efficient manner. We call those structures *aggregate tables*. As a result, we could have an aggregate table per aggregation level. In our example, we could have one table containing the customer's summary amount of money transferred, the customer's maximum amount, etc. In real use cases however, more aggregation levels might be needed for different time periods and different aggregation levels. For instance, the finance institution might be interested to retrieve the average amount of money being transferred per customer and per month or day of the week, etc.

Sticking to our example, this *aggregate* table that contains the customer's summary amount of money transferred can be depicted in Table 2. This table has as key the name that identifies the customer. Then, every time a customer is performing a financial transaction, a new data item is being added in the structure depicted in Table 1, as part of the same database transaction, and the corresponding value in Table 2 is being updated according to its key. In the aggregate table, we would update the row with the associated name, updating the value by adding in the summary the new value. This will need to occur in each data insertion of the first table.

Table 2: Aggregate table containing summary of transactions

Name	Value
Aleka	850
Pedro	250
José Maria	125
Ricardo	700

2.3 Problems Using Aggregate Tables with Traditional Approaches

As described above, our approach for solving the problem of cost efficient execution of a real time analytic operation, while data is being ingested in high rates, is based on the definition and use of the *aggregate table*. However, traditional datastore vendors still cannot benefit of such a data structure, for different reasons. This concerns both SQL databases and NoSQL data management systems. We will present why this is happening for both these two different datastore ecosystems.

We will first examine the NoSQL world. What happens if we implement the aggregate tables based on a NoSQL data store? Let's consider that we have two concurrent invocations over the same row in the aggregate table, for instance the one that can be identified by the value '*Aleka*'. It is important to mention here that in real life scenarios deployed in production, it might be the case that millions of concurrent transactions might be inserting data over the same identifier, and as a result, trying to update the same row of the *aggregate table*. As NoSQL data store does not ensure data consistency in terms of database transaction semantics, in our example, both concurrent invocations will be

allowed to proceed. However, this will create the effect that one of them will be lost in what is called the “Lost Updates” anomaly[1]. This is due to the fact that in order to update the *aggregate row* for the name ‘Aleka’, the process will first need to execute a *get* operation, using the ‘Aleka’ value as the key, retrieve the current value, and then execute a *put* operation with the same value as the key, and the aggregated value incremented with the corresponding value of the money that has been transferred.

To concretize our example, let’s get back to the *aggregate table* of Table 2. Aleka has a summarized amount of money transferred with value 850. Now, two concurrent invocations appear for Aleka, with amount of money transferred set to 150 and 200 respectively. As we are using a NoSQL datastore that does not ensure data consistency in terms of database transactions, both are allowed to be executed. As we explained, this will require for each one of the two to perform a *get* operation, so that the current value 850 can be retrieved. The first invocation will have to calculate the accumulated value incremented by 150, and the second will increment this value by 200. This means that the first invocation will have to put the aggregate value of 1000 and the second one, the aggregate value of 1050. As both of them are being processed in parallel, let’s assume that the second succeeds first, and thus updates the value in the *aggregate table* to 1050. Then eventually, the first invocation succeeds, and updates the value to 1000. Now, the value of this aggregate row for ‘Aleka’ as the key, contains the value of 1000, instead of 1200, which is the accumulated summary of 850+150+200. This is caused due to what is called a *race condition* where two concurrent operations that share the same state (i.e. the value of the accumulated value 850 in the beginning) changes the value of the state in parallel and write it to the persistent storage. As a result, the second write operation erased the effect and the value of the first one. As *put* operations are blind writes, they do not take into account the value that was written before and just overwrite the latter.

This problem can be removed by using a traditional SQL database, which provides transaction semantics and ensures data consistency when concurrent invocations operate over the same data items. This happens using a multi-statement transaction, where an insert operation takes place on table depicted in Table 1, and then an update operation modifies the value of the *aggregate table*. Concurrent transactions over the same data item (the aggregated row of Aleka) are managed by the transaction processing mechanisms of the database, and one has to be blocked until the successful commit of the other. The following pseudo code can be used to describe such a transaction.

```
INSERT INTO Transactions (Name, Value) VALUES ('Aleka', 150)
DOUBLE currentvalue = SELECT FROM Aggregates WHERE Name = 'Aleka'
UPDATE Aggregates SET Value=currentvalue+150 WHERE Name='Aleka'
COMMIT
```

