

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



D2.13 – INFINITECH Reference
Architecture - I

Lead Beneficiary	GFT
Task Reference	T2.7
Due Date	2020-06-30
Delivered Date	2020-07-17
Revision Number	3.0
Dissemination Level	Public (PU)
Type	Report (R)
Document Status	Release
Review Status	UNP, LXS, INNOV reviewed
Document Acceptance	WP Leader Accepted and Coordinator Accepted
EC Programme Officer	Pierre-Paul Sondag

HORIZON 2020 - ICT-11-2018



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

Contributing Partners

Partner Acronym	Role ¹	Name Surname ²
GFT	Lead Beneficiary	Ernesto Troiano, Alessandro Amicone, Vittorio Monferrino
INNOV	Beneficiary	John Soldatos, Nikos Kapsoulis, Antonis Litke,
FTS, DWF	Beneficiary	Jürgen Neises
LXS	Beneficiary	Patricio Martínez, Javier López Moratalla, Sandra Ebro, Boyan Kolev, Luis Miguel Garcia
HPE	Beneficiary	Alessandro Mamelli, Domenico Costantino
All Proxy	Beneficiary	

Revision History

Version	Date	Partner(s)	Description
0.1	2020-01-14	GFT	ToC Version
0.15	2020-05-13	GFT, INNOV	Updated ToC with requested contributions
0.16	2020-06-10	GFT	Updated contributions
0.17	2020-06-15	GFT	Integration of contributions
0.18	2020-06-21	GFT	Completion of RA (initial version)
0.19	2020-06-21	INNOV	Contributions and Major Updates in Section 1(Methodology), Section 2 (Finance Industry Requirements), Section 4 (Review of Industry RA and Best Practices)
0.19	2020-06-22	UPRC	Contributions in Section 2 and Section 4
0.19	2020-06-23	LXS	Contributions in
0.22	2020-07-02		Contribution from Pilot #3 and 5.7.3 Data Acquisition
0.23	2020-07-02	INNOV	Revisions in the Sample Pipelines (Section 4) and Updates on the INFINITECH-RA position to the state of the art (Section 3)
0.24-0.26	2020-07-02	GFT	Various editions to integrate coming contributions
0.27-0.29	2020-07-09	INNOV	Major Quality Control of the Document
1.0	2020-07-14	GFT	First Version for Internal Review
1.1	2020-07-14	LXS, UNP	Review Updates
2.0, 2.1	2020-07-15	GFT, INNOV	Version for Quality Assurance; Review by the Quality Manager
3.0	2020-07-17	GFT final	Version for Submission

¹ Lead Beneficiary, Contributor, Internal Reviewer, Quality Assurance

² Can be left void

Executive Summary

The document presents the first part of work done within INFINITECH project to design and build a suitable Reference Architecture (RA) for the Smart, Autonomous and Personalized Services in the European Finance and Insurance Services Ecosystem.

The work on Reference Architecture has been conducted as teamwork that involved all Consortium beneficiaries. The work is also a continuous and incremental process and the present deliverable is a stable snapshot of the status of the RA at the date of submission. An updated version of this deliverable will be produced at month 18 of the project, taking into account all the improvements and due verifications over the period. Nonetheless, the actual RA fulfils the project milestone MS02 fully specifying the first RA of the INFINITECH platform.

The INFINITECH partners have selected a methodology to work on the RA, identifying it in the “4+1” architectural view model, which is presented in the document. The methodology is based on five different views, from which the structure of the system can be analysed (logical view, process view, development view, physical view and scenarios). Moreover, it will be demonstrated that all the functionalities of INFINITECH environment are properly covered by this model.

The State-of-the-Art survey underlines that some already existing Reference Architectures provide substantial input to INFINITECH, such as the pipelined and workflow approach to support the functionalities of the different Pilots and Use Cases of the project.

Relevant inputs to the task have been considered, in particular the input coming from use-cases considered in task T2.1 and a cross reference matrix.

Finally, a layered and high-level reference view and a detailed logical view of the RA are presented. Different layers have been identified (infrastructure, data management and protection, data processing and architecture, analytics, interface, and presentation/visualization). The layers are mainly a mean of classification of the building blocks to form different workflows. The resulting RA provides a schema for building solid workflows and ensures full communication and interaction between all the building blocks, from the data source level (at the infrastructure level of the organizations) up to the Data Stores and Processing Analytics to presentation and visualization applications.

High-Performance Computing (HPC) can be distributed at nodes within the platform supporting a high degree of scalability. Moreover, RA considers for external data sources such as public and private Data Lakes, IoT networks and Blockchains. A list of identified building blocks provides the basic functionalities of the INFINITECH reference sandbox for a more general class of use cases. Building blocks will be identified where existing technologies are available while other components will be designed, implemented, and integrated during the tasks belonging to work packages WP3, WP4 and WP5 of the project.

The validity of the RA has been proved by a mapping the workflows of the pilots of the projects, ultimately proving the conceptual approach of the INFINITECH RA.

Table of Contents

1.	Introduction	12
1.1	Scope and Main Objectives of the Deliverable	12
1.2	Insights from other Tasks and Deliverables	12
1.3	Methodology	14
1.4	Input from Other Tasks	15
1.5	Document Structure	15
2.	Challenges in infrastructures for BigData, IoT and AI applications in the finance sector	17
2.1	General AI / BigData Challenges	17
2.1.1	Big Data	17
2.1.2	Data Pipelines	18
2.2	Specific Challenges for the Finance Sector	19
2.2.1	Siloed Data and Business Operations	19
2.2.2	Real Time Performance Requirements	19
2.2.3	Mobility	20
2.2.4	Omni-Channel Banking - Multiple Channels Management	20
2.2.5	Automation	20
2.2.6	Transparency	20
3.	Analysis and Alignment to State of the Art and Industrial Best Practices	21
3.1	Rationale for a Reference Architecture	21
3.2	Overview of References Architectures for Data Driven Digital Finance Systems	21
3.2.1	Industry Vendors' Architectures for BigData in Digital Finance	21
3.2.1.1	IBM RA Architecture	21
3.2.1.2	Microsoft RA Architecture	23
3.2.1.3	WSO2 Architecture	25
3.2.1.4	HortonWorks	28
3.2.2	Reference Architectures from Standardization Bodies and Industry Associations	29
3.2.2.1	BDVA Reference Model	29
3.2.2.2	NIST Big Data Reference Architecture (NBDRA)	31
3.2.2.3	Industrial Internet Reference Architecture (RA)	32
3.2.3	Architectures of Relevant EU Projects and Research Initiatives	34
3.2.3.1	H2020 BigDataStack	34
3.2.3.2	H2020 BOOST	37
3.2.3.3	H2020 FINSEC	38
3.3	Overall INFINITECH-RA Positioning	39
4.	The INFINITECH Reference Architecture	40
4.1	Methodology: Architectural View Model	41
4.2	Key Assumptions	42

4.2.1	Definitions	43
4.3	Logical View	43
4.3.1	Logical Components Grouping	45
4.4	Process View	48
4.5	Development View	48
4.6	Deployment View	50
4.7	Use of INFINITECH-RA in Common BigData and AI Scenarios – Initial “Scenarios” View	51
4.7.1	Simple Machine Learning Pipeline	52
4.7.2	Blockchain Data Sharing and Analytics	52
4.7.3	Data Ingestion, Anonymization, Analytics and Visualization Pipeline	53
4.7.4	Data Exchange and Semantic Interoperability Pipeline	54
4.8	Mapping Methodology for Use Cases in the RA	57
5.	Addressing Stakeholder’s requirements	59
6.	Pilots’ Reference Designs and Initial Alignment to the INFINITECH-RA	61
6.1	Pilot#1: Invoices Processing Platform for a more Sustainable Banking Industry	61
6.1.1	Pilot Objectives	61
6.1.2	Data Sources	61
6.1.3	Data Produced	61
6.1.4	Explainable Workflow	61
6.1.5	Logical Schema	62
6.1.6	Components	62
6.2	Pilot#2: Real-time risk assessment in Investment Banking	62
6.2.1	Pilot Objectives	62
6.2.2	Data Sources	62
6.2.3	Data Produced	63
6.2.4	Explainable Workflow	63
6.2.5	Logical Schema	63
6.2.6	Components	63
6.3	Pilot#3: Collaborative Customer-centric Data Analytics for Financial Services	64
6.3.1	Pilot Objectives	64
6.3.2	Data Sources	64
6.3.3	Data Produced	64
6.3.4	Explainable Workflow	64
6.3.5	Logical Schema	65
6.4	Pilot#4: Personalized Portfolio Management (“Why Private Banking cannot be for everyone?”)	65
6.4.1	Pilot Objectives	65
6.4.2	Data Sources	66

6.4.3	Data Produced	66
6.4.4	Explainable Workflow	66
6.4.5	Logical Schema	67
6.4.6	Data Components	67
6.5	Pilot #5b: Business Financial Management (BFM) tools delivering a Smart Business Advise	68
6.5.1	Pilot Objectives	68
6.5.2	Data Sources	68
6.5.3	Data Produced	68
6.5.4	Explainable Workflow	69
6.5.5	Logical Schema	69
6.5.6	Components	69
6.6	Pilot#6: Personalized Closed-Loop Investment Portfolio Management for Retail Customers	70
6.6.1	Pilot Objectives	70
6.6.2	Data Sources	70
6.6.3	Data Produced	70
6.6.4	Explainable Workflow	70
6.6.5	Logical Schema	71
6.6.6	Components	71
6.7	Pilot #7: Operation Whitetail – Avoiding Financial Crime	72
6.7.1	Pilot Objectives	72
6.7.2	Data Sources	72
6.7.3	Data Produced	72
6.7.4	Explainable Workflow	72
6.7.5	Logical Schema	73
6.7.6	Components	73
6.8	Pilot#8: Platform for Anti Money Laundering Supervision (PAMLS)	73
6.8.1	Pilot Objectives	73
6.8.2	Data Sources	73
6.8.3	Data Produced	74
6.8.4	Explainable Workflow	74
6.8.5	Logical Schema	75
6.8.6	Components	75
6.9	Pilot #9: Analyzing Blockchain Transaction Graphs for Fraudulent Activities	76
6.9.1	Pilot Objectives	76
6.9.2	Data Sources	76
6.9.3	Data Produced	76
6.9.4	Explainable Workflow	76
6.9.5	Logical Schema	77

6.9.6	Data Components	77
6.10	Pilot #10: Real-time cybersecurity analytics on Financial Transactions’ BigData	77
6.10.1	Pilot Objectives	77
6.10.2	Data Sources	78
6.10.3	Data Produced	78
6.10.4	Explainable Workflow	78
6.10.5	Logical Schema	79
6.10.6	Components	79
6.11	Pilot #11: Personalized insurance products based on IoT connected vehicles	80
6.11.1	Pilot Objectives	80
6.11.2	Data Sources	80
6.11.3	Data Produced	81
6.11.4	Explainable Workflow	81
6.11.5	Logical Schema	81
6.11.6	Components	82
6.12	Pilot #12: Real World Data for Novel Health-Insurance products	82
6.12.1	Pilot Objectives	82
6.12.2	Data Sources	82
6.12.3	Data Produced	83
6.12.4	Explainable Workflow	83
6.12.5	Logical Schema	83
6.12.6	Components	83
6.13	Pilot #13: Alternative/automated insurance risk selection - product recommendation for SME	84
6.13.1	Pilot Objectives	84
6.13.2	Data Sources	84
6.13.3	Data Produced	84
6.13.4	Explainable Workflow	84
6.13.5	Logical Schema	85
6.13.6	Components	85
6.14	Pilot #14: Big Data and IoT for the Agricultural Insurance Industry	86
6.14.1	Pilot Objectives	86
6.14.2	Data Sources	86
6.14.3	Data Produced	86
6.14.4	Explainable Workflow	86
6.14.5	Logical Schema	87
6.14.6	Components	88
7.	INFINITECH-RA positioning	89

7.1	Scoping the INFINITECH-RA	89
7.1.1	RA seems generic and high level. Shouldn't it be more concrete?	89
7.1.2	Where is the Innovation?	89
7.1.3	What are the specific features for the Financial and Insurance Sector?	89
7.1.4	How can the RA cope with Business Requirements, Technological Requirements, Regulatory Requirements, etc.?	90
7.1.5	Where are the building blocks described in the DoA and how can be mapped in RA?	90
7.1.6	How can the INFINITECH-RA support Pilots and Use cases, and solve real Business Problems?	90
7.1.7	I am an End User and I cannot see how my Use Case can fit in the RA.	90
7.1.8	Is the RA being implemented in the Cloud? Will my data be distributed outside my organization? How can I preserve the classification of my own data?	91
7.1.9	Blockchain Technology does not seem to fit in the RA. How can it be managed?	91
7.2	Development Roadmap and Future Outlook	91
8.	Conclusions	92
9.	Appendix A: Literature	93
10.	Appendix B: APPENDIX User requirements' mapping	95

List of Figures

<i>Figure 1</i>	: Interrelationship of INFINITECH-RA (D2.13) to other key deliverables of WP2	13
Figure 2:	High-Level Overview of the Project's Phased Approach to INFINITECH-RA Development	15
<i>Figure 3</i>	– Relations between Tasks in WP2	15
Figure 4 -	IBM Big Data and Analytics Reference Architecture	22
Figure 5	Microsoft RA Logical Banking Technology Architecture	24
Figure 6 -	Microsoft RA Anti-Money Laundering Scheme	24
Figure 7 -	A next-generation financial infrastructure requirements	26
Figure 8 -	WSO2 Ecosystem Solution	27
Figure 9	Big Data Adoption Lifecycle	28
Figure 10	BDVA – Reference Model	30
Figure 11 -	BDV – Reference Model (top) vs INFINITECH Reference Architecture (bottom)	31
Figure 12 –	NIST Reference Architecture	32
Figure 13:	NIST BigData Architecture and Infrastructure Challenges	32
Figure 14	IIRA constructs and application	33
Figure 15 -	BigDataStack architecture model	34
Figure 16 -	The Boost 4.0 Reference Architecture	37
Figure 17 -	The FINSEC Reference Architecture version 2.0	38
Figure 18 -	"4+1" views	42
Figure 19 -	INFINITECH Reference Architecture Mapping with BDVA Reference Model	44
<i>Figure 20</i>	– Example of Neural Network the IRA Analogy	44
Figure 21 -	INFINITECH Reference Architecture Logical View – Example Mapping	45
Figure 22 -	INFINITECH Reference Architecture Logical View – Example Mapping	48

Figure 23 Deployment options for an end-to-end Software System..... 51

Figure 24: Simple Machine Learning Workflow Example implemented in-line with the INFINITECH-RA52

Figure 25: Blockchain Data Sharing and Analytics Pipeline 53

Figure 26: Pipeline for a Scenario involving Semantic Interoperability Across Diverse Data Sources 55

Figure 27: Logical Architecture of Invoices Processing Pilot..... 62

Figure 28: Logical Architecture of Real-Time Risk Assessment Pilot 63

Figure 29: Workflow Elements of the Customer-Centric Data Analytics Pilot – KYC Data Sharing Process – Business Workflow 64

Figure 30: Workflow Elements of the Customer-Centric Data Analytics Pilot – KYC Data Sharing Process – Technical Workflow 65

Figure 31: Logical Architecture of the Customer Centric Analytics Pilot in-line with the INFINITECH-RA 65

Figure 32: Explainable Workflow of the INFINITECH Personalized Portfolio Management Pilot..... 67

Figure 33: INFINITECH Personalized Portfolio Management Pilot Pipeline in-line with the INFINITECH-RA 67

Figure 34: Business Financial Management Pilot Architecture in-line with the INFINITECH-RA..... 69

Figure 35: Personalized Closed-Loop Investment Portfolio Management 71

Figure 36: Logical Architecture of the Financial Crime Pilot in-line with the INFINITECH-RA 73

Figure 37: PAMLS Pilot Logical Architecture in-line with the INFINITECH-RA 75

Figure 38: Logical View of the Blockchain Transaction Graphs Analysis pilot in-line with the INFINITECH-RA components..... 77

Figure 39: Logical Architecture of the Real-time cybersecurity analytics Pilots in-line with the INFINITECH-RA..... 79

Figure 40: Logical Architecture of the Personalized insurance products based on IoT connected vehicles pilot in-line with the INFINITECH-RA..... 82

Figure 41: INFINITECH-RA Compliant Architecture for the Real World Data for Novel Health-Insurance products pilot..... 83

Figure 42: INFINITECH-RA Compliant Architecture for the Alternative/automated insurance risk selection - product recommendation for SME Pilot 85

Figure 43: INFINITECH-RA Compliant Architecture for the Pilot on Big Data and IoT for the Agricultural Insurance Industry 87

Figure 44: Workflow of the Pilot on Big Data and IoT for the Agricultural Insurance Industry..... 87

List of Tables

Table 1: List of INFINITECH Components and Technologies 46

Table 2: INFINITECH-RA vs Key Requirements for BigData and AI Applications in Digital Finance 59

Table 3: INFINITECH-RA vs Key Challenges for BigData and AI Applications in Digital Finance 60

Table 4: Mapping of INFINITECH DoA/Task KPI with Deliverable Achievements 92

Abbreviations

Acronym	Definition
ACID	Atomicity, Consistency, Isolation, Durability
AI	Artificial Intelligence
AML	Anti Money Laundering 4AMLD fourth edition of AML directive
API	Application Programming Interface
ART	Adaptive Resonance Theory
BDVA	Big Data Value Association
BFM	Business Finance Management
BI	Business Intelligence
CCTV	Closed Circuit TeleVision
CEP	Center (for) Energy Policy
CI	Critical Infrastructure
CPU	Central Processing Unit
CRISP	Complex-Reduced Instruction Set Processor
CRM	Customer Relationship Management
CRUD	Create Retrieve Update Delete - Basic Operations in DBMS
DL	Deep Learning
DLT	Distributed Ledger Technology
DM	Distributed Memory
DPO	Data Protection Officer
DPOR	Data Protection Orchestrator
DUOS	Digital User Onboarding System
DSS	Decision Support System
EO	Earth Observation
ES	Expected Shortfall
ESB	Emergency Service Bureau
ETL	Extract, Transform, Load
ETSI	European Telecommunications Standards Institute
EU	European Union
FA	Factory Automation
FAR	Federation of American Research
GDPR	General Data Protection Regulation
GPS	Global Positioning System
GUI	Graphical User Interface
HDP	High Definition Progressive
HPC	Handheld Personal Computer
HTAP	Hybrid transaction/analytical processing

D2.13 – INFINITECH Reference Architecture - I

Acronym	Definition
HW	HardWare
IBM	International Business Machines
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IEEE	Institute (of) Electrical (and) Electronic Engineers
IIRA	Industrial Internet Reference Architecture
IRA	INFINITECH REFERENCE ARCHITECTURE
ISO	International Organization for Standardization
IT	Information Technology
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
LEGO	Low End Graphics Option
MDM	Mobile device management
ML	Machine Language
MVP	Minimum Viable Product Platform
NIST	National Institute of Standards and Technology (formerly NBS)
ODS	Open Data Services Overhead Data Stream
OLAP	On Line Analytical Processing
PAMLS	Platform for Anti Money Laundering Supervision
PLM	Product Lifecycle Management
POD	Pay On Delivery Point Of Distribution
PSD2	2nd Payment Services Directive
RA	Reference Architecture
RDBMS	Relational Database Mangement System
REST	Representational State Transfer
SAP	Service Access Point
SOTA	State of The Art
SQL	Structured Query Language
TBD	To Be Determined
UI	User Interface
UML	Unified Modelling Language
VaR	Value-at-Risk
VPC	Virtual Path Connection
XML	Extensible Markup Language

1. Introduction

1.1 Scope and Main Objectives of the Deliverable

INFINITECH is developing and validating BigData, IoT and Artificial Intelligence (AI) technologies for the finance and insurance sectors. In this direction, the project is advancing the state of the art in several technological development areas such as infrastructures for integrated, incremental and real-time analytics, semantic interoperability, as well as decentralized information sharing between stakeholders of the sector. Furthermore, the project designs and implements a range of pilots and use cases that aim at validating these technologies in real-life scenarios of the sectors. One of the main objectives of the project is to define a Reference Architecture (INFINITECH-RA), which will serve as a blueprint for developing, deploying and operating BigData, AI and IoT in the finance and insurance sectors. This reference architecture should define the structuring principles of the solutions, along with their main building blocks. In this way it will facilitate solution architects in developing and documenting specific solutions. Likewise, the INFINITECH-RA will facilitate stakeholders' communications regarding BigData, AI and IoT systems for the finance and insurance sector, through providing a uniform set of terms and definitions that will be unambiguously understandable by stakeholders.

The purpose of this deliverable is to introduce and describe the INFINITECH-RA. It emphasizes on the presentation of the structuring principles of the RA, building on experience and recommendations from other general-purpose architecture for BigData applications such as the RA of the BigData Value Association (BDVA). Moreover, the deliverable will illustrate why and how the presented RA can support the main use cases of the sector, including the use cases that are designed and implemented in the scope of WP7 of the project. It will consider business requirements, application requirements and regulatory requirements documented in other deliverables of this workpackage.

The deliverable presents the initial version of the INFINITECH-RA. This initial version will drive the technological developments of the project, as well as the design and initial integration of the use cases. As part of these development and integration processes feedback on the relevance and appropriateness of the RA will be solicited. This feedback will be taken into consideration in the development of the second and final version of the RA as part of deliverable D2.13.

1.2 Insights from other Tasks and Deliverables

The INFINITECH Reference Architecture (RA) signals an important milestone for the INFINITECH work on WP2. Specifically, it provides the structuring principles that will drive the design, development and integration of BigData, AI and IoT use cases in the finance sector, taking into account the INFINITECH use cases and technologies, as well as wider requirements of the sector. As such INFINITECH RA is:

- Driven by earlier WP2 deliverables on user stories and stakeholders' requirements, namely deliverable D2.1 "User Stories and Stakeholders' Requirements (Version 1)". The INFINITECH RA shall be able to fulfil the project's BigData/AI user stories for the finance sector, along with key stakeholders' requirements. Specifically, the user stories of the project will be used to validate the functionality and the versatility of the Reference Architecture, in terms of its ability to support a broad range of applications for the sector.
- Considering the security and regulatory compliance specifications presented in deliverable D2.7 "Security and Regulatory Compliance Specifications (Version 1)". The INFINITECH-RA shall boost compliance to the main regulations of the sector such as the 2nd Payment Services Directive (PSD2) and the 4th Anti-Money Laundering (AML) directive, while considering the compliance of financial sector applications to general, broadly applicable regulations such as the General Data Protection Regulation (GDPR).

- Closely related to the deliverable that describes the INFINITECH components and technologies, namely deliverable D2.5 “Specifications of INFINITECH Technologies (Version 1)”. In practice, there has been a two-way interaction between D2.5 and D2.13. The structure and functionalities of the INFINITECH components in D2.5 have been considered in the development of the reference architecture that drives their integration in compete BigData / AI use cases and pipelines. On the other hand, the Reference Architecture is also unveiling the need for new functionalities in the INFINITECH components, as well as for entirely new software and middleware elements needed to support the integration of INFINITECH technologies in turn-key solutions.

The following figure illustrates the interaction of the present deliverable with other key deliverables of WP2. Note that this interaction will take place as part of the methodology for the development of the architecture, which is illustrated in the next section. For example, the adoption of the 4+1 views methodology will enable the validation of the INFINITECH-RA against user stories developed in D2.1 of the project.

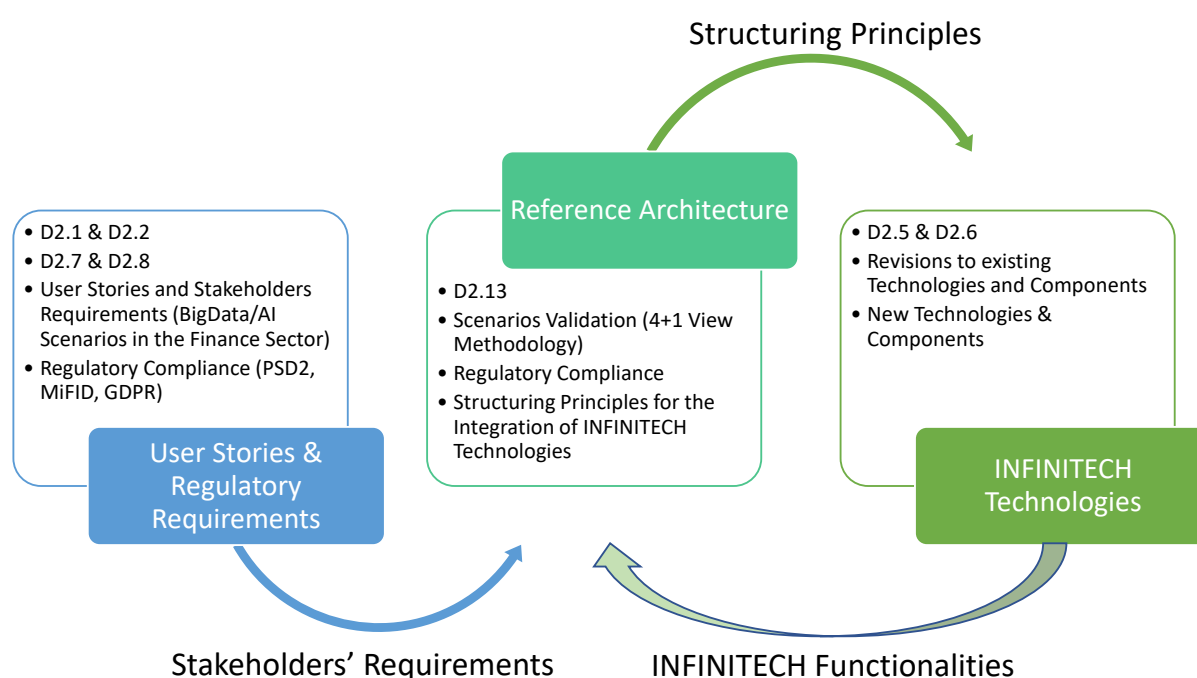


Figure 1 : Interrelationship of INFINITECH-RA (D2.13) to other key deliverables of WP2

Beyond its cornerstone role in WP2, the present deliverable is closely linked to other technical developments of the project, namely developments of the technical workpackages of the project. Specifically, INFINITECH-RA will:

- **Impact the design of sandboxes and testbeds in WP6 of the project**, as these testbeds will be destined to host INFINITECH-RA deployments.
- **Drive the integration of the project’s pilots and use cases in WP7**, as all pilots will align to the reference architecture of the project. The present deliverable provides already an early, INFINITECH-RA compliant design of the main BigData/AI pipelines that comprise the project’s pilots.
- **Provide insights for the detailed design of technical components in WP3, WP4 and WP5 of the project**, including for example data management middleware, semantic interoperability elements, blockchain components for data sharing and decentralized application development, as well as Machine Learning (ML) workflows and algorithms. While the design and development of these components is in-principle independent of their use in a specific architecture, INFINITECH-RA will have an impact on the way these components interact between them. This will have a secondary influence on their design and development.

Overall, the present deliverable has an instrumental role in the project, as it drives many tasks and deliverables in the project. That's the reason why it is linked with one of the milestones of the project's workplan.

1.3 Methodology

The methodology for the development of the INFINITECH-RA is based on the following overlapping phases:

- **Phase 1 – Review and Analysis (M2-M7):** This phase was focused on the analysis of relevant technical and scientific information, towards ensuring that the INFINITECH-RA considers recent developments in BigData architectures in general and for the finance sector in particular. It analysed the BDVA RA, along with BigData architectures for digital finance applications that have been introduced and are widely used in the industry. In this context, BigData/AI environments and tools have been reviewed as well. Likewise, this phase considered the stakeholders requirements (from T2.1 of the project), as well as regulatory compliance requirements and the specifications of INFINITECH technologies reflected in D2.5 and D2.7 respectively. The overarching goal of this phase was to make sure that the design of the INFINITECH-RA aligns to key requirements of BigData applications in digital finance and to the evolution of the state of the art in BigData in finance.
- **Phase 2 – Architecture Design (M6-M12):** This phase produced the initial design of the INFINITECH-RA, using the 4+1 view methodology [Kruchten95] for specifying software architectures. The selection of this proven and well-known methodology was motivated from the fact that INFINITECH pilot systems are essentially software systems. Hence, a methodology for specifying software architectures was applicable. In-line with the 4+1 architecture the INFINITECH-RA is described as a collection of complementary views (i.e. process, development, logical, physical) that represent dynamic behaviours and static behaviour of the systems, as well as relevant implementation aspects. Likewise, 4+1 signifies the need for confronting the architecture against a number of scenarios like the INFINITECH pilots. As part of the second phase of the INFINITECH development, a set of main views for the INFINITECH-RA were specified and described, as presented in subsequent paragraphs. Note that the phase extends beyond the delivery of the first version of the INFINITECH-RA (i.e. to M12 of the project), as several of the presented views will be revised in the coming months.
- **Phase 3 – Fine Tuning and Updates (M9-M18):** This is the third development phase of the INFINITECH-RA, which is destined to begin following the conclusion of this deliverable. It will aim at revised and fine-tune the INFINITECH-RA based on its actual use in other workpackage, including the sandboxes and pilot development workpackages. It will receive and exploit stakeholders' feedback from the practical use of the INFINITECH-RA for developing BigData/AI systems for the finance sector. The feedback will be exploited towards improving aspects of the INFINITECH-RA, but also towards validating the architectural concepts introduced in practice. This Phase will lead to the final version of the INFINITECH-RA as part of deliverable D2.14 i.e. the second and updated version/release of this deliverable.

The following figure illustrates the three development phases of the INFINITECH_RA and the main activities that they comprise.

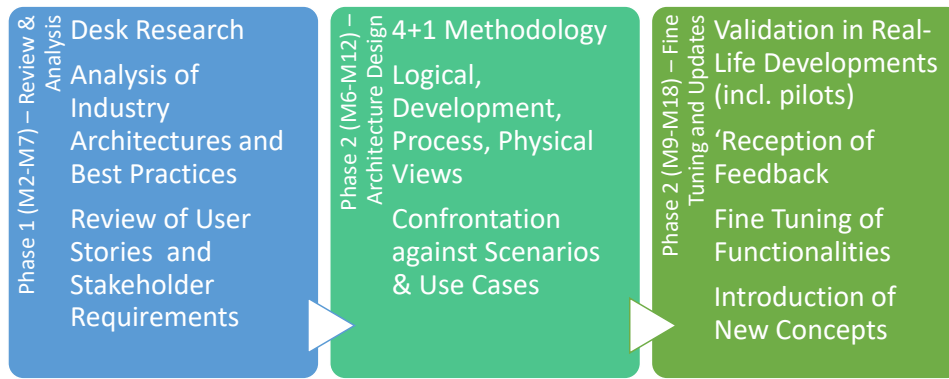


Figure 2: High-Level Overview of the Project's Phased Approach to INFINITECH-RA Development

1.4 Input from Other Tasks

The specifications included in this deliverable take into account the requirements and specifications of all the previous tasks of WP2 [1-6] according to the diagram below.

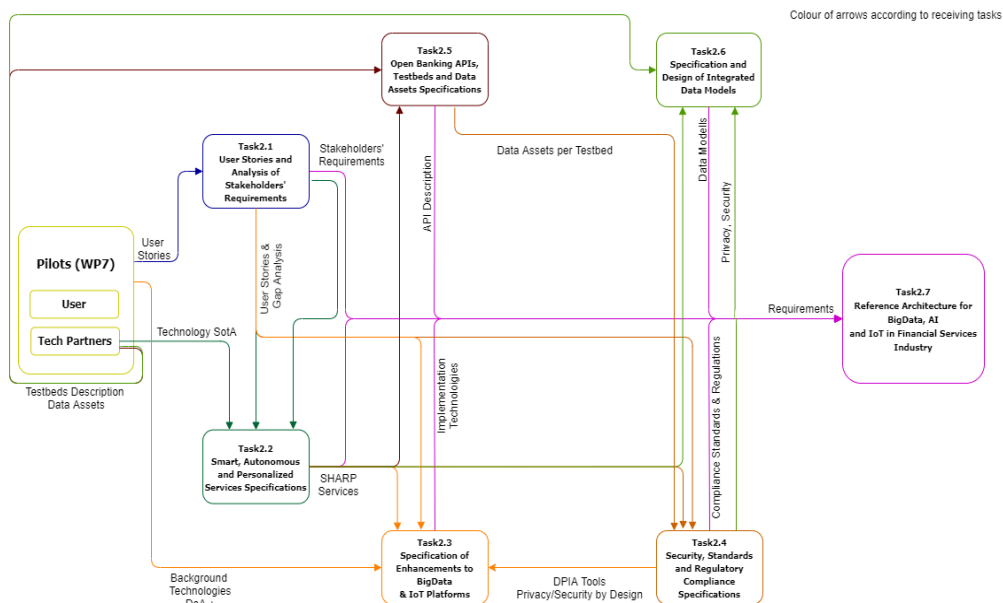


Figure 3 – Relations between Tasks in WP2

1.5 Document Structure

The deliverable is structured as follows:

- **State of the Art Analysis:** First, the earlier approaches for RA design of big data systems are reviewed. Then, utilisation of ML/AI techniques in SW engineering is discussed. Then, different architectural approaches are presented for realising ML utilisation in edge computing environments. State of the art analysis is included in Sections 2 and 3 of the document.
- **RA DESIGN:** The INFINITECH-RA design is presented with different architectural views (Section 4).
- **ANALYSIS:** The RA is analysed and discussed, while its ability to support some popular BigData/AI scenarios in Digital Finance is outlined (Section 5).
- **Pilots Mapping:** Section 6 of the deliverable provides a detailed mapping of the Pilots workflows to the RA, which validates the RA conceptual designs and building blocks. This is part of the

“Scenarios” view of the 4+1 architecture methodology that is used to introduce and describe the INFINITECH-RA.

- **Discussion and Outlook:** The unique properties of the INFINITECH-RA and its practical use is discussed in Section 7. Moreover, an outlook for the future development and evolution of the RA is presented.
- **Conclusion:** Section 8 is the concluding section of the deliverable.
- **Appendices:** Two Appendices of the deliverable are devoted to the references/citations of the document and to the mapping of stakeholders’ requirements to the architecture respectively. There are numbered as Sections 9 (Appendix A) and 10 (Appendix B) of the deliverable.

2. Challenges in infrastructures for BigData, IoT and AI applications in the finance sector

2.1 General AI / BigData Challenges

2.1.1 Big Data

Over the past few years and in the context of big data, new challenges have emerged in data analytics, turning the attention to methodologies that raise the abstraction level and facilitate the convergence rate of algorithms. To this end, several big data frameworks have been developed by researchers and engineers and respective algorithms have been compiled, mainly being domain-specific. However, the development of scalable and distributed analytics solutions for extreme-scale analytics in the finance / insurance sector remains quite a complicated process given the diversity in datasets and data sources. Data sources can be of different type (structural, semi-structural or completely unstructural), in various formats with different ways for data access. Due to this, in recent years, modern enterprises rely on a variety of different and heterogeneous data management systems in order to handle with this data diversity: relational datastores that are used to store structural data compliant with an entity-relational model where the ensurance of transactional semantics and data consistency is of major concern, document-based data management systems to store semi-structural data, key-value stores that are considered efficient to store data coming from IoT sensors or logging information (i.e. when navigating among different web pages in the network, or simply logging the details of a finance transaction that took place) and even the use of HDFS data lakes is now considered prominent and facilitates the analysis of unstructural data that can be available. However, the analysis over a superset of the available data management system of an organisation is not a trivial task. Joining data across datasets is very cost demanding and difficult to be implemented efficiently in the application or data processing level efficiently. To overcome this barrier, various analytical framework has been proposed that provides polyglot capabilities and abstract this problem from the data scientist, making the data management process transparent from her perspective.

