REMINDER

REMINDER:

pRivacy-prEserving Machine LearnIng through secure managemeNt of Data's lifecyclE in distRibuted systems

Deliverable number: D6.1

Project Periodic Report

chist-era

FWF Austrian Science Fund

UKRI Engineering and Physical Sciences Research Council

uefiscdi Unitatea Executivă pentru Finanțarea Învățământului Superior, a Cercetării, Dezvoltării și Inovării

AGENCIA ESTATAL DE INVESTIGACIÓN

Project Acronym:	REMINDER
Project Full Title:	pRivacy-prEserving Machine LearnIng through secure manage-meNt of Data's lifecyclE in distRibuted systems
Call:	Security and Privacy in Decentralised and Distributed Systems (SPiDDS). 2022
Grant Number:	PCI2023-145989-2
Project URL:	https://ants.inf.um.es/en/reminder
Editor:	UMU
Deliverable nature:	Report
Dissemination level:	Public
Delivery Date:	31/03/2025
Authors:	UMU, SIE, UWE, AIT

Table 1: Project details.

Abstract

This deliverable presents the first annual progress report of the REMINDER project, which aims to enable secure and privacy-preserving Federated Learning (FL) across distributed and heterogeneous systems. During this initial period, the consortium has made substantial advances in designing robust architectures, cryptographic protocols, and lightweight security mechanisms to ensure the confidentiality, authenticity, and resilience of decentralized collaborative learning environments.

Key achievements include the definition of a modular FL framework combining Differential Privacy, Homomorphic Encryption, robust aggregation functions, and AI-assisted detection of malicious clients. Novel contributions were also made in privacy-preserving cryptographic primitives, including structure-preserving signatures, decentralized pseudonym systems, and post-quantum secure Oblivious Pseudo-Random Functions (OPRFs), pushing the boundaries of secure FL architectures.

These technological developments were validated through two real-world use cases: one focused on medical diagnostics using cardiovascular datasets, and another targeting energy optimization in smart buildings using IoT sensor data. Both deployments confirmed the viability of the REMINDER framework in GDPR-sensitive and resource-constrained scenarios.

This deliverable also highlights the collaborative efforts across partners, the synergy between advanced cryptographic design and applied machine learning, and outlines the roadmap for future development and experimentation. The project remains on track, with solid foundations laid for higher Technology Readiness Levels in upcoming phases.

Table of Contents

List of Tables

1 Progress Report

1.1 Project objectives and activities implemented

During the first year of the project, we focused on designing and implementing security and privacy solutions aligned with the main goals proposed. All partners collaborated in a cross-cutting manner, with each team focusing on different aspects such as cryptographic mechanisms, robustnes in FL and integration with intrusion detection techniques. Below is a detailed description of the activities carried out per objective and a summary of the work performed. Finally, we outline our planned efforts, mention the difficulties encountered, and discuss potential future refinements.

1.1.1 Objective 1: Analysis and identification of the main security and privacy requirements associated with the data lifecycle in ML systems

Activities implemented:

- Analysis of relevant regulations (e.g., GDPR, HIPAA) and state-of-the-art research to establish a baseline set of obligations for every lifecycle phase.
- Systematic threat modelling of poisoning, inference, and leakage attacks against FL pipelines, producing a catalogue of attack scenarios that drive requirement derivation.
- Domain knowledge gathering from the two REMINDER demonstrators (eHealth and smart buildings) to capture both functional needs and environmental constraints.
- Creation of a requirement taxonomy that maps lifecycle phases, threats, and countermeasures.
- Integration of the requirement set into the REMINDER FL architecture, identifying where Differential Privacy, Homomorphic Encryption, robust aggregation, and signature schemes must be enforced.

1.1.2 Objective 2: Design of a privacy-preserving and decentralized FL architecture compatible with constrained devices

Activities implemented:

- Design of a modular client stack layering: (i) Differential Privacy noise addition, (ii) compact Homomorphic Encryption (partial HE by default, low-depth FHE when available), and (iii) ECC-based digital signatures to keep computation and communication costs minimal.