In this code snippet, the multi-statement transaction adds a new record in the first table with the value of 150 as the transfer, and then it reads the current value of the *aggregate table* for the corresponding table and updates its value accordingly, before commit. As the SQL databases uses transactional concurrency control imposed by their transactional processing mechanism, the “Lost Updates” anomaly disappears. However, one has to take into account that there are two different families of implementations of the transactional concurrency control: locking used by the 2PL (two-phase locking) algorithms and multi-version, used by implementations relying on the *Snapshot Isolation* paradigm. Even if they are very different, both those families are introducing an inherit contention problem. For instance, the transactional concurrency control forbids two concurrent transactions to modify the same row. This happens by allowing only one transaction to succeed and aborting (by a rollback) all other concurrent transactions that try to update the same data item. In our example, only the first database transaction that will try to eventually commit will succeed, and the second one will be aborted. This might probably be handled by the application by retry the insertion. In real use cases where hundreds or even thousands of concurrent transactions compete

each other to update a common data item, this will lead to a huge contention problem, no matter which implementation of the concurrency control algorithm the database vendor is making use of. In fact, relying on transactions to ensure data consistency and to remove the “Lost Updates” anomaly, will make our *aggregate table* proposition useless in use cases targeting Big Data needs.

2.4 Our solution: Online Aggregates

In order to solve this problem, we introduce at this phase the *online aggregates*. They are solving the problem with a new technology that is based on a novel semantic multi-version concurrency control provided by LeanXcale currently being integrated with the INFINITECH central data management layer. Relying on this new concurrency control, writes are not blind anymore. They actually carry the operation performed (i.e. sum(150) or sum(200) in our example). Since additions are commutative, they do not create write-write conflicts, as far as one keeps track that they are additions until the corresponding version of the row is written. Taking into account that this mechanism is built upon the operational mechanism of LeanXcale, integrated with the INFINITECH data management layer, this mechanism is making use of the *Snapshot Isolation* paradigm. As a result, in order to attain data consistency with respect to this paradigm, we have implemented the multi-versioning algorithm in a sophisticated way. The underlying distributed storage layer is now able to support a new kind of data structure, using *aggregate columns* (we often call them *delta columns*), which adopt this new semantic multi-version concurrency control and enable online aggregates. This can be also combined with the HTAP capabilities implemented under the scope of T3.1 which allows for *online aggregates* under data ingestion at high rates, due to the capabilities of the HTAP framework to scale linearly in hundreds of nodes, which in practice, makes it possible to serve incoming workload independently of its rate.

In order to dive deeper in how the *online aggregates* work, one has to remember that the aggregate columns are conflict-less, and thus, they do not create write-write conflicts. As a result, we can have two different kinds of rows: regular rows (often called value rows), which are rows that contain values as in any traditional database management system, and operation rows (often called delta rows) that represent operations to be performed over the columns of the rows. In our previous example, these rows will have values sum(150), sum(200) etc. Regular rows create conflicts, as any regular row in a traditional database management system, however delta rows do not, as they are conflict-less. However, since they do not create conflicts the key problem that arises is that there might be gaps in the commit order, as concurrent transactions cannot commit in an ordered manner, so there might be gasps in the commit order, and in fact, it will not be possible to generate the accumulated value for each delta, at the time a transaction, with an operational row, is committed.

To further clarify this issue, let's go back in our previous example. As we rely on the *Snapshot Isolation* paradigm, each transaction is being given a timestamp. More information on how this works can be found in section 2.3 of the D3.1. In our example, we assume that the transaction that needs to add the value 150 is called t1, and the one that needs to add the value 200 is called t2. Let's assume that t2 commits first and t1 commits second. As opposed to the multi-version concurrency control mechanism, in order to provide different snapshots of the data set, each row is being labelled with a commit timestamp (CTS), which is a numeric value that is being increased monotonically and determines the order of the commit. In our example, as t1 commits second, it gets a CTS=2 and t2 gets a CTS=1. If a new transaction t3 is started after the commit of t1 and before the commit of t2, it would get a snapshot of 0, in order to guarantee that it will not observe any gap in the commit order. If it would get the snapshot at timestamp 2, it would miss the updates from the transaction with CTS 1, which would violate the *Snapshot Isolation*.

