

Architectural Framework for Federated Learning in Aviation Maintenance

Designing and Analyzing Key Features for Data Sharing Acceptance with ArchiMate Modeling

MSc Thesis in Engineering and Policy Analysis

Bas van de Walle

September, 2024



Architectural Framework for Federated Learning in Aviation Maintenance

Master thesis submitted to Delft University of Technology

In partial fulfillment of the requirement for the degree of

MASTER OF SCIENCE

in Engineering and Policy Analysis

by

Bas van de Walle

Student Number: 4713044

To be defended in public on 25-09-2024

Chair (TU Delft):	Dr.ir. G.A. (Mark) de Reuver
First Supervisor (TU Delft):	Dr.ir. G.A. (Mark) de Reuver
Second Supervisor (TU Delft):	Dr. Ir. Rutger van Bergem
Advisor (TU Delft):	Dr. Ir. M. (Marcela) Tuler de Oliveira
Advisor (KLM):	Asteris Apostolidis PhD
Project Duration:	January 2024 – Sept 2024
Faculty:	Technology, Policy and Management



Preface

As a child, I always to be a pilot and the aviation industry has always fascinated me since. In a world where we are no longer bound by the borders of a single country, aviation has become a unifying force, helping people connect and fostering global citizenship. This belief in aviation's potential to unite people drove me to explore this topic in my thesis.

Working on this thesis posed several challenges, particularly the need to work independently without frequent feedback. The long hours of research, often without the reassurance of daily or weekly guidance, made the process demanding at times. However, it also taught me the importance of perseverance and self-motivation. One of the most valuable skills I gained from this experience is the ability to continue, even when the task at hand wasn't always pleasant or immediately rewarding.

I am incredibly grateful for the support I received throughout this journey. My professors offered invaluable insights and guidance, beside that I am deeply appreciative of the members of IDCA. Without their insights, I would not have been able to reach many of the conclusions in this thesis.

Looking ahead, I remain intrigued by the aviation industry. I firmly believe that if companies can move past their competitiveness and start collaborating more, the industry can achieve incredible efficiencies. I hope my work contributes, in a small way, to this vision of a more collaborative and innovative aviation sector.

This thesis also marks the end of my seven-year journey at Delft University of Technology. My time here has shaped me into the person I am today. I've gained not just academic knowledge, but lifelong friendships and unforgettable memories. As I close this chapter of my life, I look back with immense gratitude and look forward to future opportunities.

Executive Summary

Current Status

The aviation industry is transitioning from traditional maintenance practices, typically scheduled after a specified number of flight hours, to more advanced, data-driven approaches. These predictive maintenance techniques leverage machine learning models to enhance the accuracy of assessments. These models can reduce the frequency of unscheduled maintenance and optimize inventory management. However, such models rely heavily on large datasets, which are challenging to compile in the aviation sector due to the rarity of specific operational incidents and the diverse types of data collected by different companies. When collaborating, a tradeoff arises between the level of security measures, trust in partners, and ensuring the system still functions. Federated Learning emerges as a promising solution to these challenges. Federated Learning is a novel form of machine learning that allows multiple entities to collaboratively develop a shared model while keeping their data localized, thus maintaining privacy and data sovereignty.

Question

The primary objective of this research is to identify and validate the critical architectural features necessary for the acceptance of Federated Learning in the aviation maintenance industry. Architectural features refer to design choices such as security protocols and governance frameworks. By focusing on these aspects, this study addresses the technical and collaborative challenges that must be overcome to develop a Federated Learning system for predictive maintenance in aviation. Resulting in the following research question:

‘What architectural features should be included in the design of the ‘Federated Learning for aircrafts’ predictive maintenance system’ to be accepted by the stakeholders in the aviation industry?’

Approach

This study employs a Design Science Research methodology. The sub-questions for this research follow the steps in the Design Science Research methodology. This is not the case for the demonstration phase which was deemed not possible due to the conceptual nature of the design.

1. What are the challenges in sharing maintenance data, and how do they impact data safety and collaboration? Through a literature review and interviews with stakeholders, the study identified concerns about data privacy, security, and competition. These challenges significantly limit data-sharing initiatives. (Problem Identification and Motivation)
2. What specific technical requirements should the Federated Learning system have to address these challenges? Based on the interviews and thematic analysis, a list of requirements was developed. These include robust privacy mechanisms, transparent governance, ensuring model explainability, and clear accountability mechanisms. (Defining the Objectives for a Solution)
3. What architectural features should be included in the Federated Learning design to meet these requirements? ArchiMate modeling was used to design a system incorporating these requirements. Federated Learning, combined with encryption techniques and a consortium-managed system, was the result. (Design and Development)
4. How do stakeholders perceive the acceptability of the designed system? After the validation interviews and the expert session, three changes were implemented: Switching from

differentiating on the model version to differentiating on the usages. Improving explainability by switching to Trusted Executive Environments and removing differential privacy. Traditional contracts to increase trust towards each other were also added. (Evaluation)

5. What lessons can be learned from the development and evaluation of the Federated Learning system for future improvements? The research highlights the importance of continuously building trust among stakeholders. Furthermore, the token-based reward system based on contribution is a good incentive to be adaptable and develop long-term collaboration. (Communication)

Results

To build trust, this study employs traditional methods like legal contracts and appoints a consortium as a neutral party. Transparent governance and involving a neutral, trusted entity is necessary to gain stakeholder acceptance and ensure long-term collaboration. Blockchain technology enhances the transparency of the consortium's operations, ensuring all transactions and data exchanges are recorded on an immutable database.

Despite commercial mistrust and skepticism among stakeholders, the industry prefers encryption solutions that do not compromise model accuracy. This study recommends using a Trusted Execution Environment, which enables secure computations on encrypted data without the need for decryption by the central party, preserving privacy while supporting model explainability.

The research emphasizes the importance of equitable benefit sharing to ensure sustained motivation and collaboration. This could be achieved by implementing a token-based model access system, which distributes usage rights based on contribution scores. This approach also prevents free-riding by ensuring that participants must contribute to gain access to the system's benefits, aligning incentives with contributions, and creating a collaborative environment.

Contribution

This research contributes to understanding Federated Learning adoption by examining how trust can be established within decentralized governance models and proposes a token-based system for equitable model utilization among participants. It also presents a practical framework tailored to the aviation sector. This framework addresses the challenges in aviation maintenance, such as the need for absolute accuracy and extensive regulations. It could also serve as a starting point for adaptation in similar industries requiring accuracy and collaboration.

Future Research

Future research should address the identified limitations by expanding geographical diversity, engaging participants less inclined to share data, and involving operational technical staff. Surveys and statistical analyses could contribute to the validation of the qualitative findings and ensure they represent the industry's broader views.

Collaboration with researchers and regulatory bodies to develop standardized metrics and guidelines for explaining machine learning is necessary. Furthermore, economic impact studies could provide compelling evidence of the financial benefits of adopting Federated Learning systems, which will further support the economic viability of Federated Learning initiatives.

Overall, this research highlights the transformative potential of Federated Learning in aviation maintenance, offering a comprehensive framework for its implementation. The insights gained from this study can guide the development and implementation of Federated Learning, enhancing the aviation industry's operational efficiency and safety.

Content

Preface	3
Executive Summary	4
Content	6
List of Figures	8
List of Tables	9
Nomenclature	9
1 Introduction	10
1.1 Problem Introduction	10
1.2 Literature Review	11
1.3 Research Question	15
2 Research Approach	16
2.1 Sub-Questions and Research Methodologies	16
2.2 Research Methodology	20
2.3 Research Flow Diagram	22
3 Problem Identification	23
3.1 Context and Background	23
3.2 Interview Setup	25
3.3 Thematic Analysis	30
3.4 Insights Regulators	35
4 System Requirements	37
4.1 Overview Requirements	37
4.2 Diverging Opinions and Inclusion Criteria	38
4.3 Privacy and Security	38
4.4 Regulatory Compliance	39
4.5 Management of Intellectual Property	39
4.6 Distribution of Benefits	39
4.7 Impact of Federated Learning on Predictive Maintenance	40
4.8 Adoption and Trust	40
4.9 Conclusion	41
5 Architectural Features	42
5.1 Overview of the System	42
5.2 Strict Security Measures	45
5.3 Robust Data-Sharing Framework	47

5.4 Compliance with Existing Legal Standards-----	49
5.5 Legal Adaptability-----	51
5.6 Consortium-Based IP Management-----	52
5.7 Equitable Benefit Sharing-----	54
5.8 Transparent Value Attribution-----	58
5.9 Accuracy of Predictions-----	59
5.10 Explainability and Transparency of AI Models-----	61
5.11 Training Programs-----	64
5.12 Appointment of a Trusted Neutral Entity-----	65
5.13 Conclusion-----	67
6 Validation-----	69
6.1 Structure of the Validation-----	69
6.2 Expert Validation Session-----	69
6.3 Validation Interviews-----	71
6.4 Results Validation Interview-----	75
6.5 Changes to the system-----	81
6.7 Architectural Changes-----	87
6.8 Overview of the Validated System-----	89
7 Discussion-----	91
7.1 Key Findings-----	91
7.2 Practical Implications-----	94
7.3 Literature Implications-----	95
7.4 Comparison with Existing Literature-----	96
7.5 Contributions to the Literature-----	96
7.6 Limitations of Study-----	97
7.7 Future Study Recommendations-----	98
7.8 Concluding Remarks-----	100
8 Conclusion-----	102
8.1 Recapitulation of Research Objectives and Findings-----	102
8.2 Contribution to the Field-----	102
8.3 Practical Implications-----	103
8.4 Limitations and Challenges-----	103
8.5 Future Directions-----	104
8.6 Final Thoughts-----	104
Literature List-----	105
Appendices-----	109

Appendix A: Literature Review -----	109
Appendix B: Informed Consent-----	111
Appendix C: Interview Questions Regulators -----	113
Appendix D: Description Components ArchiMate Model -----	115
Appendix E: Relationships ArchiMate -----	117
Appendix F: Component Contributions to Requirements -----	121
Appendix G: Informative Slides Security -----	122
Appendix H: Overview Validated Model -----	123

List of Figures

Figure 1: (left) Decentralized, (right) Centralized (Aledhari et al., 2020)	12
Figure 2: Research Flow Diagram	22
Figure 3: Relations ArchiMate	43
Figure 4: Overview System	44
Figure 5: Zoom Security Measures	47
Figure 6: Zoom Data Sharing Framework	49
Figure 7: Zoom Regulatory Framework	51
Figure 8: Zoom Legal Adaptability	52
Figure 9: Zoom IP Management	54
Figure 10: Zoom Equitable Benefit Sharing	57
Figure 11: Zoom Transparent Value Attribution	59
Figure 12: Zoom Accuracy of Predictions	61
Figure 13: Highlight Model Explainability.....	63
Figure 14: Zoom Training Programs.....	65
Figure 15: Zoom Trusted Neutral Entity	66
Figure 16: Overview Validated System, Highlighting Differences.....	90
Figure 17: Slide homomorphic encryption (Roth et al., 2021)	122
Figure 18: Slide differential privacy (Devaux, 2022)	122
Figure 19: Overview Validated Model	123

List of Tables

Table 1: Abbreviations Used	9
Table 2: Demographic Information Interviewees (n=12).....	27
Table 3: Questions from the Questionnaire	28
Table 4: Thematic Analysis (n=10)	30
Table 5: Overview Requirements.....	37
Table 6: Demographic Information Interviewees Validation (n=9).....	72
Table 7: Interview Questions Validation	74
Table 8: Views Acceptance Design Choices	76
Table 9: Sources Form the Literature Revie.....	109
Table 10: Description Components	115
Table 11: Relationships ArchiMate	117
Table 12: Overview Component Contributions to Requirements	121

Nomenclature

Table 1 shows the abbreviations used in this report.

Table 1: Abbreviations Used

Abbreviation	Definition
AI	Artificial Intelligence
BD	Big Data
CAMO	Continuing Airworthiness Management Organization
DGA	Data Governance Act
DMA	Digital Markets Act
DSA	Digital Services Act
DSR	Design Science Research
EASA	European Union Aviation Safety Agency
EU	European Union
FAA	Federal Aviation Administration
FL	Federated Learning
FLI	Federated Learning Incentifier
GDPR	General Data Protection Regulation
IDCA	Independent Data Consortium for Aviation
IP	Intellectual Property
ML	Machine Learning
MRO	Maintenance, Repair and Overhaul
NARP	National Aviation Research Plan
OEM	Original Equipment Manufacturer
RFFL	Robust and Fair Federated Learning
SMPC	Secure Multi-Party Computation
TEE	Trusted Execution Environment

1 Introduction

1.1 Problem Introduction

Aviation maintenance ensures aircraft safety and reliability through regular inspections, repairs, and part replacements. This maintenance is typically carried out using Predictive Maintenance and Condition-based Maintenance approaches, which leverage data insights to optimize maintenance schedules and enhance aircraft availability. The more data available, the better the ability to identify rare events, as it helps detect patterns that may not be visible with smaller datasets. Analyzing large amounts of data allows for recognizing general deterioration trends, leading to more accurate predictions of component failures and enabling timely preventive actions.

In recent years, the aviation industry has been shifting from traditional maintenance methods to data-driven approaches. Traditional practices rely on scheduled tasks based on flight hours, prioritizing safety while often overlooking the potential of Big Data (BD) and BD analytics, which have been proven to enhance operational efficiency and cost-effectiveness (Adamopoulou & Daskalakis, 2023; Dinis et al., 2019). However, data sharing in aviation faces significant barriers. Bilateral agreements with Original Equipment Manufacturers (OEMs) frequently restrict data exchange between airlines and maintenance providers. In addition, confidentiality agreements with pilots can limit the operational data needed for comprehensive analysis. Furthermore, variations in national legislation regarding data privacy and transfer regulations complicate the establishment of a unified data-sharing framework.

Advanced predictive maintenance techniques, using machine learning (ML) models, aim to improve the accuracy of maintenance assessments, potentially reducing the frequency of unscheduled maintenance and optimizing inventory management. However, these models depend on access to large datasets, which is particularly challenging in the aviation sector due to the rarity of specific operational incidents. For machine learning models to be effective, they need to learn from a large number of events. However, within a single airline, rare occurrences such as unexpected component failures happen infrequently. Even large airlines may rarely experience enough of these events to provide sufficient data for reliable ML models. This creates a significant problem, as no airline can gather enough data on its own to develop accurate predictive maintenance models.

Additionally, airlines are often reluctant to share data due to competitive reasons. This reluctance persists even among airlines within the same group. For example, Air France and KLM, though part of the same group, have different business strategies (Financial Times, 2023). This separation demonstrates how competitive concerns hinder data sharing, further reducing the data available for ML. As a result, no single party can accumulate enough data to create a robust predictive maintenance system.

The desire to collaborate without losing competitive advantages results in a tradeoff: the balance between the level of protection mechanisms needed and the trust among stakeholders. Given the fact that the more protection implemented, the more complex it becomes to work together and maintain practicality.

Federated Learning (FL) emerges as a promising solution to these challenges. Unlike traditional ML, where a single organization works with its dataset, FL allows multiple organizations to collaboratively develop a shared model while keeping their data localized, preserving privacy and data sovereignty. However, FL involves a tradeoff: while it protects data privacy and encourages collaboration, it requires a high level of trust among stakeholders.

For FL to succeed, stakeholders must believe that the benefits of improved predictive maintenance outweigh the risks associated with data sharing, even if it is done indirectly. This research explores data sharing for predictive maintenance, focusing on engine maintenance as a case study. The thesis examines the adoption of predictive maintenance in aviation through FL, focusing on architectural design considerations such as the extra security measures necessary for stakeholder acceptance and implementation. The aim is to find the right balance between the technical need for more accurate models and the concerns of stakeholders who are cautious due to issues of privacy, competition, and regulation.

1.2 Literature Review

This literature review outlines specific challenges regarding stakeholders' acceptance of designing a 'Federated Learning for aircrafts' predictive maintenance system. The analysis of these stakeholders is described in Chapter 3.1.1 Stakeholders.

For this literature review, Google Scholar was used with the primary term 'Federated Learning' in combination with one of the following: 'Architectures', 'Acceptance', 'Aviation', 'Predictive maintenance', 'Healthcare' and 'Banking'. No papers older than 2019 were chosen due to the fast-changing technology. One exception is the reference to the original idea of FL from 2016. See Appendix A for the complete list of the sources used.

1.2.1 Federated Learning

Typically, ML is done by one organization on its dataset, trying to find a pattern or model on its data. FL is different because there is an ML environment in which many clients (such as mobile devices or entire businesses) work together to train a model jointly. This is being done while maintaining training data decentralized under the supervision of a central server (Kairouz et al., 2019). This approach allows systems to handle large numbers of clients, typically reaching hundreds or thousands (Niu et al., 2020). Within an FL framework, multiple entities work together to train models without sharing their raw data (Q. Li et al., 2023).

McMahan et al. (2016), on behalf of Google, designed the original FL framework. The intended practical applications include improving image classification and enhancing language models. The original design intended to gain insights from individuals' private data inputs while preserving privacy.

This research takes a different approach to FL. Instead of focusing on a single central server gathering small bits of information from individuals, it uses FL as a tool for data exchange between multiple organizations. This shift from the original design introduces implications such as increased complexity in coordination and potential challenges in maintaining data integrity and privacy across diverse entities.

1.2.2 Architectures

A well-designed architecture is essential for ensuring data integrity, synchronization and secure communication, which are vital for the successful deployment of FL systems in complex environments (Lo et al., 2022). The design of the underlying architecture is a critical component that often receives less attention than optimizing ML algorithms (Q. Li et al., 2023).

Architectural design involves creating a reliable and scalable framework that can support decentralized data and the collaboration of numerous clients, ranging from mobile devices to entire businesses (Kairouz et al., 2019).

In this research, reliable architecture is defined as the system's ability to efficiently handle large-scale data from diverse sources while maintaining high levels of security and privacy, even with unexpected issues such as varying data quality. Scalable architecture refers to the system's ability to manage increasing amounts of data and clients efficiently. As the number of participating clients grows, the architecture should improve its performance, ensuring quick and accurate updates to the global model.

Although numerous architectures exist for FL, few are tailored to specific industries. Aledhari et al. (2020) outline two main distinct architectures: centralized and decentralized. While these FL architectures may have broad applicability across various sectors, adapting them to specific industries like aviation can pose unique challenges. A critical decision in the design process is whether to use a central server or adopt a decentralized approach, as each choice has significant implications for system operation, security and scalability. Figure 1 provides an overview of these different architectures, highlighting the differences between centralized and decentralized approaches.

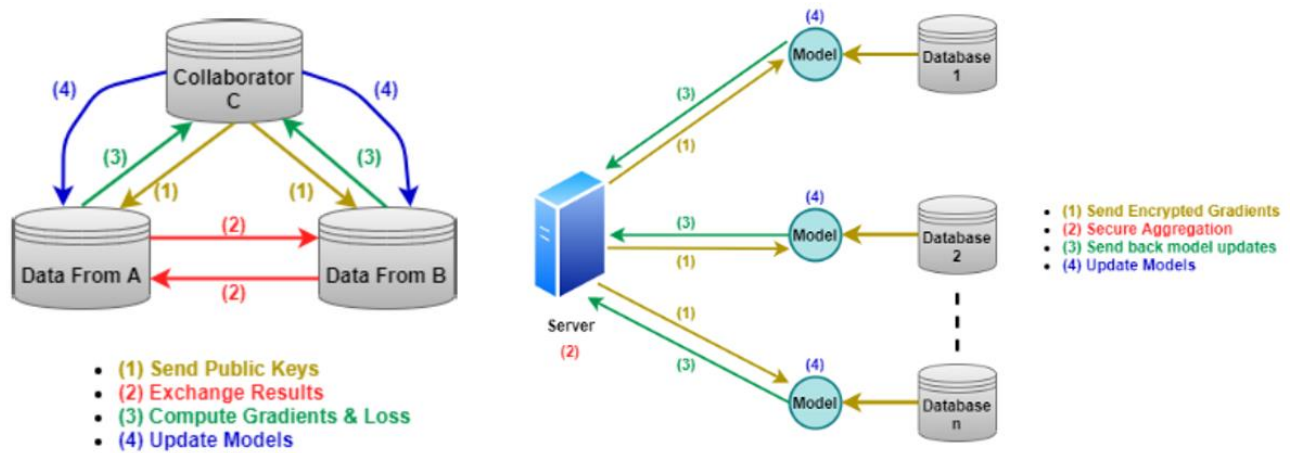


Figure 1: (left) Decentralized, (right) Centralized (Aledhari et al., 2020)

A decentralized approach enhances security and privacy by distributing the computational load and data storage across multiple nodes, eliminating the single failure point. However, this approach complicates synchronization and increases communication overhead, potentially slowing the learning process. Additionally, all participating parties must trust each other, which can be a significant challenge.

On the contrary, a central server can simplify the coordination and aggregation of model updates by acting as a hub, collecting updates from all clients, integrating them into a global model and redistributing the improved model back to the clients. This centralization can streamline processes and ensure consistency across the model. However, it introduces a single point of failure, making the system vulnerable to attacks and privacy breaches. Another critical implication is that all participating entities must trust the central server since the central server will be an aggregator. Ensuring this trust can be challenging to achieve in practice. Balancing these trade-offs for a decentralized or centralized approach is essential for designing an effective FL system for aviation predictive maintenance.

This research emphasizes architectural features because of their fundamental role in the functionality and acceptance of FL systems in aviation. Architectural features refer to design choices such as security protocols and system scalability mechanisms. By focusing on these aspects, this study aims to identify and address the technical and operational challenges that must be overcome to develop an FL system for predictive maintenance in aviation. A well-designed architecture ensures the system's technical reliability and enhances stakeholder trust and cooperation, which are important for its adoption.

1.2.3 Factors Influencing Acceptance

The acceptance of FL systems depends mainly on two factors: privacy and security. Organizations risk a potential breach of privacy when updates sent to the central server reveal details about their sensitive data. Security threats arise when malicious organizations inject incorrect updates, thereby compromising the learning process and degrading the model's performance (Domingo-Ferrer et al., 2022).

Privacy Concerns

Despite FL's inherent protection of local data, the model parameters exchanged during the process can potentially expose sensitive information. Consequently, implementing privacy preservation mechanisms is critical (Lu et al., 2021). While privacy threats in FL are less frequent than security threats, they still pose significant challenges (Mothukuri et al., 2021).

This concern is echoed by Das & Brunschweiler (2019), who identify privacy as the primary challenge within FL systems. Thus, privacy risks can be considered necessary, reinforcing the rationale for integrating advanced privacy-preserving techniques.

Generating a sense of trust among participating organizations is essential. The effectiveness of privacy mechanisms is important, but the perceived feeling of security could be equally relevant. A balance must be created between actual privacy protection and the perceived security that encourages participation in such collaborative settings.

Security Threats

Security attacks in FL predominantly come from the malicious behavior of certain peers who inject false data or manipulative updates to distort the global model. This degrades the model's performance and undermines the entire learning process. Addressing these security threats requires mechanisms to detect and mitigate malicious activities (Domingo-Ferrer et al., 2022). Adequate security measures are needed to maintain the integrity and reliability of the FL system.

In the new system, because of working with established known actors who care about reputation, malicious actors that intentionally will try to hurt the system are less of a concern. However, some participants might be less committed to providing high-quality data and updates. Ensuring all parties are motivated to contribute their best efforts is fundamental for the system's success.

Critical Analysis

This research focuses on the relationship between privacy and security in FL systems. Enhancing privacy in FL requires balancing efficiency and accuracy. One method is adding noise to the data, which means changing some data points. This improves privacy but reduces accuracy. The challenge is finding the right amount of noise. Too much noise lowers model accuracy, while too little increases privacy risks (Mothukuri et al., 2021). This nuance highlights the delicate trade-off between maintaining model performance and ensuring data privacy, and this has yet to be determined.

Furthermore, there could be a difference between the actual importance of privacy and the perceived privacy needed to join this collaboration. Finding a balance where participating organizations feel secure enough to join the FL system is therefore critical. The amount of privacy measures required will be affected depending on which organization or even which type of organization acts as the central party.

Linking this analysis to architectural choices, mainly whether to use a central server, is fundamental. Ensuring sufficient trust in the central server requires the appropriate privacy and security measures.

This trust is critical for successfully adopting the FL system, especially in this sensitive application where organizations might fear losing competitive advantages. Therefore, all measures discussed will be determined by how they contribute to building a secure, reliable and trusted FL system that organizations are willing to adopt.

1.2.4 Federated Learning in Practice

According to the article by Duc Nguyen et al. (2019) in collaboration with KLM, FL is not yet widely applied in the aviation industry. The concept is discussed as a promising future approach to address challenges in data sharing and model training across different stakeholders in aviation.

FL can improve Prognostics and Health Management systems by enabling collaborative training without compromising data privacy and security. However, FL's current application is more theoretical and exploratory than implemented in practical, operational environments within the industry. While reading about FL, banking and healthcare were recurring sectors. Therefore, these are further explored.

Federated Learning in Banking

FL is increasingly discussed in the banking sector, where data privacy and security are paramount. In banking, FL is used to improve fraud detection systems by allowing multiple institutions to collaboratively train models without sharing sensitive customer data (Myalil et al., 2021). This enables banks to enhance fraud detection accuracy while complying with strict data protection regulations. For example, financial institutions often collaborate through consortiums like the Society for Worldwide Interbank Financial Telecommunications for secure financial messaging and Anti-Money Laundering initiatives, where data sharing and joint efforts are already a norm. Using FL, banks can leverage diverse and large datasets, leading to more accurate fraud detection systems while maintaining customer trust by keeping their data private.

Federated Learning in Healthcare

In healthcare, FL is becoming more popular for gaining knowledge to improve the decision-making process (Antunes et al., 2022). Healthcare providers often participate in collaborative research networks and data-sharing agreements for the common goal of improving patient health outcomes (Voss et al., 2024). Initiatives like the European Health Information Initiative exemplify this collaborative mindset. FL allows these institutions to collaborate on training predictive models without sharing raw patient data, thus maintaining patient confidentiality.

Comparison with Aviation

The use of FL in banking and healthcare aligns well with FL's current capabilities. Both sectors benefit from a collaborative environment driven by common goals, financial security in banking and patient health in healthcare. This collaborative mindset facilitates data sharing and joint efforts. This is not the situation in aviation, where competitive commercial interests often make organizations more skeptical about sharing data.

The aviation industry operates under stringent safety regulations. Any predictive model must be highly accurate and reliable, as errors can have catastrophic consequences. In aviation, revisiting safety norms is necessary for implementing predictive maintenance (Torens et al., 2022). Unlike banks or hospitals, the aviation industry involves more diverse stakeholders, including airlines, aircraft manufacturers and maintenance organizations, each with proprietary data and systems. All these factors contribute to FL's more complex and less developed application in the aviation industry.

1.2.5 Knowledge Gap

Despite FL's promising potential to enhance predictive maintenance in aviation, significant gaps remain in the current literature. Most existing research focuses on FL's theoretical advantages and broad applications without considering a comprehensive view of its practical implementation. There is a notable lack of detailed studies addressing the unique challenges posed by the aviation industry's stringent safety standards and the critical need for data privacy and security. These safety standards require high accuracy and reliability in predictive models to prevent catastrophic failures.

Furthermore, the aviation industry involves diverse stakeholders with varying interests, including airlines, aircraft manufacturers and maintenance organizations. This diversity creates additional data-sharing and collaboration complexity, as each stakeholder may have proprietary data and different priorities. While other industries like banking and healthcare have progressed in conceptualizing and implementing FL, aviation remains behind in practical applications. This indicates a need for research that bridges this gap with tailored architectural solutions specifically designed for the aviation sector.

1.3 Research Question

In addressing the complexities of adopting predictive maintenance in the aviation sector via FL and gaining an insight into how an architectural design needs to be made, the following research question comes forth:

'What architectural features should be included in the design of the 'Federated Learning for aircrafts' predictive maintenance system' to be accepted by the stakeholders in the aviation industry?'

In chapter 2 the research approach will be discussed.

2 Research Approach

This research identifies an FL system's architectural specification for sharing maintenance knowledge across the aviation industry and emphasizes ensuring its acceptance. While improvements in code and programming can enhance FL systems, this study focuses on the architectural aspects rather than code optimization. The sub-questions in this research are aligned to ensure a comprehensive exploration of FL systems. They address the design challenges, investigate the willingness to share data and look into collaborating within the aviation industry.

2.1 Sub-Questions and Research Methodologies

The main research question guiding this study is: *'What architectural features should be included in the design of the 'Federated Learning for aircrafts' predictive maintenance system' to be accepted by the stakeholders in the aviation industry?'* The following sub-research questions have been formulated to address this question systematically.

2.1.1 Problem Identification and Motivation

Sub-question 1: *'What are the current challenges and limitations in sharing maintenance data across the aviation industry, and how do these impact data safety and collaboration?'*

Identifying these challenges is crucial because understanding the barriers to effective data sharing is important to developing an FL system that can be widely accepted and implemented in the aviation industry. These challenges directly impact the industry's safety, efficiency and collaborative efforts, making their identification a foundational step in the research process.

The research was conducted through a literature review and interviews with key stakeholders. These methods, as outlined by Galvin (2015), result in a comprehensive data collection on current challenges and their implications for data safety and collaboration. Literature research was done to uncover existing challenges documented in academic papers, industry reports, and regulatory guidelines. This established a solid theoretical foundation and provided a broad understanding of the issues. Interviews with stakeholders from various sectors within the aviation industry offered real-world insights and nuanced perspectives on these challenges.

This approach is particularly effective because the literature review provides an extensive overview of well-documented issues, ensuring that no critical aspect is overlooked. At the same time, interviews capture current, context-specific challenges and practical concerns that might not be fully addressed in the literature. Combining these methods ensures a comprehensive data collection process, integrating the breadth of literature with the depth of stakeholder experiences. This dual approach enhances the reliability and validity of the research findings, ensuring that theoretical and practical dimensions of the problem are thoroughly explored and understood.

This step aligns with the Problem Identification and Motivation phase of the Design Science Research (DSR) methodology (vom Brocke et al., 2020). DSR emphasizes the importance of understanding the problem space and its significance. The DSR methodology, which will be fully explained in section 2.2, guides the structure of this research. Interviews gather qualitative data directly from industry professionals, capturing context-specific challenges that might not be fully represented in the literature.

By integrating these methods, the research addresses both theoretical and practical dimensions of the problem. This phase ensures that the problem is well-defined and justified, setting the stage for subsequent phases where the identified challenges will guide the development of the FL system's architectural specifications. A thorough problem identification process ensures the research tackles relevant and significant issues, creating a foundation for theoretically sound and practically viable solutions.

2.1.2 Define the Objective for a Solution

Sub-question 2: *'What specific technical requirements should the Federated Learning system have to address the identified challenges?'*

Identifying specific technical requirements is fundamental because it ensures that the FL system will effectively address the challenges and limitations previously identified. Clearly defined requirements are necessary for developing a system that meets the needs of all stakeholders, ensuring its functionality, security and acceptance. These requirements form the blueprint for the FL system's architecture, guiding its design and development phases.

The research methods included a detailed study of the thematic analysis of the interviews and a comprehensive review of relevant literature. This involved looking at the coding and categorizing of the qualitative data from interviews with key stakeholders. This process helped identify recurring themes and patterns, providing insights into the stakeholders' specific technical needs and concerns. The literature review complemented this by examining existing research, industry reports and technical documents to ensure that all relevant technical considerations were covered.

The suitability of these methods lies in their ability to capture both broad theoretical insights and specific practical requirements. Analyzing the thematic analysis allows for a deep understanding of stakeholder perspectives, ensuring that the FL system addresses real-world concerns and expectations. The literature review ensures that the technical requirements are grounded in established research and industry best practices, providing a theoretical foundation. This combined approach ensures that the requirements are comprehensive and aligned with stakeholder needs, enhancing the reliability and relevance of the findings.

This step aligns with the Define the Objectives for a Solution phase of the DSR methodology, which emphasizes deriving clear, actionable objectives from the problem definition. As detailed in section 2.2, this DSR phase is crucial in creating specific goals to guide the development process. This research sets concrete goals for the FL system's development by identifying specific technical requirements, ensuring it effectively addresses the identified challenges. This phase provides a clear direction for the subsequent design and development steps, ensuring the FL system is both theoretically sound and practically viable and capable of meeting the complex needs of the aviation industry.

2.1.3 Design and Development

Sub-question 3: *'What architectural features should be included in the design of the 'Federated predictive maintenance system' to fulfill the technical requirements?'*

Identifying the architectural features is essential because they are fundamental to the FL system, ensuring it meets the technical requirements identified in the previous phase. A well-defined architecture guarantees the system's functionality, scalability and security, making it capable of addressing the complex needs of the aviation industry. The purpose is to create a conceptual architectural design, which serves as a high-level blueprint that outlines the components and their interactions within the system.

The research methods included a detailed analysis of the requirement list derived from the previous phase and a thorough review of existing literature. The requirement list from stakeholder interviews and thematic analysis provided specific technical needs and concerns that the FL system must address. This list served as a foundation for identifying architectural features. The literature review complemented this by examining existing studies, industry reports and technical documents related to FL and conceptual platform architectures. This dual approach ensured that the proposed architectural features were grounded in proven methodologies and tailored to meet the identified requirements.

The suitability of these methods lies in their targeted and evidence-based approach. The requirement list ensures that the architectural features directly address stakeholder needs and practical concerns. At the same time, the literature review provides access to a wide range of documented architectural solutions and best practices. By combining these methods, the research ensures that the proposed design is both theoretically sound and practically relevant, enhancing the reliability and applicability of the findings.

This step aligns with the Design and Development phase of the DSR methodology, which involves creating the artifact, in this case, the FL system architecture. The research develops a conceptual platform that meets the established technical requirements by identifying and integrating the necessary architectural features based on the requirement list. This phase transforms the theoretical and practical insights gathered into a concrete design, providing a clear blueprint for the system's development and implementation. A well-defined architecture ensures that the FL system is reliable and scalable, enabling it to handle large-scale data from diverse sources efficiently and efficiently manage increasing amounts of data and clients.

2.1.4 Evaluation

Sub-question 4: 'What is the perception from the stakeholders of the acceptability of the designed conceptual platform architecture for the 'Federated predictive maintenance system'?'

Understanding stakeholder perception is vital because their acceptance is fundamental for successfully implementing and adopting the FL system. Stakeholder feedback provides valuable insights into the system's practical feasibility and usability, ensuring it meets the needs and expectations of those using it.

The research methods included conducting a second round of interviews with key stakeholders from the aviation industry, who had also previously participated in the initial phase of interviews. The interviews were designed to gather qualitative data on stakeholders' opinions and views regarding the acceptability of the developed system. This method allowed for in-depth exploration of stakeholder perceptions, capturing nuanced feedback on the proposed architectural platform.

The suitability of these methods lies in their ability to provide direct and detailed insights from the stakeholders who will be directly involved with the FL system. Using follow-up interviews, the research ensured continuity and built on the established rapport with the stakeholders, leading to detailed feedback. This approach also made it possible to clarify any ambiguities and go deeper into specific concerns raised by the stakeholders. The qualitative data obtained was valuable for validating the design.

This step aligns with the Evaluation phase of the DSR methodology, which involves assessing the artifact's ability to meet its objectives and satisfy stakeholder requirements. By gathering stakeholder feedback through interviews, the research evaluates the acceptability of the FL system's architectural design. This phase ensures that the design meets technical specifications, stakeholder expectations and practical needs.

The validation process helps identify potential issues or areas for improvement, providing a solid foundation for refining the design and ensuring its successful implementation in the aviation industry.

It is important to note that the Demonstration phase of the DSR methodology is excluded in this research. The Demonstration phase typically involves implementing the designed system in a real-world setting to solve one or more instances of the problem, providing practical evidence of its utility and effectiveness. However, several factors made this phase infeasible for this study.

Firstly, there were constraints on resources and time, making it impractical to carry out a real-world implementation within the project's duration. Additionally, the FL system is still in a conceptual phase, meaning that it has not yet reached the level of development necessary for practical application. The conceptual nature of the design implies that while the theoretical framework and architectural specifications have been established, further development and refinement are required before it can be tested in a real-world environment. Consequently, the research proceeded directly to the Validation phase, focusing on gathering detailed stakeholder feedback on the proposed design. This approach allows for assessing the system's acceptability and identifying areas for improvement, ensuring that the final design will be well-aligned with stakeholder needs and practical requirements when it is ready for implementation.