Accessing heterogeneous data sources (a concept often addressed by data integration systems or multidatabases [Özsu11], [Tomasic98]) is a problem that has been widely studied in the literature and with the recent emerge of cloud database and big data processing, it has been evolved towards polystore systems. Early implementations [Minpeng11], [Ong14], [Simitsis12] relied on a single common model that the target datastores had to transform their schema to. A further improved model presented by the polystore BigDAWG [Duggan15], [Gadepally16] which defines islands of information, which makes use of a single data model and language. Nowadays, Spark SQL [Armbrust15] exploits its advantage for massive parallel processing over a federation of different and heterogeneous stores. It defines the notion of *dataframe* and provides different connectors over a variety of supported datastores. By providing a single interface and a common query language, it pushes down the query execution on the target databases, when this is possible, and retrieves data into those dataframes that are being used for further processing. That requires however the retrieval of the data in the data analysis layer, which can be memory greedy. Similarly Presto³ allows for massive parallel processing where a coordinator orchestrates the query execution on several workers which make use of corresponding connectors, all of them implementing a common interface in order to hide the details of the query execution on the target database. Towards the same direction, Apache Drill⁴ and Impala⁵ maintain the notion of the data connectors that are able to retrieve data from external and heterogeneous datastores, and transform it to their own intermediate format and model

³ Presto – Distributed Query Engine for Big Data, <https://prestodb.io/>

⁴ Apache Drill – Schema-free SQL Query Engine for Hadoop, NoSQL and Cloud Storage, <https://drill.apache.org/>

⁵ Apache Impala, <http://impala.apache.org/>

that can be used for data processing in the upper layer of the framework. All those approaches provide polyglot capabilities, however, they are considered as additional frameworks that require the retrieval of the dataset in memory, and exploit their abilities for massive parallelism to be able to scale out adequately in order to deal with these requirements. The challenge lies on exploring the attributes and statistics of each individual dataset that is stored in each store, optimize the query execution and retrieve only the minimum amount of data that is needed in each query operator, thus minimizing the need for memory consumption and data traffic in the network.

Besides the analytics level, running analytics over stale data is not ideal. For example, consider the case of a recommender system that aims to predict the behavioural patterns of a user, her preferences or dislikes, and provide personalised recommendations of relevant items. In a real-world scenario, user preferences change frequently, and new data continuously arrive in a real-time manner. A recommender system should, if possible, adapt to these changes as they happen, modifying its model to always reflect the current status, while requiring a single pass through the data. Furthermore, training new versions of the model in a continuously growing dataset is computationally inefficient, leading also to unnecessarily amplified infrastructure costs. Hence, data management approaches are required both to account for new datasets and to support these incremental analytics.

Apart from the scenarios where systems and analytical algorithms should continuously taking into account the overgrowing datasets in order to update the results of their analysis, they will also need to analyse real-time data in order to respond to events and create alerts and notifications as they events occur. An example of such a need is the fraud that must be detected when the fraud transaction occurs. However, dealing with data related with financial transaction raises a lot of concerns regarding data consistency and isolation. This is the reason why traditional relational database management systems are used, which ensures transactional (from a database theory perspective) semantics and ensure that a series of actions are executed all at once or the transaction will be aborted (Atomicity), the data will be consistent (Consistency) and they will not be stored breaking the rules that have been imposed by the database administrator (the balance of a bank account cannot be negative), parallel actions can be executed concurrently, but the result of their modifications will be likewise each action was executed in order (Isolation), and data will be preserved in terms of a system failure and all information can be recovered (Durability). Those four attributes consist of the ACID properties that traditional database management systems ensure and are of major importance in the finance sector. These databases are ideal for handling operational workloads, and they rely on the use of locks on data items that are being modified when a transaction access a data table. However, in order to perform a fraud detection, the system needs to read a variety of data by scanning the whole data table. In other words, to perform an analytical operation. Operational and analytical operations are compatible and one blocks the other, as the existence of a lock will prevent the latter from being executed. Likewise, an long-running analytical operation will pose read locks in the whole data table, thus blocking all operational workloads from being executed: the database will not be able to serve customer needs as long as the long-running operation takes place. Due to this, modern enterprises take a snapshot of the operational database and send it to a data warehouse, thus separating those two workloads: The relational datastore will serve operational workloads by ensuring the ACID-ity, while the data warehouse will serve analytical query processing. The drawback of this approach is that detection of a fraud transaction will take place the following day, as moving data from one store to the other is cost-expensive operation and relies on heavy ETLs that are being performed during the night. Even exploiting modern approaches and apply micro-batching in the data movement process (i.e. the modifications of a transaction are being sent to a data queue, and a worker periodically gets data from the queue and sends them to the analytical warehouse) this allows for near real-time analytical processing, and cannot detect events that happen on real-time. Due to this, the analytics over operational data is still of great challenge.

2.1.2 Data Pipelines

An additional aspect and challenge refers to data management that needs to be enhanced in order both to facilitate the needs at extreme-scale (in terms of efficiently for data throughput and access)

and to address the challenge of various federated and distributed systems that hold and corresponding datasets. What is more, one needs to consider the changes on the underlying systems (and datasets) and thus the required level of dynamicity, raising the challenge for approaches that perform data pipelining from heterogeneous distributed systems towards analytics frameworks, while being adaptive to the aforementioned changes. Tightly-coupled multistore systems trade autonomy for performance, typically in a shared-nothing cluster, to integrate structured (RDBMS) and HDFS data. Polybase [Minukhin18] is a feature of Microsoft SQL Server Parallel Data Warehouse to access HDFS data using SQL. It allows HDFS data to be referenced through external PDW tables and joined with native PDW tables using SQL queries. Hybrid systems are similar to tightly-coupled systems, e.g. integrating HDFS and RDBMS in a shared-nothing cluster, except that the HDFS data is accessed through a data processing framework like MapReduce. For example, QoX [Xu18] uses a dataflow approach for optimizing queries over relational (RDBMS and ETL) data and unstructured (HDFS) data, with a black box approach for cost modelling. Moreover, when dealing with analytics performed on a variety of data sources, it is no less important to focus on the data aspects. Data-intensive distributed frameworks such as Apache Spark and Apache Drill can access multiple data stores using a unified API such as SQL. However, applications running on these frameworks have direct access to specific data stores and as a result to specific datasets, while frameworks such as Apache Ranger offer standardized access authorization to data stores, but only for a limited set of supported data stores and with limited policies. What is required refers to an approach for data management that minimizes the data pipelining process and would enable a hybrid management of data, both for analytical and for transactional workloads. The latter would enable analytics to account for the different datasets made available as they are ingested in the data stores (mainly through the respective transactions).

2.2 Specific Challenges for the Finance Sector

2.2.1 Siloed Data and Business Operations

One of the most prominent challenges faced by banks and financial organizations is the fragmentation of data across different data sources such as databases, data lakes, transactional systems (e.g., e-banking) and OLAP (On Line Analytical Processing) systems (e.g., Customer Data Warehouses). This is the reason why financial organizations are creating BigData architectures that provide the means for consolidating diverse data sources. As a prominent example, the Bank of England has recently established the a “One Bank Data Architecture” based on centralized data management platform⁶. This platform facilitates the BigData Analytics tasks of the bank, as it permits analytics over significantly on larger datasets.

Note that the need for reducing data fragmentation has been also underlined by financial institutions following the financial crisis of 2008, where several financial organizations had no easy way to perform integrated risk assessments as different exposures (e.g., on subprime loans or ETFs (Exchange-Traded Fund)) were siloed across different systems.

INFINITECH-RA must therefore provide the means for reducing data fragmentation and taking advantage of previous siloed data in integrated BigData Analytics and ML tasks.

2.2.2 Real Time Performance Requirements

Real Time Computing is when an IT system must respond to changes according to definite time constraints, usually on the order of milliseconds or seconds. In the realm of financial and insurance sectors, real time constraints apply where a response must be given to provide services to users or organizations and are in the order of seconds or more. Examples range from banking application to cybersecurity. Most of real-world financial applications are NOT real time as other industrial automation (plant control) and are usually solved putting more computing resources (cpu/gpu power,

⁶ <https://www.cloudera.com/about/customers/bank-of-england.html>

memory, ...) at the problem. However, in the case of ML/DL and BigData algorithms can take significant amount of time and become useless in practical cases (responses arrive too late to be used). In those cases, a quantitative assessment of the computing time of algorithms is needed to configure resources to provide acceptable time.

2.2.3 Mobility

The digital transformation of financial institutions includes a transition to mobile-first banking [Bons12]. This refers to the interaction of customer and financial organizations through mobile channels as part of mobile banking. INIFINITECH must support mobile channels when visualizing BigData, AI and IoT applications for digital finance, but also when collecting and processing user's / customer's input.

2.2.4 Omni-Channel Banking - Multiple Channels Management

One of the main trends in banking and finance is the transition from conventional multi-channel banking to omni-channel banking [Komulainen18]. The latter refers to seamless and consistent interactions between customers and financial organizations across multiple channels. Hence, omni-channel banking/finance focuses on integrated customer interactions that comprise multiple transactions, rather than individual financial transactions. The INFINITECH-RA should provide the means for supporting omni-channel interactions through creating unified customer views and managing interactions across different channels based on integrated/consolidated information about the customer. Note that BigData analytics is the cornerstone of omni-channel banking as it enables the creation of unified views of the customers and the execution of analytical functions (including ML) that track, predict and anticipate customer behaviours.

2.2.5 Automation

Data intense applications, also in the banking and insurance sectors, are realized by specialized IT and RDBMS administrators. Recently more and more data scientists and business analysts are involved in the development of such applications at a significant cost.

The INFINITECH-RA should provide the means for orchestrating data intense application and management through easily creating workflows and data pipelines.

2.2.6 Transparency

During the last couple of year, financial organizations and customers of digital finance services raise the issue of transparency in the operation of BigData systems as a key prerequisite for the wider adoption and use of BigData analytics systems in financial, sector use cases. This is particularly important for use cases involving the deployment and use of AI/ML systems that operate as black-boxes and are hardly understandable by finance sector stakeholders. Hence, a key requirement for AI/ML financial sector application is to be able to explain their outcomes. For example, a recent paper by bank of England [Bracke19] illustrates the importance of providing explainable and transparent credit risk decisions. Overall, INFINITECH shall provide support for transparency in AI/ML workflows based on the use of Explainable Artificial Intelligence (XAI) techniques such as LIME (Local Interpretable Model-agnostic Explanations) [Ribeiro16], which develops/finds a local model around the prediction that is interpretable. XAI techniques should be supported as part of the ML techniques of the project, but also as an add-on to them as in several cases they will be used to interpret the outcomes of other ML/AI techniques (e.g., some classifiers).

3. Analysis and Alignment to State of the Art and Industrial Best Practices

3.1 Rationale for a Reference Architecture

RAs are designed for facilitating design and developments of concrete technological architectures, mostly in the IT domain, reducing risks with proven components, and improvement of communication within an organisation.

In general, the value of RAs can be summarized in the following points:

- reduction of development and maintenance costs of systems
- facilitation of communication between important stakeholders
- and reduction of risks

Typically, when a system is designed without a RA, an organisation may accumulate technical risks and end up with a complex and non-optimal implementation architecture.

In the industry, complex infrastructures for big data systems and high performance computing (HPC) have been developed and proved to sustain intensive data processing services (Netflix, Facebook, Twitter, LinkedIn etc.). The architectures and technologies of world class infrastructure have been published and RAs have been designed and proposed. However, very few solutions have been published for the Financial and Insurance sectors and this deliverable aims to partially cover the gap.

In the following sections, some relevant Reference Architectures and Models will be considered along with their relevance to the domain sector at which INFINITECH is aimed.

3.2 Overview of References Architectures for Data Driven Digital Finance Systems

3.2.1 Industry Vendors' Architectures for BigData in Digital Finance

3.2.1.1 IBM RA Architecture

The IBM Big Data RA capabilities view shown in *Figure 4* is a vertically layered architecture. The left most layer are the data sources, both traditional ones and new ones, structured (such as relational DBMSs, flat files), semi-structured (such as XML) or unstructured (video, audio, digital etc). Traditional or back office applications execute and process the business transactions. These systems include the order processing, billing, marketing, product development and sales types of systems. New sources are systems that supply information to augment the information generated by the back-office applications. Those sources might include IoT data, video surveillance data, etc. The data sources provide data or raw materials to the analytics ecosystem that will be the foundation for analysis and knowledge.

The second left most layer is the data integration and analytics layer, both for data at rest and data at motion. **Data Integration** vertical is the data acquisition and transport pipeline between the Data Source Providers and the analytic ecosystem. The **Real-Time Analytical Processing streaming layer** supports data transfer at a steady high-speed rate to support many zero latency (“business real time”) applications. The **Data Warehousing** provides raw and prepared data for analytics consumption. **Shared Operational Data** components own, rationalize, manage and share important operational data for the enterprise.

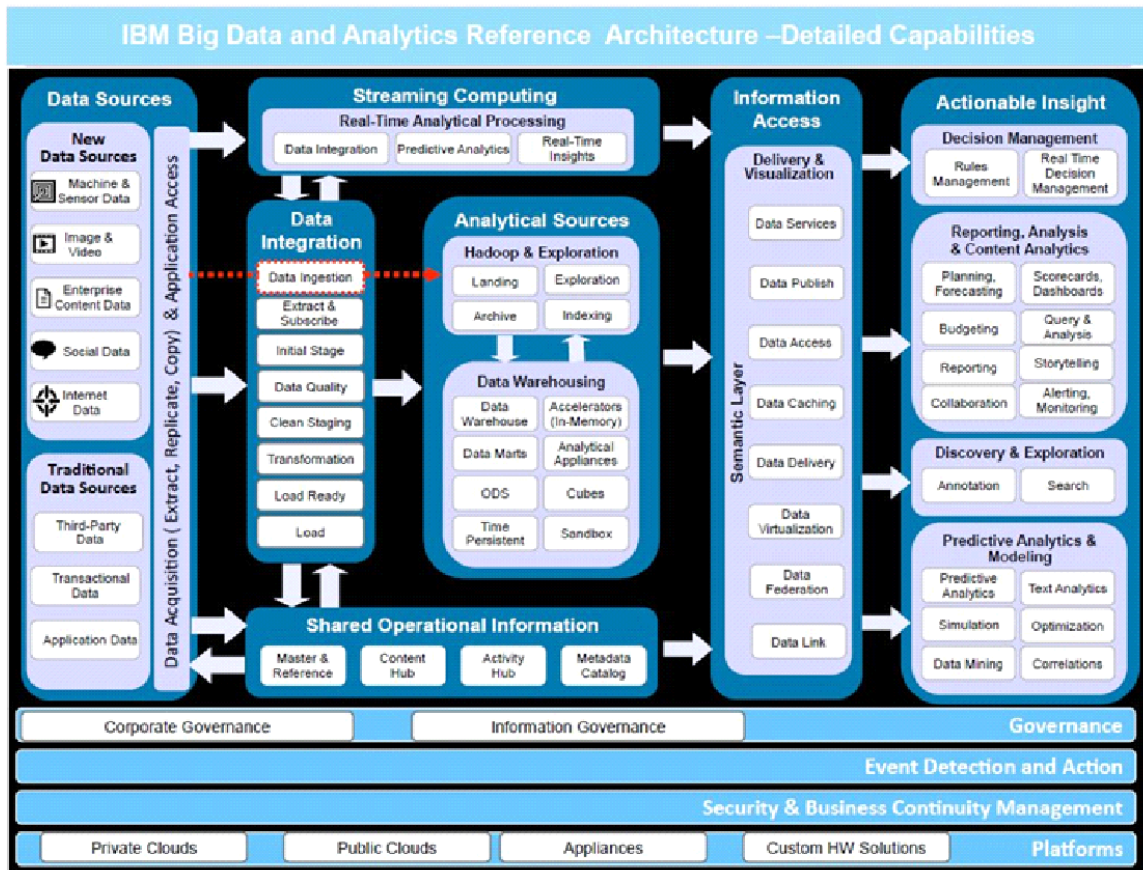


Figure 4 - IBM Big Data and Analytics Reference Architecture⁷

The third layer is the **Information Access Layer**, which is a semantic layer responsible for delivery and visualization of processed information to the next layer, the actionable insight layer. It enables the analytic applications to publish, locate and consume information without needing to understand exactly how and where it is physically stored and maintained. The goal is to provide a unified view of all information assets and can be divided into multiple components such as a database-resident Semantic Layer, data federation, data virtualization and data services.

Finally, the **Actionable Insight layer** use data and models to provide insight to guide business decisions. It includes decision management, reporting, predictive analytics, planning and forecasting applications. On top of this layer enhanced applications can be provided, such as new business models, enhanced customer experience applications, risk-assessment applications, fraud and operational efficiency applications etc.

Cross cutting across all those functional layers are the operational layers of information governance, security and business continuity and infrastructure platforms.

Every industry can leverage big data and analytics platform capabilities to its advantage, for example for banking industry it can be used to optimize offers and cross-selling, customer service and call service efficiency etc.

Relevance to INFINITECH RA

The IBM Big Data and Analytics RA suggests a collection of advanced data governance and security as well as analytics technologies. Data protection and management in a vertically layered approach

⁷ <https://www.ibm.com/developerworks/community/files/form/anonymous/api/library/e747a4bd-614d-4c5d-a411-856255c9ddc4/document/bbc80340-3bf4-4e0a-8caf-a43f64a22f05/media>

accompanied with the respective infrastructure frameworks are considered absolutely necessary for business stability and development.

INFINITECH is intended to govern, secure and analyze advanced data in the financial industry through its custom developed enablers. Thus, IBM Big Data and Analytics RA is relevant to INFINITECH RA as it addresses advanced data governance and security together with analytics technologies. The following IBM Big Data and Analytics RA elements are acknowledged as applicable and relevant for the INFINITECH RA:

- **Shared Operational Data:** Process mining is destined to be of the fundamental capabilities in the overall INFINITECH. This IBM component should be taken under consideration when designing and developing governance, management and monitoring of system processes as well as operational data processing in INFINITECH.
- **Information Access Layer:** INFINITECH will implement data distribution and management when combining and maintaining dissimilar data from different sources, such as IoT and Blockchain. Information Access layer should be certainly considered in this manner.
- **Actionable Insight layer:** INFINITECH will envelope several types of data analysis, with predictive analytics being of highly important technology focus along the other types, such as incremental, declarative and parallel analytics. The Actionable Insight layer should be taken into account when organizing and planning the INFINITECH analytics.
- **Data Integration, Real-Time Analytical Processing streaming layer and Data Warehousing:** INFINITECH will be needed such tools when implementing algorithms for efficient data integration and preparation before delivery upon the data analysis infrastructure. These IBM components and layer should be examined when determining continuous unification of INFINITECH modules for data transfer efficiency and arrangement for analytics applications.

3.2.1.2 Microsoft RA Architecture

The Microsoft Reference Architecture provides a Logical Banking Technology Architecture schema (*Figure 2*) (see also [MIRAB12]). It enables high value integration to other systems through a wide-array of industry-standard integration interfaces and techniques (such as ISO, BIAN, IFX) and reduces further the managing and business solution maintaining cost in the banking industry. Additionally, Microsoft RA offers an industry-leading set of robust functionalities defined and exploited in both the bank datacenters and cloud. This kind of functionalities extend across the overall IT stack from the crucial operations to the end-user and constitute a valuable framework for fraud detection in INFINITECH. Microsoft RA serves perfectly the Enterprise IT lifecycle of management and computing operations while simultaneously creates the next-generation customer experience through business productivity and cutting edge application and data services. Also, Microsoft RA's ability to manage and maintain huge volumes of critical transactions and processes as well as financial workloads only add value to the overall INFINITECH and how INFINITECH can exploit effectively those solutions. Diving deeper, the Microsoft RA provides Master Data Management (MDM), Data Quality Services (DQS) and pre-defined BI Semantic Metadata (BISM) which overlay BI capabilities delivered via pre-tuned datawarehouse configurations, near real-time analytics delivered through High-Performance technical Computing (HPC) and Complex Event Processing (CEP). Microsoft RA assets and commodities adopt jointly a high integration, which increases the overall architectural value and process efficiency especially when it comes to anti-money laundering. In this manner, components integration matrix costs and time costs of implementation and maintenance are efficiently reduced for any business through the highly-integrated Microsoft platform's optimal attributes exploitation. Overall, Microsoft RA is a leading provider of Data and Business Intelligence platforms providing a robust, end-to-end data, analytics and collaboration platform and constitutes a valuable part of the complete INFINITECH.

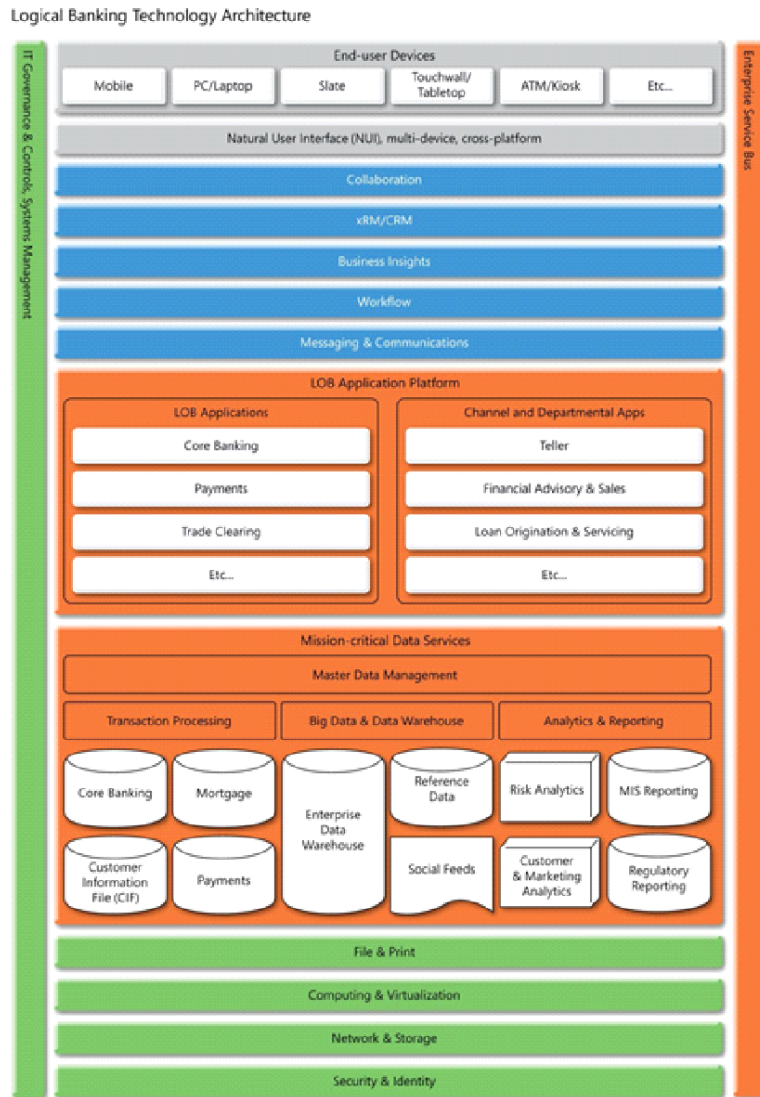


Figure 5 Microsoft RA Logical Banking Technology Architecture⁸

Anti-Money Laundering

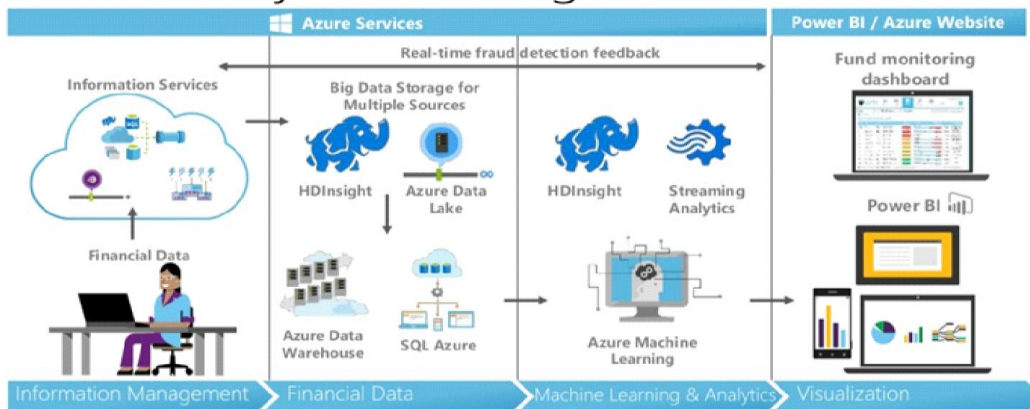


Figure 6 - Microsoft RA Anti-Money Laundering Scheme⁹

⁸ <https://news.microsoft.com/download/presskits/msfinancial/docs/MIRAB.pdf>

⁹ <https://www.slideshare.net/HadoopSummit/big-data-application-architectures-fraud-detection>

Relevance to INFINITECH RA

The Microsoft RA suggests a Logical Banking Technology Architecture and Anti-Money Laundering Scheme for the financial industry. Fraud detection as well as anti-money laundering are of utmost significance for the finance business.

INFINITECH is destined to address and eliminate fraud and money laundering in the financial sector through plenty of the specially designed pilots along with the overall solution of Testbeds and Sandboxes. Thus, Microsoft RA is relevant to INFINITECH RA as it focuses on organizing and sustaining massive volumes of transactions along with robust functionalities in bank datacenters. The following Microsoft RA elements are considered eventually related to the INFINITECH RA:

- **Information Management:** INFINITECH will implement information and data management of various financial corporations, such as banks, financial institutions, insurance-tech and fintech organizations. Process mining is destined to be of the fundamental capabilities in the overall INFINITECH. This Microsoft RA component (Anti-Money Laundering Scheme) should be taken into account in the context of designing and developing information services for the secure and scalable access of the type of data mentioned.
- **Financial Data:** In a similar manner, INFINITECH should consider thoroughly, for the same reasons, this Anti-Money Laundering layer in the context of storing this kind of financial data.
- **ML & Analytics:** INFINITECH will establish a variety of data analysis, including parallel, incremental, predictive and declarative real-time analytics. This layer of Microsoft RA should be definitely examined when preparing and constructing the analytics in Big Data, IoT and AI fields of INFINITECH.
- **Visualization:** INFINITECH will be using visualization tools and dedicated software components in order to extract the highly important information gained from the INFINITECH Analytics and conceive it in a human- and conclusion-friendly environment. This layer should be taken into account while assembling and formulating the visualization related INFINITECH modules.

3.2.1.3 WSO2 Architecture¹⁰

WSO2 offers a modular platform that enables the implementation of connected applications for the financial service industry. The philosophy of the platform is to divide complex systems into simpler individual sub-systems that can be more flexibly managed, scaled and maintained. It emphasizes flexibility given the need to comply with a rapidly changing landscape of business and regulatory requirements. From an implementation perspective, it comprises various applications, data management systems and toolkits.

The platform architecture is illustrated in the following *figure*. It comprises a number of operational systems that feed a data warehouse to enable analytical processing and data mining. On top of the data warehouse a number of enterprise applications are implemented including accounting applications and reporting applications.

¹⁰ <https://wso2.com/whitepapers/connected-finance-reference-architecture/>

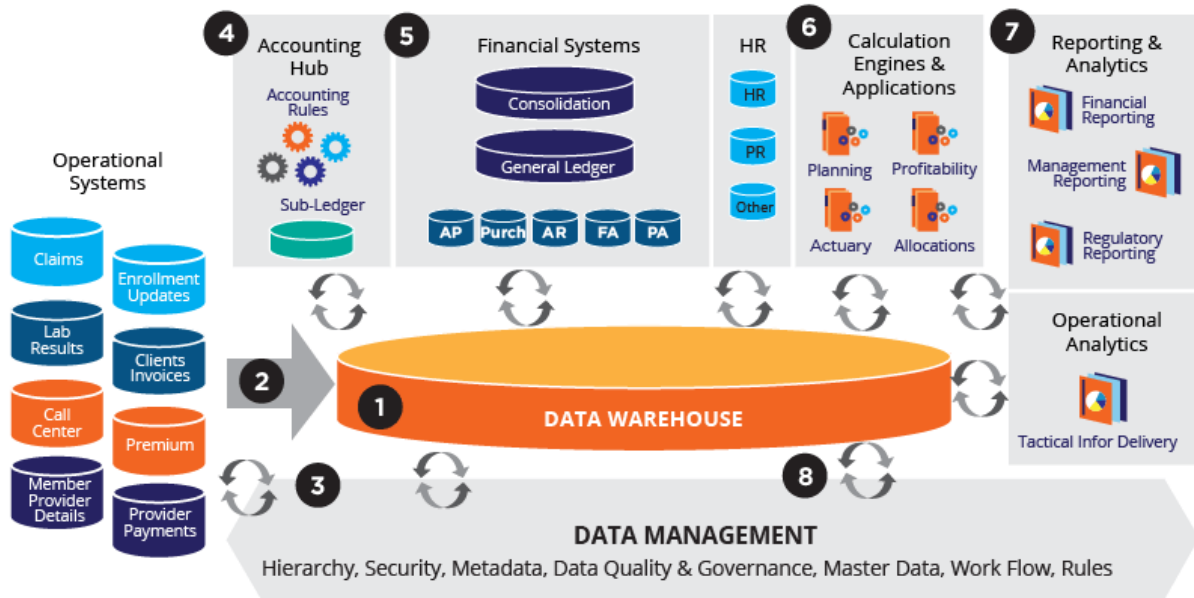


Figure 7 - A next-generation financial infrastructure requirements¹¹

WSO2 sustains connected architectures in the financial domain through its determined professional and sophisticated product. Figure 5 illustrates how state-of-the-art business needs match into a connected financial reference architecture. In particular, WSO2 Private PaaS are eligible to be used in the data layer and the identity layers as well as in many financial clouds. WSO2 User Engagement Server qualifies as a builder of analytics needs, thus, it constitutes an essential factor on the different analytics automations that are needed in INFINITECH, while WSO2 Data Analytics Server aids in real-time processing of the stored statistics and data. Moreover, WSO2 Message Broker suits perfectly into managing communications under the business-specific financial protocols while WSO2 Business Process Server handles all the workflow processes which will be extensively required in INFINITECH since the big number of businesses participating while WSO2 DSS is ideal for data integration. WSO2 Applications Server suits for service development with defined standards, and WSO2 API Manager exposes the appropriately configured and developed financial APIs. Completely, WSO2 ESB constitutes the fundamental integration component rendering both latter components while being essential for the co-existence of standard, modular and configurable testbeds and sandboxes of INFINITECH. Lastly, the WSO2 middleware platform is designed in favour of specific standards related to electronic exchange of insurance data that regulate the (insurance) industry and is developed by the Association for Cooperative Operations Research and Development (ACORD), the global industry’s nonprofit standards developer. In this manner, INFINITECH will benefit by the professional standards of the insurance industry. Overall, WSO2 is presenting a valuable solution and creates an abundance of opportunities for INFINITECH.

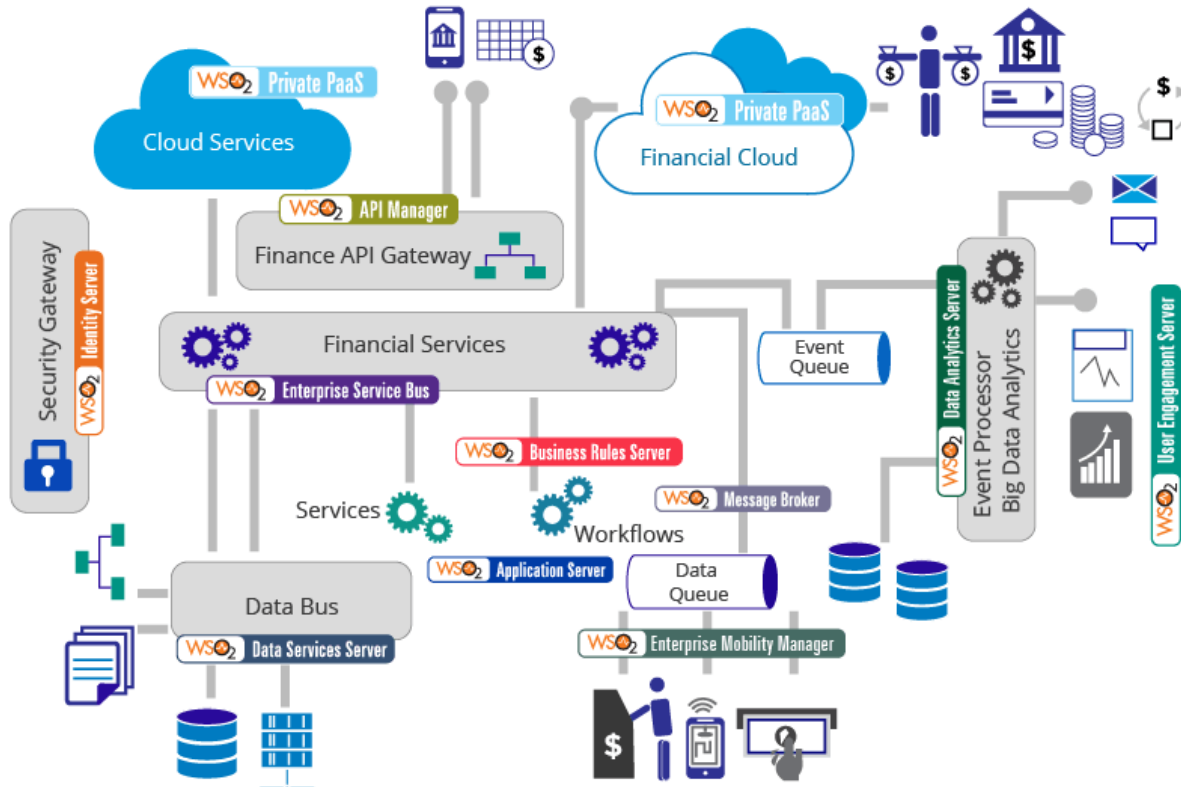


Figure 8 - WSO2 Ecosystem Solution¹²

Relevance to INFINITECH RA

The WSO2 Ecosystem accomplishes a sufficient solution for establishing an interconnected business followed by the necessary standards, components, and toolkits. The WSO2 RA upholds expertise when it comes to integration of internal components and it allows the performance of both in-house and foreign systems, while it manages the system components in an efficient and optimal way.

