D1.3

Legal and Ethical Issues and Guidelines

Project number:	101070627
Project acronym:	**REWIRE**
Project title:	Rewiring the Compositional Security Verification and Assurance of Systems of Systems Lifecycle
Project Start Date:	1st October, 2022
Duration:	36 months
Programme:	HORIZON-CL3-2021-CS-01
Deliverable Type:	Report/Other
Reference Number:	HORIZON-CL3-2021-CS-01-101070627/ D1.3 / v1.0
Work package:	WP 1
Actual Submission Date:	27/01/2025
Responsible Organisation:	8BELLS
Editor:	Christiana Kyperounta
Dissemination Level:	PU
Revision:	v1.0
Abstract:	This deliverable presents the second and final version of the REWIRE Data Management plan. In addition to detailing the DMP, it also puts forth the steps taken for ensuring its conformance to all legal regulations and guidelines as it pertains to the use and management of all extracted research data and methodologies/algorithms. Finally, D1.3 touches upon the ethical dimensions of trust reasoning on what is needed for not only making a system trustworthy but how it can also be perceived as one by the users in order to enhance user adoption and acceptance. This requires moving beyond vague claims about improving trust, to developing a comprehensive methodology for assessing trustworthiness without impeding user privacy.D1.3 defines trust explicitly, link it to ethical values, and list those steps followed as the REWIRE roadmap towards converging on this interplay between technical and ethical considerations of trust.
Keywords:	**Keywords:** Ethics, Legal Issues, Trust, Trustworthiness

The project REWIRE has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101070627.

Versioning and contribution history

Version	Date	Author	Notes
0.1	28/08/2024	Christiana Kyperounta (8BELLS)	ToC
0.2	23/09/2024	Christiana Kyperounta (8BELLS)	Updates on Sections 1 and 2
0.3	24/09/2024	Athanasios Charemis (UBI)	Updates on Sections 3,4,5
0.4	30/10/2024	Sofianna Menesidou (UBI)	Review and polishing of all Sections
0.5	13/11/2024	Sofianna Menesidou (UBI)	Updates on Sections 1 and 2
0.9	18/12/2024	Sofianna Menesidou (UBI)	Updates on Sections 3,4,5
1.0	23/01/2025	Sofianna Menesidou (UBI), Christiana Kyperounta (8BELLS), Athanasios Charemis (UBI)	Finalisation of Section 4 and Final Review

Executive Summary

The objective of this deliverable is to serve as the reference and guide for action, in legal and ethical matters within the REWIRE project. As an outcome of the Task 1.3 on the ethics and legal monitoring, aims to determine the process and measures undertaken on how research will be conducted and executed within the overall REWIRE's context. More specifically, the present deliverable aims to provide a holistic view of the REWIRE Legal and Ethical issues and establish a coherent framework for the ethical and legal dimensions of the project. To this end, the document presents the updated REWIRE Data Management Plan, offering the latest updates on the types of research data managed and processed in the context of REWIRE, both in relation to the developed artifacts and the experimentation activities relevant to the project's use cases. In addition, the deliverable elaborates on the four-step roadmap established by REWIRE to integrate the legal and ethical dimensions of trust across the Compute Continuum. Last but not least, it is important to mention that, in compliance with the Data Management policy, all the referenced research data will be open source.

Contents

List of Tables

List of Figures

1. Introduction

1.1 Scope and Purpose

The document aims to support the consortium partners in the REWIRE project in identifying and addressing ethical issues inherent in their work throughout the project. It offers a holistic view of the REWIRE Legal and Ethical issues and establish a coherent framework for the ethical and legal dimensions of the project. The objective of this Deliverable is to analyse relevant European laws and regulations relevant to the scope of REWIRE, including Privacy and Data Protection, Cybersecurity, Artificial Intelligence (AI) and General Product Safety Regulation (GDPR). It aims to establish legal and ethical guidelines and identify potential legal and ethical challenges. Moreover, it outlines procedures for assessing the need for mitigating ethical issues and provides tools for managing ethical issues in research activities.

1.2 Relation to other WPs and Deliverables

Deliverable 1.3 “Legal and Ethical Issues and Guidelines” is the primary legal deliverable in REWIRE, building on the work performed in T1.3 “Legal and Ethics Monitoring”, which addresses how research will be executed in the project regarding the ethics issues during the implementation of the REWIRE in collaboration with the project partners and the independent Ethics Committee. It is a living document, subject to updates based on the progress of project activities.