- Analysis of the adaptation of each building block into lightweight variants using fixed-point arithmetic, streaming buffers, and minimal external libraries, making the DP injector, HE module, and Keccak-512 tagger feasible for IoT gateways and microcontrollers.
- Definition of a pluggable “security pipeline” interface allowing flexible substitution of heavy cryptographic modules (e.g., PHE vs. FHE) without modifying application logic.
- Definition of secure-channel and key-management guidelines favouring low-overhead protocols (TLS/DTLS with ECC certificates) and short-lived session keys for constrained and intermittently connected devices.

1.1.3 Objective 3: Design of privacy-enhancing data and entity authentication protocols

Activities implemented:

- Development of a lightweight integrity verification pipeline using compact neural network fingerprints combined with Keccak-512 hashing to validate model updates.
- Design of a mutual authentication workflow, where clients attach digests to encrypted gradients and the server verifies authenticity using digital signatures, eliminating reliance on stable device identifiers.
- Implementation of the hashing and verification protocol for edge devices, taking into account resource constraints.
- Evaluation of Keccak-512-based verification compared to alternative schemes to balance security guarantees and computational overhead.

1.1.4 Objective 4: Management of dynamism and heterogeneity in decentralized FL environments through client selection

Activities implemented:

- Development of clustering-based algorithms to group clients based on model similarity and filter outliers.
- Proposal of a transfer learning-based client selection approach to reduce communication overhead during training rounds.
- Evaluation of multiple similarity metrics, such as parameter-space distance, KL divergence, loss stability, and hardware/network profiles, within a pluggable selection framework.
- Comparative analysis of measurement techniques to balance accuracy, cost, and timeliness.

1.1.5 Objective 5: Rapid prototyping and early validation of the REMINDER framework in target use cases

Activities implemented:

- Development of a reference implementation of the REMINDER edge architecture, integrating client-side privacy modules and a secure server pipeline (malicious client filtering + FFT-based robust aggregation).
- Deployment analysis in two representative domains—eHealth and smart buildings—demonstrating effective operation under heterogeneous conditions.

1.1.6 Future Plans and Open Issues

In the next steps, our efforts will concentrate on fortifying the federated learning architecture against adversarial clients and enhancing privacy mechanisms. First, two complementary detection schemes will be implemented to identify malicious clients: one that measures the similarity of client updates in near- real time, and another that tracks historical weight patterns to spot suspicious deviations over multiple rounds. Both methods will then be compared to determine their respective strengths and deployment trade-offs. Finally, we aim to provide advanced encryption techniques—potentially leveraging homomorphic-like operations—to safeguard model integrity and privacy, thus mitigating the risk of sensitive information leakage. Currently, no major schedule or budget issues appear. Overall, the progress so far aligns with the initial work plan. Each partner's specialized research has converged into integrated analysis, with the synergy across cryptographic techniques, policy-driven access control to avoid the access to sensitive information, and IDS functionalities demonstrating strong promise for secure and privacy- preserving FL in real-world IoT settings.

1.2 Transnational collaboration

The **CHIST-ERA REMINDER** project unites a purposeful, multidisciplinary consortium of four partners from four countries—**Spain**, the **United Kingdom**, **Romania** and **Austria**—with a balanced mix of academic excellence and industrial insight. This manageable partnership brings together long-standing expertise in cybersecurity, privacy-enhancing cryptography and machine-learning-based IoT.

University of Murcia (UMU), as coordinator, centres its work on security-aware system design and on reducing the energy footprint of distributed learning. Its research groups have pioneered privacy-enhancing protocols, threat-modelling and secure orchestration for FL, and they collaborate with leading energy-systems scientists to reduce the power consumption of distributed training and inference. This makes UMU the natural lead in the smart-building energy use case, where guaranteeing data minimisation and cost-efficient operation are equally critical.