INFINITECH promises to regulate, protect and evaluate advanced financial information and data in the finance industry over a variety of deliberately tailored services and components. WSO2 RA relates to INFINITECH RA as it focuses on integration of multiple and dissimilar modules, such as analytics, real-time data processing, financial clouds, and data exchange in the finance industry. The following WSO2 Ecosystem modules are determined highly relevant for the INFINITECH RA:

- **WSO2 Private PaaS:** INFINITECH is by definition focusing on the financial industry, thus, the WSO2 Private PaaS module should be evaluated as it contributes a good fit for many financial clouds.
- **WSO2 User Engagement Server:** INFINITECH will exploit a collection of analytics technologies. This WSO2 module should be considered for the INFINITECH analytics requirements types.
- **WSO2 Business Process Server:** INFINITECH will engage in a complex and intensive workflow due to the rich total of businesses participating. The WSO2 examined here, should be taken into account when organizing and cultivating INFINITECH workflow of processes and procedures.
- **WSO2 ESB:** INFINITECH will ultimately require essential module integration in the holistic approach of the project. WSO2 ESB could constitute a vital component when evaluating cooperation of the standard, modular and configurable testbeds and sandboxes of INFINITECH.
- **WSO2 middleware platform:** INFINITECH platform design would require specific standards related to electronic exchange of insurance data, thus, ACORD’s design of the WSO2 middleware platform should be taken under consideration in this manner.

¹² <https://wso2.com/whitepapers/connected-finance-reference-architecture>

3.2.1.4 HortonWorks

Hortonworks Data Platform (HDP)¹³ provides the Hadoop Stack for Enterprise which is a centralized, enterprise-ready platform for storage and processing of any kind of data. When combined with a NoSQL database (e.g. Couchbase), HDP creates a huge volume of business value and intelligence. The currently demanded requirements are covered by this combination of the two. They include accuracy, in the context of bringing superior and precise analytic insights to Hadoop, and scalability, in the sense of providing comprehensive support from data-to-decision in order to boost the Hadoop’s value across the business. Also, operational performance, in a manner of superior performance and high scalability data access across a range of business use cases, as well as economics, in the context of driving bottom-line benefits by boosting the value of operational and analytics infrastructure while reducing Total Cost of Ownership, are included.

At a glance, state-of-the-art big data is constituted of the operational part (mostly NoSQL logic) and the analytical one (mostly Hadoop logic). Under the co-existence of both in finance, i.e. NoSQL together with HDP, several new scenarios are born. For instance, deep analytics when pulling data from Couchbase into Hadoop or training machine learning models and then cache them in Couchbase. Also, possible is now the integration with software like Sqoop, Kafka and Storm under the umbrella of Spark. Overall, the mentioned analytics have taken the form of Big Data Analytics as shown in *Figure 6*, and consist of three main phases: Data Pooling & Processing, Business Intelligence and Predictive Analytics.

Big Data helps Enterprises drive Business Efficiency and Profits

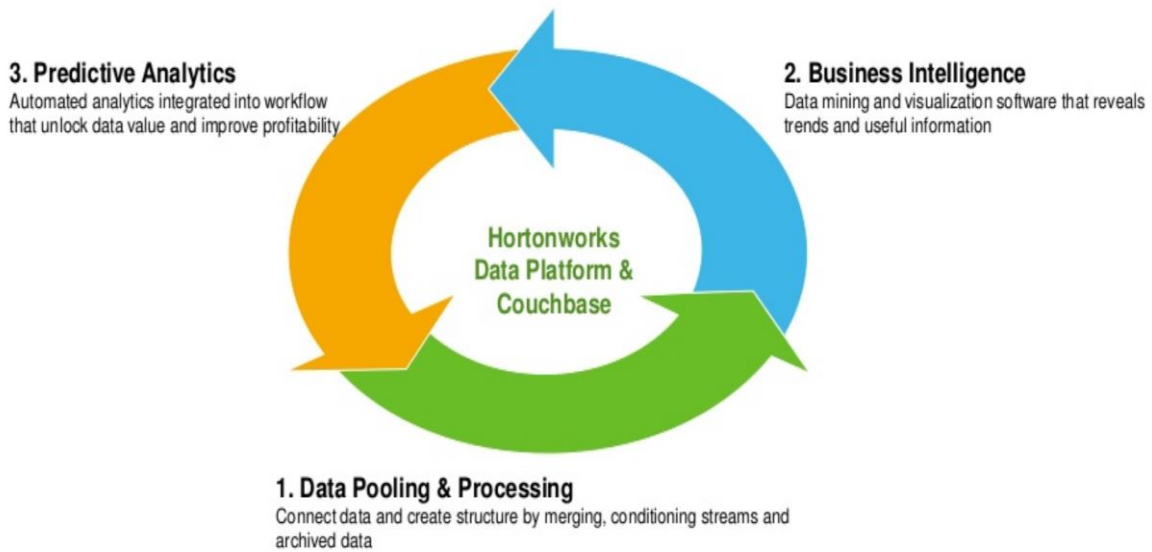


Figure 9 Big Data Adoption Lifecycle¹⁴

In the finance sector, Big Data has transformed traditional banking. Internally, where the customer is unaware of, there are dissimilar pieces that form the whole Big Data picture. Risk management, security, compliance, digital banking, fraud detection and anti-money laundering constitute the essence of state-of-the-art banking which is perfectly built upon the HDP principles when combined with NoSQL (e.g. Couchbase). For instance, in a media and entertainment use case, *Figure 7* shows how HDP and Couchbase can be used in order to integrate an Operational Data Store (ODS) and analytics capabilities in an enterprise environment.

¹³ <https://www.slideshare.net/Couchbase/couchbase-and-hadoop-couchbase-live-new-york-2015>

¹⁴ <https://www.slideshare.net/Couchbase/couchbase-and-hadoop-couchbase-live-new-york-2015>

In a similar manner, several concepts of the Hortonworks architecture could offer INFINITECH a plethora of benefits related to leading big data and artificial intelligence technologies. In particular, constant, and interoperable access to data lakes and NoSQL stored information, as well as integrated querying of data streams could be conveniently exploited in INFINITECH. Also, real-time analytics with open AI algorithms together with open finance APIs will be of immediate interest in INFINITECH as part of the regular standard mechanisms where the complete testbeds and sandboxes will be built on.

Relevance to INFINITECH RA

Hortonworks RA proposes a valuable Big Data Analytics. State-of-the-art big data analysis is crucial for any type of enterprise. Hortonworks is capable of satisfying big data needs and requirements of various use cases of dissimilar nature and prerequisites.

INFINITECH intends to make use of Big Data Analytics extensively and widely inside its core and surface structure. Hortonworks RA is absolutely relevant to INFINITECH's RA as it serves and reassures the needs and challenges which originate from big data analysis. The following Hortonworks elements are related to INFINITECH's RA:

- **Data Pooling & Processing:** INFINITECH will utilize big data for analysis, storage and distribution broadly, hence this first phase of Hortonworks Big Data Analytics should be taken into account when constructing, implementing and applying functions to INFINITECH data such as connecting, merging, conditionally streaming and archiving.
- **Business Intelligence:** In the context of big data analysis in INFINITECH, this second phase of Hortonworks Big Data Analytics should be considered when creating and establishing data operations such as result- and trend-driven mining and visualization.
- **Predictive Analytics:** INFINITECH will adopt different analytics processes among different scopes and pilot visions. Hortonworks third phase Big Data Analytics should be eligible when designing and composing automated result- and profit-driven integrated analysis into INFINITECH workflows.
- **Operational Data Store (ODS) & Data Lake:** INFINITECH will require storage capabilities and functionalities of big data scale. Hortonworks ODS and Data Lake should be taken under consideration when forming and building INFINITECH repositories and storehouses.

3.2.2 Reference Architectures from Standardization Bodies and Industry Associations

3.2.2.1 BDVA Reference Model

Big Data Value Association (BDVA) has presented a **reference model** for big data [BDVA]. The model has horizontal layers encompassing aspects of the data processing chain, and vertical layers addressing cross-cutting issues (e.g. cybersecurity and trust).

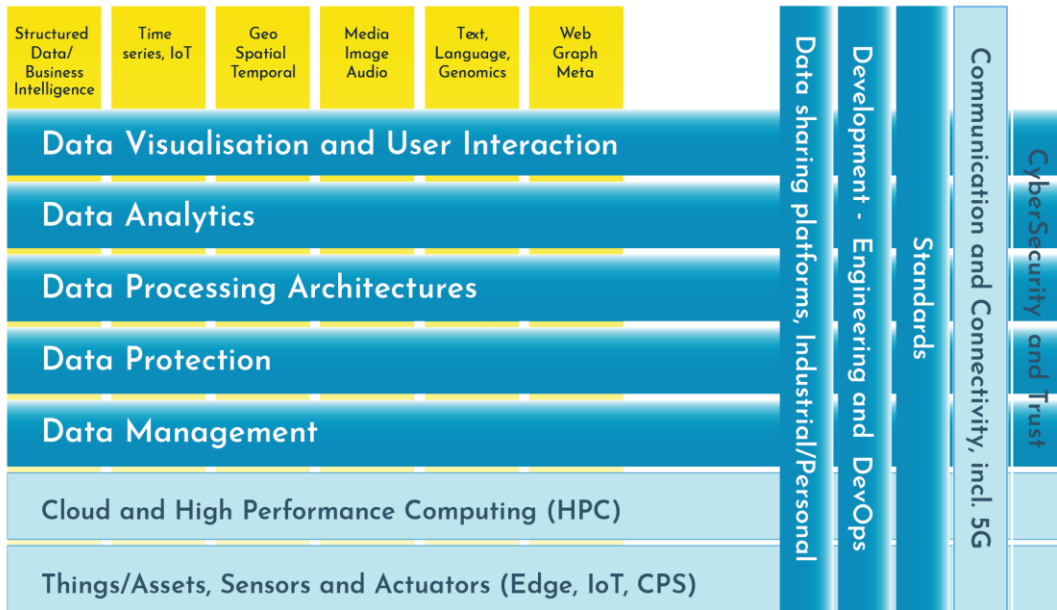


Figure 10 BDVA – Reference Model

The BDV Reference Model is structured into horizontal and vertical concerns.

- **Horizontal concerns** cover specific aspects along the data processing chain, starting with data collection and ingestion, and extending to data visualisation. It should be noted that the horizontal concerns do not imply a layered architecture. As an example, data visualisation may be applied directly to collected data (the data management aspect) without the need for data processing and analytics.
- **Vertical concerns** address cross-cutting issues, which may affect all the horizontal concerns. In addition, vertical concerns may also involve non-technical aspects.

Relevance to INFINITECH RA

BDV Reference Model (BDV RM) is not a Reference Architecture in the IT sense. However, the horizontal and vertical “concerns” of BDV RM are present in INFINITECH as layers, cross-cutting and data modelling sections.

The BDV Reference Model serves as a common framework to locate Big Data technologies on the overall IT stack. It addresses the main concerns and aspects to be considered for Big Data Value systems. In the same way, the IRA consider a 1:1 correspondence with the horizontal “concerns” treated as layers in the INFINITECH RA conceptual schema.

In the same way that the BDV RM is not a layered architecture, as it does not impose hierarchy over sections or interdependencies in elements, IRA components can be freely pipelined (orchestrated) depending on the specific problem they are called to solve along other components. The most important consideration in IRA is that the all “concerns” are taken into account on any given solution. In the world of INFINITECH sandboxes, data components of each specific layers should be present as they add not only a level of specific processing but they assure that the specific problem is addressed (e.g. anonymization in data security, or user interface).

INFINITECH RA is largely influenced by the BDV Reference Model as it can be seen by this comparison figure. This will be made more evident in following Sections that introduce the INFINITECH-RA and its alignment to BDVA layers.

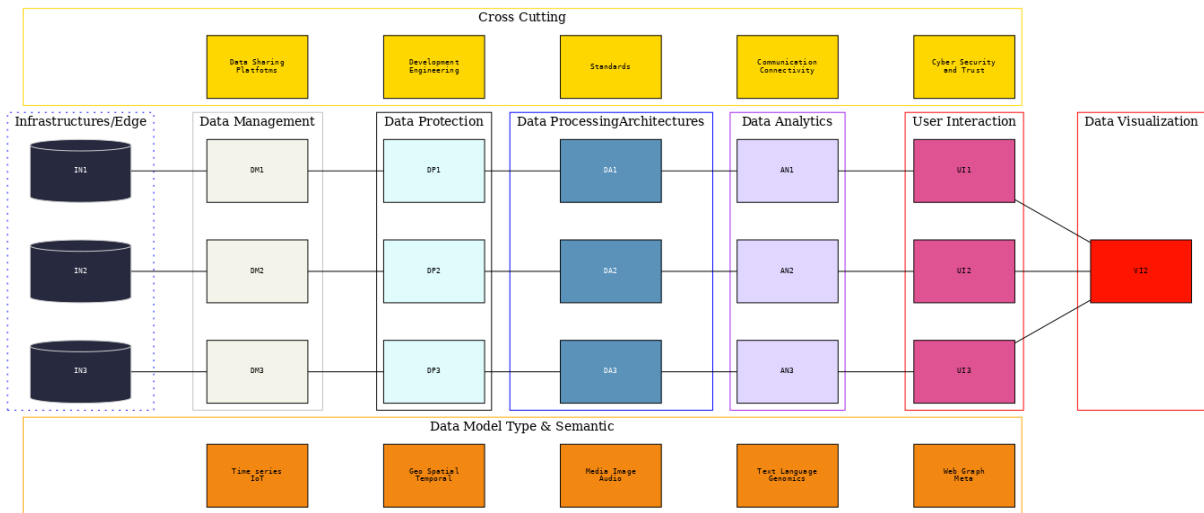
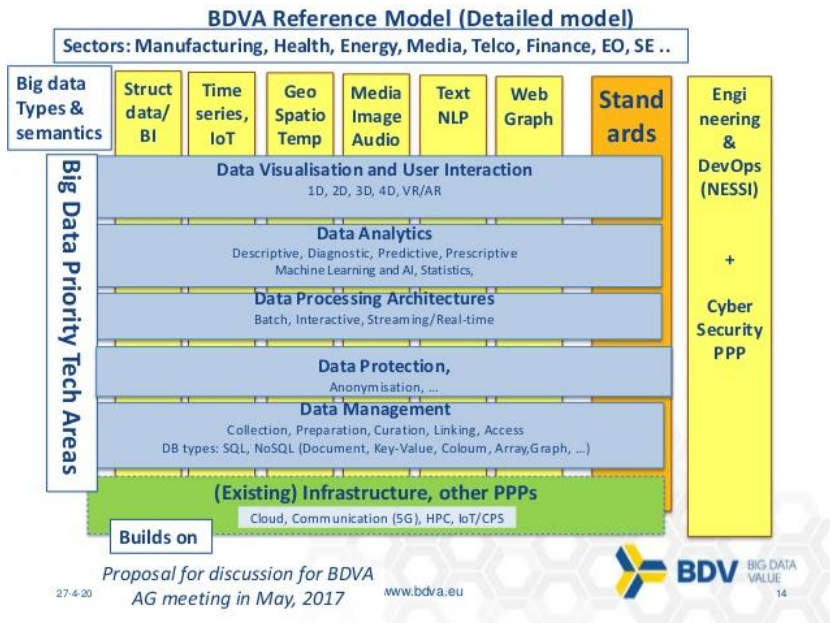


Figure 11 - BDV – Reference Model (top) vs INFINITECH Reference Architecture (bottom)

3.2.2.2 NIST Big Data Reference Architecture (NBDRA)

National Institute of Standards and Technology (NIST) Big Data Program is developing a big data RA [Boid18]. The conceptual model is comprised of five functional components: data producer/consumer, system orchestrator, and big data application/framework provider. Data flows, algorithm/tool transfer, and service usage between the components can be described with different types of arrows. Activities and functional component views of the RA can be used for describing a big data system, where roles/subroles, activities, and functional components within the architecture are identified.

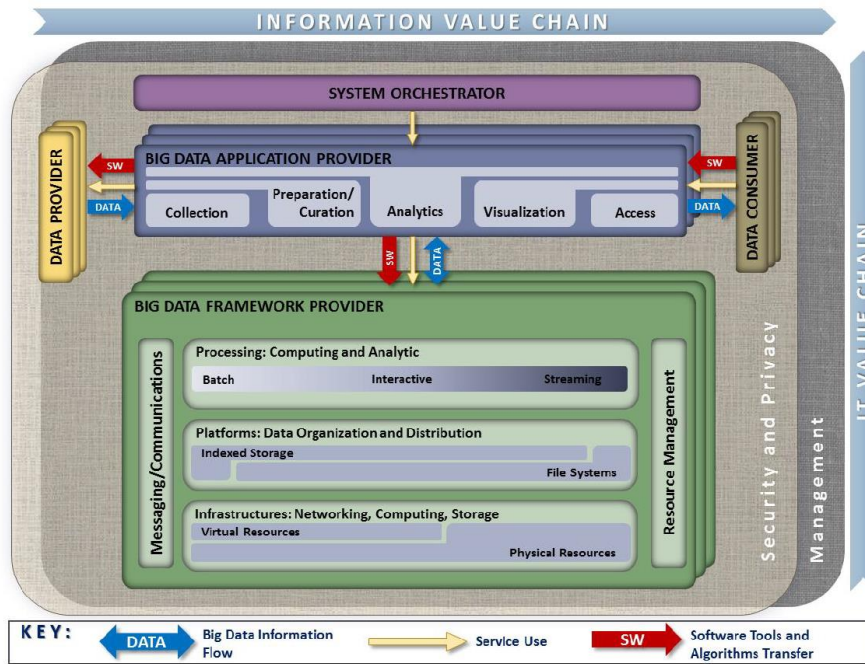


Figure 12 – NIST Reference Architecture

Relevance to INFINITECH:

NIST approach is important for the INFINITECH RA as it identifies also roles and functional components. NIST is also considering the challenges in Big Data Architectures (as depicted in Figure 13: NIST BigData Architecture) which are very relevant to INFINITECH as well.

Big Data Architecture and Infrastructure – Challenges (Computing Stack)

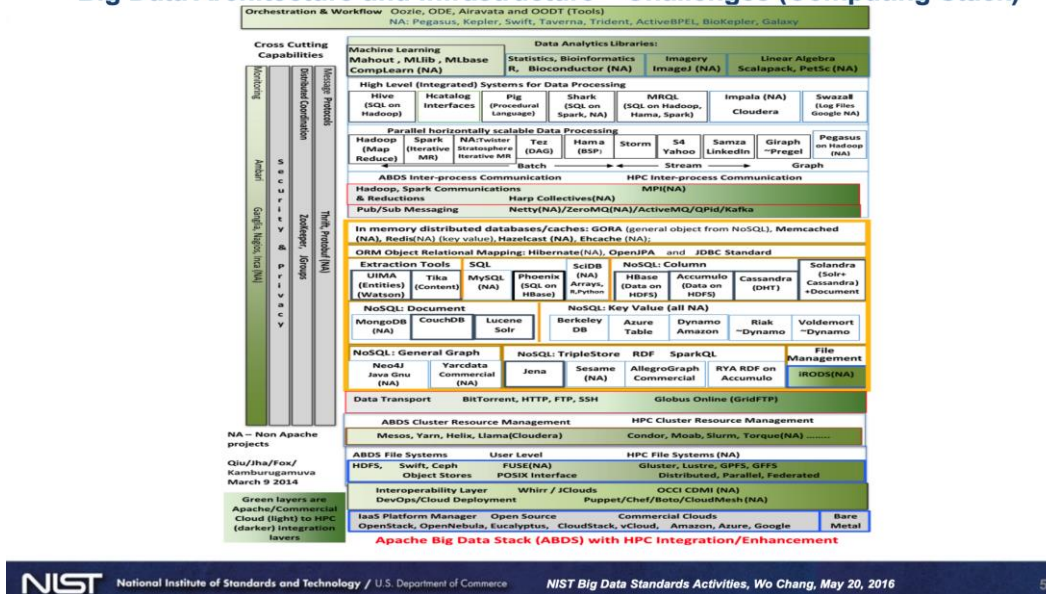


Figure 13: NIST BigData Architecture and Infrastructure Challenges

3.2.2.3 Industrial Internet Reference Architecture (RA)

Internet of Things constitutes a fundamental territory of the INFINITECH Initiative since a vital part of the focus will be invested in it; thus, the Industrial Internet Reference Architecture (IIRA) will provide important value to INFINITECH as explained later in this section. The IIRA relates directly to Industrial Internet of Things (IIoT) systems as the latter establish the basic collections of IIRA interest. In

particular, IIRA determines and describes the IIoT most critical architecture issues in the current industry while arranging them into the suitable viewpoints along with their respective stakeholders. Afterwards, the IIRA attempts to resolve these issues in the different viewpoints through efficient analysis by generating definite abstract architecture representations. With regard to INFINITECH, such an abstract structure will serve substantially in the arranging foundations of the project, as far as organization, alignment and management of the dissimilar existing formulations of IIoT are concerned. In *Figure 9*, there is an illustration of the key ideas about the IIRA’s architectural components and applications.

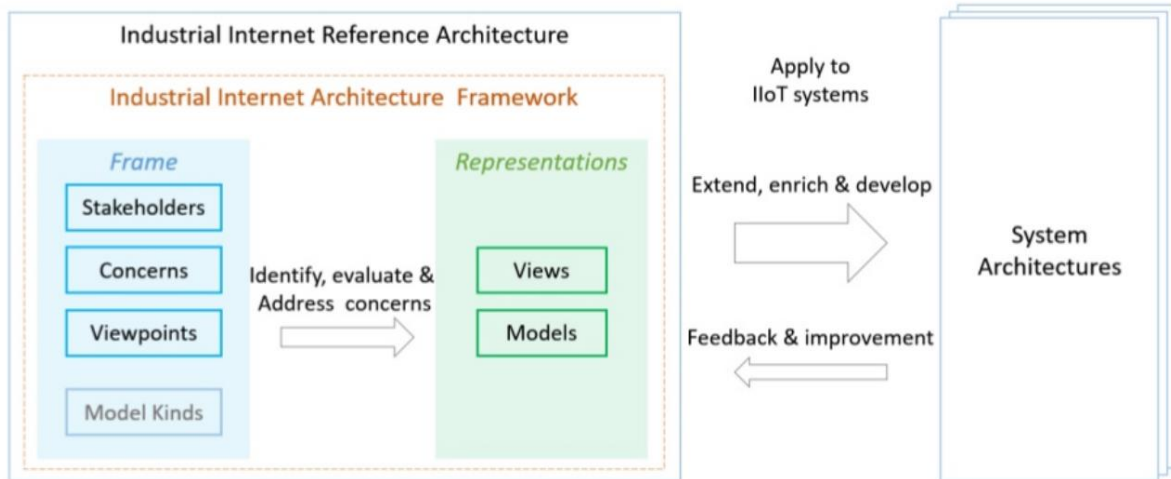


Figure 14 IIRA constructs and application¹⁵

The IIRA’s abstraction level omits architectural components who are assessing under concrete systems particularities. Thus, IIRA does not specify all the architectural components (constructs) of *Figure 9*, rather it adapts the ISO architecture specification (ISO/IEC/IEEE 42010) with minor rearrangements in two features as described in the following paragraphs.

The first feature is that IIRA does not particularly determine the model kind as a key component of the whole framework; rather it roughly uses this idea during the concern analysis in the abstract representation development.

The second feature is that IIRA does not particularly consider specific architectural components, such as Correspondence Rules, Correspondence and Architecture Rationale; it simply allows the concrete systems development to handle them.

The models and their views’ representations in the IIRA are carefully selected in this manner since they display the respective concerns at the corresponding level of abstraction and simultaneously determine the key ideas of this reference architecture. Nevertheless, they are not the only models and views for addressing concerns in the respective viewpoints or appropriate for a production system. These views can help initiate a concrete architecture, and then can be enhanced properly towards the specifications of the required IIoT system. Architecture developers that are eager about the ISO/IEC/IEEE 42010 Architecture Description for their concrete systems, can use the fundamental framework of IIRA in the beginning, then augment and upgrade the components offered that are applicable to the their individual needs and finally, develop those components that are outside the IIRA in the context of the architectural operation.

The architectural description mentioned above will certainly increase the ability to regulate and supervise plentiful modules, as well as untap the complete potential of component development and

¹⁵ <https://www.iiconsortium.org/pdf/IIRA-v1.9.pdf>

rearrangement regarding IoT in INFINITECH. With an INFINITECH vision to a large but simple architecture of numerous requirements, IIRA will help INFINITECH to inspect its IoT and other needs from a broader and wider perspective, and simultaneously deliver and address productive, beneficial and profitable solutions with competence and dexterity on the field of architectural description and design approach.

Relevance to INFINITECH RA

The IIRA recommends a fundamental framework of abstract architecture design and deployment. IIRA uses a standard way to apply the Industrial Internet Architecture Framework (IIAF) capabilities to IoT systems and extend, enrich and develop system architectures.

INFINITECH consists of various dissimilar and independent ideas while needs IoT technology integration and management. IIRA can be relevant to INFINITECH as it focuses on the conceptualization of different notions and components and is directly related to IIoT. The following IIRA aspects are suitable for acknowledged as applicable and relevant for the INFINITECH RA:

- **IIoT systems:** INFINITECH will require organization, adjustment and administration of disparate formulations of IoT. IIRA should be considered when defining and constructing the IoT fundamentals of INFINITECH.
- **Abstract architecture representations and abstract structures:** INFINITECH conceptualization should be performed with discretion, prudence, and experienced organizing. IIRA abstract structure should be taken into account when conceptualizing INFINITECH’s architecture.

3.2.3 Architectures of Relevant EU Projects and Research Initiatives

3.2.3.1 H2020 BigDataStack

The BigDataStack project focuses on the provision of an architecture that drives resource management decisions based on data aspects, such as the deployment and orchestration of services [Kyriazis18]. A relevant architecture is presented in the following figure, including the main information flows and interactions between the key components.

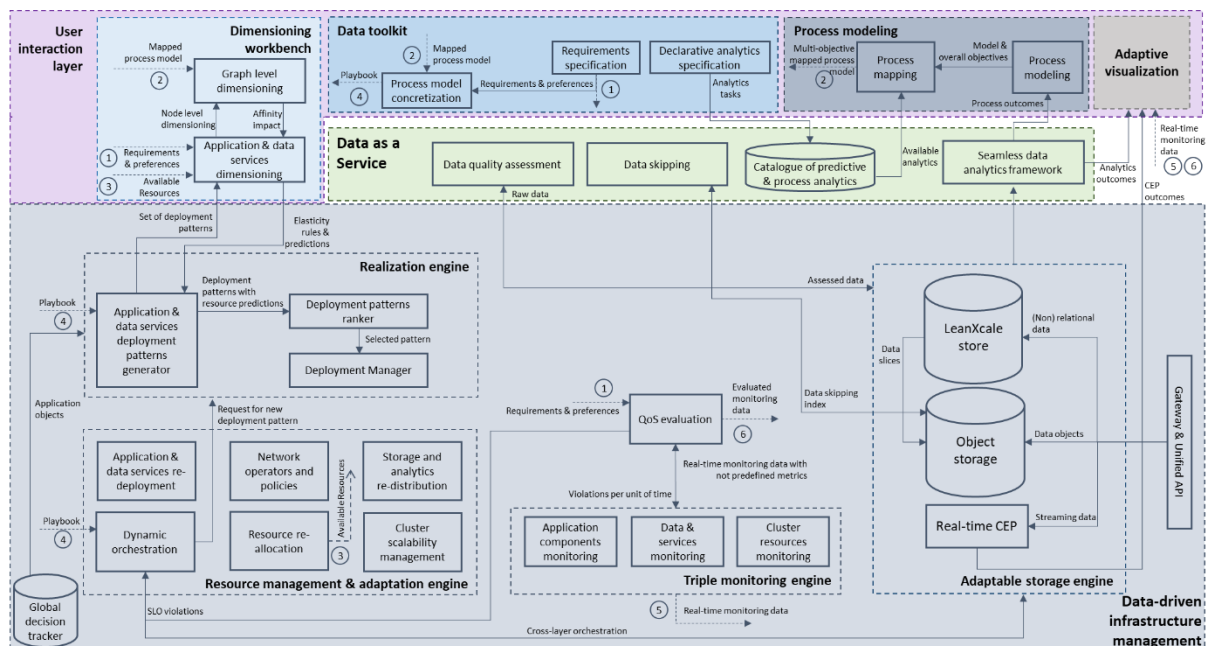


Figure 15 - BigDataStack architecture model

As cited in the figure, raw data are ingested through the *Gateway & Unified API* component to the *Storage* engine of BigDataStack, which enables storage and data migration across different resources. The engine offers solutions both for relational and non-relational data, an *Object Store* to manage data as objects, and a CEP engine to deal with streaming data processing. The raw data are then processed by the *Data Quality Assessment* component, which enhances the data schema in terms of accuracy and veracity and provides an estimation for the corresponding datasets in terms of their quality. Data stored in *Object Store* are also enhanced with relevant metadata, to track information about objects and their dataset columns. Those metadata can be used to show that an object is not relevant to a query, and therefore does not need to be accessed from storage or sent through the network. The defined metadata are also indexed, so that during query execution objects that are irrelevant to the query can be quickly filtered out from the list of objects to be retrieved for the query processing. This functionality is achieved through the *Data skipping* component of BigDataStack providing the relevant data skipping objects. Furthermore, the overall storage engine of BigDataStack has been enhanced to enable adaptations during runtime (i.e. self-scaling) based on the corresponding loads. The latter is of major importance for INFINITECH given that the LeanXcale database will also be exploited in the project and the provided self-scaling capabilities will ensure the seamless provisioning of services (i.e. analytics) towards the financial / insurance data intensive applications.

Given the stored data, decision-makers can model their business workflows through the *Process Modelling framework* that incorporates two main components: the first component is *Process modelling*, which provides an interface for business process modelling and the specification of an end-to-end optimisation goals for the overall process (e.g. accuracy, overall completion time, etc). The second component refers to *Process Mapping*. Based on the analytics tasks available in the *Catalogue of Predictive and Process Analytics* and the specified overall goals, the mapping component identifies analytics algorithms that can realise the corresponding business processes. The outcome of the component is a model in a structural representation e.g. a JSON file that includes the overall workflow, and the mapped business processes to specific analytics tasks by considering several (potentially concurrent) overall objectives for the business workflow.

Following, through the *Data Toolkit*, data scientists design, develop and ingest analytic processes/tasks to the *Catalogue of Predictive and Process Analytics*. This is achieved by combining a set of available or under development analytic functions into a high-level definition of the user's application. For instance, they define executables/scripts to run, as well as the execution endpoints per workflow step. Data scientists can also declare input/output data parameters, analysis configuration hyper-parameters (e.g. the k in a k -means algorithm), execution substrate requirements (e.g. CPU, memory limits etc.) as service level objectives (SLOs), as well as potential software packages / dependencies (e.g. Apache Spark, Flink etc.). The output of the *Data Toolkit* component enriches the output of the previous step (i.e. *Process Modelling*) and defines a BigDataStack Playbook.

The generated playbook (including operation sequence templates) is utilized by the *Application and Data Services Deployment Patterns Generator*. The component creates different arrangements (i.e. patterns / configurations) of deployment resources for each application and data service pod. To enable the latter, data from the *Global Decision Tracker* are obtained. These data refer to application objects that emerge following the registration of a new application via a playbook, resulting to the creation of an application object and a series of object templates representing the different parts of an application. Based on the above, a set of candidate deployment patterns are passed to the *Application Dimensioning Workbench*, along with an end-to-end optimization objective and the information on the available resources, in order to estimate resource usage and QoS performance prior to actual deployment. The primary output of the *Application Dimensioning Workbench* is an elasticity model, which defines the mapping of the input QoS parameters to the concrete resource needed (such as the number of VMs, bandwidth, latency etc.). These decisions are depended on data-defined models. Thus, based on the obtained dimensioning outcomes, deployment patterns are ranked by the *Deployment Patterns Ranker* and the optimum pattern is selected for deployment, making the concluding arrangement of services data-centric. The *Deployment Manager* administers the setup of the application and data services on the allocated resources. This is a key component that automated

and optimizes the deployment process and could be exploited in INFINITECH for the different deployment configuration of the pilots of the project.

During runtime, the *Triple Monitoring engine* collects data regarding resources, application components (e.g. application metrics, data flows across application components, etc.) and data operations (e.g. analytics / query progress, storage distribution, etc.). An advancement comparing to the previous version of the architecture is that these metrics are not predefined and are identified during runtime so as to optimize the metrics to be collected and thereafter their evaluation. The collected data are evaluated through the *QoS Evaluation* component to identify events / facts that affect the overall quality of service (in comparison with the SLOs set in the toolkit). The evaluation outcomes are utilised by the *Runtime adaptation engine*, which includes a set of components (i.e. cluster resources re-allocation, storage and analytics re-distribution, network operators and policies enforcement, application and data services re-deployment, and dynamic orchestration patterns), to trigger the corresponding runtime adaptations needed for all infrastructure elements to maintain QoS. It should be noted that the dynamic orchestration employs a reinforcement-based logic that leads to cross-layer orchestration and optimization addressing both the resources and the data services as depicted in the figure.

Moreover, the architecture includes the *Global decision tracker*, which aims at storing all the decisions taken by the various components. In this context, the tracker holds additional information such as application logging data, Candidate Deployment Patterns, QoS failures, etc. Thus, as a global state tracker, provides the ground for cross-component optimisation, as well as tracking the state and history of BigDataStack applications.

Finally, the architecture includes the *Adaptive Visualisation* environment, which provides a complete view of all information, including raw monitoring data (for resource, application and data operations) and evaluated data (in terms of SLOs, thresholds and the evaluation of monitoring in relation to these thresholds). Moreover, the visualization environment acts as a unique point for BigDataStack for different stakeholders, actors, thus, incorporating the process modelling environment, the data toolkit and the dimensioning workbench. These accompany the views for infrastructure operators (e.g. regarding deployment patterns).

Relevance to INFINITECH RA

H2020 BigDataStack shares similar concepts and approaches on what concerns the data lifecycle and the automation of the deployment of the components. Regarding the data lifecycle, they both treat data *at rest* and *streaming* data. Data are stored in different silos on premise of the application owner and need to be ingested inside the platform. This process takes place in BigDataStack when the data arrives to the *gateway* component and a data ingestion process takes place to migrate the data inside the data repositories of the project. Data can live inside the operational database of the platform or inside a data warehouse. BigDataStack provides a seamless analytical framework that exposes a central endpoint and a common query language to access data that might be resided in both stored, by hiding the complexity of joining the data from the two heterogeneous stores from the data user. Moreover, streaming data are also processing in the real-time using the Context Event Processing (CEP) component of the platform that also allows for querying the data at rest performing complex analytical query processing by ensuring data ensuring. This is ensured by the INFINITECH RA by exploiting the Declarative Real-Time Analytics framework that is offered by the data management layer. What is more, H2020 BigDataStack provides a set of different analytical tools that access the central repository in a seamless way. This is similar to the libraries of ML/DL algorithms that the INFINITECH RA provides that access data coming from the data management layer in a seamless manner, by additionally ensuring data interoperability between different datasets via the corresponding semantic interoperability framework of INFINITECH. Finally, the use of open source technologies for the automation of the deployment and the software lifecycle management in terms of fault tolerance and elastic scalability that the H2020 BigDataStack is relying on, is towards the same direction with INFINITECH notion of sandboxes.

3.2.3.2 H2020 BOOST

BOOST 4.0 aims to be an open standardised and transformative shared data-driven Factory 4.0, demonstrating how European industry can build unique strategies and competitive advantages through big data across all phases of the product and process lifecycle.

