Gain insights into the latest empirical Cyber Security trends and results from Horizon Europe Funded Projects

A joint White Paper from AI4CYBER, CERTIFY, CROSSCON, ENCRYPT, REWIRE and TRUSTEE

AI4CYBER REWIRE CROSSCON encrypt TRUSTEE CERTIFY

Funded by the European Union

Gain insights into the latest empirical Cyber Security trends and results from Horizon Europe Funded Projects: A joint White Paper from AI4CYBER, CERTIFY, CROSSCON, ENCRYPT, REWIRE and TRUSTEE

Executive Summary

In recent years, the significance of cybersecurity has grown exponentially in an increasingly digital world. This white paper represents a collaborative effort to consolidate insights from eight distinct European projects, all funded to address the evolving cybersecurity landscape. Within these pages, we aim to provide valuable insights into each project's objectives, use cases, high-level architecture, and the pivotal technologies that will be harnessed to safeguard our digital ecosystem. By shedding light on these initiatives, we intend to equip various stakeholders with the most current information necessary to tackle emerging cybersecurity challenges effectively.

1 Introduction

In today's interconnected world, the rapid expansion of digital data has ushered in an era of unprecedented innovation and convenience. However, it has also brought forth an array of cyber threats that jeopardize the security and privacy of sensitive information across critical domains such as healthcare, finance, and entertainment. These challenges have necessitated the development of robust solutions that not only enhance data security but also ensure compliance with stringent regulations like the General Data Protection Regulation (GDPR).

The European Union has recognized the urgency of addressing these cybersecurity challenges, leading to the initiation of eight groundbreaking projects, each committed to advancing the field of cyber security. This joint white paper aims to provide an overview of these projects and their objectives, shedding light on their significance in shaping the future of data protection and cybersecurity.

The domain of cybersecurity has evolved rapidly in response to the increasing digitization of critical sectors, exposing vulnerabilities that necessitate innovative solutions. With digital data becoming the lifeblood of industries, the need for privacy-preserving technologies has become paramount. These technologies include Fully Homomorphic Encryption (FHE), Secure Multi-Party Computation (SMPC), and Differential Privacy (DP), which form the backbone of efforts to safeguard sensitive information.

The challenges in the realm of cybersecurity are multifaceted. First and foremost, the exponential growth of digital data has placed immense pressure on existing privacy-preserving technologies, often rendering

them inefficient and non-scalable. Moreover, the diverse spectrum of cybersecurity threats demands a comprehensive approach, involving the integration of various privacy-preserving methods.

The limitations of single-key FHE schemes have spurred research into multi-key and threshold FHE, coupled with hardware acceleration, to improve computation times. Traditional security mechanisms also require seamless integration with advanced privacy-preserving technologies, necessitating methods like transciphering to address scalability issues.

To address challenges and limitations related to cybersecurity, the projects discussed in this white paper are poised to embark on several basic research directions. These include:

- Increasingly sophisticated cyber attacks by hackers and cybercriminals.
- Rapidly evolving malware and ransomware threats targeting both individuals and organizations.
- Shortage of skilled cybersecurity professionals and experts to address growing security needs.
- Difficulty in keeping up with constantly changing cybersecurity regulations and compliance requirements.
- Lack of comprehensive cybersecurity awareness and training programs for employees.
- Vulnerabilities in Internet of Things (IoT) devices leading to potential security breaches.
- Insider threats posed by disgruntled employees or careless handling of sensitive data.
- Complexity of managing and securing cloud-based services and data storage.
- Limited resources and budgets for implementing robust cybersecurity measures.
- Persistent challenges in securing critical infrastructure systems against cyber attacks.

The eight European projects highlighted in this white paper represent a collective effort to address the pressing challenges posed by the digital data revolution and the escalating cybersecurity threats. These initiatives seek to redefine the landscape of cybersecurity by advancing privacy-preserving technologies, enhancing data security, and ensuring regulatory compliance. As we delve into each project's specifics, we will uncover the innovative approaches and technologies that hold the potential to reshape the future of cybersecurity in Europe and beyond.

2 Background

The motivation behind this joint paper is rooted in the pressing need to confront the dynamic and evolving challenges of cybersecurity in Europe. The proliferation of digitalization across crucial sectors has ushered in a new era of opportunities and convenience, but it has also ushered in a parallel era of heightened risks. The exponential growth of digital data, particularly in domains like healthcare, finance, and entertainment, has led to an abundance of sensitive information, making data security and privacy preservation paramount concerns.

The digital data revolution has exposed vulnerabilities that malicious actors are quick to exploit, leading to an ever-increasing array of cyber threats. These threats pose significant risks not only to individual privacy but also to the stability and security of organizations, institutions, and nations. The motivation behind this collaborative effort is to harness the collective knowledge and expertise of eight prominent European projects. By working together and sharing insights, these projects seek to form a united front against the multifaceted challenges of data security and privacy, ultimately strengthening the cybersecurity landscape in Europe.

In addition to the eight projects featured in this joint paper, the European Union has been actively promoting research and development initiatives aimed at bolstering cybersecurity on a continental scale. A multitude of EU-wide research initiatives has been launched, encompassing a wide spectrum of activities within the cybersecurity domain. These initiatives include the development of cutting-edge cybersecurity technologies, the establishment of industry best practices, and the creation of comprehensive frameworks to facilitate regulatory compliance.

While these broader initiatives play a vital role in advancing cybersecurity awareness and capabilities across Europe, this joint paper focuses on the eight projects that represent the pinnacle of innovation in privacy-preserving technologies. These projects delve deeply into the technical nuances of cybersecurity, striving to create scalable, efficient, and user-friendly solutions capable of countering the most sophisticated cyber threats while ensuring strict adherence to regulatory requirements.

2.1 Presentation of Regulatory Bodies' Efforts

Regulatory bodies at the European Union level have taken proactive steps to address the critical challenges posed by cybersecurity in the digital age. One of the most significant regulatory milestones is the implementation of the GDPR. GDPR has set rigorous standards for the collection, processing, and protection of personal data, imposing substantial penalties for data breaches. It has become a cornerstone in data protection and privacy regulation, both within the EU and beyond.

Furthermore, the European Union has been actively engaged in shaping policies and regulations that govern data security and privacy. The EU Data Strategy, for instance, outlines a comprehensive approach to data governance, emphasizing the need for secure and trustworthy data management practices. Regulatory bodies are committed to creating an environment where individuals and organizations can trust that their data is handled with the highest standards of care and compliance.

The joint paper is designed to complement and reinforce the efforts of regulatory bodies by showcasing research initiatives that contribute to the development of privacy-preserving technologies and robust cybersecurity solutions. In doing so, the paper bridges the gap between cutting-edge research and regulatory compliance, serving as a catalyst for a more secure and privacy-conscious digital ecosystem across Europe.

3 Key Research Directions in the Projects and empirical results

3.1 ENCRYPT

3.1.1 ENCRYPT General description

The ENCRYPT project addresses pressing privacy and security challenges in sectors such as healthcare, finance, and entertainment, where the rapid growth of digital data demands robust data protection measures. Its main goal is to make privacy-preserving (PP) technologies—including Fully Homomorphic Encryption (FHE), Secure Multi-Party Computation (SMPC), and Differential Privacy (DP)—more scalable, efficient, and aligned with the General Data Protection Regulation (GDPR). To achieve this, ENCRYPT investigates cutting-edge FHE schemes, leverages hardware acceleration, and incorporates various cybersecurity approaches to enhance data confidentiality.

The project's vision is centered on developing a configurable, intelligent framework (see Fig. 1) that facilitates the secure processing of sensitive data. ENCRYPT combines state-of-the-art PP technologies

(FHE, SMPC, DP, and Trusted Execution Environments—TEEs), optimizing them through tailored configurations that balance security and performance. A notable feature is the platform's use of GPU-based acceleration to improve computational efficiency in data processing.

ENCRYPT's objectives include:

- Enhancing the scalability and practicality of PP tools for GDPR-compliant, cross-border federated data processing.
- Improving usability by supporting the identification, understanding, and adoption of PP technologies by all stakeholders involved in personal data handling.
- Promoting interoperability for privacy-preserving operations across different organizations and sectors.
- Supporting the creation of common European Data Spaces and the secure sharing of cyber threat intelligence, aligning with standardization efforts.
- Ensuring solutions are co-designed with end-users and validated in real-world scenarios, such as federated infrastructures handling personal data.
- Strengthening the open-source ecosystem by disseminating project results and supporting researcher training and development.