Siemens SRL (SIE) brings an industrial, clinically certified environment and takes charge of the arrhythmia and congenital heart defect detection use cases. It operates secure data-centre pipelines that keep raw ECG streams on-premises and provides the real-time dataset collected in healthcare contexts. By validating the REMINDER stack on sensitive cardiovascular data, SIE demonstrates that high-performance diagnostics can be achieved without compromising patient confidentiality.

University of the West of England (UWE) contributes deep expertise in adversarial machine learning and lightweight intrusion detection. UWE analyses privacy risks, hardens the federated workflow against poisoning and inference attacks, and embeds on-device anomaly detectors to flag suspicious updates in real time. These security reinforcements are applied across both

demonstrators, guaranteeing that the models remain trustworthy even when some clients may be compromised.

Austrian Institute of Technology (AIT) underpins the entire platform with post-quantum, attribute-based encryption and credential schemes. They develop authentication protocols and lightweight cryptography techniques to secure the federated learning architectures. Its cryptographers implement fine-grained, need-to-know access control and tamper-evident ledgers that protect model parameters and sensor data throughout their lifecycle—today and in a future quantum era. These primitives are reused by UMU, UWE and SIE alike, providing a cryptographic backbone that spans e-health and smart-building scenarios.

Within REMINDER's modular architecture, responsibilities align with each partner's core strengths. UMU has been leading the top-level framework design, defining system requirements, interfaces and orchestration logic together with the other partners. Furthermore, while UWE and AIT concentrate on the privacy module, combining adversarial-resilient analytics with post-quantum encryption and fine-grained access control to safeguard data in motion and at rest. To counter poisoning attacks, UMU and UWE jointly propose defence modules that provide robust aggregation to neutralise malicious clients before they corrupt global models. Once the whole framework is ready, SIE will provide enough data to test the framework in medical contexts, whereas UMU instantiates the smart-building use case with energy-aware sensor networks. This clear division of activities across design, privacy, security hardening and domain deployment ensures that each module is developed by the partner best equipped to maximise its effectiveness, yet all integrate seamlessly into a cohesive, end-to-end REMINDER stack.

1.3 Significant Events and Results

The REMINDER project has made significant progress across multiple fronts, combining extensive research, technological innovation, and prototyping activities.

1.3.1 Robust FL Architecture

One of the earliest milestones was the development of a comprehensive **State of the Art** on Federated Learning (FL), covering its evolution (link). This work summarized FL's strengths and limitations, identified major vulnerabilities—such as poisoning and inversion attacks—and laid the foundation for further research on resilience and privacy in non-iid, decentralized environments. These aspects were also explored in the context of intrusion detection systems (IDS).

A key contribution is the introduction of **FedRDF**, a robust aggregation function based on the **Fast Fourier Transform (FFT)** (link). This method enables secure aggregation without requiring prior knowledge of the number of adversaries. Ongoing work focuses on removing fixed thresholding and integrating differential privacy.

In addition, the consortium proposed an **unsupervised FL approach** for misbehavior detection in vehicular networks (link). Traditional FL systems rely on supervised methods requiring labeled datasets, which limits their ability to detect new threats. The REMINDER framework combines Gaussian Mixture Models and Variational Autoencoders in an FL setup to overcome this limitation.

Furthermore, lightweight and efficient security mechanisms have been developed to further protect FL systems in constrained environments:

- An integrity-driven defense mechanism using feed-forward neural networks and Keccak-512 hashing with >98% accuracy and minimal memory usage on real-world IoT datasets.
- A hybrid FL architecture with Differential Privacy, Zero-Knowledge Proofs (ZKP), and Median Aggregation that resists adversarial clients while reducing latency and overhead.
- A malware detection framework based on **Recursive Self-Distillation**, achieving >92% adversarial robustness in connected vehicles.
- A **blockchain-enhanced FL protocol** allowing decentralized, auditable training with sub-60ms latency and >97% detection precision in IoT intrusion detection.

These innovations contribute to REMINDER's mission of delivering secure, scalable, and privacy-preserving collaborative learning.