This document will apply the previously expressed theoretical principles to the REWIRE project and its deliverables. The work performed in under this Deliverable follows up D1.2 “REWIRE Data Management Plan” [1]. Additionally, it provides valuable legal and ethical input to activities within WP3 “Design-Time Compositional Security & Safety Assurance for SoS”, WP4 “Runtime Secure Field Devices Execution Layer”, WP5 “Secure Data Sharing & Trust Management Logic for Modular SoS” and WP6 “REWIRE Framework Integration & Use Cases Demonstration”.

1.3 Deliverable Structure

This deliverable is structured as follows:

- **Section 1** describes the purpose and scope of this deliverable in accordance with the relation with other Work Packages, as well as the structure of the deliverable;
- **Section 2** provides an updated Data Management Plan, including the research datasets per partner as part of the REWIRE project;
- **Section 3** references to the relevant EU and international legislation that the consortium will make sure to abide by;
- **Section 4** presents the four-step roadmap towards analysing ethical and legal dimensions of trust across the Compute Continuum;
- **Section 5** concludes the deliverable.

2. Management or Research data

This chapter presents an updated version of the project's Data Management Plan (DMP) documented in D1.2. Specifically, here we provide a table including the research datasets per partner as part of the REWIRE project.

2.1 Update of the DMP

In the D1.2 submitted in M6; the consortium outlined the envisioned strategy for managing research data throughout the project. This included detailing the data management life cycle for data to be collected, processed, or generated by REWIRE, with a strong emphasis on adhering to the FAIR data principles, ensuring research data is findable, accessible, interoperable, and reusable. Through publications as well as the Open-Research Development Plan (OSD). In addition, all technical and use case partners provided information about the data expected to be generated through a structured questionnaire. Such a questionnaire encompasses all the research innovations of REWIRE around trust extensions in RISC-V such as the AI-based Misbehavior Detection; the Secure Blockchain Oracles for the secure and privacy-preserving management of data sharing agreements and the investigation of lightweight crypto algorithms with storing key leakage resilience. REWIRE is the first to converge the two worlds of formal verification for enabling the target deployment and configuration of security and trust functions in next-generation safety-critical ecosystems. Also, the questionnaire includes the evaluation and benchmarking activities conducted across the three use cases. The responses resulted in a preliminary dataset list, including details on data type (e.g., about research algorithms, research datasets as well as evaluation datasets from our use cases), format (e.g., .xls, .csv, .txt, .docx and .pdf) and management tools.

The implemented data management life cycle has proven effective during the project's first year. REWIRE will disseminate this data as open-source through open-access scientific publications or open-source libraries. The same methodology will be extended to the results of benchmarking activities once the evaluation phase concludes.

2.1 Research Datasets per Technical Artifact

Following the first release of the REWIRE framework, an updated and more concrete overview of the datasets associated with each technical component is provided in the following Table 1. More specifically, Table 1 provides an update for the type of data that we have now started extracting from the evaluation in the use cases. It has to be noted, that all components are based on simulated data, which is open source and not subject to any specific restrictions. The use cases will be tested in controlled environments; thus, no real data will be generated. Instead, only synthetic data will be produced and made available as open-source. All research data related to the technical components will be published on the REWIRE project's GitLab and through open-access scientific publications.

Table 1 REWIRE research datasets per technical component

REWIRE Activity	Type of data	Format	Management Tool	WP	Partner
Formal Verification toolset pipeline	Logs to formally verify the functional correctness of the SW update over an authenticated encryption channel and the implicit attestation framework. Also, logs to evaluate the achievement of use case requirements through RESOLUTE. More specifically, we have created (open-source) RESOLUTE models for verifying	Logs	Git	WP3	Collins