2.1.5 Communication

Sub-question 5: 'Considering the features and the validation, what lessons were learned in the development and evaluation of the Federated Learning system towards knowledge sharing in the aviation industry?'

Identifying lessons learned is necessary because it provides valuable insights into the strengths and weaknesses of the developed system and the overall research process. These insights are needed to refine the FL system and improve future research and development efforts. Understanding what worked well and what challenges were encountered helps make informed decisions and adjustments.

The research methods included a thorough review of the interviews and all research conducted throughout the study. This involved a comprehensive examination of stakeholder feedback obtained during the validation phase and an analysis of the development process documented in research notes and transcripts. Reviewing these data sources allowed for the identification of recurring themes, challenges, and successes encountered during the development and validation of the FL system.

The suitability of this method lies in its ability to evaluate both the system and the design process. The research captures practical insights into the system's acceptability and effectiveness by examining stakeholder feedback. Analyzing the development process highlights areas for improvement in methodology and implementation. This dual approach ensures a well-rounded understanding of the lessons learned, enhancing the reliability and depth of the findings.

This step aligns with the Communication phase of the DSR methodology, which involves disseminating the outcomes and insights gained from the research and development efforts. The study identifies key lessons learned from the system development and validation phases by reviewing the interviews and research conducted. This reflection comprehensively evaluates the entire research journey, presenting findings and insights that can guide future efforts. Reviewing existing data and documentation ensures feasibility, although careful analysis and synthesis of conclusions are required to address potential limitations in the scope of lessons learned. Although discussion is not always formulated as a separate research question, it is a dedicated phase in the DSR methodology. Therefore, it was decided to make it a dedicated sub-question in this research.

2.2 Research Methodology

This study employs the DSR methodology, a problem-solving paradigm that develops innovative artifacts to extend human knowledge (vom Brocke et al., 2020). As outlined earlier, the DSR methodology underpins the structure of this research, guiding each phase systematically. DSR's focus on creating and evaluating artifacts makes it ideal for exploring the technological and social dimensions of FL systems. This practical approach helps develop real-world solutions, making it suitable for investigating how FL systems can be set up and operated in complex industries (Peppers et al., 2007).

DSR uses a systematic process of building and evaluating artifacts to address identified problems. As outlined by vom Brocke et al. (2020) the DSR methodology consists of six steps:

1. **Problem Identification and Motivation:** This step defines the specific research problem and justifies the value of a solution. Understanding the current challenges in aviation predictive maintenance and the potential of FL systems sets the stage for the research.
2. **Define the Objectives for a Solution:** Objectives are derived from the problem definition and what is feasible. These objectives guide the development of the FL system, ensuring it meets the needs of the aviation industry stakeholders.
3. **Design and Development:** In this step, the artifact, the FL system, is created. This involves determining its functionality and architecture based on the defined objectives and existing knowledge.
4. **Demonstration:** The artifact is then demonstrated to solve one or more instances of the problem. This can involve simulations, case studies, or real-world applications to show how the FL system can be used in aviation predictive maintenance.
5. **Evaluation:** The evaluation measures how well the artifact supports a solution to the problem. This involves comparing the defined objectives with the results obtained from the demonstration. Based on this evaluation, further refinements may be made to the artifact.
6. **Communication:** Finally, all aspects of the problem, the artifact and its evaluation are communicated to relevant stakeholders. This ensures that the findings are disseminated effectively and can be utilized by others in the field.

The DSR methodology was chosen for this study because it provides a structured and iterative framework that aligns well with the research's goals. Unlike purely theoretical or strictly empirical methods, DSR allows for both the creation of innovative solutions (in this case, the FL system's architecture) and their evaluation within real-world constraints. This balance between theory and practice is crucial for addressing the specific needs of stakeholders in the aviation industry. Furthermore, DSR's phases ensure that stakeholder input directly affects the solution, making it more likely to be accepted and implemented effectively. Other methodologies, such as traditional case studies or experimental methods, would not offer the same flexibility to develop and validate a practical system within the complex and dynamic environment of aviation predictive maintenance.

Adopting DSR is appropriate for this study for several reasons. First, DSR's iterative processes facilitate continuous improvement based on stakeholder feedback and evaluation, which is important for addressing the aviation industry's dynamic and complex requirements. Second, DSR combines theoretical knowledge with practical insights, ensuring the FL system is theoretically sound and practically viable by leveraging existing theories and real-world stakeholder feedback. Third, DSR's structured approach provides a clear framework for addressing the research question by dividing the research into distinct phases, ensuring thorough exploration and a well-defined solution.

Additionally, DSR's emphasis on the artifact and design process facilitates a deeper understanding of factors influencing the FL system's acceptance and implementation, helping to identify and address potential adoption barriers.

The DSR methodology offers a structured and iterative approach to developing an FL system for aviation predictive maintenance. This method ensures that the design is informed by theoretical knowledge and practical insights from aviation industry stakeholders. DSR facilitates effective, innovative solutions and provides a comprehensive framework for addressing the complex challenges of implementing FL systems in the aviation industry.

The Research Flow Diagram is presented on the next page.

2.3 Research Flow Diagram

See Figure 2 for the research flow diagram. This diagram distinguishes the different research phases corresponding to each sub-research question. It specifies the input required for each phase and the expected outputs, providing an overview of the research process. By illustrating the alignment of sub-questions with the Design Science Research (DSR) methodology phases, the diagram tries to ensure a systematic approach to addressing the main questions.

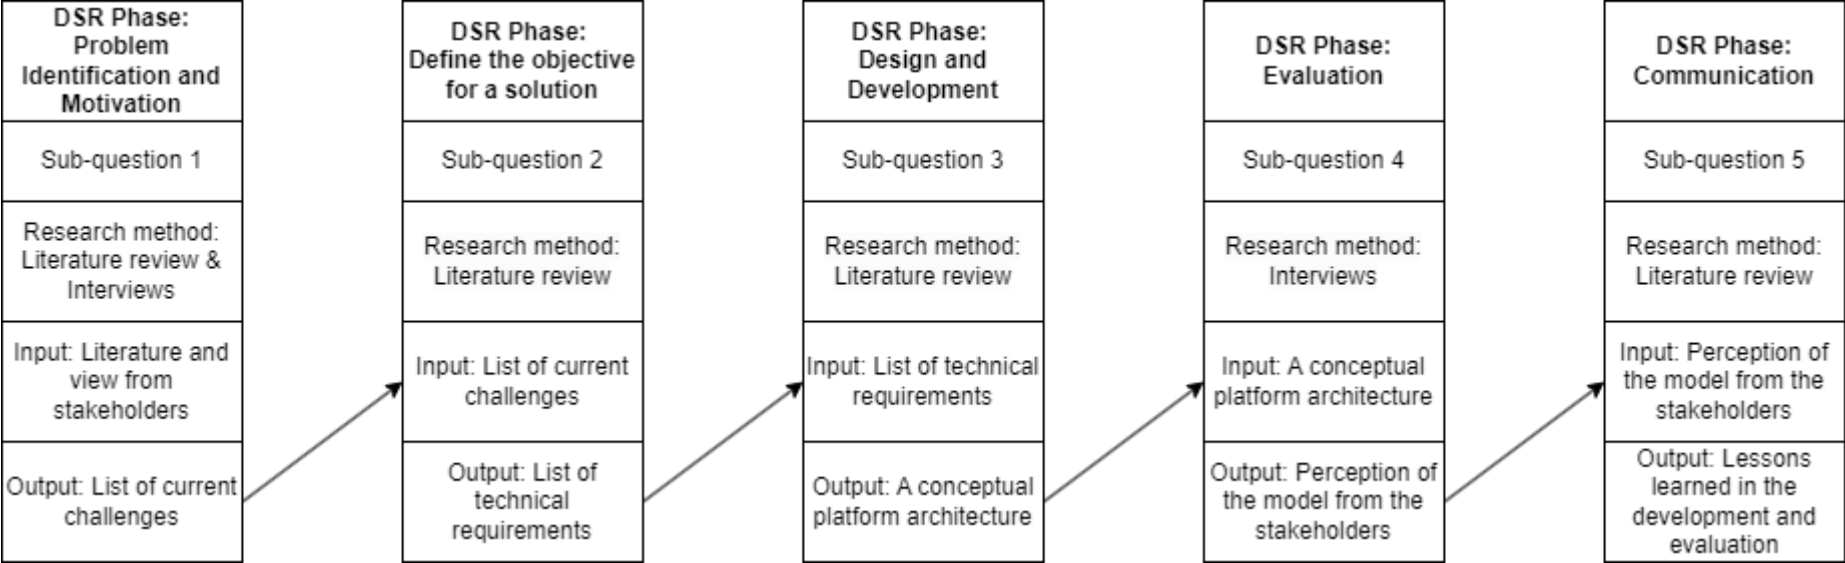


Figure 2: Research Flow Diagram

3 Problem Identification

3.1 Context and Background

Predictive maintenance aims to detect unforeseen equipment breakdowns by constantly monitoring equipment status and issuing early alerts. This early warning enables experts to schedule and execute maintenance, mitigating operational disruptions (Korvesis et al., 2018). However, the aims of different stakeholders are slightly different.

3.1.1 Stakeholders

The main stakeholders involved in predictive maintenance are presented below, based on the model developed by the Independent Data Consortium for Aviation (IDCA) (IDCA, n.d.-a):

Original Equipment Manufacturer (OEM): OEMs are responsible for building aircraft engines, airframes and components. They are interested in predictive maintenance to ensure their products meet high standards of reliability and performance. They typically collect data on engine performance, temperature, vibration and fuel efficiency. This data helps OEMs identify early signs of wear and tear, enabling them to recommend maintenance schedules and design improvements.

Operators: Airlines operate the aircraft and rely on predictive maintenance to minimize downtime and maintain fleet safety. Operators collect data from onboard sensors, such as engine performance, fuel consumption and flight data, such as altitude and airspeed. This information is used to plan maintenance activities and ensure optimal fleet availability.

Maintenance, Repair and Overhaul (MRO): MRO organizations are responsible for performing maintenance, repair and overhaul services. They collect data on component wear, historical maintenance records and service life. This data helps MROs predict the remaining useful life of components and ensure they have the necessary spare parts and workforce available for upcoming maintenance activities.

Regulators: Entities like the European Union Aviation Safety Agency (EASA) and the Federal Aviation Administration (FAA) set regulations and compliance requirements. Regulators are interested in predictive maintenance to enhance aviation safety and reduce regulatory burdens. They collect aggregate data on incidents, maintenance compliance and safety performance across the industry. Additionally, EASA is a Continuing Airworthiness Management Organization (CAMO) which is responsible for managing the continuing airworthiness of aircraft, ensuring compliance with maintenance programs and conducting airworthiness reviews (EASA, 2021).

Service Providers: Service providers offer additional software and data analytics services. They collect and analyze data from various stakeholders to provide predictive maintenance solutions tailored to individual operators and OEMs. This includes performance benchmarking, anomaly detection and predictive analytics.

Lessors: Lessors lease aircraft to operators and are interested in predictive maintenance to protect their investments and maintain fleet value. They often collect data on aircraft usage, engine cycles and maintenance records to ensure the leased aircraft remain in optimal condition.

3.1.2 Initial Challenges

The initial research reveals three primary challenges that a FL system faces. These challenges derived from the literature will also be the base for the interview questions.

Challenge 1: Data Safety, Security and Competitive Advantage

Organizations in the aviation industry are hesitant to share their data due to fears of security breaches and potential losses in competitive advantage. In an FL system, data remains decentralized, allowing each participant to retain control over their dataset. Adamopoulou & Daskalakis (2023) emphasize the importance of high security and data privacy in data integration frameworks, noting that such measures are essential for ensuring participation. This highlights that organizations are unlikely to engage in data-sharing initiatives without confidence in the system's security.

Addressing privacy risks is important because protecting participant identities and sensitive information from being exposed through reverse engineering is challenging. Even aggregated models may inadvertently disclose strategic insights or competitive intelligence, posing a possible threat to competitive advantage (Domingo-Ferrer et al., 2022).

Challenge 2: Legal and Regulatory Compliance

Legal and regulatory compliance poses significant challenges in implementing data-sharing systems, especially with multiple jurisdictions involved. The aviation industry is governed by strict regulations to ensure safety and privacy, making compliance a critical concern. Two main acts that are considered relevant are:

1. The Artificial Intelligence (AI) Act explicitly categorizes the aviation industry as high-risk, imposing strict obligations around risk management, data quality, transparency and human oversight (European Commission, 2024b).
2. The General Data Protection Regulation (GDPR) mandates specific safeguards for personal data protection (EU, n.d.-a).

While these acts apply across different sectors in the European Union (EU), their relevance lies in their extra restrictions on aviation and its global character. The AI Act specifically mentions aviation as a high-risk category, putting more regulations in place for the use of AI in the aviation sector. GDPR is stricter than other data protection laws in different parts of the world. Therefore, this is also explicitly taken into consideration. Duc Nguyen et al. (2019) also emphasize the need for future investigations of designing a standard procedure to get data.

Moreover, regulatory bodies such as EASA and the FAA impose additional layers of compliance, which complicates the implementation of FL systems across borders. Generally, the certification via EASA or the FAA follows the same structures. However, sometimes there are differences (Kurri, 2020). This situation highlights the necessity of a standardized framework that can seamlessly integrate and comply with varying regulations across jurisdictions.

Challenge 3: Proportional Value & Intellectual Property

In collaborative models like FL, ensuring proportional value distribution is challenging and essential for maintaining contributor motivation. Each participant must receive a fair share proportional to their data contribution (Liu et al., 2021). Without effective mechanisms for proportional value attribution, contributors may lose motivation, leading to reduced engagement or even withdrawal from the system (Fan et al., 2022). Therefore, transparent processes for value attribution need to be established to achieve collaboration and maximize the benefits of FL.

Managing IP rights in the developed model is another critical concern in FL. IP issues often revolve around the ownership and usage rights of the improved predictive maintenance models and the insights derived from shared data (Lansari et al., 2023). Establishing clear ownership structures is necessary to prevent disputes and ensure all participants feel adequately compensated for their contributions. The challenge is not only to ensure that the generated IP is valuable but also to protect it from misuse and unauthorized access. Effective IP management is important in sectors where technological advancements are critical competitive differentiators.

To summarize, categorizing these concerns about data safety, legal and regulatory compliance, IP and ensuring proportional value distribution into different challenges provides a logical framework. This framework helps address the multifaceted barriers to implementing FL in aviation. Each category contains distinct yet interconnected issues that stakeholders must overcome to achieve a cohesive and effective FL system. Organizing these challenges makes it easier to understand the needs and concerns of various stakeholders.

3.2 Interview Setup

Key individuals from various segments of the aviation industry were interviewed in depth. The questionnaire was divided into six sections: Organizational Context, Data Collection and Sharing, Data Interests, Federated Learning Participation, Federated Learning Benefits and Governance of the Federated Learning Platform.

- Organizational Context: establishes participants' roles and expertise, providing foundational insights.
- Data Collection and Sharing: examines current practices and barriers, identifying integration challenges.
- Data Interests: assess data availability and value, which is fundamental for predictive model feasibility.
- Federated Learning Participation: explores openness to collaboration, addressing adoption barriers.
- Federated Learning Benefits: evaluates perceived advantages and benefit distribution, which is important for motivation.
- Governance: focuses on trust and regulatory aspects, ensuring alignment with industry standards and legal requirements.

This structured approach ensures comprehensive data collection and an understanding of the factors influencing FL adoption and implementation in aviation predictive maintenance.

The data collected from the interview transcripts were then examined using qualitative thematic analysis, which allowed for identifying recurring themes and insights.

The value of these interviews comes from the detailed, qualitative insights they provide, which cannot be gained from numbers alone. By speaking directly with stakeholders in the aviation sector, the study captures the real concerns, opportunities, and challenges they face. This gives a clearer understanding of how FL could fit into current practices and reveals practical issues that theoretical models might miss. The interviews also allow for unexpected topics to come up, adding more depth to the analysis. These insights complement the technical and operational reviews of FL, giving a complete picture of how it could be adopted in predictive maintenance.

3.2.1 Recruitment

Participants were recruited from six groups: representatives from OEMs, operators, MROs, service providers, aviation experts and regulators. Contacting lessors was unsuccessful, and given their limited involvement in data collection relative to other stakeholders, they were not included in this study. The aviation experts were not linked with a specific company but had over 15 years of highly relevant experience.

Recruitment began through the Independent Data Consortium for Aviation (IDCA), an organization with the mission to gather stakeholders in the aviation sector to create consensus-driven data-sharing governance (IDCA, n.d.-b). IDCA's role must be acknowledged, as they helped facilitate participant contact. However, there was no financial reward for writing this research, ensuring the research remained independent.

The guiding professor, Marcela Tuler de Oliveira, sent the initial email under the assumption that a professor's request would gain more responses. This email introduced the project, briefly explained FL and asked recipients to reply with an available timeslot. However, after a week, only a few responses were received. Subsequently, the president of IDCA sent another request for participation, eventually resulting in twelve responses. Non-respondents were assumed to be uninterested, and no further contact was made. Despite sending emails to 98 addresses, the low response rate is likely due to IDCA's early development stage and varying levels of member commitment.

Letting participants reply to the email might lead to an overly optimistic view because they might already be interested in the topic. Structured interviews were employed to mitigate this risk. The structured nature of the interviews ensured consistency in questioning and allowed for capturing critical perspectives, reducing the likelihood of an unbalanced view. Nonetheless, it is acknowledged that despite these measures, the potential for some degree of positive bias remains.

Despite twelve being a relatively small number, significant effort was invested in achieving this participation. Locating knowledgeable individuals was challenging, but those who participated provided valuable insights. Efforts were made to include different perspectives per group by selecting two representatives. Because the response rate was low, even with the assistance of an IDCA link, the decision was made not to pursue random individuals via LinkedIn due to concerns about time and effectiveness. The contributions of these 12 participants were instrumental in enriching the research with diverse and informed viewpoints.

In this study, theoretical saturation was reached after interviewing 12 participants. Theoretical saturation refers to the point in data collection where no new themes or insights emerge and the data becomes repetitive, indicating that further data collection is unnecessary (Hennink & Kaiser, 2022). This conclusion was drawn as recurring themes consistently emerged across the interviews. Participants from different organizations shared similar concerns, underscoring the validity of these themes. For instance, all participants expressed openness towards the FL concept, highlighting its potential benefits. Additionally, concerns regarding data privacy were frequently mentioned, with participants emphasizing the need for efficient security measures. Another common theme was the necessity for a central entity to oversee and coordinate efforts. The consistency and repetition of these themes across interviews confirmed that theoretical saturation had been achieved, making additional data collection redundant.

3.2.2 Participants

All participants in this study were interviewed via MS Teams because it was logistically easier, as most participants were located on other continents. For consistency, even those in the Netherlands were interviewed online. In general, the interviews lasted approximately 30 minutes. A conversation transcription was the only type of data collected during the interview, using the transcription function of MS Teams. From these transcriptions, structured summaries were made based on the questionnaire. These summaries were shared with the guiding professors using MS Teams. All retraceable data collected was kept in a private file.

Aviation is a technical sector, and technical sectors typically have a more male-oriented workforce. This is also shown by the fact that all interview participants were male. Research done by Both (2021) on factors influencing the willingness to share data found that gender did not have an impact. Therefore, although all interviewees in this study are male, this was not considered a problem.

Potential downsides of online interviews, such as lack of personal interaction and technical issues, were addressed by ensuring stable internet connections and requiring video during interviews. Despite these challenges, the online format did not feel restrictive. Observing both the answers and the manner of delivery allowed for identifying the interviewees' true intentions.

The interviewees hold various technical roles, such as managing digital product programs, implementing blockchain solutions, providing IT and digital services, advancing aircraft studies and innovation, developing predictive maintenance services, and overseeing AI program management. Their senior positions give them significant decision-making power, enabling them to provide credible assessments of the feasibility of FL in predictive maintenance. Because of their extensive experience, they have witnessed the evolution of the aviation sector and are well aware of the common obstacles that arise with the adoption of new technologies. Their professional standing and years of experience further support the reliability of their insights into both the challenges and opportunities that FL presents. For privacy reasons, it is not specified who holds which specific role. Table 2 provides an overview of the demographic data.

Table 2: Demographic Information Interviewees (n=12)

Variables	Count, n (%)
Gender	
Male	12 (100)
Female	0 (0)
Other	0 (0)
Organization	
OEM	2 (17)
Operator	2 (17)
MRO	2 (17)
Service provider	2 (17)
Regulator	2 (17)
Aviation expert	2 (17)
Experience in aviation	
0-4	2 (17)
5-9	1 (8)
10-14	1 (8)
15-19	2 (17)
≥20	6 (50)

Region	
Europe	6 (50)
US	5 (42)
Latin-America	1 (8)

3.2.3 Data Management

The interview transcripts were converted into a summary based on the questionnaire, pseudonymized using the same identifiers and categorized into six groups. For example, 'OEM 1' is the first participant from a company acting as an OEM.

3.2.4 Data Analysis

The data was analyzed using the four steps outlined in the principles of qualitative study and systematic text condensation: Total impression, Identifying and sorting meaning units, Condensation and Synthesizing. This involves moving from initial codes to themes, then extracting meaning and generating descriptions and concepts (Malterud, 2012). See Chapter 3.3 for the Thematic Analysis.

3.2.5 Ethical Considerations

Each participant was required to give written consent after being provided with both oral and written details about the study; only those who consented were included. The study tried to avoid gathering actual flight-related or competitive data. Information collected during interviews was made untraceable. Access to these summaries was limited to the research group, which consisted of the guiding professor and the researcher.

The study's design and consent forms were evaluated by the ethics commission of TU Delft and the data protection officer of the Technology Policy and Management faculty. Their assessment deemed the ethical review sufficient since no personal data was collected. See Appendix B for the full informed consent document.

3.2.6 Interview Questions

There were three versions of the interview questions. In general, they were very similar; however, for the Aviation Experts and the Regulators, the interview had to be modified a bit. For the Aviation Experts, this resulted in an interview that went deeper into the organizational level since discussing specific company-related data collection was less relevant. The interview with the regulators was more focused on the regulation aspect. The interviews with the Aviation Experts were considered for the thematic analysis. The interviews with the regulators were evaluated separately.

There was sufficient time to discuss the 19 questions in dept, despite the questionnaire containing 19 questions. This is because the first five were short factual questions, allowing sufficient time to discuss the remaining questions.

The interview questions are presented in Table 3, each indicating to which challenge it relates.

Table 3: Questions from the Questionnaire

The questionnaire part and question	Challenge
Organizational	
1. What type of organization do you represent?	-
2. Could you describe your organization's primary activities related to aviation maintenance?	-

- 3. What is your role within the organization, and what expertise do you bring to aviation maintenance and data management? -
- 4. How many years have you worked in this position? -
- 5. Where is your company situated? -

Data Collection and Sharing

- 6. What types of engine-related data does your organization currently collect? 1
- 7. Do you share some of this engine data with other parties? 1 & 2
- If 7 is yes:
 - 8a. For what purposes is this data typically shared? 1
 - 9a. How does your organization determine a fair compensation for sharing its data? 1 & 3
- If 7 is no:
 - 8b. What are the barriers to sharing data with others? 1 & 2
 - 9b. How does your organization determine a fair compensation for sharing its data? 1 & 3

Data Interests

- 10. Does your company possess sufficient quality data to independently develop a predictive maintenance algorithm? 1 & 3
- If 10 is yes:
 - 11a. Does your company already have a service based on this algorithm? 1
 - 12a. How widely accepted and utilized is this service among your clients or industry peers? 1
- If 10 is no:
 - 11b. Does your company plan to develop its own predictive algorithm (AI solution) for commercialization as a service to industry peers? 1
 - 12b. How does your organization determine a fair compensation for sharing its data? 3
- 13. Which specific information or data type would significantly impact your organization's predictive maintenance strategies? 1

Federated Learning participation

- 14. In a scenario where it you do not share the actual raw data but you are able to learn from each other and are able to improve predictive maintenance tools. How open is your organization to participating in such a collaborative initiative? 1
- If data protection is not discussed: 1 & 2
 - 15a. Imagine the scenario again; what steps would you consider necessary to protect your data privacy?
- 16. Imagine the scenario again where you learn from each other and collaboratively develop a new model. How would you prefer to manage the co-created intellectual property? 2 & 3

Federated Learning benefits

- 17. How does your organization value its data in terms of potential to enhance predictive maintenance practices? 1 & 3
- 18. How do you prefer the benefits to be distributed among the various parties? 3

Governance

- 19. Which entities would your organization trust to initiate such an initiative? 2

3.3 Thematic Analysis

3.3.1 Overview Analysis

A thematic analysis has been performed using open coding. Six main themes are identified, each of which can be divided into two subthemes. For each subtheme, there is also an indication of how much this theme was described in the interviews. In general, the interviewees showed significant interest in the potential of FL and were keen to share their insights and feedback. Six themes were identified from the thematic analysis: Privacy and Security, Regulatory Compliance, Management of Intellectual Property (IP), Distribution of Benefits, Adoption and Trust and the Impact of Federated Learning on Predictive Maintenance. See Table 4 for the outcome of the Thematic Analysis. In the table, it can be seen that for the total number of participants, ten is used since this is the number of participants without the regulators.

Table 4: Thematic Analysis (n=10)

Themes	Count, n (%)
Privacy and Security	
• 1.1 Data privacy concerns	10 (100)
• 1.2 Importance of security measures	7 (70)
Regulatory compliance	
• 2.1 Adherence to legal standards	5 (50)
• 2.2 Adaptability to legal changes	5 (50)
Management of IP	
• 3.1 Consortium-based IP management	8 (80)
• 3.2 Equal rights and access	7 (70)
Distribution of benefits	
• 4.1. Need for equitable benefit sharing	6 (60)
• 4.2 Reflection of contribution	4 (40)
Adoption and trust	
• 5.1 Openness to collaboration	10 (100)
• 5.2 Need for a trusted entity	8 (80)
Impact FL on predictive maintenance	
• 6.1 Enhancement of predictive tools	9 (90)
• 6.2 Opportunity to share information	8 (80)

3.3.2 Privacy and Security

Data Privacy Concerns (100%)

All participants expressed significant concerns regarding their privacy and the sharing of data. These concerns emphasized the critical nature of maintaining data privacy to protect sensitive information and competitive advantage. Operator 1 described the result of this fear as: *"I am quite sure that the data that we are not mandated to share, we keep that internally."* Additionally, Service Provider 2 raised the need to be sensitive, stating: *"Even if not classified, we may have some more or less sensitive information ... this can be a little bit sensitive."*

Importance of Security Measures (70%)

Most participants (7/10, 70%) discussed the importance of strict security measures to safeguard data within FL systems. Participant 2 OEM described the need for data protection as essential for collaboration. Aviation Expert 2 reaffirmed this need by stating: *"Whatever we can do to prevent that, we will have to put in place to ensure that these databases are secure."*

The other three participants did not focus on the security measures because they believed that FL's inherent architecture already provides sufficient privacy by not sharing raw data. For instance, Service Provider 1 mentioned that he felt no concern when people got the model, not the actual data.

Analysis and Implications

The unanimous concern for data privacy (10/10, 100%) and the emphasis on security measures (7/10, 70%) highlight a general concern about data vulnerability. Within this context, an important notion is to differentiate between protection mechanisms aimed at external threats and those designed to mitigate risks from insiders. Diverging opinions exist on the feasibility and practicality of implementing these measures. Service Provider 1 expressed skepticism, stating: *"Organizations should be open to it, but you would get a lot of questions about the privacy of data."*

This skepticism contrasts with the more optimistic views of others who believe security measures can effectively protect data. The varying confidence levels in data security measures indicate a potential barrier to collaboration unless these concerns are adequately addressed.

Participants' perspectives can be explained by their organizational roles. Operators and service providers directly managing large amounts of sensitive customer data naturally prioritize data security. In contrast, OEMs and MROs, while still concerned about data security, may view FL as a strategic advantage to enhance their services and products.

3.3.3 Regulatory Compliance

Adherence to Legal Standards (50%)

Half the participants (5/10, 50%) emphasized complying with existing legal standards and regulations when sharing and managing data. For example, MRO 2 pointed out that legislative restrictions, collective labor agreements and bilateral agreements with OEMs strictly regulate data-sharing practices. OEM 2 indicated the complexity of compliance by stating: *"Our current business model does not support or necessitate data sharing externally. Data we receive is kept within the company and used by internal departments only."*

Adaptability to Legal Changes (50%)

Another half of the participants (5/10, 50%) expressed the need for adaptability in their data management practices to comply with evolving legal requirements. Aviation Expert 2 described the changing regulatory landscape: *"Governments are getting into the act, especially with AI. The Biden administration is discussing rules for accessing data, and the UK Government has conducted workshops on how AI will change data-related algorithms."*

Analysis and Implications

Half of the participants' emphasis on regulatory compliance (5/10, 50%) indicates a significant barrier to data sharing. This legalistic approach underscores a cautious attitude, likely shaped by strict industry regulations and potential legal repercussions of data misuse. The need for adaptability highlights the dynamic nature of regulations, requiring organizations to remain agile to stay compliant. The split between compliance and adaptability (5/10, 50%) indicates a tension between adhering to current regulations and the need to remain flexible in response to new laws, complicating collaborative efforts.

The diverging views reflect the different roles and responsibilities within the organizations. OEMs and MROs, which operate under strict regulatory frameworks, focus more on compliance to avoid legal repercussions. In contrast, aviation experts and service providers recognize the evolving nature of regulations and the need for adaptability to stay competitive. Balancing the need for strict regulatory compliance with the flexibility required to innovate and adapt to new technologies will be critical for collaborative efforts.

3.3.4 Management of Intellectual Property

Consortium-Based IP Management (80%)

Most participants (8/10, 80%) favored a consortium-based approach for managing co-created IP, ensuring fair access and preventing monopolization by any single entity. Aviation Expert 2 emphasized that no individual party would gain an unfair advantage if a consortium had the IP.

This sentiment was shared by OEM 2, who stated: *"Intellectual property created through such collaborations should be shared among all participants to prevent barriers to innovation and ensure fair benefits."* Aviation Expert 1 also emphasized the need for a consortium-led approach: *"A consortium-led approach is essential to tackle industry challenges."*

Equal Rights and Access (70%)

A significant portion (7/10, 70%) advocated for equal rights and access to co-created IP among consortium members to gain trust and collaboration. OEM 2 stressed this importance by stating that the management of co-created IP should be shared among all participants to avoid innovation barriers. Service Provider 1 highlighted the importance of sharing: *"Anybody who has participated in a meaningful way... should have an equal share. Moreover, that community would have equal rights and would not charge each other any fees."*

Analysis and Implications

The strong support for consortium-based IP management (8/10, 80%) and equal rights to co-created IP (7/10, 70%) reflects a desire for collaborative equity and innovation. However, concerns about potential disagreements over contributions and ownership highlight the need for clear, fair IP management guidelines within consortia to mitigate potential conflicts and promote cooperative innovation. Operator 2 expressed concern that disagreements over contributions and ownership can arise even with consortium-based IP management, which can stall progress. This concern highlights the need for clear guidelines to prevent disputes and ensure all participants feel adequately compensated for their contributions.

3.3.5 Distribution of Benefits Equitable

Need for Equitable Benefit Sharing (60%)

Over half of the participants (6/10, 60%) stressed the importance of equitable benefit sharing in collaborative projects. OEM 2 stated that the benefits should be distributed based on each party's contribution, focusing on enhancing efficiency and effectiveness within the industry rather than on monetary gains. Operator 2 also highlighted this perspective: *"The organization believes that the distribution of benefits should ensure a win-win situation for all parties involved."*

Reflection of Contribution (40%)

Some participants (4/10, 40%) suggested that the distribution of benefits should reflect the contributions made by each member, considering both data provision and involvement in model development. MRO 2 emphasized that those who provide more data or play a more significant role in model development should receive greater benefits.

Analysis and Implications

The emphasis on equitable benefit sharing (6/10, 60%) and the consideration of contributions (4/10, 40%) highlight the importance of fairness and recognition in collaborative efforts. Participants' focus on equitable distribution suggests an understanding that collaboration should enhance overall industry efficiency and effectiveness. Developing transparent, agreed-upon metrics for contribution and benefit distribution will be vital to address these challenges and maintain collaboration motivation.

The preference for equitable benefit sharing indicates that most participants prioritize long-term collaboration and inclusivity over individual gains. However, the suggestion to reflect contributions in benefit distribution shows an understanding that not all contributions are equal and should be recognized appropriately. Balancing these perspectives will be needed to maintain motivation and fairness within the consortium.

Openness to Collaboration (100%)

All participants (10/10, 100%) expressed a strong willingness to participate in collaborative initiatives like FL, provided their concerns about data privacy and security are addressed. MRO 2 underlined this: *"We are very open to participating in collaborative initiatives, particularly those that allow for Federated Learning where data privacy can be maintained while still gaining collective insights."*

Need for a Trusted Entity (80%)

A majority (8/10, 80%) emphasized the need for a trusted, neutral entity to manage collaborative efforts and ensure fair governance. MRO 2 explained that trust would be extended through a consortium of other transparent groups like EASA or the FAA. Operator 2 echoed this sentiment: *"We would trust independent bodies, particularly academic institutions like TU Delft, to lead initiatives involving data sharing and collaborative development."*

The other two participants did not mention the need for a trusted entity, possibly because they believe their existing industry relationships and agreements provide sufficient oversight and governance. OEM 2 stated: *"We have well-established bilateral agreements that already ensure fair and effective collaboration."*

Analysis and Implications

The unanimous willingness to collaborate (10/10, 100%) and the need for a trusted entity (8/10, 80%) underscore the importance of trust and governance in successful FL initiatives. This willingness indicates a recognition of the collective benefits of collaboration. However, ensuring transparent governance and selecting efficient, neutral entities to lead collaborative efforts will be critical to achieving FL's potential. The call for trusted entities, such as academic institutions or industry organizations, reflects a desire for impartial oversight to ensure fairness and transparency. This trust in neutral parties underscores the importance of neutral governance in developing innovation. Aviation Expert 2 expanded on this idea: *"Federated Learning and other collaborative models require a shift from competitive to cooperative mindsets, which is challenging but necessary for the advancement of industry practices."*

3.3.7 Impact of Federated Learning on Predictive Maintenance

Enhancement of Predictive Tools (90%)

A significant majority (9/10, 90%) of participants highlighted the potential of FL to enhance predictive maintenance tools. OEM 1: *"Federated Learning can significantly improve our predictive models by incorporating diverse data sets, leading to more accurate maintenance predictions."*

Operator 1 underlined this potential: *"We would like to have more sensor data from other parties to enhance our health monitoring. Having a bigger batch of observations could significantly improve our algorithms and predictive accuracy."*

Opportunity to Share Information (80%)

Most participants (8/10, 80%) recognized the opportunity to share information and insights through FL. Operator 2 mentioned: *"By sharing insights derived from our data without exposing raw data, we can collaboratively improve maintenance practices."* The two participants who did not emphasize this opportunity were concerned about the potential risks of sharing insights that could inadvertently reveal sensitive information.

Analysis and Implications

The potential for FL to enhance predictive maintenance tools and the opportunity to share information were recognized by almost all participants. This underscores the value of collaborative data integration for advancing predictive maintenance capabilities.

However, the success of these initiatives depends on overcoming the challenges related to data privacy, security and trust. Participants' recognition of the benefits of FL indicates a readiness to adopt innovative approaches for predictive maintenance. The challenge lies in creating a framework that addresses data privacy and security concerns while promoting collaboration. Transparent governance and fair benefit-sharing mechanisms will be important to realize the potential of FL in predictive maintenance.

3.3.8 Conclusion

From the interviews, there was a clear dual recognition: the intrinsic worth of their data and its amplified potential when integrated with external datasets. This confidence, however, also leads to a protective stance towards data sharing. Participants frequently emphasized the need for strict data privacy and security measures, reflecting concerns about maintaining competitive advantage and protecting sensitive information. This protective attitude poses a potential barrier to fully realizing FL's benefits, requiring reliable and transparent data-sharing frameworks to mitigate fears and encourage collaboration.

The participants' nuanced understanding of the importance of integrating data from various entities is evident in their responses. They recognize that while their data is valuable, achieving the highest levels of predictive accuracy is insufficient. This is reflected in their willingness to engage in FL initiatives, provided data privacy is maintained.

Moreover, the participants' emphasis on equitable benefit distribution and fair IP management underscores the importance of trust and fairness in collaborative efforts. There is a clear recognition that FL's success depends on technological integration and establishing trust among participants. By advocating for consortium-based IP management and equitable benefit-sharing models, participants highlight the need for governance structures that ensure fair access and prevent monopolization. This approach tries to balance competitive interests with the collective goal of enhancing predictive maintenance capabilities across the industry.