With our design, we solve the problem by converting operation (delta) rows into value rows only when the snapshot of the database is beyond the commit timestamp of the operation rows. This guarantees that the serialization order is gap free and therefore the generated versions are guaranteed to be consistent and according to the semantics of snapshot isolation. In the previous example, Table 2 contains row (Aleka, 850, CTS=0). When the snapshot of the database reaches timestamp 1, then the operation row from t2 adding 200 will be converted into a value row generating the value row (Aleka, 200, CTS=1). When the snapshot reaches timestamp 2, then the operation row from t1 adding 150 is applied to the latest value row (the one generated by t2), yielding the following value row: (Aleka, 350, CTS=2). In this way, all the versions needed by the different snapshots are generated in a consistent way, thus fulfilling the snapshot isolation semantics.

To conclude, we demonstrate how the proposed *online aggregates* can be used in order to return the online accumulated value of an aggregate operation, by enforcing data consistency and transactional semantics, removing the “Lost Updates” anomaly and avoiding the high contention that can be observed when using the same technique with traditional operational SQL database management systems. As a result and in combination with the HTAP framework developed under T3.1, they can be used to provide real-time analytics. At this phase of the project, the initial design has been validated and we have provided a first implementation of our solution. The next section of this document reports on the use of this functionality, focusing on how this can be invoked in a declarative fashion. In the next version of the deliverable, we will include a separate section with more technical details of the implementation of the *online aggregates*, which provide the declarative real-time analytical framework of INFINITECH, as at this phase of the project, their implementation details has been submitted by LeanXcale for a patent, and cannot be reported in a public document.

3. Using Declarative SQL to enable Real-Time Data Analytics

This section provides a basic documentation on how to configure the use of the *online aggregates* and how an aggregate operator can be invoked by the application of a data analyst. In order to configure a table to exploit the *online aggregations*, the database administrator would need to declare this online aggregation table that is related to a raw table. We will continuously refer to the table containing the raw data with the regular rows as *parent table*. The following code snippet provides an example on how to declare those tables.

```
CREATE TABLE EVENTRAW (
  ev_id integer NOT NULL,
  ev_im_id integer,
  city char(24),
  ev_price integer,
  ev_data char(50),
  CONSTRAINT pk_eventraw PRIMARY KEY (ev_id)
);

CREATE ONLINE AGGREGATE ON EVENTRAW AS AGG EVENTBYCITY (
  city,
  max_price max(ev_price),
  count_price count(*),
  min_price min(ev_price),
  sum_price sum(ev_price)
);
```

In this example, we define a *parent* table called EVENTRAW, with a standard DDL statement. This declares that this table will have the *ev_id* as its primary key, and further defines four (4) additional columns containing raw data. Additionally, it defines the AGG_EVENTBYNAME as an *aggregated table* on the parent one defined above, with the definition of the row operators. We can notice that the *ev_name* will be the key of this table, as it does not contain any function, while the 4 additional columns will calculate online the values of the *max_price*, *count_price*, *min_price* and *sum_price*. From the code snippet, it can be noticed that these columns will calculate online the result of the aggregate operation (max, min, count, sum) of the column *ev_price*, that is included in the *parent* table.

After defining the *aggregate table*, the user might want to start inserting some data in those tables. A data insertion in the *parent table* will need to be accompanied by a corresponding insertion in the *aggregate table*. In order for these two statements to be atomic and to ensure data consistency when updating both tables, these statements need to be bracketed inside a single transaction. The following code snippet provides such an example.

```
UPSERT INTO EVENTRAW VALUES (1, 11, 'London', 10, 'aabbccdd');
UPSERT INTO AGG_EVENTBYCITY VALUES ('London', 10, 1, 10, 10);
COMMIT;
```

In this transaction, a new record is being added in the EVENTRAW table, with the corresponding values. At the same transaction, the record with key 'London' is being updated, by an UPSERT in the corresponding *operational row*. In this example, three (3) columns are calculating the value based on an aggregate operation on the *ev_price* column, and as a result, the UPSERT statement adds the value 10. This will be translated in operational values as max(10), min(10), sum(10) and the actual result of this operation will be calculated online. As the *count_price* delta column is operating on a count(*), for this column we added the '1'. It is important to mention at this point that the insertion on a *parent table* needs to be followed programmatically with an UPSERT on all their relevant

aggregated tables with which it is related. It is up to the application developer or data scientist to write these lines of code appropriately, and it is true that they can be error-prone. In the next version of the prototype of the *Declarative Real-Time Analytics Framework*, this will be made automatically by the framework itself, as it has been planned to be implemented during the second phase of the project.