1.3.2 Cryptographic Advances and Data Protection

Significant progress was made in **data sharing, authentication, and cryptographic protocols:**

- **Attribute-Based Encryption (ABE)** mechanisms were reviewed to enforce need-to-know principles and hide access policies. Gaps in the literature were identified, setting a roadmap for future development.
- We initiated research into **homomorphic group signatures** that combine authenticity with anonymity. If successful, they would enable verifiable yet privacy-respecting model aggregation.
- A decentralized **pseudonym system** was developed for delegatable entity authentication. This system supports both user-side and sensor-side identity delegation, enhancing flexibility in onstrained networks. The scheme is described in detail in:
  - S. Krenn, D. Lesaignoux, S. Ramacher: *"bPk#: Delegatable Pseudonyms (And Their Applications to National eID Systems)"* (currently under submission).
- On the theoretical side, we introduced **structure-preserving cryptographic primitives** applicable to ring signatures and verifiable aggregation. This led to:
  - S. Krenn, O. Mir, and D. Slamanig: *"Structure-Preserving Compressing Primitives: Vector Commitments and Accumulators"* (currently under submission; preprint available at link.
- Furthermore, **post-quantum secure OPRFs** were developed. The new Leap scheme, accepted at EUROCRYPT 2025, enables secure private set intersection and pseudonymous FL aggregation.
  - L. Heimberger, D. Kales, R. Lolato, O. Mir, S. Ramacher, C. Rechberger: *"Leap: A Fast, Lattice-based OPRF With Application to Private Set Intersection"*, EUROCRYPT 2025. Preprint available at link.

These developments form the cryptographic backbone of the REMINDER architecture.

1.3.3 Real Use Cases and Prototyping Achievements

To validate the framework, REMINDER focused on two demanding real-world domains:

eHealth Use Case:

- Risk prediction of **coronary in-stent restenosis** using angiographic and clinical data.
- **Heart arrhythmia detection** from ECG time-series.
- Integration of public (PhysioNet) and synthetic datasets, including generation of realistic aortic coarctation data.
- All data processing adheres to GDPR by ensuring no raw data leaves its source.

Smart Buildings Use Case:

- Leveraged the **PLEIAData** dataset from the University of Murcia (https://www.nature.com/articles/s41597-023-02023-3).
- Predictive models were developed for **energy consumption** and **HVAC system control**, incorporating CO_2, temperature, humidity, and occupancy data.
- All models were trained via REMINDER's FL architecture, protecting resident privacy while optimizing building operations..

1.3.4 Software Development and Impact

We aim to extend the **Flower federated learning framework** with REMINDER's security modules. A complete prototype implementation for both eHealth and smart building use cases is in progress and will demonstrate:

- End-to-end encrypted model training,
- Differentially private updates,
- Verified aggregation using FedRDF.

This will contribute to the open-source community and serve as a reference implementation..

1.3.5 Economic Impact and Exploitation Potential

REMINDER lays the foundation for:

- New **privacy-preserving services** in healthcare and smart cities,
- **Energy-efficient management platforms** for smart infrastructure,
- **Secure FL-as-a-service** offerings for regulated sectors like finance and public administration.

With a strong focus on practical deployments, REMINDER's outputs have clear pathways for commercialization and public-sector adoption.

1.4 Technology Readiness Level (TRL)

The REMINDER project addresses key security and privacy challenges in distributed systems through Federated Learning (FL). At the beginning of the project, our technology was at Technology Readiness Level (TRL) 2, indicating that basic principles had been observed. During

this stage, the inherent privacy and security limitations of traditional centralized Artificial Intelligence systems were clearly identified, particularly concerning the exchange and storage of sensitive data in applications such as eHealth and smart buildings. The data collection strategy encompasses two cardiovascular disease use-cases: aortic coarctation prediction and ECG-based heart arrhythmia assessment. These applications were strategically selected to evaluate the framework across diverse data modalities - tabular data for aortic coarctation and time series for atrial fibrillation. The Clinical Decision Support (CDS) systems development utilizes two distinct data sources: a synthetic dataset for aortic coarctation generated by Siemens, and the publicly available PhysioNet database for heart arrhythmia analysis.