The main objectives of BOOST 4.0 are to establish 10 big data lighthouse smart connected factories, provide the RAMI 4.0 and IDS based BOOST4.0 open EU framework and governance model, for both services and data assets and put together methodologies, assets, models and communities in order to maximise visibility, mobilization, replication potential and impact.

The BOOST 4.0 Reference Architecture consists of a number of layers at its core, alongside with a Factory dimension and a manufacturing entities dimension. More detailed information about these layers can be found in the D2.5 – BOOST 4.0 Reference Architecture Specification v1 [BOOST4.0-D2.5]].

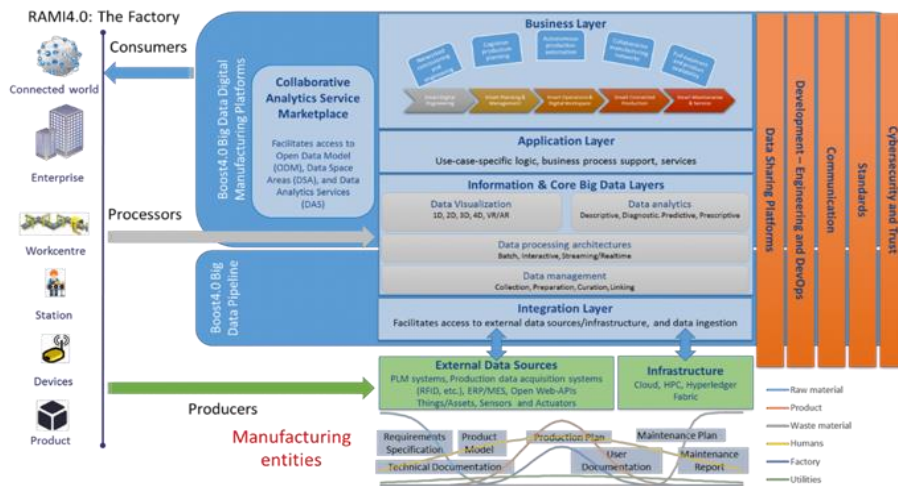


Figure 16 - The Boost 4.0 Reference Architecture

The core layers represent a collection of functionalities/components performing a specific role in the data processing chain, and consist of Integration layer, Information and Core Big Data layers, Application layer and Business Layer.

The Integration layer facilitates the access and management of external data sources such as PLM systems, production data acquisition systems, Open web API’s, and so on.

The Information and Core Big Data layer is composed of components belonging in four different sublayers:

- Data management groups together components facilitating data collection, preparation, curation and linking.
- Data processing groups together architectures that focus on data manipulation.
- Data analytics groups together components to support descriptive, diagnostic, predictive and prescriptive data analysis.
- Data visualization groups together algorithms components to support data visualization and user interaction.

The Application layer represents a group of components implementing application logic that supports specific business functionalities and exposes the functionality of lower layers through appropriate services.

The Business layer form the overall manufacturing business solution in the BOOST 4.0 five domains (networked commissioning & engineering, cognitive production planning, autonomous production automation, collaborative manufacturing networks, and full equipment & product availability) across

five process life-cycle stages (Smart Digital Engineering, Smart Production Planning & Management, Smart Operations & Digital Workplace, Smart Connected Production, Smart Maintenance & Service)

Alongside the core layers, there are a number of other cross-cutting aspects that affect all layers:

- Communication aims to provide mechanisms and technologies for reliable transmission and receipt of data between the layers
- Data sharing platforms allow data providers to share their data as a commodity, covering specific data, for a predefined space of time, and with a guarantee of reversibility at the end of the contract.
- Development-Engineering and DevOps covers tool chains and frameworks that significantly increase productivity in terms of developing and deploying big data solutions.
- Standards covers the different standard organisations and technologies used by BOOST 4.0
- Cybersecurity and trust cover topics such as device and application registration, identity and access management, data governance, data protection and so on.

Factory dimension (as defined in RAMI 4.0 (Reference Architecture Model Industry 4.0)) is used to identify the roles of various entities in the context of the Reference Architecture (producer, consumer and processor). Finally, the dimension of manufacturing entities allows the specification of such elements across the manufacturing process lifecycle stages and adds the relevant data types and ontologies for them. On the other hand, Manufacturing entities dimension provides a unified semantic model covering the multi-domain knowledge of the BOOST 4.0 pilots.

3.2.3.3 H2020 FINSEC

H2020 FINSEC project is targeted for cyber and physical security. The FINSEC platform RA aims to provide a general approach to the overall security of organizations, especially in the Financial Sector. The FINSEC RA platform is devised to be above the infrastructure (both IT and Physical) of a typical use case, e.g. of a company organization which already has its own infrastructure and applications and procedures to monitor the IT assets from a cyber-security standpoint. The following figure illustrates the actual design of the FINSEC RA¹⁶

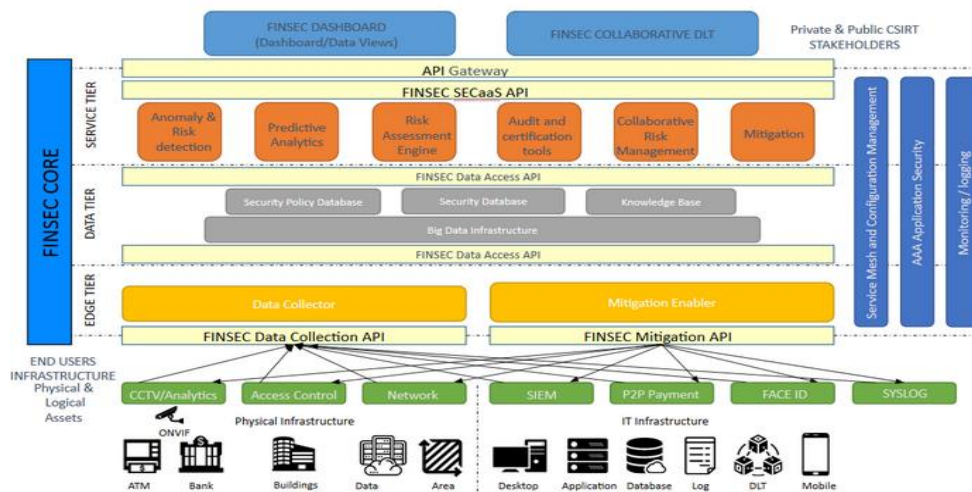


Figure 17 - The FINSEC Reference Architecture version 2.0

The FINSEC RA can be viewed as an N-tier architecture composed by several layers:

¹⁶ <https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5ce3a941d&appId=PPGMS>

- **Field Tier** - data from the physical and logical assets to be protected against threats. For example, CCTV analytics and SIEM are involved in this layer to give useful information about potential attacks to the upper tiers.
- **Edge Tier** - Data Collection module, which is necessary to filter the needed information during their flow towards the upper levels. The Actuation Enabler is responsible to allow some actions requested from the upper layers to be performed onto the probes, such as the shutdown of a server in case of threat or the closure of an automatic door of a protected room.
- The **Data Tier** is where information is stored and is organized into three different storage infrastructures, providing consisting data access API to all other modules.
- The **Service Tier** is where the kernel applications of the FINSEC and the security toolbox will be running, also exposing the functionality to external entities.
- The **Presentation and Communication tier** offer interfaces to the rest of the world. This tier is where the dashboard will be provided to monitor data and where assets and the FINSEC Collaborative Module will be available to share information with other FINSEC instances.

Relevance to INFINITECH RA

The FINSEC RA, is an example of layered bottom-to-top hierarchical IT architecture, suited for data flow from field to presentation and communication layers of applications. Data can flow also in other directions, especially from Analytics and Mitigation back to the fields with information that can transform also the data and the underlying infrastructure. Although the scenarios and the use cases of FINSEC are very different from INFINITECH, the two worlds share the same Financial and Insurance business sectors and therefore most of the constraints and regulatory frameworks apply in both architecture reference solutions.

Furthermore, the FINSEC -RA makes provisions for decentralized data exchange between stakeholders of the financial supply chain. It promotes a novel decentralized data sharing approach based on a blockchain infrastructure. As such it also provides a vision for exploiting blockchain technology in the finance sector beyond conventional blockbuster cryptocurrencies (e.g., BitCoin, Ethereum). INFINITECH could build on this approach towards supporting data sharing and data exchange in a broader range of use cases, beyond security data exchange.

3.3 Overall INFINITECH-RA Positioning

Overall, the INFINITECH-RA will be built on several concepts that have been introduced by other reference architecture models, including models developed by industrial organizations (e.g., large IT enterprises) and associations (e.g., BDVA), as well as models produced by other projects. Moreover, it considers the structuring principles of BigData and AI applications in digital finance, as these principles are reflected in industrial reference architectures. Nevertheless, the INFINITECH-RA aims at introducing a more flexible approach: Instead of providing a rigorous (but monolithic) structure of BigData/AI applications, it defines these applications as collections of data-driven pipelines. The latter are built based on INFINITECH components, notably the components specified in D2.5 of the project. Hence, the INFINITECH-RA provides a number of layered architectural concepts and a rich set of digital building blocks, which enable the development of virtually any BigData or AI application in Digital Finance. The project will offer increased flexibility in defining data-driven applications in the sector, in a ways that covers'/subsumes most of the rigorous architectures outlined in earlier paragraphs.

4. The INFINITECH Reference Architecture

The purpose of a Reference Architecture is to provide a conceptual and logical schema for solutions to a large class of problems. In the INFINITECH project the domain is as vast as the Financial and Insurance Sectors where most of the applications are data-driven. The class of problems of the project (pilots and use cases) and in general the service management of financial institutions and insurance companies are largely based on data that should be managed in the safest and protective way.

In these domains customers' enormous data sets must be processed to derive information with the purpose to provide better and more competitive services respecting the complex and sometimes conflicting regulatory frameworks such as privacy, security, interoperability, etc.

Therefore, a reference architecture should have explored in advance the specific domains in which the class of problems must find solutions providing a general model to which stakeholders (end-users, business owners, designers, data scientists, developers, maintainers etc.) can refer for best practices in the specific problem-solution space.

In information technology, a RA can be used to check solutions to a particular problem in that class against the best practices and specific technologies. The INFINITECH RA is no exception and it is the result of the analysis of the significant number of use cases in the project's pilots, their requirements (users' stories) and constraints (regulatory, sector and technological) as well as the state-of-the-art technologies and similar architectures.

It is important to state what is the RA in the INFINITECH project:

- A set of views for the Logical, Process, Development and Physical implementations
- A set of common scenarios referring to generic use cases
- A way to verify the use cases' scenarios and solutions
- A way to speak the same language among stakeholders
- A way to leverage solutions referring to best practices and building blocks
- A way to verify if constraints in requirements, regulatory, technical and logical have been addressed properly

It is also important to state what the INFINITECH RA is NOT:

- A ready-to-deploy technological IT framework
- A rigid and unmutable set of connecting building blocks
- A set of mandatory rules for development and integration
- A manual for implementation and rollouts
- A one-size-fits-all recipe for all business cases

The INFINITECH RA, is largely based on the concept of Data Pipelines used in Data Science and Data Manipulation. This model will allow to consider the blocks which make the solution by-design comply with rules and constraints. In that respect, ML and DL are emerging as key technologies to be developed, tested and evaluated in their specific domains. However, the utilization of machine learning (ML) as part of the computing infrastructure is still an area for further research [8]. Therefore, at the heart of the project as well as a pillar concept of the RA is to leave great flexibility in the data management workflow to experiment with different solutions.

Particularly, the RA and the workflows based on it should be used to understand how data can be collected, and how models and technologies should be developed, distributed, and deployed for decision making purposes in the specific use cases.

Therefore, the goal of this deliverable is to provide a RA for facilitating the design of the INFINITECH Pilots' workflows which encourages the utilization of standard data-processing components and ML techniques as well as big data architectures and infrastructure.

Another important fact, it has been addressed how development and deployment of ML models may be related to actors (e.g. Data scientist, ML engineer) at the business layer.

In general, the value of RAs can be summarized in the following points:

- reduction of development and maintenance costs of systems
- facilitation of communication between important stakeholders
- and reduction of risks

Also, when a system is designed without a RA, an organization may accumulate technical risks [11] and end up with a complex and non-optimal implementation architecture.

The presented reference architecture (RA) has been created inductively based on the 15 Pilots logical workflows. The RA has been extended based on an earlier published RA for big data systems [4]. Particularly, development and deployment views has been focused on in big data computing environments.

High level and deployment environment views of the RA may facilitate architecture design of workflows in the future.

The following sections contain the different views of the INFINITECH Reference Architecture. It is inspired by BDVA Reference Model and other logical architectures for big data presented in the section related to the State of the Art.

The INFINITECH Reference Architecture or IRA for short, is also derived from the use cases defined in the project and provides a general framework for the solutions to be developed and deployed in the pilots.

Most of the logical views of IT architectures, and not only in H2020 projects, present schemes with links and relations between the different building blocks from which it should be evident to understand the logical flow of data and the interdependencies between blocks.

However, most of these schemes impose rigid mechanisms to how data can be managed which in real world scenarios complicates the solution instead of simplifying it. What data scientists want is to experiment with data (especially big data) in a flexible way to play with algorithms and technologies to select the best combination to solve real use cases, like those proposed in INFINITECH.

The presented reference architecture (RA) has been created inductively based on the 15 Pilots logical workflows The RA has been extended based on an earlier published RA for big data systems. Particularly, development and deployment views has been focused on in big data computing environments.

High level and deployment environment views of the RA may facilitate architecture design of workflows in the future.

4.1 Methodology: Architectural View Model

The INFINITECH Reference Architecture uses the “4+1” architectural view model as the methodology to cover all the aspects of the different problems the Project and the Consortium want to address.

The “4+1” architectural view model¹⁷ is a methodology to design an architecture for a software platform, having the main capacity of describing it from 5 concurrent “views”.

These views represent different stakeholders who could deal with the platform and the architecture, from the management, development and user perspectives. The four main views which facilitate the definition and design of the architecture are the logical, process, development and physical ones,

¹⁷ Kruchten 95 Kruchten, P.: “Architectural Blueprints - The “4+1” View Model of Software Architecture”, Paper published in IEEE Software 12, November 1995, pp. 42-50

while the “+1” view is represented by the use cases or scenarios, thus making this model an abstraction of the developed solution/platform and the basis for the development.

In the following, the meaning of the different views is explained.

- **Logical view:** it represents the range of functionalities or services that the system provides to the end users, and can be shown as block diagrams
- **Process view:** it represents the system processes and data flows and how the different processes and building blocks communicate between each other, with details on the runtime behaviour of the system.
- **Development view (or implementation view):** it illustrates the software management aspects of the system, from the programmer’s point of view. It is represented by UML components and package diagrams.
- **Physical view (or deployment view):** it describes the lower levels of the architecture, dealing with the physical infrastructures where the software components run and the physical connections between them. It can be represented by deployment diagrams.
- **Scenarios:** the architecture of the system can be explained from the end user’s point of view too, with the description of use cases. The use cases or “scenarios” aim at describing some possible functioning situations of the system and interactions between components. The validation and assessment of functionalities are usually performed by this view.

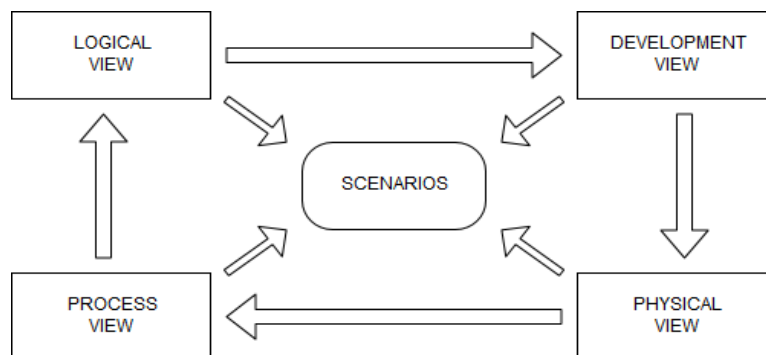


Figure 18 - "4+1" views

The present document describes in particular how the Logical view and the Process view were designed by the Consortium, and analyses afterwards the Development view and the physical view.

4.2 Key Assumptions

The INFINITECH RA is inspired by state-of-the-art solutions and technologies. The following assumptions are valid as general requirements for the IRA:

1. The BDVA Reference Model is a “reference” also for the INFINITECH project. More information are presented in the section 3.2.2.1 and can be found at SRIA BDVA [BDVA-RM]
2. The INFINITECH Reference Architecture (RA) should provide an abstraction to solve a general class of use cases and in particular those arising from the pilots.
3. INFINITECH RA should provide a mapping of SotA technologies, development methodologies, deployment and testing environments, production.
4. INFINITECH RA will exploit Microservices technologies and DevOps methodologies
5. Data Components should be mapped/encapsulated in Microservices
6. Loosely coupled microservices are preferred over tightly coupled intertwined applications
7. Data can be pipelined into different microservices to perform different type of processing
8. Data can be streamed and/or stored in data storages
9. Data Movements should be limited as much as possible (streaming in favour of storing)

10. Cross cutting services will provide added functionalities across the different layers

4.2.1 Definitions

- In the INFINITECH RA a **node** is a unit of data processing.
- A node exhibits **interfaces** (APIs) for data management in particular for consuming and producing data (IN and OUT)
- In a microservice implementation of IRA, nodes are **microservices** by preferable exposing a REST API
- IRA has layers that can be referred to the **BDVA Reference Model**
- Nodes belong to **layers** in the RA
- Nodes with IN OUT interfaces can be stacked up to form **data pipelines** (LEGO concept)
- Nodes are **loosely coupled** in the RA: i.e. they are not connected until a pipeline is created
- Data can **flow in all directions** (e.g. a ML can also push back data into a Data Layer)
- Node **stacks** with other compatible nodes (stackability depends on the IN-OUT interfaces)
- As a consequence of the previous point, RA can be segmented into **vertical layers** (called bars) to group compatible nodes.
- A node can belong to one or more **vertical bar**
- Node can be **orchestrated** to form solutions (**sandboxes**)

Therefore, the IRA can be considered:

- **Layered:** RA identifies Layers as way of grouping nodes in the same way as DBVA Reference Model has “concerns”
- **Loosely Coupled:** no rules to connect nodes in a predefined way or in a rigid stack
- **Distributed:** computing and nodes can be distributed anywhere (on-prem, on cloud, on different clouds)
- **Scalable:** nodes can distribute computing at edges, at HPC nodes (GPUs) or centrally
- **Multi workflow:** RA allows for simple pipelines and/or complex dataflows
- **Interoperable:** nodes can be connected to other nodes with compatible interfaces
- **Orchestrable:** nodes can be composed in different ways allowing creation of virtually infinite combinations (the sandbox concept)

4.3 Logical View

The Logical Views of the IRA are presented in the schemes below. Rather than identifying and specifying functional blocks and how they are interconnected, the **INFINITECH Reference Architecture** is defined by a set of “**principles**” to build pipelines or workflows. These principles constitute the guidelines that will drive any specific implementation of INFINITECH solutions (Pilots implementation, Minimum Viable Platform, Sandboxes, etc).

At a first level the INFINITECH RA can be seen as a pipelined mapping of nodes referring to the BDV Reference Model and cross cutting layers as in the schema below:

D2.13 – INFINITECH Reference Architecture - I

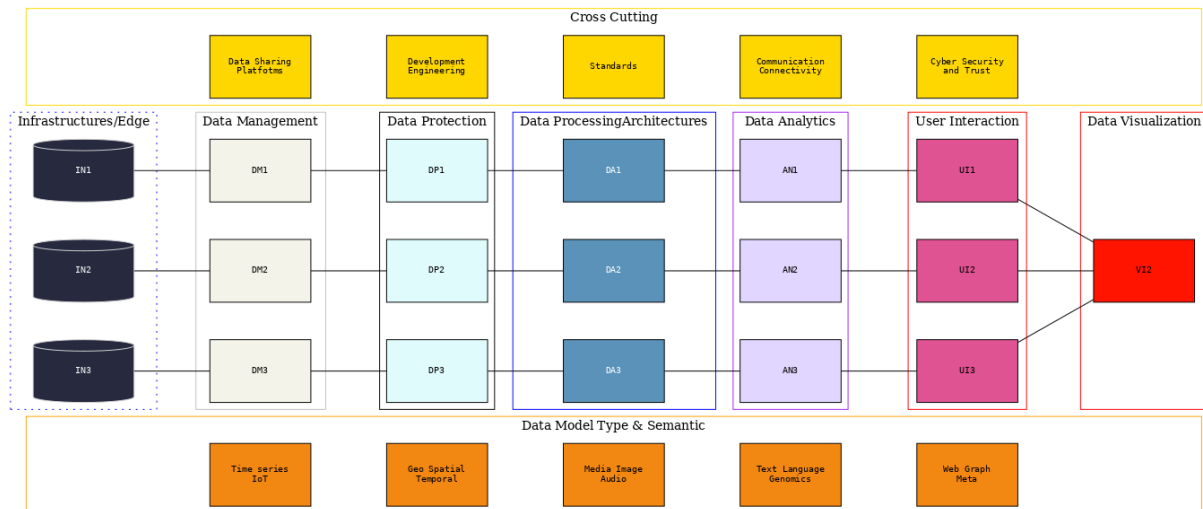


Figure 19 - INFINITECH Reference Architecture Mapping with BDVA Reference Model

Besides the correspondence with the BDVA Reference Model, the idea to solve the problems with data pipelines and workflows is close to the way data scientists work with data and resembles the powerful concept of neural networks. Nodes fire other nodes with data and transform the sources in the desirable result along the way.

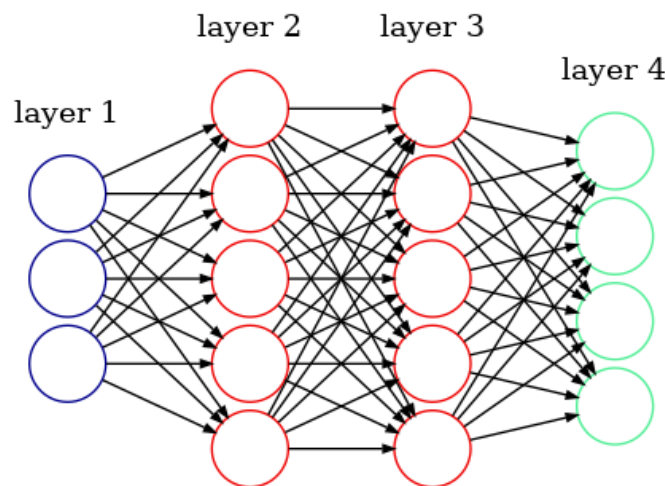


Figure 20 – Example of Neural Network the IRA Analogy

The following schema is a **High Level Logical View** that depicts the INFINITECH RA. The generic nodes depicted are only examples and do NOT refer to any specific tool or applications but to generic services that belong to general class of applications performing the functionality of the corresponding layer in the BDVA Reference Model.

Specific Nodes will be defined and mapped in specific solutions for all the Use Cases of the project. However, RA leaves to nodes a great level of flexibility depending on the implementation.

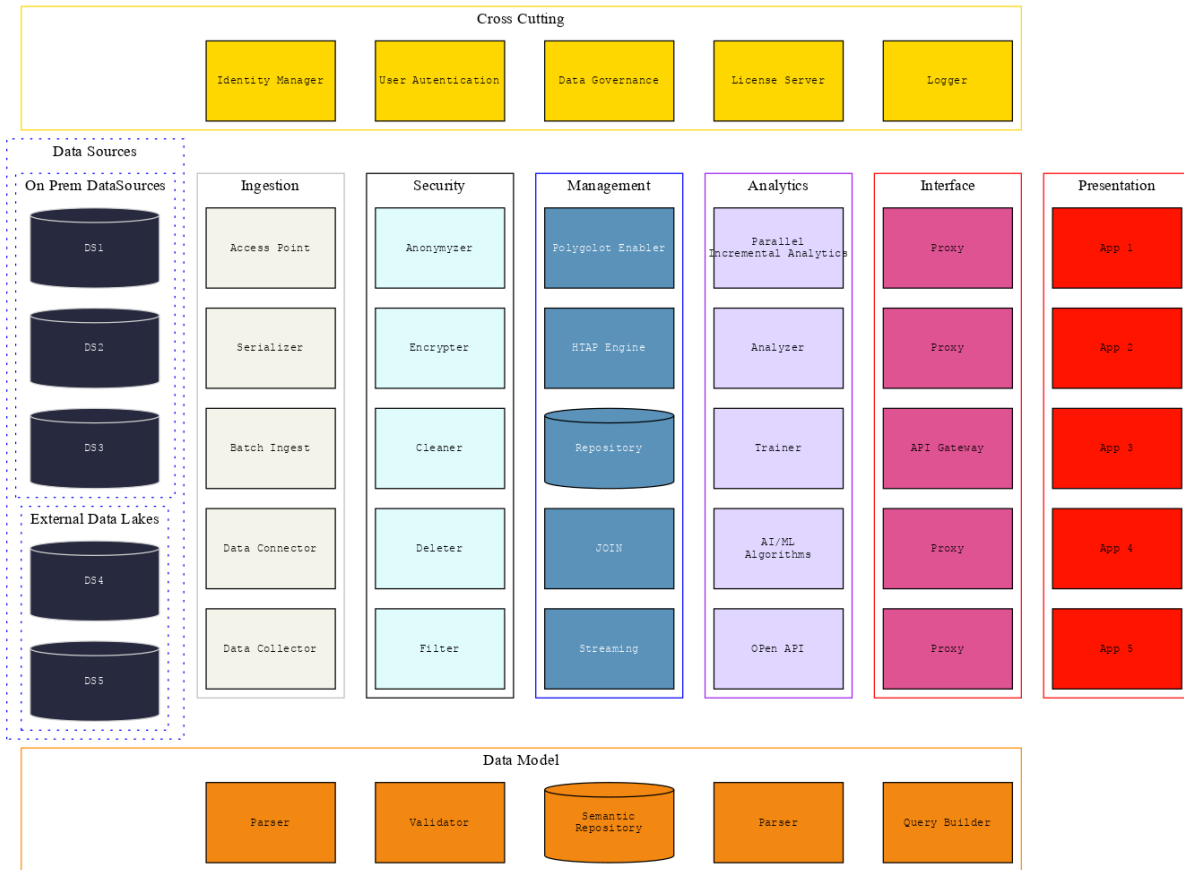


Figure 21 - INFINITECH Reference Architecture Logical View – Example Mapping

The INFINITECH RA defines *layers* as a way to logically group components. The identified layers are:

- **Data Sources:** at the infrastructure level there are the source of data (database management systems, data lakes holding non-structural data, etc)
- **Ingestion:** a layer of data management usually associated with data import, semantic annotation and filtering from data sources
- **Security:** a layer for management of the clearance of data for security, anonymization, cleaning of data before any further storing or elaboration
- **Management:** a layer responsible for the data management aspects, including the persistent storage in the central repository and the data processing enabling advanced functionalities such as Hybrid Transactional and Analytical Processing (HTAP), polyglot capabilities, etc
- **Analytics:** a layer for the AI/ML/DL components
- **Interface:** a layer for the definition data to be produced for user interfaces
- **Cross Cutting:** a layer with service components that provides functionalities orthogonal to the data flows (e.g. Authentication, Authorization, ...)
- **Data Model:** a cross cutting layer for modelling and semantics of data in the data flow
- **Presentation/Visualization:** a layer usually associated with the presentation applications (desktop, mobile apps, dashboards and the like)

It should be noted that the IRA **does not impose** any pipelined, or sequential composition of nodes. However, it is recommended to consider each different layer and the relative components to solve specific problems of the use case.

4.3.1 Logical Components Grouping

Deliverable D2.5 of the INFINITECH Project has made a preliminary list of components grouped in categories which can be referred to the IRA presented. These components provide the building blocks

that can be combined and used together to create BigData, IoT and AI pipelines for Digital Finance applications in-line with the logical view of the INFINITECH-RA. These components are listed in the table below to facilitate understanding of the INFINITECH-RA examples and scenarios given in subsequent section. For the documentation of the functionalities and the use (inputs/outputs) of these components, readers should consult deliverable D2.5. Note that the presented list of components is extensible i.e. it is likely to be augmented during the evolution of the technical developments of the project.

Table 1: List of INFINITECH Components and Technologies

Group	INFINITECH Components
Data Sources	Data Bases (RDBMS, SQL, NoSQL, ...) Data Lakes (Internet and Intranets sources,...) Stored files with text, images, audio/video IoT sources (mqtt, ...) Blockchains (Bitcoin, Ethereum, others, ...)
Data Ingestion	Data collector Data connector Serializer
Data Management	Relational Database Polyglot Query Processing Incremental Analytics Filtering Joiner OneHotEncoder Visualization Preparation Stream Processor
Data Security and Privacy	eIDAS integration Data Anonymization Digital User Onboarding System (DUOS) Data Protection Orchestrator (DPOR)
Blockchain and Information Sharing	Blockchain Reader Blockchain Writer Smart Contract Executor Blockchain Data Visualizer Blockchain Authenticator Blockchain Encryptor Blockchain Decryptor Blockchain Transaction Dataset Preparation Component Scalable Transaction Graph Analysis Component User Interface for Blockchain Transaction Reports and Visualization Component
Data Model and Semantics	Semantic Stream Analyzer Semantic Reasoner Ontology Mapping Preprocessing Anomaly Analysis Pattern Analysis Stream Story

Internet of Things	IoT Registry IoT Stream Handler
Analytics and Machine Learning	Fraud Detection Service Training Fraud Detection Service Execution Pay As You Drive Service Training Pay As You Drive Service Execution Recommendation Engine Training Recommendation Engine Execution Association Rules Training Association Rules Execution Unexpected transaction detector Account transaction forecaster Overdraft forecaster Automation proposer Transaction sentinel Bank’s product recommender Smart expense advisor Savings proposer Recommender Cash Flow Prediction Model Budget Prediction KPI Engine Transaction (Txn) Monitoring Transaction (Txn) Categorization Invoice Processing Benchmark(s) KMeans Random Forest (Model) Random Forest (Predict)
Interface	API Gateway Proxy

Appendix B contains a matrix mapping of RA modules versus User Requirements of Deliverables D2.1 and D2.2, while following figures depicts the mapping of components into INFINITECH RA layers.

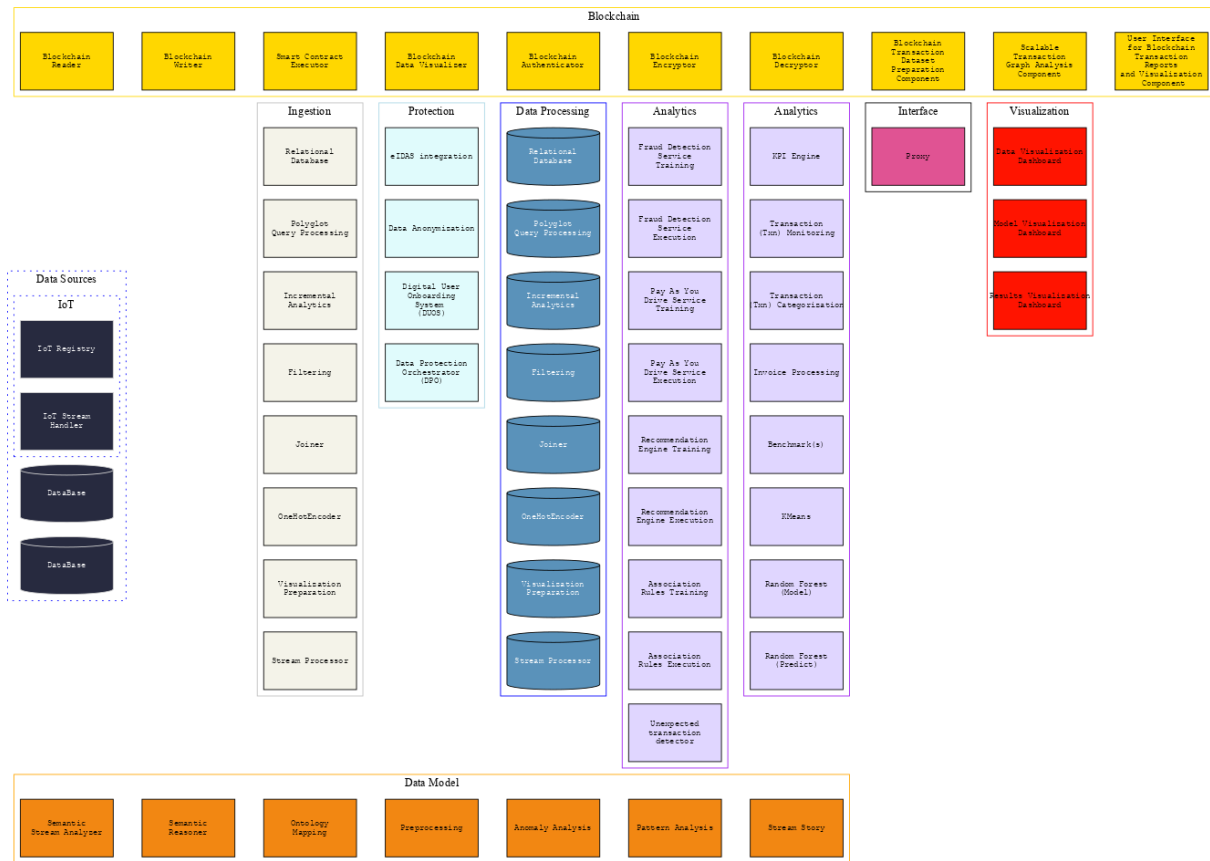


Figure 22 - INFINITECH Reference Architecture Logical View – Example Mapping

4.4 Process View

In the context of the “4+1” methodology, the runtime behaviour of the workflow will be analysed, to give a representation of the Process View in Section 6. In the specific case of INFINITECH project, the data-driven architecture to design the proposed solutions suggests to identify the Process View with the flow of the data through the different tiers already exposed for the Logical View. These data can be referred to simple events, or more complex situations such as security, analytics, data processing, and can be exchanged among the data components such as data collector, databases, services, dashboard, etc.

A schematic representation of the pilots at runtime is given in the Section on PILOTS, through a number of RA diagrams, each representing a different relevant workflow of the proposed use cases.

4.5 Development View

The INFINITECH RA is designed considering a microservices architecture implementation, with services interacting among them through REST APIs. In that respect, all microservices run in Docker containers on the Kubernetes platform. The following explains the technology background and the deployment view adopted in the project.

In recent years the DevOps (Development and Operations) methodology has been developed in a sense that is trying to bring, as much as possible, the world of those who develop applications (Developer) and the world of those who manage them (Operations). By bringing these two worlds closer it allows both to understand the problems of each other. It allows them the possibility of being faster and more efficient in the development of an application and its software lifecycle management,

because the developers specialize in the environment in which the software will be deployed and the Operation people know best its characteristics before still to put it into production. This type of approach requires not only a technological change but above all a cultural change.

The practical implementation of DevOps goes through the CI/CD. The CI/CD stands for the combined practices of Continuous Integration (CI) and Continuous Delivery (CD).

In order to support the development and testing activities, we have adopted a Continuous Integration approach, using open source tools such as Jenkins. Every time a developer pushes changes to the source code repository, Jenkins automatically triggers a new build of the component and deploys the updated container to an integration environment on Kubernetes. This will allow the team to continuously test their components against an up-to-date environment, speeding up development and avoiding painful integration problems at the end of the cycle. Continuous Delivery is an extension of that process. It's the automation of the release process so that new code can be deployed to target environments, typically to test environments, in a repeatable and automated fashion.