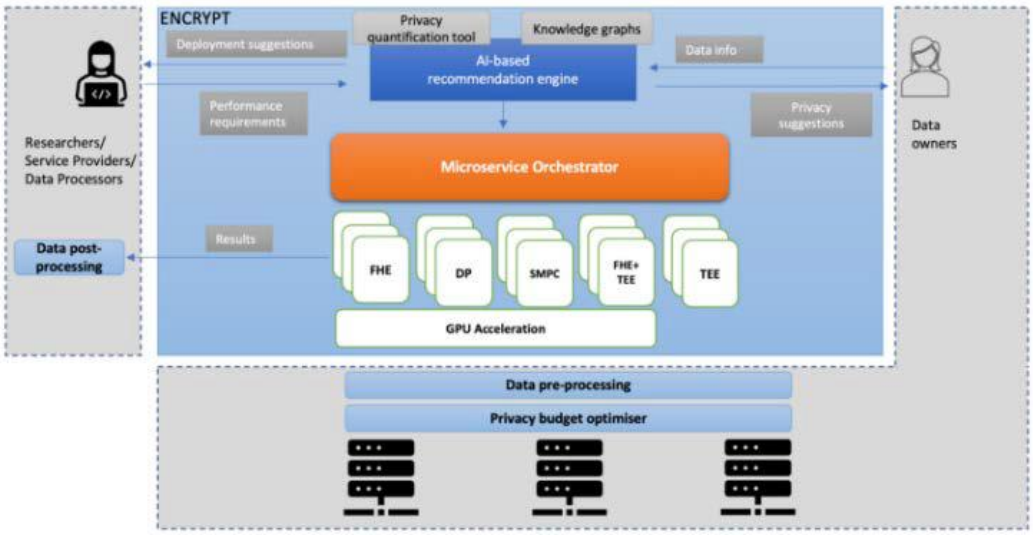

Figure 1: ENCRYPT architecture

The platform integrates AI-based tools to guide users in selecting and configuring appropriate PP technologies. The recommendation system adapts to user roles and data requirements, offering personalized privacy insights. It also includes knowledge graph-based tools that facilitate GDPR compliance through standardized metadata handling. Additional features support privacy quantification, allowing users to evaluate and align their data privacy needs, and include pre- and post-processing tools to reduce complexity and computational overhead.

3.1.2 ENCRYPT empirical use case results

The ENCRYPT project is designed to tackle the challenges of data privacy and secure data processing in critical fields like healthcare. In the medical domain, ENCRYPT aims to protect sensitive information while improving how we connect and understand health data. One of the key technologies used in this work is the knowledge graph, which supports secure, meaningful, and interoperable data usage across different systems and countries.

A knowledge graph is a way of organizing data so that each item such as a symptom, diagnosis, or treatment is connected to related items through well-defined relationships. These graphs are machine-readable, allowing computers to interpret and use the information effectively. In healthcare, this is especially valuable, as clinical data often comes in complex forms, from electronic health records to diagnostic codes. By structuring this data into a graph format, it becomes easier to integrate, search, and analyze. In ENCRYPT, medical data was mapped to a widely used medical vocabulary called SNOMED CT. SNOMED provides standardized codes and terms for diseases, procedures, findings, and other clinical concepts. By aligning local terms from hospital data with SNOMED CT, ENCRYPT ensures that different systems understand each other, even if they originally use different formats or labels. This mapping is essential for building knowledge graphs that are consistent and useful across settings.

The project tested whether large language models (LLMs) advanced AI systems trained on massive amounts of text could help with this mapping. Instead of manually linking hundreds of terms to SNOMED, the team used LLMs like GPT-4o, Claude 3.5, and Gemini 1.5 to automatically suggest the most suitable matches. The goal was to see whether these models could improve both the speed and quality of mapping, while reducing human effort. To evaluate the results, ENCRYPT used a dataset of 108 real-world medical terms collected from a hospital setting. These included both simple items like “weight” or “age” and more complex ones like “withdrawal of therapy” or “diagnostic question.” Six systems were tested: five LLMs and one older method known as BERTMap. Each system attempted to map the terms to SNOMED CT, and their outputs were compared to expert-reviewed mappings.

The evaluation used three standard metrics: precision (how many suggested matches were correct), recall (how many correct matches were found), and F1-score (a balance of both). The results showed that GPT-4o outperformed all other systems by a clear margin. It achieved 93.75% precision and 98.90% recall, resulting in an F1-score of 96.26%. This indicates that GPT-4o not only found the right mappings but did so with very few errors. Other systems, like Gemini and Claude, performed reasonably well, but none reached the same level of accuracy. Llama 3.3 and DeepSeek R1, which are open-source models, showed weaker performance, especially on complex terms. The following table presents the key performance scores for each system tested.

Method	Precision	Recall	F1-Score
GPT-4o	93.75	98.90	96.26
Claude 3.5 Sonnet v2	53.75	69.35	60.56
Gemini 1.5 Pro	60.27	66.67	63.31
BERTMap	48.84	71.19	57.93
Llama 3.3 70B	**19.19**	**70.37**	**30.16**
DeepSeek R1	**25.76**	**32.69**	**28.81**

Table 1: Comparison of method performance metrics

In addition to performance scores, the number of correct and incorrect matches was also tracked. These figures are shown in the following table.

Method	True positives	False positives	True negatives	False negatives
GPT-4o	90	6	11	1
Claude 3.5 Sonnet v2	43	37	9	19
Gemini 1.5 Pro Latest	44	29	13	22
BERTMap	42	44	5	17
Llama 3.3 70B	19	80	1	8
DeepSeek R1	17	49	7	35

Table 2: Contingency table

The following figure provides a visual comparison of how each method performed, clearly showing GPT-4o's lead.

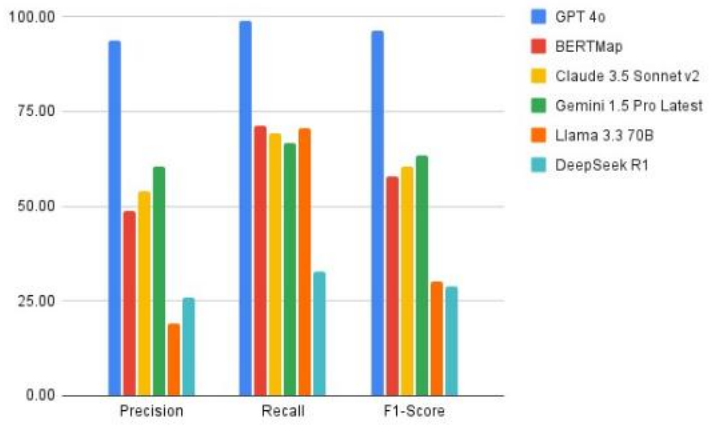

Figure 2: A depiction of the performance for all methods

Looking more closely, GPT-4o performed especially well on complex clinical language. It was able to correctly handle detailed terms involving procedures, diagnostic questions, or treatment decisions. This is important because real-world medical records often involve these kinds of terms. The other models were generally more reliable when dealing with simpler data, like basic measurements or demographic details. By using knowledge graphs powered by SNOMED and supported by advanced AI models, ENCRYPT has shown how structured, standardized data can be built efficiently from clinical sources. The use of large language models simplifies the process of connecting hospital records to a universal medical vocabulary. This has clear benefits for patient care, research, and system interoperability. It also ensures that privacy and legal requirements like GDPR are respected, as ENCRYPT integrates secure processing tools alongside knowledge graphs.

In summary, the healthcare use case within ENCRYPT demonstrates the value of combining knowledge graphs with advanced AI to better understand medical data. The strong results from GPT-4o point to new possibilities in automating complex tasks that previously required expert input. With further development, these techniques could become standard tools for hospitals and research centers aiming to make better use of the information they already have securely, accurately, and efficiently.

3.1.3 ENCRYPT summary and future steps

The ENCRYPT project is set to finish in June 2025, but the technologies and tools developed will not stop there. Work will continue after the project ends, aiming to expand and improve the results. One focus will be on enhancing the main privacy-preserving technologies used in ENCRYPT, such as Homomorphic Encryption, Trusted Execution Environments, and Differential Privacy. These will be further tested and developed to handle more demanding tasks and larger datasets, while also improving their speed and efficiency.

Future activities will also look at applying ENCRYPT's methods in new use cases. Building on the current applications in oncology, cybersecurity, and finance, the goal is to explore additional sectors where secure data sharing and processing are important. This will help show how the platform can adapt to different needs and environments. The continuation of the project's work will also support improvements in tools like the recommendation system, data preprocessing services, and the risk assessment framework. These components will be refined to become easier to use and more flexible. By doing so, ENCRYPT's approach to privacy and data protection will remain relevant and useful beyond the official end of the project.

3.2 AI4CYBER

3.2.1 AI4CYBER general description

The AI4CYBER framework was conceptualised as a collection of Artificial Intelligence (AI) capabilities that improve and automate different cybersecurity services. As it can be seen in the figure below, AI4CYBER framework offers eleven components that leverage AI and they can work independently and in collaboration if needed.

