Ref. Ares(2021)5310710 - 27/08/2021

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

D3.5 – Integrated (Polyglot) Persistence – II

Revision Number	3.0
Task Reference	T3.2
Lead Beneficiary	UBI
Responsible	Konstantinos Perakis
Partners	UBI, LXS, UPRC
Deliverable Type	Report (R)
Dissemination Level	Public (PU)
Due Date	2021-04-30
Delivered Date	2021-08-27
Internal Reviewers	RB, HPE
Quality Assurance	INNOV
Acceptance	WP Leader Accepted and Coordinator Accepted
EC Project Officer	Pierre-Paul Sondag
Programme	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

© INFINITECH Consortium

Contributing Partners

Partner Acronym	Role[1]	Author(s)[2]
UBI	Lead Beneficiary	Konstantinos Perakis, Dimitris Miltiadou
LXS	Contributor	Pavlos Kranas Ricardo Jiménez-Peris, Boyan Kolev, Patricio Martinez, Francisco Ballesteros, Rogelio Rodriguez
UPRC	Contributor	Ioannis Kranas, Christos Doulkeridis
RB	Internal Reviewer	Victoria Michailidou
HPE	Internal Reviewer	Alessandro Mamelli
INNOV	Quality Assurance	Filia Filippou

Revision History

Version	Date	Partner(s)	Description
0.1	2021-06-24	UBI	ToC Version and updated initial input of D3.4
0.2	2021-06-28	UBI	Change formatting
0.3	2021-07-09	UBI, LXS, UPRC	Contributions in Section 6
0.4	2021-07-16	UBI, LXS, UPRC	Updated contributions in Section 6
0.5	2021-07-20	LXS	Finalisation of Section 6
0.6	2021-07-22	UBI, UPRC	Updating the overall document
1.0	2021-07-28	UBI	First Version for Internal Review
1.0_HPE	2021-08-02	HPE	Internal review
1.0_RB	2021-08-03	RB	Internal review
2.0	2021-08-17	UBI	Version for Quality Assurance
3.0	2021-08-27	UBI	Version for Submission

[1] Lead Beneficiary, Contributor, Internal Reviewer, Quality Assurance

[2] Can be left void

Executive Summary

The goal of task T3.2 "Polyglot Persistence over BigData, IoT and Open Data Sources" is to provide a common and integrated way to access data that is stored in a structured, semi-structured or even unstructured fashion over a variety of heterogeneous data stores, in a unified manner. The rationale behind this task is that in modern enterprises, especially in organizations coming from the finance and insurance sectors, it is a necessity to access and process data coming from multiple sources. As there is no *one size fits all* data management system, different databases have been proposed to server different types of needs and workloads. Due to this, a finance organization may use on one hand traditional relational datastores that ensure transactional semantics and provide data consistency, which stores data in a structured schema: the relational model. On the other hand, there is always the need for using NoSQL databases that can scale more adequately, sacrificing the support for transactional semantics, and often using less strict data models, so that the data can be considered as semi-structured. Document-based datastores are a prominent example, where the schema of the data can be easily modified and extended in order to be adapted in rapidly changed software. Key-value stores are also often used in order to store information coming from IoT devices, a pattern that is a norm in the insurance companies where sensor data from smart devices or agricultural sensors are being continuously ingested, or logging information tracking the web traffic or the finance transactions of a customer, which is of great interest for the finance sector. Key-value stores are candidate where there is a huge write-intensive workload with rare data modifications, where the data management system needs to scale out easily. Even if they do not provide any guarantees over the schema of the *value,* they are considered as semi-structured, as the data tends to preserve the schema. Finally, such organizations need to process data coming from external sources, like feeds from social media networks, or newspaper articles, in order to extract trends and opportunities. This type of data can be considered as unstructured and are often being imported into a data lake and by using specific tools for data analysis, this type of information can be extracted.

Accessing and processing data coming from such a variety of heterogeneous data sources is a very complicated task in various levels. Firstly, each of the data stores provides distinct and specific methods for data connectivity, while on the same time, makes use of different query languages. Even different relational databases might provide different connectivity (i.e. JDBC vs ODBC) and even if they all rely on the SQL query language, its dialect might differ significantly. The data scientist must be aware and familiar with a variety of different tools and languages. Moreover, not all of the data base management systems return data in the same model. Relational databases return data as table with relationships, which can be easily transformed into an entity-relational model. However, key-value and NoSQL stores do not comply with specific schemas and the application developer or the data scientist must maintain different models and create a common layer on top. Most importantly, when there is the need to combine data coming from different sources, where joining data sources is mandatory, this must be done in the application or data processing level. However, joining data sources is a very challenging task and there is a lot of literature on how to do it effectively, which has been implemented by the database management systems. In fact, this operation should be transparent to the data scientist and the application developer and should be implemented in the data management level, that can do it more effectively. Due to this, various data analytical frameworks have been used during the recent years that provide a common way to address these challenges. However, they still require to fetch the majority of data into the data processing layer, which can become very resource consuming.

Task T3.2 "Polyglot Persistence over BigData, IoT and Open Data Sources" has as its main objective to provide a solution for efficient and unified integrated access over a variety of heterogeneous data stores that provide data in all structured, unstructured or semi-structured fashions. As an integral part of the INFINITECH data management layer, its goal is to provide a common API that can be used in order to access data seamlessly, hiding the internal complexities from the data scientists and application developers, therefore provide *polyglot* capabilities to the platform. In order to achieve this, a state-of-the-art analysis of the existing frameworks that can be considered as *polystores* has been included first, aiming to identify the best practices that have been proposed in the literature and are widely used by modern enterprises today,

and to highlight their weaknesses and current challenges that are open to research to further improve their solutions.

Moreover, as it was mentioned before, different datastores often provide their own query language for accessing and retrieving data. One of the requirements for the polyglot data management system that we are presenting in this deliverable, is to provide a seamless way to access data coming from a variety of heterogeneous sources. Due to this, we have defined the INFINITECH Common Query Language that will be used by the data scientists in order to request data for analysis in a unified manner. The INFINITECH Common Query Language has a lot of similarities with the standard SQL, allowing the data scientists to write queries in well-known standards. It also allows writing expressions in the native language of a data store, in cases the user wants to exploit some unique characteristics of the database that cannot be expressed with standard SQL.

At the first phase of the project, we concluded on the principle architectural design pattern that will drive the whole design of the solution, and the main building blocks and components of the solution have been also designed. We followed the mediator-wrapper approach, where one orchestrator drives the whole execution of the query, and different wrappers implement the data connectivity and query execution on the target data stores, hiding the internal complexity and becoming transparent from the orchestrator. The orchestrator is now part of the central data repository and data management layer of INFINITECH, also known as the INFINISTORE, and it extends in order to provide the polyglot capabilities. It has been incorporated in the query engine of INFINISTORE, in order to take advantage of the already available functionalities that are crucial for the effective data retrieval: the query planner, the query optimizer and the query executor. All of them have been further extended in order to take into account the polyglot capabilities that the engine now offers.

A complementary solution that has been implemented during the second phase of the project, is what we call the *Real-Time Data Warehousing*. Its objective is to solve the inherit technological barriers that modern system integrators have to deal with when having vast amounts of data that become obsolete after a period of time and thus can be considered as historical. Having to deal with both operational and historical data introduces several restrictions and different architectural solutions can be used, each one of those having its own drawbacks however. Our approach provides a holistic query processing framework for accessing both operational and historical data that can be split from the operational datastore to a data warehouse. Our implementation is built upon the polyglot query engine that had been provided at the first phase of the project.

Finally, it is important to highlight that this is the second version of this deliverable and reports the work that have been done in the scope of task T3.2 until M19. At this phase of the project, apart from the necessary initial state-of-the-art analysis of other polystore systems, the biggest part of the effort was spent in the definition of the INFINITECH Common Query Language, which is the basis for this task. A great effort had also been spent in order to design the basic architecture of the solution and the fundamentals that have driven the implementation and the more advanced functionalities that have been planned for the next periods. At this phase, we also provide a concrete implementation of the solution, allowing for a query execution over federated databases: the central data repository of INFINITECH, a relational JDBC-compatible database, a MongoDB instance or a data lake. In the third and final version of this document, an extensive evaluation of our solution will be provided.

Table of Contents

1. Introduction 7
1.1. Objective of the Deliverable 7
1.2. Insights from other Tasks and Deliverables 7
1.3. Updates from the previous version (D3.4) 8
1.4. Structure 8
2. State-of-the-Art Analysis on Polystores 9
3. INFINITECH Common Query Language 11
3.1 Query Language 11
3.2 Bind Join 12
3.3 MFR Extensions 13
4. Design of Integrated (Polyglot) Persistence 15
4.1 Basic Architectural Principal: The mediator/wrapper paradigm 15
4.2 Parallel Integrated Processing using the INFINITECH Common Query Language 18
5. Implementation of the Integrated (Polyglot) Persistence 21
5.1 Abstract Wrapper Implementation 21
5.2 RDM Wrapper Implementation details 22
6. Real-Time Data Warehouse 26
6.1 Motivation 26
6.2 Design Principles and Requirements 28
6.3 The Query Federator 30
6.4 Ensuring data consistency while concurrently moving data across the datastores 32
7. Conclusions 36
8. Appendix A: Literature 37

List of Figures

Figure 1: The Mediator/Wrapper paradigm in INFINITECH Polyglot Persistence 17
Figure 2: RDBWrapper class diagram 21
Figure 3: Lambda Architecture 27
Figure 4: Current-Historical Data Splitting Architectural Pattern 28
Figure 5: Simple full scan 32
Figure 6: Moving data from INFINITORE to the data warehouse 33
Figure 7: Second read transaction while data is being moved 34
Figure 8: First read operation finishes 34
Figure 9: Dropping the data slice 35

Abbreviations/Acronyms

API	Application Programming Interface
BDVA	Big Data Value Association
DoA	Description of Action
ETL	Extract, Transform, Load
HDFS	Hadoop Distributed File System
HTAP	Hybrid Transactional and Analytical Processing
IoT	Internet of Things
JDBC	Java DataBase Connectivity
JSON	JavaScript Object Notation
MFR	Map, Filter, Reduce operations
NoSQL	Not Structured Query Language
ODBC	Open DataBase Connectivity
OLAP	Online Analytical Processing
RDBS	Relational DataBase System
SQL	Structured Query Language
WP	Work Package

1 Introduction

This deliverable summarizes the work that has been done in the scope of task T3.2 "Polyglot Persistence over BigData, IoT and Open Data Sources" at the second phase of the project (M19). The goal of this task is to provide common and integrated access over structured, unstructured or semi-structured data that can be stored in a variety of different and heterogeneous data management systems. These systems include traditional relational databases, NoSQL stores, Hadoop data lakes, etc. Task T3.2 will provide a unique interface which will be based on the JDBC specification in order to make use of a native API and a common query language that is very close to SQL, in order to make use of a native scripting language for query processing, as well. By providing both a common API and scripting language for transparent data retrieval from heterogeneous data stores, INFINITECH has extended its data management system with polyglot capabilities, thus providing an integrated data layer that can be used transparently by the Open APIs and the semantic framework. Moreover, built upon the polyglot query engine, we have also implemented the Real-Time Data Warehousing that solves main technological barriers when dealing with both operational and historical datasets.

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 T3.2 at this phase of the project (M19). This task lasts until M27, and therefore, an additional version will be released, extending and modifying, when necessary, the content of this document, following the agile approach for system development, in order to update the solution and implementation with the current trends of the environment as the project progresses. The work done during this phase (M03-M19) was mainly focused on the definition of the INFINITECH Common Query Language, which is the basis of the whole solution. Moreover, the basic architectural design of the solution and details about the implementation have been provided in order to validate our approach. In the second version of this document, we have extended this document with the details of the implementation of the **real-time data warehousing.** In the third and final version, it is planned to provide an extensive evaluation of our approach.