We intend to enhance this process by adopting a DevSecOps approach and including the related tools in our CI/CD pipeline.

DevSecOps aims at including security in the software development life cycle since the beginning, following the same principles of DevOps. Since security is considered throughout the process and not just as an afterthought at the end of it, products get a safer, more solid design and teams avoid costly release delays and rework due to non-compliance detected only late in the process.

DevSecOps project teams include security experts, who work with developers and operations to make sure that security requirements are properly addressed and best practices followed, in addition to validating product design and architecture.

Furthermore, based on the DevOps core principle of automation, DevSecOps introduces several security tools in the CI/CD pipeline, so that different kinds of security checks are executed continuously and automatically, giving to developers quick feedback whether their latest changes introduce a vulnerability that must be corrected.

If the DevOps methodology has tried to unify the world of the developer and that of the operation, the MLOps¹⁸ is focusing on:

- Facilitate communication and collaboration between teams
- Improve model tracking, versioning, monitoring and management
- Standardize the machine learning process to prepare for increasing regulation and policy

Putting this into practice is often very complicated, because ML works in heterogeneous environments, for example the models are often on Data Scientist's notebooks, its training is done in the cloud to take advantage of available resources, to finish with the execution of the software in production on premises.

Therefore, the first step towards MLOps requires standardizing these environments as much as possible and in order to achieve this, the Kubernetes along with the Docker containers provide the abstraction, scalability, portability, and reproducibility required to run the same piece of software in all these environments. As a second step it is necessary to make standard the workflow used for construction and building of the ML models.

¹⁸ MLOps is a compound of Machine Learning and IT Operations. MLOps is a practice for collaboration and communication between data scientists and operations professionals to help manage production ML lifecycle (source Wikipedia)

In this sense, software like Kubeflow¹⁹ provides an infrastructure to build models and it enables the portability of these models and workflows. In particular ML workflows are defined as Kubeflow pipelines. A pipeline consists of various steps:

- Data preparation
- Training
- Testing
- Serving

Each step is a container and the output of each step is the input of the following step. Once compiled, this pipeline is portable across environments. However, the possible usage of MLOPs technologies will be assessed in detail within the WP6 tasks.

4.6 Deployment View

In the last few years there has been a strong transformation, due to container spread which led to rethinking both how to manage own infrastructure and how to design and build the applications. As we said before the INFINITECH platform has been devised on the concept of a microservice architecture.

The nature of the microservices that is characterized by atomic components, makes them easier to develop, update, and deploy the entire solution, but if on the one hand they allow to have applications that quickly scale according to the requests and that could be easily updated, on the other hand they consider that a software has been previously managed as a single piece indivisible that is split into several dozen of microservices (containers), making it difficult to manage them.

In this context, arose the necessity to develop a tool that was able to manage the life-cycle of the microservices (deployment, scaling, and management), this tool was developed in Google with the name of "Project Seven of Nine "and released as open source software in 2014. Today that tool is known as Kubernetes.

In particular, it will use two main concepts available in Kubernetes in order to implement "Sandbox" in INFINITECH, these concepts are:

1. **Namespaces:** They are a logical grouping of a set of Kubernetes objects to whom it' possible to apply some policies, in particular:
 - **Quota** sets the limits on how many HW resources can be consumed by all objects
 - **Network** defines if the namespace can be accessed or can access to other Namespace, in other word if the Namespace is isolated or not

Different namespaces can be given different policies

2. **POD** is the simplest unit in the Kubernetes object. A Pod encapsulates one container, but in some cases (when the application is complex) a POD can encapsulate more than one container. Each POD has its own storage resources, a unique network IP, access port and options related to how the container/s should run.

The Kubernetes Namespace allows us to isolate logically the objects (mainly PODs) inside it from other Namespace, in other words each Namespace will be one use case of a specific pilot.

¹⁹ Kubeflow is a free and open-source machine learning platform built by developers at Google, Cisco, IBM, Red Hat, CoreOS and CaiCloud, and first released at Kubecon North America in 2017 (Wikipedia)

Regarding the “Testbed”, the set of hardware resources of each Data Center (Storage, Compute and Network) or VPC (Virtual Private Cloud) on the cloud provider will be considered a Testbed.

Therefore each dedicated Testbed will have only one cluster Kubernetes with many Namespace as many as are the use case to implement, instead for a shared testbed will have many clusters Kubernetes as are the pilots to manage.

Moreover in order to simplify development and integration by each partner, a blueprint will be created on a cloud provider. In particular on the cloud provider will be created a cluster Kubernetes on which will be implemented tools for CI/CD like:

- Gogs(Go Git Service) as software version
- Jenkins as pipeline manager
- Harbor as Docker registry

Finally, we planned to create a proper development environment in order to test the INFINITECH services/components.

After that, all partners will choose to use the same cloud provider we planned to provide the possibility to replicate a similar environment using tools like Terraform/Ansible that allow us to create the infrastructure as code. Regarding the pilot on-premises environment we will try to automate as much as possible the recreation of the blueprint environment.

Further details and specifications about the previous concepts will be reported in the upcoming WP6 deliverable D6.4 “Tools and Techniques for Tailored Sandboxes and Management of Datasets I”.

The following questions should be answered in the deployment phase:

- Which elements comprise RA of an INFINITECH development and deployment sandbox of ML algorithms in a distributed computing environment (testbed)?
- Where should the main architectural elements of the RA be deployed?

The possible options for a sandbox in any INFINITECH testbed are depicted in Figure 23:

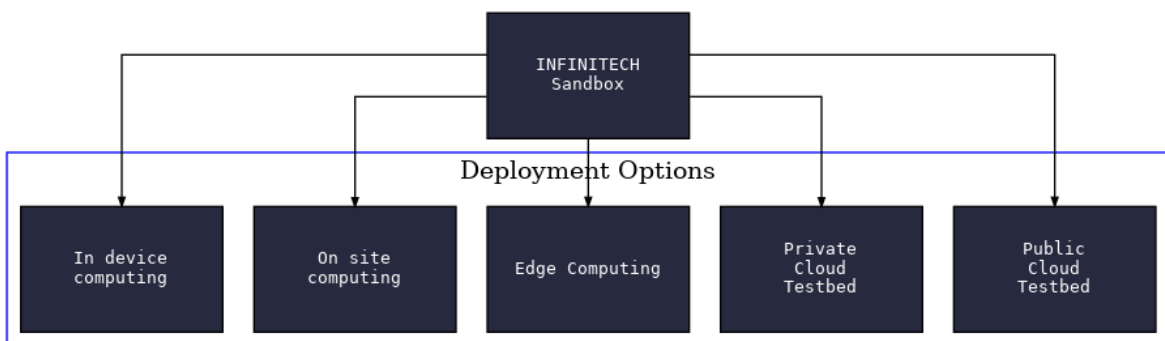


Figure 23 Deployment options for an end-to-end Software System

4.7 Use of INFINITECH-RA in Common BigData and AI Scenarios – Initial “Scenarios” View

The following subsections present some sample reference solutions for common use case scenarios in the sector.

4.7.1 Simple Machine Learning Pipeline

The reference architecture enables the design and execution of classical machine learning pipelines. Typical machine learning applications are implemented in-line with popular data mining methods like the CRISP-DM (Cross Industry Standard Process for Data Mining) [Shearer00]. CRISP-DM includes several phases, including the phases of data preparation, modelling and evaluation. In typical Machine Learning pipeline, data are acquired from various sources and prepared in-line with the needs of the target Machine Learning model. The preparation of the data includes their segmentation into training data (i.e. data used to train the model) and test data (i.e. data used to test the model). The model is usually evaluated against the requirements of the target business application. **Figure 24** illustrates how a typical machine learning pipeline can be implemented/mapped to the layers and functionalities of the INFINITECH-RA. Specifically, the following layers/parts are envisaged:

- **Data Sources:** At this level reside the data sources of the BigData/AI application. These may be of different types including databases, datastores, datalakes, files (e.g., spreadsheets) and more.
- **Data Ingestion:** At this level of the pipeline data are accessed based on appropriate connectors. Depending on the type of the data sources, other INFINITECH components for data ingestion can be used such as data collectors and serializers. Moreover, conversions between the data formats of the different sources can take place.
- **Data Management:** At this level data are managed. Filtering functions may be applied and data sources can be joined towards forming integrated data sets. Likewise, other preprocessing functionalities may be applied, such as partitioning of datasets into training and test segments. Furthermore, if needed, this layer provides the means for persisting data at scale, but also for accessing it through user friendly logical query interfaces. The latter functionalities are not however depicted in the figure.
- **Analytics:** This is the layer where the machine learning functions are placed. A Typical ML applications entails the training of Machine Learning model based on the training datasets, as well as the execution of the learnt models based test dataset. It may also include the scoring of the model based on the test data.
- **Presentation:** This is the layer where the model and its results are visualized in-line with the needs of the target application.

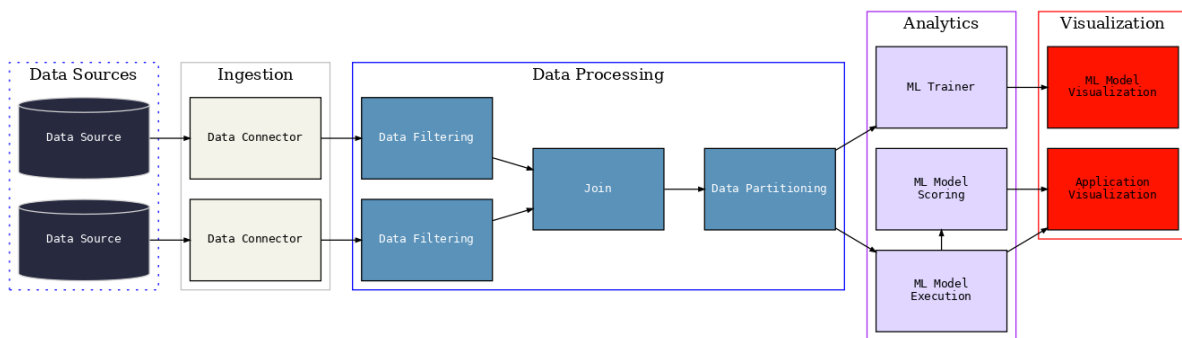


Figure 24: Simple Machine Learning Workflow Example implemented in-line with the INFINITECH-RA

4.7.2 Blockchain Data Sharing and Analytics

INFINITECH platform can also facilitate the execution of blockchain empowered scenarios for the financial sector, where e.g. several discrete entities (including banks, insurance companies, clients etc.) may be engaged in data sharing activities, thus empowering these financial organizations to update customer information as needed, while being able to access an up-to-date picture of the customer’s profile at all times. This is extremely significant for a plethora of FinTech applications, including yet not limited to Know Your Customer / Business, Credit Risk Scoring, Financial Products Personalization, Insurance Claims Management and more. Let’s consider a use case in which a

Financial Organization A (e.g. a bank, or an insurance company) wishes to be granted access to a customer of another financial organization (e.g. another bank within the same group, or a different entity altogether). Supposing that legal constraints are properly handled (e.g. consents for the sharing of the data has been granted, and all corresponding national and/or European regulations such as GDPR are properly respected and abided by), and that all data connectors (required for the retrieval of the data from the raw information sources) and data curation services (e.g. the Cleaner, required for cleaning the raw data sources, the Anonymizer, required for anonymizing the same information sources, and the Harmonizer, required for mapping the cleaned, and properly anonymized information sources to the Common INFINITECH Information Model) are in place, enabling the collection and pre-processing of the raw information sources, the following flow is envisaged:

- 1) The Blockchain Encryptor encrypts the properly pre-processed data so that they can be saved in the private collection of Organization A.
- 2) The Blockchain Authenticator authenticates (the user of) Organization A, so that access to update the ledger is granted.
- 3) Once authenticated, the Blockchain Writer inserts the encrypted data to the private collection of Organization A and the hash of the data is submitted to the ledger.
- 4) Organization B requests access to the data of a customer from Organization A.
- 5) The Blockchain Authenticator authenticates (the user of) Organization B that initiates the request so as to grant or decline the request to access to private-actual data.
- 6) Once (the user of) Organization B is authenticated, the Smart Contract Executor translates the Query submitted and queries the private collection of Organization A.
- 7) The Blockchain Authenticator authenticates (the user of) Organization B, so that access to read the ledger is granted.
- 8) The Blockchain Reader retrieves the queried information from Organization’s A private collection.
- 9) The Smart Contract Executor generates a new transaction and triggers the Blockchain Authenticator so as to authenticate (the user of) Organization A, in order to grant access to update the ledger.
- 10) Once authenticated, the Blockchain Writer submits the encrypted data to the ledger, and a new record on the same ledger is created, containing the metadata of the contract (organizations involved, data created, metadata of the encrypted data transmitted, validity of the contract) and the actual encrypted data.
- 11) Organization A sends out of band the decryption key to the Organization B, and the Blockchain Decrypter decrypts the queried data.

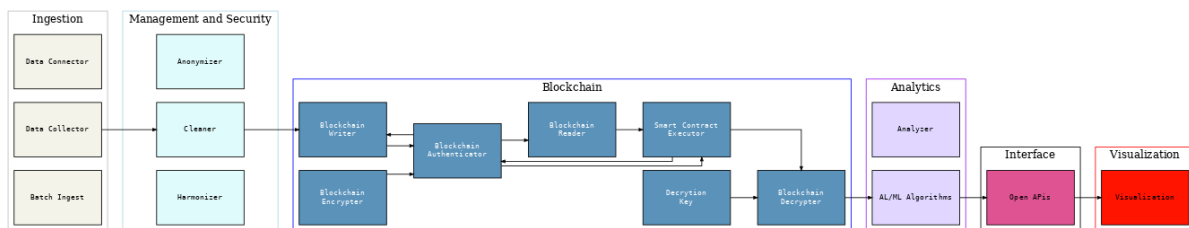


Figure 25: Blockchain Data Sharing and Analytics Pipeline

4.7.3 Data Ingestion, Anonymization, Analytics and Visualization Pipeline

One of the most common BigData and AI pipelines involves data acquisition, anonymization and processing of the data using some data mining or machine learning technique. Data acquisition is usually the starting point as a daily basis for a Data Scientist, or an AI/ML developer. Cleaning, homogenization, study, transformations of data until you get the final set of features (the real raw data feeding your model). But something is usually missed, in ML scenarios (when learning, courses, research, PoCs, experiments, etc), is the importance of keeping data privacy and being compliant with

GDPR. Something crucial as you develop into the real world, and specially in concrete scenarios like insurance, banking, and financial services. Where the data privacy is a pillar, and the black-box, that is usually related to AI, needs to be reliable by users (as data owners).

Therefore, these kinds of scenarios remark the importance of anonymization data as a part of AI pipeline development, during the data preparation phase. But anonymization is not only about hiding personal data in obvious scenarios. More complex situation would infer personal data from data in principle not so directly related to a person.

For example: enough information from GPS positioning would study your movement patterns, in some regions, this would easily point to a very concrete person. In this case, if GPS information is crucial for the model, GPS coordinates need to be converted to not so accurate information. For example, converting to GPS coordinates to just knowing the city where the event happened.

With the data prepared, privacy respected and features ready, the AI developer will start creating new models to infer new information from the current existing ones. ML phase could be divided, at least, in two components the model training and the model deployment. The model training uses the input data to learn from the past, looking for patterns, that would infer or predict future situations. With the model ready, this is deployed in order to provide a concrete service: *Is it going to rain today? Are my stocks going to fall? Am I a good driver?*

The visualization shows the results of querying these deployed models to show the responses for a final user. Visualization would be a very specific GUI for that model, or part of a more complex existing dashboard. In this point, INFINITECH fosters REST APIs and JSON data formats for quick and robust integrations. INFINITECH WP4 is in charge of providing the (semantic) data models making these integrations scalable and replicable in different situations and scenarios.

4.7.4 Data Exchange and Semantic Interoperability Pipeline

The INFINITECH reference architecture addresses the need for defining the methods and conventions alike the recommendations for naming, identification, categorising about the use, manipulation and inclusion of data sets in the INFINITECH project. There are regulatory aspects about the data that are not addressed in this document but in data management plan, rather this section focus on the operational part of the data, looking closely at semantic and interoperable aspects within the project.

This section related to data exchange and semantic interoperability pipeline also identifies the different important reference use cases addressing general INFINITECH Pilot aspects as reference use cases across different Pilots. The data exchange and interoperability aspects are according to the WP4 implementation work in accordance with the signed Grant Agreement of the INFINITECH project and with respect to EC Horizons 2020 recommendations for operating data.

The information provided in this section is meant to define common grounds in relation to data identification, data annotation, data sharing, data exchange and data sharing and must be used for the foundations of successful internal and external Pilot data interoperability aspects in order to motivate participation and collaboration within the INFINITECH consortium partners at the mainly at the pilot level but also between and amongst external partners or participants in the project by directing the developments and integration of semantic interoperability and data exchange aspects.

The following pipeline is similar to the classical data collection / data analytics pipeline, yet it includes INFINITECH components and data transformation that facilitate semantic interoperability.

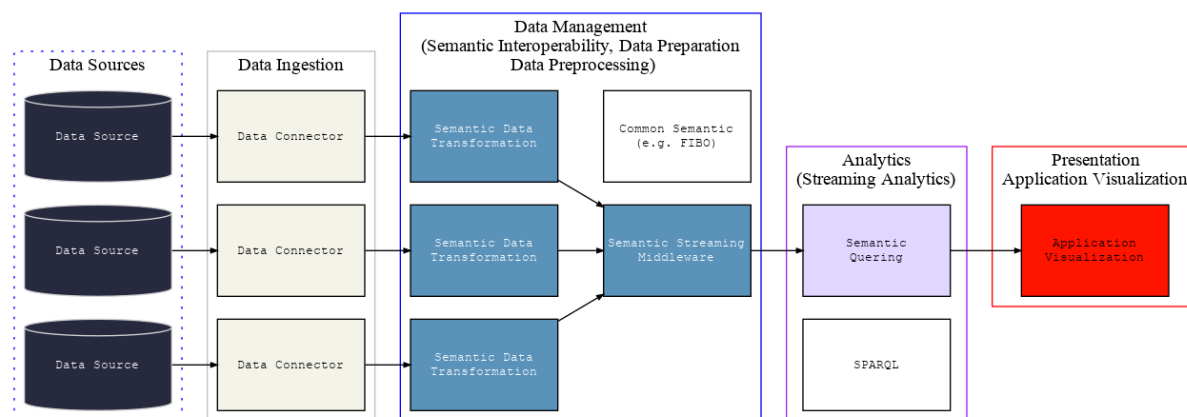


Figure 26: Pipeline for a Scenario involving Semantic Interoperability Across Diverse Data Sources

The Semantic Interoperability pipeline and its pre-conceived components (specifically semantic data transformation components, semantic streaming middleware, semantic querying and visualisation of semantic data) are under development in WP4 and are designed and intended to be used to unify the semantics of diverse data sources. This approach can accordingly enable the use of the INFINITECH Semantic Streaming Analytics components in the pipeline enabling multiple general semantic interoperability operations and its particular use cases, some of them are described as follow:

- **Data Identification:** The data identification points in INFINITECH are where data will be used, a part of any generalization and/or categorization the data is a representation of a fact (virtual or physical) and thus its identification and intended usage modify its way for communication. Identification is a process where the data is not only identify but categorised and classify with the purpose of:
 - Outline the discoverability of data (metadata provision)
 - Outline the use of identifiers of data and refer to standard identification mechanism. e.g. Do you make use of persistent and unique identifiers such as Digital Object Identifiers?
 - Outline naming conventions used
 - Outline the approach towards search keyword
 - Outline the approach for clear versioning

Example use case or application use: Identification of the data in INFINITECH can be useful when the process of joining specified datasets with the updated new datasets version is required, by maintaining versions of this datasets multiple benefits of interoperability are activated. For example main categories of data assets can be identified by INFINITECH pilots and comprise both the data created previous any operation via INFINITECH solution, data sets used during the operation of the INFINITECH Pilots solutions and the data produced as result of the INFINITECH services process can be also identified.

- **Data Annotation:** In INFINITECH annotation procedures shall be used to coordinate at the different deployment sites on how to collect the data locally and cross pilot, if required. In some cases, it is a questionnaire presented to the user before and after the pilot; in other cases, is the result of tracking the usage of the system. Data annotation is a process for gathering the information according to reference vocabularies or conventions from domain professionals to enhance the available data and make it machine understandable; etc. Therefore, it is expected that in INFINITECH each deployment site will have its own data format and collection method that will be connected INFINITECH Open Data Model as a collection online available tool, and if desired all sites will share the same tool for the same data format, as well as the same procedure to annotate the data. Annotation is a process where the data is enhanced with the purpose of:
 - Data consumed and/or handled within the INFINITECH project by means of electronic and

non-electronic forms.

- Data Registered using external information to the project that will be derived to the INFINITECH project by means of electronic and non-electronic forms.
- Have control over the processes following particular data lifecycle.
- Data Exposition by means of data that has been registered and or produced in the INFINITECH project and could be distributed by means of electronic and non- electronic forms.
- Data Management by means of data and information that can be registered and produced within INFINITECH mediums, by means of, electronic and non-electronics forms that is used within the consortium for processes, operations and or simply usage for the normal activity within the project.

The following are the envisage use cases involved in data capture annotation: When data sets that will be produced as result of management and operations with other sources of information. When data sets that will be provided to the INFINITECH project to store, handle and or processing. This is the primary activity in INFINITECH. And finally if data sets need to be make them accessible via interfaces to get the information used and/or generated by INFINITECH providing that those data sets are used for informative process.

- **Data Sharing:** Sharing data sometimes is equally treated as the process of exchanging data, however in INFINITECH and following dataflows best practices, this is a process to guarantee a durable threat for the available (stored or not) data. In a wider definition data sharing of data defines the process of formulating and defining the ways that guarantee the data will remain during the time in different instances or data holders with all the guarantees of identification, traceability and ownership and to ensure that the data will be useful for the same purpose which it was created. Data Sharing is a process where the data is treated with the purpose of:
 - Migrate the data to best format
 - Migrate the data to another medium or a more suitable medium
 - Back-up stored data to preserve the information
 - Create metadata and documentation for the stored data
 - Archive data in a defined physical or virtual medium

Example use case or application use: In INFINITECH, pilot-generated data and/or external data to the project that will be necessary used to derivate to INFINITECH project services by means of electronic and non-electronic data records. Data sets that will be provided by INFINITECH Pilot(s) to be stored, handled and or processed. Data captured by the INFINITECH pilot infrastructure defined on each pilot with the particularity that in the case of the data generated by the INFINITECH Pilot solutions must follow a defined semantic description approach based on INFINITECH standard-adopted formats.

- **Data Exchange:** Enabling data exchange defines the adequate ways to access the data but also the treatment of the data by its different involved entities or subjects to operate the data. Data Exchange is one of the main challenges in INFINITECH project. Data Exchange is a complex process because based on the demand for the data is the way to facilitate data access. If the demand is coming from multiple parts the access process is multiplied, first because there are activities associated that are specific for every part demanding access to the data and second because if the data is not prepared to be used by multiple parts it has to be pre-processed to repeat parts of the lifecycle in generating new data. Additional to that data exchange imply Data Access making necessary to define data values against the data. Data exchange is a process where the data is not only identified, categorised, classified and annotated with the sole but not limited purpose of :
 - Distribute Re-distribute the data
 - Share and format of the data
 - Mechanisms for control and accessibility

- Establish copyright for the data
- Promote data

The following are the envisaged use cases for data exchange that can be applicable in INFINITECH, Users of the INFINITECH pilot looking at data sets for public evidence e.g. profile building, liability aspects, ownership of data etc. INFINITECH users and professionals can benefit from Data Exchange by being informed about the data exchanged and shared with evidence and impact rather than just a reporting or informative element. INFINITECH users at the Pilot level using data dashboard(s) which will provide the view and access to data.

- **Data Discovery:** INFINITECH data discovery focuses on providing technological means for uncover aspects of the data and the data itself by using semantic annotation and descriptions that are not considered evident or as part of the data itself. Discover relies in the INFINITECH lifecycle simply because the demands for the information is to generate a meaningful outcome that helps to inform, observe, visualize features of the data. Data Discovery will ensure the re-usability of the data and the ways to transform it into meaningful results. In INFINITECH Data Discovery provides mechanism for data transformation, correct data collection and data re-usability and that the data can be repurposed. Data Discovery is also process where the data is treated with the purpose of:
 - Follow-up data along other data places or data sources
 - Identify or define new Data lifecycles
 - Undertake information about the produced data
 - Find points for storing/improving data
 - Learn from the shared data
 - Identify better mechanisms for sharing data
 - Define post process beyond re-using and transforming data

Example use case or application use: Discovery can look at the data processes in the INFINITECH lifecycle in order to enrich the capacity to look at those data characteristics and operations. Discovery is one of the main objectives in INFINITECH for enabling analysis having more implications based on a specific demand of the data discovery following the original representation and basic information provided when the data was created.

4.8 Mapping Methodology for Use Cases in the RA

The goal of INFINITECH RA is to create a RA for facilitating the realization of Big Data pipelines dealing also with IoT and Blockchain data sources utilizing ML techniques in distributed computing environments.

As stated before, the IRA is NOT a one-size-fits-all solution but rather a schema to compare the solutions against best practice and standard flows and components. A methodology to check the Use Case with the IRA is suggested as follows:

1. A Use Case will be considered as a transformation from some Data Sources to some Data Destinations. This is called the “stretching phase” where a pipeline or workflow is defined from the different sources to the resulting data.
2. The INFINITECH RA is the total transformation of the sources into destinations once the appropriate workflow is defined and implemented.
3. Each Data Sources are ingested into the INFINITECH via a first layer of Data Management components (serializers, database connectors, ...)
4. Data Sources can be accessible from different types of data management technologies (SQL, NoSQL, IOTs streams, Blockchains’ data but also raw data such as Text, Images, Videos, etc.) and a connector must be provided in order to support all the variety of the supported target datastore management systems.

5. Data should not be stored into the INFINITECH platform data layers, unless they are “cleared”. Clearing can involve cleaning, filtering, deletions, anonymization, encrypting of raw data etc. These transformations are different among the various datasets. This “clearing” can be managed via the security layer of the RA model or the cross cutting services.
6. The Data Model provides reference for ontologies and semantics across the nodes.
7. In most use cases, most of the elaborations, requires to store data (filtered by the processes beforehand) in a data store. The datastore of the IRA can be a very complex infrastructure able to manage big data and should be able to scale out accordingly in order to support the diversity of data and work load.
8. A generic data component should provide the basic CRUD (create, retrieve, update, delete) endpoints along with advanced endpoints (like streaming in Real Time, HTAP capabilities) with standard API. When standard RESTful APIS are provided components can be stacked (piled up as the LEGO concept).
9. Analytics work on top of the Data (stored or streamed), consume and might produce data, which is stored into the data repository as well.
10. Ideally, out of the analytic process, an interface to the data should be defined. This is usually provided by an intermediate Interface layer with definitions for APIs used to build applications. A fundamental component like an API gateway can be inserted into the chain before any presentation/visualization of data.
11. Finally, data can be presented or even ingested into another pipeline for further data elaboration. The presentation layer defines a way to present or visualize data and is related to building applications on top of the interfaces provided by the interface and workflow.

However, many other points must be considered. Most of the real-world use cases can be much more complicated than the methodology proposed above. They are intertwined with existing infrastructures and data processing components and cannot be easily stretched into a “pipelined” workflow as suggested. In those cases, it should be possible to identify the boundaries of “computational nodes” that provide basic interfaces.

A complex infrastructure can be more than one microservice with many endpoints distributed among the different microservices in order to keep the functionalities homogeneous. On the other hand, a huge infrastructure can be encapsulated (wrapped) into a microservice interface to exhibit basic functions. When this exercise is completed, a map of a sandbox to the IRA can be provided as a reference schema to guide (or at least to benchmark) the implementation.

5. Addressing Stakeholder’s requirements

The INFINITECH-RA provides the means for addressing the identified requirements and challenges of BigData and AI applications in the target sector. This is illustrated in the following tables where the RA is confronted to the earlier listed requirements and challenges. Note however that several of these requirements are not addressed by the RA per se, but rather based on some of the INFINITECH components and technologies such as the components presented in D2.5. The latter provide the technical instruments that substantiate the high level concepts of the RA in concrete solutions.

Table 2: INFINITECH-RA vs Key Requirements for BigData and AI Applications in Digital Finance

Requirements	INFINITECH RA
Siloed Data and Business Operations	The INFINITECH-RA is versatile in supporting data collection and ingestion from diverse sources (e.g., databases, data lakes). Furthermore, it offers a rich set of data pre-processing and data preparation functionalities, range from simple join of relational/tabular data to semantic interoperability functionalities.
Real Time Performance Requirements	Based on the INFINITECH-RA and the project’s data management technologies it is possible to support integrated data management (ingestion, access) pipelines that improve data access performance. Additional performance gains are obtained based on the use of INFINITECH analytics technologies like incremental analytics.
Mobility and Multiple Channels	INFINITECH-RA decouples the visualization and application development parts of an application from the data management parts. As such it facilitates the implementation and delivery of multi-channel, mobile first applications.
Automation	INFINITECH-RA Automates the development, deployment, and operation of BigData and AI applications in Digital Finance based on the use of pipelines. The latter eliminate manual steps and process, through enable seamless flow of data from the sources to the applications. INFINITECH technologies have been structured as pipelines’ building blocks to facilitate this approach to automation. The pipelines of the INFINITECH-RA provide a sound basis for further automation based on the use of AutoML frameworks (e.g., MLBOX, AutoSklearn or TPOT)
Transparency	Pipelines provide transparency on the data processing. They include explainable rule-based algorithms (e.g., random forests) that can be used when transparency at the level of ML algorithms is required. Likewise, INFINITECH-RA allows for the implementation of transparent and explainable AI techniques.
Security and Regulatory Compliance	The INFINITECH-RA specifies dedicated layers where security and regulatory compliance functions can be placed. This is for example the case with the anonymization functionalities that have been exemplified earlier. Furthermore, a group of INFINTIECH Security and Privacy Components and technologies is available for use in pipelines with stringent security and data protection requirements.
Ethical and Responsible AI	The above-listed transparency, security and privacy friendliness features boost the project’s ability to support Ethical and Responsible AI applications, such as applications that are secure, trustworthy and transparent in terms of their AI operations.

Table 3: INFINITECH-RA vs Key Challenges for BigData and AI Applications in Digital Finance

Challenge	INFINITECH RA
Limited Integration	The INFINITECH RA and components attack integration challenges from multiple complementary perspectives, including: (i) Integration of data and ML operations through workflows and pipelines; (ii) Integrated data modelling, including semantics; (iii) Continuous integration in-line with a DevOps and MLOps paradigm.
Poor Stakeholders Collaboration in Financial Services	The INFINITECH RA specifies a module/approach for decentralized and transparent movements of data, which is a foundation for more efficient collaboration between stakeholders of the financial services supply chain. Specifically, the blockchain and information sharing components that support the RA will boost the collaboration of stakeholders in the Financial Services supply chain.
Compliance with Stringent Regulatory Requirements	The INFINITECH-RA pipelines are built based on components that support security, data protection and regulatory compliance requirements. Relevant provisions are made in the layers of the INFINITECH-RA. Furthermore, INFINITECH specifies the security and data privacy groups of components that facilitate compliance at the technical level.
Lack of Flexibility	INFINITECH RA leverages industry best practices for the development of flexible data driven pipelines based on DevOps/MLOps approaches. It leverages microservices and other design principles that boost automated testing and continuous integration. It will be also augmented with constructs that will facilitate its evolution to MLOPs that enable continuous updates of Machine Learning and AI related pipelines. Also, INFINITECH technologies provide flexibility at the levels of data ingestion and data sources integration, including support for relevant ontologies.

6. Pilots' Reference Designs and Initial Alignment to the INFINITECH-RA

The RA is grounded by the Pilots workflows as references.

The following sections present the mapping of the Pilots' Use Cases in the logical view and components identification of the INFINITECH Reference Architecture. For every pilot a basic description of the objectives, sources and workflow along with the expected results and interfaces are described.

The Reference Designs presented are not intended be exhaustive and complete and in most cases do not present the different phases of the deployments, like the training and the run phase of the Machine Learning algorithms. They are intended to demonstrate the viability and versatility of the INFINITECH-RA approach. Hence, the mappings of the pilot systems to the INFINITECH-RA concepts and components are at this stage by no means exhaustive and complete. Following releases of the INFINITECH-RA (i.e. deliverables D2.14 and D2.15) will provide more complete mappings of the pilot systems, both in terms of the mapping of different viewpoints and in terms of the use of more INFINITECH components and technologies.

6.1 Pilot#1: Invoices Processing Platform for a more Sustainable Banking Industry

6.1.1 Pilot Objectives

The main objective of pilot #1 is to develop, integrate and deploy a data-intensive system to extract information from notary invoices, in order to:

- Establish the sustainability index of each notary based on the number of physical copies that are issued.
- Provide financial institutions the information (properly indexed) about the documents that are finally generated by notarial services required by the bank.
- Promote notarial services from those with the higher sustainability score.

6.1.2 Data Sources

Data will be extracted from 32.300 real invoices documents from 3.000 different notaries extracted from Bankia systems. Invoice documents will be digitalized in PDF format or may also arrive already digitalized from other channels (email attachments, bulk sftp, etc.). Data type will be: PDF/Image/Text. Data format will be: PDF/ PNG/ TXT.

Dataset TableBank: Table Benchmark for Image based Table Detection and Recognition.

6.1.3 Data Produced

Digitization of contracting and invoicing processes will allow an automated analysis of the digitized documents enabling a smart and autonomous scoring of notary services. Rating notaries based on a sustainability index will provide a new criterion to be applied when contracting these services impacting positively in the short and long-term in the amount of paper used and the economic fees applied.

6.1.4 Explainable Workflow

Invoice documents will be securely storage in a data lake. The system will parallelize different jobs to pre-process, process and post-process the documents and the outcomes. For instance: Image pre-processing (cropping, adjusting brightness, contrast, etc.); converting PDF to Text; OCR; text

correction. A computer Vision system will identify and extract tables from invoices that will allow extracting sensible information to establish a sustainability scoring. And using machine learning we will extract information from the identified and extracted tables.

The extracted information will be displayed so it can be validated and re-introduced to the system. The AI models will be trained (offline process) with a combination of public huge datasets and specific invoices samples. Trained models will be published to the runtime processing time after an expert evaluation.

6.1.5 Logical Schema

The following figure illustrates a preliminary workflow of the pilot with a first mapping with INFINITECH-RA constructs. More data components are introduced that needs to be mapped in RA.