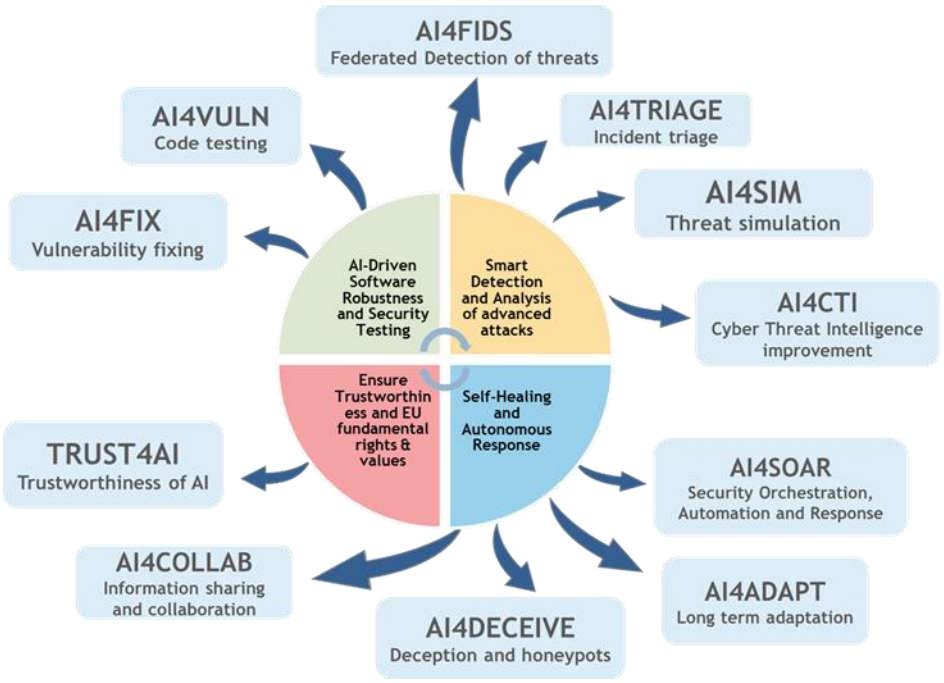

Figure 3: AI4CYBER Framework components and domain of application in cybersecurity

AI4CYBER is a collection of models and algorithms that enhance cyber resilience capabilities of IT systems in multiple ways and stages, as follows:

AI-DRIVEN CODE TESTING:

- AI4VULN – Code testing
  Open-source solution to automatic identification and verification of vulnerabilities in the code applying symbolic execution and the use of AI. The tool is specialised in java code.
- AI4FIX – Vulnerability fixing in code
  Open-source vulnerability fixing solution supporting Java, bringing automatic unit testing of proposed fixes, which enables to fix faster and earlier in the SecDevOps cycle. The tool is integrated with AI4VULN for a complete robustness of the code.

AI-POWERED THREAT DETECTION:

- **AI4FIDS - Federated Detection of threats**
  A high-performance and accuracy solution for Advanced and AI-powered attacks detection in distributed environments where privacy of data processed by detection agents need to be kept. The solution adopts Federated Learning and it is enhanced with TRUST4AI.XAI (see below)
- **AI4TRIAGE – Incident triage**
  AI-based root cause analysis and alert triage to prioritize events to focus on the response.
- **AI4SIM – Threat simulation**
  An Advanced cyberattacks simulation solution capable to simulate advanced and AI-powered attacks against IT, OT and IoT systems depending on the customer needs.
- **AI4CTI - Cyber Threat Intelligence improvement**
  An advanced solution that offers latest AI-powered Cyber Threat Intelligence (CTI) to detection and threat simulation tools for raising their efficiency.

AI-POWERED RESPONSE:

- **AI4SOAR – Security Orchestration, Automation and Response**
  AI-powered SOAR to deploy multiple security controls at different layers of the system.
- **AI4ADAPT – Long term adaptation**
  AI-based service that enriches the AI4SOAR with long-term response based on self-learning the system status and the efficiency of the security controls deployed
- **AI4DECEIVE – Deception and honeypots**
  The intelligent deception mechanisms that will enrich the response of the AI4SOAR.
- **AI4COLLAB –Information sharing**
  Automatic sharing of anonymized incident information.

TRUSTWORTHY AI:

- **TRUST4AI - Trustworthiness of AI**
  A set of highly innovative methods and models ensuring trustworthiness of AI systems. The methods include explainability (TRUST4AI.XAI), fairness (TRUST4AI.Fairness) and security (TRUST4AI.Security), and therefore, the solution is in fact three solutions in one.

3.2.2 AI4CYBER empirical use case results

As part of the AI4CYBER framework, the AI4FIDS component is a Federated Intrusion Detection System (FIDS) solution (A comprehensive survey of Federated Intrusion Detection Systems: Techniques, challenges and solutions - ScienceDirect) that utilises an AI-powered multimodal approach to recognise potential

cyberattacks based on different data types. To do so, AI4FIDS consists of four detection systems, namely: (a) L-FIDS: Log-based FIDS, (b) O-FIDS: Operational data-based FIDS, (c) N-FIDS: Network flow-based FIDS and (d) V-FIDS: Visual-based FIDS that work in a collaborative manner. The training procedure of such detection systems adopts Federated Learning (FL) to facilitate the collaborative training of models across several decentralised units, thereby eliminating the necessity for direct data exchange. This is useful when there are privacy concerns or when data is dispersed, for example, when elements of the system under protection are deployed in different networks.

The research on AI4FIDS was focused on solutions that overcome the challenges in FL training for IDS. One of the main challenges in FL training is the inherent data heterogeneity across datasets from clients in the federation. This non-IID (non-independent and identically distributed) nature of the data can significantly impact model convergence. As a result, FL models may exhibit biased learning, favouring clients with dominant data distributions while underperforming on minority patterns. This issue can degrade the performance of FL-based IDS, rendering them unreliable for real-world deployment. Addressing this challenge requires robust aggregation methods, adaptive learning techniques, and strategies to balance the contributions of different clients to ensure a fair and efficient global IDS which is able to accurately detect intrusions for different scenarios of dataset sizes and heterogeneity among clients.

As part of AI4CYBER's algorithms to enhance the detection accuracy of FL-based IDS while ensuring they are reliable for practical use, the StatAvg FL aggregation technique was developed to address feature shift in client data distributions. StatAvg introduces a global data scaling mechanism to improve model consistency across decentralized datasets. Our approach preserves privacy while significantly enhances the detection accuracy of FL-based IDS, outperforming various state-of-the-art FL aggregation techniques as shown in Figure 2. A preprint version of the technical specifications of StatAvg can be found in https://arxiv.org/pdf/2405.13062

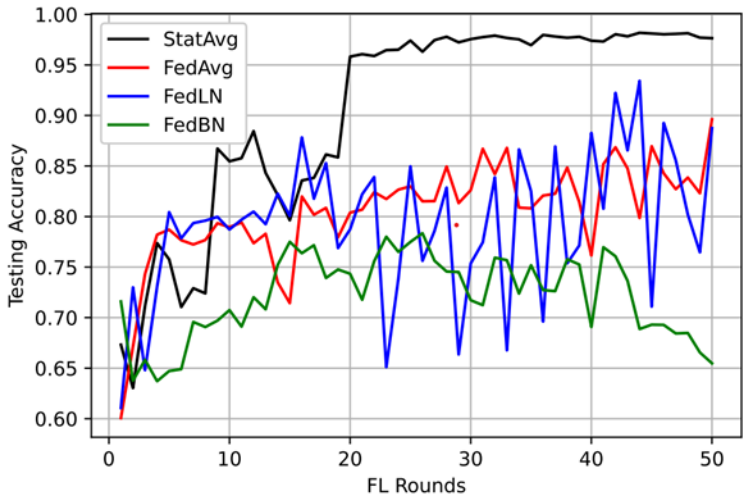

Figure 4: StatAvg's performance evaluation compared to other FL aggregation techniques

In line with AI4CYBER's commitment to reproducible and open research, AI4CYBER continues to push the boundaries of what's possible in AI-powered cybersecurity, ensuring that our solutions are not only effective but also accessible to researchers and practitioners worldwide. We are excited to announce that StatAvg has been integrated as a baseline technique in the repository of the Flower Framework. Flower is

one of the most widely recognized frameworks for Federated Learning, adopted by both academia and industry, with a strong and active community continuously enhancing its capabilities. The open-source code for StatAvg is available at this link. The journey embarked upon by AI4CYBER aligns with our commitment to advancing the field of cybersecurity through innovative, privacy-preserving cybersecurity solutions.

3.2.3 AI4CYBER summary and future steps

The AI4CYBER framework was evaluated at the end of the summer of 2024 in the project use cases (energy, health and banking systems) with very positive results. After evolving the components and refining their capabilities in a second version, the project is currently preparing the final evaluation of the final version of all of its components integrated into the use cases. We expect that the final integrated and evaluated framework is available this summer when the project ends (31/08/2025).

As main conclusion of the project, the research done has demonstrated very good results of the use of Artificial Intelligence and Generative Artificial Intelligence for different cybersecurity purposes when preparing the systems to be resilient, when protecting the deployed systems and when responding to cyber incidents. The future plans of the partners include the follow up of the path initiated in AI4CYBER to explore the potential of AI, and particularly Generative AI, for aiding in cybersecurity services. Furthermore, exploring the trustworthiness of the AI and the Generative AI is seen as core research by the Consortium. Trustworthy AI and Generative AI applied to cyber resilience will ensure not only advanced cybersecurity services, but also trustworthiness in their decisions, their automation processes, and even in their own explainability.