1.2 Insights from other Tasks and Deliverables

The work that is reported in this deliverable is based on the overview description of the corresponding task T3.2, which has been further specified at WP2 level, which is the fundamental work package that defines the overall requirements of the whole platform. More precisely, task T2.1 with the corresponding D2.1 deliverable refers to the user stories of the pilots that will be accommodated by the platform, and reports their user requirements. To this direction, task T2.3 with the corresponding D2.5 deliverable defines the specification of the technologies that INFINITECH provides, and translates the user stories and requirements to specific technical requirements that must be addressed by the technologies. Moreover, task T2.5 and deliverable D2.9 provide a list of data asset specifications, where the available target data store that need to be taken into account, has been defined. Last but not least, the work that is reported in this document is also related with the Reference Architecture of the INFINITECH, as the solution provided by the Integrated Polyglot Persistence is an integral part of the platform. It is worth mentioning that even if there is no direct connection with task T3.2 and WP7, there have been discussions with all pilots in order to further clarify their needs and their proposed solutions, as the pilots are getting in more technical details in the work that is being currently done in WP7. Finally, T3.2 is at the lower layer of the Reference Architecture which follows the guidelines of the BDVA, and therefore, its output will be taken into account by the system components that are located in the upper layers, and more precisely the semantic interoperable engine in WP4, and the analytical tools that will be developed in the scope of WP5.

1.3 Updates from the previous version (D3.4)

In this version of the report we have added section 6 which is built upon the polyglot query engine that was delivered in the first phase of this task. Section 6 provides the details of the **real time data warehousing**, whose objective is to provide a holistic architectural design that can allow for federated query processing over the operational and the historical part of a common dataset, allowing the data movement from the one to the other in real time, with no downtimes and ensuring the consistency of the data in terms of database transactions when data is being moved. We provide the motivation behind our solution and then the basic principles, requirements and architectural decisions of our architecture. Finally, we give analytical details on how we implement the federated query processing and how we ensure the consistency of the data, which means on-going database transactions return equivalent results when data is being moved concurrently from the operational datastore to the data warehouse.

1.4 Structure

This document is structured as follows: Section 1 introduces the document, putting the work reported in this deliverable under the context of the project, highlighting its relationship with other tasks of the DoA. Section 2 provides a state-of-the-art analysis of existing solutions and frameworks in the wider technological and scientific area of *polystores.* Section 3 introduces the INFINITECH Common Query Language that will be the basis for accessing data across different data stores, while section 4 presents the overall design of the solution. Section 5 provides details on the initial implementation that has been provided as a validation of the architecture of this component, targeting one external datastore that is compatible with JDBC. The additional section 6 of this second version of the deliverable provides the details of the real time warehousing, while finally, section 7 concludes the document.

2 State-of-the-Art Analysis on Polystores

Accessing data from diverse and heterogeneous data sources has been long studied and various solutions have been proposed, usually called multidatabase or data integration systems [1][2]. The most typical approach that can be found in a variety of those proposals involves the definition of a common data model and query language that can be used in order to transparently access data sources via the mediator-wrapper paradigm, which aims to hide the details of the diverse data connectivity and distribution. In the latest years, with the emergence of cloud databases and big data processing frameworks, the multidatabase solutions became into what we call nowadays polystores systems. The latter, enable integrated access to traditional relational database management systems, NoSQL solutions and Hadoop data lakes via a common query engine.

A first category of polystores can be considered as loosely-coupled, and resembles much of the traditional multidatabase systems that can deal with autonomous datastores, being accessed via a common interface, following the mediator-wrapper paradigm, where the access is being granted via the common interface exposed by the mediator, while the wrapper implements the details on how to connect, access and retrieve data from the source. Most of the loosely-coupled systems can only support read-only operations. In this category, BigIntegrator[1] integrates data from cloud-based NoSQL big data stores, such as Google's Bigtable, and relational databases using its own query language, which does not support however Hadoop data lakes, while QoX [3] integrates data from RDBMS and HDFS data stores. SQL++ [4] mediates SQL and NoSQL data sources through a semi-structured common data model. Its query engine is capable of translating the subqueries of a statement to native queries that will be executed in the target datastores. BigDAWG [5][6] on the other hand, instead of translating the different datastores into a common data model, it defines the *islands of information*, where its island corresponds to a specific data model and its language and provides transparent access to the target database. On top, it enables cross-island query execution by exchanging and moving intermediate results between the different islands.

Another category of polystores have an opposite approach, and they are considered tightly-coupled polystores. Their goal is to integrate Hadoop or Spark for big data analysis with traditional relational database management systems. They tend to trade autonomy for performance, and they provide massive parallelism, using shared-nothing nodes in a cluster, and can benefit from the use of high-performance computing. One example is Odyssey [7] which enables storing and querying data within HDFS and RDBMS, using opportunistic materialized views. MISO [8] aims to tune the physical design of a multistore system in order to improve the overall performance when used in big data processing. JEN [9] aims at joining data coming from two different datastores such as Hadoop and traditional relational database management systems, parallelising join algorithms and minimizing data movement when executing these algorithms. Polybase [10] enables HDFS access using SQL scripting language. Moreover HadoopDB [11] provides MapReduce access to several RDMS systems that are supported, it establishes data connectivity and then execute SQL queries that return key-value pairs to be further processed by the MapReduce. Teradata IntelliSphere [12] provides an integrated data access over heterogeneous data sources, that are SQL compatible though.

Apart from these two categories of polystore systems, in the latest years, hybrid solutions that have recently been evolved and are widely used by the industry, are proposed in the literature. They support data source autonomy as loosely-coupled systems do, while on the other hand, they preserve parallelism by exploiting the local datastores, as in tightly-coupled systems. They can be seen as parallel query engines that provide several different connectors to external sources that can be executed in a parallel fashion as well. One representative of this hybrid category is Spark SQL [13] which is a parallel SQL engine that provides tight integration between relational and procedural processing via a declarative API and takes advantage of massive parallelism when executing the query statements. It defines the notion of DataFrames that map arbitrary object collections that are being retrieved from the local datastores into relations, and thus, it enables relational operations. Presto [14] is a distributed SQL query engine running on a cluster of machines and can process analytical queries against big data sources via massively parallel

processing. This is achieved by using a coordinator process, and multiple workers that make use of connectors that provide the interface with the external data sources and provide all kind of metadata information to the coordinator to optimize the query execution. Apache Drill [15] is also a distributed query engine for large-scale datasets that makes use of massive parallel processing. *Drillbit* services run at each node and receive the query, compile it to an optimized execution plan and exploit data locality in order to further extend the level of parallelism. Myria [16] is yet another recent polystore, built on a shared-nothing parallel architecture, which efficiently federates data across diverse data models and query languages. Finally, Impala [17] is an open-source SQL engine with massive parallelism capabilities, operating over Hadoop data processing environment. As opposed to typical batch processing frameworks for Hadoop, Impala provides low latency and high concurrency for analytical queries.

The main difference with INFINITECH Integrated Polystore engine is that not only does it enable parallel integration with external data sources, but it also combines massive parallelism with native queries that enables the exploitation of the unique characteristics of the target data management system. Moreover, all aforementioned solutions either retrieve the majority of the data in an intermediate layer or process there the query in a parallelized fashion, or push down the query execution to the node, which implies a lot of data movement when joining data sources from different stores.

One innovative aspect of the INFINITECH Integrated Polystore engine is its ability to efficiently execute joint operations using *bind joins.* A query statement can be analysed in subqueries, each one of those is targeting an external datastore. These are handled by table functions, which can be considered as query operators that can be considered by the query planner and query optimizer of the central data management layer of INFINITECH. Once the integration of the Polystore engine of the platform with the data management is achieved, then bind joins can be proposed automatically when selecting the optimal query execution plan and data execution over integrated heterogeneous datastores, that can be more effective. Combining this aspect with the ability to express a subquery by native language or scripts, allows to fully exploit the power and unique characteristics of the target data management system, as opposed to static mappings to a common data model used to execute query statement across datastores.

3 INFINITECH Common Query Language

As mentioned in the previous section, the Integrated Polyglot Persistent component of the platform fits into the hybrid category of polystores, being both loosely-coupled, thus providing autonomy on the external datastores, and on the same time tightly-coupled, providing a mechanism for massive parallelism, having a common data model. The main difference with the aforementioned solutions is that the data model is not static but instead, it can be created dynamically fitting to the needs of the submitted query and the target datastores. To do this, we need a novel query language that can allow this. Towards this direction, we introduce the INFINITECH Common Query Language that will be based on the CloudMdsQL [18]. This Common Query Language is an SQL-based scripting language with extended capabilities that allows embedding subqueries to be expressed in terms of each data store's native query interface. The common data model respectively is table-based, with support of rich datatypes that can capture a wide range of the underlying data stores' datatypes, such as arrays and JSON objects, in order to handle non-flat and nested data, with basic operators over such composite datatypes.

3.1 Query Language

In this subsection, the design and the basic principles of the INFINITECH Common Query Language will be presented. The latter requires a deep expertise and knowledge by the data scientist and application developers regarding the specifics of the underlying datastores in order to exploit their unique characteristics, as well as awareness about how data are organized across them. It is important to highlight that the Integrated Polyglot component takes into account only read-only operations, allowing the data ingestion and modification to happen at the data store level. As a result, the integrated queries will make use of *projection* over several *selections* that consist of native subqueries. On the component level, the results of the *selections* are *joined* and the *projection* is being applied. In the level of the scripting language, that will involve a SELECT statement over native subqueries. The latter are defined as *table functions*, also called *named table expressions*. This means that a *function/expression* is being submitted and returns a virtual table, which has a name and a signature, consisted by the names and types of the columns of that virtual table. The *function/expression* can be either a regular SQL SELECT statement or a native statement expressed in the query interface of the target datastore. It is important to notice that the Integrated Polyglot component includes a query compiler, which can analyse the SQL statement and re-write it in order for the engine to execute it more efficiently. To highlight the difference between SQL and native query, let's assume that we have an integrated query targeting a traditional relational SQL compatible database and MongoDB[3], which is a document-based datastore with its own query interface. The integrated query expressed in the INFINITECH Common Query Language will need to join two subqueries, one expressed in standard SQL and the other in a native way. The query will be the following:

```
T1(x int, y int)@rdb = (SELECT x, y FROM A)
T2(x int, z array)@mongo = {*
  return db.A.find( {x: {$lt: 10}}, {x:1, z:1} );
*}
SELECT T1.x, T2.z FROM T1, T2
WHERE T1.x = T2.x AND T1.y <= 3
```

In this example, we define two named table expressions, T1 and T2 which target datastore rdb (an SQL compatible database) and datastore mongo (a MongoDB database). T1 is defined by standard SQL while T2 holds a native expression (it is included in the bracket symbols {* *}). The result will be the *projection* of T1.x and T2.z from the *joint* operation of the results of those two subqueries that will be sent independently to the two target datastores. At this point, we need to focus on the fact that the query includes a *filter* condition (T1.y <= 3) that is applied on T1, which holds a standard SQL statement. The query compiler of the Integrated Polyglot component can identify an optimization and can push down this

[3] https://www.mongodb.com/

filtering in the T1 itself, in order for the results of the latter to be much less, as the filtering will be applied at the external datastore level.

Apart from being able to send native or SQL queries, the INFINITECH Common Query Language also allows table functions to be defined as expressions in a scripting language (i.e. Python, JavaScript). This is useful when datastores offer only API-based query interface and do not simply allow connections where the user submits a query statement. These scripting expressions can either return the returned tuples into a table-like result set or return an *iterable* object representing the result set, as the example above with MongoDB. Moreover, the INFINITECH Common Query Language can use as input the result of other subqueries, thus allowing for nested queries.

Last but not least, in the same way it is possible for traditional data management systems to create views or stored procedures, our language allows to create *named expressions* via the corresponding command. These expressions are defined during the execution of this command and stored in the global catalogue and can be later re-used and referenced by other queries. This can be of great importance for the data analysts who do not understand the deep details and unique characteristics of each of the underlying datastores, rather than they are interested in executing statements and make their analysis on the results. The data scientist or administrator who fully understands the underlying data technologies and the specifics of the data organization can prepare those named expressions to be frequently re-used.