The analysis also revealed a gap in participants' awareness of the specific dangers associated with FL. Many are not fully aware of the possible attacks. This gap highlights the need for comprehensive education and discussion about these vulnerabilities to ensure informed participation in FL initiatives.

In conclusion, while there is a high level of confidence in the value of individual data sets, there is also a clear understanding of the potential of combined data. This balance needs to be addressed to advance with FL in aviation maintenance. To build the necessary trust and safeguard data effectively, addressing the lack of awareness about FL's specific security threats is important. This can be achieved through targeted education and efficient security measures. These insights underscore the need for strong governance and transparent collaboration practices to drive innovation and improve maintenance practices through FL initiatives.

3.4 Insights Regulators

The interview with the regulators, distinct in its focus on regulations, provides unique insights. This section presents the outcomes and analysis of that interview, serving as a valuable complement to the thematic analysis of all other interviews. The interviews with Regulators 1 and 2 from EASA shed light on the challenges and opportunities of implementing FL in aviation maintenance. The complete interview questions are in Appendix C: Interview Questions Regulators.

3.4.1 Cautious Yet Progressive Stance

The regulators emphasized EASA's cautious approach to AI adoption, prioritizing safety and compliance while recognizing the promising potential of FL in predictive maintenance. This perspective reflects the agency's focus on balancing innovation with regulatory strictness. According to Regulator 1: *"I would not say we are conservative, but we will always work on the safe subsets of all of what the technology can bring."*

This cautious approach ensures that technological advancements do not compromise safety, which is paramount in aviation. A balanced perspective is needed for sustainable innovation, as it allows for gradually integrating new technologies while mitigating risks. By focusing on safe subsets of technology, EASA aims to maintain high safety standards while progressively incorporating AI advancements.

3.4.2 Potential of FL

The potential of FL to enhance predictive maintenance while maintaining data privacy is recognized as a significant advancement. The regulators highlighted FL as a solution to the data-sharing challenge in aviation. Regulator 2 noted: *"Federated Learning could be quite an interesting concept in predictive maintenance. We all know that sharing these types of data between industry players remains a challenge, so each of them benefiting from the outcome of a federated model without sharing the data could be very interesting."*

This recognition underscores the strategic importance of FL in the aviation industry. FL addresses privacy concerns and stimulates collaboration among industry players by enabling predictive maintenance without requiring direct data sharing. This approach can improve maintenance models, enhance operational efficiency and reduce downtime. However, the success of FL depends on the development of predictive models and the establishment of data governance frameworks.

3.4.3 Shifting to Condition-Based Maintenance

Regulator 1 highlighted the importance of predictive models that could adapt maintenance regulations toward a more condition-based approach: *"If predictive models are reliable, there is potential to rethink the schedule of obligatory checks to adapt to actual conditions and usage rather than fixed intervals."*

This shift from fixed interval to condition-based maintenance could revolutionize aviation maintenance practices. Condition-based maintenance allows for more efficient and cost-effective practices by tailoring maintenance schedules to aircraft parts' actual condition and usage. This can reduce unnecessary maintenance, lower costs and increase aircraft availability. The transition, however, hinges on the accuracy and reliability of predictive models, which remains a significant challenge.

3.4.4 Challenges in Implementation

The significant challenges identified include ensuring the explainability of complex AI models, designing effective human-AI interfaces and managing uncertainty. These challenges highlight the complexities involved in implementing FL:

- Explainability is fundamental for gaining regulatory approval and industry trust. Without clear explainability, stakeholders may be reluctant to adopt AI technologies due to many AI models' "black box" nature.
- Human-AI Interfaces: Effective interfaces are needed for seamless integration into existing workflows. Poorly designed interfaces can lead to user resistance or errors in interpretation and action.
- Managing Uncertainty: Maintaining safety and reliability in autonomous systems is essential. Autonomous systems must be able to handle unexpected situations and provide predictable outcomes to ensure safety.

3.4.5 Conclusion

The most critical aspect presented was the importance of explainability. Ensuring that AI models are transparent and understandable is needed for regulatory approval, industry trust and effective integration into aviation systems. Explainability forms the foundation for addressing challenges such as ethical considerations, human-AI interfaces and uncertainty management. Adopting AI technologies in aviation without improved explainability will face significant resistance and barriers.

The interview with EASA regulators provides valuable insights into the cautious yet optimistic approach to AI adoption in aviation. The regulators' focus on safety, compliance and ethical considerations aligns with broader industry trends. This highlights the complexity of integrating AI into regulated environments. This interview's findings complement the other interviews' thematic analysis, offering a comprehensive view of the challenges and opportunities in implementing FL in aviation maintenance.

4 System Requirements

4.1 Overview Requirements

This chapter provides an overview of all the system requirements derived from the literature analysis and interviews. First, an overview of the requirements will be provided. Then, they will be discussed individually. In Table 5 the requirements are presented based on themes derived from the interviews.

Table 5: Overview Requirements

ID	Requirement	Purpose	Theme
R1	The system must have strict security measures.	To protect sensitive data and ensure data safety and security.	Privacy and Security
R2	The system must have a reliable data-sharing framework.	To facilitate efficient and secure data sharing.	Privacy and Security
R3	The system must ensure compliance with existing legal standards.	To meet legal and regulatory requirements.	Regulatory Compliance
R4	The system must be adaptable to evolving regulations.	To remain compliant with future legal changes.	Regulatory Compliance
R5	The system must develop a consortium-based approach for managing co-created IP.	To ensure fair access and avoid monopolization of intellectual property.	Management of intellectual property
R6	The system must establish mechanisms for equitable benefit sharing.	To maintain motivation and collaboration among participants.	Distribution of Benefits
R7	The system must develop transparent processes for value attribution.	To build trust and engagement among participants.	Distribution of Benefits
R8	The system must be more accurate than current predictions.	To provide superior predictive capabilities.	Impact of FL on Predictive Maintenance
R9	The system must ensure that AI models are explainable and transparent.	To gain regulatory approval and industry trust.	Adoption and Trust
R10	The system must have training programs for stakeholders.	To enhance stakeholder knowledge and engagement.	Adoption and Trust
R11	The system must appoint a trusted, neutral entity for collaborative management and fair governance.	To ensure continuous monitoring and fair governance.	Adoption and Trust

4.2 Diverging Opinions and Inclusion Criteria

This chapter outlines the system requirements for developing the FL platform, gathered through a literature review, interviews, and thematic analysis. The aim is to ensure the system addresses stakeholders' needs and concerns while balancing overall system goals.

Different opinions emerged among stakeholders regarding priorities and necessary features during the requirements-gathering process. For example, while all stakeholders agreed on the importance of data privacy and security, there were varying views on the extent of additional measures needed beyond what FL offers. A balanced approach is required to address these critical concerns while remaining technically feasible.

There were also differing opinions on the distribution of benefits. Some stakeholders favored equal distribution for all participants, while others emphasized measuring individual contributions. The goal is to develop equitable benefit sharing and proportional value attribution mechanisms to ensure fair compensation, balancing complexity with transparency.

Another significant divergence concerned who should serve as the central coordination entity. Some stakeholders trusted academic institutions or regulators to manage collaborative efforts, while others preferred a consortium-based approach. The requirement must ensure that the chosen entity is trusted, neutral and capable of effectively managing governance and regulatory complexities.

Functional requirements specify the system's behaviors and functions, focusing on what it should do. Non-functional requirements describe the system's qualities, focusing on how it performs (Eckhardt et al., 2016). Careful judgment is required to distinguish between these requirements and prioritize urgency and importance, ensuring functionality and quality. Since the platform is in a developing phase, it is important to focus on fundamental aspects such as privacy, security, compliance and equitable benefit sharing. For each requirement, a statement will be made to determine how the purpose of the requirement will be fulfilled. One of the main objectives will be to build industry trust and encourage collaboration.

4.3 Privacy and Security

R1: The system must have strict security measures. (Functional)

This requirement protects sensitive data and ensures data safety and security. It is derived from all participants' unanimous concern for data privacy. Seventy percent of participants emphasized the need for strong security measures, highlighting the importance of safeguarding data to protect sensitive information and competitive advantage. Ensuring data privacy will protect the information and build trust among stakeholders.

Purpose Fulfilled: Sensitive data remains protected and is used to train locally, and a secure aggregation scheme is used.

R2: The system must have a reliable data-sharing framework. (Functional)

This requirement facilitates efficient and secure data sharing. It stems from participants' emphasis on security measures. A secure data-sharing framework will develop trust and encourage active participation from all stakeholders by establishing data-sharing protocols and ensuring participant acceptance. By providing clear guidelines and secure channels for data sharing, the system can mitigate risks associated with unauthorized access and data breaches.

Purpose Fulfilled: Stakeholders share data confidently, knowing that protocols are in place and participant satisfaction regarding data security is high.

4.4 Regulatory Compliance

R3: The system must ensure compliance with existing legal standards. (Functional)

This requirement is to meet legal and regulatory requirements. Ensuring compliance with regulations will help gain trust and support from regulators and other stakeholders. This includes adhering to data protection laws such as GDPR and industry-specific regulations. Legal compliance is needed to avoid penalties and maintain the integrity and credibility of the system.

Purpose Fulfilled: The system successfully passes all audits and consistently adheres to relevant regulations.

R4: The system must be adaptable to evolving regulations. (Non-Functional)

This requirement is to remain compliant with future legal changes. Participants expressed the importance of adapting to legal changes to stay compliant. Monitoring emerging regulatory trends and proactive planning will be critical to maintaining compliance. Adapting to new regulations swiftly will help the system remain ahead in a dynamic regulatory environment.

Purpose Fulfilled: The system swiftly integrates regulatory changes, maintaining uninterrupted compliance.

4.5 Management of Intellectual Property

R5: The system must develop a consortium-based approach for managing co-created IP. (Functional)

This requirement is to ensure fair access and avoid monopolization of IP. Participants favored a consortium-based approach to ensure fair access to co-created IP and prevent monopolization. This approach involves creating agreements that define the ownership, use and distribution of IP among consortium members. Establishing a governance structure within the consortium will help manage IP rights effectively. By developing a collaborative environment, the system can ensure that all contributors benefit pretty from the co-created IP.

Purpose Fulfilled: IP is managed collaboratively, ensuring equitable access and preventing monopolization.

4.6 Distribution of Benefits

R6: The system must establish mechanisms for equitable benefit sharing. (Functional)

This requirement is to maintain motivation and collaboration among participants. Ensuring equitable benefit sharing will motivate participants to contribute actively, promoting sustained engagement and collaboration. These mechanisms could include typical profit-sharing agreements, recognition programs or other forms of compensation. By aligning incentives with contributions, the system can maintain high levels of participation and collaboration.

Purpose Fulfilled: Benefits are distributed fairly, ensuring all participants agree with the sharing process, leading to high collaboration.

R7: The system must develop transparent processes for value attribution. (Functional)

This requirement aims to build trust and engagement among participants. Establishing transparent processes for evaluating and compensating contributions ensures that participants understand how they are rewarded. This transparency increases trust and encourages continuous contributions and innovation.

Purpose Fulfilled: Contributions are transparently rewarded, enhancing participant trust and engagement.

4.7 Impact of Federated Learning on Predictive Maintenance

R8: The system must be more accurate than current predictions. (Functional)

This requirement is to provide superior predictive capabilities. If the models developed through FL are not significantly better than the current ones, the effort and resources invested in implementing FL would not be justified. Participants highlighted the potential of FL to improve predictive maintenance. By leveraging data from multiple sources, FL can develop more comprehensive models than those using a single dataset. This approach allows for capturing a more extensive range of conditions and scenarios.

Purpose Fulfilled: Predictive accuracy is significantly improved, reducing unexpected maintenance issues.

4.8 Adoption and Trust

R9: The system must ensure that AI models are explainable and transparent. (Functional)

This requirement is to gain regulatory approval and industry trust. Explainable and transparent AI models are necessary to achieve these goals. Providing stakeholders with clear explanations of how models work will build confidence in their use. Transparency will also help identify and mitigate biases in AI models, ensuring fair and ethical outcomes.

Purpose Fulfilled: AI models receive regulatory approval and are trusted by users.

R10: The system must have training programs for stakeholders. (Non-Functional)

This requirement enhances stakeholder knowledge and engagement. Participants emphasized the need for training to promote informed participation and collaboration. These programs should cover technical aspects of FL, data security and the operational benefits of predictive maintenance. Educated stakeholders are likelier to engage actively and contribute meaningfully to the system.

Purpose Fulfilled: Stakeholders are well-informed and actively participating, leading to effective contribution and collaboration.

R11: The system must appoint a trusted neutral entity. (Functional)

This requirement ensures collaborative, continuous management and fair governance. Trust is the start of successful collaboration, and participants emphasized the need for a trusted entity to manage collaborative efforts and ensure fair governance. The trusted entity is also necessary for aggregating the parameters from the different locally trained models, a critical function in an FL system. Fair governance will promote a collaborative and trustworthy environment.

Purpose Fulfilled: A neutral entity effectively manages governance and parameter aggregation, resulting in high trust and smooth collaboration among participants.

4.9 Conclusion

This chapter has outlined the critical functional and non-functional requirements for developing an FL platform for predictive maintenance in aviation. These requirements address the needs and concerns of various stakeholders while balancing technical feasibility and system goals.

Most requirements are functional since they focus on the system's behavior and functions, which are essential for successful operation.

Two requirements are classified as non-functional: the system must be adaptable to evolving regulations (R4) and must have training programs for stakeholders (R10). These are considered non-functional because the system could be developed without them. However, they are included due to their critical importance for long-term compliance, usability and acceptance.

The next step involves critically evaluating how to achieve the desired purposes of these requirements. This is necessary to ensure the system meets all specified goals.

5 Architectural Features

5.1 Overview of the System

5.1.1 Architecture

This chapter presents an ArchiMate model developed to address the specific requirements identified in the research for a system encompassing multiple organizations and enterprises. ArchiMate is an open and standardized modeling language. ArchiMate is designed to provide a comprehensive framework for describing, analyzing and visualizing the architecture of an organization (Jonkers et al., 2011). It is widely used for enterprise architecture because it supports a clear and structured way of representing complex systems, making it easier to understand and manage the interactions and dependencies among different components (Silhavy et al., 2017). Creating an ArchiMate model was ideal for this research because it effectively represents and integrates the diverse technology aspects and interactions across multiple organizations.

The ArchiMate framework in this model is divided into four primary layers: Technology, Application, Business Processes and Industry Players. The Technology layer covers the infrastructure and hardware components. The Application layer deals with software applications and data flows. The Business Processes layer focuses on the organizational operations and services. The Industry Players layer represents the various stakeholders and external entities interacting with the system. This structured approach ensures that the architecture aligns with business objectives and technical requirements, providing a comprehensive view of the interconnected organizations.

Throughout the development of an ArchiMate model, several implicit choices need to be made about the level of abstraction and details. More abstract models can simplify communication and highlight key relationships across organizations but may overlook critical operational information. On the other hand, highly realistic models can capture more details of organizations but may be harder to understand and use. Additionally, decisions regarding the level of detail must be carefully considered to ensure the model is comprehensive and manageable. For the designed model, the chosen level of abstraction aims to balance clarity and detail. It focuses on essential interactions and dependencies while excluding less critical elements to ensure operability.

5.1.2 Types of Relations in ArchiMate

A fundamental feature of the ArchiMate modeling language is its ability to represent various relationships between elements within a model. These relationships clarify how different components interact and depend on each other, providing a structured and comprehensive view of the system (ArchiMate, n.d.). The different relations used in the designed system can be found in Figure 3 on the next page.

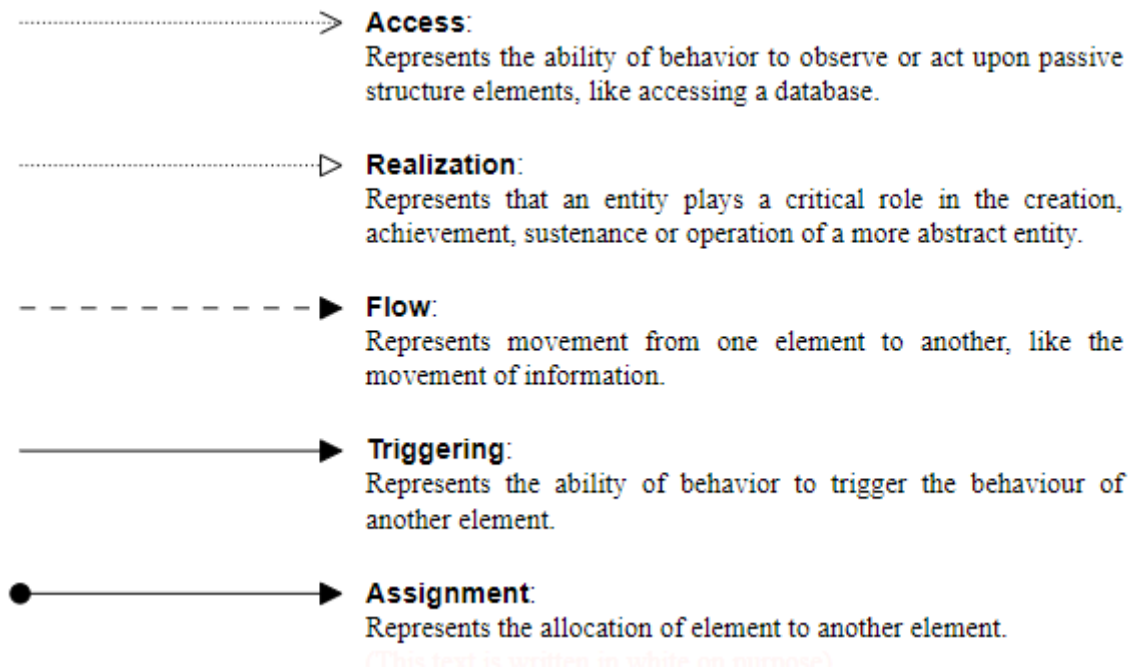


Figure 3: Relations ArchiMate

5.1.2 Overview System

See Figure 4 on the next page for an overview of the system. The description of each component and its relationships with others is textually explained in this chapter. However, this information is also systematically described in the appendices:

- Appendix D: Contains a list describing each component.
- Appendix E: Provides an overview of all relationships.
- Appendix F: Provides an overview of the component contribution to the requirements.

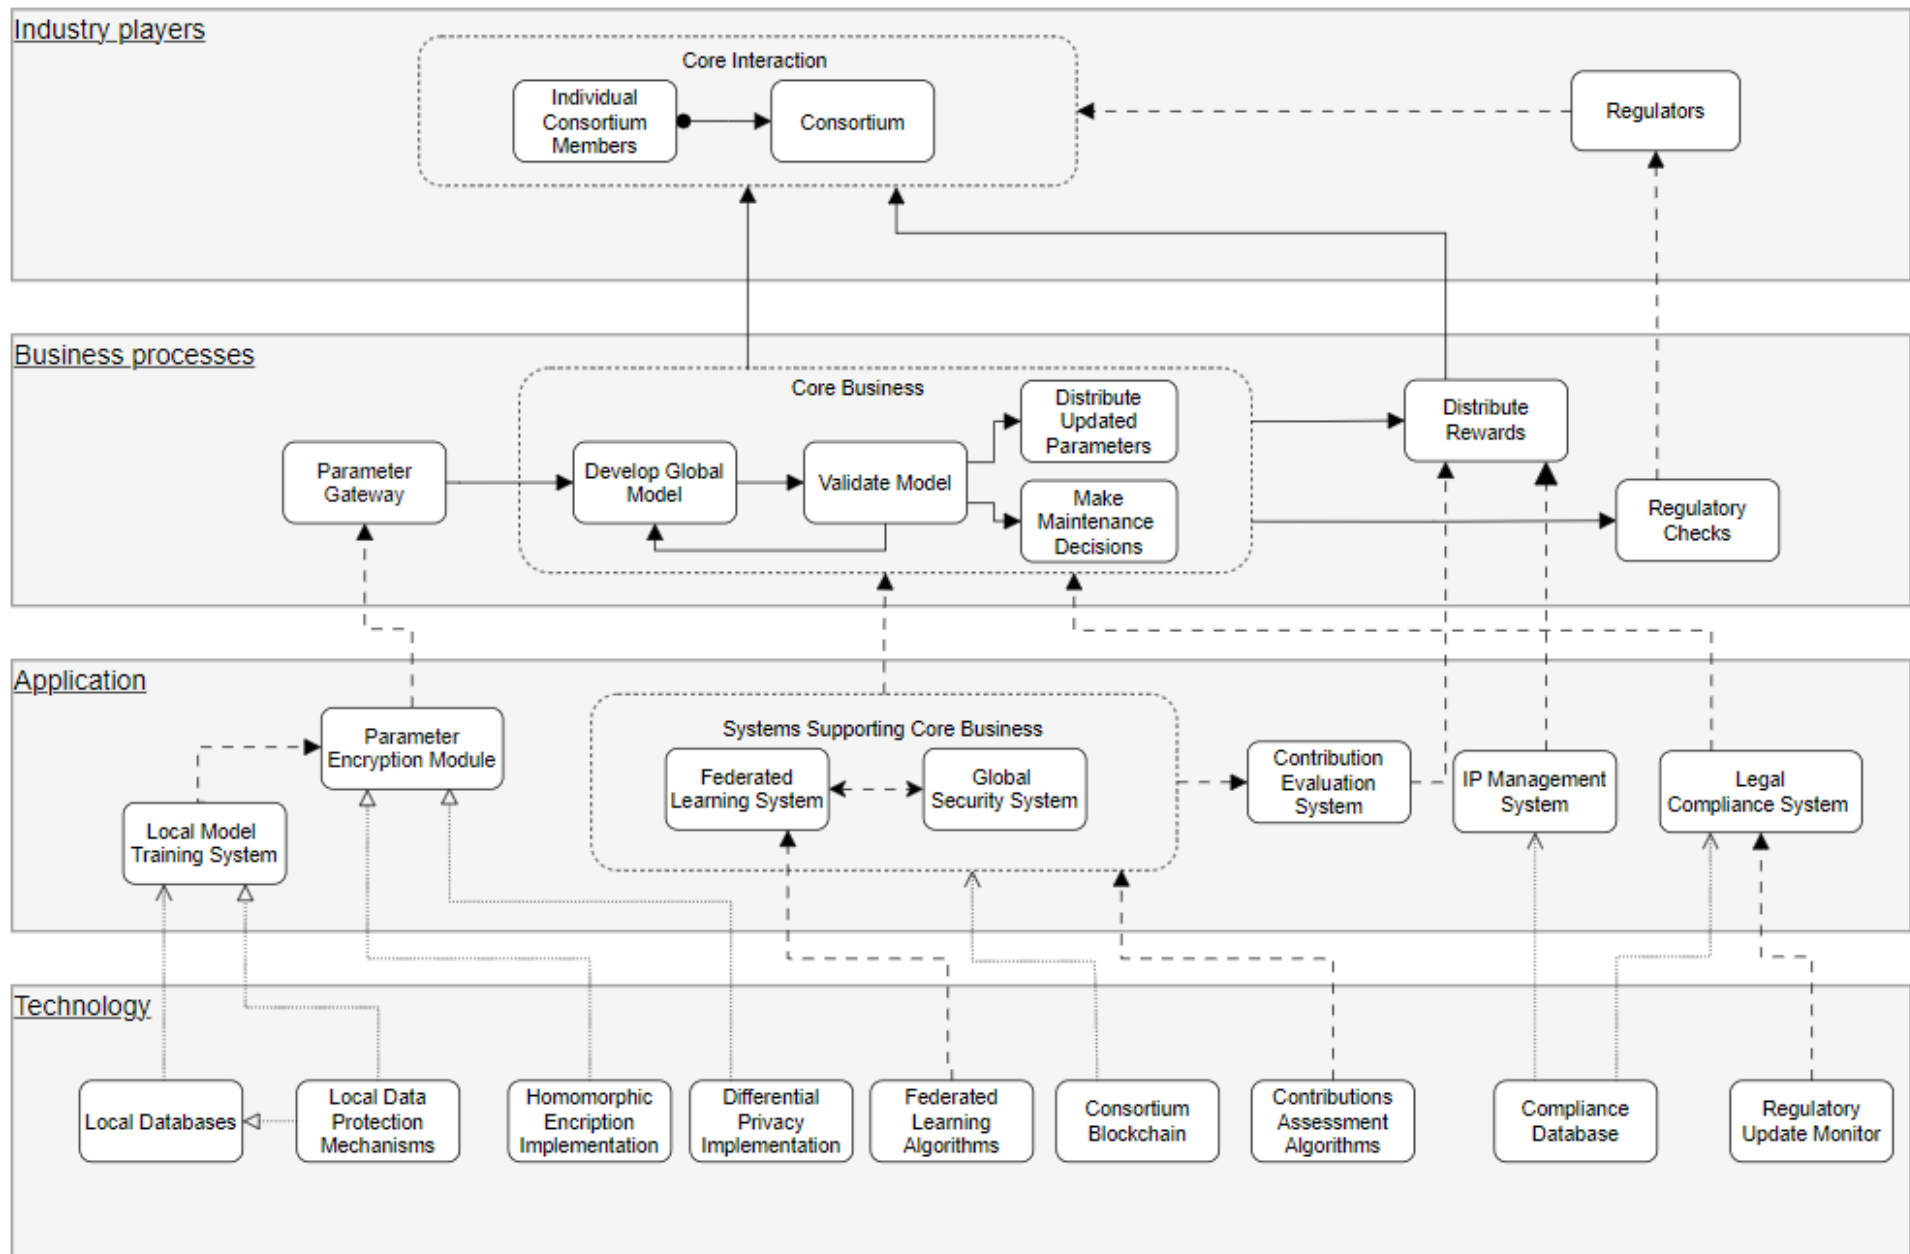


Figure 4: Overview System

5.2 Strict Security Measures

5.2.1 Solution

R1: The system must have strict security measures

This requirement enhances data safety and security by safeguarding sensitive information and parameters during the FL process from both external threats and internal risks, utilizing homomorphic encryption and differential privacy.

The raw data stays within the organization, where they fully control their local security and privacy measures. However, sensitive information regarding the data may still leak from the shared model parameters, even when the local data is not directly visible in FL. Numerous attacks on ML models have demonstrated that it is possible to infer raw data from model parameters (Fredrikson et al., 2015). Therefore, additional data protection strategies are necessary. Homomorphic encryption protects the developed model parameters, while differential privacy focuses on protecting individual data points.

Homomorphic Encryption for Secure Model Aggregation: Homomorphic encryption is a type of encryption that enables operations on ciphertext to produce an encrypted result that, upon decryption, corresponds with the outcome of actions taken on the plaintexts (Tseng et al., 2017). This means that data can be encrypted and shared securely while still allowing it to be used in calculations without exposing the raw data. This method is particularly suitable in the context of FL because it enables the central server to perform model aggregation on encrypted developed model parameters without accessing the raw data. Organizations can encrypt their model updates before sending them to the central server. The server then performs necessary computations on these encrypted updates to generate an improved model, which is sent back to the organizations in its encrypted form. This ensures that the data remains private and secure throughout the entire process, specifically protecting against internal threats as the central server cannot view the input data.

Differential Privacy for Noise Addition: Differential privacy enhances data privacy by adding carefully calibrated noise to the data or model parameters. This process ensures that the presence or absence of a single data point does not significantly affect the outcome of any analysis and provides strong privacy guarantees (Dwork & Roth, 2013). In practical terms, clients can add random noise to their model updates before sending them to the central server. Even if an attacker gains access to the server, they would only see the noisy data, which makes it extremely difficult to infer any specific details about individual data points. For example, differential privacy can be used to add noise to aircraft sensor data before it is shared with a central server. This measure is particularly focused on mitigating external threats, as it prevents attackers who might compromise the server from accessing clear, usable data, thereby safeguarding sensitive operational information. Differential privacy is beneficial in FL because it allows the aggregation of meaningful insights from the data without compromising the privacy of individual users. The added noise ensures that sensitive information about any user's data remains obscured, thus providing a protective layer.

5.2.2 Considerations and Selection Criteria

Other considerations included secure multi-party computation (SMPC), which provides data security by allowing multiple parties to compute a function over their inputs while keeping those inputs private. SMPC ensures that no single party can access the entire dataset, maintaining the privacy of individual data inputs. This technique uses complex cryptographic protocols, like secret sharing, to keep data secure during the computation.

However, SMPC can be very computationally intensive, requiring multiple rounds of communication between parties, leading to increased latency and performance issues, especially in large-scale deployments (Hiwatashi et al., 2021). This makes SMPC less suitable for real-time or large-scale applications due to its higher computational complexity and communication overhead.

Homomorphic encryption was selected for its ability to perform computations on encrypted model parameters, thereby maintaining privacy without sacrificing functionality. Differential privacy was chosen because its privacy guarantees prevent extracting individual data points. These methods, combined, provide a comprehensive security solution that balances efficiency and protection. The choice of these methods is reinforced by repeated recommendations in academic literature, which emphasize their effectiveness and practicality in real-world applications (Domingo-Ferrer et al., 2022).

Implementing homomorphic encryption to secure model parameters and differential privacy to protect individual data points fulfills the purpose of achieving strict security measures. This combination balances security and computational efficiency, ensuring comprehensive protection of sensitive data within the FL system. Therefore, these solutions were selected as the most appropriate for the proposed system.

5.2.3 Connection ArchiMate

The proposed solution for strict security measures involves using homomorphic encryption and differential privacy to protect data and shared parameters. Components contributing to this requirement include Local Databases, Local Data Protection Mechanisms, Homomorphic Encryption Implementation, Differential Privacy Implementation, the Parameter Encryption Module and the Parameter Gateway. These elements work together to ensure data remains secure locally and during transmission. The inclusion of the Global Security System further reinforces the multi-layered security strategy, ensuring comprehensive protection against unauthorized access and breaches.

In Figure 5 the implementation of these security measures is illustrated. The Homomorphic Encryption Implementation and Differential Privacy Implementation are inputs for the FL System and Global Security System, providing additional security on top of using FL. In the Technology layer, Local Data Protection Mechanisms support the Local Database and Local Model Training System, varying per organization to ensure tailored protection. The Global Security System and Local Data Protection Mechanisms are abstract terms used to encompass the variety of security strategies that Individual Consortium Members may employ, such as firewalls, intrusion detection systems and data encryption methods, reflecting the diverse security needs and policies across organizations.

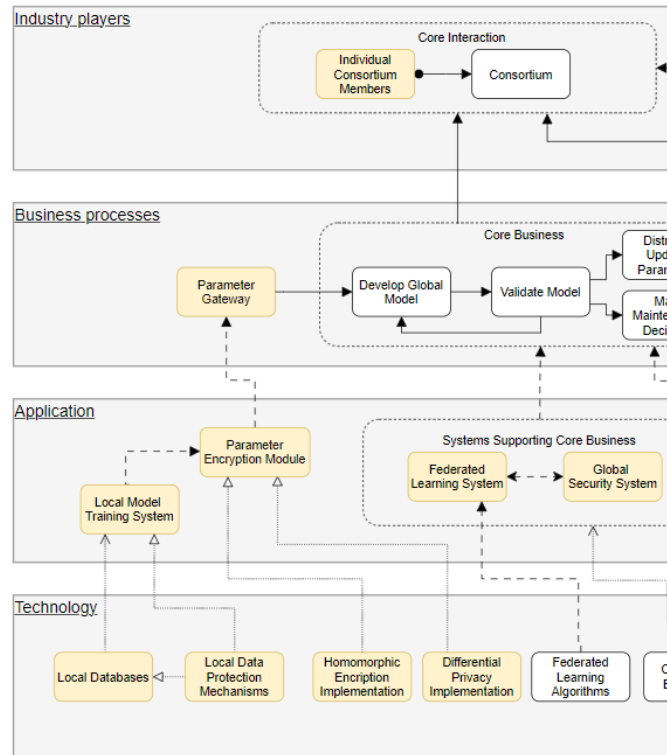


Figure 5: Zoom Security Measures

5.3 Robust Data-Sharing Framework

5.3.1 Solution

R2: The system must have a reliable data-sharing framework.

The purpose of this requirement is to facilitate efficient and secure data sharing among participating entities. The proposed solution utilizes a consortium blockchain to enhance transparency and security in addition to the security provided by an FL system.

FL inherently supports secure data sharing by allowing multiple entities to train an ML model collaboratively without sharing their raw data. Each participant trains the model on their local data and only shares the model updates with the central server. This process ensures that sensitive data remains within the local environment of each participant, significantly reducing the risk of data breaches and ensuring privacy. The central server aggregates the updates from all participants to improve the global model, which is then shared back with the participants. This method maintains data privacy and supports collaborative model development, making organizations feel secure about their data.

A consortium blockchain is integrated into the FL system to enhance security and transparency further. Although the central entity managing the FL process is expected to be trustworthy, implementing a consortium blockchain adds an extra layer of oversight and accountability. Consortium blockchains establish a secure and transparent record of participant data exchanges and model updates. This unalterable, permanent record mitigates the risk of data tampering and leaks. Consortium blockchains allow for a clear and secure record of all participant data exchanges and model updates. This secure record is permanent and cannot be altered, which helps protect against data tampering and reduces the risk of data leaks (Zhang & Lin, 2018). Moreover, consortium blockchains control who can see and use specific data, ensuring that only authorized participants can access sensitive information. This feature helps build trust among all the participants, making consortium blockchain an adequate security and trust-building tool for FL systems (Chen et al., 2020).

In this system, the metadata related to model updates is stored on the consortium blockchain. This ensures all participants can verify which model version is the most up-to-date, ensuring that transparency is maintained. Since blockchain creates an immutable and decentralized record, it prevents any disputes or confusion about model versions and their integrity, providing all participants with confidence in the model's accuracy and timeliness. This further increases trust and transparency among participants.

5.3.2 Consideration and Selection Criteria

Other considerations included the use of public and private blockchains for data sharing. Public blockchains, while highly secure and transparent, can be inefficient and unsuitable for sensitive data due to their open nature. Private blockchains restrict access to a single organization, offering better control and privacy but potentially lacking the collaborative advantages for FL systems. Consortium blockchains strike a balance by allowing multiple trusted entities to participate, providing both transparency and controlled access, making them an ideal choice for secure and efficient data sharing in FL systems (Chen et al., 2020).

Leveraging the combination of FL and a consortium blockchain for secure and transparent data exchanges fulfills the purpose of achieving a reliable data-sharing framework. This approach effectively balances security and transparency, providing comprehensive protection and efficient data sharing while allowing for greater control and oversight of the central entity within the FL system. Therefore, these solutions were selected as the most appropriate for the proposed system.

5.3.3 Connection ArchiMate

The system leverages Federated Learning Algorithms and a Consortium Blockchain to ensure a reliable data-sharing framework. The Federated Learning System, Consortium Blockchain and Parameter Encryption Module work together to enable secure and transparent data sharing. This approach allows for decentralized data processing, enhancing security and maintaining the integrity of shared data across consortium members.

Local Model Training Systems are also needed to support FL's decentralized nature, ensuring that data remains localized while contributing to global model development.

In Figure 6 the implementation of these components is illustrated. The Federated Learning System and Consortium Blockchain are central to the secure data-sharing framework, with the Parameter Encryption Module ensuring that model parameters are shared safely. The Local Model Training Systems contribute to decentralized processing. At the same time, the Global Security System and Local Data Protection Mechanisms provide additional layers of security and privacy tailored to different organizations' diverse needs and policies.

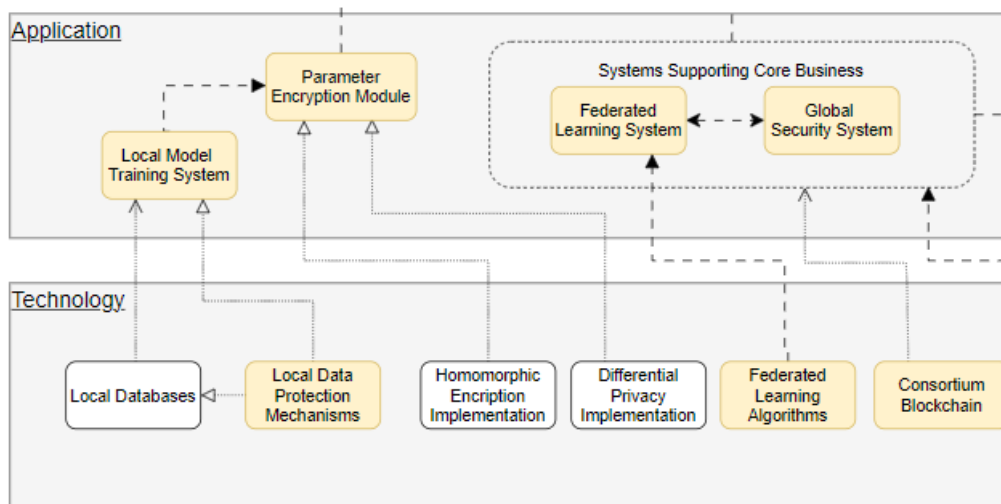


Figure 6: Zoom Data Sharing Framework

5.4 Compliance with Existing Legal Standards

5.4.1 Solution

R3: The system must ensure compliance with existing legal standards.