3.3 CERTIFY

3.3.1 CERTIFY general description

CERTIFY aims to provide IoT stakeholders with the tools and mechanisms needed to achieve a guaranteed level of security by enabling the detection, evaluation, and response to cyberattacks in a collaborative and decentralized manner throughout the entire lifecycle of IoT systems. The project adopts a comprehensive approach that integrates security by design, continuous security assessment, timely threat detection and mitigation, secure device updates, and continuous information sharing. By managing IoT cybersecurity holistically, CERTIFY seeks to enhance resilience while reducing costs.

To achieve its mission, CERTIFY defined a set of specific objectives:

- Cybersecurity situational awareness for IoT-enabled environments through a multi-stakeholder sharing of threats and mitigations.
- Secure reconfiguration and maintenance of customizable embedded devices by means of open hardware primitives and services.
- Perform security operational management based on bootstrapping and monitoring of attacks and malicious behaviors.
- Runtime security compliance and continuous certification methodology via empirical and objective metrics.
- Foster knowledge delivery via wide dissemination, capacity building and supporting standardization activities.

- Industrial validation of the CERTIFY framework in IoT ecosystems.

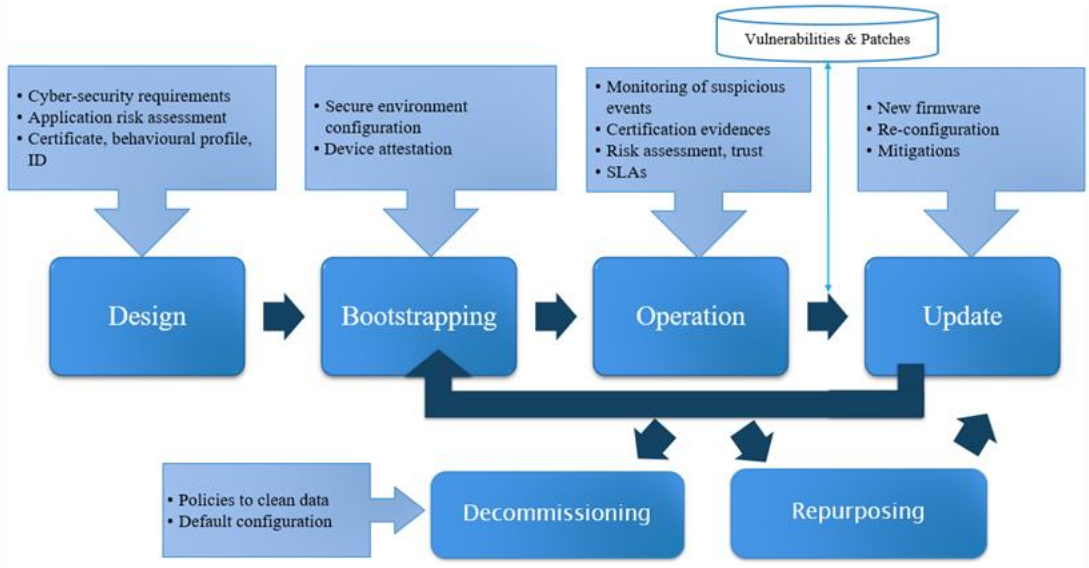

Figure 5: The enhanced cybersecurity lifecycle in CERTIFY

The CERTIFY framework is a three-layer architecture for IoT security. The foundational layer focuses on security enablers for RISC-V and ARM hardware, the middle layer develops software for device and network monitoring, and the top layer provides security services and data sharing with external partners. CERTIFY introduces a cybersecurity lifecycle management framework that enables continuous security assessment, rapid threat response, secure over-the-air updates, and information sharing with manufacturers and ISACs. This approach aims to enhance security resilience and support agile certification processes. Components and solutions part of the CERTIFY architecture are regrouped in several categories:

- CERTIFY secure node: it includes the design and low-level security enablers for the MCUs (Micro Controller Units).
- CERTIFY security services: it includes the agents running on the embedded device offering, secure bootstrapping, runtime attestation, authentication, and reconfiguration.
- CERTIFY enforcement domain: services deployed in the domain to provide a secure device update and enrolment.
- CERTIFY orchestration domain: supported by the inventorying and registry, it orchestrates the security operations performed on the device and in the domain.
- CERTIFY runtime sensors and monitoring domain collects and processes logs and row data to protect against cyberthreats and vulnerabilities, and enforces the best security responses considering external knowledge and current configuration of the domain.
- Cyber-threat sharing domain includes services for the secure configuration of the IoT devices according to device specifications and latest threats, and the sharing of the local knowledge.

3.3.2 CERTIFY empirical use case results

Internet of Things (IoT) connected devices are being deployed more and more in the cabin to enhance passenger experience and improve airlines' operations. Main benefits span from better remote prognostics and health management (PHM) to reduced maintenance time and support to a continuous (re-)certification process. Moving towards a reduction in cost and an increment in flexibility, modifiable off-the-shelf (MOTS) devices and wireless (in place of wired) connections are used where possible. Such extended flexibility must comply with cybersecurity requirements issued by the FAA (Federal Aviation Administration) and EASA (European Union Aviation Safety Agency) to protect on-board electronic networks and systems against cybersecurity threats throughout the whole lifecycle of the on-board devices.

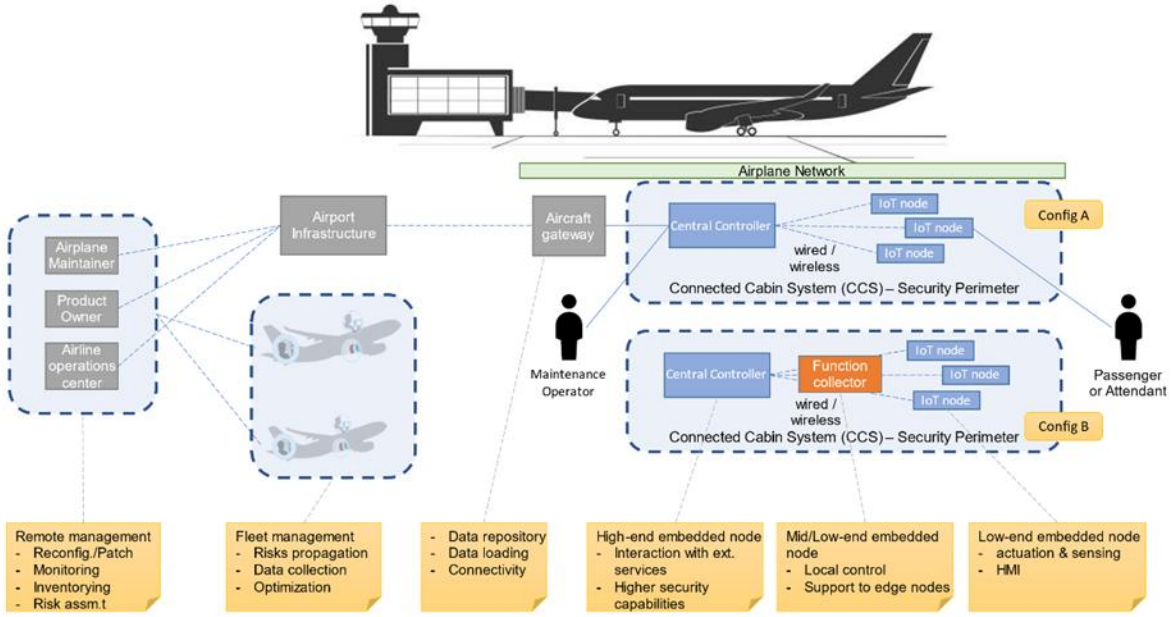

Figure 6: CERTIFY use case

Risk assessment

In recent years, researchers have studied the cybersecurity risks affecting the aviation industry. The analysis covers a range of threats, including data injections, software corruption, and the compromise of remote maintenance interfaces (Habler, Bitton, & Shabtai, 2022). Cyber incidents over the past two decades reveal that Advanced Persistent Threat (APT) groups, often working with state actors, target IT infrastructure through unauthorized access, ransomware, and data breaches, while the increasing use of commercial off-the-shelf (COTS) equipment and wireless communications presents additional challenges (Ukwandu et al., 2021). The shift to IP-based networks in aviation introduces threats like message manipulation, impersonation, and sensitive data leakage, which can be mitigated through multiple communication protocols and cloud computing (Shaikh et al., 2019). Establishing a secure state for wireless networks before integrating them into aircraft systems is crucial, with trusted platform modules playing a key role (Markantonakis, Akram, & Holloway, 2016). Vulnerabilities in embedded systems, such as unauthorized access, information leaks, and malicious code deployment, further exacerbate the security landscape (Papp, Ma, & Buttyan, 2015). Additionally, secure software development,

communication protection, access control, and forensic analysis support are vital to managing the expanded attack surface of e-enabled aircraft (Wolf, Minzlaff, & Moser, 2014). The NIST Risk Management Framework (NIST, 2016) provides a comprehensive set of cybersecurity guidelines, addressing areas like access control, data protection, and system integrity, essential for safeguarding aviation systems.