3.2 Bind Join

As we have mentioned in section 2, one of the problems of the polystore systems is their inefficient implementations of *join* operations when they need to combine data coming from several data sources. They either retrieve all intermediate results of the involved subqueries in memory or apply the operation using massive parallelism, or they tend to move the datasets across the data stores in order to perform the operation at the lower level, without though avoiding the movement of the large amount of data containing the majority of the intermediate results. The Integrated Polyglot component of INFINITECH instead, makes extensive use of *bind joins*, a technique firstly proposed by IBM that has been later described thoroughly in the literature [19]. It allows performing semi-joins across datastores efficiently, by rewriting the subquery statement in order to push down the join conditions. This means that the distinct values of the attribute(s) that are involved in the join condition that are retrieved by the left operat (the subquery that is on the left side of the *join* clause) is passed as a filter to the right operat, through an *in* clause. The following example aims to clarify this:

```
A(id int, x int)@DB1 = (SELECT a.id, a.x FROM a)
B(id int, y int)@DB2 = (SELECT b.id, b.y FROM b)
SELECT a.x, b.y FROM b JOIN a ON b.id = a.id
```

In this example, we have two subqueries projecting their ids and values from a selection over two tables (one per subquery), and an *equity join* operation on those subqueries over their ids. The query optimizer will make use of the *bind join* technique and will bound the results of the left operat to the right-hand side of the operation. During the query execution, the relation B will be retrieved from the DB2 datastore using its own query mechanism. Then, the Integrated Polyglot Persistent component will make use of the list of the distinct values of B.id and will push them into an *in* operation down to DB1 to filter the values of the selection of the table A. Assuming the list of values of B.id are [b1, b2 …. bn], the query A will be transformed as follows:

```
SELECT a.id, a.x FROM a WHERE a.id IN (b1, …, bn)
```

In case of native queries, the compiler cannot re-write the query in the native subquery, and therefore, the data scientist or application developer has to explicitly declare its use in the script, making use of a *JOINED ON* clause in the signature of the named table, as the example below that involves a native script compliant with the MongoDB interface.

```
A(id int, x int JOINED ON id
```

```
    REFERENCING OUTER AS b_keys)@mongo =
{*  return db.A.find( {id: {$in: b_keys}} );   *}
```

In this example, the *in clause* of the native statement will take as parameters the values of the outer intermediate result that is being passed into the *b_keys* reference. By doing this, the query executor is enforced to submit firstly the subquery A in order to retrieve the values to be further used in the reference parameter, to then provide this reference to the native query.

Even if the bind join technique seems very promising and reduces a lot the execution time and data movement, thus making the execution of the query much more efficient, it is not a panacea and is restricted by several causes which can create an important overhead. Firstly, the bind join operation demands that the query executor waits for the completion of the execution of one of the two operats, in order to retrieve all required values, before pushing them down for filtering out the results of the second operat. Secondly, if the number of distinct values of the join attribute that will be pushed down is large, the bind join operation may slower down the performance as it will require to send a lot of data via the network, while the selectivity of the data on the second operat will not be adequate enough to filter out values, and as a result, neither the query execution time on the second subquery will be reduced, nor the amount of retrieved data will be lowered. Due to this, it is important for both external datastores to expose a cost model in order for the query optimizer to predict the number of rows and distinct values that each of the subqueries is expected to return. However, if a native query is involved (thus, the query compiler cannot understand what this is about, as it sees it as a black box) or the cost information is not available, the query can still take this decision, but on the runtime: it will attempt to perform a bind join, it will start collecting the intermediate results of one of the two operats, and if the number of the distinct join keys exceeds a certain threshold, then execution will fall back to an ordinary *hash join* (as bind joins can be used over *equity joins*). In any case, the data scientist and application developer have the best of the knowledge of the data distribution and can explicitly request the execution of a bind join by using the reserved word BIND in the statement (i.e. FROM b BIND JOIN a).

3.3 MFR Extensions

To address distributed processing frameworks (such as Apache Spark) as data stores, the INFINITECH Common Query Language introduces a formal notation that enables the ad-hoc usage of user-defined map/filter/reduce (MFR) operators as subqueries to request data processing in an underlying big data processing framework (DPF) **Error! Reference source not found.**. An MFR statement represents a sequence of MFR operations on datasets. In terms of Apache Spark, a dataset corresponds to an RDD (Resilient Distributed Dataset – the basic programming unit of Spark). Each of the three major MFR operations (MAP, FILTER and REDUCE) takes as input a dataset and produces another dataset by performing the corresponding transformation. Therefore, for each operation there should be specified the transformation that needs to be applied on tuples from the input dataset to produce the output tuples. Normally, a transformation is expressed with an SQL-like expression that involves special variables; however, more specific transformations may be defined through the use of lambda functions. Let us consider the following simple example inspired by the popular MapReduce tutorial application “word count”. We assume that the input dataset for the MFR statement is a text file containing a list of words. To count the words that contain the string ‘cloud’, we write the following composition of MFR operations:

```
T4(word string, count int)@spark = {*
   SCAN(TEXT, 'words.txt')
  .MAP(KEY, 1)
  .REDUCE(SUM)
  .FILTER( KEY LIKE '%cloud%' )
*}
```

For defining map and filter expressions, the special variable TUPLE which refers to the entire tuple, can be used. The variables KEY and VALUE are thus simply aliases to TUPLE[0] and TUPLE[1] respectively.

To optimize this MFR subquery, the sequence is subject to rewriting according to rules based on the algebraic properties of the MFR operators, as explained in **Error! Reference source not found.**. In the example above, since the FILTER predicate involves only the KEY, it can be swapped with the REDUCE, thus allowing the filter to be applied earlier in order to avoid unnecessary and expensive computation. The same rules apply for any pushed down predicates, including bind join conditions.

4 Design of Integrated (Polyglot) Persistence

This section will illustrate the basic architectural principles that the Integrated Polyglot component of INFINITECH follows and will focus on the details of the overall design of the proposed solution. Moreover, it will give some technical insights on how the integrated query processing can be executed in parallel and how the high expressivity of the INFINITECH Common Query Language fits together with the design of the implementation.

4.1 Basic Architectural Principle: The mediator/wrapper paradigm

During the state-of-the-art analysis on the area of the polystore systems, it was clear that most of the proposed solutions rely on the mediator/wrapper paradigm, or modifications of it, which tend to be widely used when there is the need to access data and information coming from several external data sources that are both diverse and heterogeneous. This heterogeneity is highlighted by the fact that they not only provide different means and protocols for connectivity, exposing different APIs and making use of different types of drivers required to establish a connection, but they are also compatible with different query scripting languages. To make things worse, each of the data store may have its own data model to store data and return results (i.e. a relational model in traditional relational database management systems and a JSON schema in document-based datastores), and according to the data source, data might be structured, semi-structured or totally unstructured. Due to this, it is very complex to implement a unique way to execute query statements and access integrated data spanned among different stores.

In order to facilitate the implementation of a polystore, the mediator/wrapper paradigm that aims to hide the implementations details of the data connectivity and data access at a lower level, is introduced providing generic interfaces of which the query engine can make use in order for the execution of the query to become transparent regardless the target datastore. Following this approach, a central building block in terms of a component diagram (i.e. it does not have to be a centralized component, rather that it can be implanted in a distributed manner) it is the *mediator* that is responsible of the orchestration of the execution of the query statement. In order to access data and retrieve the results, it makes use of the *wrappers*, which hide the complexity and the internal details of how to access data and retrieve the intermediate results. Wrappers themselves can also be implemented in a distributed manner, thus allowing for intra-operation parallelism, if we consider that the data retrieval from an external datastore is a single operation in the query execution tree, no matter the complexity of that operator.

Each *wrapper* is responsible to handle a specific type of a datastore. The polystore supports as many different types of external datastores given the wrapper implementations for those datastores are available and supported by the former. The *wrapper* provides the following functionalities:

- It retrieves a subquery in a predefined format (i.e. it can be a standard SQL statement, a native query, or an agreed model of the subquery in a structured way)
- If not a native query, it translates the input to a query that is compatible and equivalent to the supported dialect and query interface that the target datastore accepts
- It can establish a connection to the target datastore. Internally, it makes use of a *data connector* subcomponent that holds the specific driver to the target datastore, and has implemented all the details on how to open, maintain, close a connection, send statements and receive results over that connection according to the protocol with which the datastore is compatible. The internal details are irrelevant from the wrapper's point of view, so it might use a pool of connections and do whatever type of optimization it considers necessary, as long as it does not violate the agreed protocol with the datastore.
- It is able to retrieve data over that connection, iterate over the return result and close the connection (via the interfaces exposed by the *data connector*) once the data has been returned.

- It is responsible for transforming the data coming from the target datastore into the common data model that the whole solution is using, and for returning the results back. The retrieved data are being transformed to tuples yield to a *named table expression* that the core of the Integrated Polyglot makes use of.
- (Optional) It may implement a cost estimation model. A requirement for this to happen is that ability of the wrapper to request statistics from the target datastore, which is possible only if the latter is able to expose this kind of information. This includes types of indexes, number of rows per table, information about the histogram of the tuples over the indexes, number of hits, history of submitted query statements, etc.
- (Optional) It may be able to transform the incoming query, explore different query execution plans by applying various transformation rules, and re-write the query to an equivalent one that can improve the overall execution. In order for this to be possible, it is required that the *wrapper* can perform a cost estimation model that would be taken into account by its query optimizer when investigating the cost of the explored query execution plans.

On the other hand, the *mediator,* as the central building block in this architectural paradigm orchestrates the whole query execution of the integrated statement and makes use of the available *wrappers.* It receives as an input the submitted integrated statement that consists of one or more subqueries, each of those targeting a specific external datastore. Internally, it contains a query compiler that is capable of transforming the script into a structured query plan that can be parsed by the query planner. The latter explores and suggests alternative and equivalent query plans for execution. For instance, in case a subquery is a standard SQL statement, it can be decided to push some *filter* operations down to the external datastore engine in order to reduce the amount of data that will be transmitted across the network and the amount of memory that will be required to process the intermediate result. In case that the external datastore can provide statistics that can be taken into account by the cost estimation model of the mediator, then further improvements can be decided on the preparation phase of the execution, as it was highlighted in the case of the *bind joins* in section 3.

At the end of the preparation phase, the query plan to be executed has been decided, which consists of the subqueries that need to be executed in the target datastores. The *mediator* then sends these queries into the corresponding wrapper and retrieve the results in the form of a *named table expression,* or a *virtual table.* To do so, it provides a common interface that all *wrappers* must implement, in order for the overall execution to be transparent from the *mediator* point of view. As a result, it uses this interface to submit the corresponding subquery and retrieve the result in the agreed format, hiding all the complexities for data connectivity and data access to the *wrapper,* following the *separation of concerns* concept that is of major importance in the development of system and software solutions. Finally, it merges the results and returns the result set back to the data scientist or application developer.

The bird-eye-view of this architecture principle that the INFINITECH Integrated Polyglot component follows can be depicted in Figure 1.