The important innovation of this approach is that we rely on the semantic concurrency control to ensure data consistency while executing these two statements, while at the same time removing the high contention that they can introduce when being executed within a traditional SQL relational database management system. Let's have a look at the following code snippet, which illustrates how the multi-statement transaction defined above should be re-written using a typical SQL database.

```
UPSERT INTO EVENTRAW VALUES (1, 11, 'London', 10, 'aabbccdd');
UPDATE AGG_EVENTBYCITY SET
  MAX_PRICE=MAX(MAX_PRICE, 10),
  COUNT_PRICE=COUNT_PRICE+1,
  MIN_PRICE=MIN(MIN_PRICE, 10),
  SUM_PRICE=SUM_PRICE+10
WHERE CITY = 'London';
COMMIT;
```

In the case above, according to the type of database, it can cause several problems. As explained in the previous section, NoSQL databases will suffer from the “Lost Updates” anomaly that will be revealed when executing the update statement in the AGG_EVENTBYCITY. In particular, the relevant {attribute} = {attribute} + value expressions will require a *get* operation, followed by a *put*, which will cause the race condition when two or more concurrent invocations occur, which is actually the cause of the “Lost Updates”. Relational SQL databases on the other hand, will suffer from huge contention, as concurrent updates will target the same operational row identified by the key ‘London’. On the contrary, by managing these types of operations using the semantic concurrency control, there is no contention. The only impact will be on data ingestion, as it will require two operations instead of one, however, the overhead is very low and the overall impact on getting the pre-calculated values of the aggregated operations is significantly much more important.

Let's investigate now how the application developer or the data analyst will have to perform an aggregate operation on the *parent table* that contains the original raw data. In order to get the max and average price per city, he / she should submit the following SQL statement:

```
SELECT city, MAX(ev_price), AVG(ev_price) FROM EVENTRAW GROUP BY city
```

This statement will indeed return the maximum and minimum prices, grouped by the city. However, this is a cost-demanding operation, as the database management system will have to traverse the whole data table, in order to retrieve all *ev_prices* for the cities, group them by the city column, and then apply the operation. As this will require the traverse of the whole table, its latency is significant high, and as a fact, it cannot be considered as online, neither can it be used in use cases where the online calculation of these values is crucial. Additionally, in traditional operational datastores based on the 2PL (two-phase locking) protocol for ensuring transactional semantics, it would require from the database management system to acquire *shared locks* in each data items that is being accessed. This would prevent any other write operation to be performed in parallel, as all data modification operations need to acquire an *exclusive lock*, which will be forbidden, as *shared locks* have already been put by the aggregate operation.

Using the *online aggregates* proposed by INFINITECH's *Declarative Real-Time Analytics* framework, that statement will be internally translated to the following:

```
SELECT city, MAX_PRICE, SUM_PRICE/COUNT_PRICE FROM AGG_EVENTBYCITY
```

This second query runs in much lesser time compared to the previous one, whose execution time is much higher due to the need to traverse a lot more rows. Even if we apply a filter condition, in order to get the maximum and average price of a specific city (i.e. adding a `WHERE city='London'`) in the first example, it will require the traverse/scan of the data table, which has a complexity of $O(n)$ or $O(\log(n))$ (according to whether or not the column `city` is indexed). On the contrary, using online aggregates, the *online aggregate* operation costs $O(1)$ according to the theory of complexity.

Our approach can be compared to calculating those values in memory, so that you can have the results already pre-calculated and have the results returned immediately. However, as we are dealing with persistence, doing those operations in memory cannot be done. There are several issues regarding data persistency and fault-tolerance in cases of unexpected shutdowns or crashes, concurrency control when having to deal with high rates of parallel ingestion and most importantly, memory is not infinite, and therefore, there is a limited size of records that such a memory-based implementation could handle. Instead of using an in-memory implementation, we push all those issues to be solved by the data management layer, which has been designed to solve those and provides high availability, fault-tolerance and concurrency control mechanisms, while it uses the persistent storage volumes to scale out.