Based on this, we have categorized several threats in the following table (see figure 7):

Lifecycle Phase	Des	Des	Boot	Oper	Boot/Oper	Des/Boot/Oper	Oper	Oper	Oper	Boot/Oper	Oper
Threat scenario Connected cabin	**TEE**	**SE**	**Secure enrolment**	**Device runtime attestation**	**Network bootstrapping monitor**	**Secure device (re-)configuration/Extended MUD**	**Privacy Preserving CTI and MISP**	**SIEM-SOAR**	**Network IDS**	**Inventorying & registry**	**Secure upgrading**
TS.01 (Config. / OTA package integrity compromise)	Y			Y	Y	Y	Y		Y	Y	Y
TS.02 (Credentials leaked)	Y	Y	Y		Y	Y				Y	
TS.03 (Sensitive data leaked due to malicious/inadvertent misconfiguration)	Y					Y		Y	Y		
TS.04 (Compromise of outward-facing services interface)	Y			Y		Y	Y	Y		Y	Y
TS.05 (Injection of counterfeit data)				Y				Y	Y		
TS.06 (DoS affecting availability)							Y	Y	Y		
TS.07 (Data leakage due to rogue IoT/WiFi node in the cabin network)	Y	Y	Y		Y	Y		Y	Y	Y	
TS.08 (Malicious sw or LRU injection through the supply chain)		Y		Y						Y	
TS.09 (On board attack to an IoT node through HMI or wireless interfaces)		Y		Y		Y	Y		Y		Y

Figure 7: Risk categorization

Risk mitigation

These scenarios provide a realization of a threat, i.e., describe a concrete way in which a threat can be realized. For each threat scenario we have gone through the process of describing an attack vector (the primary entry point for the threat at the security perimeter), then the intermediate steps through which a threat can traverse the system under analysis, in particular describing if there are ways to circumvent standard security measures, and finally to achieve a threat condition. These threat scenarios include the current risk scoring of the specific threat scenario, mitigation plan, and status of implementation of the identified security measures. Threat scenarios are important to provide concrete elements to evaluate the feasibility of a specific threat and the associated level of risk.

For the sake of conciseness, we will not include the complete description of each threat scenario, but based on them we have reached the technologies and solutions that conform CERTIFY to offer appropriate mitigations to the identified threats, which are the following:

1. Strong authentication of the nodes and remote counterpart [direct anonymous attestation]
2. Robust certificates/key distribution protocol and related infrastructure [secure enrollment]
3. Protection and isolation by using trusted environments.
4. Continuous attestation of integrity and behavioral monitoring to identify unexpected deviations/anomalies [attestation/integrity monitors].
5. Package signature for (integrity and) authenticity.
6. Protection during the network authentication [network bootstrapping monitor].

7. Isolation/separation of frontend APIs from backend data collection services guaranteeing non-bypassable security profiles and corresponding access control policies. Isolation/separation of authorization logic with guaranteed enforcement to treat data and services according to established security domains.
8. Security configuration of the devices [advanced device bootstrapping/extended MUD]
9. Network analysis and protection [network intrusion detection]
   Remote network analysis and mitigation distribution [CTI, SIEM]

The threat scenarios have been assigned a score before and after the mitigations. We now consider (in Table/figure 8) the residual risk before and after the CERTIFY solutions are deployed.

	Technical difficulty					Impact	Old risk	Residual risk	New priority
Threat ID	expertise	knowledge	equipment	time					
TS.01	3	2	3	3	**11 -> 16**	**13**	1.18	0.81	Mod -> Low
TS.02	3	3	2	3	**11 -> 15**	**11**	1	0.73	Mod -> Low
TS.03	2	3	1	3	**9 -> 16**	**10**	1.11	0.69	Mod -> Low
TS.04	3	3	2	3	**11 -> 14**	**15**	1.36	1.07	High -> Mid
TS.05	2	2	2	3	**9 -> 11**	**13**	1.44	1.18	High -> Mod
TS.06	1	2	1	2	**6 -> 8**	**11**	1.83	1.38	High -> High
TS.07	3	2	3	3	**11 -> 14**	**9**	0.82	0.64	Low -> Low
TS.08	4	3	3	4	**14 -> 15**	**18**	1.29	1.2	Mod -> Mod
TS.09	3	3	2	3	**11 -> 14**	**11**	1	0.79	Mod -> Low

Figure 8: Residual risks after the adoption of CERTIFY in the connected cabin use case.

All in all, the introduction of the CERTIFY mitigations can significantly reduce the risk by increasing the technical difficulty in carrying out an attack (while the impact remains unaltered). The results presented here should be considered qualitative only with the purpose of coarsely estimate benefit of mitigations and assign priorities in addressing the threat scenarios.

3.3.3 CERTIFY summary and future steps

CERTIFY enhances IoT security through a comprehensive framework that integrates security-by-design principles, continuous monitoring, and compliance assessment. By focusing on continuous certification, it ensures that security evolves to counter emerging threats through empirical security metrics and automated compliance mechanisms. The framework strengthens cybersecurity resilience by incorporating open hardware security enablers, secure bootstrapping, and real-time threat mitigation strategies. Additionally, CERTIFY promotes decentralized threat intelligence sharing and security orchestration, making it a scalable solution for various IoT environments.

One key application of CERTIFY is in the secure management of connected cabin systems in aviation. With the growing reliance on wireless communications and modifiable off-the-shelf devices, ensuring cybersecurity compliance with aviation regulations is critical. CERTIFY addresses risks such as unauthorized access, device authentication vulnerabilities, and remote maintenance threats, providing robust security

mechanisms that enhance both operational efficiency and regulatory adherence. By implementing structured risk assessment and mitigation strategies, CERTIFY significantly reduces attack feasibility while maintaining system integrity, demonstrating its effectiveness in securing IoT ecosystems.

Ultimately, CERTIFY represents a step forward in ensuring that IoT devices can achieve and maintain a verifiable and resilient security state in the face of evolving cyber threat.

CERTIFY is expected to have a major impact on the different sectors identified in relation to the application of cybersecurity and IoT. To maximize this impact, the project results will be communicated to a wide scientific community, air transport authorities and operators, cybersecurity agencies, transport IoT and software providers and players, industrial companies, the European and international community, as well as the general public.

3.4 CROSSCON

3.4.1 CROSSCON general description

CROSSCON (Cross-platform Open Security Stack for Connected Devices) addresses a central challenge in IoT: billions of heterogeneous devices – from bare-metal microcontrollers to AI-capable processors – lack consistent **Root-of-Trust (RoT)** and **Trusted Execution Environment (TEE)** features, creating weak links exploitable by attackers. CROSSCON filled this gap by designing an **open-source, modular, vendor-independent security stack** designed for the fragmented IoT landscape. It achieved that by adopting the following design principles:

- **Layered modular design**: CROSSCON provides a customizable stack that can be tailored to underlying hardware. It supports both **bare-metal TEEs** (using software instrumentation or MPU when hardware protection is absent) and **hypervisor-based TEEs** leveraging TrustZone, RISC-V PMP, or virtualisation primitives on complex SoCs.
- **Unified APIs & abstraction**: CROSSCON adopted and extended the Global Platform APIs to abstract TEE operations across architectures, promoting both *portability* and *interoperability*
- **Enhanced hypervisor**: Based on the open-source Bao hypervisor, CROSSCON adds capabilities like per-VM TEEs, dynamic VM instantiation, cache-coloring, secure shared-memory channels, and the capacity to run multiple TEE models in parallel – even host Intel SGX enclaves or Arm TrustZone environments on RISC-V hardware.

In short, CROSSCON delivers a unified security infrastructure that adapts to device constraints, brings strong isolation via hypervisors or software TEEs, ensures correctness through formal methods, and offers a toolbox for trusted services, enabling robust, interoperable IoT systems across a variety of hardware.

3.4.2 CROSSCON empirical use case results

CROSSCON's use cases are designed to demonstrate how its open, modular, and vendor-independent IoT security stack can address real-world security challenges across a range of device types and application domains. These use cases highlight the practical relevance of CROSSCON's architecture and its capacity to support essential IoT security services across heterogeneous platforms.

The use cases focus on three primary areas:

Device Multi-Factor Authentication

This use case tackles the growing need for strong authentication mechanisms in IoT. Rather than relying solely on Physical Unclonable Functions (PUFs), which are often difficult to implement and vulnerable to certain attacks, CROSSCON explores a robust **multi-factor authentication (MFA)** approach. The goal is to combine PUFs with additional authentication methods to mitigate risks such as Man-in-the-Middle (MITM) attacks and unauthorized access. This use case ensures only trusted devices can join and communicate within a networked environment, strengthening the first line of defense for IoT deployments.

Secure Firmware Updates

Recognizing that insecure firmware update processes are one of the most common vulnerabilities in IoT device lifecycles, CROSSCON develops a secure, manifest-based update mechanism. The approach supports both **full and partial updates**—the latter reducing bandwidth and resource use by transmitting only differences (diffs) between versions. Importantly, CROSSCON embeds validation of the firmware and manifest signatures within a Trusted Execution Environment (TEE), ensuring updates cannot be tampered with, even if the device's main operating system is compromised. This use case emphasizes **resilience, integrity, and efficiency** in keeping devices securely up to date.