As the project progressed, REMINDER advanced to TRL 3, formulating a specific technological concept: integrating advanced cryptographic techniques with FL. This innovative approach enables participants to collaboratively train machine learning models without directly exposing their data, significantly enhancing privacy and reducing associated information security risks. At this stage, all experiments and validations are conducted within simulated environments, utilizing the initial design and implementation of the REMINDER framework. This framework includes modules for data obfuscation via differential privacy, cryptographic techniques ensuring the integrity and authenticity of shared models, and robust strategies for detecting malicious clients. The current validations in simulated scenarios include applications such as eHealth, where sensitive data from medical devices are protected during collaborative training, and smart buildings, aimed at optimizing energy consumption without compromising user privacy.

The primary objective remains to continuously advance the REMINDER framework, experimenting with novel techniques and thoroughly validating their effectiveness. Over the coming months, the goal is to transition from simulated tests to more practical environments, enabling direct application and evaluation in real-world use cases. In the long term, our aim is to progressively scale through higher TRL levels, thereby demonstrating the feasibility of transferring innovations from the laboratory to the marketplace.

1.5 Meetings

Table 2: Consortium meetings since project start

N°	Date	Location	Attending partners	Purpose
1	01/03/24	Audioconference	UMU, SIE, UWE, AIT	Initial project description
2	03/06/24	Audioconference	UMU, SIE, UWE, AIT	Project Status Update Meeting
3	07/10/24	Audioconference	UMU, SIE, UWE, AIT	Project Status Update Meeting
4	03/02/25	Audioconference	UMU, SIE, UWE, AIT	Project Status Update Meeting
5	07/04/25	Istanbul	UMU, UWE	CHIST-ERA meeting

1.6 Deliverables

1.7 Free comments

The consortium remains fully aligned with the Description of Action. Deliverable D1.1 defined the two pilots (e-Health, Smart Buildings) and the functional requirements that guide every technical work package. The eHealth domain, led by Siemens, leveraging their extensive

N°	Title	Nature	Delivery date (month) Contractual / Actual	Partner in charge
D1.1	Use case definition and initial requirements analysis	R, PU	M5 / M5	SIE
D2.1	First version of REMINDER architecture	R, PU	M6 / M6	UMU
D4.1	Detailed report on vulnerabilities of existing approaches to an array of attacks	R, PU	M9 / M9	UWE
D5.1	Retrospective clinical data collection	R, PU	M6 / M6	SIE
D6.1	First Annual Project Report	R, CO	M12 / M12	UMU

Table 3: Cumulative list of deliverables from project start

healthcare expertise, focused on two critical cardiovascular applications: heart arrhythmia prediction through ECG signal analysis and congenital heart defect detection through anatomical and physiological measurements assessment. The Smart Buildings domain was successfully addressed by the UMU. These carefully selected use cases serve to demonstrate the broad applicability and potential impact of the REMINDER framework in addressing significant societal challenges while advancing the state of the art in their respective fields. D2.1 translated those requirements into a first system architecture that embeds differential privacy, homomorphic encryption and robust aggregation. D4.1 then benchmarked the architecture against the complete threat taxonomy foreseen in WP4, proving that the selected countermeasures cover each attack surface identified in the proposal. Thus, every compulsory milestone for Year 1 has been reached.

UMU and SIE co-led Deliverable D1.1, with UMU drafting the Smart-Buildings pilot and SIE detailing the e-Health use case; for D2.1 every beneficiary contributed, with UMU and UWE implementing the security modules (robust aggregation, authentication), AIT specifying the privacy stack (differential-privacy and FHE pipelines), and SIE ensuring architectural coherence with both use cases; in D4.1 UMU produced the complete taxonomy and experiments on poisoning attacks while UWE authored the inference-attack analysis and counter-measures; this distributed effort is kept coherent through a shared GitLab and Overleaf workspace.