This requirement ensures adherence to all relevant legal and regulatory standards across various jurisdictions. The proposed solution, explicitly conceptualized for this requirement, involves using a Compliance Database and a Legal Compliance System. The Compliance Database serves as a central repository for all regulatory requirements, ensuring that the system operates within the legal frameworks of various jurisdictions. This database is an input for the Legal Compliance System, which utilizes this information to maintain ongoing compliance. All relevant regulations are included in the database. However, GDPR and the AI Act are highlighted explicitly due to their significant impact on the system's operations and the stringent requirements they impose.

GDPR

The EU enforces far stricter data protection standards than the rest of the world in the form of the GDPR, which focuses on personal data. These restrictions introduce two main concerns: the potential for pilot data to be traceable back to individual performance and the stringent security measures required when sharing data with countries that do not meet the EU's standards.

Pilot data can be traced back to assess pilot performance, posing significant privacy risks. Data transfers within the EU, such as from the Netherlands to France, are straightforward as both countries comply with GDPR (European Data Protection Board, n.d.). However, transfers to the US are complicated because the European Commission has not recognized the US as having sufficient data protection standards. Organizations must use Standard Contractual Clauses or other mechanisms under Article 46 of the GDPR for these transfers (EU, n.d.-a).

To address these issues, the system ensures that any data shared within the FL framework is pre-cleaned by operators, removing personal identifiers. Additionally, sharing only model parameters instead of raw data enhances privacy protection. Parameters do not contain personal data but still allow collaborative model training. As the EU states: *"When the shared parameters are anonymous, Federated Learning facilitates the training of models with data coming from different jurisdictions."* (EU, n.d.-b) This approach mitigates traceability risks and ensures compliance with GDPR, protecting data and reducing privacy violations. Compliance with GDPR mitigates legal risks, ensuring that all data-sharing activities within the system are legally sound.

AI Act

The AI Act is a regulation that ensures the safe and ethical use of AI within the European Union. This regulation categorizes AI systems based on their risk levels, with high-risk systems such as those used in aviation requiring strict regulatory compliance to ensure safety and transparency (European Commission, 2024b). The AI Act mandates several critical requirements for high-risk AI systems, including detailed documentation to trace datasets and algorithms used in development. This ensures that the AI system can be audited and assessed for compliance with safety standards, thereby preventing accidents and ensuring the reliability of maintenance predictions.

By maintaining a comprehensive record of all applicable regulations, the Compliance Database aligns the FL process with legal standards, mitigating the risk of legal infractions. Compliance with the AI Act is critical, as it provides a framework for ensuring AI systems' safety, transparency, and accountability. Non-compliance could lead to severe legal and operational consequences, including fines and operational disruptions. Therefore, adhering to the AI Act ensures legal compliance and enhances the predictive maintenance system's trustworthiness and reliability.

5.4.2 Consideration and Selection Criteria

Other regulations were considered for inclusion in the highlighted regulations due to their potential importance and suggestive names. These include the Digital Services Act (DSA), Digital Markets Act (DMA) and Data Governance Act (DGA). However, they were ultimately not included in the final Compliance Database as primary highlights. To ensure a clear understanding of the regulatory scope, it is needed to clarify why these regulations are not relevant to the system:

- Digital Services Act: Focuses on consumer rights and content moderation, which do not apply to this B2B platform (European Commission, n.d.-b).
- Digital Markets Act: Targets platforms with a significant market presence, defined as those with over 45+ million monthly active users. This threshold does not apply to the predictive maintenance data-sharing platform (European Commission, n.d.-a).
- Data Governance Act: Primarily concerns public sector data reuse and data intermediation services, which are not relevant to the system (European Commission, 2024a).

This requirement ensures adherence to all relevant legal and regulatory standards across various jurisdictions. The proposed solution achieves this adherence by presenting an abstract approach and providing detailed specifics. This abstract design is necessary because the complexity of aviation regulations makes it impractical to include every regulatory detail. Therefore, the solution focuses primarily on a European perspective, offering a working framework that acknowledges global compliance challenges.

5.4.3 Connection ArchiMate

Compliance with existing legal standards is managed through the Compliance Database, Legal Compliance System, Regulatory Checks and Regulators. The Compliance Database acts as the central repository for all regulatory requirements, ensuring the system operates within the legal frameworks of various jurisdictions. The Legal Compliance System utilizes this database to maintain ongoing compliance. Regulatory Checks ensure that the system's activities align with specified legal standards. Finally, the involvement of Regulators ensures the system remains compliant with current regulations and adapts to any changes in legal requirements. These interconnected components work together to facilitate adherence to legal standards and maintain regulatory compliance throughout the system's operations. Figure 7 provides a visual representation of these connections.

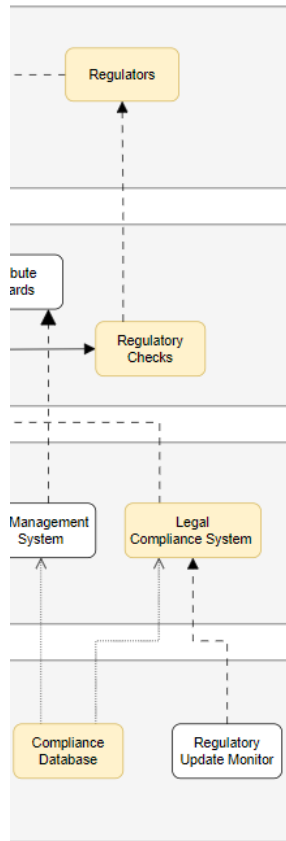


Figure 7: Zoom Regulatory Framework

5.5 Legal Adaptability

5.5.1 Solution

R4: The system must be adaptable to evolving regulations.

This requirement aims to ensure the system can quickly and efficiently adapt to new and changing regulations, maintaining compliance as legal standards evolve. The proposed solution, explicitly conceptualized for this requirement, involves using the Compliance Database and a Regulatory Update Monitor. These tools are integral parts of the Legal Compliance System. The Regulatory Update Monitor monitors regulatory sources for updates, ensuring that any changes in the legal landscape are promptly identified and integrated into the Compliance Database. This database then serves as a reliable resource for maintaining the legal standards of the system's everyday operations.

5.5.2 Consideration and Selection Criteria

The only other approach considered was manual monitoring of regulatory changes. Manual monitoring was deemed impractical due to the high likelihood of human error and the substantial time required to track and implement changes.

The dynamic solution of combining the Compliance Database with the Regulatory Update Monitor was selected for its ability to provide real-time updates and automated monitoring of regulatory changes. This combination ensures continuous compliance and reduces the risk of legal infractions. By integrating these tools within the Legal Compliance System, the proposed solution offers a scalable and responsive framework for monitoring and updating compliance measures. This conceptualized approach acknowledges the abstract nature of the tools but is necessary to address the global nature of aviation regulations effectively. Therefore, this dynamic approach was deemed the most logical and practical solution for the proposed system.

5.5.3 Connection ArchiMate

The Compliance Database and Regulatory Update Monitor are integral parts of the Legal Compliance System, ensuring that the system stays within the boundaries of the law. Regulatory Checks conducted by Regulators verify that the system's activities align with the updated standards stored in the Compliance Database.

These components work together to maintain real-time compliance with evolving regulations. Figure 8 illustrates how the Compliance Database and Regulatory Update Monitor are integrated within the Legal Compliance System and how Regulatory Checks by Regulators ensure ongoing compliance and system integrity.

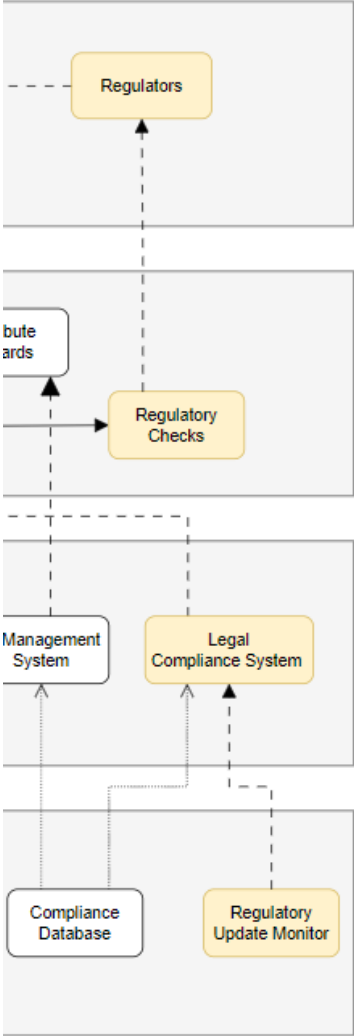


Figure 8: Zoom Legal Adaptability

5.6 Consortium-Based IP Management

5.6.1 Solution

R5: The system must develop a consortium-based approach for managing co-created IP.

This requirement ensures fair access to and avoids monopolization of IP. The proposed solution involves a consortium-based approach, where the IP of the new model is collectively owned by the consortium.

This approach ensures that no single entity can monopolize IP, resulting in a collaborative environment where innovations and improvements benefit all consortium members. Licensing agreements will allow the model to be used by the participating organizations, promoting innovation within the consortium.

Licensing offers several advantages, including setting specific terms for the model's use, controlling its distribution and generating revenue through licensing fees. These agreements can define the scope, duration and financial arrangements for using the model, ensuring that all members have a stake in the collective success of the IP. Additionally, external parties who wish to use the model can access it through a fee-based licensing structure, which regulates access and provides a revenue stream that can be reinvested into the consortium's activities and further development. This structured approach to IP management helps prevent conflicts of interest and ensures a fair distribution of benefits derived from the IP (Yang & Maskus, 2001).

5.6.2 Consideration and Selection Criteria

Another option considered was co-ownership of IP, where multiple companies jointly own the IP. Co-ownership can facilitate collaborative technology development and resource sharing. However, it presents significant challenges due to inconsistent regulations and patent laws across different jurisdictions, leading to potential legal complications and disputes (Gorbatyuk, 2020). Co-ownership also grants participating companies substantial freedom to use the model as they see fit, which complicates management and regulation, potentially resulting in misuse or overuse of the IP (Belderbos et al., 2014).

The consortium-based IP management system addresses these issues by centralizing ownership and control. This model ensures IP management under a unified set of rules agreed upon by all members, facilitating smoother management and reducing legal and operational complexities. Utilizing a licensing approach allows the consortium to set clear usage terms, control distribution and generate revenue through licensing fees. This structured approach mitigates co-ownership risks and provides a reliable framework for managing IP.

The proposed solution centralizes IP ownership to ensure fair access to IP and avoid monopolization by implementing a licensing system. External parties can access the model through a fee-based licensing arrangement, ensuring equitable access and generating additional resources for the consortium.

5.6.3 Connection ArchiMate

Several vital components support the IP handling: IP Management System, Distribute Rewards, Consortium and Core Interaction. The IP Management System ensures centralized control and licensing of the IP, while Distribute Rewards handles the equitable distribution of benefits derived from the IP. The Consortium block represents the participating organizations' collective ownership and collaborative environment. Core Interaction integrates these components, facilitating seamless interaction and management within the consortium. These interconnected elements work together to manage, distribute and ensure fair access to the co-created IP, promoting collaboration and innovation within the consortium. This approach streamlines IP management and enhances trust and cooperation among members, making it an effective solution for managing co-created IP. Figure 9 illustrates how these components are integrated within the overall system architecture.

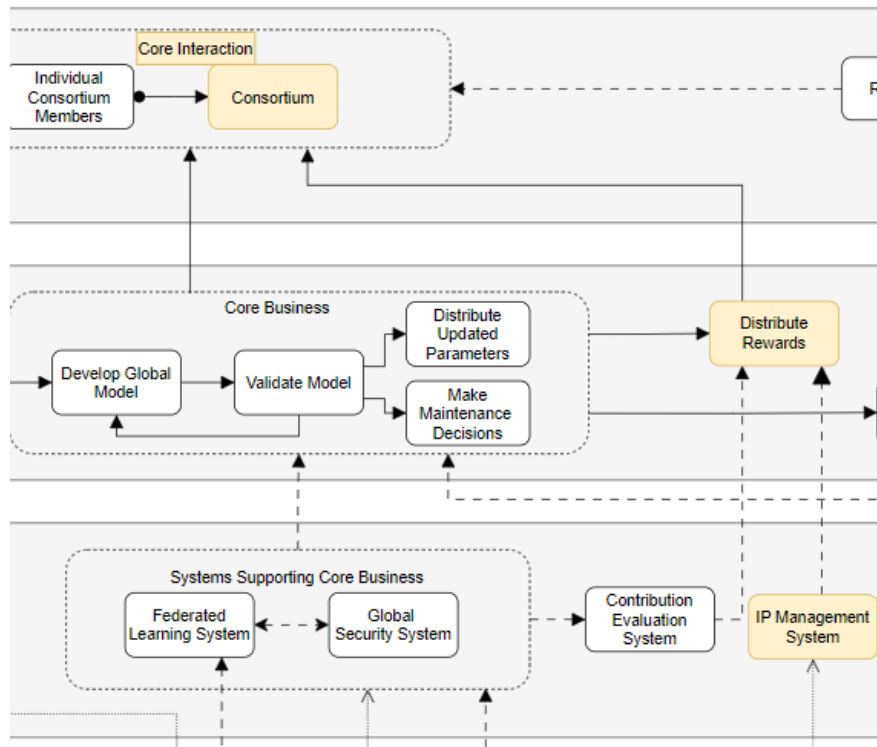


Figure 9: Zoom IP Management

5.7 Equitable Benefit Sharing

5.7.1 Solution

R6: The system must establish mechanisms for equitable benefit sharing.

This requirement is intended to maintain motivation and collaboration among participants. The chosen approach provides different model versions to organizations based on their data quality, quantity and timeliness. This encourages them to share high-quality data and work together effectively. This choice is a fundamental system aspect underpinning participation and data sharing.

Without these incentives, there is a significant risk that participants might withhold valuable data, undermining the system's overall effectiveness and the quality of the predictive models. Commercial competitiveness also plays a role in this decision, as equitable benefit sharing can drive collaboration by providing rewards that justify the investment of time and resources.

To achieve this, the system is designed based on Robust and Fair Federated Learning (RFFL) principles proposed by (Xu & Lyu, 2020). RFFL uses a reputation system to calculate participants' contributions. They receive a specific version of the model based on their reputation score. This iterative process ensures dynamic and ongoing assessment rather than a one-time measurement.

Fairness and robustness

The RFFL framework ensures both high performance and fairness, protecting against behaviors such as free-riding and data poisoning. Free riding refers to benefit from others' contributions without providing valuable data, while data poisoning involves injecting false data to avoid sharing actual data or degrading the model's performance. By dynamically adjusting to the quality and quantity of contributions, the system remains effective and equitable, even if some participants attempt to exploit it.

Key aspects that need to be considered are the Distribution of Model Versions, the Purpose of Model Differentiation, Initial Consortium Formation, New Members Joining the Consortium and Companies Interested in Purchasing the Model Without Joining.

Distribution of Model Versions

Best model availability: The 'best' version of the model incorporates updates from all participants, weighted by their contribution through reputation scores. This global model represents the most comprehensive and accurate state of learned patterns from collective data.

Access to the best model: High-contributing participants, identified by higher reputation scores, receive versions of the model closer to this global model. Their significant contributions justify access to the most accurate and comprehensive version, promoting ongoing high-level contributions.

Differentiation of model access: Participants with lower contributions receive models less representative of the global update. These models incorporate fewer parameters or adjustments from the latest global learning round. All models must meet the safety standards required by regulatory entities. The RFFL system's protocols for quantifying and integrating contributions determine the degree of differentiation.

Purpose of Model Differentiation

Incentivization: Differentiation in model access motivates participants to strive for higher contributions by linking the quality of the model they receive to their involvement and data quality. This encourages continuous improvement and active participation.

Fairness: The system ensures fairness by preventing participants from disproportionately benefiting from others' contributions without providing equivalent contributions. This helps prevent free-riding and maintains a balanced contribution-reward dynamic.

Security and robustness: Controlling model access based on contributions mitigates risks associated with data poisoning and other security threats. High-reputation participants are less likely to introduce malicious data, ensuring accurate models aligned with security best practices.

Company Size

Small companies: Small companies have less data to offer, resulting in an older model version. However, they would still benefit from it since they are getting a better model than they could have developed.

Large companies: The larger companies have more data to offer, resulting in newer model versions. It is also interesting for the bigger companies to join because this setup does not endanger their competitive position.

Initial Consortium Formation

Mechanism: The initial consortium forms with companies contributing data to train a predictive model collaboratively. Each organization gains access to the global model, with versions based on their reputation scores. This ensures an inclusive yet competitive environment where contributions are valued and rewarded.

Fairness: Active contributors are incentivized to maintain or increase their data contributions by receiving better, more accurate model versions. The reputation score is a transparent measure of their engagement and data quality, promoting fairness.

New Members Joining the Consortium

Integration: New companies can integrate into the consortium by contributing data and building their reputation scores. They gain model access according to their contribution level, aligning incentives for ongoing data sharing and collaboration.

Concerns: Late entrants might feel disadvantaged if earlier contributions are weighted more heavily. The reputation scoring mechanism adjusts to reflect evolving data quality and quantity over time, ensuring fairness for all participants.

Collaboration incentives: New members are incentivized to contribute high-quality data to build their reputation scores and gain better model access. The dynamic adjustment of reputation scores establishes a collaborative and inclusive environment.

Companies Interested in Purchasing the Model Without Joining

Feasibility: Companies wanting to access the predictive maintenance model without contributing data could purchase it. However, they would only get access to the current version without ongoing updates unless they joined the consortium and contributed data, maintaining the incentive for data sharing.

Ownership and stocks: The developing organization holds stock in the current model, determining each contributor's share of the benefits. This stock correlates with their reputation score, ensuring proportional benefit sharing.

Risk: Companies that wait to purchase a well-developed model rather than contribute to its development may de-incentivize participation. Agreements between consortium parties can define access rules for available models, mitigating this risk.

5.7.2 Consideration and Selection Criteria

One considered alternative was providing the same model to all participants. However, this approach would undermine the system's effectiveness due to a lack of incentivization, leading to several negative consequences such as:

Minimal Contributions: Without differentiation in model versions, organizations are little motivated to contribute high-quality data. This lack of incentive results in minimal contributions, diminishing overall data quality and hampering collaboration efforts.

Commercial Interests: In a competitive environment, organizations need tangible benefits to justify participation and data sharing. Equitable benefit sharing through differentiated model versions aligns with commercial interests by offering rewards proportional to contributions. This encourages organizations to invest time and resources, creating a collaborative environment while allowing participants to gain competitive advantages based on their input.

Security and Robustness: Providing the same model to all participants increases the risk of data poisoning and other security threats. Controlling and monitoring data integrity is challenging without differentiating based on contributions. Differentiated model access ensures that high-reputation participants, less likely to introduce malicious data, receive more accurate models. This alignment with security best practices enhances system robustness.

Fairness: Uniform model distribution could lead to participants disproportionately benefiting from others' contributions without providing equivalent input. This lack of fairness encourages free-riding behavior, where participants exploit the system by benefiting from others' data without offering valuable data themselves. Differentiation maintains a fair and balanced contribution-reward dynamic, ensuring benefits align with the level of input and discouraging free-riding.

This requirement aims to maintain motivation and collaboration among participants. The RFFL framework ensures high performance and fairness through dynamic adjustment based on contributions. By offering differentiated model access, the system incentivizes high-quality contributions, aligns with commercial interests and maintains security and fairness, making RFFL the chosen solution for equitable benefit sharing.

5.7.3 Connection ArchiMate

The equitable benefit-sharing mechanism is a core function of the system architecture, necessitating the involvement of multiple components to ensure its effectiveness. The key components include Federated Learning Algorithms, Contributions Assessment Algorithms, a Federated Learning System, Systems Supporting Core Business, a Contribution Evaluation System, a Global Model, Core Business, Distributing Updated Parameters, Distributing Rewards, Consortium and Core Interaction.

These components collaborate to develop a global model, evaluate contributions and distribute rewards based on participant contributions. For instance, the Federated Learning Algorithms support the Federated Learning System by providing the necessary computations for model training. The Contributions Assessment Algorithms feed data into the Contribution Evaluation System, which assesses each participant's contributions. The Develop Global Model component is used to create a comprehensive model. Finally, the Distribute Updated Parameters and Distribute Rewards components ensure that participants receive model updates and rewards proportionate to their contributions. Figure 10 illustrates how these components are integrated within the system architecture.

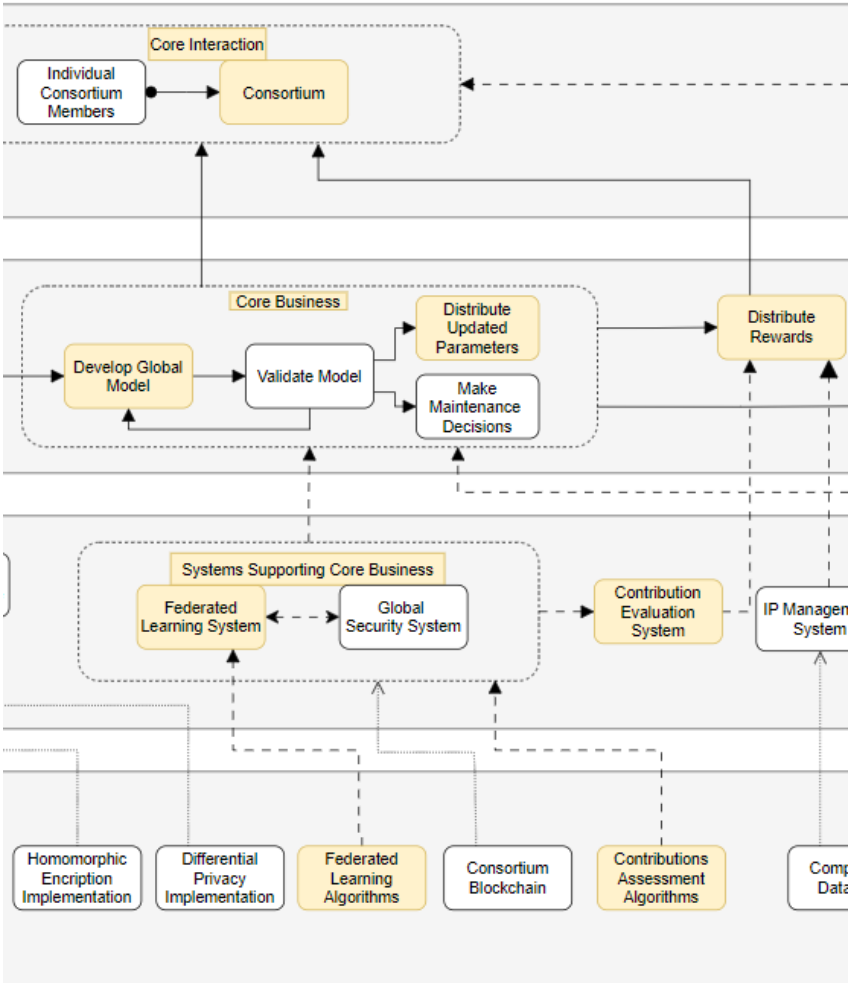


Figure 10: Zoom Equitable Benefit Sharing

5.8 Transparent Value Attribution

5.8.1 Solution

R7: The system must develop transparent processes for proportional value attribution.

This requirement aims to build trust and engagement among participants. The chosen solution calculates contribution scores using the RFFL algorithm. These scores are based on the quality, quantity, timely delivery of data and potentially other aspects. By using a calculated score, the system ensures proportional value attribution.

A vital feature of this system is the transparency of the scoring process. Participants can see how their scores are calculated. The RFFL algorithm's calculations are recorded on the consortium blockchain, providing an immutable and transparent record of each participant's contribution. This transparency lets participants understand precisely how their contributions translate into their scores and subsequent rewards, enhancing trust in the system. Participants can also view the calculations for others, showing a sense of fairness and accountability.

Another significant benefit of the contribution score system is its ability to incentivize timely and high-quality data contributions. Knowing that their efforts will be fairly assessed and rewarded motivates participants to provide better data consistently. This continuous incentive structure ensures that the overall quality of data within the consortium remains high, which is crucial for developing accurate predictive models. This method quantitatively assesses each participant's input and allocates rewards based on their contributions' quality, quantity, and timing.

The consortium blockchain's role in recording and verifying each participant's contribution score ensures that all assessments are transparent and auditable by all consortium members. This helps prevent disputes over contribution evaluations by providing an immutable, shared record. Blockchain's decentralized nature ensures that no single participant can manipulate the scoring process, further enhancing fairness.

5.8.2 Consideration and Selection Criteria

One considered alternative for value attribution was a peer review-based assessment system. In this approach, participants would evaluate each other's contributions. While this method promotes accountability and mutual oversight, it was deemed less practical for several reasons. Firstly, peer review can introduce subjectivity and potential biases, as participants might not always provide impartial evaluations. This could lead to disputes and reduce trust in the system. Secondly, the peer review process can be time-consuming and complex, reducing the system's efficiency and scalability.

This requirement aims to build trust and engagement among participants. The RFFL contribution score, enhanced by the transparency of a consortium blockchain, provides the necessary mechanisms to fulfill this purpose. This approach ensures fairness and transparency, making it the chosen solution for value attribution within the system.

5.8.3 Connection ArchiMate

The Federated Learning Algorithms provide the necessary computations for model training within the Federated Learning System. Contributions Assessment Algorithms feed data into the Contribution Evaluation System, which assesses each participant's contributions. The Consortium Blockchain records these calculations, ensuring transparency and immutability.

The individual consortium members interact with these components, seeing how their contributions are calculated and rewarded. This establishes trust and accountability within the system. The overarching Consortium framework integrates these elements, facilitating seamless collaboration and value attribution. Figure 11 illustrates how these components are integrated within the system architecture.

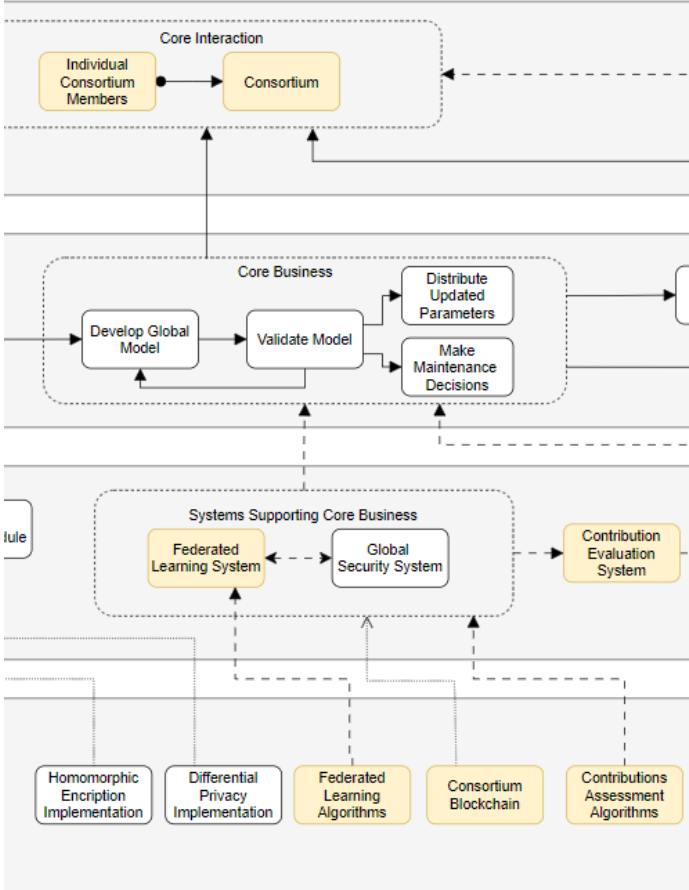


Figure 11: Zoom Transparent Value Attribution

5.9 Accuracy of Predictions

5.9.1 Solution

R8: The system must be more accurate than current predictions.

This requirement ensures that the system surpasses the accuracy of existing predictive models. Without achieving superior accuracy, the system would not offer a compelling advantage, making it less attractive for organizations to join. Higher accuracy is critical to show the system's value and meet consortium members' operational and financial needs.

The proposed solution integrates FL to use different datasets to achieve this, improving predictive maintenance tools and reducing unexpected maintenance issues. The key benefits of these more accurate predictions are presented on the next page and include:

- **Reduction of Unplanned Maintenance:** By utilizing predictive models that leverage the shared data of consortium members, the system can better predict the status of aircraft components. Maintenance teams can use these insights to anticipate which components will likely fail and schedule preemptive maintenance activities. This proactive approach reduces the likelihood of unexpected breakdowns, decreasing costly downtime and minimizing the impact of sudden component failures on operations.
- **Optimized Maintenance Planning:** Predictive analytics will initially enhance maintenance planning by providing more accurate assessments of component status. This adjusts maintenance schedules based on actual needs rather than fixed intervals, reducing downtime and improving resource utilization. Maintenance can be scheduled immediately to prevent failure if a component shows signs of degradation sooner than expected. Conversely, if a component remains in good condition beyond its expected lifespan, its replacement can be delayed, thus optimizing resource utilization. This dynamic and needs-based scheduling reduces unnecessary downtime and efficiently uses maintenance resources.
- **Inventory Cost Reduction:** A significant financial advantage for participants in the FL system is the reduction of inventory costs. By leveraging shared data to improve predictive accuracy, stakeholders can maintain a smaller spare parts inventory. Accurate predictive models allow participants to minimize excess stock, reducing capital investment in unnecessary spare parts. Instead of maintaining a large inventory due to uncertainties in demand, participants can predict when and which components will need replacement. This streamlined inventory approach reduces storage costs and potential losses from unused or outdated spare parts.

5.9.2 Consideration and Selection Criteria

Achieving superior predictive capabilities is a mandatory requirement for this system. The proposed solution, utilizing FL to integrate diverse datasets, is designed to produce the best possible model. High predictive accuracy makes the system attractive and beneficial to consortium members, justifying the effort and resources invested. Without these enhanced capabilities, the system would fail to provide the compelling advantage necessary for widespread adoption and collaboration. Given these factors, achieving this requirement is fundamental and cannot be compromised. The current design aims to meet this requirement effectively.

5.9.3 Connection ArchiMate

Achieving superior predictive capabilities requires a cohesive integration of various interconnected components. Local Databases within each consortium member's infrastructure store data, which Local Data Protection Mechanisms protect. Federated Learning Algorithms operate on this distributed data within the Local Model Training System, enabling local model training while maintaining data privacy. These models share parameters securely via the Parameter Encryption Module and the Parameter Gateway, which feeds into the Develop Global Model. This global model is validated for accuracy and reliability before making maintenance Decisions. This process is integrated into the Core Business operations. Together, these components create a system that enhances predictive maintenance capabilities. The involved components can be seen in Figure 12.

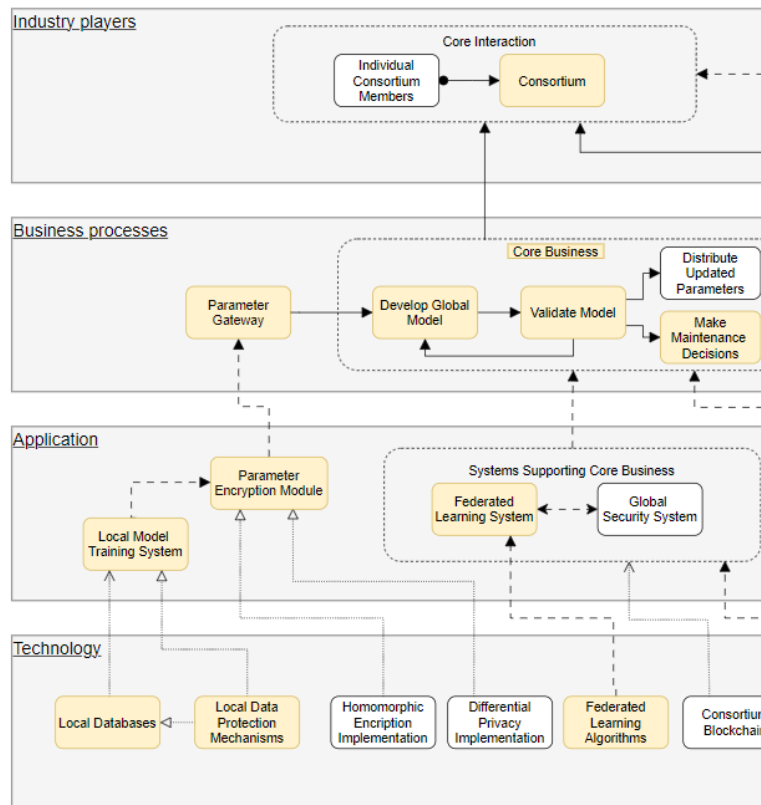


Figure 12: Zoom Accuracy of Predictions

5.10 Explainability and Transparency of AI Models

5.10.1 Solution

R9: The system must ensure that AI models are explainable and transparent.

This requirement aims to gain regulatory approval and industry trust by making AI models transparent and explainable. This is achieved through model transparency and Shapley values to explain model decisions.

The FAA and EASA share similar requirements outlined in the National Aviation Research Plan and the AI Act. For this discussion, the focus will be on EASA regulations (EASA, 2023; FAA, 2024). From the interview with the regulators, it is known that the EASA Concept Paper is a delivery of the exploration phase and is utilized as an initial set of guidelines for projects. It has been made available for public consultation and is currently used to guide the initial AI-related projects. This paper is used to make the system future-proof.

To clarify, FL is a form of ML, and they both fall under the umbrella of AI. EASA regulates the integration of AI and ML into aviation to ensure these technologies are safe, reliable and effective. FL alters traditional compliance with these regulations, impacting crucial areas such as explainability, learning assurance and security. The classification of the AI/ML model determines the applicable regulations, and this FL system can be classified into the lower categories since human intervention can override it, serving as a tool for decision-making (EASA, 2023).

Explainability and Transparency in AI Models Explainability and transparency are fundamental for AI/ML systems used in aviation. These attributes ensure that human operators understand how AI systems make the necessary decisions to maintain trust and facilitate necessary interventions.

Explainability in both federated and traditional learning models can be categorized into global explainability and operational explainability.

- **Global Explainability:** This comprehensive view is intended for those overseeing operations, such as engineers and regulatory compliance officers. It includes a model Overview, Algorithm Details and Performance Indicators.
- **Operational Explainability:** This focuses on specific decisions made by the model, which is needed for personnel like maintenance crews. It involves Data Points Considered, Threshold Comparisons and Historical Comparisons.

FL Impact on Explainability

Despite the distributed nature of data handling in FL, the resulting model can be interpreted similarly to models trained via conventional methods. However, this requires extra caution due to the complexities involved. The final model integrates learned features from various organizations, which can make it more challenging to trace back decisions to specific data points. This complexity requires thorough documentation and transparency throughout the model training process. Operators must be particularly thorough in understanding how features are combined and weighted. Ensuring this interpretability ensures trust and allows human operators to understand and validate the model's decisions.

Use of Shapley Values

Shapley values provide a method to explain individual predictions by assigning importance to each parameter, showing how much a specific parameter influences the outcome. This method ensures transparency and allows participants to understand how their data impacts the model's decisions. Shapley values are calculated by considering all possible parameter combinations and determining each parameter's contribution to the prediction. This approach provides a way to measure the impact of each data point on the model's decisions. The integration of Shapley values adds an extra layer of explainability, making it easier to understand complex AI models and their decision-making processes (Nohara et al., 2022).

Learning Assurance

Learning assurance in AI/ML systems demands that the data used for training these systems is comprehensive, representative and free from biases. This ensures that the systems can perform reliably across all expected operational scenarios without errors that could lead to unsafe outcomes. FL can improve learning assurance by using a broader range of data from various participating organizations without centralizing the data. This diversity helps build a more advanced model by naturally incorporating various scenarios and operational data, potentially reducing bias more effectively than single-dataset models.