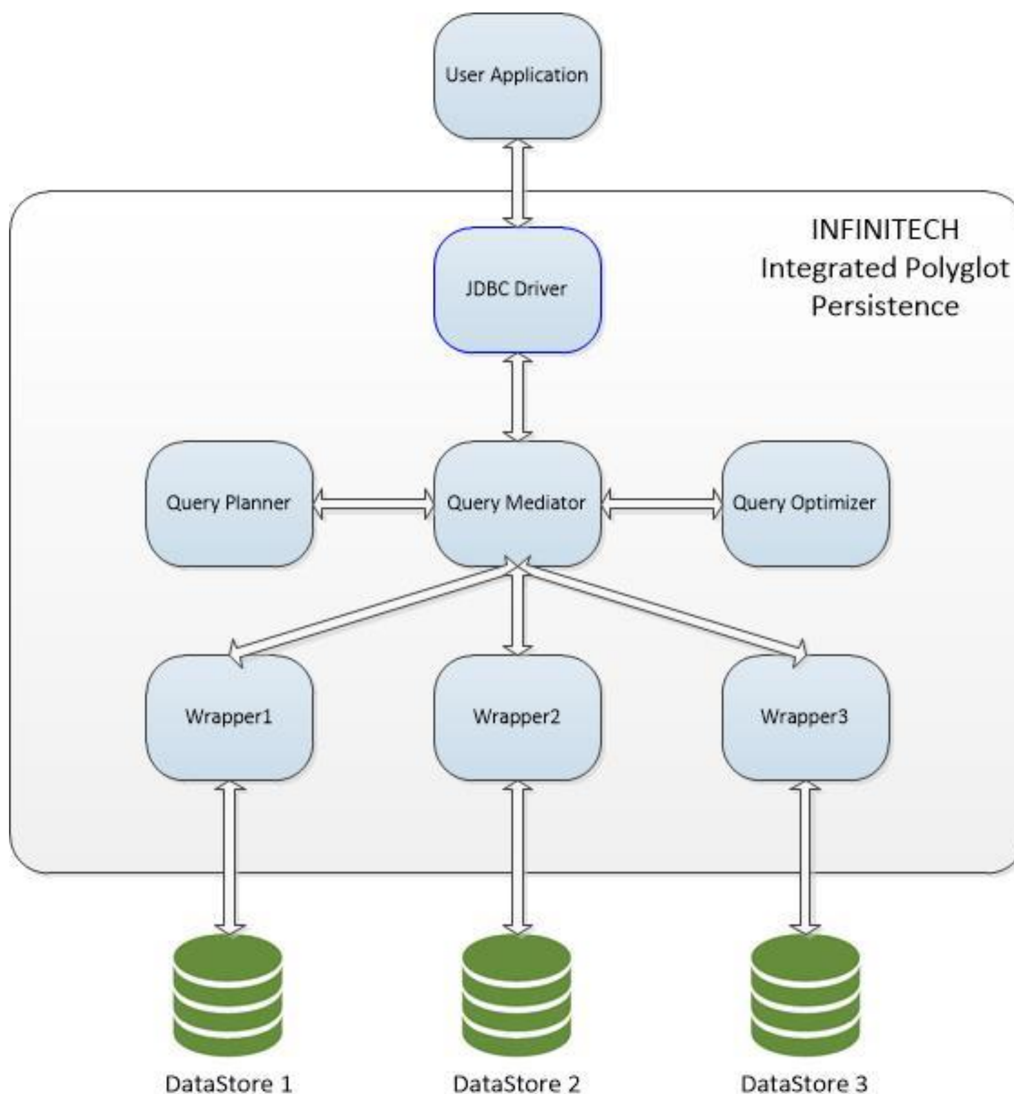

Figure 1: The Mediator/Wrapper paradigm in INFINITECH Polyglot Persistence

The data scientist or application developer submits an integrated query in the system for execution. This is being done in a seamless way, using a unique interface which is the standard JDBC. The integrated query arrives to the *mediator* component that starts the preparation phase. As explained previously, it compiles the query and transforms it in a structured way for further processing and it communicates with the *planner* that explores the space of possible query transformations into equivalent ones. Then it communicates with the *optimizer* which takes the responsibility to estimate the most cost effective one, taken into account its own cost estimation models (i.e. pushing down a *filter* and *projection* operations before the *selection* is always more efficient than applying those operations after the data has been retrieved from the datastore). When the *optimizer* decides about the plan to be executed, then the *mediator* sends the several subqueries down to the corresponding *wrappers* by invoking the common interface that each of the wrapper implements. At that point, it is totally transparent to the *mediator* how the data will be retrieved, as this logic is hidden by the *wrappers* themselves. They are responsible to establish the connection to the target datastore, send the query for execution, retrieve the results and return them back to the *mediator* in a form of a *virtual table.* The *latter* receives the intermediate results, applies the *join operation* and return the data in the form of a *ResultSet* back to the user via the JDBC connection.

According to the capabilities of each of the target datastores, the intermediate result might be returned as a whole where the execution of the subquery has been finished, or might return via batches. In the first case, the *mediator* has to wait for the execution to be completed before continuing on the execution of the operation that is placed in the upper node of the query tree. This is not efficient enough, as it will block the

overall execution waiting for this operation (the intermediate data retrieval through the wrapper) to be finished first. If the external datastore is capable of returning the results in batches, the overall execution can be much more accelerated, as the data pipeline that is being established by the *mediator* can be executed on the fly, as data are received, without the need to block the execution and have to maintain all data in memory. Taking into account that the *mediator* has to perform one of the most expensive operations, the *join*, this can be of great benefit: This will allow the execution of *merge joins* in case the join attributes are ordered over an index, and will also improve the execution of *nested-loop joins* and *hash joins,* as the nested loop will be sent to the right-hand operat the moment the data arrives from the left-hand, while the *hash* in the hash join can release resources much more easily.

It is important to highlight at this point the fact that the *mediator* along with the *planner* and *optimizer* are internal parts of the INFINITECH central data repository, with the appropriate extensions that are planned in order to take into account the additional requirements that the polyglot engine introduces. As a result, the INFINITECH Integrated Polyglot component does not have to re-create the *join* operations, but instead, it can rely on the core of the query engine of the data repository. As explained in more details in the corresponding deliverables of T3.1 “Framework for Seamless Data Management and HTAP” (D3.1, D3.2 and D3.3), its query engine provides massive parallelism processing and provides intra-query and intra-operation parallelism, making the execution of these types of operations much more efficient. The scope of this deliverable is to provide technical information on the polyglot extensions and not to give more insights on the query engine as a whole, which is part of T3.1.

4.2 Parallel Integrated Processing using the INFINITECH Common Query Language

As it has been analysed in D3.1, the distributed query engine of the INFINITECH central data repository is designed to be integrated with arbitrary data management clusters, which can store data in their natural format, without the need for pre-processing in order to be transformed to a compatible data schema, and can be retrieved in a parallel fashion by either executing declarative queries or by running specific scripts. The query engine supports a variety of diverse data management clusters, from distributed raw data files, parallel SQL database management systems, sharded NoSQL databases and parallel processing frameworks such as Apache Spark. To that sense, the query engine of the data repository of the platform can be integrated with the Integrated Polyglot component and thus transform it into a powerful *big data lake* polyglot engine that is capable of taking the full advantage of both expressive scripting and massive parallelism. Moreover, as it was described in the previous subsections that the integral query engine of the INFINITECH data repository supports the efficient execution of a variety of implementation of the *join* operator, and by providing intra-operator parallelism; these operations can be executed in a distributed manner, in parallel. As a result, joining data coming from native datasets being stored in external datastores along with the internal data tables of the repository itself can be applied exploiting the most suitable implementation that will exploit efficient parallelism. To highlight that fact, we will illustrate in the following examples how the execution of a parallel join operation across a relation table and the result of a JavaScript subquery to a document-based datastore like MongoDB is being done, along with the join of a data table with an MFR query, using the INFINITECH Common Query Language and the Integrated Polyglot Persistence.

Let’s assume that we have a table (collection in MongoDB terminology) called *orders* with the following data schema:

```
{order_id: 1, customer: "ACME", status: "O",
 items: [
  {type: "book", title: "Book1", author: "A.Z.",
          keywords: ["data", "query", "cloud"]},
  {type: "phone", brand: "Samsung", os: "Android"}
] }, ...
```

This is an example of semi-structured data as each of the records in the *orders,* contains an array of *items*, whose attributes differ according to the type. We need to return the title and the author of all books in the *orders* by a given customer. This can be expressed with a *flatMap* operation in JavaScript and a MongoDB *find()* operation. This can be depicted in the following code listing, where this expression is being wrapped into a *named table expression* of the INFINITECH Common Query Language:

```
BookOrders(title string, author string,
           keywords string[])@mongo =
{*
  return db.orders.find({customer: "ACME"})
  .flatMap( function(v) {
    var r = [];
    v.items.forEach( function(i){
      if (i.type == "book")
        r.push({title:i.title, author:i.author,
                keywords:i.keywords});
    } );
    return r; });
*}
```

We can see from the above code listing that we define a *virtual table* called *BookOrders* with the specified signature (title column as a String, author column as a String, etc.) which will be executed in the *@mongo* external datastore, using a native script (as indicated by the {* *} brackets) where we place the *flatMap* inside the *find()* MongodB operation. Furthermore, we need to join the result of this with a table named *authors* that is being stored inside the central data repository. The integrated statement according to the INFINITECH Common Query Language will be the following:

```
SELECT B.title, B.author, A.nationality
FROM BookOrders B, Authors A
WHERE B.author = A.name
```

This involves an *equity join* operation over a *scan* operation on the right-hand and a polyglot operation on the left. Let's assume for simplicity reasons that the query optimizer decides to use a *nested-loop join* operation. In what concerns the scan operation over the internal data table, this can be executed in parallel due to the distributed data management clustering that supports this. Regarding the polyglot operation, this is related to whether the wrapper can be executed in parallel or not. In any case, the *nested-loop join* operation also supports intra-operator parallelism. This means, that the data pipeline of the integrated query plan can be configured and once data is retrieved by the polyglot operation, then the *nested loop* implementation of the *join,* can use the hashcode of the join attribute and send the tuple to the related *worker* to execute the right-hand of the *join,* as explained in details in D3.1.

Additionally, we can have a more sophisticated data transformation logic needs to applied over the unstructured data before being able to be processed by relational operators. In the following example, we need to analyse the logs of a scientific forum in order to identify the top experts for particular keywords, assuming that the most influencing user for a given keyword is the one who mentions the keyword most frequently in their posts. We assume that the application keeps the log data in the non-tabular structure that is depicted below:

```
2014-12-13, http://..., alice, storage, cloud
2014-12-22, http://..., bob, cloud, virtual, app
2014-12-24, http://..., alice, cloud
```

There are text files where a single record corresponds to one post which contains a fixed number of fields about the post itself (timestamp, link to the post, and username in the example) followed by a variable number of fields storing the keywords mentioned in the post. The unstructured data needs to be transformed into the following tabular format:

```
KW        expert   frequency
cloud     alice    2
storage   alice    1
virtual   bob      1
app       bob      1
```

Such transformation requires the use of programming techniques like chaining map/reduce operations that should take place before the data is involved in relational operators. This can be expressed with the following MFR subquery with embedded Scala lambda functions to define custom transformation logic:

```
Experts(kw string, expert string)@spark = {*
  SCAN( TEXT, 'posts.txt', ',' )
  .MAP( tup=> (tup(2), tup.slice(3, tup.length)) )
  .FLAT_MAP( tup=> tup._2.map((_, tup._1)) )
  .MAP( TUPLE, 1 )
  .REDUCE( SUM )
  .MAP( KEY[0], (KEY[1], VALUE) )
  .REDUCE( (a, b) => if (b._2 > a._2) b else a )
  .MAP( KEY, VALUE[0] )
*}
```

Skipping the details of(?) the MFR query, we can see that we define a *named table expression* called *Experts,* which provides its signature and makes use of a native MFR query. We further join this table with the *BookOrders* that will be defined earlier and whose data are persistently stored in the MongoDB. We can write the integrated query in the following way:

```
SELECT B.title, B.author, E.kw, E.expert
FROM BookOrders B, Experts E
WHERE E.kw IN B.keywords
```

We illustrate in this example how the *bind join* supported by the Integrated Polyglot component can be used for optimal execution of this query. In this case, the bind join condition (which involves only the kw column) can be pushed down the MFR sequence as a FILTER operator. As per the MFR rewrite rules, this would take place immediately after the FLAT_MAP operator, thus reducing the amount of data to be processed by the expensive REDUCE operators. To build the bind join condition, the query engine flattens B.keywords and identifies the list of distinct values.

By processing such queries, the distributed query engine of the INFINITECH central repository can take advantage of the expressivity of each local scripting mechanism, enabled via the use of the Common Query Language of the platform, yet allowing for results of subqueries to be handled in parallel by the query engine itself and be involved in operators that utilize the intra-query parallelism. The query engine architecture is therefore extended by the implementation of the Integrated Polyglot Persistence to access in parallel shards of the external data store through the use of DataLake distributed wrappers that hide the complexity of the underlying data stores' query/scripting languages and encapsulate their interfaces under a common DataLake API to be interfaced by the query engine.

5 Implementation of the Integrated (Polyglot) Persistence

This section illustrates the initial implementation of the Integrated Polyglot Component, giving some high-level overview of how the code is organized along with code examples in the form of pseudo code. As it has been already mentioned, at this phase of the project, a wrapper that is capable of accessing data that is stored in a traditional relational database management system has already been implemented. The purpose of this section is to focus on the implementation related to the Integrated Polyglot component, and not to give more details on the overall query engine of the INFINITECH central repository. For that, a class diagram regarding how the code of the *wrapper* has been organized, giving analytical details of the functionalities provided by the code, along with a more detailed explanation provided with a pseudo code, that can be used as a roadmap for the implementation of the additional *wrappers* that need to be implemented in order to access code that is stored in other types of datastores.