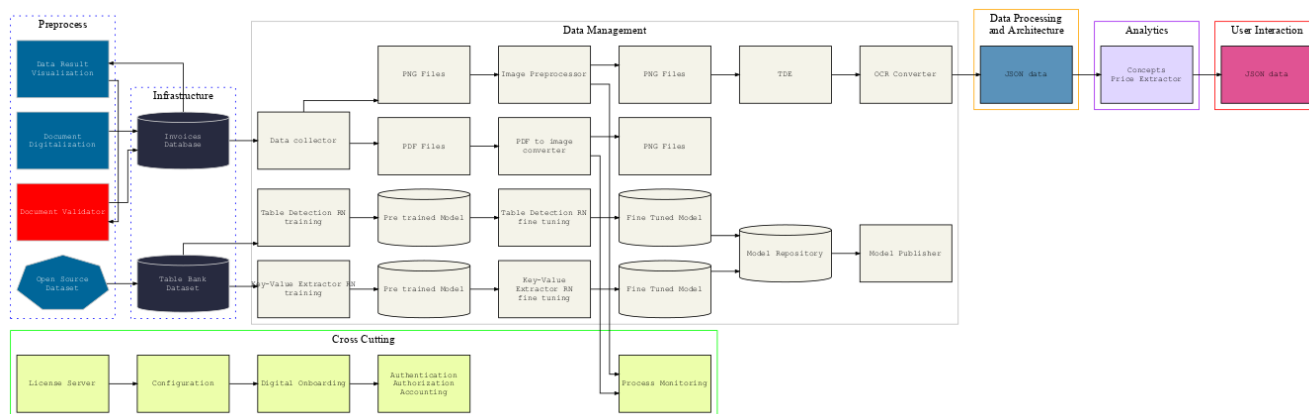


Figure 27: Logical Architecture of Invoices Processing Pilot

6.1.6 Components

The main components that will be part of this pilot are:

- Invoices and TableBank databases.
- Data Collector.
- Data converters/formatters: Image pre-processor, PDF to Image converter and OCR Converter.
- DL models (neural networks): Table Detector and Concepts & Prices Extractor.
- Data Results Visualization and Process Monitoring.
- Data Validator.
- Data Tagger.
- Training and fine-tuning pipelines of DL models.
- Models Repository and Models Publisher.

6.2 Pilot#2: Real-time risk assessment in Investment Banking

6.2.1 Pilot Objectives

The Pilot#2 will implement a real-time risk assessment and monitoring procedure for two standard risk metrics: VaR (Value-at-Risk) and ES (Expected Shortfall). The main outcome is the measurement of market risks of assets portfolios. In addition, the pilot will evaluate what-if scenarios allowing pre-trade analysis.

6.2.2 Data Sources

Data will be extracted from several data sources: real-time market data, historical market data, electronic order platform (derived analysis data, risk measures), and news article data.

6.2.3 Data Produced

The main data computed and produced include the VaR (Value-at-Risk) and ES (Expected Shortfall) parameters.

6.2.4 Explainable Workflow

Data from the real-time market database are extracted and managed using a data windows component, creating segments of time series. Data from the electronic order platform are managed firstly using a data streamer, the portfolio-input stream, capable of executing order and trade data, and secondly using a data extractor. These data will also serve as input for the scenario specifications component.

The historical market data will then feed the correlation matrix component together with the processed data from the real-time market database and the electronic order platform database. The correlation matrix processes and calculates the ingested data, merging the different data sources. The output will then serve as input, together with the scenario specifications component, for the scenario generation, the basis for the Montecarlo simulation. The processed data will then go into the Analytics component.

On the other side, data from the news article database are processed using the text extraction component and then market sentiment extraction one. Therefore, sentiment and behavioural analysis will be performed, serving as well as input for the Analytics component.

The analytics component will perform calculations on the data from the above-described flows and from the inputs of the configurator. The latter involves interaction of the user, in order to configure specifications for the scenario generation.

The results will be viewed in the User interface, showing numeric value(s) and graph(s).

6.2.5 Logical Schema

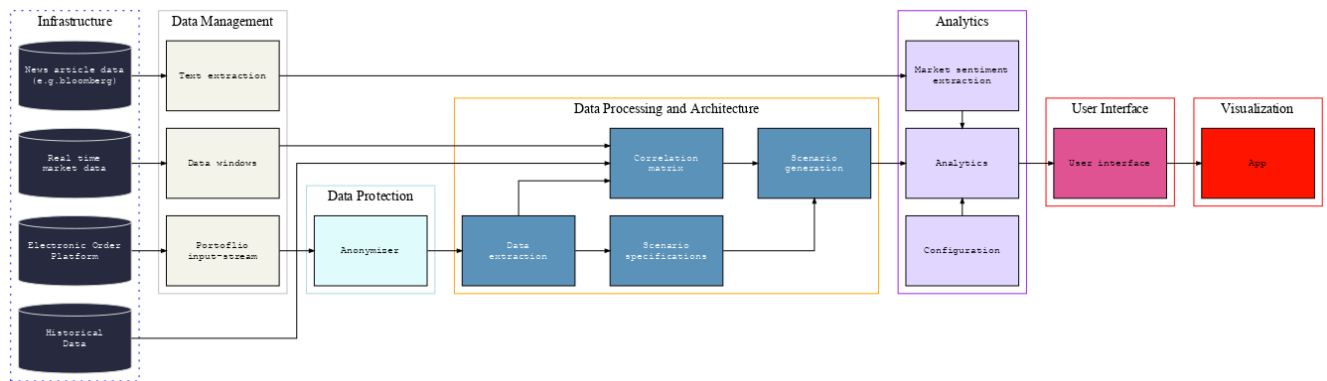


Figure 28: Logical Architecture of Real-Time Risk Assessment Pilot

6.2.6 Components

The workflow leverages the JOIN and Anonymizer components and others depicted in the logical schema (see figure above)

6.3 Pilot#3: Collaborative Customer-centric Data Analytics for Financial Services

6.3.1 Pilot Objectives

The Irish pilot will need to mimic a data sharing ecosystem by mimicing participants in that ecosystem and provide rules of engagement and highlighting the value exchanges between participants. A digital ecosystem framework is described here to articulate testbed components required.

6.3.2 Data Sources

Pilot #3 will consider two sources of data:

- **Operational Data Sources** – We will not using existing BOI Ops. Data sources, because of confidentiality issues even if anonymized and also data consistency issues. Instead Proof of concept data sources are ‘synthetic’ customer, account and transactions data designed to mimic real world data scenarios from financial services.
- **Captured data** from data entry in application including consent or metadata exhaust from sharing process.

6.3.3 Data Produced

For Pilot 3 a better representation of the data lifecycle might be. 1.Data utilized/transferred (E.g. data sharing payload – Customer/KYC, Account or Transaction Data), 2.Data transformation (e.g. any data changes) , 3. Data produced (e.g. new data) & 4. Data deletion (e.g. revoked consent) etc.

6.3.4 Explainable Workflow

The whole premise of the pilot purposed is to enable unlimited use cases between any participants via a single application creating a single ecosystem. Specific back end data services might be built to support a particular use case e.g. KYC. Below is a data flow of KYC use case in terms of business process/data flow and technical data flow.

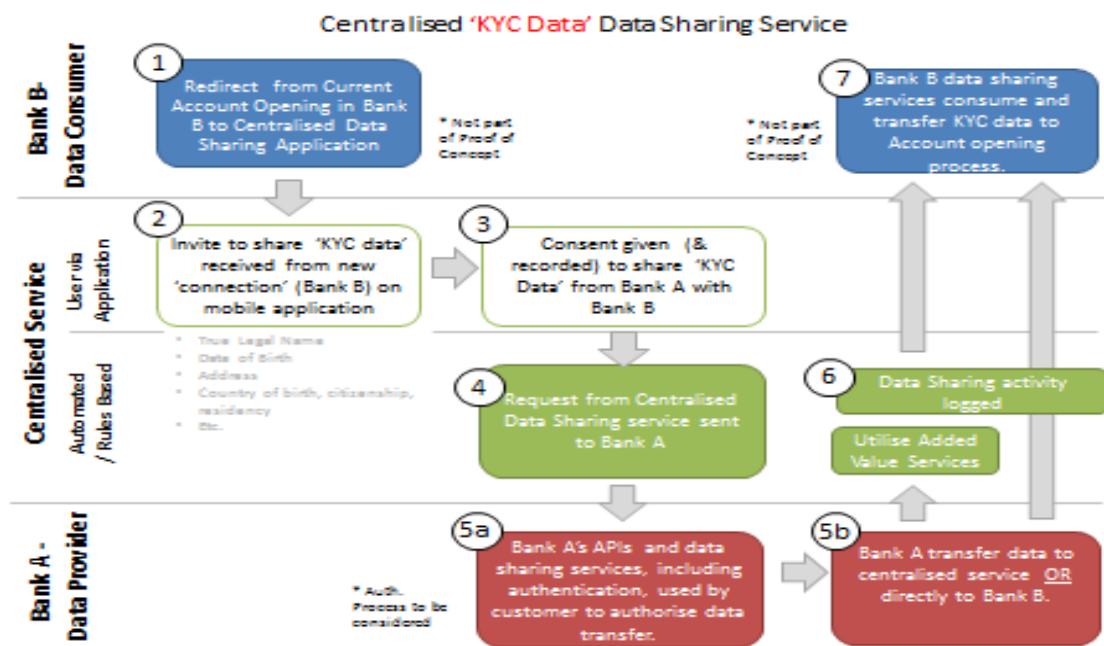


Figure 29: Workflow Elements of the Customer-Centric Data Analytics Pilot – KYC Data Sharing Process – Business Workflow

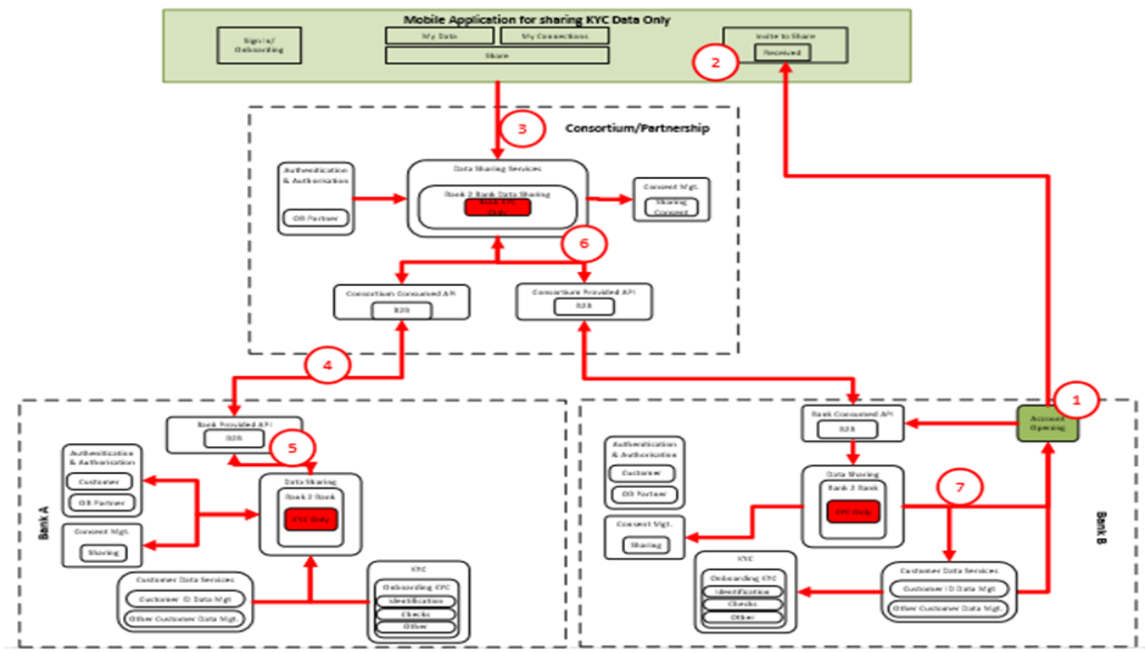


Figure 30: Workflow Elements of the Customer-Centric Data Analytics Pilot – KYC Data Sharing Process – Technical Workflow

6.3.5 Logical Schema

The following figure refactors the components of the above-listed workflows towards illustrating the pilot logical architecture in-line with the INFINITECH-RA.

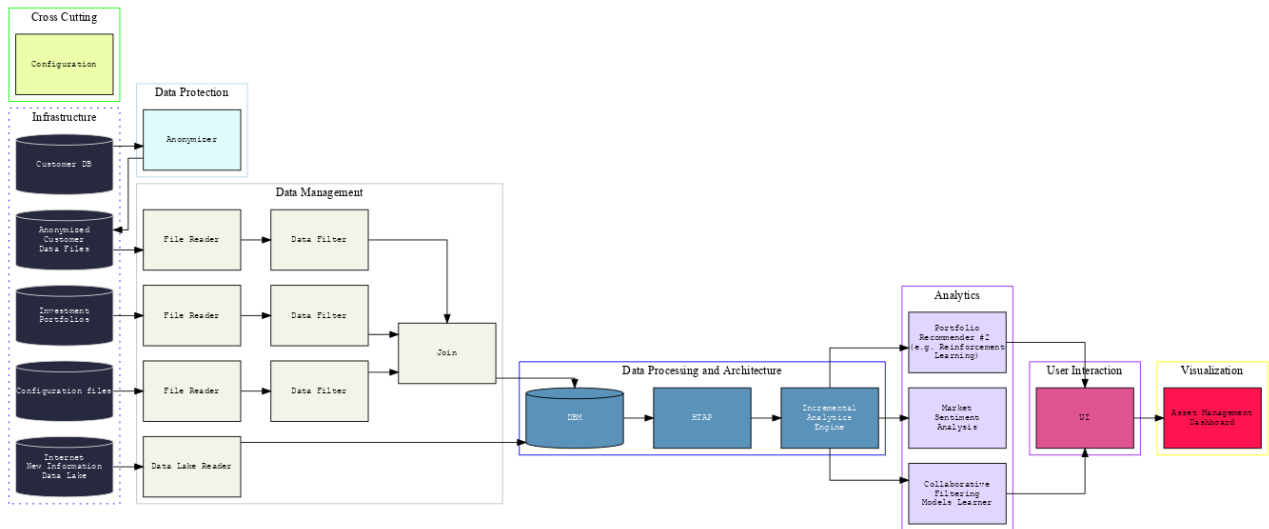


Figure 31: Logical Architecture of the Customer Centric Analytics Pilot in-line with the INFINITECH-RA

6.4 Pilot#4: Personalized Portfolio Management (“Why Private Banking cannot be for everyone?”)

6.4.1 Pilot Objectives

The main goal is to develop and adapt within SaaS based Privé Managers Wealth Management Platform a Portfolio Optimization algorithm (further on called Privé Optimizer or “AIGO”), as well as

improving and expanding its capabilities as an artificial intelligence engine to support better investment propositions for retail clients.

This pilot will explore the possibilities of AI-Based Portfolio construction for Wealth Management processes, regardless of the amount to be invested (therefore the slogan “Private Banking could be for everyone”). The AI-Based Portfolio Construction will enable advisors and/or end-customers, to use the existing Wealth Management Platform “Prive Managers” and make use of its risk-profiling and investment proposal capabilities, starting from his/her personal risk-awareness. AIGO allows for a variety of use cases which cater to the needs of financial advisors, end-clients and financial services companies. The innovative AIGO genetic algorithm can be used for proposing investments and evaluating them given an easy-to-use, personalizable set of criteria, in the form of so-called fitness factors. These fitness factors will be used to generate “health” scores for portfolios, which are used to define the “fittest” investments.

6.4.2 Data Sources

The data to be used by this pilot will be Customer Transactions Data, Financial Market Price Data, Financial Market Asset Master Data, Customer Risk Profile Data, Mutual Fund, ETF and Structured Products Breakdown Data, Customer Economic Outlook as well as single Account & Investors Data. All datasets will be stored within Privé SaaS solution in a cloud setup. Asset data and Client data are fetched from 3rd party databases and partially from selected market-data providers.

Risk metrics are calculated in the historical backtesting component for each single portfolio. A Genetic Algorithm component evaluates different Fitness Factors and generates a customised portfolio proposal.

6.4.3 Data Produced

JSON files will be produced from an API provided by Privé, that will also allow other 3rd party solutions to address this Portfolio Optimisation functionality. Based on this API PDF files can be generated for UI display and customer documentation.

The output data consists of the single portfolio holdings, their weights and amounts to decide about the Proposed Portfolio. Fitness Factors Scores and Total Fitness Score will be output for both the current and proposed (optimised) portfolio. For both Portfolios also Risk and Return metrics will be shown: 5 year annualized return, volatility and sharpe ratio.

6.4.4 Explainable Workflow

Starting from a client’s cash pool or current investments/portfolios, the user will select the fitness factors and constraints or preferences to perform the portfolio construction, basing themselves on the client’s risk profile and his preferences. The tool will run on a pre-set universe of assets taking into account all the input data and constraints. The AI genetic algorithm will generate a new proposal, where the selected preferences and risk parameters have been recognised. The optimisation tool can be run multiple times, after the necessary changes in initial parameters are made. This process can result in a UI proposal or a PDF generated investment proposal.

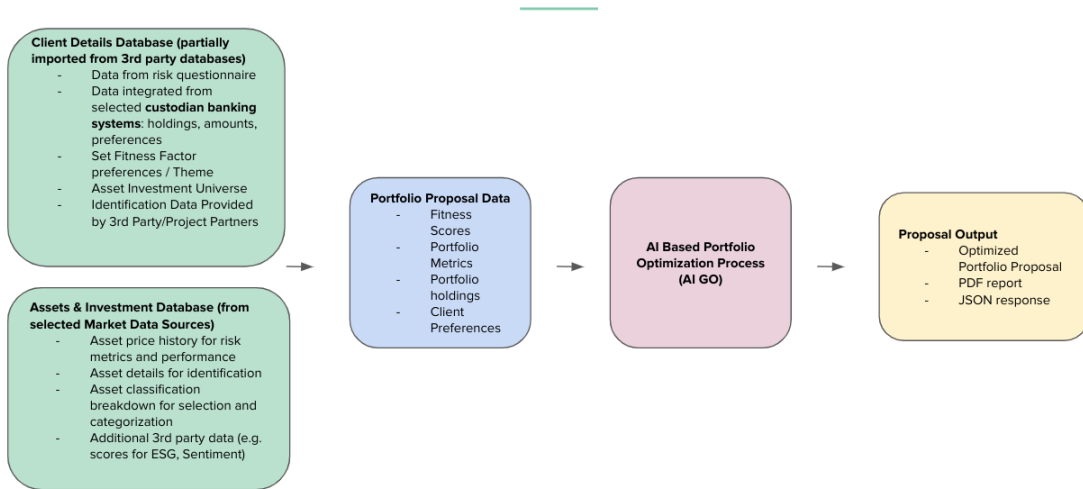


Figure 32: Explainable Workflow of the INFINITECH Personalized Portfolio Management Pilot

6.4.5 Logical Schema

An initial mapping of the explainable workflow of the pilot to the INFINITECH-RA layers and constructs is depicted in the following figure.

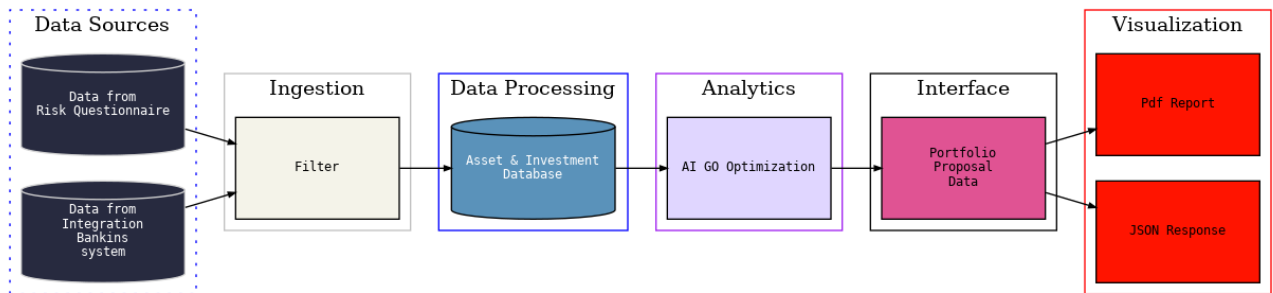


Figure 33: INFINITECH Personalized Portfolio Management Pilot Pipeline in-line with the INFINITECH-RA

6.4.6 Data Components

- Data from Risk Questionnaire (Risk Classification).
- Data from integration to banking system: holdings, amounts, preferences.
- Customer Risk Profile. Set Fitness Factor Preferences/Theme.
- Asset Price History for Risk Metrics and Performance.
- Asset Details for Identification.
- Asset Classification Breakdown for Selection and Categorization.
- 3rd Party Data (eg. Scores for ESG, Sentiment,...)
- Customer Transactions Data fetched directly from the Bank or an Asset Manager.
- Financial Market Price Data fetched from several Market Data Providers.
- Financial Market Asset Master Data fetched from several Market Data Providers.
- Customer Risk Profile Data fetched directly from the Bank or an Asset Manager.
- Mutual Fund, ETF and Structured Products Breakdown fetched from several Market Data Providers.
- Customer Economic Outlook fetched directly from the Bank or an Asset Manager based on questionnaires and Customer Profiles.
- Asset data and Client data are fetched from respective databases or custodian banks Risk metrics are calculated in portfolio historical backtesting component
- Genetic Algorithm component evaluates Fitness Factors and generates the portfolio proposal

6.5 Pilot #5b: Business Financial Management (BFM) tools delivering a Smart Business Advise

6.5.1 Pilot Objectives

Most of today's Financial Management tools for Small Medium Enterprises (SMEs) are geared towards analysing only past transactions, making such tools inadequate in today's world. Today, SMEs and their customers alike demand just-in-time processing, transparency and personalized services to assist SME owners not only in understanding better their SME business/financial health but also to be able to decide on the next best action to take.

Thus, Pilot#5b aims to assist SME clients of Bank of Cyprus (BOC) in managing their financial health in the areas of cash flow management, continuous spending/cost analysis, budgeting, revenue review and VAT provisioning, all by providing a set of AI powered Business Financial Management tools and harnessing available data to generate personalized business insights and recommendations. Machine learning algorithms, predictive analytics and AI-based interfaces will be utilized to develop a kind of smart virtual advisor with the aim to minimize SME business admin effort, to focus on growth opportunities and to optimize cash flows performance.

6.5.2 Data Sources

The following data sources will be integrated and used in the pilot:

- Transaction Data from BOC
- Transaction Data from Open Banking (i.e. PSD2 data)
- Transaction Data from SMEs (optional)
- Other Data (Market)
- Other Data from SMEs (optional)
- Accounts Data from BOC
- Accounts Data from Open Banking
- Customer Data from BOC
- Direct Input from SMEs

6.5.3 Data Produced

The pilot will combine the abovementioned diverse datasets in order to produce personalized business insights and recommendations for SME customers of BOC. Output data, as shown below, will be generated by the various engines in relation to cash flow predictions, budgeting, KPIs, benchmark(s) and transaction monitoring and categorization. The output data include:

- A JSON/CSV/XML file containing the obtained insights and recommendation to be provided to the respective SME
- A JSON/CSV / XML file containing the obtained cash flow related data to be provided to SME directly or indirectly
- A JSON/CSV / XML file containing the derived budget target for each category used by the respective SME
- A JSON/CSV / XML file containing results on Financial Health and Performance matrix
- A JSON/CSV / XML file containing results on abnormal transactions and suspicious expenses
- A JSON/CSV / XML file containing Matrix with invoice information and payment prioritization
- A JSON/CSV / XML file containing benchmarks that allow the SME to compare to likewise businesses

The data will be stored in the common datastore and be available to the end user (SME) via the Infinittech Reference Architecture (IRA) gateway and the banks middleware.

6.5.4 Explainable Workflow

Some of the available datasets require real time data collection, while in others historical data collection is sufficient to provide actionable business insights. In detail, transaction and account data related to the respective SME will be drawn from BOC’s repository by a real time/historical data collector as well as transaction and account data from Open Banking (PSD2), as well as BOC customer data, will utilize a historical data collector. In addition, an external data collector will also be used in order to integrate other related Open Banking/macroeconomic data. The SMEs data source (e.g. ERP/Accounting system) utilization remains optional as consent is required for the collection and processing of such data and its cloud availability being required. All data except external macroeconomic data will be pseudoanonymized (by tokenization) before being uploaded to the IRA.

The cloud Data Repository (within IRA) will then store all collected data, along with the generated insights, past SME financial actions (to measure at what degree the SME actions reflect the recommended insights), as well as minimum user input which is required. A continuous data streaming will connect the Data Repository with the various deployed BFM tools (machine learning algorithms), which would allow the retraining of the respective AI models and the generation of useful insights and recommended actions. A reverse data pseudoanonymization will then be applied before the processed data move to the bank middleware component that contains composite APIs and produces push notifications, all which will be offered to the SMEs via Android, iOS and web apps.

Upon SME user login the IRA is also accessed, insights/recommendations picked up from the cloud data repository and provided to the SME user.

6.5.5 Logical Schema

The following figure illustrates a logical view of the pilot system architecture in-line with the INFINITECH-RA.

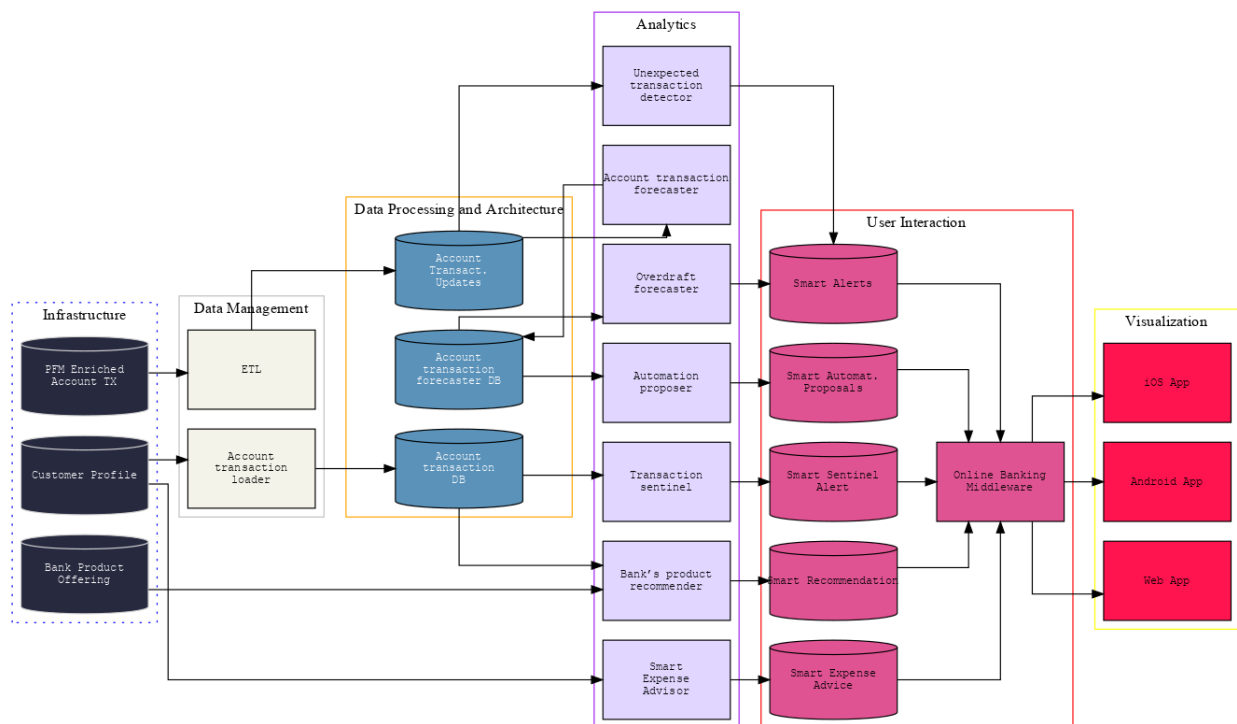


Figure 34: Business Financial Management Pilot Architecture in-line with the INFINITECH-RA

6.5.6 Components

The following components will be deployed and used in the pilot pipelines:

- Cash Flow Prediction

- Budget Prediction
- KPI Engine
- Transaction Monitoring
- Transaction Categorization
- Invoices Processing
- Benchmark(s)
- Recommender
- Data Streaming
- Generated Insights
- Generated Actions
- Collected Enriched Data
- Real Time Data Collector
- Historical Data Collector
- Gateway (Data Access Services)
- Orchestrator

6.6 Pilot#6: Personalized Closed-Loop Investment Portfolio Management for Retail Customers

6.6.1 Pilot Objectives

The goal of Pilot#6 is to create a system for personalized investment recommendations for the retail customers of the bank. NBG will leverage large customer datasets and large volumes of customer-related alternative data sources (e.g., social media, news feeds, on-line information) in order to make the process of providing investment recommendations to retail customer more targeted, automated, effective, as well as context-aware (i.e. tailored to state of the market).

6.6.2 Data Sources

Data that will be used for this pilot will be extracted from NBG Datawarehouse and several data sources: CRM Data, Deposit Account Transactions, Cards Transactions, Instruments Historical Prices, Investment Related Transactions, Instruments Characteristics and Instruments News feeds & Blogs sentiment analysis data.

6.6.3 Data Produced

Personalized investment recommendations for the retail customers of the bank, based on their Risk and transactions profiles. Banks relationship managers based on each customer risk and transactions profile, will be able to propose the possible alternatives of financial instruments that a customer will be interested to invest, with the relative prioritization. The proposed recommendations will be based on the instruments available from the bank, with the necessary input data for sentiment analysis for each financial instrument, based on the news & social feed, for the specific instrument (e.g. stock, bond, etc), or the relative instrument category.

6.6.4 Explainable Workflow

Data from NBG Datawarehouse related to Investment Products Clients (CRM Data, Deposit Account Transactions, Cards Transactions, Investment Related Transactions,) will be extracted in CSV files and utilizing the relative tools for data processing, anonymization and quality checking and cleansing will be imported to Leanxcale Datastore. Based on the data extracted for NBG Clients, through the Customer Risk Profile engine using data analysis tools, will be able to divide all customers in specific

profile clusters based on the investment & banking behaviour (MIFID questionnaires), deposit and card transactions, as well as investments transactions. Also, NBG will provide for each investment customer cluster profile the relative instruments that will be suitable for investment.

Based on the Risk profile of each customer, the relative financial instruments mapping with each customer cluster profile, using the Personalised Investment Recommendation Engine, will be possible for each relationship manager, thru the relative visualisation application, to provide a personalized investment proposal for each individual customer that will be interested for investment. The personalized investment proposal will be refined based on news & social feed data as sentiment analysis data, in order to take in consideration in real time any factors that may affect the relative mixture or financial instruments. The investment proposal will provide all the alternative financial instruments that are suitable for the specific customer, with specific grading (percentage) that fulfil both customer’s expectations, in combination with bank’s risk profile analysis.

6.6.5 Logical Schema

A first approach to mapping the pilot architecture to the INFINITECH-RA layers and pipelines approaches is illustrated in the following figure.

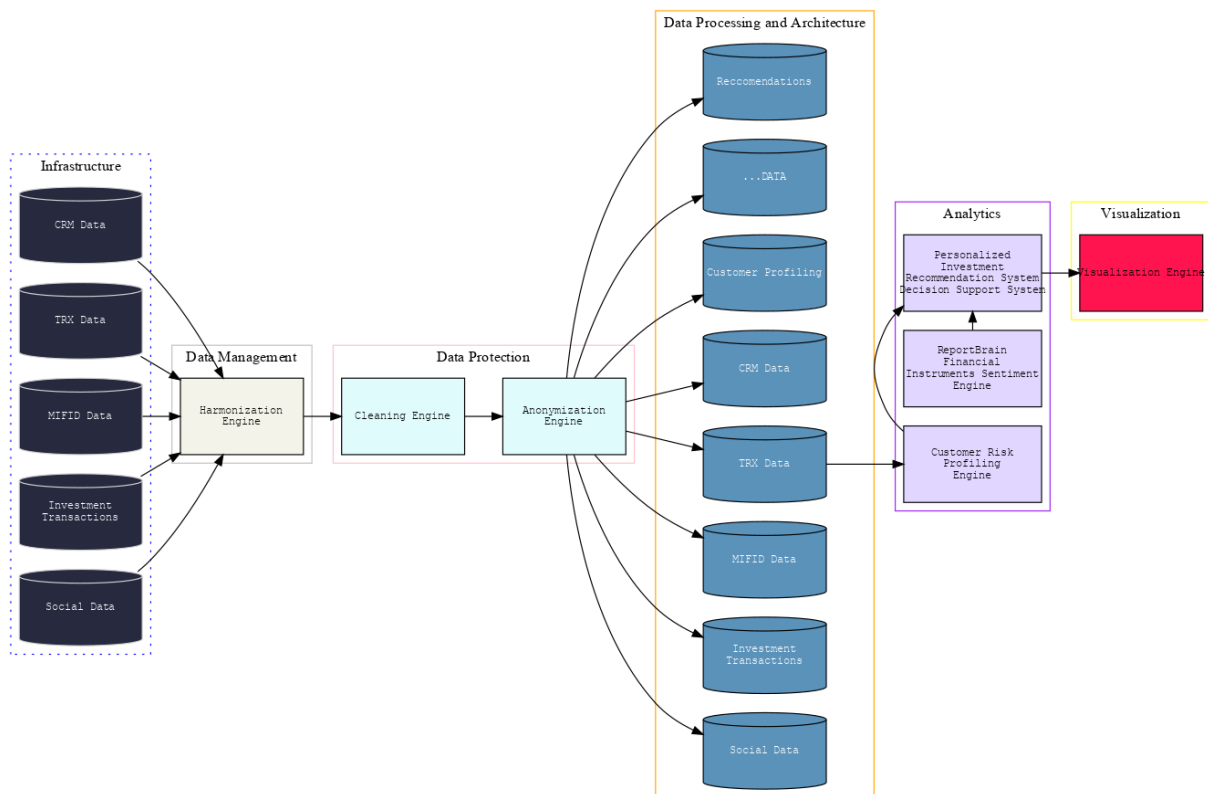


Figure 35: Personalized Closed-Loop Investment Portfolio Management

6.6.6 Components

The following components will be deployed and used in the pilot:

- Data Management & Anonymization Tool (Icarus)
- Customer Risk Profiling Engine
- Personalised Investment Recommendation Engine & Decision Support System
- Financial Instruments News & Social Feed Sentiment Analysis.

6.7 Pilot #7: Operation Whitetail – Avoiding Financial Crime

6.7.1 Pilot Objectives

Financial Crime like money laundering, terrorist financing, fraud, tax evasion, and identity theft, costs society billions and facilitates illegal operations. In addition to the extreme cost to society, Financial Crime costs the business in terms of reputational and fines. Over the last few years, financial institutions have been fined billions.

The Bank's customers make millions of transactions a day. The bank needs to be able to identify transactions that are moving the proceeds of crime (money laundering). Using advanced computing techniques could start creating a more up-to-date, accurate, versatile and complete view of the customer's profile and behaviour based on the near real-time analysis of transactions. This level of information could reveal abnormal behaviour earlier, and more accurately and enable the business to intercept criminal behaviour more effectively.

Criminals could be tracked down earlier, banks could avoid fines and a more precise alarming with less false positive reports of suspicious transactions will enable the bank investigating the critical alerts.

6.7.2 Data Sources

Functional Services of this pilot comprise the near real-time data access of transactions and external customer related data sources. The latter includes public sources or sources actively shared by the customer. Information from external sources will be obtained traditionally, e.g. credit reference agencies; sanctions lists.

These sources are used to update the customer risk profile based on ML/DL tools for the behavioural analysis and an advanced financial crime risk model for risk scoring and forwarding suspicious transactions for further investigation.