	the achievement of the security requirements of the use cases through the REWIRE framework.				
SW update over an authenticated encrypted communication channel	Desing and implementation of LRBC algorithm and performance metrics and evaluation for the leakage resilience is ASIC environments.	Algorithm	Git	WP3	UCL
Threat modeling and SW/FW vulnerability analysis	Vulnerability Analysis Tool and detailed logs of the vulnerability analysis in context of use case application binaries.	Logs	Git	WP3	SECURA
Trusted Computing Base (RISC-V) - SW Update and migration protocol	SW Update and migration protocol for secure lifecycle management.	JSON	RISC-V keystone acting as the underline RoT of StarFive experimentation RPi	WP4	RUB
Trusted Computing Base (RISC-V) - Attestation Enablers	Attestation enablers and SW abstractions for supporting the self-issuance of security claims Example of the trustworthiness claims output is provided in Appendix 5.1. It has to be state here that REWIRE has two types of TCB instantiations the Genesys 2 board and to Vision Five board.	Source code	RISC-V keystone acting as the underline RoT of StarFive experimentation RPi	WP4	UBI
Trusted Computing Base (RISC-V) - Key Management	Crypto-extensions for managing complex key hierarchies as part of the Security Monitor of a TEE instantiated in a RISC-V architecture	Source code	RISC-V keystone acting as the underline RoT of StarFive experimentation RPi	WP4	NEC
Risk Assessment	Risk assessment services for continuous calculation of required trust level (RTL) based on zero-day threats and exposure of adequate interfaces to consume or interact with deployed tools (e.g., AI-based Misbehavior Detection). Example of the output is provided in Appendix 5.2.	Streams of timestamp values as part of a JSON message format	Git	WP4	UBI
Policy-compliant BC Infrastructure	Management of the data-sharing agreements and the auditability for the trust-related information supported through secure oracles (i.e., TownCrier, FPC). Especially with the SGX-enabled secure oracle (i.e., FPC) and their enrichment with advance crypto.	Source code	Git	WP5	TUD/UBI
AI-based Misbehavior Detection	Set of classification models for managing/processing time series evidence extracted from devices in the AI-based Misbehavior Detection engine. The output of the engine will be used as a trust source.	Source code, misbehavior report and logging results of smart satellites and smart cities misbehavior indicators in JSON format transmitted sent	Git and a program to process the streams of timestamp generated device evidence	WP5	S5

		to the risk assessment			
Dataset in the context of Smart City	Emulation of the traffic-lights operation throughout a specific time period	Simulated dataset for evaluating traffic light management in smart cities intersection	Simulation tools providing wide range of application traces supporting the project experimentation activities	WP6	ODINS
Data from in-vehicle network	In the context of an eIDAS system	Simulated dataset for evaluating SW update and migration controls in in-vehicle networks	Simulation tools providing wide range of application traces supporting the project experimentation activities	WP6	KEN
Dataset in the context of Smart Satellites		Simulated dataset capturing the characteristics of low orbit satellites and the process of SW update through the consecutive transmission of multiple packets	Simulation tools providing wide range of application traces supporting the project experimentation activities	WP6	LSF

3. Legal & Ethical Considerations of REWIRE

The REWIRE consortium will ensure that the project adheres to all relevant legal frameworks and directives, protecting personal data, maintaining cybersecurity standards and ensuring the ethical use of the data within the European Union (EU). This section highlights the key legislative frameworks that will guide the REWIRE consortium's practices throughout the project's lifecycle. The primary regulations governing research data are derived from the General Data Protection Regulation (GDPR) - for completeness see Section 3.1 where all details are put forth on how REWIRE takes all necessary steps to ensure GDPR conformance. In a nutshell, REWIRE complies with these regulations by ensuring that no Personally Identifiable Information (PII) or sensitive data is collected during its research and evaluation activities. Instead of using real use cases, REWIRE employs simulated and emulated "hardware-in-the-loop" methodologies. This approach, with realistic yet synthetic data, not only ensures compliance to data protection regulations but also provides the flexibility needed for testing and evaluation across diverse scenarios and conditions. Recall that one of endmost goals of REWIRE is to provide all core trust enables for the establishment of a generic Trust Assessment Framework (TAF) focusing on the RISC-V enabled far edge devices and to analyze its impact on the trust levels indicated by the TAF for a given service. This strategy, based on simulated data, facilitates the comprehension of how stakeholders can potentially exploit these trust calculations to improve the trustworthiness and effectiveness of their services without any privacy breach.