To conclude, the *online aggregates* are a really powerful mechanism because they allow you to have pre-computed data immediately available to serve your application. Note that even if the data analyst needs a complex KPI, this may be composed of pre-computed aggregates. That's the case of the standard deviation statistic for example. Other common scenarios relevant to the finance and insurance institutions can be also foreseen. For instance, *online aggregates* can be used in cases where there is the need for an immediate statistical aggregate in order to provide real-time results for an application. This is the case of the risk assessment use cases of INFINITECH. Moreover, working with multi-resolution data is also a good example. There might be a system whose source devices send information every second, but most statistics can be calculated much more easily at minute resolution. Aggregates can be used to store aggregated data at minute resolution and at 15 minute resolution will have second raw data. This is the case of the insurance pilot#11 of INFINITECH, which relies on IoT data coming from sensors installed in the vehicles, and they are pre-processed before being sent to the central application component deployed within the sandbox. Finally, together with multidimensional partitioning, online aggregates can be seen as pre-computed Online Analytical Processing (OLAP) cubes. This is a very powerful data to have available with little latency to get it. The following section provides a hands-on example on how to make use of this framework.

4. Real-Time Data Analytics in practice

In this section we will provide a hands-on example on how to make use of the *Real-Time Data Analytics* framework, developed in the context of T5.3, based on a hypothetical scenario of an application that monitors the web traffic generated by end-users visiting a website. The requirement is for the business analyst to have a dashboard that he / she can be able to analyse the user visiting the website every second. In case the system architect does not want to implement a typical but complex lambda architecture [4], but instead, to rely only on a single database management system, then it would be required to periodically execute count operations over a group-by clause at the time of the insertion. However, this would cause the loss of performance due to the fact that *reads* are competitive with *writes* in an operational datastore. As there are usually many more *insertions* than *reads*, a solution could be to cache the results of the analytics in memory, yet with all the drawbacks explained in the previous chapter. Instead, the system architect could decide to exploit the *online aggregates* provided by our framework in order to serve those aggregate operators in real-time.

Online aggregates are built on top of the transactional and analytical processing (HTAP) provided by the data management layer in the scope of T3.1, which is the fundamental pillar for those operations. Our approach makes use of the *delta column*, which is capable of providing the result of an aggregation query, such as the one described above, pre-calculated at the time of insertion, and persistently stored in the storage medium. This way, getting the aggregate for a value requires simply reading a row from the relevant aggregate table, which is already pre-calculated, instead of doing a scan to find the right row and calculate the aggregation. This in turn, means aggregations in real-time. They enable the calculation of aggregates of any kind over the data management layer, without executing a heavy group-by query and without losing performance on an insertion. This example will demonstrate this, which simulates a monitoring application.

As explained above, when implementing these types of applications, it is very important to take into consideration the balance between two things: the performance of the data insertions and the query execution. In the majority of the cases, the number of insertions is dominant, as the rate can be hundreds of thousands of records per second, while a read operation is executed periodically. As a result, it seems reasonable to prioritize the insertion performance. However, as our hypothetical application would need to report various things that would need to execute aggregations over the end-user's cookie IDs, when having high volumes of data, those aggregations are very low, and they can block the insertions. This means that the system is either going to lose currently inserted data or that the reported data are not going to be real.

4.1 Setting up the application

The application that will demonstrate the use of the *online aggregates* consists of the following components:

- The INFINITECH data management layer.
- The Declarative Real-Time Analytics framework, which makes use of the data layer.
- A program³ that simulates data insertion and executes queries periodically.
- An SQL client, to visualize the results of the queries (SQuirreL [2] or DBeaver [3], are recommended).

Both the INFINITECH data management layer and the Declarative Real-Time Analytics framework are available in the project's private repository and will be retrievable by the INFINITECH Marketplace.

³ https://gitlab.com/leanxcale_public/onlineaggregation

For our hands-on demonstrator, we will generate some controlled data to simulate the insertion of thousands of cookies per second. An example table containing information of the user when visiting a website can be depicted as follows:

Table 3: Table containing information from a visitor's cookies

Cookie ID	Post Date	Other fields
c034868a-2e3e-11eb-adc1-0242ac120002	2020-11-17 15:48:52	{value_1}
c0348a36-2e3e-11eb-adc1-0242ac120002	2020-11-17 15:48:53	{value_2}
c0348b44-2e3e-11eb-adc1-0242ac120002	2020-11-17 15:48:53	{value_3}
c0348c16-2e3e-11eb-adc1-0242ac120002	2020-11-17 15:48:54	{value_4}

The architecture of the solution of this hypothetical scenario consists of i) a loader process and the program that performs the query executions, ii) the INFINITECH data management layer, with the addition of the Advanced Analytical Capabilities provided by the Declarative Real-Time Analytics framework and iii) an SQL client, which can be one of the proposed ones above. We can benefit from the versatility of the data management layer itself, which provides various ways for data connectivity, using an SQL or a No-SQL interface. In our example, we rely on the No-SQL interface for both insertion and query execution, in order to benefit from its improved performance, as it bypasses the query engine of the datastore, however, there is always the possibility of executing SQL queries, and we can make use of the SQL client for that. Having this versatility to make use of the *online aggregates*, it makes it possible to use them with whatever Business Intelligence (BI) tool the data analyst is familiar with or is better for his / her purposes.