5.1 Abstract Wrapper Implementation

In order to be able to implement wrappers for the corresponding external datastores, various interfaces have been defined that can be used by the *mediator* component, which is part of the query engine of the central data repository. As already explained, having generic interfaces makes it possible for the *mediator* to handle all different types of external datastores, in a unified manner. To test our design, we implemented a *wrapper* as a proof-of-concept, and we assume that it can be used to establish connections and access data in a traditional SQL-compatible relational datastore. We call our mock datastore as *rdb.* The class diagram is depicted in Figure 2.

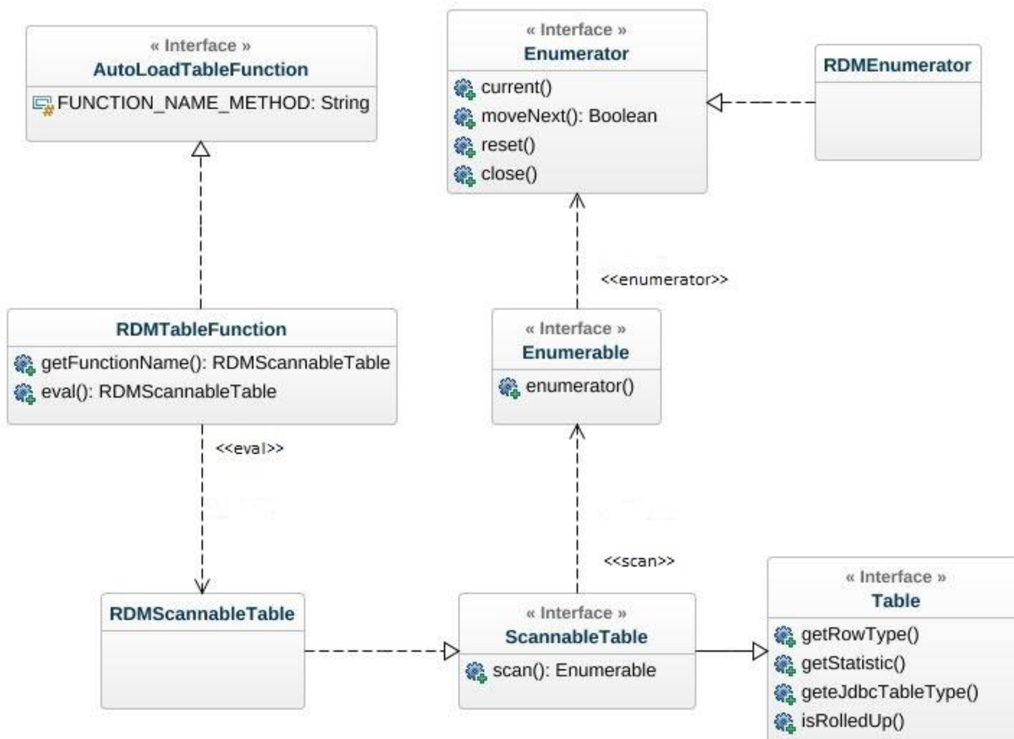

Figure 2: RDBWrapper class diagram

From the figure above, we can see that each wrapper must implement the *AutoLoadTableFunction.* As it has already been mentioned, a subquery is being submitted either in declarative language, or via a scripting

expression, and it formulates the *table function.* For our implementation of the *rdm* wrapper, the *RDMTableFunction* class implements this interface. Upon initialization of the query engine, this class is loaded via Java Reflection. It implements two methods: *getFunctionName* which returns the name of the wrapper, which in our case will be *'rdb'*, and the *eval* which returns a *ScannableTable* object, which implements the business logic that the wrapper is actually doing. To make things clearer, let's assume we have the following statement in the INFINITECH Common Query Language:

```
A(id int, x int)@RDB = (SELECT a.id, a.x FROM a)
```

This illustrates a subquery that will be executed in the *@RDB* wrapper. In that case, the query engine holds a map with all the wrappers, along with their names, so this code will be passed to the *RDMTableFunction* to be executed. The *eval* method expects two string arguments: the subquery to be executed and the signature of the result. In this example, the query will be the (SELECT a.id, a.x FROM a) statement, and the signature the (id int, x int).

The query engine invokes the *eval* method to get the results, and the latter returns an implementation of the *ScannableTable.* The latter extends the *Table* interface, and it will be used by the query engine as it consists of the polyglot operation that is part of the data pipeline established by connecting the various operations needed to retrieve data according to the tree of the query plan. It defines various methods that are relevant to the query engine itself for creating the list of equivalent query trees and for providing information to the query optimizer, along with information needed for the real-time execution of the operator. The functionality and rationale of these methods are beyond the scope of this deliverable, apart from the *scan.* This is used by the query engine to retrieve the *Enumerable* object that will return the row data transformed to the common model, as retrieved by the external datastore.

The *Enumerable* object returns an implementation of the *Enumerator* interface, which implements the methods that actually access data in the external datastore. As an *enumerator*, it defines methods for closing the object, which will imply to close the connection to the external database, check if there are more data to be returned, reset the pointer to the first tuple of the retrieved data, and to actually return the tuple, via the *current* tuple. In our implementation, the *RDMEnumerator* implements this logic, and the next subsection goes deeper into the details, with pseudo code snippets.

5.2 RDM Wrapper Implementation details

From the previous subsection, the main interfaces that the developer of a wrapper should provide to the INFINITECH Integrated Polyglot component in order to grant access to an external datastore, are the following:

- AutoLoadTableFunction
- RDMEnumerator

All other interfaces are being provided by the core of the query engine of the INFINITECH central repository and its polyglot extensions. Regarding the *AutoLoadTableFunction,* this interface defines two methods. The first one, returns the name of the wrapper so that the query engine can be informed that there is an additional polyglot operator to be taken into account and might be addressed by a subquery of an integrated statement. The following code illustrates its implementation:

```
public static String getFunctionName(){
        return "rdb";
}
```

We have named our wrapper as *rdb*, in the sense that it is targeting an external relational database management system, therefore, this term will be used in the integrated statement in order to drive the component to use this implementation when addressing subqueries with the *@rdb* indicator.

Moreover, the *eval* method returns an implementation of the *ScannableTable* interface, which extends others that are necessary from the query engine to manipulate these types of objects. A pseudo code snippet for implementing this method can be the following:

```
public ScannableTable eval(final String query, final String schemaDefinition) {
         //parse the schemaDefinition to grab the field names and types
        String [] fieldNames = getFieldNamesFromSignature(schemaDefinition);
        SqlTypeName [] fieldTypes = getFieldTypesFromSignature(schemaDefinition);

         return new ScannableTable() {
                @Override
                public Enumerable<Object[]> scan(DataContext dc) {
                        return new AbstractEnumerable<Object[]>() {
                               @Override
                               public Enumerator<Object[]> enumerator() {
                                       return new RDMEnumerator(query, fieldNames, fieldTypes);
                               }
                        };
                }

                @Override
                public RelDataType getRowType(RelDataTypeFactory relDataTypeFactory) {
                        return getRowType(relDataTypeFactory, fieldTypes);
                }

                @Override
                public Statistic getStatistic() {
                        return Statistics.UNKNOWN;
                }

                @Override
                public Schema.TableType getJdbcTableType() {
                        return Schema.TableType.TABLE;
                }

                @Override
                public boolean isRolledUp(String string) {
                        return false;
                }

                @Override
                public boolean rolledUpColumnValidInsideAgg(String string, SqlCall sc, SqlNode sn,
CalciteConnectionConfig ccc) {
                        return false;
                }
         };
 }
```

In this pseudo code, it is depicted that a new instance of the *ScannableTable* interface is being created and its methods are overridden by the implantation needed by this specific wrapper. It illustrates that at the beginning, the String containing the schema definition is being parsed because this information will be needed during the scanning phase, in order to transform the data to the common model that has been defined by the integrated statement. Then, the *scan* method creates a new *Enumerator* object, which is the *RDMEnumerator* which has been provided for this wrapper. It is important to be highlighted at that point that custom enumerator defines a constructor whose arguments must be the column names and types of the signature. By doing that, the instantiation object of this class will have all this information available in order to proceed for the proper transformation of the retrieved data.

Regarding the *RDMEnumerator* class itself, there are several methods that need to be implemented by the interface. An additional one with *private* visibility which manages to open and establish a connection with the target datastore, has been introduced.

```
private void init() {
        if(it is already opened) {
```

```
                return;
        }

        try {
                this.connection = createConnection(connection arguments)
                this.stametement = this.connection.createStatement();
                this.resulSet = this.stametement.executeQuery(query);
        } catch(SQLException | SAFFederatorException | CQEException ex) {
                throw new RuntimeException(ex);
        }

        setAlreadyOpen();
}
```

This method checks if the connection to the external datastore is already open, and if not, it establishes a connection, creates the statement object according to the JDBC standard, and executes the query, storing the result set that needs to be parsed later in order to retrieve and transform the data. It is worth to be mentioned that this pseudo code is illustrating a very basic implementation of a relational wrapper which executes the query and waits for the result. The execution of the query might be a blocking operation or not, depending on the type of the query and the implementation of the JDBC driver of the specific datastore. A more sophisticated approach would be to open a new thread that will be responsible to execute this query in parallel with the preparation of the query engine to settle the data pipeline and begin to request data. This thread could retrieve the data in parallel and feed the result in a *blockingqueue* that could be used as the pipeline of this thread and the thread opened by the query engine to execute these lines of code. As a result, the data will be available when the query engine will start fetching them, as the corresponding method would pick them from this *blockingqueue* that would have been already started to be filled with row data from the external datastore.

Regarding the *moveNext* method, it has to check whether or not there are more data to be retrieved from the external datastore. A code snippet could be the following:

```
@Override
public boolean moveNext() {
        init();
        return this.iterator.moveNext();
}
```

It firstly checks if the connection is already opened, and if not, it establishes the connectivity. As a result, the very first time that this method will be invoked, the wrapper will open the connection to the external datastore and will execute the query. As we are mocking a traditional relational database management system that provides a JDBC interface, this code relies on its implementation to check this. In the more sophisticated approach with data being retrieved in a parallel thread, this code will have to fetch data from the *blockingqueue* until an *END_FLAG* is received, that will signal the end of the dataset, and the code will be unblocked and return false.

The *current* method returns the current tuple of the retrieved data. A pseudo code could be the following:

```
@Override
public Object[] current() {
        return this.iterator.current();
}
```

As the *RDMEnumerator* mocks a JDBC compliant datastore, this code relies on its implementation to return the current data. It is important to be mentioned here that as both stores, the external one and the polyglot, are relational data stores, there is no need for data transformation. However, in case of a document-based datastore or a Hadoop Data lake, the code snippet would have to transform the raw data into the corresponding format, before sending back the array of objects to the query engine and the *mediator.*

The *reset* method has to set the pointer in the first row of the retrieved dataset. In our case, this operation is not supported by the JDBC standard, so the code snippet must not allow this, and indeed, it throws a runtime exception as follows:

```
@Override
public void reset() {
        throw new UnsupportedOperationException("Reset operation is not supported by the " +
this.getClass().getSimpleName() + ".");
}
```

Finally, when all data has been retrieved by the query engine, the latter closes the *enumerator*. The implementation of the corresponding method must close all open connections and release all resources that have been reserved for the execution of this subquery.

```
@Override
public void close() {
        try {
                if((this.resultSet!=null)&&(!this. resultSet.isClosed())) {
                        this. resultSet.close();
                }
        } catch(IOException ex) {
                Log.warn("Could not close iterator {}. {}", this.
resultSet.getClass().getSimpleName(), ex);
        }
        try {
                if((this.statement!=null) && (!this.statement.isClosed())) {
                        this.statement.close();
                }
        } catch(SQLException ex) {
                Log.warn("Could not close statement for retrieving tuple. {}", ex);
        }
        try {
                if((this.connection!=null) && (!this. connection.isClosed())) {
                        this. connection.close();
                }
        } catch(SQLException ex) {
                Log.warn("Could not close connection to datastore. {}", ex);
        }
}
```

6 Real-Time Data Warehouse

In this section, we describe a novel architecture that we propose. It is called **real-time data warehouse** and it can be used to overcome inherit technological barriers of modern approaches when having to deal with a vast amount of data that becomes obsolete after a certain period of time and is widely characterized as historical. This data needs to be moved to a data warehouse to offload the operational datastores that have different types of query processing capabilities and the focus is mainly to ensure database transactions. Our implementation is based on the polyglot engine that has been developed during the first phase of task T3.2 ("Polyglot Persistence over BigData, IoT and Open Data Sources") and has been presented in the previous sections. Here, at first the motivation behind our solution is described, then the basic principles, requirements and architectural components are also described and finally, the details on how the algorithms behind the query processing are implemented and how this approach ensures data consistency in terms of database transactions.