Commissioning and Decommissioning of Devices

The process of securely onboarding (commissioning) and offboarding (decommissioning) IoT devices is critical, especially in industrial and multi-stakeholder environments where sensitive credentials and configurations are at stake. This use case enables secure provisioning of devices—including credentials, certificates, and configuration data—at the start of their lifecycle, and ensures that when devices are retired or reassigned, all sensitive data is wiped or reconfigured appropriately. CROSSCON's stack provides the tools to automate and secure these transitions, supporting data protection and regulatory compliance.

These use cases serve both as validation of the CROSSCON stack and as blueprints for industry adoption, showcasing how modular and secure IoT systems can be developed and deployed even on highly constrained hardware platforms.

3.4.3 CROSSCON summary and future steps

CROSSCON has successfully demonstrated that a **unified, open-source IoT security stack** can operate effectively across widely diverse hardware—from resource-constrained microcontrollers to sophisticated multi-core APUs—while maintaining strong security assurances through formal methods. The integration of both bare-metal and hypervisor-based TEEs, along with trusted services like enriched multi-factor authentication using PUFs and context-aware methods, control-flow integrity, and secure update mechanisms, offers a flexible yet consistent infrastructure for building secure IoT ecosystems.

Looking forward, future steps include:

- **Extending hardware support**, notably RISC-V custom SoCs and FPGA-accelerated architectures, enabling TEE capabilities even in GPU and NPU or domain-specific platforms
- **Enhancing trusted services**: expanding the range and sophistication of services, including biometric TEE modules, ML-model protection.

- **Industrial deployments and impact**: through active contribution to the RISC-V specification process, with the availability of a mature SoC version of the stack, with testbed deployment CROSSCON is poised to transition from lab to real-world systems.
- **Long-term standardization**: promoting CROSSCON specifications and API standards within open-source communities and emerging IoT protocols, catalyzing broader adoption and ecosystem compatibility.

3.5 TRUSTEE

3.5.1 TRUSTEE general description

The European Commission's data strategy envisions the development of a unified European Data Space to foster a single market for data. This initiative seeks to ensure sovereign control over data produced by individuals and organizations, while promoting its wide-scale reuse across key economic and societal domains such as healthcare, agriculture, energy, and smart cities. By enhancing Europe's digital competitiveness, the strategy aims to unlock data-driven innovation and support a resilient digital economy. The construction of the European Data Space is ongoing and involves multiple layers, including sector-specific data spaces, enabling legislation, and governance frameworks aligned with EU values. A critical challenge remains the fragmentation of data across silos, as well as legal, social, and technological barriers to data interoperability and reuse. To address these challenges, the strategy is underpinned by technological initiatives focused on secure, reliable, and interoperable data infrastructures. Key developments include standards for data quality and interoperability, privacy-preserving technologies, secure cloud and edge computing, federated learning, and blockchain-based trust mechanisms. These technologies support secure, privacy-respecting, and legally compliant data exchanges across borders and sectors. In alignment with this vision, the TRUSTEE project plays a pivotal role by advancing a secure and trusted technological architecture for data sharing in compliance with EU data protection legislation. TRUSTEE's framework emphasizes privacy, trust, and accountability, leveraging FAIR (Findable, Accessible, Interoperable, Reusable) principles, homomorphic encryption, and self-sovereign identity mechanisms. The platform aims to enable secure data access and exchange within and across various European data spaces through user-centric, environmentally sustainable ICT solutions.

The TRUSTEE architecture is built on a layered approach:

- Federated Access Layer: Provides a user interface and service APIs for seamless data interaction and integration with European open data platforms such as GAIA-X and EOSC.
- Accountable Transaction Layer: Uses Hyperledger Fabric to ensure verifiable, privacy-preserving user transactions.
- Self-Sovereign Identity and Security API Layer: Implements privacy impact assessments (PIAs) and legal safeguards to manage identity and privacy across federated domains in compliance with EU law.
- Cloud Continuum Service Layer: Supports dynamic resource management from edge to cloud, accommodating mobility and heterogeneity.
- Core Application and Service Layer: Delivers cross-sectoral analytics, governance, and AI/ML services, including anomaly detection and workflow orchestration.
- Knowledge Repository Layer: Enables semantic harmonization and fusion of heterogeneous datasets, facilitating knowledge discovery and decision-making.

- Common Data Space Sources: Integrates structured and unstructured data from diverse domains with persistence and evaluation mechanisms.

TRUSTEE also addresses the growing need for responsible data governance by integrating legal, ethical, and social considerations. It aligns with regulatory requirements such as GDPR and sector-specific regulations (e.g., HIPAA, CCPA), and promotes trust through transparent practices and impact assessments. By combining cutting-edge digital tools with a commitment to social and environmental responsibility, TRUSTEE exemplifies the EU's vision for a secure, interoperable, and trusted European Data Space.

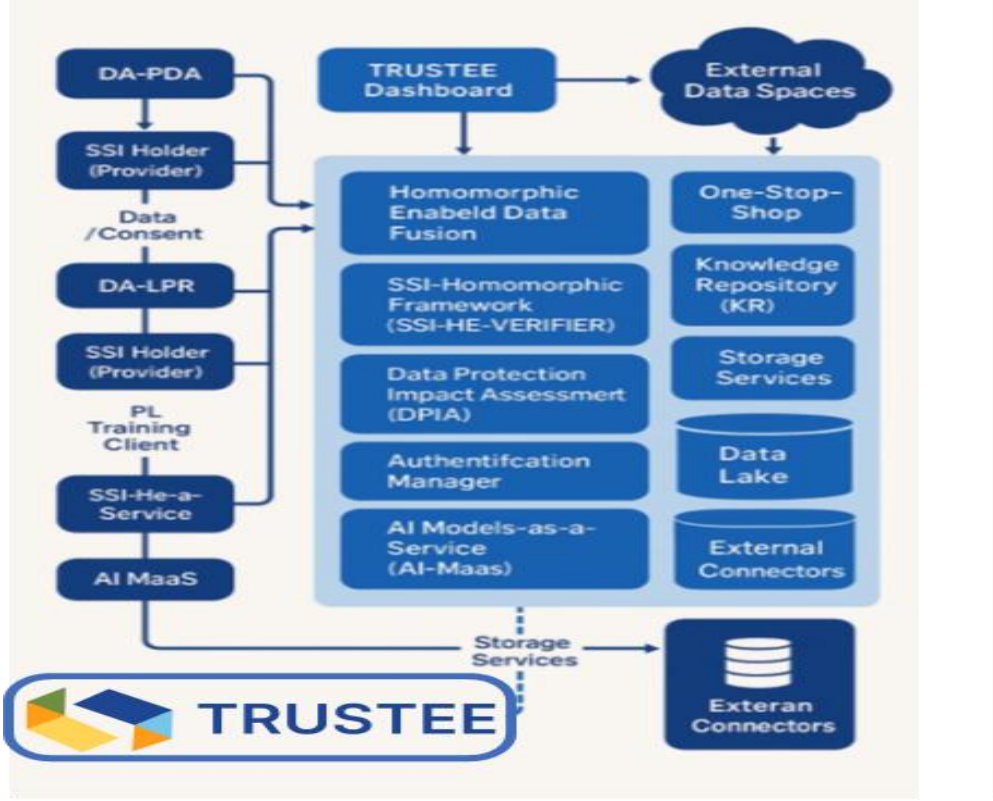

Figure 9: TRUSTEE project

3.5.2 TRUSTEE empirical use case results

Case study: a statistical seismic indicator used in the assessment of driving styles

As part of the Space use-case of TRUSTEE the consortium has focused its efforts on implementing the computation of the so-called waiting times in the encrypted domain, with the goal of providing a broad category of users with a versatile statistical indicator for time-series analysis. In a nutshell, given a generic time series, we define a waiting time as the shortest time interval needed to find an entry of value of at least $A + \delta$, with δ a known threshold, after a certain entry of value of at least A was observed. Analyzing the entire time series, entry by entry, we obtain the set of waiting times for a given value of δ and can determine their distribution. This distribution can be seen as a dynamic fingerprint of the process to which the time series pertains and – more importantly - is particularly useful to directly compare time-series pertaining to otherwise very different systems. Following a series of extensive computations with un-

encrypted datasets, we have seen that the distribution of waiting times exhibits common features for time-series as different as those pertaining the quakes on the moon and exchange rates of cryptocurrencies, including synthetic automotive datasets on position, velocity and acceleration of vehicles. The common aforementioned features refer to the distribution being scale-free like in the limit of small values of δ and Pareto-Tsallis like in the limit of large values of δ. Observing that the aforementioned results hold also for relatively small datasets, not only large ones, we are currently exploring the possibility of using the distributions of waiting times computed in the encrypted domain to help classify driving individual driving styles.