2 Dissemination of results, exploitation, impact

2.1 Scientific Publications

Reference	Same Country Partners	Different Country Partners	Open Access	DOI	URL
Hernandez-Ramos, J., Karopoulos, G., Chatzoglou, E., Kouliaridis, V., Marmol, E., Gonzalez-Vidal, A., & Kambourakis, G. (2025). *Intrusion Detection based on Federated Learning: a systematic review*. ACM Computing Surveys.	No	No	Yes	10.1145/3731596	https://dl.acm.org/doi/10.1145/3731596
Campos, E. M., Gonzalez-Vidal, A., Hernández-Ramos, J. L., & Skarmeta, A. (2025). *Federated learning for misbehaviour detection with variational autoencoders and Gaussian mixture models*. IJIS, 24(2), 1-16.	No	No	Yes	10.1007/s10207-025-01000-8	https://link.springer.com/article/10.1007/s10207-025-01000-8
Laidi, R., Merabtine, N., Djenouri, D., Latif, S., Qadir, H. A., Djenouri, Y., Balasingham, I. (2025). *Federated Learning in IoT Environments: Examining the Three-way See-saw....* IEEE Comm. Surveys & Tutorials.	No	Yes	Not yet	10.1109/COMST.2025.3557475	https://ieeexplore.ieee.org/abstract/document/10948452
Latif, S., Djenouri, D., Hernandez-Ramos, J. L., Skarmeta, A., Ahmad, J. (2024). *A lightweight integrity-driven FL approach to mitigate poisoning attacks in IoT*. IEEE FNWF.	Yes	Yes	Not yet	not available yet	https://uwe-repository.worktribe.com/output/12895988
Latif, S., Djenouri, D., Adamatzky, A. (2025). *An Integrated Approach to Mitigate Poisoning Attacks in FL Frameworks*. IJCNN.	No	No	Not yet	not available yet	https://uwe-repository.worktribe.com/output/14326974
Latif, S., Louadj, R., Djenouri, D. (2025). *FL Meets Recursive Self-Distillation: A Scalable Malware Detection Framework for IoVs*. IEEE VTC-Spring.	Yes	No	No	not available yet	https://uwe-repository.worktribe.com/output/14326990
Krenn, S., Lesaignoux, D., Ramacher, S. *bPk#: Delegatable Pseudonyms (And Their Applications to National eID Systems)* (under submission).	No	No	Not yet	not available yet	
Krenn, S., Mir, O., Slamanig, D. *Structure-Preserving Compressing Primitives: Vector Commitments and Accumulators* (under submission).	No	No	Yes (green)	not available yet	https://eprint.iacr.org/2024/1619
Heimberger, L., Kales, D., Lolato, R., Mir, O., Ramacher, S., Rechberger, C. (2025). *Leap: A Fast, Lattice-based OPRF With Application to PSI*. EUROCRYPT 2025.	No	No	Yes (green)	not available yet	https://eprint.iacr.org/2025/333

Table 4: Scientific publications resulting from the REMINDER project

2.2 Exploitation Plan

The main goal of REMINDER is an **end-to-end framework for federated learning** that (i) detects malicious participants before their updates reach the model, (ii) aggregates the remaining updates with a robust, attack-resilient algorithm, (iii) shields every client's data through differential-privacy noise and on-device encryption, and (iv) runs comfortably even on resource-constrained edge hardware.

The framework will be exploited first by the academic partners that created it, who will maintain the codebase and drive early pilots in healthcare analytics and smart-building energy management. Immediately after the internal proof-of-concept phase it will be pushed to a public GitHub repository, so that hospitals, building-management companies, cloud-edge providers and independent researchers can download, fork and extend it without having to negotiate licences.

Although some parts of the software will be openly available, ownership of the original code and documentation stays with the individual universities. This arrangement lets commercial actors build paid services or support contracts on top of the kernel while ensuring that improvements to the core remain accessible to the whole community.