However, all data sources are pre-processed and synthesized in advance of being processed in INFINITECH. Thus, the pilot data sources are synthesized data sets of internal and external customer data and transaction data.

6.7.3 Data Produced

The workflow will produce data giving insight to the financial crime risk score. This may include a risk score, customer data, transaction patterns and details. The detailed data, which shall be produced, are yet to be specified depending on the advice of Financial Crime

6.7.4 Explainable Workflow

Within the pilot processes on these synthetic data address

- (i) KYC, screening available internal and external data sources (including next gen. information sharing capabilities) updating and enhancing the KYC data;
- (ii) KYT, extraction of transaction patterns from the synthetic transaction data;
- (iii) Advanced Financial Crime risk scoring, based on the transaction-based behavioural profile data and KYC results by an advanced financial crime risk
- (iv) Visualization of the risk score in a project related dashboard.

6.7.5 Logical Schema

The following figure presents an initial mapping of the pilot architecture to the INFINITECH-RA components and layers:

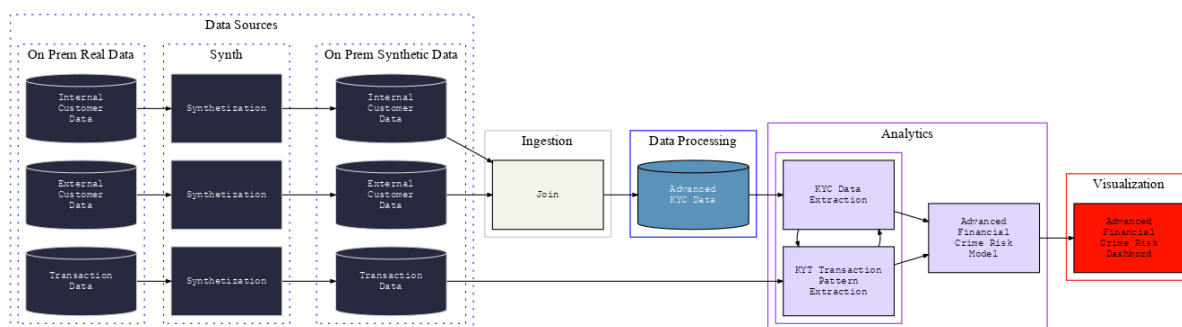


Figure 36: Logical Architecture of the Financial Crime Pilot in-line with the INFINITECH-RA

6.7.6 Components

A List of the main components to be deployed and used in the pilot follows:

- Preprocessing in advance of piloting: Synthesizing data
- Data enhancement (join) for advanced KYC Data
- Advanced KYC Data Lake
- KYC data extraction
- KYT transaction pattern extraction
- Advanced FC Risk Model
- Advanced FC Risk Dashboard

6.8 Pilot#8: Platform for Anti Money Laundering Supervision (PAMLS)

6.8.1 Pilot Objectives

The objective of the Pilot, is to develop a Platform for anti-money laundering Supervision (PAMLS), which will improve the effectiveness of the existing supervisory activities in the area of anti- money laundering and combating terrorist financing (AML/ CTF) by processing large quantity of data (Big Data) owned by the Bank of Slovenia (BOS) and other competent authorities (FIU).

PAMLS will contain the following functionalities:

- **Risk assessment tool:** to assess the risk of the financial institutions (its inherent risks and control environment) and a whole sector on a risk based approach
- **Screening tool:** for screening of payment transactions to detect potential suspicious transactions
- **Search engine:** allowing supervisor to look for a specific transaction or a sample of transactions
- **Distribution channel:** for gathering data from other departments within BOS, other supervisory authorities (domestic and foreign) and the financial institutions.

6.8.2 Data Sources

Relevant datasets that will be analysed within PAMLS:

TARGET2 transactions:

⇒ transactions executed by the Slovenian payment institutions

Confidential data.

Personal data will be

within TARGET2 (Trans-European Automated Real-time Gross Settlement Express Transfer System)	anonymized by the source, prior data delivery to PAMLS.
⇒ high value (above 50.000 EUR), urgent transactions in EUR	
⇒ transactions processed through BOS payment systems (responsible BOS Payment Settlement and Systems department - PPS)	
SEPA transactions:	Confidential data.
⇒ transactions executed by the Slovenian payment institutions within SEPA (Single Euro Payments Area)	Personal data will be anonymized by the source, prior data delivery to PAMLS.
⇒ domestic and international transactions within SEPA area in EUR under 50.000EUR value	
⇒ transactions processed through payment systems by third party provider	
FIU transactions (public data):	Public data.
⇒ transactions related to high risk countries above 15.000 EUR reported to the Slovene Financial Intelligence Unit (FIU)	
FI identification data	Confidential data.
⇒ identification information about Financial Institution (FI)	
⇒ aggregated statistical data on the FI inherent risk and control environment (number of clients, number of Suspicious transactions reports (STR) etc.)	
⇒ FI reports to the BOS (reports are confidential)	
ePRS data	Public data.
⇒ Slovenian Business Register (public data on legal entities)	
eRTR data	Public data.
⇒ Slovenian Transactions Accounts Register (public data on legal entities)	
High risk country list	Public data.
⇒ List of countries defined as high risk due to lack of or not effective AML/CTF system	
⇒ List is managed and published by the Slovene FIU (public data)	

6.8.3 Data Produced

Ongoing risk assessment for the purpose of the Anti-Money Laundering and Combating Terrorist Financing Supervision over the FI and FI sector.

6.8.4 Explainable Workflow

PAMLS will use various data sources, which we can divide in to three groups.

- ⇒ The first group consists of transactional data process through payment service providers in Slovenia (TARGET and SEPA transactions). This group will first be anonymized, then joined with High risk country List and stored in PAMLS internal data storage.
- ⇒ The second set of Data sources represents public data (FIU transactions and data from public registers) which will be joined with High risk country List and stored in PAMLS internal data storage.

⇒ Third group of data sources represents FI identification data, which will be joined with High risk country List and data from public registers and then stored in PAMLS internal data storage.

After data is ingested in PALMS platform, it needs to be preprocessed in a way, that information is properly enriched and it needs to be provided in a suitable data format (vectors, graphs).

Process of feature engineering, tailored to specific goals, will follow. PAMLS will develop and test novel approaches for detecting suspicious behaviours, which could be labelled as high risk later in the process and will have an effect on final FI risk assessment.

Part of the PAMLS is also Risk Calculation engine. There the risk calculation will be continuously calculated on a level of a sector or a particular FI, using predefined Risk Assessment methodology.

To empower bank analysis, to develop and test novel approaches, PAMLS provides three components: Stream story, Pattern discovery & matching, Anomaly detection & prediction. With introduction of enriched graph topologies and a hierarchical Markov chain models, PAMLS will capture the qualitative behaviour of the systems’ dynamics and enable analyst to discover new regularities and correlations on a larger scale.

These components will enable iterative development of potentially new upgrades to existing Risk Assessment methodology and discovery of novel or additional money laundering and terrorist financing typologies.

PAMLS will also provide three different user interfaces, which corresponds to 3 different use cases.

6.8.5 Logical Schema

The following figure illustrates an initial logical mapping of the pilot components to the layers and pipelines approach of the INFINITECH-RA.

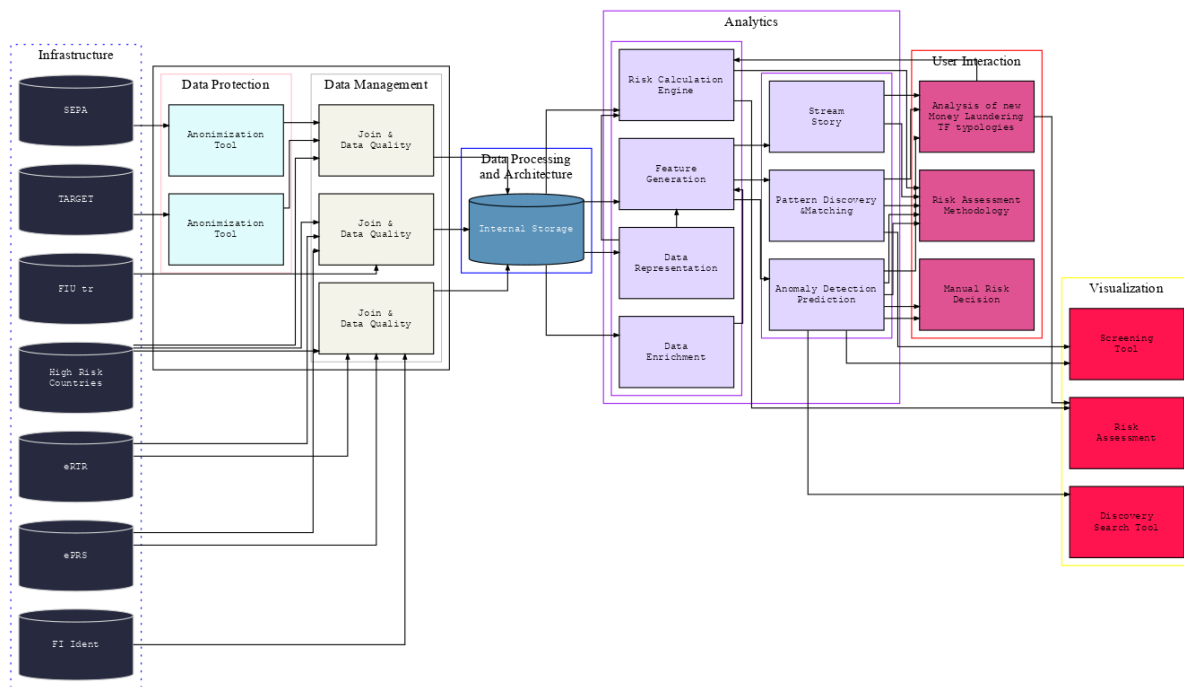


Figure 37: PAMLS Pilot Logical Architecture in-line with the INFINITECH-RA

6.8.6 Components

Anomaly Analysis provides two main functionalities:

- Anomaly detection
- Anomaly prediction

Pattern Analysis provides two main functionalities:

- Pattern matching
- Discovery

Stream Story is a component for the analysis of multivariate time series

6.9 Pilot #9: Analyzing Blockchain Transaction Graphs for Fraudulent Activities

6.9.1 Pilot Objectives

Blockchain crypto currencies and tokenized assets that are obtained fraudulently can go through various transfers on the blockchain and also end up as stable coins (e.g. USD, EUR, TRY tokens) in different jurisdictions. As a result, it is possible that a company may be deposited crypto currencies and stable coins that can be traced to addresses involved in fraudulent activities. Pilot #9 will develop a parallel and scalable system that can construct the massive Bitcoin and Ethereum blockchain transaction graph with distributed dynamic data structures on an HPC cluster. The pilot will perform analysis using graph and machine learning algorithms on the transaction graph. The pilot will also provide a user interface producing reports and visualization of results using graph drawing package.

6.9.2 Data Sources

Pilot #9 will use the following data sources: (i) Public Bitcoin blockchain transactions (ii) Public Ethereum blockchain transactions (ether as well as several stable coins transactions such as USD, EUR, TRY ERC20 tokens) (iii) blacklisted blockchain addresses collected from the Internet.

6.9.3 Data Produced

Pilot #9 will generate paths and subgraphs tracing a given blockchain address to blacklisted addresses. Transaction graph analysis may also produce additional illicit addresses which can be added to blacklisted address list.

6.9.4 Explainable Workflow

Pilot #9's Blockchain Transaction Dataset Preparation Component will parse raw blockchain data and extract Bitcoin, Ethereum and major ERC20 token transactions (such as Gemini USD (GUSD), Tether USD (USDT), Tether Gold (XAUT), Stasis Euro (EURS) and Turkish BiLira (TRYB)) that come from the Bitcoin and Ethereum Mainnet blockchains. After retrieving all the blocks up until now, this component will be run periodically to retrieve newly generated blockchain blocks during the period.

Scalable Transaction Graph Analysis Component of the pilot will take the massive bitcoin and Ethereum public transaction dataset. Since the transaction graph size is massive and dynamically growing, it will construct distributed and partitioned transaction graph in parallel using MPI message passing libraries in order to achieve scalability. It will also input blacklisted addresses and utilize graph and machine learning algorithms to analyse fraudulent transactions.

Finally, the User Interface for Blockchain Transaction Reports and Visualization Component of the pilot will be responsible for providing interaction services with the Scalable Transaction Graph Analysis component. It will collect and manage annotated blacklisted addresses from the Internet. It will submit queries consisting of blockchain addresses and blacklists to the transaction graph analysis component. It will then generate web-based reports and visualization based on received results which will be in the form of paths and subgraphs tracing a given blockchain address to blacklisted addresses.

6.9.5 Logical Schema

The following figure illustrates how the pilot architecture can be expressed in terms of the layers and the pipelines approach of the INFINITECH-RA.

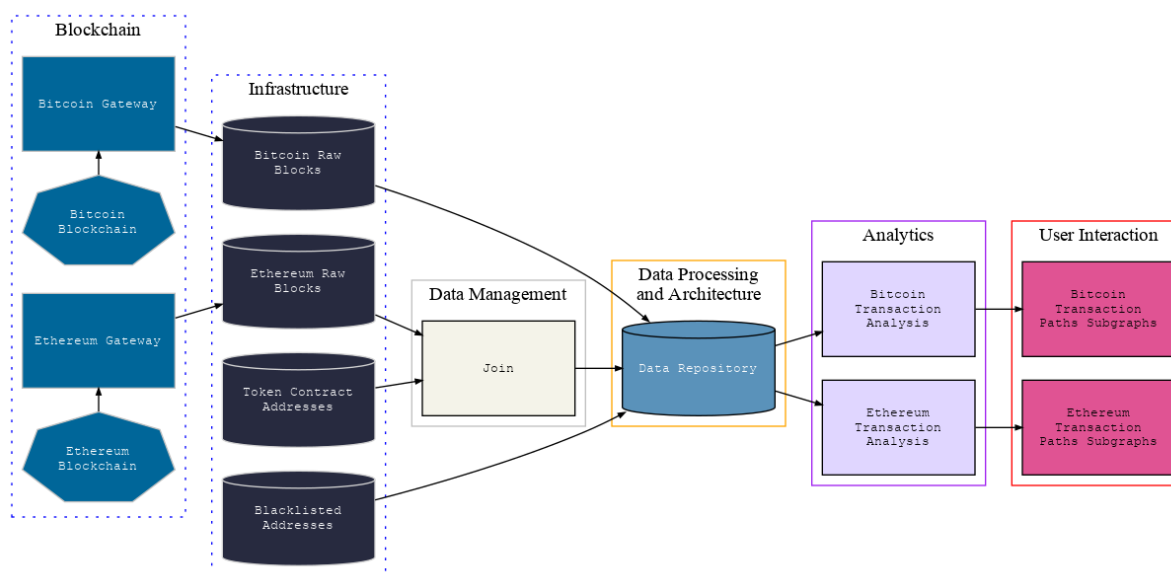


Figure 38: Logical View of the Blockchain Transaction Graphs Analysis pilot in-line with the INFINITECH-RA components

6.9.6 Data Components

The following components will be developed, deployed and used in the pilot:

- Blockchain Transaction Dataset Preparation Component
- Scalable Transaction Graph Analysis Component
- User Interface for Blockchain Transaction Reports and Visualization Component

6.10 Pilot #10: Real-time cybersecurity analytics on Financial Transactions' BigData

6.10.1 Pilot Objectives

The main objectives of this pilot use case are:

- The early detection of new and subtle types of frauds. Since fraudsters keep innovating novel ways to scam people and online systems, it becomes crucial to apply AI/ML methods to detect outliers in large transactional datasets and be robust to changing patterns.
- The reduction of the number of false positives which are usually analysed to understand if they are real fraud attempts or not. To this aim, it is very important to be able to train, validate and test ML models to make the most accurate ones operational.
- adaptivity is the key to follow ever changing fraud schemas: software modular approach, allows to implement and simply "plug-in" new algorithms to move with the times.

For both goals ML algorithms and BDA (Big Data Analytics) stream and batch services will be used on data sets made at disposal.

- 1) **Internal transfer of funds** (enabling transfer of funds from the Bank current account of one payer to that of another Bank account holder)
- 2) **SEPA transfer** The Single Euro Payments Area is a payment-integration initiative of the European Union for simplification of bank transfers denominated in euro
- 3) **Foreign transfer** is a tool used in the interbank system to transfer money between countries

- 4) **SMWCA (Sending money to people who do not have a current account)** is a means of payment, through which you can send a sum of money even to those who do not have a bank account
- 5) **PCTU** Top-ups phone credit
- 6) **STFTS (Secure telematic fund transmission system)** The procedure requires that a beneficiary is selected, the way he/she prefers to receive the money and the amount to be sent. You then have to pay the amount due, enter and check the personal data and send)

6.10.2 Data Sources

The data sets in input of the workflow are related to several types of transactions:

- *Bank Transfer SEPA* (The Single Euro Payments Area (SEPA): a payment-integration initiative of the European Union for simplification of bank transfers denominated in euro. SEPA covers predominantly normal bank transfers.
- *Foreign Bank Transfers*: international transfer of funds between banking accounts held in financial institutions located in different countries
- *Internal transfers of funds*: transactions involving transfer of funds from the Bank current account of one payer to that of another Bank account's holder
- *PCTU*: Top-ups phone credit
- *SMWCA (Sending Money to people Who do not have a Current Account)*: is a means of payment for transferring money to those not holding a bank account
- *STFTS (Secure telematic fund transmission system)* Neither the sender nor the receiver of the money being transferred need to hold a banking account

For each data set, when appropriate, data as name, date of the operation, IBAN of the beneficiary will be considered.

Even though the datasets that will be used in Pilot #10 are Fully Synthetic ones, therefore no specific regulation could be applied to them, we are considering our data as real: that's why we identified some macro requirements to get compliance of Pilot #10's tools to applicable regulations.

6.10.3 Data Produced

Data on transactions, properly preprocessed, will be used to train a prediction model that in turn will be adopted to predict cases of suspected fraudulent transactions. Such model is expected to be periodically retrained through new transaction data that will have been stored.

On the other hand, data on suspicious frauds detected by a real time prediction algorithm will be delivered to fraud analysts by a friendly user interface, in order to allow them to adopt proper countermeasures just-in-time and not when it is not too late and fraud has been monetized.

6.10.4 Explainable Workflow

A preliminary stage of the batch workflow subsequently loads the data sets related to several types of transactions (bank transfer SEPA, foreign bank transfers, internal transfers of funds, PCTU, SMWCA, STFTS) into the HDFS storage by means a certain ingestion job.

Then batch data sets are properly filtered (to remove some columns and rows unneeded for the ML) and joined to get only one unlabeled data set.

In this phase the goal is to cluster such data, that is to create labelled samples to feed the supervised machine learning classifier of the next phase. We are supposing that it will be needed to analyze datasets made of mixed-type data, where numeric and nominal features coexist.

For this aim, the OneHotEncoder algorithm has been adopted, in order to convert categorical variables in numerical ones and to get the proper input numerical data set to feed the KMeans algorithm (unsupervised clustering technique).

After analyzing the resulting clusters, PI (as a domain expert) determines which clusters present a fraudulent behavior and assign proper labels to such clusters.

Then, a supervised machine learning technique is used to predict new unlabeled transaction data high-level information. A typical machine learning approach includes two phases: training phase and testing phase. In the training phase, the Random Forest (RF) classifier is trained using labelled data sets. In this approach, the classifier is trained by the labelled data subset taken from the previous phase (clustering-based labelling).

After learning the mapping, the Random Forest (RF) classifier can map new real time unlabeled transaction data to their corresponding high-level information (i.e. label) on the basis of the RF model trained in the batch layer, so to detect financial fraud events while they happen and visualize such events.

It is worth noticing that in order to process such transactions, a data model will be defined so that transaction data can be properly processed by the Random Forest predictor.

6.10.5 Logical Schema

An initial mapping of the pilot’s components and modules to the INFINITECH-RA pipelines approach is illustrated in the following figure.

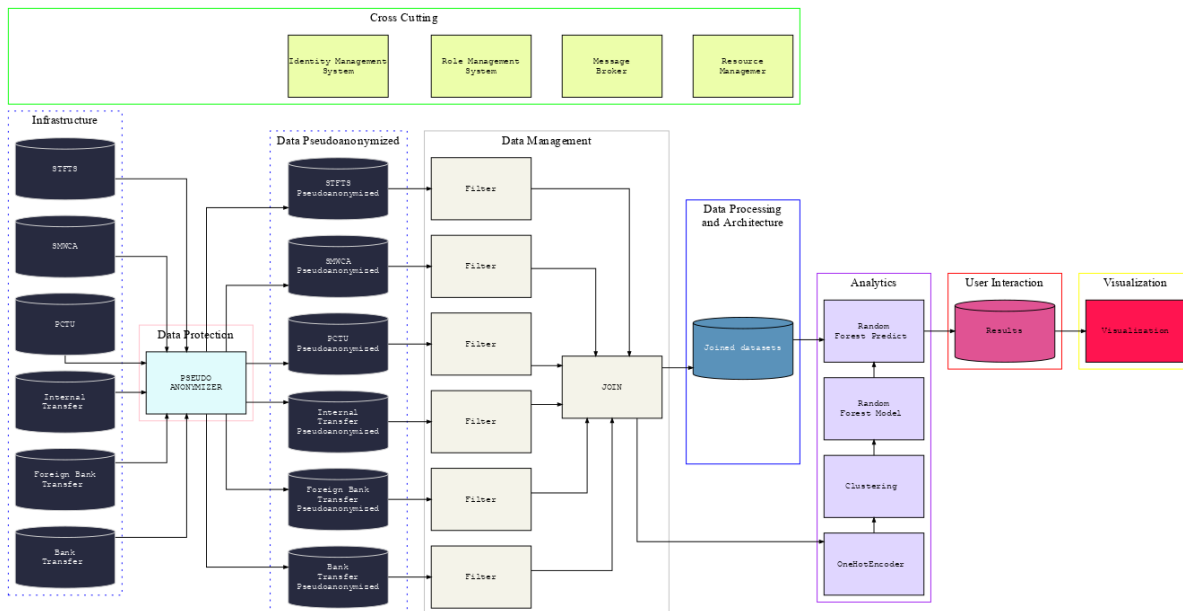


Figure 39: Logical Architecture of the Real-time cybersecurity analytics Pilots in-line with the INFINITECH-RA

6.10.6 Components

The list of main components to be deployed and used in the pilot includes:

- **Identity Management System:** It is a cross cutting system that guarantees user authentication, authorization and management. It allows or denies the access to the federated services that run within the architecture.
- **Role Management:** It implements and handles the roles and the privileges that can be associated to the users. It is often tightly coupled with the Identity Management System.
- **Message Broker:** works as an intermediary software that allows system components to communicate each other effectively, implementing a common communication protocol over message buses.

- **Resource Manager:** It is a lower level software that consents to handle and use the infrastructure resources seamlessly and dynamically, according to the number of requests received per time interval.
- **Pseudoanonymizer:** tool to pseudonymize personal or sensitive data at source, in order to preserve privacy according to GDPR regulation
- **Filter:** Filtering component to remove specific rows and columns
- **Join:** Service to join two or more datasets where at least one column must be the same
- **OneHotEncoder:** Service to transform categorical variables into numerical ones
- **Clustering (Kmeans):** Given a set of observations (x_1, x_2, \dots, x_n) , where each observation is a d -dimensional real vector, k -means clustering aims to partition the n observations into k ($\leq n$) sets $S = \{S_1, S_2, \dots, S_k\}$ so as to minimize the within-cluster variance.
- **Random Forest:** An ensemble learning method for classification, regression and other tasks that operate by constructing a multitude of decision trees at training time and outputting the class that is the mode of the classes (classification) or mean prediction (regression) of the individual trees. Random decision forests correct for decision trees' habit of overfitting to their training set
- **Results:** Storage that contains all the processed data elaborated by the workflow.
- **Visualization:** *The service that gathers the resulting datasets to be delivered to the visualization clients.*

6.11 Pilot #11: Personalized insurance products based on IoT connected vehicles

6.11.1 Pilot Objectives

In a few words, this pilot aims to develop new services for driving insurance companies, based on the information gathered from a connected vehicle, as an IoT ecosystem. Current driving insurance services try to reward good drivers against the “bad one”, but based on very static or historical information: your age, colour of your car, incidents by year, etc. A new approach, more dynamic, adapted and custom services are needed. You pay as you drive, in a similar approach to a cloud word, where you pay as you consume.

Complementary to this, a second service will help to detect possible fraud's situation. Fraud causes not fair costs to the company that would affect indirectly to the good/honest drivers.

In both use cases the underline technology is based on connected vehicles, IoT and BigData, because of the expected amount of data to be managed. The business analysis part, which determines how good driver you are, and the detection of possible frauds, will be based on AI and ML techniques. Due to the personal data managed in the pilot, security and privacy will be also a technology challenge to achieve.

6.11.2 Data Sources

The main data source in the pilot is produced by the connected vehicle, with about 20 vehicles. It is under study, the inclusion of some historical data provided by vehicles from other previous project; if legally possible. The data produced by the connected vehicle includes: CAN data, traffic events, gps, speed, etc. Complemented with data provided by the city of Vigo.

Finally, the data will be complemented with some synthetic/simulations of vehicles trips. Based on an opensource tool, SUMO²⁰ and a custom developed adaptor to integrate and transform the data, according to expected data pipeline and data standards.

About data standards, the connected vehicle platform layer, in charge of gathering, homogenizing, filtering, etc, is based on a FIWARE platform. Therefore, FIWARE Data Models²¹ will be used during the

²⁰ SUMO tool: <https://sumo.dlr.de/docs/index.html>

project. In that point, it is expected to contribute back to these standardization efforts fostered by the FIWARE Foundation. Some models would be adapted, or new ones would be created.

6.11.3 Data Produced

Two main business services will be produced during the pilot's implementation. Therefore, it is not so focused on producing data, but, to provide services use. These services will be used, internally, by the insurance company. In any case, the data produced (or the results) by these services would be considered as data produced, that can be stored in a database, to feed new chains/workflows.

Pay-As-You-Drive service:

- Input: driver's trip info
- Output: a value from 0 to 100 about the driver's behaviour.

Fraud detection:

- Input: driver's trip info
- Output: a kind of driver.

It would be used to compare the kind of driver against an historical register.

Example of usage in case of an accident: it would check if the kind of driver differs from previous days (stolen vehicle, identity theft).

6.11.4 Explainable Workflow

The data collected from the vehicle is transmitted to the INFINITECH Testbed, where the data is pipelined into a workflow with a set of steps. Before going to the Connected Vehicle platform, data is prepared about regulation and anonymized to protect the driver's privacy. With the data prepared to be managed, the Connected Vehicle Platform homogenize, filter, clean, and standardize the data (based on FIWARE Data models). Here the data is prepared as time series for real time management, or, it is stored as historical information.

Looking at the AI Platform, it is expected to develop/train two different ML models for the two business services. Once the models have been implemented and these are available in the platform, these models will be trained, supported by the previous data gathering workflow. Getting the training data from the Connected Vehicle Platform.

It is important to clarify that the training process is not a matter of getting data, training and finish. The model will be constantly trained according to specific scheduling. The model will be always updated with the new data that constantly is generated by the connected car:

- (1) Data Management: data produced -> pre-prepared -> gathered -> streamed or store
- (2) (scheduling time raises)
- (3) ML Model training: data features extraction -> training model -> store the model

The usage of the model, or inference service, or business service, it is an independent workflow. It just deploy a service that will exploit the previously trained model. These are interconnected, the first time the services are deployed with the model, this is linked to future training. When new training succeeds with more accurate models, the inference service will update the resulting model automatically.

6.11.5 Logical Schema

An INFINITECH-RA compliant architecture of the pilot is depicted in the following figure:

²¹ FIWARE Data Models: <https://github.com/FIWARE/data-models>

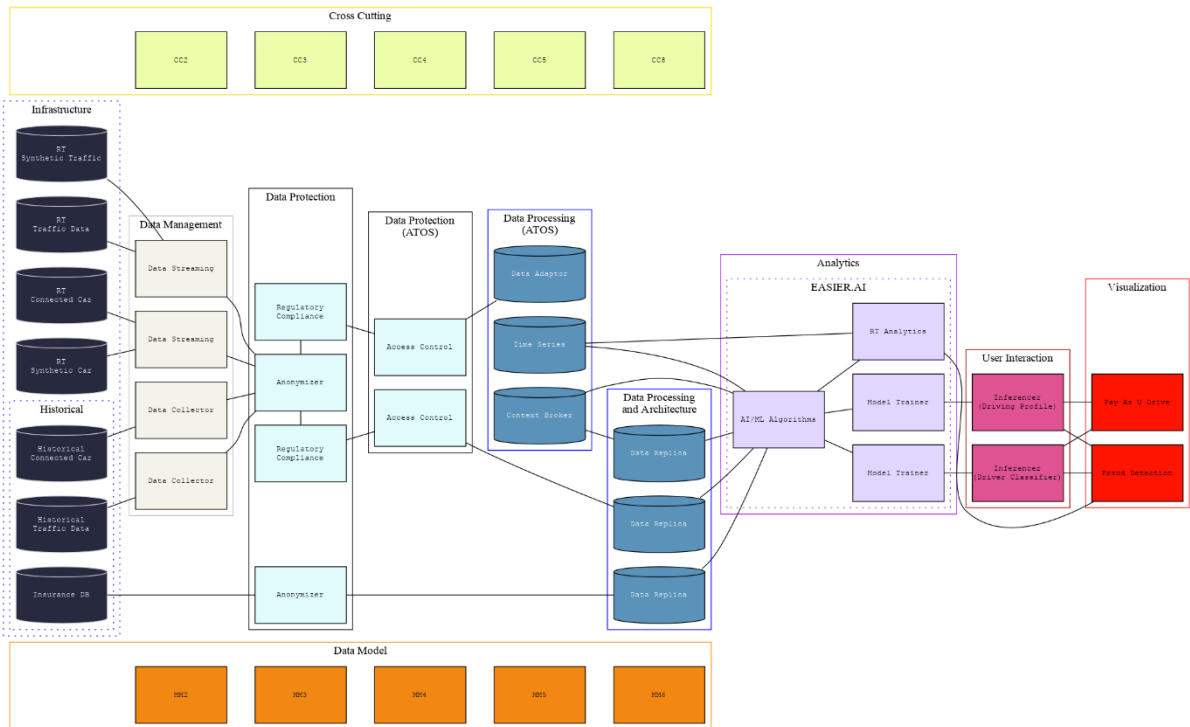


Figure 40: Logical Architecture of the Personalized insurance products based on IoT connected vehicles pilot in-line with the INFINITECH-RA

6.11.6 Components

The main components to be used in the pilot include:

- Anonymization tool.
- Connected Data processing.
- AI models, trainers and inferencers.

6.12 Pilot #12: Real World Data for Novel Health-Insurance products

6.12.1 Pilot Objectives

This pilot aims to develop new services for health insurance companies, based on the information gathered from current or prospective clients. Current health insurance services take into account clients’ medical history in a very static way. A new approach, more dynamic, adapted and custom services are needed. Premiums are calculated dynamically based on clients’ habits.

Complementary to this, a second service will help to detect possible fraud’s situation. Fraud causes not fair costs to the company that would affect indirectly to the good/honest clients.

In both use cases the underline technology is based on analysing Real-World Data (RWD) of the clients. The business analysis part, which determines how healthy a client’s lifestyle is, and the detection of possible frauds, will be based on ML techniques. Due to the personal data managed in the pilot, security and privacy will also play an important role.

6.12.2 Data Sources

The main data source in the pilot is the RWD collected by Healthentia. Healthentia is a platform for measuring and reporting RWD. Measurements are based on sensors on smartphones or IoT wearable devices. Reports employ questionnaires that the clients periodically answer utilising the Healthentia app. A secondary source of data is the records of the clients of the health insurance companies. Finally, the data will be complemented with synthetic/simulated data.

6.12.3 Data Produced

Two main business services will be produced during the pilot’s implementation. Therefore, it is not focused on producing data, but on service provision. These services will be used by the insurance company. To have these services, the ML module will be producing models which can be considered as data produced, that can be stored in a database, to feed new chains/workflows.

Risk assessment service:

- Input: client’s long term behavior.
- Output: a value from 0 to 100 enumerating client’s lifestyle.

Fraud detection:

- Input: client’s long term behavior, models of all clients.
- Output: probability of fraud, enumerating mismatch of current behavior from past behavior of client and other clients.

6.12.4 Explainable Workflow

The RWD collected from the client using Healthentia and the secondary sources is transmitted to the INFINITECH Testbed, where they are aggregated together, anonymised for protection and stored. Stored data are either used to (re)train the risk and fraud assessment models, or are used by the models to give the services’ outputs. These outputs are offered via the presentation layer of the pilot, together with all collected RWD for human insights/verification, to the system users, employees of the insurance company.

6.12.5 Logical Schema

An initial mapping of the pilot architecture to the INFINITECH-RA is depicted in the following figure.

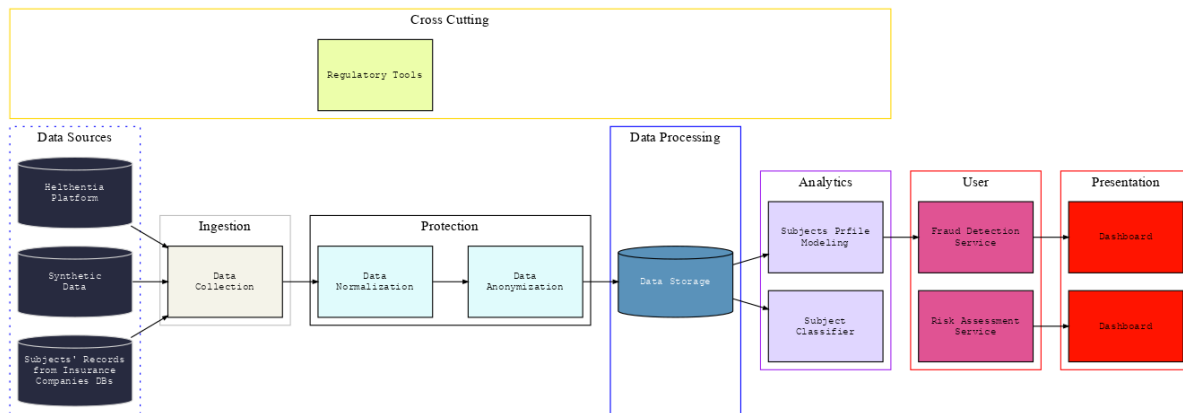


Figure 41: INFINITECH-RA Compliant Architecture for the Real World Data for Novel Health-Insurance products pilot

6.12.6 Components

The following components will be deployed and used as part of the pilot:

- Anonymization tool
- AI models, trainers and inferencers
- Visualization tools

6.13 Pilot #13: Alternative/automated insurance risk selection - product recommendation for SME

6.13.1 Pilot Objectives

The pilot will implement an automation of the subscription process helps the insurance company reduce costs. In addition, being able to verify that the data entered is correct with a double verification avoids possible errors in the cost of the insurance premium

The monitoring and identification of real-time risk changes allows the company to know if the insurance cost really corresponds to the real risk of the SME or if it should increase or decrease it to adapt it to its current situation.