6.1 Motivation

As modern enterprises in finance and insurance sectors currently have to deal with a vast amount of data, the majority of this information eventually becomes obsolete. The result is that data kept by these enterprises can be categorized as current and historical. The data management systems are continuously being ingested with fresh data coming from various sources. Examples can be found in various domains. For instance, IoT sensors installed in a vehicle are continuously feeding with such information an insurance company, which is the scenario of pilot#11 ("Personalized insurance products based on IoT connected vehicles"). Moreover, finance currencies are being ingested in a per-second frequency to the management system of finance companies that consult about risk assessment, which is the scenario of pilot#2 ("Real-time risk assessment in Investment Banking"). What is more, online finance transactions are being stored, so that can be later used in order to detect anti-money laundering activities and to avoid financial crime. This is the scenario of pilot#7 ("Avoiding Financial Crimes") and pilot#8 ("Platform for AML supervision (PAMLS)").

In all these scenarios, there is the need to deal with current data that are useful for real time detection or other operations that require the assurance of data consistency, but they also require to effectively manage the vast amount of data that has become obsolete and can now be considered as historical. The importance of keeping historical data is evident, as they feed AI algorithms that can be combined with current data. In order to cope with both current and historical data, two common design patterns are widely used nowadays.

The first pattern is commonly referred as *lambda architecture*, which combines techniques from batch processing with data streaming to be able to process data in a real-time manner. The lambda architecture is motivated by the lack of scalability of operational SQL databases, which are used to store current data. The architecture consists of three layers:

- <u>Batch layer</u>: It is based on append only storage, typically a data lake, such as the ones based on HDFS. Then, it relies on MapReduce for processing new batches of data in the forms of files. This batch layer provides a view in a read-only database. Depending on the problem being solved, the output might need to fully re-compute all the data to be accurate. After each iteration, a new view of the current data is provided. This approach is quite inefficient but solves a scalability problem that used to have no solution, the processing of tweets in Twitter.
- <u>Speed layer</u>: This layer is based on data streaming. In the original system at Twitter, it was accomplished by the Storm data streaming engine. It basically processes new data to complement

 © INFINITECH Consortium

the batch view with the most recent data. This layer does not aim accuracy, but instead aims at providing more recent data to the global view achieved with the architecture.

- Serving layer: The serving layer processes the queries over the views provided by both the batch and speed layers. Batch views are indexed to be able to answer queries with low response times and combines them with the real-time view to provide the answer to the query, combining both real-time data and historical data. This layer typically uses some key-value data store to implement the indexes over the batch views.

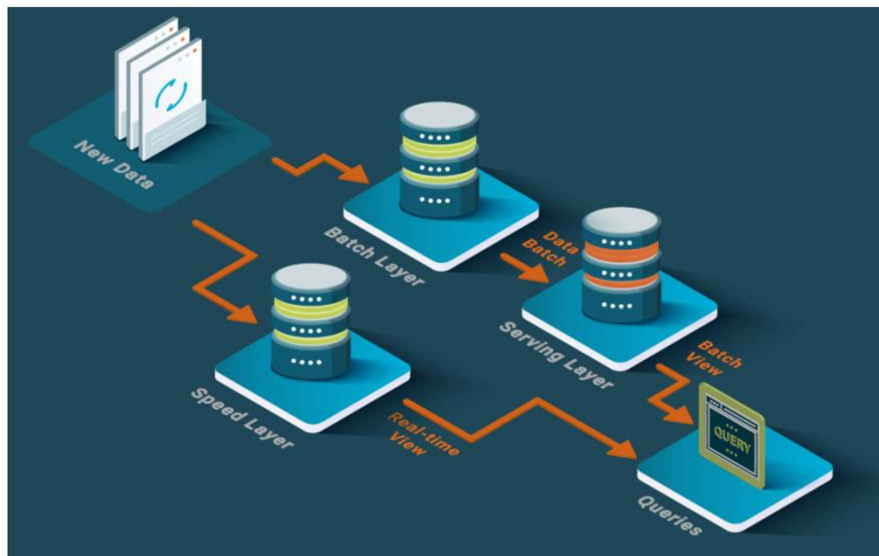

Figure 3: Lambda Architecture

The main shortcoming of the lambda architecture is its complexity and the need to have totally different code bases for each layer that have to be coordinated to be fully in sync. The maintenance of the platform is very hard since debugging implies understanding the different layers with totally different natures, technologies and approaches.

Other more traditional architectures are based on combining an operational database with a data warehouse. The operational database deals with more recent data while the data warehouse deals with historical data. In this architecture, queries can only see either the recent data or historical data, but not a combination of both as it was done in the lambda architecture. In this architecture there is a periodic process that copies data from the operational database into the data warehouse. This periodic process has to be performed very carefully since it can hamper the quality of service of the operational database. This periodic process is most of the time achieved by ETL tools. This process is typically performed over the weekends in businesses where their main workload comes during weekdays.

Figure 4: Current-Historical Data Splitting Architectural Pattern

Another problem with this architecture is that the data warehouse typically cannot be queried while it is being loaded, at least the tables that are being loaded. This forces to split the time of the data warehouse into loading and processing. When the loading process is daily, finally the day is split into loading and processing. The processing time consumes a fraction of hours of the day that depends on the analytical queries that have to be answered daily. It leaves a window of time for loading data that is the remaining hours of the day. At some point, data warehouses cannot ingest more data because the loading window is exhausted. We call this architectural pattern **current-historical data splitting**.

In this pattern, data is split between an operational database and a data warehouse or a data lake. The current data is kept on the operational database and historic data in the data warehouse or data lake. However, queries across all the data are not supported with this architectural pattern. With the achievements of this task T3.2 ("Polyglot Persistance over BigData, IoT and Open Data Sources"), a new pattern, called **Real-Time Data Warehousing**, will be used to solve this problem. This pattern will be solved by a new innovation that will be introduced in INFINISTORE, namely, the ability to split analytical queries over the operational datastore and an external data warehouse. Basically, it will copy older fragments of data into the data warehouse periodically. INFINISTORE will keep the recent data and some of the more recent historical data. The data warehouse will keep only historical data. Queries over recent data will be solved by INFINISTORE, and queries over historical data will be solved by the data warehouse. Queries across both kinds of data will be solved using a federated query approach leveraging the polyglot capabilities to query across different databases and innovative techniques for join optimization. In this way, the bulk of the historical data query is performed by the data warehouse, while the rest of the query is performed by INFINISTORE. This approach enables to deliver real-time queries over both recent and historical, data giving a 360° view of the data.

6.2 Design Principles and Requirements

Our approach has been designed to allow datasets to be split over the operational datastore and a data warehouse, over a specific column or a group of columns. This means that we are able to split on one or more data tables that are part of the target datasets, using a column (or a group of columns) that belong to these specific data tables. An import requirement is that the column (or the group of columns) must contain values that are monotonically incremental. Such an example might be columns containing auto-incremented values that are part of a primary key, timestamps, etc. This requirement might seem as an important constraint, but in reality, this covers the majority (if not all) of the use cases. The reader must

remember that we split a dataset that is becoming historical, leaving the current part in the operational datastore, while moving the historical part to the data warehouse. A dataset becomes historical over the time, and as a result, it will always contain a field that keeps the current timestamp at the time when the data row was added, or a field that is part of the primary key of the data table on which the row has been added. Getting back to our pilots that have been used as an example behind the motivation of our solution, all information contain a timestamp field: from the IoT datum coming from the deployed sensor, which provides information at a specific point in time, to the finance currency or finance transaction that takes place in a specific point in time.

Another requirement is that the historical dataset should not allow updates or other data modification operations. However, this is the same requirement and constraint as the other architectural designs have: when moving data to a data warehouse, data is being dropped from the operational datastore and is moved to the data lake or warehouse. In those data management systems, updates are not feasible. We keep data that can be modified in the operational datastore of INFINISTORE, and we move the data after a point in time when the data can be considered as obsolete. The difference between our approach and the **current-historical data splitting** is that using the **Real-Time Data Warehousing**, query processing can be done live, getting into consideration both its current and historical part, while data is being moved and migrated from one data store to the other, ensuring database transactions and data consistency at the same time. This is due to the transactional engine of the INFINISTORE itself. More information will be given in the next subsections.

Our implementation consists of the following main architectural pillars:

- <u>The INFINISTORE</u>: This is the operational datastore of our solution. It can be continuously ingested with data at very high rates, exploiting its HTAP capabilities and Kafka connectors implemented under the scope of task T3.1 ("Framework for Seamless Data Management and HTAP"). It can hold the current part if the target dataset provides efficient query processing via its parallel OLAP engine, currently being implemented under the scope of the same task. What is also important to be mentioned, is that by exploiting its HTAP capabilities, it is now feasible to move data (which will imply the execution of an analytical query that involves a *scan* of a data table) while leaving the support for operational workload unaffected. This means, that we can move data without any downtime, which is the technological main constraint of the **current-historical data splitting** architecture.
- <u>The data warehouse</u>: This is the data management system that will store the historical part of the dataset. Any kind of a data warehouse or data lake can be used as this part of our integrated solution. The only constraint at the point when this report was written, is for the data warehouse to expose a JDBC interface, in order to allow that type of data connectivity from the polyglot engine implemented in this task. In case that other types of data connections are supported, then the polyglot engine needs to be extended by implementing the corresponding *wrapper,* as explained in more detail in section 5.
- <u>The data mover</u>: This component is responsible for moving data from INFINISTORE to the data warehouse. As it has already been mentioned, a data table is split according to a column. Let's take the example of a timestamp. The data mover, periodically and according to the configuration put by the database administrator, requests a *data slice* from INFINISTORE to be stored to the data warehouse. It executes an analytical query on the INFINISTORE, on the specific timestamp, and when it retrieves the data, it stores it to the data warehouse. The transactional engine of INFINISTORE ensures that the result of a query execution to our integrated solution will always be equivalent as the dataset is stored in a single datastore, even if the data is being moved concurrently. In case of failures during the data migration process, a *re-do* pattern will be applied

that ensures that data will be eventually moved, stored in the data warehouse and dropped in the INFINISTORE.

- The query federator: This component is responsible of executing a query over a dataset that has been split between the two datastores, returning an equivalent result as the datastore was kept in a single data management system, ensuring data consistency on the same time while data is being concurrently moved. It uses the polyglot engine that was developed during the first phase of task T3.2 (“Polyglot Persistance over BigData, IoT and Open Data Sources”). As the polyglot engine extends the query engine of INFINISTORE, federated queries targeting datasets split between the two datastores can be submitted via a standard JDBC interface. The query federator implements the logic for receiving the two parts of the dataset, taken into account that they have been split over a column (or group of columns).

In the following subsections, we will provide detailed information on how the query federator executes the common SQL operators, and how our approach ensures data consistency in terms of database transactions while data is being moved from the operational datastore, the INFINISTORE, to the data warehouse.

6.3 The Query Federator

As was mentioned in the previous subsection, this component is responsible for executing a submitted query over a dataset that has been split between two different datastores. It is assumed that the data tables have been split over a column that contains a monotony incremental timestamp, without our assumptions to lose the general applicability of our solution.

The query federator supports all standard SQL operations. In the following part of this section, we will provide details about the implementation of the most standard operations. It is important to mention that the query federator makes use of the polyglot engine of the INFINISTORE and in fact, it consists of a specific implementation of a *wrapper*, following the *mediator-wrapper* paradigm introduced and explained in section 4, based on the code presented in section 5.