The framework is aimed primarily at two domains: **healthcare**, where hospitals and research clinics need to train diagnostic models (e.g., from ECG or imaging data) without exposing patient records, and **energy-efficient smart buildings**, where facilities managers optimise HVAC and lighting schedules while ensuring that both the optimisation models and the underlying consumption data remain fully protected.

The short-term exploitation work, carried out during the lifetime of the project, focuses on hardening the prototype's security and privacy layers, refining the malicious-client selection logic, and tailoring the code so it can run smoothly on devices with limited compute or memory. The long-term phase will broaden the horizon: the partners will explore new use cases and test the same protection techniques in more complex, domain-specific deployments beyond healthcare and smart-building energy optimisation.

2.3 Exploitation Overview (Software, Products, Spin-offs, etc.)

Period	Planned Goals	Actual Exploited Results
Year 1	Our original objective was to outline the overarching architecture of the REMINDER scenario so it could withstand both inference and poisoning attacks in a federated-learning context. At this stage we aimed only to sketch, at a high level, how client-side protections and server-side defences would interact, and to specify the two reference use-cases that will drive the evaluation: (i) an e-health application focused on cardiac arrhythmia detection, and (ii) an energy-efficiency optimisation task that leverages sensor data from smart-building deployments.	We have finalised the two-tier framework. On the client side, three dedicated modules add privacy guarantees and shield local data and models. On the server side, a first module filters out easily identifiable malicious clients, while a second module applies a robust aggregation strategy to dampen the influence of any adversaries that evade the initial filter. Concretely, we devised a novel aggregation rule—**FedRDF**—that resists outliers when combining model weights, and we built a lightweight integrity-verification pipeline that couples compact neural-network fingerprints with **Keccak-512** hashing so that only untampered updates from genuine clients are accepted. Regarding the use cases, the arrhythmia-detection pipeline for e-health is now fully specified, and for the energy-optimisation track we have collected and curated the building-sensor dataset from University of Murcia (UMU) facilities.

Period	Planned Goals	Actual Exploited Results
Year 2	Goals for the second year expanded the scope significantly, aiming to create a sophisticated detection framework capable of identifying malicious clients based on historical data from model updates, while also planning to further improve the privacy of federated learning environments and implement an efficient client selection method intended to optimize system resources and enhance overall model security and convergence.	
Year 3 (if applicable)	This year's objective is to operationalise the REMINDER framework within the two reference use cases—cardiac arrhythmia detection in e-health and energy-optimisation in smart buildings—by running end-to-end simulations that evaluate not only the resulting model accuracy but also the framework's robustness against a spectrum of poisoning and inference attacks, and additionally deploy the system on resource-constrained devices.	
Project end + 1 year	The forthcoming phase focuses on turning REMINDER from an offline research prototype into a real-time defence stack. We will plug the framework directly into live data streams—heart-rate telemetry, building-management buses, and event logs—so that each federated round is executed under strict latency budgets.	n/a

Period	Planned Goals	Actual Exploited Results
Project end + 2 years	With real-time operation validated, attention shifts to the extremes of edge computing: we will port the client layer to ultra-constrained hardware such as 32-bit Arduino-class microcontrollers, introducing aggressive compression, quantisation and intermittent-connectivity scheduling to fit within kilobytes of RAM and milliwatts of power; in parallel, we will explore an orthogonal frontier—clients whose local model is a parameter-efficient LLM—and analyze the effectiveness of REMINDER in these models.	n/a
Project end + 3 years	The final year broadens the framework's horizon beyond the original use cases. REMINDER will be instantiated in new domains, such as agriculture-sensor networks, and collaborative autonomous vehicles, to verify portability across data modalities and regulatory regimes. Success criteria include seamless integration with domain-specific toolchains.	n/a

Spin-offs, Products, Standards, and Future Impact

The REMINDER project specializes in developing innovative frameworks designed to enhance security and privacy within federated learning environments. A central aspect of the project's dissemination strategy is the commitment to making these new frameworks openly available as open-source software. This open access approach will facilitate broad adoption and integration into diverse research and commercial platforms, significantly expanding REMINDER's impact and practical utility across various sectors.