Regarding the data model, we will create a table similar to one depicted in Table 3 in order to store the cookie information and two additional tables to store the delta aggregates that will be pre-calculated at the time of the insertion.

4.2 Implementing the application using online aggregates

Firstly, we will implement the data inserter program, which will be built using Java and Spring Boot. It will respond to the following three types of request:

- Run: to start the insertion.
- Clean: to clear the database.
- Query: to start the periodic query executor.

It will make use of a CSV file where it will start reading its lines and it will ingest them to the database.

Before starting, the application developer needs to make use of the No-SQL interface to enable the data connectivity with the data management layer. Using maven, it can be locally installed in the maven repository and then it can be used as a dependency. The following code snippet depicts this process:

```
mvn install:install-file -Dfile=kivi-api-1.6-direct-client.jar -
DgroupId=com.leanxcale -DartifactId=kivi-api -Dversion=1.6 -Dpackaging=jar

<dependency>
  <groupId>com.leanxcale</groupId>
  <artifactId>kivi-api</artifactId>
  <version>1.6</version>
```

```
</dependency>
```

Now it is the time to create the main table containing the row data from the end-users visits, and the corresponding *aggregate tables*. The main table will have an alphanumeric ID, and a timestamp, which will form a composite primary key. In standard SQL this can be done by executing the following statement:

```
create table INFO (
  ID VARCHAR,
  POSTDATE TIMESTAMP,
  OTHERFIELD VARCHAR,
  CONSTRAINT PK_INFO PRIMARY KEY (ID, POSTDATE)
);
```

Programmatically via the No-SQL interface, this can be achieved by executing the following code snippet:

```
private static final String TABLE_NAME = "INFO"

(...)

Settings settings = new Settings();
settings.credentials(new Credentials()
  .setDatabase(databaseName)
  .setUser(user)
  .setPass(password.toCharArray()));
settings.transactional();

try (Session session = SessionFactory.newSession(URL, settings)) { // New session
// Table creation:
if (session.database().tableExists(TABLE_NAME)) {
  // Primary key fields
  List<Field> keyFields = Arrays
    .asList(new Field[]{new Field("id", Type.STRING),
      new Field("postdate", Type.TIMESTAMP)});
  // Rest of fields
  List<Field> fields = Arrays
    .asList(new Field[]{
      new Field("otherfield", Type.STRING)
    });
  // Table creation
  session.database().createTable(TABLE_NAME, keyFields, fields);
}
}
```

In order to retrieve the information regarding how many different users have visited the monitored web page and how many times, or how many visits per day our page has had, then the database administrator will have to execute SQL queries such as the following:

```
select id, count(id) from info group by id;

select FLOOR(POSTDATE to DAY) as fecha, count(*) AS total FROM info group by
FLOOR(POSTDATE to DAY);
```

In order to take advantage of the Declarative Real-Time Analytics framework and execute these queries online in order to get the pre-calculated aggregated value with the same cost of accessing a

row by its primary key, we need to use the *delta* columns to holds operational values. Those operational values will pre-calculate the result of the aggregate operation. For that, we are going to create an *aggregate table*. We would need to define two: one for the ID count and one for the date count. The field from which we are going to aggregate must be the aggregate table PK, and the aggregator must be defined as delta in the field creation. The following code snippet indicates how to do it:

```
private static final String ID_COUNT_DELTA_TABLE_NAME = "INFO_ID_DELTA";
private static final String DATE_COUNT_DELTA_TABLE_NAME = "INFO_DATE_DELTA";

(...)

Settings settings = new Settings();
settings.credentials(new Credentials()
    .setDatabase(databaseName)
    .setUser(user)
    .setPass(password.toCharArray()));
settings.transactional();

try (Session session = SessionFactory.newSession(url, settings)) { // New session

// Aggregation table creation:
if (!session.database().tableExists(ID_COUNT_DELTA_TABLE_NAME)) {
    List<Field> keyFields = Arrays
        .asList(new Field[]{new Field("id", Type.STRING)});
    List<Field> fields = Arrays
        .asList(new Field[]{
            new Field("count", Type.LONG, DeltaType.ADD)});
    session.database().createTable(
        ID_COUNT_DELTA_TABLE_NAME, keyFields, fields);
}
// Aggregation table creation:
if (!session.database().tableExists(DATE_COUNT_DELTA_TABLE_NAME)) {
    List<Field> keyFields = Arrays
        .asList(new Field[]{new Field("postdate", Type.DATE)});
    List<Field> fields = Arrays
        .asList(new Field[]{
            new Field("count", Type.LONG, DeltaType.ADD)});
    session.database().createTable(
        DATE_COUNT_DELTA_TABLE_NAME, keyFields, fields);
}
```