6.13.2 Data Sources

Data will be extracted from open sources such as company websites, official registers, social networks, opinion forums, etc.

SMEWIF	SMEs website information and functionalities
ROPS	Review and opinions platforms
EUBD	European SMEs Business Directories
GIO	SMEs geolocation information and characteristics
SMSIP	Social media SMEs information and presence

The Pilot will also use synthetic data

P13-Alternative/automated insurance risk selection - product recommendation for SME SMEs synthetic raw data

6.13.3 Data Produced

The output will be the ERP (Enterprise Risk Profile) and EIAU (Enterprise Insurance Automated Underwriting).

With these two outputs, not only the risk profile and its levels of any SME company are obtained, but also the information for the automation of the subscription and the application of rules to obtain the price automatically.

6.13.4 Explainable Workflow

Pilot #13 is a "Big Data" data analysis platform applying ML (Machine Learning) and AI (Artificial Intelligence) technologies to better predict the insurance needs of SMEs.

Well, this system must be prepared to offer a commercial use to the companies, so it must have a user interface so that they can manage the information (Frontend) and a management layer at a logical level (Backend).

The companies (enterprises) will access our platform through a registration process and subsequent validation by assigning a package of number of customers, the basic and commercial information will be recorded in Amazon Cognito and the logical information of the company will be recorded in a table of DynamoDB called Enterprises.

With regard to the use of the information by the companies, the user must load the information they have stored in their systems in our platform, this will receive the name of raw data (crude-data). The raw data will be uploaded to the platform as structured information in CSV format. The companies

that use our service will have a limited amount of clients loaded in crude-data, for this, the fields of the Enterprises table, limit, clients_uploaded, total_clients_uploaded will be used in a monitored way.

Each row of this document will identify a client, which can be target in different sources of information on the Internet and other open sources in real time, depending on the information available (the quality of information depends on the company), which will be recorded in the DynamoDB Targets table (Infrastructure).

Once the robots can obtain the least updated target for a particular source, they will proceed to obtain the information and subsequent storage in big data, this information container will be in Amazon S3 in a loop called big-data as well as stored in the folder with the name of the company's identifier. The information obtained from the source will be stored in the folder mentioned above in a JSON document whose name will be the user's identifier. (Data Management)

From this layer the analysis algorithms will be applied and the results of the analyzed companies will be shown with the indicators of risk levels and the configuration and automation of the underwriting obtaining the ERP (Enterprise Risk Profile) and EIAU (Enterprise Insurance Automated Underwriting) (Data Processing and Visualization).

It is important to note that the quotes to services provided by AWS are for illustrative purposes and provided they have the same technical and technological characteristics they can be replaced by another supplier, as in the case of the project under consideration may be NOVA and LeanXcale.

6.13.5 Logical Schema

The following figure provides an INFINITECH-RA compliant logical view of the logical architecture of the pilot.

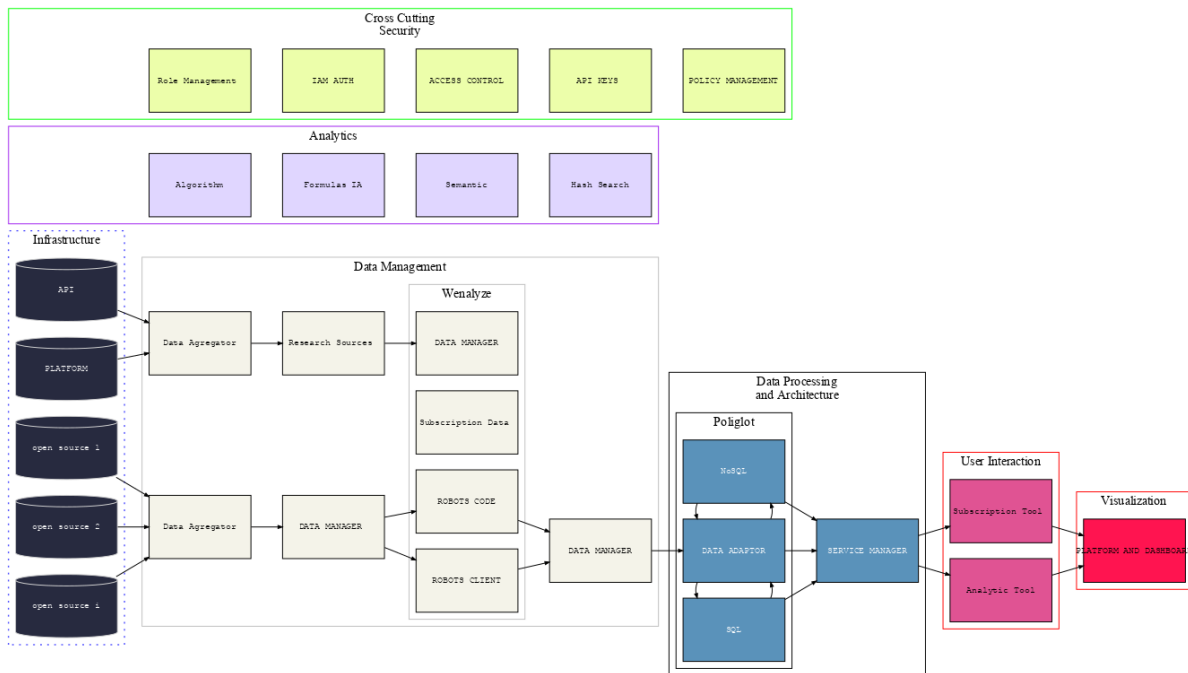


Figure 42: INFINITECH-RA Compliant Architecture for the Alternative/automated insurance risk selection - product recommendation for SME Pilot

6.13.6 Components

The following INFINITECH technologies and component will be used as part of this pilot:

- Polyglot Query Processing
- Relational data base (SQL and NoSQL)

6.14 Pilot #14: Big Data and IoT for the Agricultural Insurance Industry

6.14.1 Pilot Objectives

The objective of Pilot #14 “Big Data and IoT for the Agricultural Insurance Industry” is to deliver a commercial service module that will enable insurance companies to exploit the untapped market potential of Agricultural Insurance (AgI), taking advantage of innovations in Earth Observation (EO), weather intelligence & ICT technology. EO will be used to develop the data products that will act as a complementary source to the information used by insurance companies to design their products and assess the risk of natural disasters. Weather intelligence based on data assimilation, numerical weather prediction and ensemble seasonal forecasting will be used to verify the occurrence of catastrophic weather events and to predict future perils that could threaten the portfolio of an agricultural insurance company. The INFINITECH AgI-module derived indices will allow and enable the agricultural insurance industry to enlarge its market, while delivering a larger portfolio of products at lower costs and serve areas, where classical insurance products could not be delivered.

6.14.2 Data Sources

The main data source for the pilot is produced by satellite and a weather intelligence engine. The Earth Observation (EO) data will be derived from the satellites Sentinel-1,2,3, LandSat-8, MODIS and PROBA-V). Also, numerical weather predictions for the pilot areas (gridded data) are generated each day and will replace the previous prediction. Lastly, gridded climate indices based on ERA-5 Land and ERA-5 Reanalysis Data will be used for the pilot.

6.14.3 Data Produced

The data produced will result in a solution for Agricultural Insurance companies allowing them to efficiently couple EO satellite data and weather/climate data with any type of complementary data (from separated drone shots to ultra-high-resolution SAR imagery). The INFINITECH AgI module will enable Insurance companies to alleviate the effect of weather uncertainty when estimating risk for AgI products, reduce the number of on-site visits for claim verification, reduce operational & administrative costs for monitoring of insured indexes and contract handling, & design more accurate & personalized contracts.

By deriving impartial indices on top of a multitude of data, the module will allow insurers to reduce significantly the time needed for handling and verification of claims and the costs imposed by fraud, moral hazard and adverse selection.

6.14.4 Explainable Workflow

The data produced by the Octopush EO Service (Crop Monitoring, Pest & Disease Services, Damage Assessment Services) and the Agro Apps Weather Intelligence Service (Weather Forecast Services, Climate Services) is pipelined into a Data integrator. The integrator feeds back Metadata and information on the Area of Interest (AOI) to the data producing services.

The integrator itself is retrieving the AOI & Metadata from a Business DB storage layer. The geospatial data storage and data persistence mechanisms allows the storage of the geometries and zonal statistics and provides the essential functionality for querying and retrieving data via an API (alerts) or WMS server components (vectors source). The WMS server is then responsible for rendering and serving of the GIS layers to the User Interface.

The restful API will act as a communication and data exchange bridge, that allows the platform to share processed and structured content internally, between the different components. The front-end user interface is the gateway responsible to present all the system data through user-friendly controls and web mapping interfaces.

6.14.5 Logical Schema

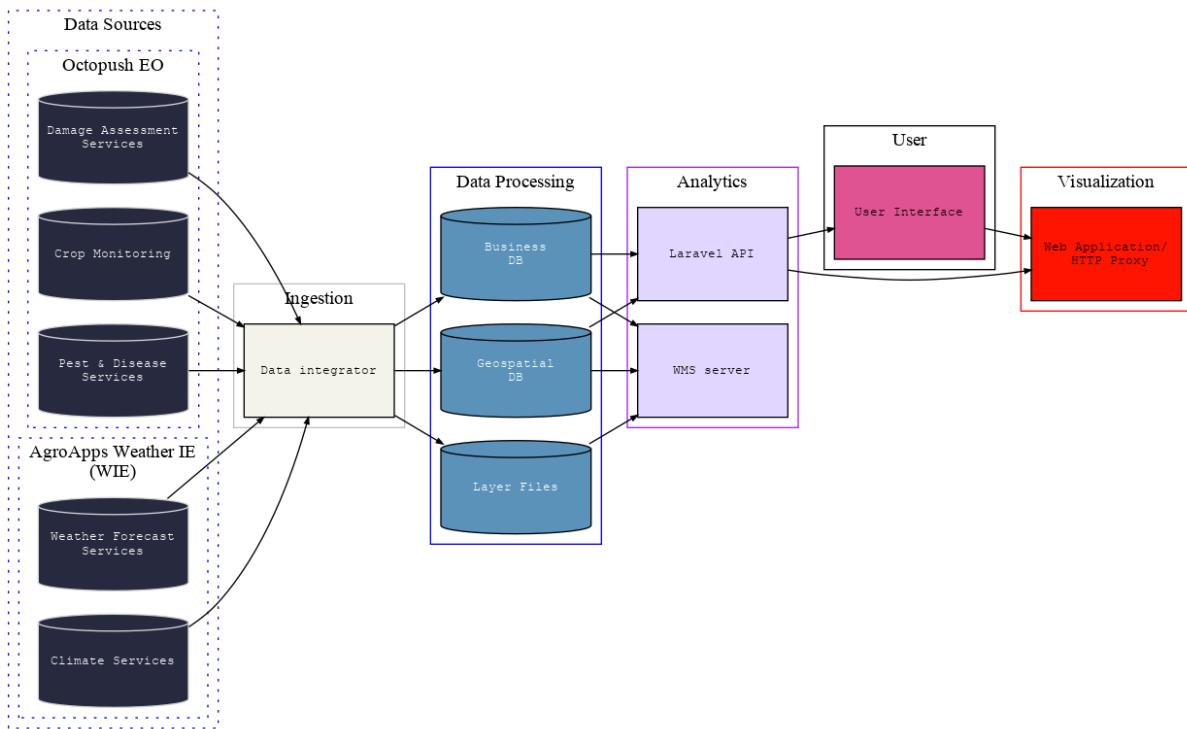


Figure 43: INFINITECH-RA Compliant Architecture for the Pilot on Big Data and IoT for the Agricultural Insurance Industry

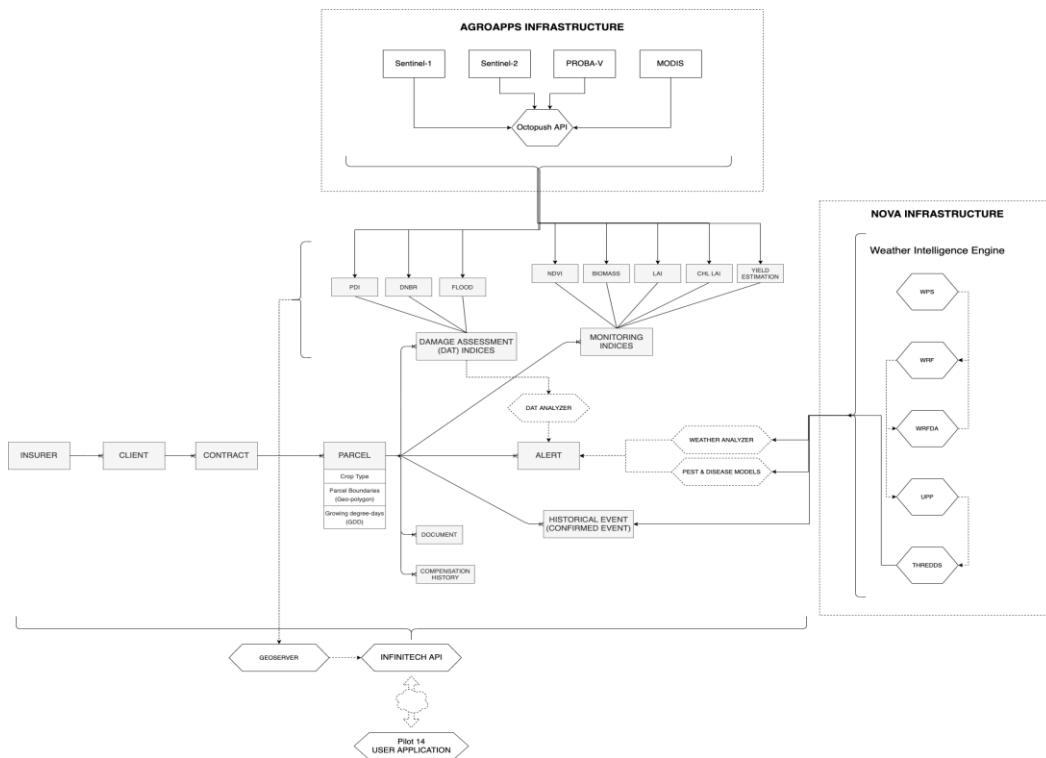


Figure 44: Workflow of the Pilot on Big Data and IoT for the Agricultural Insurance Industry

6.14.6 Components

The following components will be deployed and used as part of the pilot:

- Octopush EO Service.
- AgroApps Weather Intelligence Engine.
- Data integrator.
- Business and Geospatia DB.
- WMS Server.
- RESTful API.
- User interface.

7. INFINITECH-RA positioning

7.1 Scoping the INFINITECH-RA

Following paragraphs present the scope and limitations of the INFINITECH-RA, based on a set of plausible questions.

7.1.1 RA seems generic and high level. Shouldn't it be more concrete?

The INFINITECH RA described in the document is described in High Level on purpose. It does not specify which node does what and does not put constraints on connections between nodes. It describes the principles upon which nodes with specific functions can be integrated in the INFINITECH reference implementations to create a large class of solutions. It allows a classification of nodes based on their functions to be mapped with a well established Reference Model (BDVA RM) and allows for generalization of use cases.

A specific implementation, however, can impose strict constraints on nodes interfaces (like REST APIs) to make nodes interoperable (interoperable, stackable, orchestrable).

7.1.2 Where is the Innovation?

The RA is a model and a set of principles on how to integrate specific functionalities in any instantiation of the INFINITECH platform. However, the principles are quite innovative and powerful as they imply the possibility to build solutions out-of-the-box with standard components taken out from a library of services (nodes).

INFINITECH nodes which comply with the INFINTECH RA, will be interoperable, stackable, exchangeable, orchestrable. It allows the integration of nodes that performs functionalities like Data Collector to different data sources, Data Anonymization, Data Ingestion, Data Processing, Data Streaming, Data Abstractions and Ontologies, Data Analytics, Data Manipulation and Presentation.

An INFINITECH implementation based on the DevOps, Microservice Architecture and Kubernetes Deployment is highly powerful and innovative as it allows the creation of specific solutions to solve a virtually infinite class of use cases.

The RA contains the principles to build Library of Standard Components, to populate a MarketPlace, to work with an Orchestrator Tool, to experiment with Sandboxes, to create powerful solutions out of the box.

This has currently very few competitors in the market.

7.1.3 What are the specific features for the Financial and Insurance Sector?

The IRA is a general model and not specific to any sector. However, it is modeled after the INFINITECH requirements from use cases and pilots and complies with the Financial and Insurance Sector regulations and best practices. The specific features for banks and other financial organizations are in the solutions for data manipulation (from ingestion to analytics and to presentation) developed to solve the use cases of the project. The project will develop Ontologies and Machine Learning Algorithms which will be vertical. Offering a general framework that will be reusable in large classes of financial sector's scenarios.

7.1.4 How can the RA cope with Business Requirements, Technological Requirements, Regulatory Requirements, etc.?

Requirements of all types form a checklist against which the IRA and its implementation will be checked. For example, a business requirement will be assessed in a specific pipeline/sandbox implementation. A Regulatory Constraints will be checked both in the General Framework (there is a specific layer that complies with privacy or security in IRA) and in the specific implementation (is the solution compliant with Regulatory X?)

Moreover, cross cutting services at different layers of the RA can implement automatic checking for requirements as stackable/pipelined nodes. This will be quite innovative and will represent a unique selling point of the INFINITECH project, as for instance to Anonymization as-a-service, GDPR as-a-service, etc.

7.1.5 Where are the building blocks described in the DoA and how can be mapped in RA?

Specific tools and applications can be mapped in the IRA layered model. LeanXcale powerful database management solution, for example, can have a central role in the Data Processing Architecture, providing a generalization for Data Repository with streaming endpoints for Analytics.

Anonymization is also very easy to understand and maps beautifully into the IRA as a pipeline service. Anonymization can be implemented as a node which gets data from one source, transforms and produces data sent to another node (a data repository such as LeanXcale repo). It is as simple as two endpoints and it can be configured by cross cutting services.

There could be many types of anonymization services specific to the particular data sets but the endpoints could be the same allowing for exchangeability and interoperability of the service.

7.1.6 How can the INFINITECH-RA support Pilots and Use cases, and solve real Business Problems?

The IRA provides a general framework and each use case will be different. A mapping of a use case is basically the operation to identify which functions are needed at different layers, which tools and applications will solve them and to define a pipeline of services (nodes) able to manipulate the datasets of the use case to provide the expected results.

Some examples of Use Cases are provided and this document will not be complete until all the use cases are mapped to the RA. No use case left behind.

7.1.7 I am an End User and I cannot see how my Use Case can fit in the RA.

Most of the use cases can map easily in the pipelined approach of the IRA. For instance, a simple case is to ingest all the data from different data sources (silos in data lakes) into a central repository after collection, anonymization, categorization. Then a level of abstraction (polyglot and ontologies) can allow analytics to work on this aggregated set of data. Applications like dashboards or Visualization tools will work out of the produced data.

Another case is to have intelligent processing as close to the datasource as possible. In this case analytics will query data directly and produce data for presentation.

Distributed Ledgers (blockchains) pose a challenge to the IRA which can be solved in many different ways (see below).

However, the IRA does not contain any specific solution but a general framework to create different ways to build them up.

7.1.8 Is the RA being implemented in the Cloud? Will my data be distributed outside my organization? How can I preserve the classification of my own data?

As far as the INFINITECH Project is concerned, the principle is that Data that belongs to Financial Organizations stays in the Financial Organizations. The IRA is not imposing anything like exporting data outside any organization. It is not the specification of a Cloud Service.

The project will organize on-premise testbeds and the pilots will deploy an instance of the IRA in sandboxes in the infrastructure of the organizations.

There are cases where synthetic data can be used which will not require on-premise infrastructure but this is specific to Pilots.

However, the IRA, can be implemented as Sandbox-As-A-Service on a public Cloud using for instance synthetic data, preserving the confidentiality and security of the data of the tenants.

7.1.9 Blockchain Technology does not seem to fit in the RA. How can it be managed?

Indeed, a distributed ledger poses challenges to the central data processing layer of the IRA and to the underlined pipeline concept. However, this is only apparent as a blockchain node can be accessed to gather the relevant information and produce data that can go as input to other nodes. A Blockchain node can be considered as any other datasource at the infrastructure layer and accessed through a specific Data Collector (see Web3.js APIs).

DLT can also find a place in the central layer of the IRA and even as a cross cutting service, for peer-to-peer no-man-in-the-middle authorization, and even for licensing services.

DLT is central to INFINITECH project and the IRA provides many ways to address them. More will be defined in the specific use cases.

7.2 Development Roadmap and Future Outlook

The next development phase of the INFINITECH-RA, will be a Fine Tuning and Updates (M9-M18). It will aim at revised and fine-tune the INFINITECH-RA based on its actual use in other workpackage, including the sandboxes and pilot development workpackages. It will receive and exploit stakeholders' feedback from the practical use of the INFINITECH-RA for developing BigData/AI systems for the finance sector. The feedback will be exploited towards improving aspects of the INFINITECH-RA, but also towards validating the architectural concepts introduced in practice. This Phase will lead to the final version of the INFINITECH-RA as part of deliverable D2.14 i.e. the second and updated version/release of this deliverable.

8. Conclusions

A reference architecture for BigData systems in digital finance can greatly facilitate stakeholders in structuring, designing, developing, deploying and operating BigData, AI and IoT solutions. It serves as a stakeholders' communication device, while at the same time providing a range of best practices that can accelerate the development and deployment of effective systems. This deliverable has introduced the first version of such a reference architecture, namely the INFINITECH-RA. The latter adopt the concept and principles of data pipelines, which are in-line with the state of the art in BigData and Artificial Intelligence systems. It is also in-line with the principles of the reference architecture of the BDVA. In practice, it extends and customizes the BDVA RA with constructs that permit its use for digital finance use cases. The INFINITECH-RA defines the structuring principles that drive the integration of the INFINITECH technical components and technologies, which have been documented in deliverable D2.5.

As part of the deliverable, we have illustrated how INFINITECH-RA can be used to support the development of common BigData/AI pipelines for digital finance applications. In this direction the deliverable has provided some examples of popular pipelines. It has also provided an initial mapping of most pilots of the project to the INFINITECH-RA, using INFINITECH components and technologies. The presented RA will serve as a basis for the implementation of the initial versions of the MVP (Minimum Viable Product) of the pilots. These implementations will allow the collection of stakeholders' feedback regarding the RA. Based on this feedback, the next versions/released of the INFINITECH-RA will be produced, including updates in the different views of the architecture. Furthermore, future releases of the INFINITECH-RA will consider new components that will be made available as part of D2.6 (i.e. the next version of D2.5) of the project i.e. it will consider a broader range of BigData, IoT and AI related capabilities developed in INFINITECH.

The completion of the first version of the INFINITECH-RA as part of this deliverable is a milestone of the project. It signifies the accomplishment of several of WP2 objectives as listed in the following table.

As per the specific KPIs set for the project the following table addresses the specific indexes listed in the DoA.

KPI	Description	Comment
KPI 1	RA to be specified ≥ 1	This Deliverable documents the RA of the project and the RA logical schema of the pilots of the project. Therefore the KPI is fully achieved.
KPI 2	Supported types of different types of databases, data sources and data stores ≥ 12	The RA in principle manage different kind of Data sources (RDBMS SQL and NoSQL, Data Lakes, IOT and Blockchain datasources). The number is far more than the 12 basic KPI.

Table 4: Mapping of INFINITECH DoA/Task KPI with Deliverable Achievements

9. Appendix A: Literature

- [Armbrust15] M. Armbrust, R. Xin, C. Lian, Y. Huai, D. Liu, J. Bradley, X. Meng, T. Kaftan, M. Franklin, A. Ghodsi, M. Zaharia, "Spark SQL: relational data processing in Spark", in ACM SIGMOD, 2015, pp. 1383-1394.
- [BDVA-RM] Big Data Value Association. BVA SRIA—European big data value strategic research and innovation agenda. 2017. http://bdva.eu/sites/default/files/BDVA_SRIA_v4_Ed1.1.pdf. Accessed 18 Feb 2020
- [Boid18] Boid D, Chang W. NIST Big Data Interoperability Framework: Volume 6, RA Version 2. NIST Big Data Program. 2018. <https://bigdatawg.nist.gov/uploadfiles/NIST.SP.1500-6r1.pdf>
- [Bons12] Bons, R.W.H., Alt, R., Lee, H.G. et al. Banking in the Internet and mobile era. *Electron Markets* 22, 197–202 (2012). <https://doi.org/10.1007/s12525-012-0110-6>
- [FINSEC19-D2.5] FINSEC Reference Architecture II – Available at <https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5ce3a941d&appId=PPGMS>
- [BOOST4.0-D2.5] BOOST4.0, D2.5 – BOOST 4.0 Reference Architecture Specification v1, Available at: <https://cordis.europa.eu/project/id/780732/results> [Accessed: 15-May-2020]
- [Bracke19] Philippe Bracke, Anupam Datta, Carsten Jung and Shayak Sen, "Machine learning explainability in finance: an application to default risk analysis", Bank of England, Staff Working Paper No. 816, August 2019
- [Duggan15] 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.
- [Gadepally16] 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)*, 2016, pp. 1-6.
- [Kruchten95] Kruchten, Philippe (1995, November). Architectural Blueprints — The "4+1" View Model of Software Architecture. *IEEE Software* 12 (6), pp. 42-50.
- [Komulainen18] Komulainen, H., Makkonen, H. Customer experience in omni-channel banking services. *J Financ Serv Mark* 23, 190–199 (2018). <https://doi.org/10.1057/s41264-018-0057-6>
- [Kyriazis18] D. Kyriazis, et al, "BigDataStack: A Holistic Data-Driven Stack for Big Data Applications and Operations", *BigData Congress*, San Francisco, CA, USA, 2018: 237-241
- [Minpeng11] Z. Minpeng, R. Tore, "Querying combined cloud-based and relational databases", in *Int. Conf. on Cloud and Service Computing (CSC)*, 2011, pp. 330-335.
- [Minukhin18] S. Minukhin, V. Fedko and D. Sitnikov, "SQL-On-Hadoop Systems: Evaluating Performance of Polybase for Big Data Processing," 2018 International Scientific-Practical Conference Problems of Infocommunications. Science and Technology (PIC S&T), Kharkiv, Ukraine, 2018, pp. 591-594.
- [MIRAB12] Microsoft Industry Reference Architecture for Banking (MIRA-B), Microsoft Corporation Whitepaper, May 2012
- [Ong14] 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.
- [Özsu11] T. Özsu, P. Valduriez, *Principles of Distributed Database Systems*, 3rd ed. Springer, 2011, 850 pages.
- [Ribeiro16] Ribeiro, Marco & Singh, Sameer & Guestrin, Carlos. (2016). "Why Should I Trust You?": Explaining the Predictions of Any Classifier. 1135-1144. 10.1145/2939672.2939778.

[Shearer00] Shearer C., The CRISP-DM model: the new blueprint for data mining, *J Data Warehousing* (2000); 5:13—22.

[Simitsis12] A. Simitsis, K. Wilkinson, M. Castellanos, U. Dayal, "Optimizing analytic data flows for multiple execution engines", in *ACM SIGMOD*, 2012, pp. 829-840.

[Tomatic98] A. Tomatic, L. Raschid, P. Valduriez, "Scaling access to heterogeneous data sources with DISCO", *IEEE Trans. On Knowledge and Data Engineering*, vol. 10, pp. 808-823, 1998.

[Xu18] H. Xu, M. Chen, Y. Zhou, B. Du and L. Pan, "A Novel Comprehensive Quality Index QoX and the Corresponding Context-aware System Framework," 2018 IEEE 4th International Conference on Computer and Communications (ICCC), Chengdu, China, 2018, pp. 2415-2419.

10. Appendix B: APPENDIX User requirements' mapping

Functional Requirements	Data Sources					Data Ingestion			Data Security And Privacy				Interface		Visualiaztion		
	Data Bases	Data Lakes	Stored Files	IoT Sources	Block Chains	Data Collector	Data Connector	Serializer	eIDAS	Data Anonymization	DUOS	DPO	API Gateway	Proxy	Data	Model	Results
REQ-P1-F-B-M-O-001			X														
REQ-P1-F-B-M-O-002		x															
REQ-P1-F-B-M-O-003										X							
REQ-P1-F-B-M-O-004																	
REQ-P1-F-B-M-O-005						X	X	X									
REQ-P1-F-B-M-O-006																	
REQ-P1-F-B-M-O-007																	X
REQ-P2-F-B-M-O-008																	
REQ-P2-F-B-M-O-009									X	X							
REQ-P2-F-B-M-O-010																	
REQ-P2-F-B-M-O-011															X		X
REQ-P3-F-B-M-O-012																	
REQ-P3-F-B-M-O-013																	
REQ-P3-F-B-M-O-014																	
REQ-P3-F-B-M-O-015																	
REQ-P3-F-B-M-O-016									X								
REQ-P3-F-B-M-O-017									X	X							
REQ-P3-F-B-M-O-018																	
REQ-P3-F-B-M-O-019									X								
REQ-P3-F-B-M-O-020																	
REQ-P4-F-B-M-O-021																	
REQ-P4-F-B-M-O-022									X								
REQ-P4-F-B-M-O-023																	
REQ-P4-F-B-M-O-024																	X
REQ-P6-F-B-M-O-043	X					X	X	X									
REQ-P6-F-B-M-O-044		X	X			X	X	X									
REQ-P6-F-B-M-O-045																	
REQ-P6-F-B-M-O-046																	
REQ-P6-F-B-M-O-047																	

D2.13 – INFINITECH Reference Architecture - I

REQ-P9-F-B-M-O-061	X					X	X	X										
REQ-P9-F-B-M-O-062									X	X								
REQ-P9-F-B-M-O-063																		
REQ-P9-F-B-M-O-064																X		X
REQ-P10-F-B-M-O-065	X							X										
REQ-P10-F-B-M-O-066						X			X	X								
REQ-P10-F-B-M-O-067						X												
REQ-P10-F-B-M-O-068																		
REQ-P10-F-B-M-O-069																X		
REQ-P10-F-B-M-O-070																		
REQ-P10-F-B-M-O-071						X	X	X	X	X								
REQ-P10-F-B-M-O-072					X	X	X	X										
REQ-P10-F-B-M-O-073					X													
REQ-P10-F-B-M-O-074	X																	
REQ-P10-F-B-M-O-075																		
REQ-P10-F-B-M-O-076																X	X	X
REQ-P11-F-B-M-O-077				X		X	X	X										
REQ-P11-F-B-M-O-078		X	X			X	X	X										
REQ-P11-F-B-M-O-079		X	X			X	X	X										
REQ-P11-F-B-M-O-080										X								
REQ-P11-F-B-M-O-081																		
REQ-P11-F-B-M-O-082																		
REQ-P12-F-B-M-O-083				X		X	X	X										
REQ-P12-F-B-M-O-084									X	X								
REQ-P12-F-B-M-O-085																		
REQ-P13-F-B-M-O-086		X	X			X	X	X										
REQ-P13-F-B-M-O-087																		
REQ-P13-F-B-M-O-088																		
REQ-P13-F-B-M-O-089																		
REQ-P14-F-B-M-O-090		X	X			X	X	X										
REQ-P14-F-B-M-O-091																		
REQ-P14-F-B-M-O-092		X	X			X	X	X										
REQ-P14-F-B-M-O-093																		
REQ-P14-F-B-M-O-094																		
REQ-P14-F-B-M-O-095																		
REQ-P14-F-B-M-O-096																		
REQ-P14-F-B-M-O-097																		
REQ-P14-F-B-M-O-098																		
REQ-P14-F-B-M-O-099																		

D2.13 – INFINITECH Reference Architecture - I

A1:O34FA1:Q60 Functional Requirements	Data Management							Data Model and Semantic Engine							Block chain	Analytics
	Relational DB	HTAP Enabler	Polyglot Processing	Incremental Analytics	OneHot Encoder	Visualization Preparation	Stream Processor	Analyzer	Reasoner	Mapping	Pre-processing	Anomaly Analysis	Pattern Analysis	Stream Story		
REQ-P1-F-B-M-O-001																
REQ-P1-F-B-M-O-002																
REQ-P1-F-B-M-O-003																X
REQ-P1-F-B-M-O-004	X															X
REQ-P1-F-B-M-O-005	X					X										X
REQ-P1-F-B-M-O-006	X					X										X
REQ-P1-F-B-M-O-007						X										
REQ-P2-F-B-M-O-008	X	X					X									
REQ-P2-F-B-M-O-009																
REQ-P2-F-B-M-O-010	X	X		X			X									
REQ-P2-F-B-M-O-011						X										X
REQ-P3-F-B-M-O-012								X	X	X	X					
REQ-P3-F-B-M-O-013																
REQ-P3-F-B-M-O-014														X		
REQ-P3-F-B-M-O-015			X													X
REQ-P3-F-B-M-O-016																
REQ-P3-F-B-M-O-017																
REQ-P3-F-B-M-O-018														X		
REQ-P3-F-B-M-O-019																
REQ-P3-F-B-M-O-020																
REQ-P4-F-B-M-O-021	X	X	X													
REQ-P4-F-B-M-O-022																
REQ-P4-F-B-M-O-023																X
REQ-P4-F-B-M-O-024																
REQ-P6-F-B-M-O-043	X															
REQ-P6-F-B-M-O-044																
REQ-P6-F-B-M-O-045																X
REQ-P6-F-B-M-O-046																X
REQ-P6-F-B-M-O-047																X
REQ-P9-F-B-M-O-061	X						X									
REQ-P9-F-B-M-O-062																
REQ-P9-F-B-M-O-063						X										
REQ-P9-F-B-M-O-064						X										
REQ-P10-F-B-M-O-065	X															

D2.13 – INFINITECH Reference Architecture - I

REQ-P10-F-B-M-O-066																				
REQ-P10-F-B-M-O-067																				
REQ-P10-F-B-M-O-068																				X
REQ-P10-F-B-M-O-069								X												
REQ-P10-F-B-M-O-070	X							X												
REQ-P10-F-B-M-O-071																				
REQ-P10-F-B-M-O-072																				
REQ-P10-F-B-M-O-073																				X
REQ-P10-F-B-M-O-074	X	X						X												
REQ-P10-F-B-M-O-075	X	X		X																X
REQ-P10-F-B-M-O-076																				
REQ-P11-F-B-M-O-077	X							X												
REQ-P11-F-B-M-O-078	X																			
REQ-P11-F-B-M-O-079	X																			
REQ-P11-F-B-M-O-080																				
REQ-P11-F-B-M-O-081	X																			
REQ-P11-F-B-M-O-082																				X
REQ-P12-F-B-M-O-083								X												
REQ-P12-F-B-M-O-084																				
REQ-P12-F-B-M-O-085																				X
REQ-P13-F-B-M-O-086	X																			
REQ-P13-F-B-M-O-087								X												
REQ-P13-F-B-M-O-088	X	X		X																
REQ-P13-F-B-M-O-089			X																	
REQ-P14-F-B-M-O-090																				X
REQ-P14-F-B-M-O-091																				X
REQ-P14-F-B-M-O-092																				
REQ-P14-F-B-M-O-093			X																	
REQ-P14-F-B-M-O-094																				X
REQ-P14-F-B-M-O-095																				X
REQ-P14-F-B-M-O-096																				X
REQ-P14-F-B-M-O-097																				X
REQ-P14-F-B-M-O-098																				X
REQ-P14-F-B-M-O-099																				