The overarching objective of REWIRE is to accommodate the definition of a generic trust assessment methodology tailored to the unique needs of the Compute Continuum environment. To this end, as elaborated in the following chapter, REWIRE investigates not only the technical aspects but also considers the social and ethical dimensions of trust. The aim is to understand how these dimensions interact and influence end-user and stakeholder acceptance of such a framework. Addressing the social and ethical facets of trust, REWIRE examines newly introduced standards and regulations, particularly drawing on the EU AI Act. This Act identifies seven key areas for trustworthy AI which REWIRE adopts as a foundation for establishing trustworthiness principles within the Compute Continuum domain. A detailed discussion of this roadmap is provided in Chapter 4. Table 2 below provides a summary of all the consulted legal frameworks and directives.

Table 2: EU Legislations, directives and other instruments consulted in REWIRE

Document Title	Date	Reference
Council of Europe Convention No. 108	28/01/1981	https://rm.coe.int/1680078b37
Charter of Fundamental Rights of the European Union (CFR)	12/12/2007	https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:12012P/TXT
The General Data Protection Regulation (GDPR) - (EU) 2016/679	27/04/2016	https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679
ePrivacy Directive - 2002/58/EC	12/07/2002	https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32002L0058
The Data Protection Law Enforcement Directive (EU) 2016/680	27/04/2016	https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016L0680
The Data Act Regulation - (EU) 2023/2854	13/12/2023	https://eur-lex.europa.eu/eli/reg/2023/2854
The Data Governance Act - (EU) 2022/868	30/05/2022	https://eur-lex.europa.eu/eli/reg/2022/868/oj
The Open Data Directive -	20/06/2019	https://eur-lex.europa.eu/eli/dir/2019/1024/oj

(EU) 2019/1024		
NIS 2 Directive - (EU) 2022/2555	14/12/2022	https://eur-lex.europa.eu/eli/dir/2022/2555
The Cybersecurity Act – Regulation (EU) No 526/2013	17/04/2019	https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32019R0881&qid=1726557928072
EU Artificial Intelligence Act - (EU) 2024/1689	13/06/2024	https://eur-lex.europa.eu/eli/reg/2024/1689/oj
General Product Safety Regulation - (EU) 2023/988	10/05/2023	https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32023R0988&qid=1726558246152

To sum up, the research data, including also those that stem from the evaluation of our use cases, is based on simulated and emulated environments, thus, they do not include any PII or sensitive information. Also, our testing and validation of REWIRE's work with key stakeholders (e.g. advisory board and use case partners) not only allows us to check if our findings are valid, but also for allowing REWIRE to achieve of the endmost goal of helping identify and develop a trust assessment methodology that can technically provide the necessary means for establishing trust that it also worthies of human trust.

3.1 Overview of Personal Data Protection and Privacy Legislation at EU level

Data protection legislation has been thoroughly examined, considering that REWIRE involves various stakeholders, including the project partners' personnel, the External Experts Advisory Board (EEAB) and external people participating mostly at the REWIRE workshops, events and meetings. This places privacy and data protection at the forefront of the consortium's concerns. The protection of data has evolved significantly over the years, shaped by various European instruments, including a convention, a charter and major legal instruments that specifically address data protection, privacy, and governance. This section provides an overview of the relevant framework.

The Council of Europe Convention No. 108, also known as the Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data, was the first binding international legal instrument dedicated to data protection. Signed in 1981, it seeks to protect individuals' fundamental rights and freedoms in relation to personal data processing, particularly focusing on the cross-border flow of data. It applies to both the public and private sectors, making it a pioneering document in the global discussion on data protection. The modernized version, Convention 108+, updates these protections to address challenges posed by modern technologies, ensuring continued relevance. The convention has influenced data protection laws globally, including the GDPR. The strengths and limitations of the convention have also been examined in the context of United Nations work [2].

Articles 7 and 8 of the Charter of Fundamental Rights of the European Union (CFR) provide specific guarantees concerning privacy and data protection within the European Union. Article 7 mirrors the protections of Article 8 of the ECHR, ensuring the right to private and family life, while Article 8 explicitly establishes the right to the protection of personal data. It stipulates that data must be processed fairly, for specified purposes, and based on consent or other legitimate grounds laid down by law. Individuals have the right to access their data and the right to rectification. These provisions are binding on EU institutions and Member States when implementing EU law and have served as the backbone for the development of EU data protection legislation, particularly the GDPR.