We have defined two tables, one per aggregate:

- The first table, INFO_ID_delta, corresponds to query select id, count(id) from info group by id;. Since we want to group it by ID, the ID is going to be the PK of this delta table. The aggregate will be the count, and we will add 1 to the pre-calculated aggregate (please note the delta field is declared as DeltaType.ADD).
- Similarly, the second table, INFO_DATE_delta, corresponds to query select FLOOR(POSTDATE to DAY) as fecha, count(FLOOR(POSTDATE to DAY)) AS total FROM info group by FLOOR(POSTDATE to DAY); the table PK is going to be the postdate field, and the aggregate will again be the count.

After the schema and the *aggregated tables* definition, we are going to write a code to insert data from the CSV file. It is important to highlight at this point that an insertion in the main table INFO must be followed by a corresponding insertion in the *aggregated tables* that we previously defined. The following code snippet illustrates this:

```

try (Session session = SessionFactory.newSession(url, settings)) { // New session
    session.beginTransaction();
    Table infoTable = session.database().getTable(TABLE_NAME);
    Table infoIdTable = session.database().getTable(ID_COUNT_DELTA_TABLE_NAME);
    Table infoPostdateTable =
        session.database().getTable(DATE_COUNT_DELTA_TABLE_NAME);

    // Insert into info table
    Tuple tuple = infoTable.createTuple();
    tuple.putString("id", id);
    SimpleDateFormat dateFormat = new SimpleDateFormat(
        "dd-MM-yyyy HH:mm:ss.SSS");
    Date date = dateFormat.parse(postdateString);
    tuple.putTimestamp("postdate", new Timestamp(date.getTime()));
    tuple.putString("otherfield", otherfield);
    infoTable.insert(tuple);

    // Insert into info id delta
    Tuple tupleIdDelta = infoIdTable.createTuple();
    tupleIdDelta.putString("id", id);
    tupleIdDelta.putLong("count", 1L); // We add 1 to our DeltaType.ADD field
    infoIdTable.upsert(tuple);

    // Insert into info date delta
    Tuple tupleDateDelta = infoPostdateTable.createTuple();
    tupleDateDelta.putDate("postdate", new java.sql.Date(date.getTime()));
    tupleDateDelta.putLong("count", 1L);
    infoPostdateTable.upsert(tupleDateDelta); // We add 1 to our DeltaType.ADD
                                                field

    session.commit();
}

```

Finally, in order to retrieve data using the *online aggregates*, the following code snippet can be used, that relies on the No-SQL interface that we provide.

```

try{
    session= SessionFactory.newSession(url, settings);
    infoTable = session.database().getTable(Constants.TABLE_NAME);
    infoIdTable = session.database().getTable(Constants.ID_COUNT_DELTA_TABLE_NAME);
    infoPostdateTable = session.database().getTable(Constants.DATE_COUNT_DELTA_TABLE_NAME);

    while (executions < 1000) {
        log.info("-----");
        log.info("LX: Querying info table for number of rows...");
        session.beginTransaction();
        // Build TupleIterable, execute find with count aggregation and iterate
        // over the result (select count(*) from info)
        TupleIterable res = infoTable.find().aggregate(Collections.emptyList(),
        Aggregations.count("numRows"));
        res.forEach(tuple -> log.info("Rows: " + tuple.getLong("numRows")));
        log.info("LX Querying info table done!");

        log.info("-----");
        log.info("LX: Querying id delta table...");
        long t1 = System.currentTimeMillis();
        // select * from info_id_delta