In particular, REMINDER aims to openly release robust aggregation functions, sophisticated malicious client detection frameworks, and advanced privacy-preserving encryption methods developed during the project's lifespan. By offering these resources open-source, REMINDER not only promotes transparency and trust but also actively encourages global research communities and industry practitioners to implement and adapt these state-of-the-art tools to enhance their own federated learning infrastructures.

Additionally, in instances where the REMINDER project achieves particularly novel technological breakthroughs, such as a highly efficient encryption technique or an innovative approach to client selection or malicious actor detection, the consortium will actively pursue intellectual property protection through national and international patents. These patents could subsequently form the basis for future commercial exploitation, including licensing arrangements and potentially the creation of spin-off ventures.

The strategy of open-source dissemination greatly enhances the potential for establishing new partnerships and collaborative projects. By making advanced tools publicly accessible, REMINDER positions itself as an attractive and highly collaborative entity, effectively inviting external academic and industry partners to engage in joint research, pilot projects, or technology transfers.

Finally, the REMINDER project's outcomes and innovations, disseminated through open-source platforms and potential patent portfolios, could significantly contribute to standardization actions, ensuring interoperability and influencing best-practice guidelines for privacy-preserving machine learning applications. This proactive role in standardization and collaboration is likely to generate sustained long-term impact and facilitate further investment and funding opportunities.

2.4 Other Dissemination of Results

During the first reporting period, the REMINDER project has engaged in a variety of communication and dissemination activities to ensure broad visibility and impact across different audiences.

One of the key actions has been the creation and continuous maintenance of the official REMINDER project website, available at https://ants.inf.um.es/en/reminder. The website serves as a central platform to communicate the project's goals, consortium members, work packages, and progress updates. It is regularly updated with news items, publications, events, and deliverables, providing transparent access to the wider research and industrial communities.

The website targets a diverse audience, including:

- Academic researchers in federated learning, cybersecurity, cryptography, and AI.
- Industry stakeholders interested in secure and privacy-preserving machine learning.
- Public institutions and policy-makers concerned with data protection and GDPR-compliant technologies.
- General public and students, with simplified descriptions of project objectives and expected impact.

Beyond the website, dissemination efforts have included:

- Scientific publications and conference submissions, including accepted and under-review papers in major venues such as IEEE Communications Surveys & Tutorials, ACM Computing Surveys, IEEE Vehicular Technology Conference (VTC), EUROCRYPT, and others.
- Internal and external project presentations to stakeholders and at academic workshops.
- Collaborations with related European research projects, particularly in the fields of trusted computing and secure data sharing.
- Participation in the CHIST-ERA Projects Seminar in Istanbul.
- Participation in the ICTürkiye2025 International Brokerage Event.
- Future dissemination plans, including participation in academic conferences, co-organization of thematic workshops, and the publication of open-source tools and datasets when appropriate.

3 Resources and Funding

3.1 Project Level (From Project Start)

The table below summarizes the use of resources per partner since the beginning of the project, including person-months, total costs, and the percentage of the requested budget used:

N°	Partner	Person-Months	Total Costs	% of Requested Budget
1	Universidad de Murcia	10/16	31.278,57 Euros Retained in total 78.044,37 Euros	15%/33%
2	SC Siemens SRL	13	70,000 €	35%
3	AIT	5.6	36,000 €	18%
4	UWE	4	143,110.42 €	34.7%

Table 6: Budget used and effort spent per partner since project start

Comments on Expenses

- SC Siemens SRL reported expenditures of 70,000 €, which cover 13 person-months of personnel costs and the corresponding indirect costs (20%).
- The budget for AIT are approximate figures, as the reporting period of the national funding agency (FWF) is not based on project years, but calendar years.
- AIT has only spent about 18% in the first 33% of the project duration. This is due to challenges in hiring skilled personnel, resulting in the planned PhD student only starting in September 2024 (i.e., M07). However, at the current point in time, we do not foresee any difficulties for the overall project due to this delay.