Full scan

In case of full scans, the query federator opens two data connections to both INFINISTORE and the target data warehouse. By firstly retrieving the current value of the timestamp that has been split, and making use of the compiler of the CloudMdSQL that has been mentioned in section 3, it translates the query to a scan with a filter condition over the timestamp, with the value of the current timestamp that has just been retrieved. As a result, it will only retrieve the logical dataset that is part of the *historical* from the data warehouse, and the logical dataset of the *current* from INFINISTORE, no matter if the physical data co-exists due to a data movement process that is being performed at the same time. Data is being returned back to the upper layers of the query tree, as it is being received from the two datastores, unordered as the underlying executions is happening in parallel.

Scan with conditions

In case of conditions, the query federator still opens two data connections to both INFINISTORE and the target data warehouse and executes the same algorithm as in the case of a full scan. The difference is that the timestamp condition is being added as an additional *AND* condition over the original one. Data is being returned back to the upper layers of the query tree unordered, as it is being received from the two datastores, since the underlying executions are happening in parallel.

Projections

In case of projections, they are being pushed down to the two datastores, so that we can minimize the amount of data that is being sent from both to the query federator.

Order by

In case that such an operator is enforced by the query federator (this implies that it has been pushed down for execution to this layer), then the latter still opens two data connections to both INFINISTORE and the target data warehouse and executes the same algorithm. The difference is that data is not being returned concurrently, rather than the two iterators are checking which value is greater (or smaller) and return back accordingly. It is evident that pushing down to the query federator such an operator, makes sense only when the ordering is over an index.

Limit

Whatever the overall submitted query, the query federator executes the corresponding algorithm, but keeps internally a counter to keep track of the number of rows that have been returned. When the counter reaches the value of the *limit* operation, it stops.

Aggregations

Firstly, it is important to highlight that the aggregation operators are the following: min, max, sum, count and avg. When an operation of this group is received for execution by the query federator, the latter pushes down such operations to both datastores and merges the intermediate results. This is feasible as all these operators can be executed in a distributed manner: The min of the overall dataset is the minimum of the two intermediate minimums of the two-split dataset. The max is the overall max of the two intermediate maximums, the sum is the overall sum of the two intermediate summaries, while the overall count is the sum of the two intermediate-count of the historical and current datasets. The only tricky point here is the avg, as the overall avg is not the avg of the two intermediates, so this operator cannot be executed that way. However, the overall avg is the result of the overall sum divided by the overall count. Those two operators can be executed in a distributed manner and as a result, the query federator transforms internally the avg to such the sum / count.

Group by

The group by operation is used when there are corresponding aggregations. As a result, the query federator executes such operations as above. The difference is that instead of returning the overall result, it keeps in an internal cache the intermediate results: the key of the cache is the byte concatenation of the columns involved in the *group by* clause, and the value is the intermediate values. When both intermediate values that are associated with a specific key have been retrieved, the query federator calculates the overall value for that key and returns the result. When both iterators related with the two database connections finish, then it returns all remain values, as it might have been the case that a concatenate key exists only in one of the two logical datasets.

Joins

Here, the general case when both of the tables involved in a join operation have been split between the INFINISTORE and the data warehouse is going to be used. Then, the overall join is the union of four intermediate ones: one where data from both data tables is stored in the INFNISTORE, one where data from both data tables is stored into the data warehouse, and two additional ones that the data exists in both. As a result, the overall join is the algebraic production of those four, which are not overlapping and can be executed in parallel. For the local joins, they are being pushed down to the corresponding datastores. For the other two, we first execute a *scan* on INFINISTORE, to grab the data that participates in the join, and then the query federator retrieves the list of values that are involved in the equity of the joint

operation. Having that list, it transforms the operation into an *in* clause that is being sent to the data warehouse. In that case, the latter executes its part locally. Then, the query federator receives this data and merges it with the one that has been already received from the INFINISTORE. This technique has been already described in more details in section 3, and is widely used as *bind join*. Its advantage is that only the minimum amount of data that is required for the execution of the join is transferred. As a result, we minimize data transfer among the two datastores and the query federator. More information will be given in the next version of this report.

6.4 Ensuring data consistency while concurrently moving data across the datastores

In this subsection, we will justify how the implementation can ensure data consistency in terms of data base transactions: this means that the result of the execution of concurrent transactions while data is being moved from one store to the other, is equivalent to the result of the execution of the same query of over the same logical dataset that is being stored physically in a single datastore.

Let's assume there is a data table, split between the INFINISTORE and a data warehouse. It is also assumed that the split has been done over a timestamp column, whose current value is 2021-01-01. Without losing the general applicability of the proposed implementation, it is finally assumed that a *full scan* is performed over this data table as depicted in Figure 5.

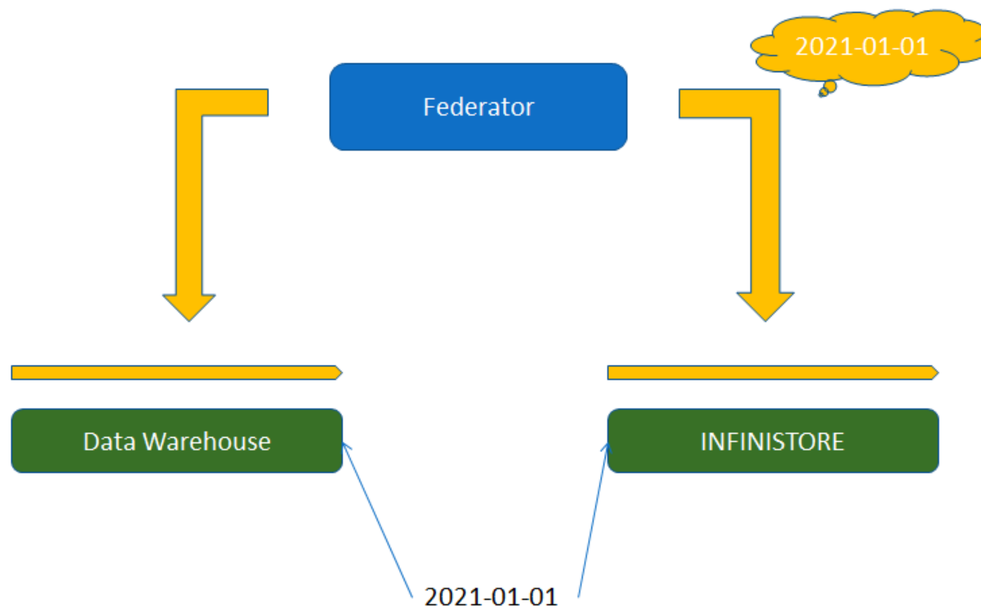

Figure 5: Simple full scan

Here it is displayed the dataset split in 2021-01-01, and the *query federator* submits the query to be executed in both. According to the previous subsection, it will take the current split timestamp, whose value is 2021-01-01 and will transform the query accordingly. With the yellow line, the part of the dataset that will be scanned from both tables is displayed.

Now, it is assumed that data will be moved and a *data slice* will have to be migrated from the INIFNISTORE to the data warehouse, as depicted in Figure 6.

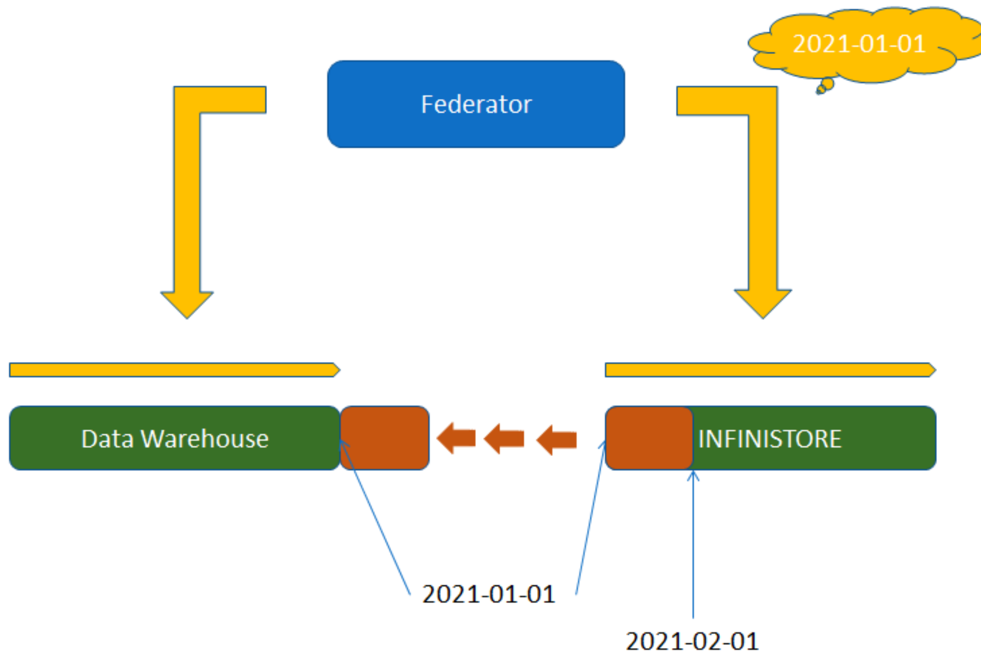

Figure 6: Moving data from INFINITORE to the data warehouse

The new split point will be now 2021-02-01 and the corresponding *data slice*, depicted with orange colour is being now moved to the data warehouse. It is also displayed that the *data slice* remains in INFINISTORE as this action takes place. This means that data now co-exists in both stores. However, the concurrent read operation marked in yellow, will only read the corresponding logical database, according to the initial timestamp that has been added by the query federator. It will never read any other data items that might be concurrently added, neither will it miss some data items, as the *data slice* still exists in the INFINISTORE.

At this point, the scenario can be made even more interesting. Data has been moved, but has not been dropped yet from the ININISTORE, as the read transaction is being executed. In addition, another read transaction starts that needs to perform the same *full scan.* This transaction is depicted in Figure 7 indicated in orange. As the new split timestamp is now 2021-02-01, the query federator will transform this operation accordingly by applying the filter condition with this value. From Figure 7 we can see, represented with the orange lines, the part of the physical dataset that will be accessed in both stores. Even if the yellow and orange transactions are accessing different physical datasets, the logical dataset is the same and the result of both operations is equivalent.

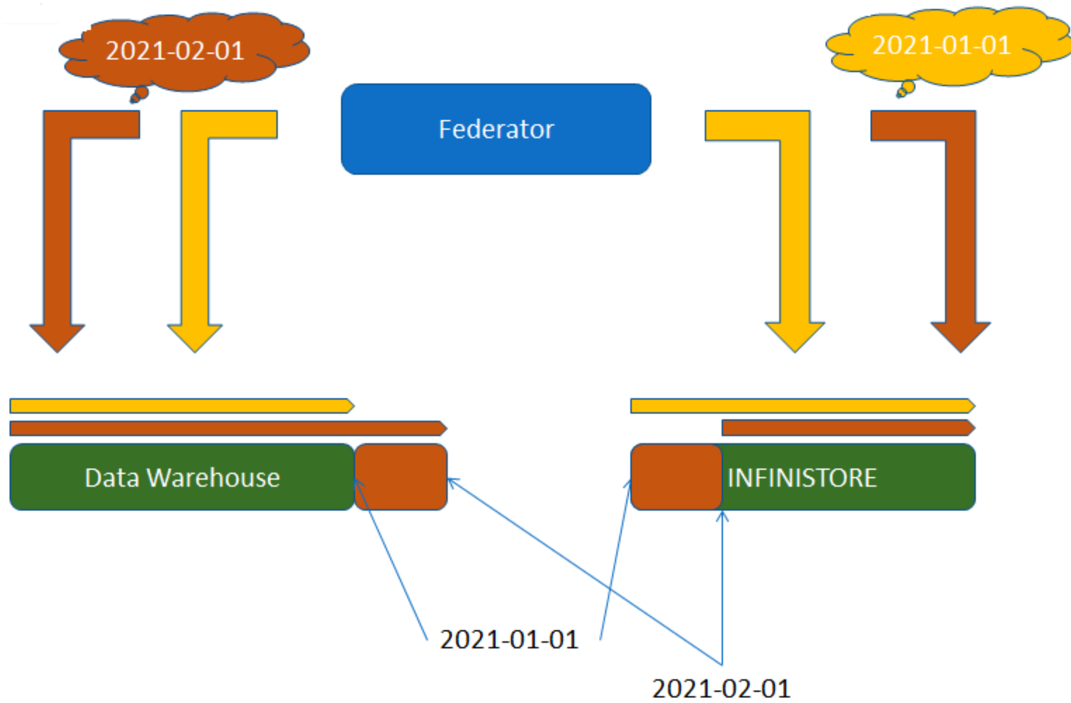

Figure 7: Second read transaction while data is being moved

Eventually, the yellow read transaction finishes after scanning the full logical data table that is split between the two datastores. At this point, as depicted in Figure 8, data still co-exists in both stores, however now there is only one read operation that access the logical dataset.