Security and Data Protection

Security and data protection are important due to the vulnerabilities of centralized AI/ML systems to cyber threats. EASA enforces strict regulations for these centralized systems. FL allows each organization to retain control over its data, avoiding exposure to a central repository that could be a target for cyber-attacks. This approach enhances data security by reducing the risk of large-scale breaches. Each organization in the federated network focuses on local data security, potentially increasing the system's overall resilience against external threats.

5.10.2 Consideration and Selection Criteria

EASA has acknowledged the difficulties associated with ensuring the explainability of the AI models in its AI Act concept paper. The paper highlights the lack of clear guidance and standardized methods for achieving transparency in AI systems. This regulatory uncertainty challenges organizations seeking to implement FL while following explainability requirements. EASA's ongoing efforts to develop more concrete guidelines underscore the complexity of the issue and the need for clearer solutions that can adapt to evolving standards.

Despite these challenges, Shapley values offer a promising approach to enhancing the explainability of AI models within the FL framework. Shapley values provide a method to explain individual predictions. This method helps to demystify the uncertain nature of FL by providing insights into how different data inputs contribute to the final model predictions. While Shapley values alone may not fully resolve the complexities of regulatory compliance, they represent a significant step toward greater transparency and accountability in AI models used in aviation. By enhancing explainability and transparency, Shapley values contribute to aligning AI models with evolving regulatory standards, fulfilling the requirement for regulatory compliance and industry trust.

5.10.3 Connection ArchiMate

Ensuring the explainability and transparency of AI models involves using many different system parts. These parts include those needed to build the model, keep data secure and meet regulatory standards. The system must handle data locally, distribute parameters securely, carry out FL processes and continuously update and check regulations. Additionally, it includes steps to validate the model and ensure it complies with legal standards. Figure 13 shows how these parts work together to achieve transparency and explainability in the AI system.

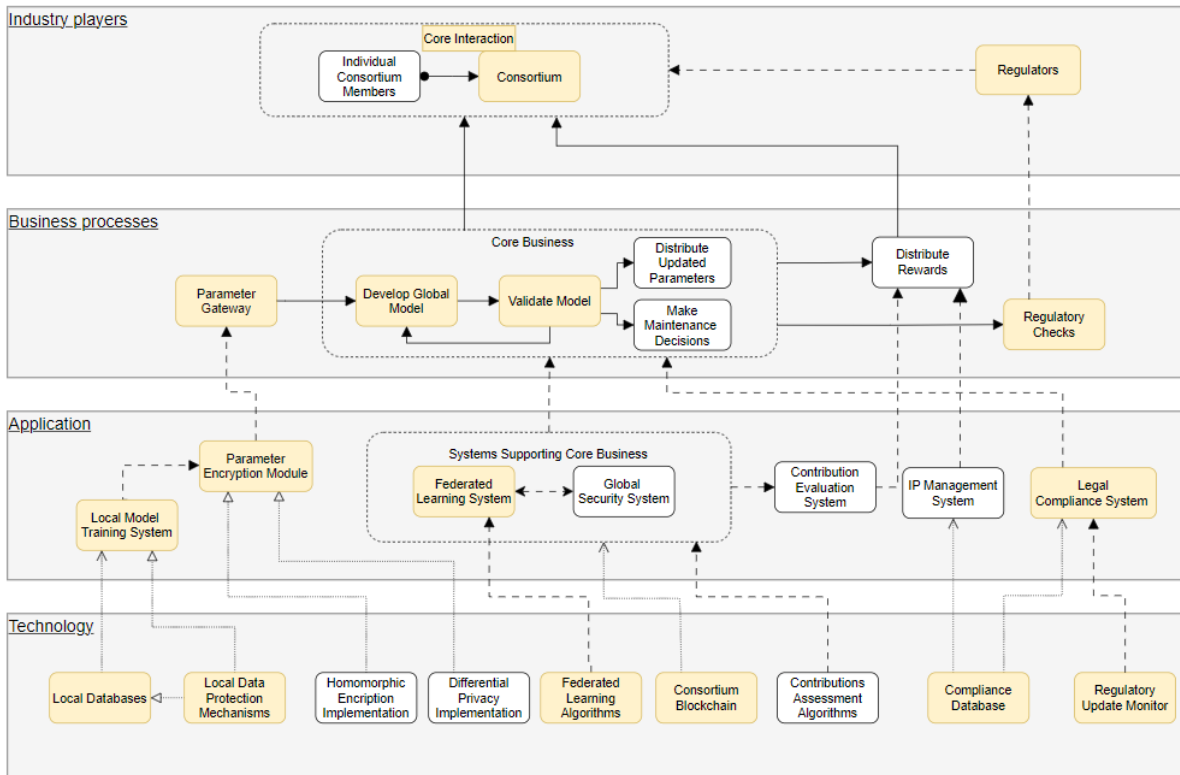


Figure 13: Highlight Model Explainability

5.11 Training Programs

5.11.1 Solution

R10: The system must have training programs for stakeholders.

The purpose of this requirement is to enhance stakeholder knowledge and engagement. FL is a novel concept with significant potential, but stakeholders are unaware of its full capabilities and benefits. Ensuring that all participants know the benefits of FL is needed to maximize its potential and ensure effective collaboration within the consortium. These programs are designed to cover a comprehensive range of topics, including:

- The principles and benefits of FL
- Practical applications and case studies
- Regulatory requirements and compliance

Training will be delivered through a combination of hands-on sessions and workshops. These face-to-face meetings will have an interactive component, which enhances learning. Within the consortium, numerous individuals are experts on specific aspects. These experts will regularly take turns providing training. This approach ensures that all stakeholders receive a consistent and thorough education, maintaining a high standard of knowledge within the consortium.

5.11.2 Consideration and Selection Criteria

One alternative considered was providing self-paced e-learning modules that stakeholders could complete at their convenience. This approach offered flexibility, accommodating different time zones and schedules. However, there is a significant risk that participants might not perceive the urgency of completing these modules. This inconsistent engagement could lead to varying levels of understanding among consortium members.

This requirement aims to enhance stakeholder knowledge and engagement. Therefore, face-to-face training programs were deemed the most effective solution. They ensure that all stakeholders are adequately informed and actively engaged.

5.11.3 Connection ArchiMate

The components involved in achieving these training programs are the Systems Supporting Core Business, Individual Consortium Members and the Consortium. The Consortium and Individual Consortium Members play central roles in these programs, supported by the Systems Supporting Core Business. These components work together to create an educational framework with many participants. Figure 14 shows how these blocks are present in the system.

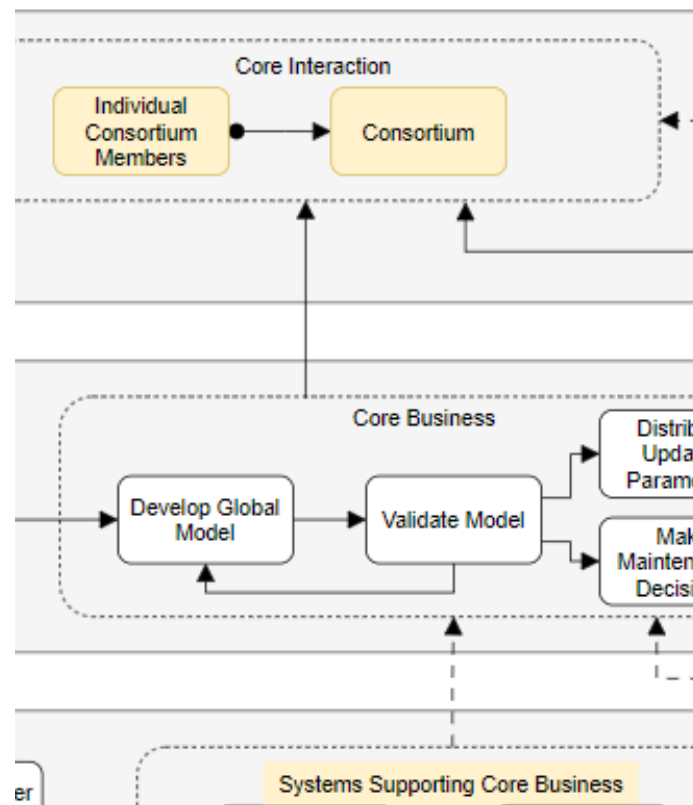


Figure 14: Zoom Training Programs

5.12 Appointment of a Trusted Neutral Entity

5.12.1 Solution

R11: The system must appoint a trusted neutral entity.

This requirement ensures continuous monitoring of fair governance and generates participant trust, which is required to successfully implement and operate the FL system. A consortium has also been chosen for this entity. This consortium will oversee the activities, manage compliance and resolve disputes, trying to maintain the integrity and impartiality of the FL process.

The consortium will perform several vital functions:

- **Dispute Resolution:** Mediating conflicts between participants and providing unbiased resolutions to maintain harmony within the consortium.
- **Transparency and Trust Building:** Enhancing transparency by ensuring that all processes are documented and accessible to participants, thus building trust and encouraging active engagement.

5.12.2 Consideration and Selection Criteria

One alternative considered was appointing a single commercial entity to manage the FL system. While this might streamline management processes, it could not provide the same level of safety and trust. Other participants will be hesitant to share their data due to concerns about impartiality and potential conflicts of interest. Another option was to let other neutral entities, such as regulators, lead the system. Although regulators are trusted and impartial, managing such a system is not aligned with their main activities.

This requirement ensures fair and unbiased governance and establishes trust among participants. Therefore, appointing a trusted, neutral entity as a consortium was deemed the most effective solution. This approach ensures impartial oversight, consistent rule application and a trustworthy environment for all participating members.

5.12.3 Connection ArchiMate

The appointment of a trusted, neutral entity is supported by three key blocks: Consortium, Core Interaction and Consortium Blockchain. The Consortium block embodies the collective governance structure, providing a framework for collaboration and decision-making. The Core Interaction block facilitates communication and coordination among consortium members, ensuring seamless and transparent interactions. The Consortium Blockchain ensures transparency and immutability of transactions and data sharing, reinforcing participant trust. These blocks collectively enable the trusted, neutral entity to maintain integrity, fairness and trust within the system. Figure 15 illustrates how these blocks are integrated within the system architecture.

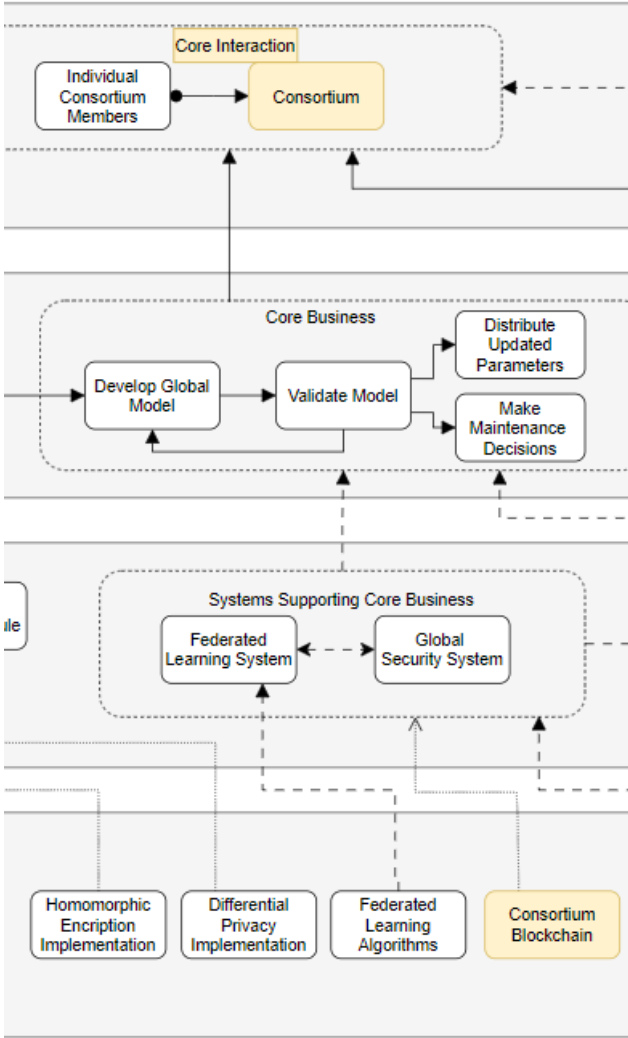


Figure 15: Zoom Trusted Neutral Entity

5.13 Conclusion

This chapter explored the architectural features needed to successfully implement and operate the FL system within a consortium of diverse organizations. The proposed architectural solutions effectively meet the various requirements identified in the research.

5.13.1 Trust-Building Elements

Trust-building is a central theme throughout the system's design. Ensuring trust among participants is fundamental for the success of the FL system. This is mainly due to their diverse commercial interests. Therefore, each design choice enhances trust and security to create a cooperative and secure atmosphere.

Sharing data through homomorphic encryption allows computations on encrypted data, ensuring sensitive parameters remain secure during processing. This guarantees data privacy and security at all stages, making participants more comfortable sharing their data. Differential privacy adds noise to data, ensuring individual data points cannot be identified. This method protects user privacy and ensures personal data cannot be derived from shared model updates, further enhancing trust.

The consortium blockchain adds transparency and immutability to data-sharing and governance processes. It records all transactions and updates in an unalterable database, allowing participants to verify the integrity of the data. This transparency ensures that all actions are documented and traceable. The blockchain facilitates accountability, as any deviations from agreed-upon protocols can be easily identified and addressed. The decentralized nature of FL minimizes the risks of data breaches and misuse, thus contributing to building trust. The decentralized approach also lets participants control and own their data.

Every system design aspect, from data encryption to transparent governance, aims to maximize trust and security. This comprehensive approach ensures that all participants feel secure and confident in their collaboration, knowing their data is protected, and their contributions are valued.

5.13.2 Role of the Consortium

The consortium plays an important role in the governance and operational success of the FL system. The consortium acts as the central coordinating body that aligns the interests and activities of its members. This collective governance structure ensures that all participating organizations agree to standards, resulting in a unified approach.

The consortium is also instrumental in maintaining the integrity and fairness of the FL process. By acting as a neutral entity, the consortium helps prevent conflicts of interest and ensures that all members are treated fairly. It also guarantees that no single entity can dominate the decision-making process. The consortium's neutrality and commitment to fairness create an inclusive and equitable environment where all participants feel valued and respected.

By managing the structured licensing system for IP, they ensure that all members benefit from co-created innovations. This approach prevents the monopolization of IP and encourages innovation and active participation from all members.

5.13.3 Overall Impact

This chapter demonstrates how a balanced architectural design can effectively meet the requirements. The architectural design outlined significantly contributes to the system's success. The architecture establishes a secure and effective FL system by incorporating trust-building elements. This approach addresses technical requirements and enhances them with the consortium's relational aspects.

Balancing centralized and decentralized elements within the system reveals a delicate balance. The centralized elements, such as the consortium's governance and the blockchain for transparency, provide a cohesive framework. This framework ensures consistency and trust. These trust-building components facilitate unified decision-making and apparent oversight needed to maintain order and accountability. On the other hand, decentralized responsibilities, like individual data security measures, allow flexibility and adaptability to specific organizational needs. This balance is needed to address the diverse requirements of the different organizations. It combines the strengths of both centralized and decentralized approaches.

The general security aspects managed by the consortium and the individual security aspects by the individual members provide a strong foundation for the FL system. The strategic incorporation of trust-building elements and the separation of responsibilities ensures security and promotes a collaborative and trustful environment.

The differences in abstraction in the model reflect the level of detail possible for each part of the system. Where detailed implementations are feasible, the model is less abstract, showing specific elements and their relationships. However, for broader or more complex components, the model adopts a more abstract approach. Given the wide scope of this research and the involvement of multiple global organizations, varying levels of abstraction are unavoidable. This approach ensures the model remains both comprehensive and adaptable, focusing on key interactions.

6 Validation

6.1 Structure of the Validation

The validation chapter consists of the following parts: expert session, individual interviews and improvements to the system. One of the main design choices was how to distribute the benefits, so the validation session focused on that. This extra focus led to a group session trying to reach a mutual consensus on the best way to distribute the benefits. The design choices were further discussed through individual interviews with participants, ensuring a comprehensive evaluation of the most critical design aspects. Lastly, this new information was evaluated and analyzed to see how the system could be improved.

6.2 Expert Validation Session

The expert validation session thoroughly examined the distribution of benefits. This fundamental and critical design choice needed a dedicated session. The session focused on determining whether distributing different versions of models would be a suitable option and discussed fair incentives among different contributors. The expert session included Aviation Expert 1, Aviation Expert 2, MRO 2 and Operator 1. This selection was logistically feasible as all were in the Netherlands then, and a group session in person seemed more productive. Professor Tuler de Oliveira, the guiding professor, was also present at the discussion but did not participate.

Circular reasoning could be risky when using the same set of interviewees for validation, but this was not the case in this scenario. The interviewees had different views on various aspects of the model, and designing the model required consensus on certain aspects. Therefore, the validation process was essential to align the design with these varied perspectives.

6.2.1 Initial Discussions

Evaluating the Distribution of Models

The session began by focusing on distributing different model versions, where fairness and contribution assessment were central themes. Participant 1 Aviation Expert, emphasized the importance of data quality: *"Or incidents present in the data ... If I flip a coin and it lands on the side, instead of the top or bottom, that is important data. That should be valued."* This highlights an inherent challenge: ensuring high-quality data contributions without incentivizing the mere volume of data. Valuing unique and impactful data points rather than sheer quantity is key for model development.

Participant 2 Aviation Expert initially supported a tiered model distribution system, reflecting an industry practice where better contributors receive more advanced benefits. He affirmed: *"Absolutely, that is what this guy has been saying forever. (points to Participant 1 Aviation Expert)."* This approach encourages more significant contributions by offering better rewards but could also lead to some participants receiving less advanced models. This creates a challenge in balancing the incentive to contribute with the need to ensure fair model access.

Emerging Concerns

As discussions progressed, concerns about the practicality and fairness of a multi-model approach became evident. Participant 1 Operator pointed out the commercial implications of not contributing data, suggesting that non-contributors could still benefit from the model. He stated his concerns: *"If you are not contributing data, you could still make use of it as a commercial model."* This argument highlights the potential risk of free-riding, where entities benefit from others' contributions without offering their data, thereby undermining the collaborative effort.

To help the discussion, Professor Tuler de Oliveira repeated a solution previously mentioned in the introduction by the researcher: a unified model with differentiated access based on contributions. She clarified one of the options: *"A unified model that could still differentiate access based on contribution, where everyone who contributes should get stocks linked to their input level, maintaining a scoring and reputation system."* This model balances fairness and incentivization by allowing all contributors access to the model while varying costs based on contribution levels. This approach addresses concerns of unequal access while maintaining incentives for higher contributions, aligning with principles of equitable collaboration and shared benefits.

Participant 2 Aviation Expert thought about it and changed his opinion towards a single model, highlighting a significant shift driven by regulatory and ethical considerations. He expressed this concern: *"I think that is a better model than different kinds of models, now I am changing my mind. I think it is better to have the same model that gets updated all at once. Everybody gets it, but you pay for it with either data or money."* This transition underscores the importance of a unified approach in ensuring regulatory alignment and ethical standards.

Consensus Building

Regulatory and ethical concerns drove the session toward a consensus on a single model. Aviation expert 2 emphasized: *"It is crucial to have one model. You are going to have big trouble with the regulators if you are going to give 20 different models to different operators."* This point highlights the potential risks and complications of justifying multiple versions to regulators, which could lead to fragmented and inconsistent safety standards.

Participant 2 MRO repeated these ethical concerns, suggesting multiple models could be perceived as intentionally sabotaging someone. This perspective aligns with ethical principles of fairness in high-stakes industries like aviation.

The session ended with a focus on ensuring continuous participation and equitable access to the model's benefits. Concerns about potential disincentives for early contributors were addressed by emphasizing the benefits of early participation and continuous model improvement. Participant 2 MRO countered this concern by highlighting the continuous improvement and benefits of early participation: *"I see two reasons why this is not an issue; the first reason is that those who delay their participation are penalized as they miss out on immediate benefits while others gain from using the model."* The second argument highlights: *"While the model continuously improves, joining at any later point still places participants ahead of others who may join even later."* This approach promotes a culture of continuous improvement and participation, ensuring that the model evolves and benefits all contributors over time.

6.2.2 Analysis

Influence and Bias Considerations

The role of Professor Tuler de Oliveira's clarification in shaping the final consensus must be addressed. While her input could potentially introduce biases, it is important to note that she mainly reiterated points already discussed earlier. The interviewees, being highly educated and experienced professionals, made decisions based on their extensive knowledge and expertise. This ensures that external influences, including a professor's, do not easily sway them.

Additionally, the long-standing relationship between the two aviation experts could introduce some bias, as they are familiar with each other's ideas. However, this familiarity also means they understand and respect each other's expertise, contributing to a broader discussion. Despite this, each expert made independent decisions, reflecting their professional judgment and experience.

The researcher's primary role in this process was facilitating the discussion and providing an environment for the experts to express their views. This involved organizing the session, providing the discussion topics and ensuring that all participants had the opportunity to contribute their insights. Maintaining an impartial stance allowed the experts to have the discussion based on their expertise and perspectives.

Key Outcomes

The expert validation session comprehensively explored model distribution strategies and incentive structures. The evolution from a multi-model to a single-model approach reflects a thoughtful consideration of various factors. The initial support for a differentiated distribution model underscored the importance of incentivizing substantial contributions. However, the concerns about fairness, regulatory challenges and ethical implications highlighted the complexities of implementing such a system.

The shift towards a unified model with differentiated access based on contributions represents a pragmatic solution to balance these competing interests. This approach addresses free-riding risks, ensures regulatory compliance and upholds ethical standards.

The discussion underscored the critical need for incentivizing substantial contributions and active participation. The session highlighted the benefits of early involvement, framing it as beneficial immediately and in the long term. This perspective aligns with collaborative innovation principles, ensuring the model remains dynamic and relevant.

Overall, the consensus reached reflects a deeper understanding of the industry's regulatory environment and the ethical considerations of providing equal access to the safest models. This approach ensures that all contributors, regardless of their entry point, benefit from a safe and evolving model, developing a culture of collaboration and shared progress.

6.3 Validation Interviews

6.3.1 Participants

The validation interviews were conducted with most participants (9 out of 12). Aviation Expert 2 and MRO 2 attended the validation session but could not schedule individual interviews due to availability constraints. Service Provider 2 did not respond to the invitation for an individual validation interview.

The validation interviews followed a similar setup to the first round. For a detailed overview of the methodology and considerations, refer to Section 3.2. Interviews were conducted online via MS Teams, each lasting about 30 minutes.

Transcriptions were made using MS Teams and then rewritten as structured summaries based on the questionnaire shared with the guiding professors. All retraceable data was securely stored. Ethical considerations remained significant, with participants having already provided written consent. They were reminded and asked for verbal consent to transcribe the sessions.

It is still notable that all participants were male. However, as previously indicated, research has shown that gender does not impact the willingness to share (Both, 2021). Making it, therefore, not a problem to have only male participants.

An overview of demographic data from participants for the validation interviews is shown in Table 6.

Table 6: Demographic Information Interviewees Validation (n=9)

Variables	Count, n (%)
Sex	
Male	9 (100)
Female	0 (0)
Intersex	0 (0)
Organization	
OEM	2 (22)
Operator	2 (22)
MRO	1 (11)
Service provider	1 (11)
Regulator	2 (22)
Aviation expert	1 (11)
Experience in aviation	
0-4	2 (22)
5-9	0 (0)
10-14	1 (11)
15-19	2 (22)
≥20	4 (44)
Region	
Europe	5 (55)
US	3 (33)
Latin-America	1 (11)

6.3.2 Validation Interview Setup

The requirements selected were chosen because their proposed solutions are critical and potentially modifiable based on stakeholder feedback. These requirements share common characteristics: their solutions are fundamental to the system’s integrity, security and compliance. Given the complexity and importance of these solutions, it is logical to question and validate them to ensure they meet the needs and expectations of participants. These requirements are important because they address the system's core functionalities and trust mechanisms, making it essential to confirm their adequacy and effectiveness through stakeholder input.

Requirements Validated

The following requirements have been used for the validation: R1, R3, R5, R6, R9 and R11. Below is an explanation of why R2, R4, R8 and R10 have not been used.

R1: The system must have strict security measures. The proposed solution involves using homomorphic encryption and differential privacy to protect data. Validating this requirement assesses participants' confidence in these advanced security measures and their effectiveness in safeguarding sensitive data.

R3: The system must ensure compliance with existing legal standards. Compliance is managed through a Compliance Database System and Legal Compliance System. Validating this requirement evaluates how well the system's design adheres to regulatory standards, ensuring proper handling of personal information and adherence to legal requirements by focusing the question on GDPR.

R5: The system must develop a consortium-based approach for managing co-created IP. The proposed solution includes a consortium-based IP management system using licensing and reputation scores. This validation examines participants' views on the fairness and effectiveness of the IP management approach, ensuring it protects the interests of all parties involved.

R6: The system must establish mechanisms for equitable benefit sharing. The solution involves the RFFL framework with a dynamic reputation system. Validating this requirement determines participants' opinions on the fairness and practicality of the distribution of different models.

R9: The system must ensure that AI models are explainable and transparent. The solution uses Shapley values to explain model decisions, ensuring transparency and compliance with regulatory requirements. This validation confirms whether participants believe these methods adequately meet the explainability requirements.

R11: The system must appoint a trusted, neutral entity for collaborative management and fair governance. The proposed governance structure involves a consortium using a consortium blockchain to ensure transparency and fairness. Validating this requirement assesses participants' trust in the proposed governance model.

Requirements Not Validated

The remaining requirements, although important, were not questioned. This decision aimed to ensure that the most impactful aspects of the system were thoroughly evaluated. These requirements were not validated due to their foundational nature or the abstract nature of their solution:

R2: The system must have a reliable data-sharing framework. This is fundamental to the design of using FL. Although using a consortium blockchain for transparency and security could have been questioned, there are currently no significant objections to this approach.

R4: The system must be adaptable to evolving regulations. The solution involves the Compliance Database System and Regulatory Reporting and Analysis Tool. This requirement is more abstract and complex to validate as it involves speculation about future regulatory changes.

R8: The system must be more accurate than current predictions. While accuracy is critical for the system's success, it is a necessary component that cannot be altered. Therefore, questioning this requirement was deemed unnecessary at this stage.

R10: The system must have training programs for stakeholders. Although beneficial, this requirement is less fundamental to the system's core design. Thus, immediate validation was not considered necessary.

Even though these requirements are important, limited time is available to focus on validating the solutions for the other critical requirements.

Other Considerations

Given the available interview time, a decision was made not to discuss the designed ArchiMate model with the interviewees. Explaining the model's workings would likely consume valuable time that could be better used in exploring the design choices regarding the different system components. This approach aimed to leverage the interviewees' expertise to gain insights into critical design decisions rather than model comprehension.

The interviewees held various technical roles, such as managing digital product programs, implementing blockchain solutions, providing IT services, developing predictive maintenance services and overseeing AI program management. To maintain privacy, specific roles were not linked to individual participants. An informative slide was prepared to facilitate understanding of technical concepts such as homomorphic encryption and differential privacy. This slide was used to explain these concepts at a high level if the interviewee was unfamiliar with them. These slides can be found in Appendix F. Similarly, the concept of Shapley values was explained to the interviewees to ensure clarity. Given their technical background, interviewees generally understood these concepts well.

Most interviewees had over ten years of working experience, making them well-suited to answer business-oriented questions. Since the interviewees were not experts in regulatory details, the questions about regulation were focused on GDPR compliance, a topic assumed to be within their general knowledge.

Furthermore, the primary goal of the interviews was to understand the acceptance level and readiness for adoption by the different organizations. The feedback gathered from these interviews will be instrumental in refining the system to ensure it meets both technical requirements and industry expectations.

6.3.3 Validation Interview Questions

The interview questions used for the individual validation interviews are presented in Table 7.

Table 7: Interview Questions Validation

Subject (Requirement)	Question
Security measurements and privacy (R1)	Your company's data will remain stored locally on your database to preserve confidentiality. For model-building, only derived model parameters are shared, not the raw data itself. The system utilizes advanced security measures: homomorphic encryption allows computations on these encrypted parameters without revealing the underlying data, and differential privacy involves adding noise to these computations to protect individual privacy. Given these security measures, how confident are you in the system's capability to safeguard your organization's sensitive data? (Along with this question, both homomorphic encryption and differential privacy were explained using the informative slides)

Regulatory compliance / GDPR (R3)	In this system, it is assumed that all pilot data, which could be personally identifiable, is removed by operators before any processing occurs. Additionally, the system mainly shares model parameters and not the actual data. Considering this setup and knowing that GDPR focuses on protecting personal data within the EU, how effectively do you believe this system complies with such data protection regulations?
IP management (R5)	In the proposed system, intellectual property created from the model is owned by the consortium. Participating companies are given a license to use this model, which regulates how they can use it and restricts them from selling it to third parties. This licensing arrangement is designed to protect the consortium's and participants' interests while ensuring fair use. What are your views on this approach to IP management and licensing?
Reward mechanism (R6)	The system employs a Robust and Fair Federated Learning framework that uses a reputation system to calculate each participant's contributions. Based on these contributions, participants are granted access to different versions of the model. More active contributors receive the most current version, while others may receive older versions. What are your thoughts on this method of distributing model access among participating organizations?
Explainability of the model (R9)	The EASA AI Act and EU AI Act require AI/ML models to be explainable to ensure they are understandable and transparent. Shapley values are suggested, a method in ML that assigns a numerical value to each feature in a model, indicating how much each feature contributes to the model's predictions. Do you think using Shapley values will adequately meet the explainability requirements mandated by regulatory entities like EASA?
Trust in the governance (R11)	The system is managed by an independent consortium, such as IDCA, which acts as the central entity. This structure is intended to ensure neutrality and fairness when managing the Federated Learning system. How does this governance approach affect your trust in the system's overall integrity and effectiveness?

6.4 Results Validation Interview

6.4.1 Overview Results

The participants are divided into the ones who agree and those who have some concerns. The regulators have not asked questions regarding IP management and trust in governance. These questions seemed more relevant for the commercial parties. The total number of respondents for these questions is thus seven instead of 9. An overview of the perceptions of the design choices can be found in the Table 8; the detailed findings are presented below.

Table 8: Views Acceptance Design Choices

Subject	# agreeing	# concern	Agreeing	Having concern
Security Measurement and Privacy (R1)	4	5	OEM 1 Regulator 1 & 2 Service Provider 1	Aviation Expert 1 MRO 1 OEM 2 Operator 1 & 2
Regulatory Compliance / GDPR (R3)	9	0	Everyone	None
IP Management (R5)	4	3	MRO 1 OEM 1 Operator 2 Service Provider 1	Aviation Expert 1 OEM 2 Operator 1
Reward Mechanism (R6)	3	6	MRO 1 OEM 1 OEM 2	Aviation Expert 1 Operator 1 & 2 Regulator 1 & 2 Service Provider 1
Explainability of the Model (R9)	3	6	OEM 1 OEM 2 Service Provider 1	Aviation Expert 1 MRO 1 Operator 1 & 2 Regulator 1 & 2
Trust in the Governance (R11)	7	0	Everyone	None

6.4.2 Security Measurements and Privacy

Agreement (4 Participants)

Several participants expressed cautious optimism about the proposed security measures. Service Provider 1 stated: *"It sounds like it would be acceptable to the community,"* implying a cautious acceptance rather than confidence. OEM 1 highlighted the security design's potential: *"The steps you have taken provide a maximum level of safety and protection ... the design you are considering with homomorphic encryption and differential privacy indeed adds a significant layer of security."* This response indicates a more certain approval of the system. Regulator 1, an expert in policy and AI, emphasized the importance of the security framework: *"Assessment, mitigation and verification of security measures are crucial to addressing potential vulnerabilities introduced by Federated Learning systems."* Given that he is an expert, this suggests a good focus on procedural thoroughness.

Concerns (5 Participants)

Other participants, particularly from technical and operational roles, expressed significant concerns about the practical challenges of implementing differential privacy. Especially regarding its impact on data precision. OEM 2 mentioned: *"Differential privacy introduces noise to protect data, but in aviation, where accuracy is critical, this could potentially lead to erroneous analyses."* This indicates a serious worry that the privacy measures could undermine the accuracy and reliability, which are critical in aviation.

MRO 1 emphasized the delicate balance needed: *"There is a fine balance between enhancing privacy and maintaining the integrity of our data. Too much noise can render the data useless for our precise needs."* This underscores a fear that privacy measures could render the data operationally ineffective.

Analysis and Implications

The divide in opinions regarding security measures underscores a fundamental tension in balancing data privacy and precision. The cautious optimism from some participants suggests that while they recognize the theoretical benefits of advanced encryption techniques, they remain skeptical about their practical implementation. This skepticism is particularly spoken by those directly responsible for operational accuracy. This could indicate a potential barrier to adoption if these measures cannot be proven to maintain data integrity. The implications are clear: Adopting FL in aviation could face significant resistance unless the security measures can be used effectively without compromising data quality.

6.4.3 Regulatory Compliance / GDPR

Agreement (9 Participants)

All participants agreed that the design complies with GDPR due to the anonymization of pilot data and the sharing of model parameters instead of raw data. Aviation Expert 1 noted: *"By focusing on model parameter sharing rather than raw data, we adhere to GDPR's stringent requirements for data protection."* Operator 1 clarified: *"The anonymization processes we have implemented ensure that GDPR concerns are comprehensively addressed, securing personal data against misuse."* This indicates confidence in the compliance strategy regarding GDPR, though it might overlook broader concerns.

From the conversation with the regulators, Regulator 1 urged consideration for the future: *"GDPR forms the baseline of our data protection strategies and we are extending these protections to cover broader ethical concerns that transcend traditional data privacy."* This suggests a comprehensive data protection approach beyond regulatory compliance, emphasizing future ethical considerations.

Analysis and Implications

The unanimous agreement on GDPR compliance reflects strong confidence in the system's design to protect personal data. Anonymization and model parameter sharing effectively meet privacy concerns but might not fully address all regulatory concerns. The regulators' emphasis on broader ethical issues highlights a potential gap in participants' understanding of comprehensive data protection.

It is important to note that the interviews addressed GDPR compliance and did not cover all relevant regulations. A full regulatory assessment would require input from specialized professionals and more time. Validating GDPR compliance is done because it is an impactful regulation. As the model evolves and potentially makes autonomous decisions, ethical considerations will become more relevant, necessitating new regulatory and ethical standards. Failure to adopt this comprehensive approach could lead to challenges if broader ethical considerations become more prominent in future regulatory standards.

6.4.4 IP Management

Agreement (4 Participants)

Over half of the participants supported the consortium retaining IP ownership and providing licenses to contributors. MRO 1 expressed this as a logical framework. Service Provider 1 also appreciated this approach, stating, *"This model of IP management, where contributors get access without extra costs, seems very fair and likely to encourage participation."* This reflects a belief in the fairness and practicality of the proposed IP management strategy, suggesting it could facilitate broader participation and collaboration among stakeholders.

Concerns (3 Participants)

Others raised concerns about the complexities of IP management, fearing it might restrain innovation. Operator 1, focused on operational efficiency regarding budgets and business impact, pointed out: *"You need to have an upfront investment for people working on it and, knowing budgets, not getting IP rights feels not okay ... Everything that we develop should be our intellectual property right."* This reveals a strong concern for ensuring that contributions are adequately compensated, potentially viewing the current proposal as insufficiently rewarding. OEM 2 shared this sentiment: *"Intellectual property rights are incredibly sensitive in our field and this system needs to carefully balance protection with innovation."* This underscores a fear that the proposed IP management could stifle innovation by not providing enough incentive for participation.

Analysis and Implications

The divided opinions on IP management reveal a tension between protecting individual interests and shared benefits. The concerns raised suggest that the current proposal might not provide sufficient incentives for all participants, particularly those unfamiliar with this way of working. This could hinder stakeholders' participation and limit the system's effectiveness and innovation capacity.

Ensuring the IP management framework is fair and sufficiently rewarding will create a collaborative and innovative environment. It is essential to clarify that contributors can use the IP without incurring additional costs intended to benefit all participants. The hesitation among some participants might stem from unfamiliarity with this model. Highlighting the design's benefits and ensuring clear communication could remove these concerns.

Failure to address these issues could result in reduced participation and slower innovation progress. Therefore, demonstrating how the IP management strategy is designed to be equitable and advantageous for all involved ensures that contributors feel adequately rewarded and incentivized to participate actively.