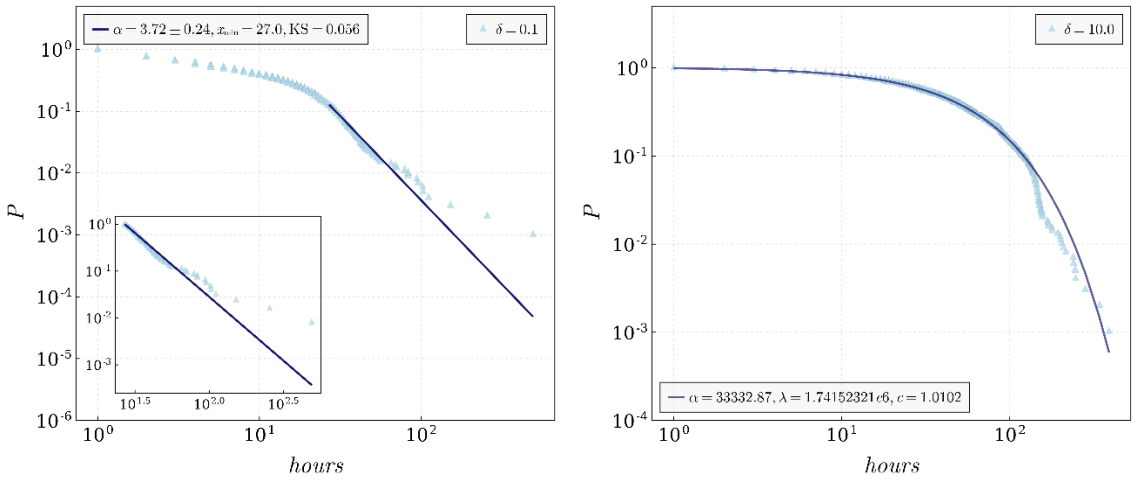

Figure 10: Results from TRUSTEE project statistical seismic indicator

Distribution of waiting times computed for the *x*-acceleration time series for δ=0.1 m/s^2 (left panel) and δ=10 m/s^2 (right panel). The distribution in the left panel can be approximated using a scale-free (or power-law) distribution, while the one in the right panel can be approximated using a Pareto-Tsallis distribution.

Assessing driving styles is a subject of intense research with studies covering disciplines as diverse as psychology and data science, sociology and biometrics, medicine and law. The concept of a driving style itself is heavily investigated from very different perspectives and no universal definition can be presented. With around 1 million people dying or being injured every year (according to World Health Organization, Road Traffic Injuries (2018)), the subject is of outmost interest which explains the aforementioned multiplicity of approaches and perspectives. Bridging between the psychological and sociological factors on one hand, and the purely technical aspects on the other hand is particularly problematic, more so as the sampling of the large-volume available data seems to bias one approach to the detriment of the others.

We consider the driving style to be a "habitual way of driving, which is characteristics for a driver or a group of drivers" (see F. Sagberg *et al.*, *A review of research on driving styles and road safety*, Hum. Factors, **57**, 1248–1275 (2015)), which is to say that we assume some robustness of a specific driving style across time (*i.e.*, times of the day, seasons, multi-annual patters, etc.) and landscapes (*i.e.*, urban and rural settings, on-road and off-road conditions, etc.) as well as some level of independence with respect to the vehicle that is driven. While we accept that a specific driving style is shaped by both individual characteristics (*i.e.*, gender, age, medical conditions, etc.) as well as a broader sociocultural context, we

restrict the analysis to what can be directly inferred from basic data like position, speed, and acceleration using the aforementioned distribution of waiting times as our indicator of choice. We therefore focus on direct processing of basic kinematic data excluding a variety of complementary data such as that pertaining to the positions of other vehicles and/or objects in traffic, that pertaining to the environment (*e.g.*, weather conditions and types of roads), that pertaining to the emotional status of the driver, that pertaining to his biometrical signals, and so on.

Results

Risk assessment

The protection of data concerning the position, velocity and acceleration of vehicles is significant for a number of reasons, which range from the vehicle health data relevant for insurance purposes to driver profiling at large based on the areas he visits, speed violations, etc. Despite the practical impact that automotive data has and the need to analyze it from the perspective of privacy and security, these implications are usually treated lightly, which, in turn, only increases the risks associated with automotive data processing.

From a different perspective the position, velocity and acceleration of vehicles are sensitive data that could be easily misused by potential attackers. For example, such data could be potentially misused for activity tracking, targeted marketing, cybercrimes, or burglaries. Protecting such data is therefore necessary for several reasons. First, sensitive data is protected by laws and regulations. Second, these data should also be protected for ethical reasons to protect the basic rights of individuals. Third, protecting automotive is important also to ensure public safety by ensuring that these data are used properly in emergency response situations and not exploited for malicious purposes.

Risk mitigation

The solution envisaged in TRUSTEE is that or computing the waiting times pertaining to automotive data using the privacy-preserving computing techniques developed in the project such that we circumvent by design at least some of the risks mentioned above.

3.5.3 TRUSTEE summary and future steps

TRUSTEE project thus aims to contribute to the next generation of secure, explainable, and interoperable data ecosystems. Building on its layered architecture and privacy-preserving foundations, future developments will aim to building beyond the current technological and methodological advancements, and the broader European Data Space initiatives, focusing on **scaling and operationalizing federated data sharing** across heterogeneous domains in real-world settings. This includes the **deployment of interoperable, privacy-preserving analytics at scale**, integration with **AI-driven decision support systems**, and **adaptive models** capable of dynamically enforcing context-aware data policies. Emphasis will be placed on **cross-sectoral validation**, particularly in high-impact areas such as education, space, **personalized healthcare, and energy management**, ensuring that solutions remain **environmentally sustainable, legally compliant, and socially responsible**. Additionally, further research will explore **hybrid AI techniques** combined with secure computation to enhance explainability and trustworthiness in sensitive data-driven applications.

3.6 REWIRE

3.6.1 REWIRE general description

The REWIRE project's present contribution emphasizes the final operational landscape of the REWIRE holistic Framework, with a brief, though coherent description of the Conceptual Architecture and Functional Components. REWIRE Architecture follows a seven-pillar framework as defined by its objectives, establishing core innovations and an end-to-end framework for addressing the main cybersecurity challenges of IoT ecosystems while safeguarding their operations holistically. The updated framework manages to enhance the cybersecurity posture during the Design and Runtime phases for the next-generation smart connectivity "Systems-of-Systems" (SoS) by safeguarding the entire lifecycle tested and validated in three heterogenous Use Cases (Smart Cities, Smart Automotive, and Smart Satellites).

REWIRE Final Conceptual Architecture

The final version of the REWIRE Conceptual Architecture is represented and described in Figure 11. Following the flow of the initial architecture release, the framework illustrates the Design-time phase (on the left-hand side) of the Figure, and the Runtime Phase (on the right-hand side), including the REWIRE-enabled edge device and the cloud-based backend infrastructure of the framework.

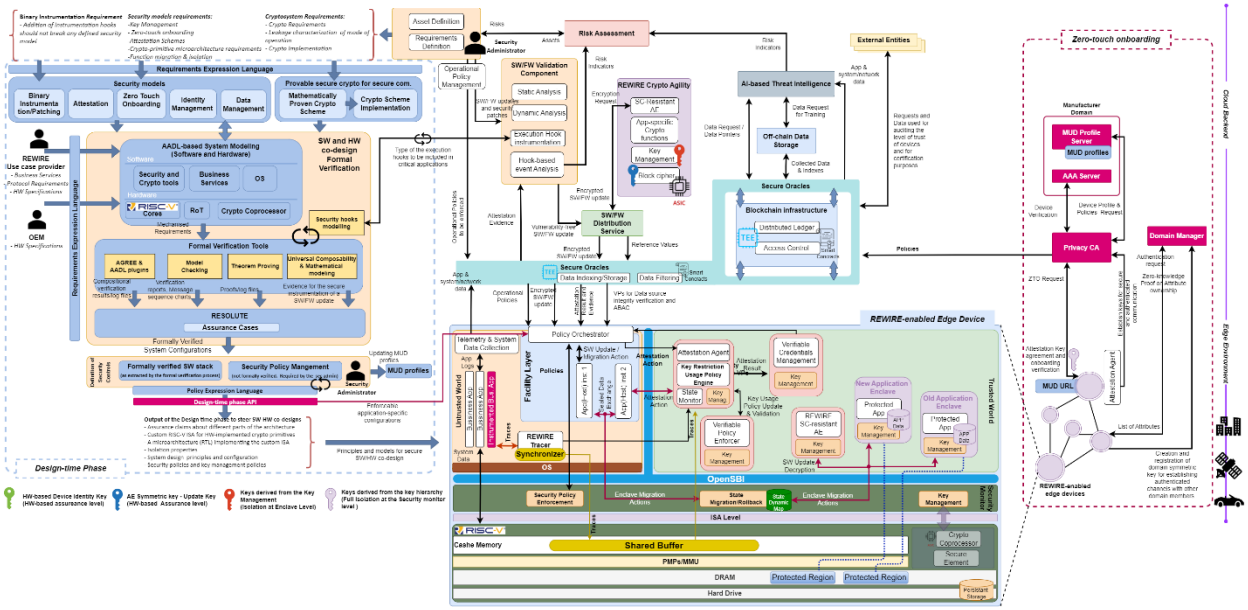

Figure 11: REWIRE Final Conceptual Architecture

Firstly, the Design Phase entails the initial instantiation and deployment of the REWIRE framework, as well as the expression of the overarching requirements of the system by the Security Administrator, the Use Case Providers, and the OEM. The main goal is essentially the identification of the trust boundary of the system, i.e., the part of the system that can be considered trustworthy, for which the Security Administrator can have the necessary guarantees that it will behave as expected. Parts of the system that remain outside this trust boundary cannot be formally verified during design-time, and need to be attested during runtime, based on the attestation of the respective properties linked to the correct and expected behaviour of the device.