```

```

    TupleIterable res2 = infoIdTable.find();
    res2.forEach(tuple -> log.info("Id: " + tuple.getString("id") + " Count: " +
tuple.getLong("count")));
    long t2 = System.currentTimeMillis();
    log.info("LX Query time: {} ms", t2 - t1);
    log.info("LX Querying id delta table done!");

    log.info("-----");
    log.info("LX: Querying postdate delta table...");
    long t3 = System.currentTimeMillis();
    // select * from info_date_delta
    TupleIterable res3 = infoPostdateTable.find();
    res3.forEach(tuple -> log.info("Date: " + tuple.getDate("postdate") + " Count: " +
tuple.getLong("count")));
    session.commit();
    long t4 = System.currentTimeMillis();
    log.info("LX Query time: {} ms", t4 - t3);
    log.info("LX Querying id delta table done!");

    executions++;
    Thread.sleep(5000);
}
}
finally {
    if (session == null) {
        session.close();
    }
}
}

```

To conclude, we demonstrated how we can implement a hypothetical application that could benefit from the use of *online aggregates* and we demonstrated with this hands-on chapter the capabilities of the Declarative Real-Time Analytical framework of INFINITECH. Currently, there is ongoing work that is performed to extend the SQL syntax and which will provide automation in the query engine level, so that the definition of *table aggregations* and insertions in tables that trigger insertions on *aggregated tables* can be automatically executed in the query engine level. In the second version of this deliverable, we will include the improvements and extensions of the work that has been currently presented with those aspects.

5. Conclusions and next steps

This report documented the work that has been carried out in the scope of task T5.3 “Declarative Real-Time Data Analytics” at this phase of the project. The main objective of this task is to provide a framework that will enable the execution of data analytics in real-time (online) in a declarative fashion. Towards this, we first identified the need for BigData applications in finance and insurance sector to execute data analytic operations in real-time, while their datasets are being continuously updated by a data ingestion process at high rates. We presented the main technological barriers that are currently being tackled by the system architectures and what solutions exist and are commonly used. Then we investigated how we could benefit by pre-calculating the aggregated value beforehand and by storing this value in a separate field, inside a particular column of a data table. For this, we proposed the notion of *aggregate tables*. We verified the feasibility of our proposal when being used by both SQL and NoSQL database management systems, and we realized that both sides of the data ecosystem deal with significant issues: in NoSQL world data consistency is broken due to the *Lost Updates* anomaly, while in SQL solutions, efficient execution of the statements is not feasible due to the high contention that the database transaction create. In order to overcome this, and exploit the hybrid transactional and analytical processing (HTAP) capabilities of the INFINITECH data management layer, we introduced the *online aggregates*. Those can in fact assist in overcoming the barriers explained before by relying on the semantic concurrency control mechanism of the database.

After introducing the *online aggregates*, we provided a documentation on how these can be used. The important thing here is the definition of the *aggregated tables* which are needed to store the operational (delta) rows along with the tables themselves that holds the raw data. Additionally, each insert operation in a raw table will require for the application developer or data scientist to also take care of the fact that he / she needs to insert the corresponding operational rows in the *aggregated tables*. Then, by executing a standard SQL statement, the query engine can transform it and take benefit from the additional structures that contain the pre-calculated values. Finally, a hands-on demonstrator was presented that can be used as a guideline for application developers and data scientists on how to integrate their solutions with our framework.

To conclude, the progress of the task T5.3 can be considered ahead of the plan at this phase of the project, as all its main objectives have been already mostly addressed. This was important as this task provides the outcomes that will be used as the main pillars for the development of the parallelization of the algorithms, which is under T5.2, and which will be exploited by the query processing framework, developed in T3.3. In the next phase of the project, additional extensions and improvements will be implemented, in order to provide a seamless way on the definition of the *aggregated tables*, and data ingestion and the query execution using the *online aggregates*. Additional technical details regarding the implementation of the delta tables will be reported, along with information regarding the protocol of the semantic concurrency control that this framework is based on. Finally, a benchmark evaluation will be provided in the second version of the document via experimentation with other solutions, comparing the execution time of our framework with other approaches used by other database vendors (i.e. PostgreSQL triggers).

6. References

- [1]. Tamer Özsu, Patrick Valduriez. Principles of Distributed Database Systems, 4th Edition, Springer, 2020
- [2]. Squirrel, SQL UI client, <http://squirrel-sql.sourceforge.net/>
- [3]. DBeaver, SQL UI client, <https://dbeaver.io/>
- [4]. What is a lambda architecture? <https://databricks.com/glossary/lambda-architecture>