6.4.5 Reward Mechanism

Agreement (3 Participants)

Some participants appreciated the structured incentive system designed to reward data contributions. MRO 1 liked the idea of being rewarded for participation. This sentiment was shared by OEM 1, who stated: *"I think that that is a novel approach. I think that that design creates some incentive ... wherein if you are willing to share or be able to contribute the most ... then you get the most current version to be able to operate with."* This reflects a recognition of the value of incentivizing contributions, though it also unconsciously highlights a significant discussion about whether all participants can share 'the most' data.

Concerns (6 Participants)

Significant concerns were about ensuring fairness, particularly regarding smaller or less resource-rich participants. Operator 2, focusing on operational realities, questioned: *"How do we ensure that the reward system does not inadvertently penalize those who might not have as much data to contribute?"* Even though the idea behind this was that not the best model would still be better than their own developed model, this still could be a potential flaw in the system. This way could disadvantage smaller organizations that may not be able to contribute as much data as larger entities.

Service Provider 1 raised regulatory and safety concerns: *"Imagine a situation where different parties are using different versions of the model because they have contributed at different levels. It becomes a regulatory nightmare to ensure that all versions meet safety and compliance standards, not to mention the potential public safety implications if an outdated model fails to perform adequately."*

This underscores a significant risk in the proposed mechanism, where varying model versions could lead to safety and regulatory compliance issues. Regulator 1, while understanding from a business perspective, expressed concerns from a safety and regulatory viewpoint: *"Distributing different versions of the model based on contributions could lead to inconsistencies that are problematic from a safety regulation perspective."*

Analysis and Implications

The reward mechanism's fairness and practical implications drew significant concern. The potential for disadvantaging smaller participants and the risk of regulatory and safety issues due to varying model versions are critical challenges that must be addressed. The current proposal could lead to unequal benefits and operational inconsistencies, undermining the system's credibility and effectiveness.

To mitigate these risks, the reward system must be redesigned to recognize contributions while equitably ensuring uniform safety and regulatory standards. It is important to emphasize that while the system effectively encourages initial participation, incentivization should focus on rewarding both larger and smaller organizations. Failure to address these concerns could reduce trust, participation, and potential regulatory hurdles.

6.4.6 Explainability of the Model

Agreement (3 Participants)

Some participants supported the use of Shapley values to enhance model explainability. OEM 1 noted: *"I think it makes sense ... You know, in the sense of you never really know until you get to the point of first contact with a regulator around such things."* This reflects an unsure acceptance of Shapley values and a cautious attitude towards regulatory feedback. Service Provider 1 was more confident: *"Utilizing Shapley values to clarify how decisions are made enhances our ability to audit and trust the model, which is crucial for wider acceptance."*

Concerns (6 Participants)

There were significant concerns about whether Shapley values alone could meet regulatory explainability requirements. Operator 2 expressed this uncertainty: *"I am unsure if Shapley values alone will satisfy our regulatory requirements. There is a gap between theoretical explainability and practical application."* MRO 1 echoed this sentiment: *"While Shapley values are a good start, they may not comprehensively address all the nuances required by regulators in our industry."*

Both regulators discussed the necessity of explainability in AI models, mandated by regulations like the EASA AI Act. Regulator 2 acknowledged the value of Shapley values but pointed out their limitations in complex scenarios: *"While Shapley values add explainability, they are not alone sufficient, especially in complex use cases."* Regulator 1 emphasized the need for a broader set of tools to enhance transparency: *"Shapley values are useful, but we must integrate them with other tools to ensure comprehensive explainability from both development and operational perspectives."*

Analysis and Implications

With the currently available techniques, explaining exactly how a complex model works is impossible. However, the skepticism regarding Shapley values alone meeting regulatory requirements indicates a need for a broader set of tools to enhance transparency. The cautious support from some participants and significant concerns from others suggest that relying solely on Shapley values may not be sufficient for regulatory compliance.

There remains considerable uncertainty surrounding the regulations of ML models, an issue that extends beyond aviation. As OEM 1 mentioned, unpredictability is expected until an actual regulatory review occurs. This underscores the lack of current regulations and the evolving nature of regulatory expectations. One effective strategy to navigate this uncertainty is to engage regulators early in the development process to shape explainability standards collaboratively.

The limitations of HE must be critically assessed. Initially, the focus was predominantly on safety, but given the current objections regarding explainability, a more critical examination is required. The broader context of ongoing struggles to explain ML models and the significant concerns raised even with the current proposal suggest that relying on HE could complicate transparency too much. Therefore, it is essential to reconsider the role of HE and explore alternative approaches that better align the intention to make the model as transparent and understandable as possible. Addressing these concerns is crucial to achieving regulatory compliance and maintaining trust in the model's outputs.

6.4.7 Trust in the Governance

Agreement (7 Participants)

All participants unanimously supported the governance of the system by an independent consortium. This structure is needed to ensure neutrality and fairness when managing the FL system. Service Provider 1 stated: *"Having an independent party manage the system seems the best way to maintain neutrality and trust among all stakeholders."* This sentiment was echoed by Aviation Expert 1 who echoed this statement, emphasizing: *"A neutral overseer can ensure that no single entity can dominate or skew the model to their advantage."*

Analysis and Implications

The strong support for governance for an independent consortium underscores the importance of neutrality and fairness in managing the system. This consensus indicates that participants believe such governance will promote trust and collaboration, which are critical for the system's success.

Ensuring transparent and effective governance structures will maintain trust and cooperation. Consortium blockchain technology within the governance framework further enhances transparency and trust. By providing an immutable and transparent ledger of all transactions and changes within the FL system, blockchain ensures that all actions are recorded and can be audited. This additional layer of accountability increases the reliability of the system.

The consortium's role should include clear communication, regular audits and inclusive decision-making processes to ensure all stakeholders' interests are represented and protected. The integration of blockchain technology supports these processes by providing a transparent record of all activities.

6.4.8 Conclusion

The interviews' analysis has revealed critical insights into the perspectives and concerns of various stakeholders regarding implementing the FL system. While there is generally considerable support for the system, two aspects must be reconsidered: the distribution mechanism and differential privacy.

One of the pressing issues identified is the distribution mechanism of the models. Initially, distributing multiple versions of the model based on contributions was considered a promising solution. However, feedback from the expert sessions revealed significant regulatory and ethical concerns with this approach. In aviation, where safety is paramount, distributing different versions of the model that could be perceived as inferior is neither regulatory compliant nor ethically acceptable.

This underscores the need for a new distribution strategy that ensures all participants can access the most current and safest model, regardless of their contribution level. The revised system should focus on creating a fair and equitable access mechanism that still incentivizes participation but does not compromise on regulatory and safety standards.

Another critical area requiring attention is the implementation of differential privacy. While it is a good method for protecting individual data points within aggregated datasets, it inherently introduces some level of distortion or noise to prevent the identification of individual data sources. This can slightly reduce the accuracy of data analysis. Given the high stakes in aviation, where the accuracy of predictions is crucial for safety and operational efficiency, this trade-off is particularly problematic. The participant feedback highlighted a consensus that, in aviation analytics, even a minor compromise on accuracy for privacy is not desirable. The non-negotiable stance on maintaining predictive accuracy has been underestimated. Therefore, there needs to be a critical reevaluation of differential privacy within the system. The goal should be to reassess the level of distortion in the outcomes and decide whether to implement this approach.

In summary, the proposed FL system has gained support for many aspects. However, it is important to address the challenges of the distribution mechanism and differential privacy. Ensuring regulatory compliance, ethical acceptability, and accurate predictions will be essential for the system's success. By focusing on these areas, the various stakeholders can accept the system.

6.5 Changes to the system

6.5.1 New Distribution Mechanism

Requirement R6 mandates establishing mechanisms for equitable benefit sharing to ensure that all participants are fairly rewarded for their contributions. Ensuring motivation, collaboration and trust within the consortium results in long-term engagement and addresses contribution disparities. The analysis of validation interviews revealed significant regulatory and ethical concerns with distributing multiple versions of the predictive maintenance model, as this could compromise safety standards. Experts highlighted the need to provide all participants with a single, updated model to ensure safety compliance and maintain trust and collaboration within the consortium.

6.5.1.1 New Design: Uniform Model Distribution with Differentiated Access

The new design consists of distributing the latest version of the model to all consortium members. This ensures that every participant has access to the most advanced and safest tool available, maintaining a consistent safety standard across the board. However, the extent to which participants can utilize the model will depend on their contribution score. This score is still determined based on the data's quality, quantity, and timeliness. Previously, the contribution score was used to determine which version of the model each participant received. In this new design, the score determines the number of times participants can use the model.

Each participant earns tokens based on their contribution score, which can be spent to run the model. Tokens are a measurable and flexible mechanism, making it easy for participants to understand the direct link between their contributions and model access, creating a transparent incentive structure. Tokens represent the right to use the model, providing a tangible incentive for participants to contribute more data.

6.5.1.2 How the Contribution Score Translates to the Incentive Model Usage

Limit on Model Runs: The contribution score dictates the number of times a participant can run the model. Organizations with higher scores gain more access, allowing them to optimize processes, improve safety, and make better-informed decisions. Regardless of the organization's size, contributing more data remains a key incentive, as it directly increases their model usage.

Practical examples for the three most prominent groups are provided below:

- **OEMs:** The contribution score could determine the number of tests they can run on their equipment using the predictive maintenance model. Higher scores enable OEMs to use the model to test and improve their developed equipment, enhancing the design and reliability of their products through extensive predictive insights.
- **Operators:** The contribution score indicates how many aircraft in their fleet can benefit from the model's predictive maintenance capabilities. Operators with higher scores can apply the model to a bigger portion of their fleet, enhancing overall operational efficiency and safety.
- **MROs:** The contribution score could determine the number of maintenance schedules or components that can be analyzed using the model. Higher scores allow MROs to optimize more maintenance processes, improve turnaround times and reduce costs through predictive insights.

6.5.1.3 Ensuring Equitable Benefit Sharing

This approach ensures that benefits are distributed equitably, rewarding those who contribute more and higher-quality data with greater access. Smaller companies with lower data contributions receive fewer tokens, aligning their benefits with their capabilities. In comparison, larger companies with extensive data contributions receive more tokens to meet their more significant operational needs.

- **Smaller Companies:** These entities may be unable to provide as much data due to their smaller fleet size, resulting in fewer tokens. However, their operational needs are naturally lower, which aligns their benefits with their capabilities. This ensures that smaller companies remain motivated to participate, knowing their efforts are valued and their contributions are fairly compensated.
- **Larger Companies:** Larger organizations with more extensive fleets can provide more data, earning them a higher contribution score and more tokens. This increased access is necessary due to their greater operational needs. Their substantial contributions will be fairly rewarded with greater usage rights. The ability to run the model more frequently and apply its predictive maintenance capabilities across a larger scope of operations becomes a critical incentive for these organizations to continue sharing data at a high level.

6.5.1.4 Smart Contracts

Smart contracts can be employed to ensure the equitable distribution of tokens and the transparent execution of consortium rules. These self-executing contracts automatically trigger actions, such as allocating tokens, when predefined conditions are met (Zheng et al., 2020). These conditions can include factors like the volume, quality, timeliness, originality, and cleanliness of the data contributed by participants. For example, data that is highly refined, free from errors or duplicates, and entirely original could be weighted more favorably in the token allocation process. When a participant uploads their data, the smart contract verifies the contribution score based on these multifaceted criteria and immediately allocates the corresponding number of tokens, ensuring a fair and transparent system.

Smart contracts operate on blockchain technology, therefore every transaction is securely recorded on a distributed ledger. This means that once a smart contract allocates tokens, the transaction is immutable and visible to all consortium members, ensuring complete transparency and eliminating the risk of tampering (Zheng et al., 2020). The system operates without the need for intermediaries, reducing administrative work for the consortium. Each consortium member can trust that their contribution is rewarded fairly and that their tokens are securely managed through this decentralized, tamper-proof system.

In addition, smart contracts enable dynamic adjustments. A change in data contributions or the usage of tokens results in automatic updates on the token balance because of the dynamic nature of the smart contracts. This dynamic adjustment ensures that participants are consistently rewarded for their real-time contributions, creating a culture of ongoing incentives and fairness. The system can also scale to accommodate new participants or changing rules within the consortium, as smart contracts can be updated with new conditions without disrupting the entire process (Zheng et al., 2020).

6.5.1.5 Impact on Consortium Phases and Participant Engagement

The principles regarding different ways of joining the consortium still apply:

- **New Participants:** Initially, new participants start with a lower contribution score, which increases as they contribute more data, encouraging continuous improvement. This system ensures that new entrants are motivated to enhance their data-sharing practices.
- **Established Members:** Established members with higher scores enjoy greater access, reflecting their significant data contributions. This recognizes their long-term involvement and incentivizes them to maintain their high level of data sharing.
- **External Organizations:** External organizations can purchase tokens for temporary access, introducing a financial dimension that benefits all consortium members. This flexibility allows external entities to utilize the model without a long-term commitment.

Dynamic Adjustments and Transparency: Similar to the earlier-designed system, contribution scores are dynamically adjusted based on ongoing data contributions, resulting in a culture of continuous participation. Transparent reporting mechanisms will show how tokens are earned and spent, ensuring participants understand the value of their contributions.

Additional Token Purchases: Consortium members may purchase extra tokens if they need to perform more calculations than their tokens allow, providing further flexibility and ensuring all operational needs are met.

6.5.1.6 Conclusion New Distribution Mechanism

The enhanced distribution mechanism ensures that the most current and safest model is uniformly distributed to all participants, addressing regulatory and ethical concerns. By differentiating access based on contribution scores, the system promotes fairness and sustained engagement, fulfilling the purpose of Requirement R6 (equitable benefit sharing). This approach suits organizations of all sizes, from smaller companies with fewer operational needs to larger companies with extensive fleets. Smaller companies remain motivated as their contributions are valued and fairly compensated, while larger companies benefit from greater usage rights that match their substantial contributions.

By leveraging smart contracts, the token-based usage system ensures an automated, transparent, and secure way to manage both the allocation and use of tokens. Smart contracts facilitate fair token distribution by autonomously executing transactions based on predefined rules, while blockchain technology ensures that all transactions are immutable and auditable.

This automation minimizes human intervention, reduces the likelihood of disputes, and promotes trust among consortium members. Furthermore, the system's dynamic adjustments allow it to scale and adapt, ensuring sustained participation and continuous incentives from all consortium members, regardless of their size or operational scope.

6.5.2 Trusted Execution Environment instead of Homomorphic Encryption

Throughout the research, HE was considered a feasible solution for ensuring data privacy during computations. HE's ability to perform operations on encrypted data without decrypting it provided a robust mechanism for preserving confidentiality. However, during the validation interviews, concerns about the explainability of models using HE were repeatedly raised. The complex nature of encrypted computations made it challenging to interpret and explain the decision-making processes of machine learning models, which is crucial for regulatory compliance and stakeholder trust. These concerns required further exploration into alternative technologies that could offer both security and transparency.

Trusted Execution Environments (TEEs) showed as a promising solution, especially due to their practical possibilities since this is already used by the likes of Google and Microsoft (Google Cloud, 2024; Microsoft, 2023). Other cryptographic methods, such as Multi-Party Computation, were also considered. While MPC offers strong security guarantees, it performed poorly in terms of computational efficiency and added significant complexity when explaining model outcomes (Du & Atallah, 2001). The widespread adoption of TEEs in industries requiring both high security and transparency further supported their suitability for addressing the dual challenge of privacy and explainability in the FL system explored in this paper.

6.5.2.1 How TEE Protects Raw Data from Client Visibility

A TEE is a secure area within a computer's main processor that isolates sensitive data and code during processing. It ensures the confidentiality, integrity, and protection of data from both unauthorized access and even from the operating system. TEEs create a secure enclave within the processor, preventing data from being accessed or tampered with by external threats. TEEs achieve this isolation using hardware mechanisms that enforce strict separation between the enclave and the rest of the system. This ensures that sensitive operations within the TEE remain secure, even if the rest of the system is compromised (Kohlbrenner et al., 2020).

In an FL system where the central server employs a TEE, clients can send their data to the server without worrying about the server accessing or compromising their raw data. The TEE ensures that data remains protected during aggregation and computation. This setup allows the central server to compute the client data while ensuring that the raw data remains inaccessible to even the server itself. This architecture addresses privacy concerns, providing enhanced guarantees that sensitive data is shielded during the learning process.

6.5.2.2 Comparison Between Homomorphic Encryption and TEE

HE allows computations to be performed on encrypted data, preserving data privacy throughout the process. However, HE introduces significant challenges, particularly around explainability in machine learning. Since the data and computations are encrypted, it becomes difficult to interpret and explain the model's decisions in a manner accessible to regulators or stakeholders (Kohlbrenner et al., 2020).

In contrast, TEEs provide a balance between security and transparency. Data is protected from unauthorized access during computation without being encrypted, allowing for clearer insights into how models process the data (Geppert et al., 2022).

While HE excels in preserving data privacy, TEEs offer strong protection while maintaining transparency, making them more suitable for the context where both security and explainability are necessary.

In the FL context, HE can secure the data transmitted from clients to the central server, but it complicates the ability to explain model outcomes due to the complex nature of encrypted computations. In contrast, by using TEEs on the central server, computations can be secured while preserving the necessary transparency to explain and understand the model's decisions.

6.5.2.3 Transparency of Computation with Remote Attestation

A key feature of TEEs is their support for remote attestation, which allows external parties to verify that computations within the TEE were executed securely and as intended. Remote attestation generates cryptographic proofs that attest to the integrity and authenticity of operations within the TEE (Kohlbrener et al., 2020).

This feature is particularly valuable in machine learning, where ensuring the integrity of computations is as crucial as the accuracy of the model's output. Remote attestation enables stakeholders to trust that the model aggregation and computation were conducted within a secure environment, enhancing system credibility (Geppert et al., 2022). By providing verifiable evidence of secure execution, TEEs address concerns about the complexity of machine learning models, offering transparency and accountability in the computational process.

In the FL system, where only the central server uses a TEE, remote attestation allows the clients to verify that their model updates are processed securely within the central server's TEE. This builds trust that the central server is executing the aggregation of models without compromising data integrity.

6.5.2.4 Auditable Logs of Operations

TEEs also provide the capability to maintain auditable logs of all operations conducted within the secure environment. These logs serve as a detailed record of data access, computations, and significant events, providing an audit trail that can be reviewed by regulators and auditors (Geppert et al., 2022).

The use of auditable logs enhances the ability to monitor and verify machine learning models' operations, especially in high-stakes environments like aviation. In contrast to HE, which obscures the computational process, TEEs provide auditable records that regulators and stakeholders can use to ensure the system's compliance with privacy and security standards. These logs provide transparency and independent verification, offering improved regulatory oversight and operational accountability. In the FL system, the server can generate auditable logs of its operations ensuring compliance with privacy and security standards. These logs reassure authorities that the data was handled correctly, even if the server itself cannot access the raw data.

6.5.2.5 Relevance of the Previous Work

The research and interviews remain relevant, even with the shift from HE to TEEs. During interviews, participants expressed support for HE, but their main interest seemed to be the idea of preventing the central server from accessing raw data. While some interviewees understood the specifics of HE, many focused more on the privacy guarantees it offered. Despite explanations through PowerPoint, it was clear that their preference centered on security rather than the technical details.

In this research, HE was initially used to ensure that the central server could compute without seeing the clients' actual data. Although we now use TEEs, the primary goal remains the same: keeping client data secure from the central server during computation. TEEs provide a different method to achieve this by using secure hardware enclaves, but the objective of protecting privacy has not changed.

While conducting another validation round would be ideal, time constraints did not allow for this. However, the design of such systems is an ongoing process. Future iterations can include new rounds of validation to further refine and improve the approach, ensuring it continues to meet the original goals of security and transparency.

6.5.2.6 Conclusion Trusted Execution Environment

Switching from HE to TEEs offers a better solution for balancing privacy and understanding how the model works. HE provides strong security by allowing data to remain encrypted during computations, but it makes the model's decisions hard to explain, which is important for building trust and meeting regulations. TEEs, on the other hand, keep data secure without encryption and allow clearer insights into the model's decision-making. Features like remote attestation and auditable logs further build trust by allowing stakeholders to verify that computations were performed securely. TEEs provide a better balance of security and transparency, making them a more practical and sustainable choice for the consortium's long-term goals.

6.5.3 Removing Differential Privacy Implementation

6.5.3.1 Initial Selection and Evaluation of Differential Privacy

Differential privacy was initially chosen to meet the consortium's Requirement R1: The system must have strict security measures. This requirement aims to protect sensitive data and ensure data safety and security. Differential privacy was selected for its ability to protect sensitive information while allowing meaningful data analysis, creating a collaborative environment based on trust. Differential privacy provided strong data confidentiality by ensuring that individual data entries could not be re-identified. The focus was on generating trust among participants, accepting some loss of accuracy as a trade-off to achieve this goal.

The initial analysis revealed concerns about the impact of differential privacy on model accuracy. While introducing noise led to a decrease in predictive accuracy, it was initially considered workable. The primary objective was to ensure that differential privacy would generate trust and safety among consortium members, accepting reduced accuracy to maintain data confidentiality and participant confidence.

6.5.3.2 Reassessment of Differential Privacy's Long-term Viability

After analyzing the interviews, additional research was conducted on differential privacy. This research revealed several issues that needed reconsideration.

Additional evidence from the literature indicated that differential privacy might cause too much distortion after many iterations. The study by Wei et al. (2020) illustrates the trade-offs involved. Their findings show a consistent decline in accuracy over multiple iterations of FL models employing differential privacy. In simple terms, as the number of iterations increases, the model's performance degrades due to the added noise to ensure privacy.

The reinvestigation highlighted additional concerns about the long-term impact on the model. Over time, the cumulative effect of noise addition was found to degrade the model's performance progressively. This deterioration raised serious questions about the sustainability and reliability of the predictive maintenance model under a differential privacy framework. The long-term implications included reduced model trustworthiness and effectiveness, which could compromise the consortium's ability to maintain high safety standards. It became evident that, in the long term, differential privacy is not a viable solution due to its impact on the model's accuracy and performance.

6.5.3.3 Alternatives to Differential Privacy

The primary purpose of differential privacy was to generate trust and ensure data safety. Given the identified issues, exploring alternative solutions to achieve these goals without compromising model accuracy is needed.

One alternative is the implementation of traditional contractual agreements among participating companies. These agreements legally bind companies to not attempt access to each other's data, with significant financial penalties for breaches. This contractual approach aims to develop a trust-based environment similar to what differential privacy intended to create but without compromising data accuracy. By relying on legally enforceable contracts, companies can be assured of the confidentiality and integrity of their data. This method not only provides a clear legal framework for data protection but also reinforces the commitment of each participant to uphold the consortium's data security and privacy standards.

6.5.3.4 Conclusion Differential Privacy

The exploration of differential privacy as a security measure was initially aligned with Requirement R1: The system must have strict security measures to protect sensitive data and ensure data safety and security. Differential privacy was chosen for its potential to anonymize individual data entries, generating a trust-based collaborative environment. Initial analyses indicated a trade-off between privacy and accuracy, deemed acceptable to maintain data confidentiality and participant confidence. In the short term, differential privacy provided immediate data privacy and trust generation benefits. However, long-term use revealed that it is not a workable solution due to its harmful effects on model accuracy and performance.

An alternative is implementing traditional contractual agreements among participating companies. These agreements would legally bind companies to protect each other's data, with severe financial penalties for breaches. Such contractual measures can establish a trust-based environment similar to what differential privacy aims to achieve but without sacrificing data accuracy. This approach, coupled with technical solutions provided by FL and homomorphic encryption, ensures that Requirement R1 is still met.

In conclusion, while differential privacy offers a short-term solution for data privacy and trust, its long-term viability is compromised by its negative impact on accuracy. The consortium can achieve secure and trustworthy data management through legally binding agreements and advanced technical solutions, ensuring high data safety and trust standards.

6.7 Architectural Changes

The proposed changes in the distribution mechanism and the use of contracts instead of differential privacy necessitate specific adjustments to the overall system architecture. This chapter outlines the necessary changes, focusing on implementing a token-based system to limit model usage and establishing data-protection contracts within the consortium.

6.7.1 Token-Based System for Limiting Model Usage

Introducing a token-based system requires adjustments to the system's architecture to manage the allocation, usage, and transparency of tokens effectively. As previously outlined, tokens are generated based on contribution scores, which reflect the quality, quantity, and timeliness of the data contributed by each participant. These scores will determine the allocation of tokens across the consortium.

A secure token management system must be implemented to track token allocation, transfer, and usage. This system will be managed on the consortium blockchain, leveraging its inherent security

features, including encryption and immutable transaction logging, to prevent unauthorized access or tampering.

Smart contracts play a central role in automating these processes. They handle the generation and allocation of tokens, ensuring fairness and transparency without manual intervention. Additionally, smart contracts will automate token transfers and usage deductions, ensuring that each transaction, including model runs and token balances, is immutably recorded on the blockchain, thus maintaining system integrity.

In the new architecture, the existing mechanism to "Distribute Rewards" must be replaced by a system that distributes tokens based on contribution scores. This change ensures that participants are compensated with tokens, creating a measurable and scalable incentive structure that directly reflects their contributions. The token-based system serves as an ongoing incentive for organizations to actively contribute high-quality data, keeping them motivated to engage with the consortium continuously. By linking contributions to model usage, the system ensures that participants remain incentivized to maximize their involvement.

Smart contracts are integrated into the blockchain as a key addition to the technology layer. By leveraging the established blockchain infrastructure, smart contracts enhance security and eliminate the need for intermediaries, ensuring that all participants can trust the system's integrity and fairness.

6.7.2 Trusted Execution Environment

The TEE replaces the HE and is positioned within the technology layer of the architecture. It serves the same purpose as HE and has similar interactions, making it easy to substitute in the architectural design. As a secure enclave within the central server, the TEE ensures that data processing is protected from unauthorized access, including from the server itself. This allows sensitive computations to be securely performed without exposing raw data, aligning with the consortium's requirements for data privacy and transparency.

The TEE integrates seamlessly with the existing blockchain and smart contracts, which manage token distribution and model usage. Remote attestation, a key feature of TEEs, provides cryptographic proof that computations were securely executed. These proofs are stored on the blockchain, ensuring transparency and trust in the system. Therefore there is a connection that represents a flow of information from the TEE to the Consortium Blockchain. Audit logs generated by the TEE offer a detailed record of all computations, these are also stored on the Blockchain. This ensures compliance with privacy regulations and provides an additional layer of verification, complementing the transparency of the blockchain.

6.7.3 Legal Contracts within the Consortium

Establishing legal agreements between the consortium and its members is essential to ensure compliance and protect data integrity. These agreements will focus on preventing unauthorized access to each other's data and maintaining high data privacy and security standards.

First, specific data protection clauses need to be incorporated. These clauses will ensure that all members uphold high data privacy and security standards by outlining the measures they must take to protect their data and the data they receive from others. Furthermore, these clauses will explicitly prevent attempts to deduce sensitive information from shared parameters, thereby safeguarding the confidentiality and integrity of the data.

Second, defining clear penalties for misconduct is necessary. These penalties will address unauthorized access, data breaches, or attempts to tamper with the token system. By establishing these penalties, the consortium can deter potential misconduct and ensure all members are accountable for their actions. Additionally, a framework for addressing breaches promptly and effectively will be established, ensuring that violations are managed transparently and fairly.

These contracts will be established between the consortium and individual members, ensuring direct accountability and protecting against data misuse. The implementation of these legal agreements can be seamlessly integrated into the current architecture. These measures can be quickly adopted by leveraging existing governance and compliance frameworks without significantly changing the overall system structure. This ease of integration ensures that the consortium can maintain operational efficiency while enhancing data security and trust among its members.

6.8 Overview of the Validated System

The TEE is positioned within the technology layer of the architecture, replacing the previously used HE. A flow of information between TEE and Consortium Blockchain is added. This information consists of cryptographic proofs from remote attestation and audit logs generated by the TEE are sent to the blockchain, ensuring transparency and compliance. The Consortium Blockchain is connected with a newly added element Smart Contracts. This connection ensures seamless interaction with smart contracts, which manage token distribution and model usage.

Additionally, the block previously labeled 'Distributing Rewards' is now 'Distributing Tokens,' to emphasize the implementation of tokens in the redesigned system.

Although the architecture has been updated, the use of traditional legal agreements between consortium members is unseen. These contracts play a role in ensuring compliance with data protection standards and accountability. However, their implementation is part of the core interaction and does not impact the architectural design.

Figure 16 underscores the differences between the previous model and the updated version, highlighted in green. The new model, without highlights, can be found in Appendix H.

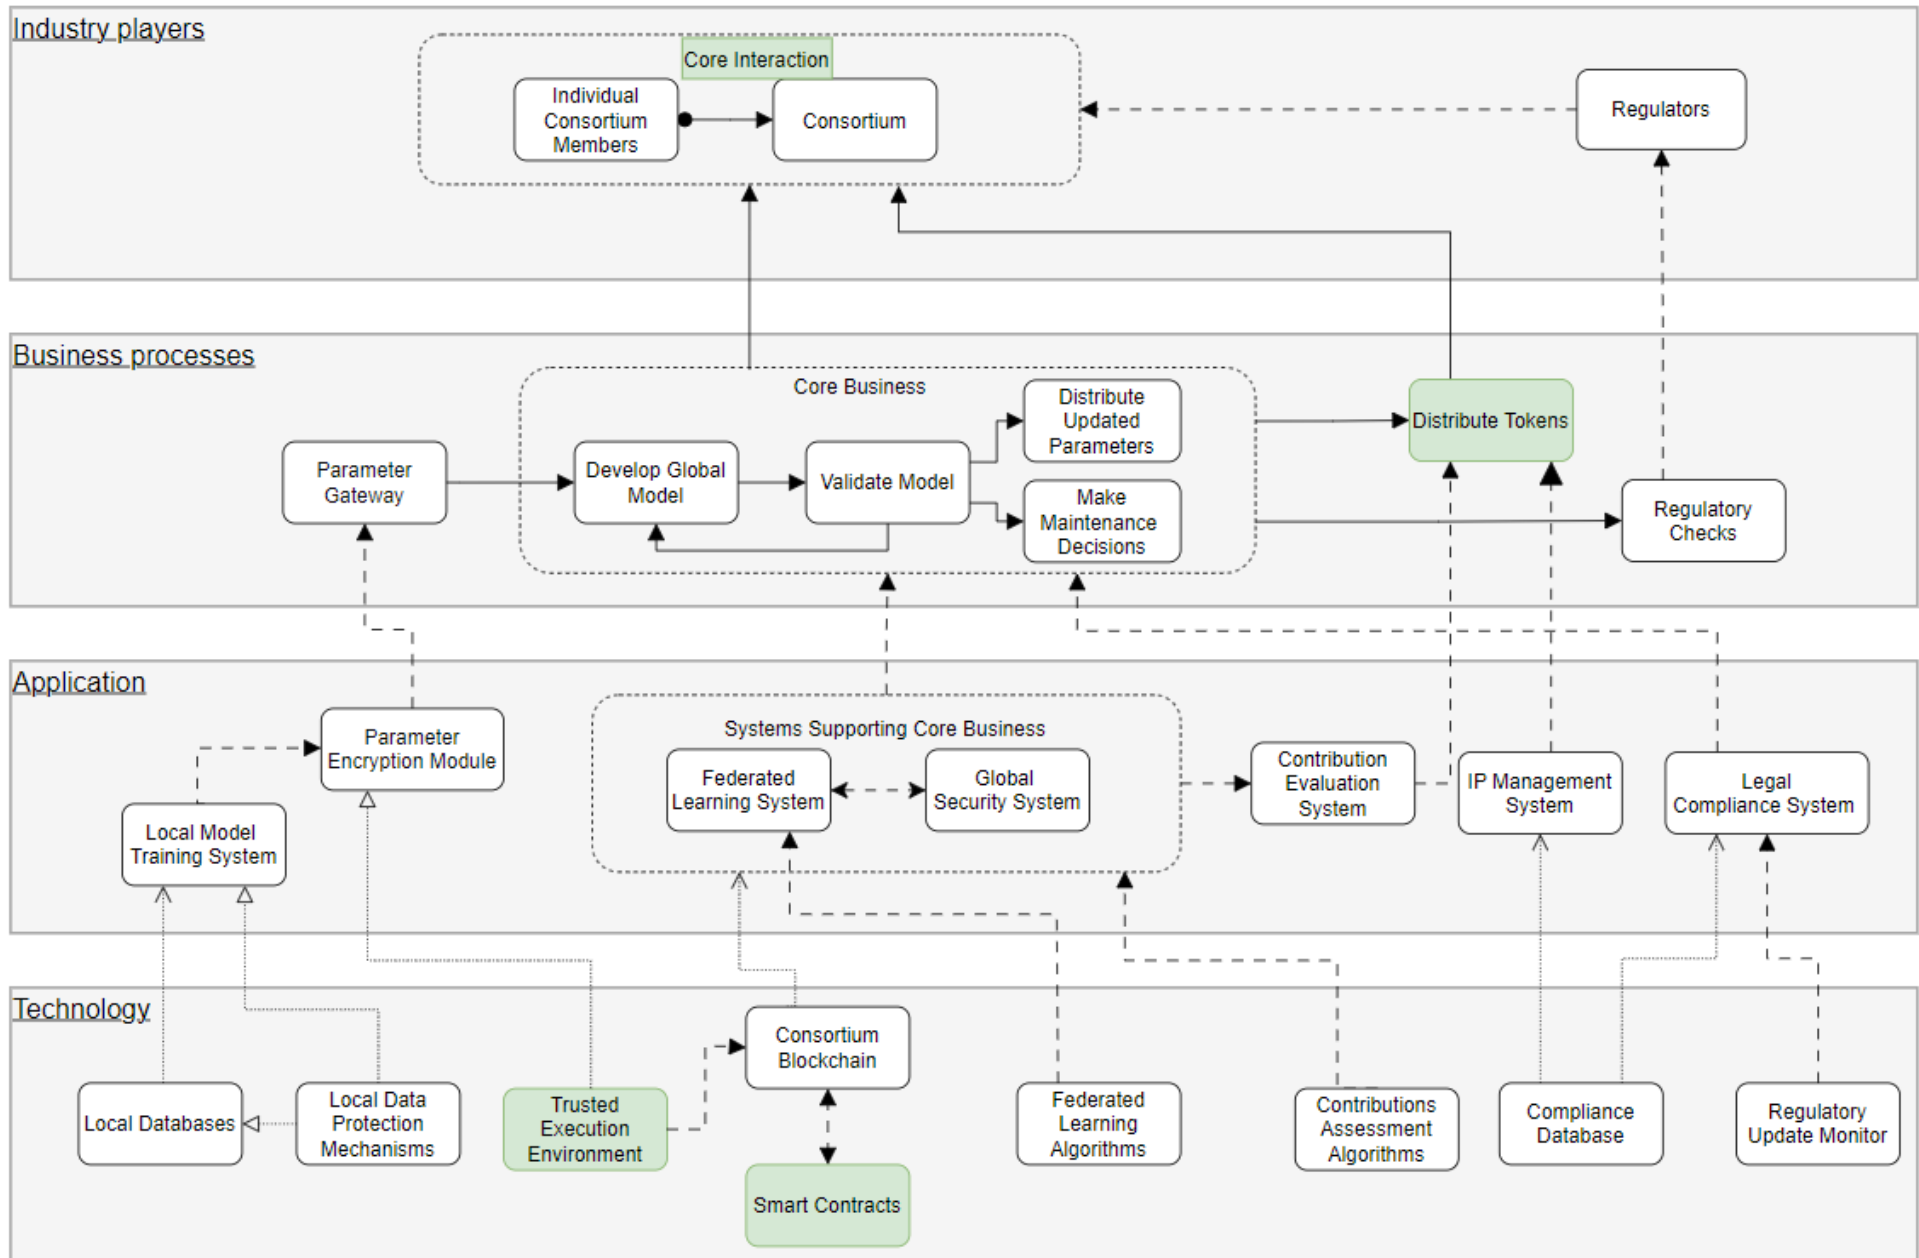


Figure 16: Overview Validated System, Highlighting Differences

7 Discussion

7.1 Key Findings

7.1.1 Necessity of Trust-Building Throughout the Design

Insight

Trust-building is essential throughout the design process, such as leveraging traditional methods like legal contracts and appointing a consortium as a neutral party.