More specifically, REWIRE follows a structured process to ensure the security and trustworthiness of devices within a given infrastructure. The first step involves defining requirements by identifying assets, their interconnectivity, and security needs. These requirements guide the system's architecture verification using AADL-based System Modelling, ensuring functional correctness before formal verification.

The Formal Verification component is crucial toward ensuring the security lifecycle of devices. It verifies protocols, mechanisms, and applications, defining the expected behavior of devices. This process refines the set of critical properties that must be attested during runtime to ensure devices operate securely. The first REWIRE framework version emphasizes using theorem-proving tools to verify LRBC, a lightweight block cipher for secure software updates in IoT devices. Future work will extend mathematical analysis to the Zero Touch Onboarding (ZTO) scheme and runtime cryptographic primitives using the Universal Composability (UC) method. The REWIRE Configuration Integrity Verification (CIV) scheme will also undergo formal verification to establish device trustworthiness.

When security updates or patches are needed, they are either integrated into the AADL-based system modeling process or sent directly for formal verification. This ensures patches maintain system integrity while refining trust boundaries and identifying new attestation tasks. Once verified, the next step is defining security controls, which are mitigation actions ensuring safe system operation. These controls form the basis of security policies, written in the Medium Security Policy Language (MSPL). REWIRE considers two types of policies (a) Operational policies, which guide software update tasks for mitigating identified risks and (b) Security policies, which dictate security enablers, such as attestation processes, to verify system properties at runtime.

The Formal Verification component also determines the SW/HW co-design system configuration and constructs a policy hash for key restriction usage policies. The deployment and enforcement of configurations occur through Manufacturer Usage Descriptions (MUDs). REWIRE extends this approach with eMUDs, ensuring synchronization with the latest formally verified system state. Updates to MUD profiles trigger automatic notifications and enforce security policies accordingly.

In case of a risk occurrence during the runtime, either through security enablers or new SW updates, the design phase is reinitiated to update formal verification processes. This iterative approach ensures that REWIRE remains adaptable to emerging threats, continuously maintaining the security and trustworthiness of devices in dynamic environments.

3.6.2 REWIRE empirical use case results

The REWIRE platform has been tested and validated in three complex and heterogeneous pilot IoT ecosystems (a) Smart Cities, (b) Automotive and (c) Smart Satellites. In the present section, we provide a first iteration of the results of the Smart Cities experiment (demonstrated by Odins), detailing improvements observed after the REWIRE framework implementation.

More specifically, the REWIRE's Smart Cities for Empowering Public Safety Use Case (UC) presents the challenges encountered in urban scenarios when deploying large-scale IoT networks. By aligning with the growth expectations of IoT networks and the innovations we find with technological paradigms incorporating critical human sectors (e.g., I2V, V2V, IIoT, remote surveillance), REWIRE manages to

mitigate multiple security gaps found in common IoT networks in the ”as-is” scenario. The demo showcases how different execution flows of an IoT system work and the shortcomings they present in the case of being used in a real critical system, represented in a lab environment with devices that simulate a real deployment. In this lab, the flows of “remote distribution of SW updates” and the “IoT network joining process” of the end devices have been specifically tested.

The evaluation goal of the Smart Cities for Empowering Public Safety Use Case is centred around the SW update user story. More specifically, this evaluation is focused on the ”one-to-one” scenario. In this regard, the first step is to provide the required environment setup for the Traffic control management scenario, including a set of electric adapters, a miniature scale, and a StarFive VisionFive2 device handling, both the REWIRE innovations and the business logic. Delving into details regarding this use case, the information about the traffic roundabout is reported in real-time to the monitoring and remote control application, showcased as a web platform.

In addition, the Smart City administrator is in charge of initiating the SW update process for the end device. The VisionFive2 board represents the IoT device installed on-site for the traffic control management. In the environment setup, the VisionFive2 board acts as the handler of all communications between our cloud services and the miniature of the traffic roundabout. In the same way, the VisionFive2 manages the cloud downlink communications to enable actions such as remote updates or execution of commands. The connection between the VisionFive2 and the roundabout miniature is wired, guaranteeing stable and reliable communication between the two systems. This approach ensures fast and uninterrupted data transmission to the different components of the miniature. The interaction with the components, such as lights, sensors and actuators, is done via the sysfs interface.

In the demo case, only the VisionFive2 board is used for the consolidation of all of the workload, meaning that the functions, previously distributed across different devices, are now concentrated on a single unit, simplifying the system architecture, but also requires enhanced processing power and resource optimization. In this new configuration, remote communications, the execution of the traffic system control logic, and the hardware-level control of the sensor system are all carried out by the VisionFive2. This includes managing data transmission between the traffic roundabout model and other external systems, making real-time decisions for traffic control, and directly overseeing the connected sensors and actuators.

The current evaluation is also focused on the IoT device joining phase to the network and management platform using five user stories, including a SW update on devices, the onboarding of new devices, efficiently attesting to device’s integrity and functional behaviour, the misbehaviour detection using AI, and the auditing of security posture. For the evaluation, quantitative measures were employed to assess performance and outcomes. The findings illustrate the implied real-world advantages of REWIRE, highlighting key insights gained throughout the implementation.

More analytically, the UC aimed to enhance operational efficiency and decision-making in complex smart cities environments through advanced data collection, analytics, and automation. Initially, the project faced challenges related to inefficient data management, delayed decision-making processes, and limited predictive capabilities. The evaluation involved a Baseline Assessment for measuring key performance indicators (KPIs) before implementation, an Implementation Phase for deploying the REWIRE framework and integrating it into operational workflows, and the Post-Implementation Evaluation with specific dedicated KPIs.

The REWIRE Framework was seamlessly integrated into the Use Case infrastructure providing automation and decision-making advancement. By leveraging REWIRE, the implementation process addressed key challenges and delivered tangible benefits, significantly enhancing smart cities operations. During the implementation process, the REWIRE Framework managed to enhance interoperability and system integration, since the REWIRE framework facilitated seamless integration of heterogeneous data sources, enabling real-time data collection and sharing across systems. AI and Automation mechanisms for data analysis allowed for faster and more informed decision-making, reducing reliance on manual processes, with AI-driven predictive maintenance leading to a 25% decrease in unplanned downtimes, preventing costly system failures and increasing the operational efficiency and error reduction. In addition, the framework's automation capabilities reduced manual intervention by 20%, freeing up resources for high-value tasks. As far as the scalability and future-proofing are concerned, the REWIRE's modular architecture ensured that the Use Case could scale seamlessly as new technologies and requirements emerged.

3.6.3 REWIRE summary and future steps

In the present contribution, REWIRE showcases the final conceptual architectural framework and the initial results from Smart Cities Use Cases, while several key actions will be evaluated to enhance the security and monitoring of devices within all the planned deployments. The REWIRE framework provides multiple advantages as a framework and from the initial complete results of the implementations. The REWIRE Framework was instrumental in ensuring the success of the ODINs use case, delivering improved efficiency, automation, and decision-making while ensuring scalability and ease of adoption. These benefits illustrate REWIRE's critical role in modernizing operations and processes, making them more resilient, efficient, and future-proof. For this specific Use Case, the next steps include the finalization of the pending integration activities (e.g., with LRBC) and the evaluation of the functionalities of interests such as zero-touch onboarding and the monitoring hooks. On top of that, we will focus on the one-to-many SW updates.

4 Future Directions and Joint Activities

As most of the contributing projects near their completion, the results and collaborations established will serve as a springboard for continued joint activities. A primary goal is to promote open science by ensuring that key outputs such as datasets, technical components, and reproducible methodologies are made openly accessible. This will support broader research and innovation efforts, enabling external stakeholders to validate findings and extend their use in new domains.

The collaboration among partners has fostered a shared vision and complementary expertise that will be leveraged in future initiatives. Efforts are already underway to co-develop new proposals and secure follow-up funding that will allow for the extension of empirical work, enhancement of core technologies, and the application of solutions in additional high-impact use cases. These may include areas such as healthcare, mobility, critical infrastructure, and beyond.

In parallel, there is a strong focus on contributing to European and international discussions on standardization, data governance, and cybersecurity. Joint work will continue through community engagement, technical contributions, and aligned development activities to ensure that outcomes remain relevant, adaptable, and compliant with evolving legal and ethical frameworks. This sustained collaboration aims to support the creation of resilient, trustworthy, and interoperable digital ecosystems.

EU projects details

encrypt	https://encrypt-project.eu/
AI4CYBER	https://ai4cyber.eu/
CERTIFY	https://certify-project.eu/
CROSSCON	https://crosscon.eu/
REWIRE	https://rewireproject.eu/
TRUSTEE	https://www.linkedin.com/company/horizon-trustee/