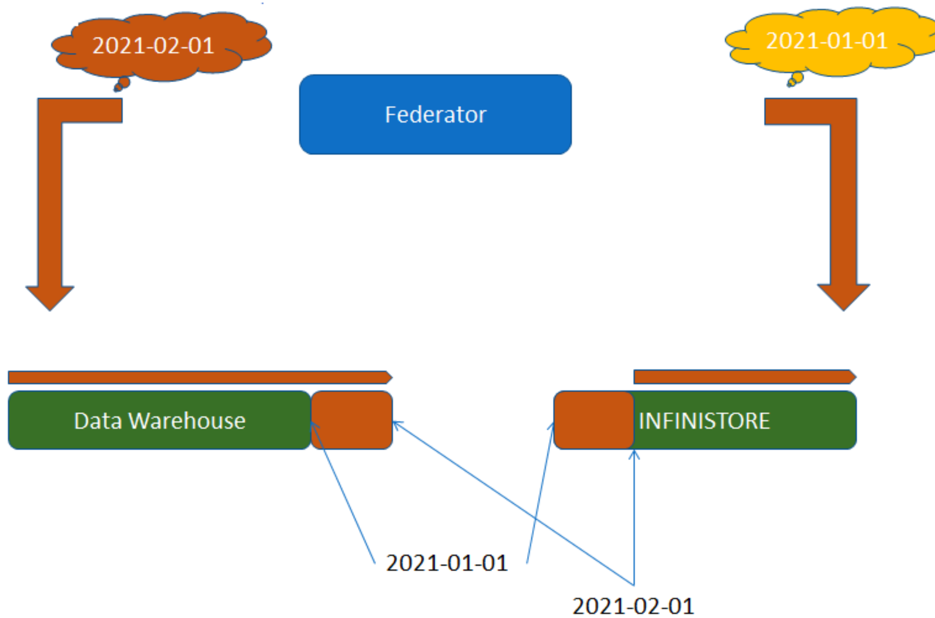

Figure 8: First read operation finishes

Finally, the query federator identifies that there is no pending or on-going transaction that requires access to the *data slice* in the INFINISTORE side, neither there will be any future transactions with a similar need. Future transactions will always take a value equal or greater than 2021-02-01, so this *data slice* that co-exists, can be now safely dropped from the INFINISTORE, as depicted in Figure 9.

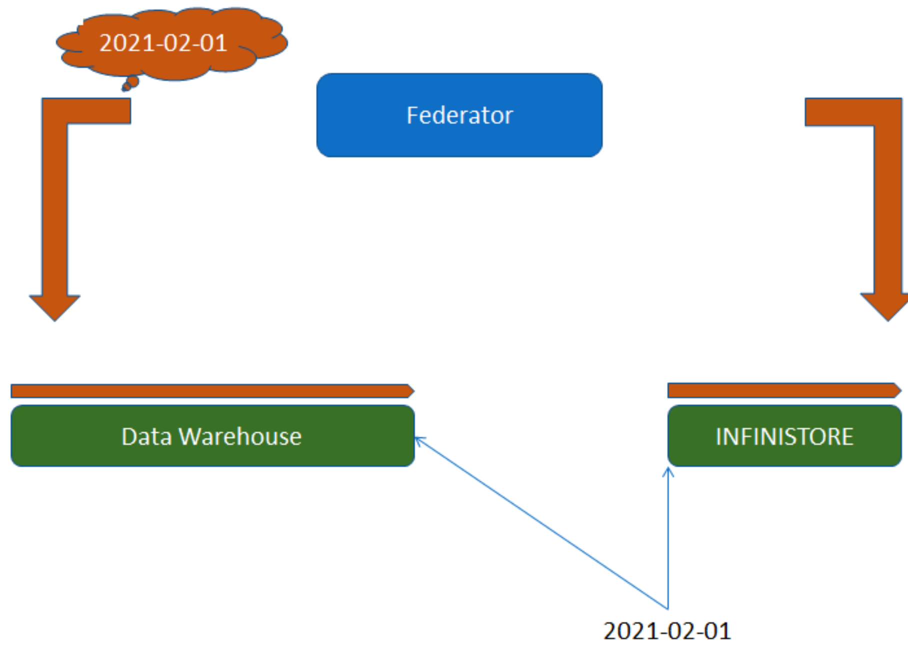

Figure 9: Dropping the data slice

7 Conclusions

This document reported the work that has been done in the scope of the task T3.2 "Polyglot Persistence over BigData, IoT and Open Data Sources" whose goal is to provide a common and integrated way to access data that is stored in a structured, semi-structured on even unstructured fashion over a variety of heterogeneous data stores, in a unified manner.

Towards this direction, firstly a state-of-the-art analysis has been made on the topic of polystore management systems. This revealed the fact that there are two major categories of polystore systems: loosely-coupled and tightly-coupled ones, each one of those focusing either on the autonomy of the external datastores, or on the efficiency of performance when processing data into a common model, using massive parallelism processing. Apart from those, nowadays a hybrid approach is widely used and combines the benefits from both approaches. It became obvious that the INFINITECH Integrated Polyglot component will make use of the mediator/wrapper architectural paradigm, widely used by the majority of the examined solutions, while it will need to provide even greater expressivity in order not to ignore the unique characteristics of the target databases, as most of the proposed solutions do.

The result of this analysis is the definition of the INFINITECH Common Query Language. This deliverable presents the basic principles of this language, which makes use of SQL language, but additionally gives the possibility to write native queries compatible with the target datastores, in a declarative way or via scripting expressions. An integrated statement written in the INFINITECH language consists of several subqueries that are targeting heterogeneous datastores. The common language abstracts this heterogeneity, while also giving the ability to exploit the datastore's unique characteristics. Moreover, the importance of the *bind join* and the way this is supported by the common language has been presented, along with an example on how we can use the latter to access Hadoop data lakes with MFR functions.

As the basis for the Integrated Polyglot component has been defined, being the INFINITECH Common query language, the general architecture design of this component has been presented. The component diagram highlighted how polyglot extensions are being designed in order to be incorporated with the INFINITECH central data repository, while the way this integration allows for the parallel execution of integrated queries across different datastores was widely presented.

After having the overall design of this component, we progressed with the implementation of a *wrapper* that can be used as the first polyglot extension that access data from an external relational database management system. More technical details have been provided in order to give the system developers the overview of the design with a class diagram, along with the code snippets to highlight what will need to be further implemented for the remaining of the wrappers. This can be used as a guideline for further development.

Finally, and based on the implementation of the polyglot query engine that was described in the previous sections, we implemented in this phase of the project what we call **real time data warehousing**, whose purpose is to provide a unified and seamless framework for data analytics over a logical dataset that has been physically split between an operational datastore and a data warehouse.

To conclude, the progress of task T3.2 is in plan with what has initially been planned, as already, the analysis of the competition has been fulfilled, the definition of the INFINITECH Query Language has been delivered, the design of the overall component has been published, while there is already an implementation that can be used as a reference for the development of the remaining of the wrappers. It is worth to be mentioned that this is the second report of the work to be done in the scope of T3.2, and there will be an additional iteration, where an extensive evaluation of the polyglot query processing framework will be provided.

Appendix A: Literature

[1] Z. Minpeng, R. Tore, "Querying combined cloud-based and relational databases", in Int. Conf. on Cloud and Service Computing (CSC), pp. 330-335 (2011)
[2] T. Özsu, P. Valduriez, Principles of Distributed Database Systems, 4th ed. Springer, 700 pages (2020)
[3] A. Simitsis, K. Wilkinson, M. Castellanos, U. Dayal, "Optimizing analytic data flows for multiple execution engines", in ACM SIGMOD, pp. 829-840 (2012)
[4] K. W. Ong, Y. Papakonstantinou, and R. Vernoux, "The SQL++ semi-structured data model and query language: a capabilities survey of SQL-on-Hadoop, NoSQL and NewSQL databases", CoRR, abs/1405.3631 (2014)
[5] J. Duggan, A. J. Elmore, M. Stonebraker, M. Balazinska, B. Howe, J. Kepner, S. Madden, D. Maier, T. Mattson, S. Zdonik, "The BigDAWG polystore system", SIGMOD Record, vol. 44, no. 2, pp. 11-16 (2015)
[6] V. Gadepally, P. Chen, J. Duggan, A. J. Elmore, B. Haynes, J. Kepner, S. Madden, T. Mattson, M. Stonebraker, "The BigDawg polystore system and architecture", in IEEE High Performance Extreme Computing Conference (HPEC), pp. 1-6 (2016)
[7] H. Hacigümüs, J. Sankaranarayanan, J. Tatemura, J. LeFevre, N. Polyzotis, "Odyssey: a multi-store system for evolutionary analytics", PVLDB, vol. 6, pp. 1180-1181 (2013)
[8] J. LeFevre, J. Sankaranarayanan, H. Hacıgümüs, J. Tatemura, N. Polyzotis, M. Carey, "MISO: souping up big data query processing with a multistore system", in ACM SIGMOD, pp. 1591-1602 (2014)
[9] T. Yuanyuan, T. Zou, F. Özcan, R. Gonscalves, H. Pirahesh, "Joins for hybrid warehouses: exploiting massive parallelism in hadoop and enterprise data warehouses", in EDBT/ICDT Conf., pp. 373-384 (2015)
[10] D. DeWitt, A. Halverson, R. Nehme, S. Shankar, J. Aguilar-Saborit, A. Avanes, M. Flasza, J. Gramling, "Split query processing in Polybase", in ACM SIGMOD, pp. 1255-1266 (2013)
[11] A. Abouzeid, K. Badja-Pawlikowski, D. Abadi, A. Silberschatz, A. Rasin, "HadoopDB: an architectural hybrid of MapReduce and DBMS technologies for analytical workloads", PVLDB, vol. 2, pp. 922-933 (2009)
[12] K. Awada, M. Eltabakh, C. Tang, M. Al-Kateb, S. Nair, G. Au, "Cost Estimation Across Heterogeneous SQL-Based Big Data Infrastructures in Teradata IntelliSphere", in EDBT, pp. 534-545 (2020)
[13] M. Armbrust, R. Xin, C. Lian, Y. Huai, D. Liu, J. Bradley, X. Meng, T. Kaftan, M. Frank-lin, A. Ghodsi, M. Zaharia, "Spark SQL: relational data processing in Spark", in ACM SIGMOD, pp. 1383-1394 (2015)
[14] Presto – Distributed Query Engine for Big Data, https://prestodb.io/
[15] Apache Drill – Schema-free SQL Query Engine for Hadoop, NoSQL and Cloud Storage, https://drill.apache.org/
[16] J. Wang, T. Baker, M. Balazinska, D. Halperin, B. Haynes, B. Howe, D. Hutchison, S. Jain, R. Maas, P. Mehta, D. Moritz, B. Myers, J. Ortiz, D. Suciu, A. Whitaker, S. Xu, "The Myria big data management and analytics system and cloud service", in Conference on Innovative Data Systems Research (CIDR) (2017)
[17] Apache Impala, http://impala.apache.org/
[18] B. Kolev, P. Valduriez, C. Bondiombouy, R. Jiménez-Peris, R. Pau, J. Pereira, "CloudMd-sQL: querying heterogeneous cloud data stores with a common language", Distributed and Parallel Databases, vol. 34, pp. 463-503. Springer (2015)
[19] L. Haas, D. Kossmann, E. Wimmers, J. Yang. Optimizing Queries across Diverse Data Sources. Int. Conf. on Very Large Databases (VLDB), pp. 276-285 (1997)

 © INFINITECH Consortium