Implications

Trust is fundamental for FL adoption, necessitating transparent governance and the involvement of a neutral, trusted entity to oversee the system. Establishing trust through legal frameworks and transparent operations is needed to gain stakeholder engagement and ensure long-term collaboration.

Findings

Consortium Blockchain: Enhances transparency in distribution allocation and consortium operations. Blockchain technology ensures that all transactions, data exchanges and updates are recorded on an immutable ledger. This transparency allows stakeholders to audit the processes, verify data integrity and ensure that resources and benefits are allocated fairly.

Decentralized FL Nature: This system maintains individual data sovereignty while enabling collective intelligence. Each participant retains control over their data, which is processed locally. Only aggregated insights are shared, ensuring compliance with data privacy regulations like GDPR.

Legal Contracts: These contracts provide a familiar and understandable method of ensuring data protection and adherence to agreed-upon terms. They can specify data usage limitations, confidentiality obligations, and penalties for breaches, providing a legal framework that reinforces trust.

7.1.2 Model Accuracy Prioritized Over Privacy

Insight

The aviation industry prioritizes model accuracy over enhanced privacy measures despite commercial mistrust and skepticism among stakeholders.

Implications

FL systems must ensure that privacy measures do not degrade model performance, maintaining the high standards required in aviation. Achieving this results in successfully adopting FL technologies in environments where accuracy is critical.

Findings

Rejection of Differential Privacy: Stakeholders rejected solutions like differential privacy that could reduce model accuracy, highlighting the industry's commitment to safety and operational efficiency. Differential privacy introduces noise to the data to protect individual privacy, but this can degrade the accuracy of predictive models, which is unacceptable in an industry where precision is critical.

TEE: The focus remains on measures that do not compromise accuracy while preserving privacy. TEE allows computations on encrypted data without decrypting, preserving privacy while maintaining model accuracy. This method ensures that sensitive parameters remain secure throughout the analysis process.

Legal Agreements: Building trust through non-technical solutions like legal agreements is vital. These agreements can enforce strict data protection protocols and outline each party's responsibilities, providing a legal safeguard that complements technical measures.

7.1.3 Equitable Benefit Sharing Through Contribution-Based Access

Insight

Equitable benefit sharing is crucial, and an ethical aviation-proof way to achieve this is by distributing usage rights based on contribution scores.

Implications

Implementing a token-based system for model access ensures sustained motivation and collaboration, crucial for the system's long-term success and fairness. This method aligns incentives with contributions, creating a collaborative environment where all participants benefit from their engagement.

Findings

Differentiated Access: Differentiated access to the best model based on contribution scores incentivizes high-quality data contributions and maintains fairness in benefit distribution. Participants earn tokens based on their data contributions' quality, quantity, and timeliness. These tokens can be redeemed to access the predictive models, ensuring that those contributing more receive more significant benefits.

Preventing Free-Riding: This approach prevents free-riding by ensuring that participants must contribute to gain access to the system's benefits. It aligns with commercial interests by providing tangible rewards proportional to contributions, encouraging continuous engagement and high-quality data sharing.

Transparency and Fairness: The token-based system for model access also supports transparency and fairness. All transactions and token allocations can be recorded on a consortium blockchain, allowing participants to verify their contributions and rewards. This system ensures that benefits are distributed equitably, maintaining motivation and collaboration among participants.

7.1.4 Collaborative IP Management Under a Consortium Model

Insight

Centralized IP management under a consortium-based model mitigates legal complexities and promotes fair usage but requires clear and enforceable agreements.

Implications

Effective IP management frameworks must be established with clear guidelines, and the advantages of collaboration must be communicated to ensure broad acceptance. By centralizing IP management, the consortium can drive collaboration and innovation and ensure that all participants benefit equitably from their contributions.

Findings

Unified Management: This approach is a departure from traditional practices in the industry, necessitating clear communication of its benefits to gain stakeholder support. Traditionally, companies manage their IP independently, leading to fragmented and sometimes conflicting management practices. A centralized consortium model ensures that all IP is managed under a unified set of rules, reducing legal complexities and ensuring fair usage.

Clear Agreements: Clear and enforceable agreements are critical to this model's success. These agreements should outline the terms of IP sharing, usage rights and revenue distribution from licensed IP. By setting these terms upfront, the consortium can prevent disputes and ensure that all participants benefit fairly from shared IP.

Collaborative Mindset: The industry's transition from competitive to cooperative practices requires a shift in mindset towards a more collaborative attitude. This shift requires stakeholders to see the long-term benefits of collaboration, such as shared innovation, reduced costs and enhanced competitiveness through collective intelligence.

7.1.5 Enhancing Predictive Maintenance with Federated Learning

Insight

FL can significantly enhance predictive maintenance tools by integrating diverse datasets, potentially reducing costs.

Implications

Substantial investments in collaborative frameworks are necessary to realize FL's full potential in predictive maintenance. These investments and frameworks are needed to overcome the initial challenges of developing collaboration and integrating diverse datasets. However, if these collaborations are in place, they could greatly benefit the participating organizations.

Findings

Stakeholder Willingness: As highlighted in the interviews, stakeholders across the aviation maintenance industry are willing to engage with FL systems. By sharing derived insights rather than raw data, stakeholders were inclined to participate. However, actual participation and collaboration are difficult to establish. This collaborative innovation leverages shared insights to improve predictive maintenance without compromising competitive advantages.

Improved Accuracy: Integrating diverse datasets can improve predictive models' accuracy, significantly benefiting maintenance practices. By aggregating data from multiple sources, FL can develop more advanced models that accurately predict maintenance needs, reducing the frequency of unscheduled repairs and improving operational reliability.

Operational Benefits: Adopting FL greatly enhances maintenance scheduling and asset management. By accurately assessing the health status of aircraft components, FL significantly reduces the frequency of unscheduled repairs, improving operational reliability and reducing downtime. Furthermore, precision in maintenance forecasting allows for a more streamlined inventory management approach, leading to significant cost savings by reducing excess inventory and storage requirements.

7.1.6 Advancements Needed in Explainability and Transparency

Insight

While Shapley values provide an initial framework for understanding contributions to model outcomes, there remained a substantial need for advancements in explainability to enhance trust in FL systems. This advancement has been partly solved by using TEE instead of HE, however, there is still a need for further explainability.

Implications

Further research and development are needed to refine explainability methods, ensuring they meet the operational needs of the aviation maintenance industry and comply with emerging regulations regarding AI/ML transparency. Enhancing explainability is crucial for building trust in FL systems and ensuring their successful adoption in high-stakes environments.

Findings

Decision-Making Clarity: The lack of clarity in AI/ML decision-making processes poses significant challenges in high-stakes environments like aviation. Current regulatory frameworks lag, providing little guidance on how explainability should be integrated into ML systems.

Shapley Values: Using techniques such as Shapley values can enhance transparency by quantitatively attributing the prediction outcome to its various inputs. This method helps stakeholders understand how different data points contribute to the model's predictions.

Need for Standards: Shapley's values alone may not suffice. Additional standards and metrics for AI explainability tailored to complex decision-making systems like FL are strongly needed. These standards should ensure all stakeholders can easily understand and trust the model's decisions.

Regulatory Involvement: Involving regulatory bodies from the start in the development process of AI models could help shape AI systems to meet explainability standards without restraining innovation. This involvement would ensure the models are developed with compliance in mind, addressing regulatory concerns from the outset.

7.2 Practical Implications

7.2.1 Establishing Trust Through Leadership and Governance

The consortium should act as the central coordinator for updating FL models and overseeing the integration and alignment of FL technologies. However, it should also mediate between stakeholders and resolve conflicts. Clearly defining roles, responsibilities and authority within the governance structure is needed to build transparency and trust. Comprehensive guidelines should support the deployment of the initial setup and ongoing management of the FL system. Performance metrics should be established to evaluate the FL system's effectiveness and guide decisions on system improvements. A phased integration plan could ensure a smooth transition and integration of FL technology across the consortium.

7.2.2 Prioritizing Model Accuracy While Ensuring Privacy

The consortium should implement TEE in the FL system to secure data during aggregation and analysis. This could ensure that sensitive data remains encrypted throughout the learning process. Integrating a consortium blockchain within the FL architecture will provide a secure, transparent and unchangeable record of all operations, benefiting security and trust. The blockchain will ensure all records are tamper-proof and allow consortium members to trace all changes and transactions. This will also reduce the risk of a single point of failure by distributing control among members.

7.2.3 Ensuring Fair Benefit Sharing Through Token-Based Access

Implementing a token-based system for model access will ensure sustained motivation and collaboration. Participants will earn tokens based on their data contributions, which can be redeemed to access the predictive models. This approach will prevent free-riding and align with commercial interests by providing rewards proportional to contributions. Recording all transactions and token allocations on a consortium blockchain will support transparency and fairness, ensuring benefits are distributed equitably and maintaining motivation among participants.

7.2.4 Managing Intellectual Property Through a Consortium Model

The consortium should establish clear IP guidelines where the consortium holds IP rights. Data contributors should be granted the ability to use the developed models through a licensing system based on their contribution levels. This approach will ensure equitable access to IP and promote fair usage. A centralized IP management approach will reduce legal complexities and ensure fair usage. The consortium needs to make clear and enforceable agreements outlining the terms of IP usage rights. Communicate the benefits, gain stakeholder support, and create a cooperative environment encouraging ongoing participation and innovation.

7.2.5 Leveraging Federated Learning for Predictive Maintenance

The consortium should encourage stakeholders across the aviation maintenance industry to engage with FL systems. Actively facilitating participation and collaboration will leverage shared insights for improving predictive maintenance. Integrating diverse datasets will improve predictive models' accuracy, reduce the frequency of unscheduled repairs and enhance operational reliability. Adopting FL will enhance maintenance scheduling and asset management, significantly saving costs by reducing excess inventory and storage requirements. Additionally, the consortium could keep parties engaged by providing ongoing training and support, ensuring they are well-equipped to utilize FL technologies effectively.

7.2.6 Enhancing Explainability and Transparency in AI/ML Models

The consortium should prioritize research to improve the explainability of AI/ML models used in FL, focusing on developing methods that provide clear insights into how models process data and arrive at predictions. This includes developing tools and methodologies that clarify AI/ML decision-making processes. Additionally, the consortium should actively follow AI/ML explainability advancements to stay updated with the latest developments and best practices. Establish new standards and metrics for AI explainability tailored to FL systems to ensure stakeholders can easily understand and trust the model's decisions. Engage regulatory bodies early in the AI model development process to proactively ensure compliance and address regulatory concerns.

7.3 Literature Implications

7.3.1 Trust and Governance in Decentralized Systems

This research provides valuable insights into the mechanisms for building trust in decentralized systems, a critical aspect of collaborative frameworks. It offers evidence of the effectiveness of governance structures and blockchain technology in maintaining trust and transparency within such systems. The study advances theoretical models of trust and governance by demonstrating the practical implementation of these mechanisms. This contribution is particularly significant as it addresses establishing and maintaining trust in environments where control and data are distributed across multiple parties. The findings suggest that well-designed governance structures and blockchain are essential to ensure accountability and transparency while developing a trustworthy collaborative environment.

7.3.2 Equitable Benefit Sharing in Collaborative Frameworks

Implementing a token-based benefit-sharing system in this research presents a novel contribution that addresses critical issues such as free-riding and fair resource distribution. The study introduces a practical mechanism for equitable benefit sharing that aligns incentives with contributions, enhancing collaboration and engagement. This system promotes sustained participation by ensuring that the benefits of shared efforts are distributed fairly among all contributors in model utilization.

This research fills a gap in the literature on collaborative frameworks where accuracy is paramount by providing a framework for incentivizing high-quality data sharing and active participation. This approach mitigates the free-riding problem and ensures that contributions are recognized and rewarded proportionately.

7.4 Comparison with Existing Literature

Research on FL has centered mainly on theoretical models and algorithmic efficiencies, emphasizing data privacy and system optimization. While most studies focus on solving specific aspects of FL, they often overlook integrating these aspects. Moreover, sector-specific FL research typically targets healthcare or finance, frequently neglecting the comprehensive picture necessary for industries like aviation.

7.4.1 Alignment with Current Research

This study supports several similar points in existing FL literature and predictive maintenance in aviation. Adamopoulou & Daskalakis (2023) and Dinis et al. (2019) show that advanced analytics can improve operational efficiency and cut costs, a conclusion that matches insights gathered from interviews. Daily & Peterson (2017) highlight benefits like better reliability and supply chain effectiveness from using predictive maintenance in aviation, which this research also confirms.

Privacy concerns in FL, noted by L. Li et al. (2020) and Mothukuri et al. (2021), were also observed in this study. The importance of privacy in FL systems is emphasized by Das & Brunschweiler (2019) and the need for privacy-preserving techniques, noted by Lu et al. (2021), were supported by the findings.

7.4.2 Contrasts with Literature

Balancing efficiency and accuracy in privacy protection is a significant challenge in FL systems. Traditional FL research often recommends adding noise to data to enhance privacy. However, this study suggests avoiding any noise addition, emphasizing the need for high accuracy to meet the strict demands of the aviation industry. The study highlights the necessity for precision in predictive maintenance and other safety-critical operations. This approach contrasts with conventional views, underscoring the unique need for precision in the aviation sector.

7.5 Contributions to the Literature

7.5.1 Advancing Federated Learning Applications

This research extends the application of FL beyond commonly studied fields like healthcare and finance, demonstrating its potential in the aviation industry. By addressing sector-specific challenges and requirements, the study contributes to a broader understanding of how FL can be adapted to meet diverse industry needs. The detailed analysis of implementing FL in predictive maintenance within the aviation sector provides a new perspective on FL's versatility and effectiveness in enhancing operational reliability and reducing maintenance costs.

7.5.2 Trust and Governance in Decentralized Systems

The research provides detailed insights into the mechanisms for building and maintaining trust in decentralized systems. By implementing and evaluating governance structures and blockchain technology, the study contributes to theoretical models of trust and governance, offering practical solutions for ensuring accountability and transparency in collaborative frameworks. These findings are significant as they demonstrate how governance and transparent operations can benefit trust, even when control and data are distributed across multiple stakeholders.

7.5.3 Equitable Benefit Sharing in Collaborative Frameworks

Introducing a token-based benefit-sharing system that aligns model usages with contributions represents a novel contribution to the literature. This approach addresses the issue of free-riding and ensures fair distribution of resources, enhancing collaboration and sustained participation in collaborative frameworks. The study's framework for equitable benefit sharing is particularly relevant for sectors prioritizing accuracy and reliability, such as aviation. This research effectively advances the theoretical understanding of managing and incentivizing participation in collaborative systems by providing a clear and practical mechanism for rewarding contributions.

7.6 Limitations of Study

7.6.1 Qualitative Data Collection

Using qualitative data collection methods introduced several limitations to the study:

- **Subjectivity in Data Interpretation:** The primary data for this study was collected through qualitative methods such as interviews and expert sessions. While these methods provide in-depth insights and understanding of contextual nuances, they inherently carry the risk of subjectivity. Interpretation of responses may be influenced by personal biases or preconceived notions about expected outcomes, potentially skewing the data. To mitigate this, a structured interview guide was used to ensure consistency in questions and reduce interviewer bias.
- **Selection Bias in Participant Sampling:** Participants selected for interviews and expert sessions were from a specific subset of the aviation industry (IDCA), which may not represent the broader industry spectrum. This selection bias could limit the generalizability of the findings. Participants who responded to invitations were already active within IDCA, which indicates a positive attitude toward data sharing. Structured interviews were used to challenge and rigorously verify their perspectives to address this bias.
- **Limited Diversity Professional Levels:** The study primarily engaged with senior professionals and decision-makers within organizations. This focus may overlook the perspectives of mid-level technicians and other operational staff interacting with FL systems more directly and frequently. Their insights are needed to understand the system's practical implementation and challenges comprehensively. Future research should aim to include a more diverse range of respondents to capture these perspectives.
- **Limited Diversity Geographic Representation:** The geographic distribution of respondents was mostly limited to Europe and the US. This narrow focus may not capture the diverse challenges and perspectives from other regions, potentially limiting the findings' applicability to a global context. Expanding the geographic scope in future studies could provide a more comprehensive understanding of regional differences and enhance the generalizability of the results.

7.6.2 Design Science Research Methodology

The DSR methodology involves systematically designing, developing and analyzing innovative solutions through prototyping, testing and refinement cycles. This method suits engineering and IT projects that aim to create functional and efficient systems.

Problem Identification: This study's initial phase of DSR involved identifying specific challenges within data sharing in aviation maintenance that could be addressed with FL. This included issues like data privacy, willingness to participate and accuracy. One limitation of using DSR is the complexity of accurately identifying and defining problems in highly specialized fields like aviation maintenance. Misidentification can lead to designs that do not fully address core issues or meet user needs.

Design and Development: Based on the identified problems, conceptual designs of FL solutions were developed. These designs considered aviation maintenance's unique regulatory, operational and technical contexts. Several versions of FL models were then developed, incorporating features such as privacy-preserving algorithms and data-sharing protocols. A limitation is that these designs remain theoretical without real-world testing and might not address all practical concerns effectively. Additionally, the conceptual nature of the artifact meant that the evaluation was limited to theoretical and expert assessments, which may not capture all the practical challenges and usability issues.

Testing and Evaluation: Prototypes ideally undergo testing in simulated environments or limited real-world settings with participating organizations. However, this research was more conceptual and experts conducted the evaluation. This approach, while valuable, does not provide the same level of practical insight as real-world testing. Additionally, the evaluation might miss out on unforeseen practical issues that only real-world applications can reveal. Future research should aim to move beyond conceptual designs to actual implementations and real-world testing for more practical evaluations.

Iteration: Feedback from the testing phase was used to refine the model. This iterative process continues until the solutions meet stakeholder requirements, ensuring practical applicability and effectiveness. A significant challenge in this research is that the iterative nature of DSR requires substantial resources such as time and expertise. In an ideal situation, the preference would have been to go through another round of validation for the validated model. Future research should allocate more resources to allow for extended iterative cycles and address the time constraints that can limit the thoroughness of the iterative process.

7.6.3 Balance Between Overview and Details

Finding the right balance between offering an overview and detailed specifics proved challenging in this research. This difficulty came from the need to cover a broad scope, including technological, cooperation and legal aspects. Addressing these elements in detail was impractical due to resource constraints and the complexity involved. Consequently, the study adopted a generalized European viewpoint to develop a practical, broadly applicable solution. This approach enabled the creation of a framework relevant across multiple contexts without becoming stuck in the impracticalities of numerous individual cases.

As a result, the study focused on delivering a tangible framework that could serve as a foundational tool to initiate conversations between organizations and potentially lead to industry-wide application. This research measured the openness of various organizations to adopt FL solutions and established a solid groundwork for future implementation across the sector. While this broad scope was necessary, it may have restricted detailed regional or technical insights. Thus, achieving a comprehensive overview while creating a usable, generalized framework presented limitations. Although beneficial for broad applicability, this approach might overlook specific challenges and nuances needed for local adaptations and precise implementations.

7.7 Future Study Recommendations

7.7.1 Expand Geographical Diversity

Future studies should include participants from a broader range of geographical regions to understand the global implementation challenges and diverse perspectives in FL systems. The current study's focus on Europe and the US may have limited its findings.

Expanding the geographical scope to include participants from Asia, Africa, South America and other regions will provide insights into regulatory environments, operational practices, and cultural attitudes toward data sharing and technology adoption. This broader perspective allows the development of FL systems that are globally applicable and effective.

7.7.2 Engage Participants Less Inclined to Share Data

Future research should target and include participants who may be less inclined to share data to address the potential positive bias toward data sharing. This includes organizations and individuals with reservations or resistance to data sharing due to privacy, security, or competitive advantage concerns. By engaging these reluctant participants, researchers can better understand the barriers to data sharing and develop strategies to overcome them. This approach will ensure that the findings represent the broader industry's attitudes and behaviors, making the solutions more generalizable and widely accepted.

7.7.3 Quantitative Validation

Conducting a quantitative study could test if the perspectives offered by the professionals in this study are broadly shared across the industry. This can be achieved using surveys and statistical analyses to validate the qualitative findings, ensuring they represent the industry's more comprehensive views. Quantitative validation could provide evidence to support the qualitative insights, enhancing the credibility and generalizability of the findings. This approach will help confirm that the proposed solutions and recommendations are widely applicable and beneficial across the aviation industry.

7.7.4 Developing Explainability Standards

Collaborating with researchers and regulatory bodies to develop standardized metrics and guidelines for AI/ML explainability in FL systems must be prioritized. Explainability is critical for building trust and ensuring compliance with regulatory requirements. Developing and implementing standardized explainability metrics and guidelines will make FL systems more transparent and understandable to all stakeholders. This collaboration will help ensure that the explainability standards are practical and widely accepted, enhancing the credibility and trustworthiness of FL systems in high-stakes industries like aviation.

7.7.5 Economic Impact Analysis

Conducting rigorous economic impact studies can provide compelling evidence of the financial benefits of adopting FL systems in aviation maintenance. Detailed cost-benefit analyses should assess the cost implications and potential financial gains from implementing FL systems. These analyses will help persuade stakeholders of the economic viability and benefits of participating in FL initiatives. Additionally, studies projecting the long-term financial outcomes of FL adoption can support decision-making processes within organizations by highlighting the potential for significant cost savings and efficiency gains over time. Understanding the economic impact is relevant for gaining stakeholder engagement and ensuring sustained investment in FL systems.

7.7.6 Engage Technical Staff

Involving mid-level technicians and operational staff in developing and refining FL systems is helpful. These individuals interact directly with maintenance processes and can provide practical insights into system design and functionality. Their feedback can highlight real-world challenges, usability issues, and potential improvements that may not be apparent to senior professionals or decision-makers. Future research should include these perspectives to create FL systems that are theoretically sound, practically viable, and user-friendly.

7.7.7 Real-World Testing

Future studies should progress from conceptual designs to implementing and testing FL systems in real-world environments. Deploying FL systems in aviation maintenance settings to observe their performance under natural operational conditions will help identify practical challenges, usability issues, and unforeseen obstacles that may not be apparent in simulated environments or theoretical assessments. The feedback from real-world testing will be invaluable for refining and improving the system, ensuring its effectiveness and reliability in practical applications.

7.8 Concluding Remarks

7.8.1 Contribution to the Field

This study has explored the potential of FL systems in the aviation industry, emphasizing the importance of trust-building, equitable benefit sharing, centralized IP management, and enhancing predictive maintenance tools. The findings highlight the need for transparent governance, reliable privacy measures, and collaborative frameworks for successful FL adoption. This research contributes to the literature by advancing theoretical models of trust and governance in decentralized systems, proposing practical mechanisms for equitable benefit sharing, and demonstrating FL applications in a high-stakes industry.

7.8.2 Influence on Future Research

Despite the limitations inherent in qualitative data collection and the use of the DSR methodology, this study provides a solid foundation for future research. Recommendations for future studies include:

- **Expanding Geographical Diversity:** Participants from more diverse geographical regions must be included to gain a global perspective on FL implementation challenges and perspectives. This will ensure that the developed systems are globally applicable.
- **Quantitative Validation:** Conducting surveys and statistical analyses to test if the professionals' perspectives in this study are broadly shared across the industry. This will provide evidence supporting the qualitative insights.
- **Developing Explainability Standards:** Collaborating with researchers and regulatory bodies to develop standardized metrics and guidelines for AI/ML explainability in FL systems. Standardized explainability metrics and guidelines will make FL systems more transparent and understandable to all stakeholders, enhancing their credibility and trustworthiness in high-stakes industries.

7.8.3 Implications for Policy and Practice

The findings have implications for policy-making and practice, particularly in calling for action regarding developing regulations supporting FL's secure and ethical use in aviation. The study emphasizes the need to establish regulations around the explainability of AI/ML models. Collaborative efforts are needed with regulatory bodies to develop standardized metrics and guidelines for AI/ML explainability in FL systems. These standards will enhance transparency and make FL systems more understandable to all stakeholders, boosting their credibility and trustworthiness in high-stakes industries like aviation.

7.8.4 Broader Impact

The broader impact of this study lies in creating a more collaborative and data-driven aviation maintenance industry. Optimizing resource usage and maintenance scheduling improves the reliability and safety of aviation operations and enhances the industry's overall sustainability. Moreover, the principles and systems designed for FL in aviation can serve as a model for other sectors. This can facilitate data sharing among competing companies, leading to more efficient and sustainable industry practices.

In conclusion, this research enlightens on the complex interplay between technology, policy, and industry practices that drive progress in aviation maintenance. By advancing the discussion on FL and its potential, this study provides a foundation for developing more efficient, reliable, and collaborative practices within the aviation industry. The integration of FL into aviation maintenance is at an early stage, but the findings and recommendations from this study offer a clear path forward. This promises significant improvements in safety, operational efficiency, and industry collaboration in the coming years.

8 Conclusion

8.1 Recapitulation of Research Objectives and Findings

The primary objective of this thesis was to determine and define the architectural features required for accepting an FL system within the aviation maintenance industry. Through a DSR, this investigation integrated literature and qualitative insights from various stakeholders across the aviation sector. The main research question was:

‘What architectural features should be included in the design of the Federated Learning for aircraft predictive maintenance system to be accepted by stakeholders in the aviation industry?’

Key findings included:

- **Trust-Building:** Establishing trust through transparent governance, legal frameworks, and the involvement of a neutral overseeing entity is essential for FL adoption.
- **Model accuracy over privacy:** The aviation industry prioritizes accuracy over privacy measures, requiring solutions that maintain high performance without compromising data security.
- **Equitable benefit sharing:** Implementing a token-based system for model utilization ensures fairness and incentivizes high-quality data contributions.
- **Collaborative IP management:** Centralizing IP management under a consortium model simplifies legal complexities and promotes fair usage.
- **Enhancing predictive maintenance:** FL can improve predictive maintenance by integrating diverse datasets, which enhances model accuracy and operational reliability.
- **Explainability and transparency:** Advancements in explainability methods are necessary to build trust and ensure compliance with emerging AI/ML regulations.

These findings underscore the need for transparent governance, effective privacy measures, and collaborative frameworks to ensure the successful adoption of FL technologies in aviation maintenance.

8.2 Contribution to the Field

This study has explored the potential of FL systems in the aviation industry. It emphasizes the importance of trust-building, equitable benefit sharing, and centralized IP management while prioritizing model accuracy to enhance predictive maintenance tools. This research contributes to the literature by looking into how trust is established by governance in decentralized systems. At the same time, it proposes practical mechanisms for equitable benefit sharing and demonstrates FL applications in a high-stakes industry.

The architectural framework developed in this study facilitates collaboration and adapts to the unique challenges of the aviation maintenance sector. By directly addressing these collaborative challenges, the research provides a pragmatic model that can serve as a starting point for implementations in similar industries where accuracy and collaboration are important. This foundational framework enables stakeholders to initiate conversations and start a dialogue about adopting and integrating FL systems within their organizations and across the industry.

By examining trust and governance, the study offers insights into the practical implementation of FL systems in a decentralized environment. It proposes mechanisms for equitable benefit sharing through a token-based system, ensuring that contributors are fairly rewarded and motivated to participate. This approach mitigates the risk of free-riding and creates a collaborative environment.

8.3 Practical Implications

The practical implications offer a forward-looking perspective on how aviation maintenance can evolve by adopting FL technologies. By creating a data-sharing environment that enhances predictive maintenance capabilities, the FL system can significantly improve the efficiency and reliability of maintenance operations.

This study provides strategic recommendations for implementing FL systems in aviation maintenance. The consortium should act as the central coordinator for updating FL models and overseeing the integration and alignment of FL technologies. This involves mediating between stakeholders, resolving conflicts and ensuring transparent governance. Comprehensive guidelines should support the deployment of the initial setup and ongoing management of the FL system.

TEE can secure data during aggregation and analysis, protecting sensitive data parameters. Integrating a consortium blockchain within the FL architecture will provide a secure, transparent, and immutable record of all operations, enhancing security and trust. The blockchain will also reduce the risk of a single point of failure by distributing control among members.

Furthermore, establishing clear IP guidelines where the consortium holds IP rights is required. Data contributors should be granted the ability to use the developed models through a licensing system based on their contribution levels. This approach will ensure equitable access to IP and promote fair usage. Drafting clear and enforceable agreements outlining the terms of IP sharing, usage rights, and revenue distribution from licensed IP will protect interests and encourage ongoing collaboration.

The research suggests that a token-based system for model access will ensure sustained motivation and collaboration. Participants will earn tokens based on their data contributions, which can be redeemed to access the predictive models. Recording all transactions and token allocations on a consortium blockchain will support transparency and fairness, ensuring benefits are distributed equitably and maintaining motivation among participants.

Implementing the FL system could improve maintenance strategies by allowing data analysis and sharing information without risking losing your data. This could drastically reduce unscheduled maintenance, minimize aircraft downtime, and reduce the overstocking of parts, leading to significant cost savings.

8.4 Limitations and Challenges

The research acknowledges several limitations due to qualitative interviews and the DSR methodology. The qualitative approach faced challenges such as limited diversity in geographic representation and professional level. Furthermore, due to the way the participants were collected, there could be a bias in the positive attitude towards data sharing. These factors raise concerns about the generalizability of the findings across different regions, potentially affecting the perceived neutrality and objectivity of the data.

The DSR methodology, while systematic, introduced complexities in the highly specialized field of aviation maintenance. Its iterative nature requires significant time and expertise, which can be a substantial burden. This need for substantial resources highlights a constraint in the study, as experts' required time and availability limit the thoroughness and depth of the iterative process.

The balance between keeping an overview and going into details was difficult. The study had to balance the difficulty of covering a broad scope, including technological, cooperative, and legal aspects, while needing to address specific issues in detail.

Despite these challenges, the research aimed to develop a practical and actionable framework that could promote a deeper understanding and encourage wider adoption of FL in aviation maintenance and potentially other sectors.

8.5 Future Directions

Expanding the diversity of participants in future studies will enhance the generalization and reliability of the research. This could ensure that the findings reflect a broader spectrum of perspectives. Engaging participants who are less inclined to share data will address potential positive biases and provide a more accurate understanding of barriers to data sharing. Future research should also involve mid-level technical staff, who interact directly with maintenance processes and can provide practical insights into system design and functionality.

Quantitative validation through surveys and statistical analyses could test if the professionals' perspectives in this study are broadly shared across the industry. This could provide evidence supporting the qualitative insights and enhance the credibility and generalizability of the findings.

Developing standardized metrics and guidelines for AI/ML explainability in FL systems is essential. Collaboration with researchers and regulatory bodies could ensure that the explainability standards are practical and widely accepted, increasing the credibility and trustworthiness of FL systems.

Conducting economic impact analyses could provide convincing evidence of the financial benefits of adopting FL systems in aviation maintenance. These analyses will help persuade stakeholders of the economic viability and benefits of participating in FL initiatives. Additionally, studies projecting the long-term financial outcomes of FL adoption will support decision-making processes within organizations by highlighting the potential for significant cost savings and efficiency gains over time.

8.6 Final Thoughts

This study has explored the potential of FL systems within the aviation maintenance industry, identifying vital architectural features necessary for their adoption and acceptance. By prioritizing trust-building, model accuracy, equitable benefit sharing, centralized IP management, and the enhancement of predictive maintenance tools, the research provides a comprehensive framework for FL implementation.

The findings underscore the importance of transparent governance, reliable privacy measures, and collaborative frameworks to ensure the successful adoption of FL technologies. This research uses trust-building features in governance and decentralized systems, proposing a practical mechanism for equitable benefit sharing, and demonstrates the application of FL in aviation. Despite limitations related to qualitative data collection and the DSR methodology, the study offers a foundation for future research.

Future investigations should focus on expanding geographical diversity, engaging participants who are less inclined to share data, and involving mid-level technical staff. Quantitative validation, real-world testing, and the development of explainability standards are crucial for advancing FL systems' practical application and acceptance. Conducting rigorous economic impact analyses will further demonstrate the financial viability of FL initiatives.

This research highlights FL's transformative potential in aviation maintenance. The insights gained can guide the development and implementation of FL systems, enhancing the aviation industry's operational efficiency and safety. As the journey of integrating FL into aviation maintenance begins, the findings and recommendations from this study offer a clear path forward, promising significant improvements in safety, operational efficiency and industry collaboration in the coming years.

Literature List

- Adamopoulou, E., & Daskalakis, E. (2023). Applications and Technologies of Big Data in the Aerospace Domain. In *Electronics (Switzerland)* (Vol. 12, Issue 10). MDPI. <https://doi.org/10.3390/electronics12102225>
- Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020). Federated Learning: A Survey on Enabling Technologies, Protocols, and Applications. In *IEEE Access* (Vol. 8, pp. 140699–140725). Institute of Electrical and Electronics Engineers Inc. <https://doi.org/10.1109/ACCESS.2020.3013541>
- Antunes, R. S., Da Costa, C. A., Küderle, A., Yari, I. A., & Eskofier, B. (2022). Federated Learning for Healthcare: Systematic Review and Architecture Proposal. *ACM Transactions on Intelligent Systems and Technology*, 13(4). <https://doi.org/10.1145/3501813>
- ArchiMate. (n.d.). *ArchiMate Relationships*. Retrieved June 9, 2024, from <https://pubs.opengroup.org/architecture/archimate31-doc/chap05.html>
- Belderbos, R., Cassiman, B., Faems, D., Leten, B., & Van Looy, B. (2014). Co-ownership of intellectual property: Exploring the value-appropriation and value-creation implications of co-patenting with different partners. *Research Policy*, 43(5), 841–852. <https://doi.org/10.1016/j.respol.2013.08.013>
- Both, L. E. (2021). *WILLINGNESS TO SHARE PERSONAL INFORMATION*. 438–440. <https://doi.org/10.36315/2021inpact096>
- Chen, Y., Chen, S., Liang, J., Feagan, L. W., Han, W., Huang, S., & Wang, X. S. (2020). Decentralized data access control over consortium blockchains. *Information Systems*, 94, 101590. <https://doi.org/10.1016/j.is.2020.101590>
- Daily, J., & Peterson, J. (2017). Predictive Maintenance: How Big Data Analysis Can Improve Maintenance. In *Supply Chain Integration Challenges in Commercial Aerospace* (pp. 267–278). Springer International Publishing. https://doi.org/10.1007/978-3-319-46155-7_18
- Das, A., & Brunschwiler, T. (2019). Privacy is What We Care About. *Proceedings of the First International Workshop on Challenges in Artificial Intelligence and Machine Learning for Internet of Things*, 39–42. <https://doi.org/10.1145/3363347.3363365>
- Devaux, E. (2022, December 21). *What is Differential Privacy: definition, mechanisms, and examples*. <https://www.staticice.ai/post/what-is-differential-privacy-definition-mechanisms-examples>
- Dinis, D., Barbosa-Póvoa, A., & Teixeira, Â. P. (2019). Valuing data in aircraft maintenance through big data analytics: A probabilistic approach for capacity planning using Bayesian networks. *Computers and Industrial Engineering*, 128, 920–936. <https://doi.org/10.1016/j.cie.2018.10.015>
- Domingo-Ferrer, J., Blanco-Justicia, A., Manjon, J., & Sanchez, D. (2022). Secure and Privacy-Preserving Federated Learning via Co-Utility. *IEEE Internet of Things Journal*, 9(5), 3988–4000. <https://doi.org/10.1109/JIOT.2021.3102155>
- Du, W., & Atallah, M. J. (2001). Secure multi-party computation problems and their applications. *Proceedings of the 2001 Workshop on New Security Paradigms*, 13–22. <https://doi.org/10.1145/508171.508174>
- Duc Nguyen, V., Kefalas, M., Yang, K., Apostolidis, A., Olhofer, M., Limmer, S., & Bäck, T. (2019). *A Review: Prognostics and Health Management in Automotive and Aerospace*. <https://www.klm.com/corporate/en/publications/2015>
- Dwork, C., & Roth, A. (2013). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–487. <https://doi.org/10.1561/04000000042>

- EASA. (2021). *CAMO (Continuing Airworthiness Management Organisation)*. <https://www.easa.europa.eu/en/the-agency/faqs/camo-continuing-airworthiness-management-organisation>
- EASA. (2023). *EASA concept paper guidance for level 1and2 machine learning applications proposed issue 02 feb2023*.
- Eckhardt, J., Vogelsang, A., & Fernández, D. M. (2016). Are non-functional requirements really non-functional? an investigation of non-functional requirements in practice. *Proceedings - International Conference on Software Engineering, 14-22-May-2016*, 832–842. <https://doi.org/10.1145/2884781.2884788>
- EU. (n.d.-a). *Art. 46 GDPR*. Retrieved April 29, 2024, from <https://gdpr-info.eu/art-46-gdpr/>
- EU. (n.d.-b). *Federated Learning*. Retrieved April 28, 2024, from https://www.edps.europa.eu/press-publications/publications/techsonar/federated-learning_en
- European Commission. (n.d.-a). *Digital Market Act*. Retrieved April 29, 2024, from https://digital-markets-act.ec.europa.eu/about-dma_en
- European Commission. (n.d.-b). *Digital Service Act*. Retrieved April 29, 2024, from https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act_en
- European Commission. (2024a, January 25). *Data Governance Act*. <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act-explained>
- European Commission. (2024b, March 6). *AI act*. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- European Data Protection Board. (n.d.). *International Data Transfers*. Retrieved April 29, 2024, from https://www.edpb.europa.eu/sme-data-protection-guide/international-data-transfers_en#toc-4
- FAA. (2024). *National Aviation Research Plan (NARP)*. <http://www.faa.gov/go/narp>
- Fan, Z., Fang, H., Zhou, Z., Pei, J., Friedlander, M. P., Liu, C., & Zhang, Y. (2022). Improving Fairness for Data Valuation in Horizontal Federated Learning. *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, 2440–2453. <https://doi.org/10.1109/ICDE53745.2022.00228>
- Financial Times. (2023, December 14). *Air France-KLM 2023 Investor Day*.
- Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence information and basic countermeasures. *Proceedings of the ACM Conference on Computer and Communications Security, 2015-October*, 1322–1333. <https://doi.org/10.1145/2810103.2813677>
- Galvin, R. (2015). How many interviews are enough? Do qualitative interviews in building energy consumption research produce reliable knowledge? *Journal of Building Engineering*, 1, 2–12. <https://doi.org/10.1016/j.jobe.2014.12.001>
- Geppert, T., Deml, S., Sturzenegger, D., & Ebert, N. (2022). Trusted Execution Environments: Applications and Organizational Challenges. In *Frontiers in Computer Science* (Vol. 4). Frontiers Media S.A. <https://doi.org/10.3389/fcomp.2022.930741>
- Google Cloud. (2024, May 9). *Confidential VM overview*.
- Gorbatyuk, A. (2020). The Allocation of Patent Ownership in R&D Partnerships: Default Rules v. Contractual Practices. *SCRIPT-Ed*, 17(1), 4–53. <https://doi.org/10.2966/scrip.170120.4>
- Hennink, M., & Kaiser, B. N. (2022). Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social Science and Medicine*, 292. <https://doi.org/10.1016/j.socscimed.2021.114523>

- Hiwatashi, K., Ogura, K., Ohata, S., & Nuida, K. (2021). Accelerating Secure (2+1)-Party Computation by Insecure but Efficient Building Blocks. *ASIA CCS 2021 - Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, 616–627. <https://doi.org/10.1145/3433210.3453109>
- IDCA. (n.d.-a). <https://www.dataforaviation.org/chapter-3>. Retrieved May 28, 2024, from <https://www.dataforaviation.org/chapter-3>
- IDCA. (n.d.-b). <https://www.dataforaviation.org/mission>. Retrieved May 28, 2024, from <https://www.dataforaviation.org/mission>
- Jonkers, H., Proper, E., Lankhorst, M. M., Quartel, D. A. C., & Iacob, M. E. (2011). ArchiMate® for integrated modelling throughout the architecture development and implementation cycle. *Proceedings - 13th IEEE International Conference on Commerce and Enterprise Computing, CEC 2011*, 294–301. <https://doi.org/10.1109/CEC.2011.52>
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., Rouayheb, S. El, Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2019). *Advances and Open Problems in Federated Learning*. <http://arxiv.org/abs/1912.04977>
- Kohlbrener, D., Shinde, S., Lee, D., Asanovic, K., & Song, D. (2020). Building open trusted execution environments. *IEEE Security and Privacy*, 18(5), 47–56. <https://doi.org/10.1109/MSEC.2020.2990649>
- Korvesis, P., Besseau, S., & Vazirgiannis, M. (2018). Predictive maintenance in aviation: Failure prediction from post-flight reports. *Proceedings - IEEE 34th International Conference on Data Engineering, ICDE 2018*, 1423–1434. <https://doi.org/10.1109/ICDE.2018.00160>
- Kurri, T. (2020). *EASA Approval Process for Aircraft Modifications*.
- Lansari, M., Bellafqira, R., Kapusta, K., Thouvenot, V., Bettan, O., & Coatrieux, G. (2023). *When Federated Learning meets Watermarking: A Comprehensive Overview of Techniques for Intellectual Property Protection*. <http://arxiv.org/abs/2308.03573>
- Li, L., Fan, Y., Tse, M., & Lin, K. Y. (2020). A review of applications in federated learning. *Computers and Industrial Engineering*, 149. <https://doi.org/10.1016/j.cie.2020.106854>
- Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., Liu, X., & He, B. (2023). A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. *IEEE Transactions on Knowledge and Data Engineering*, 35(4), 3347–3366. <https://doi.org/10.1109/TKDE.2021.3124599>
- Liu, Z., Chen, Y., Yu, H., Liu, Y., & Cui, L. (2021). *GTG-Shapley: Efficient and Accurate Participant Contribution Evaluation in Federated Learning*. <http://arxiv.org/abs/2109.02053>
- Lo, S. K., Lu, Q., Zhu, L., Paik, H. Y., Xu, X., & Wang, C. (2022). Architectural patterns for the design of federated learning systems. *Journal of Systems and Software*, 191. <https://doi.org/10.1016/j.jss.2022.111357>
- Lu, Z., Kazi, R. H., Wei, L. Y., Dontcheva, M., & Karahalios, K. (2021). A Systematic Literature Review on Federated Machine Learning: From A Software Engineering Perspective. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1). <https://doi.org/10.1145/1122445.1122456>
- Malterud, K. (2012). Systematic text condensation: A strategy for qualitative analysis. *Scandinavian Journal of Public Health*, 40(8), 795–805. <https://doi.org/10.1177/1403494812465030>
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2016). *Communication-Efficient Learning of Deep Networks from Decentralized Data*. <http://arxiv.org/abs/1602.05629>
- Microsoft. (2023, October 26). *TEE Microsoft*. <https://learn.microsoft.com/nl-nl/azure/confidential-computing/trusted-execution-environment>

- Mothukuri, V., Parizi, R. M., Pouriye, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115, 619–640. <https://doi.org/10.1016/j.future.2020.10.007>
- Myalil, D., Rajan, M. A., Apte, M., & Lodha, S. (2021). Robust Collaborative Fraudulent Transaction Detection using Federated Learning. *Proceedings - 20th IEEE International Conference on Machine Learning and Applications, ICMLA 2021*, 373–378. <https://doi.org/10.1109/ICMLA52953.2021.00064>
- Niu, C., Wu, F., Tang, S., Hua, L., Jia, R., Lv, C., Wu, Z., & Chen, G. (2020). Billion-scale federated learning on mobile clients: A submodel design with tunable privacy. *Proceedings of the Annual International Conference on Mobile Computing and Networking, MOBICOM*, 405–418. <https://doi.org/10.1145/3372224.3419188>
- Nohara, Y., Matsumoto, K., Soejima, H., & Nakashima, N. (2022). Explanation of machine learning models using shapley additive explanation and application for real data in hospital. *Computer Methods and Programs in Biomedicine*, 214. <https://doi.org/10.1016/j.cmpb.2021.106584>
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Silhavy, R., Silhavy, P., Prokopova, Z., Senkerik, R., Kominkova, Z., & Editors, O. (2017). *Advances in Intelligent Systems and Computing 575 Software Engineering Trends and Techniques in Intelligent Systems* (Vol. 3). <http://www.springer.com/series/11156>
- Torens, C., Durak, U., & Dauer, J. (2022). Guidelines and Regulatory Framework for Machine Learning in Aviation. *AIAA Science and Technology Forum and Exposition, AIAA SciTech Forum 2022*. <https://doi.org/10.2514/6.2022-1132>
- Tseng, Y. F., Fan, C. I., Kung, T. C., Huang, J. J., & Kuo, H. N. (2017). Homomorphic encryption supporting logical operations. *ACM International Conference Proceeding Series, 2017-October*, 66–69. <https://doi.org/10.1145/3145777.3145789>
- vom Brocke, J., Hevner, A., & Maedche, A. (2020). *Introduction to Design Science Research* (pp. 1–13). https://doi.org/10.1007/978-3-030-46781-4_1
- Voss, E. A., Blacketer, C., van Sandijk, S., Moinat, M., Kallfelz, M., van Speybroeck, M., Prieto-Alhambra, D., Schuemie, M. J., & Rijnbeek, P. R. (2024). European Health Data & Evidence Network—learnings from building out a standardized international health data network. *Journal of the American Medical Informatics Association*, 31(1), 209–219. <https://doi.org/10.1093/jamia/ocad214>
- Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Vincent Poor, H. (2020). Federated Learning with Differential Privacy: Algorithms and Performance Analysis. *IEEE Transactions on Information Forensics and Security*, 15, 3454–3469. <https://doi.org/10.1109/TIFS.2020.2988575>
- Xu, X., & Lyu, L. (2020). *A Reputation Mechanism Is All You Need: Collaborative Fairness and Adversarial Robustness in Federated Learning*. <http://arxiv.org/abs/2011.10464>
- Yang, G., & Maskus, K. E. (2001). *Intellectual Property Rights and Licensing: An Econometric Investigation*.
- Zhang, A., & Lin, X. (2018). Towards Secure and Privacy-Preserving Data Sharing in e-Health Systems via Consortium Blockchain. *Journal of Medical Systems*, 42(8), 140. <https://doi.org/10.1007/s10916-018-0995-5>
- Zheng, Z., Xie, S., Dai, H.-N., Chen, W., Chen, X., Weng, J., & Imran, M. (2020). An overview on smart contracts: Challenges, advances and platforms. *Future Generation Computer Systems*, 105, 475–491. <https://doi.org/10.1016/j.future.2019.12.019>

Appendices

Appendix A: Literature Review

Table 9 contains the sources and keywords used in the literature review.

Table 9: Sources Form the Literature Revie

Main Authors	Title of the Article	Year	Key Topic 1	Key Topic 2	Key Topic 3
Aledhari et al.	Federated Learning: A Survey on Enabling Technologies, Protocols and Applications	2020	Federated Learning	Architecture	-
Antunes et al.	Federated Learning for Healthcare: Systematic Review and Architecture Proposal	2022	Federated Learning	Architecture	Healthcare
Das & Brunschwiler	Privacy is What We Care About	2019	Federated Learning	Acceptance	Privacy
Dinis et al.	Valuing data in aircraft maintenance through big data analytics: A probabilistic approach for capacity planning using Bayesian networks	2019	Federated Learning	Predictive Maintenance	-
Domingo-Ferrer et al.	Secure and Privacy-Preserving Federated Learning via Co-Utility	2022	Federated Learning	Acceptance	Privacy & Security
Duc Nguyen et al.	A Review: Prognostics and Health Management in Automotive and Aerospace	2019	Federated Learning	Predictive Maintenance	Aviation
Kairouz et al.	Advances and Open Problems in Federated Learning	2019	Federated Learning	-	-
L. Li et al.	A review of applications in Federated Learning	2020	Federated Learning	Acceptance	Privacy & Security
Lo et al.	Architectural patterns for the design of Federated Learning systems	2022	Federated Learning	Architecture	-
Lu et al.	A Systematic Literature Review on Federated Machine Learning: From A Software Engineering Perspective	2021	Federated Learning	Acceptance	Privacy

McMahan et al.	Communication-Efficient Learning of Deep Networks from Decentralized Data	2016	Federated Learning	-	-
Monthukuri et al.	A survey on security and privacy of Federated Learning	2021	Federated Learning	Acceptance	Privacy & Security
Myalil et al.	Robust Collaborative Fraudulent Transaction Detection using Federated Learning	2019	Federated Learning	Banking	-
Niu et al.	Billion-scale federated learning on mobile clients: A submodel design with tunable privacy	2020	Federated Learning	Privacy	-
Q Li et al.	A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection	2023	Federated Learning	Privacy	-
Torens et al.	Guidelines and Regulatory Framework for Machine Learning in Aviation	2022	Federated Learning	Predictive Maintenance	Aviation
Voss et al.	European Health Data & Evidence Network—learnings from building out a standardized international health data network	2024	Federated Learning	Healthcare	-

Appendix B: Informed Consent

You are being invited to participate in a research study titled ‘Stimulating data sharing in the aviation industry by using Federated Learning. This study is being done by Bas van de Walle from the TU Delft.

Background and objective of the study:

In the rapidly evolving landscape of aviation maintenance, traditional methods are increasingly being complemented by data-driven approaches. Our research research Federated Learning, a decentralized approach to machine learning that promises to revolutionize predictive maintenance by enabling collaborative yet privacy-preserving data analysis across various stakeholders in the aviation sector. Specifically, our study seeks to understand the requirements, preferences and challenges for creating a platform facilitating information sharing among various stakeholders in the aviation industry, while preserving the confidentiality of raw data.

Some of the questions in the interview will be:

- What types of data are you able to offer?
- What type of data are you interested in from others?
- What is your opinion on using Federated Learning as a tool to share data?

Your participation will involve approximately 40 minutes for an audio recorded interview, with a textual transcript produced afterward. The gathered data will be utilized for research, analysis, user needs assessment and system performance evaluation. We will inquire about your needs, preferences and current challenges regarding information sharing with other entities in the aviation sector and eventually seek validation for the developed system.

As with any online activity the risk of a breach is always possible. To the best of our ability your answers in this study will remain confidential stored on an institutional storage solution at TUD, accessible the research team only. The data will be stored in TUD, Netherlands, governed by GDPR.

We will produce a summary, which we will send to you for review before publication. The summary will be made publicly available as part of the appendices of the MSc thesis, unless you raise any concerns. The summary will include company name and years of experience.

All personal data (transcript + recording) will be deleted at the latest 1 month after the completion of the project.

Your participation in this study is entirely voluntary and you can withdraw at any time. You are free to omit any questions. The data provided can later also be removed and/or not included into the research if you would like to retract your information, as long as the result has not been published.

Responsible researcher: Bas van de Walle

Signatures		
_____	_____	_____
Name of participant [printed]	Signature	Date

I, as researcher, have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

Bas van de Walle

Signature

Date

Appendix C: Interview Questions Regulators

Organizational Context:

- Could you describe your organization's activities related to aviation maintenance?
- What is your role within the organization and what expertise do you bring to aviation maintenance and data management?
- How many years have you worked in this position?
- Where is your company situated?

EASA Concept Paper - Implementation Status:

1. What is the current status of the implementation of the EASA Concept Paper on guidance for Level 1 & 2 machine learning applications?
2. How similar do you expect the final version of the guidelines to be compared to this concept version?

Data Utilization - Data Collection for Predictive Maintenance:

3. What types of data do you collect that could be utilized to develop a predictive maintenance model?
4. In a scenario where it you do not share the actual raw data but you are able to learn from each other and are able to improve predictive maintenance tools. How open is your organization to participating in such a collaborative initiative?

Classification and Application - Classification of Federated Learning Tools:

5. Based on the classification factors of human involvement, decision-making autonomy and task criticality, how would you classify and assess the suitability of a Federated Learning tool for predictive maintenance?

Validation and Safety Standards - Model Validation:

6. How should Federated Learning models be validated to align with EASA's safety standards?
7. What specific performance metrics would you recommend for this validation process, especially considering their role in maintenance decision-making?

Monitoring and Continuous Improvement - System Monitoring:

8. What mechanisms should be established to continuously monitor the efficiency and safety of Federated Learning systems in aviation?
9. How does EASA propose to oversee these systems to ensure they remain reliable and accurate over time?

Infrastructure Requirements:

10. From EASA's perspective, what are the essential elements of technical infrastructure needed to support the secure and efficient operation of Federated Learning systems in aviation?

Evolving Regulatory Landscape - Regulatory Evolution:

11. How do you anticipate the regulatory landscape will evolve to accommodate Federated Learning in aviation?

If obligatory checks are not discussed: Given the scenario where predictive models accurately determine the condition of aircraft parts, how do you foresee the adaptation of current regulations that mandate obligatory checks after a specified amount of flight hours?

12. What future guidelines might EASA consider to support the integration of this technology?

Other - Other objections:

13. In the context of designing Federated Learning systems for aviation maintenance, are there any critical considerations or potential challenges that you foresee which have not yet been addressed in our discussion?

Appendix D: Description Components ArchiMate Model

In Table 10 an overview of all the components in the ArchiMate model can be found.

Table 10: Description Components

Component	Layer	Description
Local Databases	Technology	Stores data locally to ensure privacy and control over sensitive information.
Local Data Protection Mechanisms	Technology	Implements security measures to safeguard data at the local level.
Homomorphic Encryption Implementation	Technology	Uses encryption techniques to enable computations on encrypted data without decryption.
Differential Privacy Implementation	Technology	Adds noise to data to protect individual privacy while preserving the usefulness of the data.
Federated Learning Algorithms	Technology	Provide the algorithm for the Federated Learning system.
Consortium Blockchain	Technology	A form of blockchain technology to create a secure and transparent database.
Contributions Assessment Algorithms	Technology	Provides the algorithm to assess the contributions of each participant.
Compliance Database	Technology	Maintains records of legal and regulatory compliance information.
Regulatory Update Monitor	Technology	Tracks and updates the system with changes in regulations to ensure ongoing compliance.
Local Model Training System	Application	Facilitates the training of local machine learning models using local data.
Parameter Encryption Module	Application	Make the parameters developed by the organizations secure.
Federated Learning System	Application	Manages the overall Federated Learning process.
Global Security System	Application	Implements extensive security protocols to safeguard data and operational processes across the system.
Systems Supporting Core Business	Application	Assists in the execution and maintenance of core business activities within the system.
Contribution Evaluation System	Application	Assesses and quantifies the contributions of each participating entity.

IP Management System	Application	Oversees the management and protection of intellectual property rights within the consortium.
Legal Compliance System	Application	Ensures the system complies with legal standards and adapts to regulatory changes.
Parameter Gateway	Business Processes	It collects the model parameters from the different organizations and provides them as input for developing the model.
Develop Global Model	Business Processes	Aggregates insights from local models to develop a comprehensive global model.
Validate Model	Business Processes	Tests and verifies the global model to ensure its accuracy and reliability.
Core Business	Business Processes	Refers to the core business activities that support the system's operation.
Distribute Updated Parameters	Business Processes	Distribute updated model parameters to all participating entities for synchronized improvements.
Make Maintenance Decisions	Business Processes	Utilizes model insights to inform and optimize maintenance decisions.
Distribute Rewards	Business Processes	Distributes rewards based on the contributions of each participant.
Regulatory Checks	Business Processes	Performs audits to verify regulatory compliance across the system.
Individual Consortium Members	Industry Players	Identifies the individual entities participating in the consortium.
Consortium	Industry Players	Manages the collaborative group of entities operating under shared governance.
Core Interaction	Industry Players	Refers to the essential interactions and communications among consortium members.
Regulators	Industry Players	Monitors compliance and issues guidelines to ensure legal and safe operations.

Appendix E: Relationships ArchiMate

In Table 11 an overview of all the relationships in the ArchiMate model can be found

Table 11: Relationships ArchiMate

ID	Start Block	Layer Start Block	End Block	Layer End Block	Relationship	Explanation
1	Local Databases	Technology	Local Model Training System	Application	Access	The data is accessed from the local database and used in the local model training system.
2	Local Data Protection Mechanisms	Technology	Local Databases	Technology	Realization	The Local Data Protection Mechanisms play a critical role in the operation of the Local Databases.
3	Local Data Protection Mechanisms	Technology	Local Model Training System	Application	Realization	The Local Data Protection Mechanisms play a critical role in the operation of the Local Model Training Systems.
4	Homomorphic Encryption Implementation	Technology	Parameter Encryption Module	Application	Realization	The Homomorphic Encryption Implementation is needed for the Parameter Encryption Module to handle encrypted data securely.
5	Differential Privacy Implementation	Technology	Parameter Encryption Module	Application	Realization	The Differential Privacy Implementation ensures data privacy within the Parameter Encryption Module.
6	Federated Learning Algorithms	Technology	Federated Learning System	Application	Flow	The algorithm representing the information flows from the Federated Learning Algorithms to the Federated Learning System.
7	Consortium Blockchain	Technology	Systems Supporting Core Business	Application	Access	The Systems Supporting Core Business access the Consortium Blockchain for secure and verified transactions

8	Contributions Assessment Algorithms	Technology	Systems Supporting Core Business	Application	Flow	The Contributions Assessment Algorithms provide results for evaluation by the Systems Supporting Core Business.
9	Compliance Database	Technology	IP Management System	Application	Access	The IP Management System accesses compliance data from the Compliance Database to manage IP-related compliance.
10	Compliance Database	Technology	Legal Compliance System	Application	Access	The Legal Compliance System accesses compliance data from the Compliance Database to ensure regulatory compliance.
11	Regulatory Update Monitor	Technology	Legal Compliance System	Application	Flow	Compliance information flows to the Regulatory Update Monitor for analysis and reporting.
12	Local Model Training System	Application	Parameter Encryption Module	Application	Flow	Parameters flow from the Local Model Training System to the Parameter Encryption Module.
13	Parameter Encryption Module	Application	Parameter Encryption Gateway	Business Processes	Flow	Secured parameters are distributed to the Parameter Encryption Gateway for collection from the different organizations.
14	Federated Learning System	Application	Global Security System	Application	Flow	Security information flows between the Federated Learning System and the Global Security System to ensure data integrity and security.
15	Global Security System	Application	Federated Learning System	Application	Flow	Security information flows to the Federated Learning System to provide necessary security measures.
16	Systems Supporting Core Business	Application	Core Business	Business Processes	Flow	The information from the System Supporting Core Business flows to the Core Business.

17	Systems Supporting Core Business	Application	Contribution Evaluation System	Application	Flow	Information flows to the Contribution Evaluation System for assessment and evaluation of contributions.
18	Contribution Evaluation System	Application	Distribute Rewards	Business Processes	Flow	The Contribution Evaluation System provides results to the Distribute Rewards process.
19	IP Management System	Application	Distribute Rewards	Business Processes	Flow	The IP Management System provides data to the Distribute Rewards process for reward distribution based on IP management.
20	Legal Compliance System	Application	Core Business	Business Processes	Flow	Compliance data flows to the Core Business processes to ensure legal and regulatory compliance.
21	Parameter Gateway	Business Processes	Develop Global Model	Business Processes	Triggering	The Parameter Encryption Gateway triggers the development of the Global Model by providing the secured parameters.
22	Develop Global Model	Business Processes	Validate Model	Business Processes	Triggering	The development of the Global Model triggers the need for validation of the model.
23	Validate Model	Business Processes	Develop Global Model	Business Processes	Triggering	The Validate Model triggers further development of the Global Model if necessary.
24	Validate Model	Business Processes	Distribute Updated Parameters	Business Processes	Triggering	Validation triggers the distribution of updated parameters for further use.
25	Validate Model	Business Processes	Make Maintenance Decisions	Business Processes	Triggering	The validation results trigger maintenance decisions based on the validation outcomes.
26	Core Business	Business Processes	Core Interaction	Business Processes	Triggering	Core business activities trigger interactions within the consortium.

27	Core Business	Business Processes	Distribute Rewards	Business Processes	Triggering	Core business results trigger the distribution of rewards based on business achievements.
28	Core Business	Business Processes	Regulatory Checks	Business Processes	Triggering	Core business processes trigger regulatory checks to ensure compliance
29	Distribute Rewards	Business Processes	Core Interaction	Industry Players	Triggering	The distribution of rewards triggers interactions within the consortium
30	Regulatory Checks	Business Processes	Regulators	Industry Players	Flow	Feedback from regulatory checks flows to the regulators to inform and adjust regulatory standards.
31	Individual Consortium Members	Industry Players	Consortium	Industry Players	Assignment	The Individual Consortium Members are assigned to the Consortium to fulfill specific roles and responsibilities.
32	Regulators	Industry Players	Consortium	Industry Players	Flow	Regulatory requirements and checks flow from the Regulators to the Consortium to ensure compliance with regulations.

Appendix F: Component Contributions to Requirements

In Table 12 an overview of the contribution of each component to the requirements can be found.

Table 12: Overview Component Contributions to Requirements

		Local Databases	Local Data Protection Mechanisms	Homomorphic Encryption	Differential Privacy Implementation	Federated Learning Algorithms	Consortium Blockchain	Contributions Assessment Algorithms	Compliance Database	Regulatory Update Monitor	Local Model Training System	Parameter Encryption Module	Federated Learning System	Global Security System	Systems Supporting Core Business	Contribution Evaluation System	IP Management System	Legal Compliance System	Parameter Gateway	Develop Global Model	Validate Model	Core Business	Distribute Updated Parameters	Make Maintenance Decisions	Distribute Rewards	Regulatory Checks	Individual Consortium Members	Consortium	Core Interaction	Regulators
R1	The system must have strict security measures.	X	X	X	X	-	-	-	-	-	X	X	X	X	-	-	-	-	X	-	-	-	-	-	-	-	X	-	-	-
R2	The system must have a robust data-sharing framework.	-	X	-	-	X	X	-	-	-	X	X	X	X	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
R3	The system must ensure compliance with existing legal standards.	-	-	-	-	-	-	-	X	-	-	-	-	-	-	-	-	X	-	-	-	-	-	-	X	-	-	-	X	
R4	The system must be adaptable to evolving regulations.	-	-	-	-	-	-	-	-	X	-	-	-	-	-	-	-	X	-	-	-	-	-	-	X	-	-	-	X	
R5	The system must develop a consortium-based approach for managing co-created IP.	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	X	-	-	-	-	-	-	-	X	-	-	X	X	-
R6	The system must establish mechanisms for equitable benefit sharing.	-	-	-	-	X	-	X	-	-	-	-	X	-	X	X	-	-	-	X	-	X	X	-	X	-	-	X	X	-
R7	The system must develop transparent processes for proportional value attribution.	-	-	-	-	X	X	X	-	-	-	-	X	-	-	X	-	-	-	-	-	-	-	-	-	-	X	X	-	-
R8	The system must be more accurate than current predictions.	X	X	-	-	X	-	-	-	-	X	X	X	-	-	-	-	-	X	X	X	X	-	X	-	-	-	X	-	-
R9	The system must ensure AI models are explainable and transparent.	X	X	-	-	X	X	-	X	X	X	X	X	-	-	-	-	X	X	X	X	-	-	-	-	X	X	X	X	X
R10	The system must have educational and training programs for stakeholders.	-	-	-	-	-	-	-	-	-	-	-	-	-	X	-	-	-	-	-	-	-	-	-	-	-	X	X	-	-
R11	The system must appoint a trusted, neutral entity for collaborative management and fair governance.	-	-	-	-	-	X	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	X	X	-	

Appendix G: Informative Slides Security

The slides used during the validation interview can be seen here in Figure 17 and Figure 18.

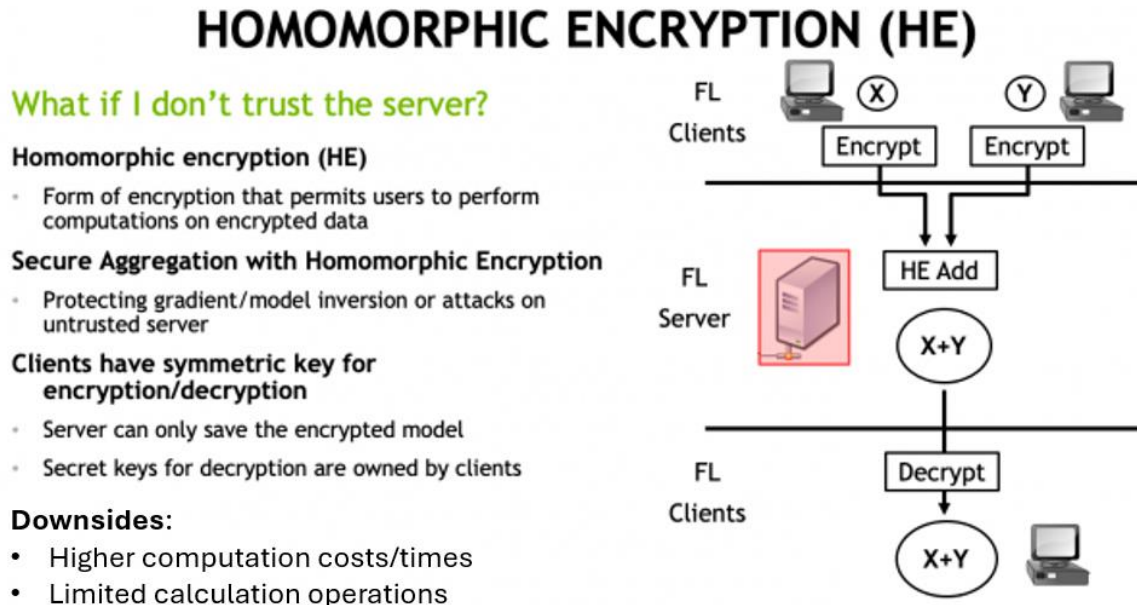


Figure 17: Slide homomorphic encryption (Roth et al., 2021)

Differential Privacy

Benefits:

Noise Addition for Data Protection:

Differential privacy adds controlled noise to data, ensuring individual privacy while allowing statistical analyses.

Feasibility of Analysis: It enables meaningful data analysis without compromising individual privacy, maintaining both research integrity and confidentiality.

Downsides:

Trade-off Between Privacy and Utility:

More noise for greater privacy reduces data accuracy, affecting the usefulness of analytical results.

Lack of Standardization:

The absence of standardized protocols complicates the implementation of differential privacy, leading to inconsistent adoption and potential security risks.




Figure 18: Slide differential privacy (Devaux, 2022)

Appendix H: Overview Validated Model

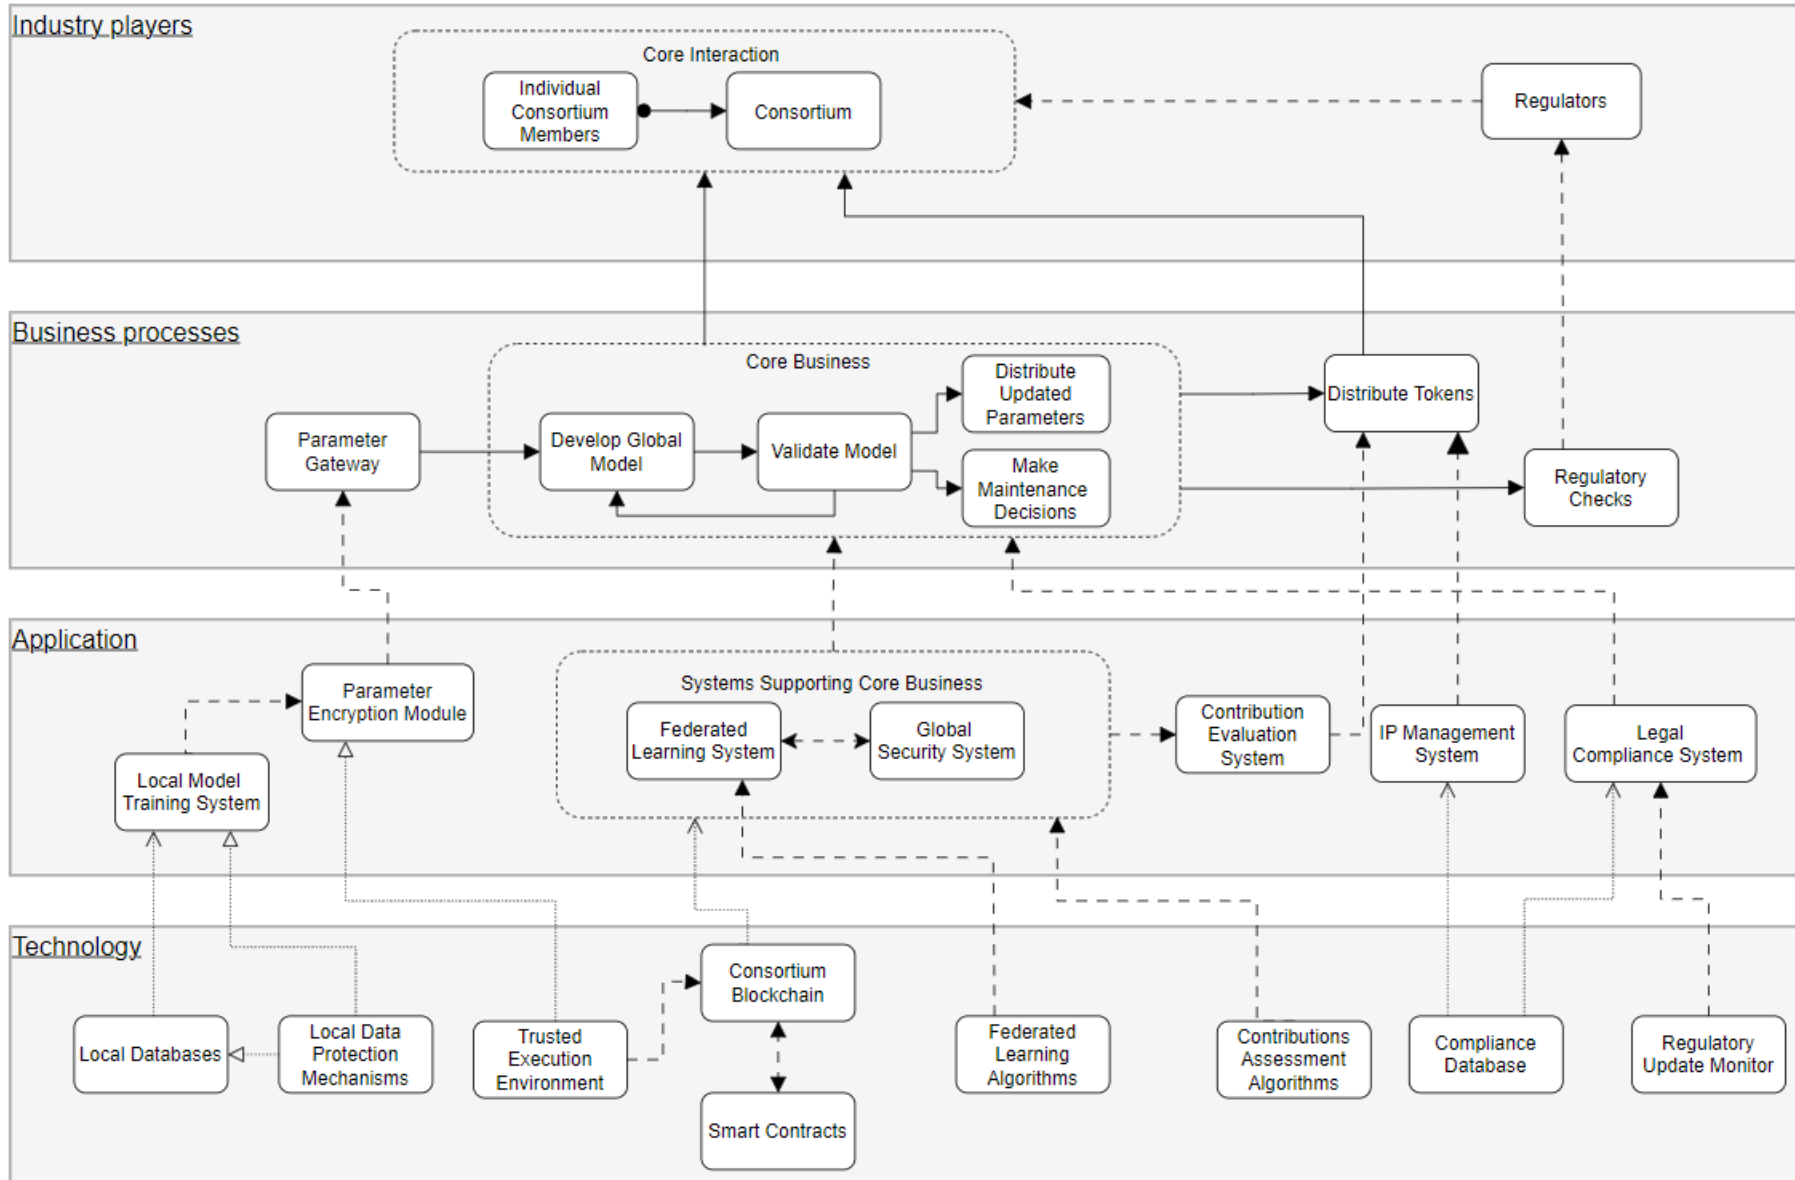


Figure 19: Overview Validated Model