The General Data Protection Regulation (GDPR) (EU) 2016/679 is the most comprehensive and widely recognized data protection regulation in the world. Adopted in 2016 and enforced since May 25, 2018, harmonizes data protection laws across all EU Member States. The GDPR establishes core principles, such as lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, and the integrity and confidentiality of personal data. It also emphasizes the principles of accountability and data protection by design and by default. Under the GDPR, organisations must implement appropriate technical and organisational measures to ensure data security and mitigate risks

to the rights and freedoms of individuals.

The ePrivacy Directive 2002/58/EC, also known as the Directive on Privacy and Electronic Communications, complements the GDPR by focusing on privacy in the digital environment. It governs the processing of personal data in electronic communications, including rules for cookies, spam, and the confidentiality of communications. Websites must obtain user consent before storing or accessing information, such as cookies, on their devices. The EU is currently working on a proposed ePrivacy Regulation, which is expected to modernize these rules to reflect new digital communication technologies and align more closely with the GDPR.

The Data Protection Law Enforcement Directive (EU) 2016/680 applies to the processing of personal data by law enforcement agencies. This directive ensures that personal data used for the prevention, investigation, detection, or prosecution of criminal offenses is adequately protected. It grants individuals rights similar to the GDPR, including access to their data, the right to rectification, and restrictions on the transfer of data to third countries. The directive recognises the unique needs of law enforcement while safeguarding the fundamental rights of individuals.

The Data Act Regulation (EU) 2023/2854, adopted in 2023, promotes the sharing of non-personal data between businesses, consumers, and public bodies to stimulate innovation. It ensures greater control for users over their data and facilitates access to data generated by connected devices. While focused on non-personal data, it complements existing personal data protections and stimulates fair competition in the data economy.

The Data Governance Act (EU) 2022/868 establishes a framework for data sharing across sectors and Member States, with an emphasis on trust, transparency, and ethical data use. It introduces the concept of data intermediaries and encourages data altruism, where individuals or businesses voluntarily share their data for public interest purposes.

Finally, the Open Data Directive (EU) 2019/1024 aims to make public sector information widely available for reuse. It encourages the release of high-value datasets that can foster innovation, particularly in areas such as artificial intelligence, while ensuring that the reuse of public sector data does not infringe on privacy rights.

3.1.1 Confidentiality and Data Protection Measures within REWIRE

As initially documented in the REWIRE Data Management Plan [1], the REWIRE project takes robust confidentiality and data protection measures to ensure that all data, particularly Personally Identifiable Information (PII), is handled securely and ethically. Given the nature of the data collected (e.g., simulated/emulated without any PII), it is anticipated that most data will not be made publicly available unless the necessary ethical clearances and participant agreements are secured. In cases where data can be shared, it will be hosted through the project's file repository located at the coordinator's premises, ensuring that access is controlled and secure.

For any data release, a formal procedure is followed. Interested parties must submit a request to the Project Coordinator, outlining their intended use of the dataset. They will then receive a "terms and conditions" document, which must be signed and returned before any data is shared. All data released will be accompanied by relevant documentation to ensure transparency and facilitate proper use. These measures ensure that data sharing aligns with ethical standards and participant consent, and that confidentiality is maintained throughout.

Additionally, the project is committed to following the European Commission's Guidelines on Open Access to Scientific Publications and Research Data. A combination of Gold and Green Open Access strategies will be employed, ensuring scientific publications are accessible while maintaining appropriate protections for personal and sensitive data. Shared data will be anonymized and stored in formats such as Word, PDF, or Excel to ensure accessibility while safeguarding participant confidentiality.

In terms of data storage and management, all data will be securely deposited on Microsoft SharePoint within the project's repository. Additional security measures include storing a backup instance of all data on the coordinator's account, ensuring that only authorized personnel have access. Personal data will be processed in full compliance with the EU and international regulations mentioned in Section 3.1, with

strong emphasis on adherence to the GDPR. This includes ensuring that data processing is adequate, relevant, secure, and only retained for the necessary duration of the project, following the principles of data quality and participant rights.

The data subjects in REWIRE are classified as the following:

- **External People participating in Operative Tasks of REWIRE:** participants in interviews, Q/As sessions, webinars, workshops or meetings organised by REWIRE. This category includes individuals who engage in interviews, Q&A sessions, webinars, workshops, or meetings organized by REWIRE. These participants are involved in operational activities that support the project's objectives, but no data release related to their participation is scheduled. All data collected from these participants will be securely stored and handled according to the project's confidentiality protocols, ensuring their personal information remains protected and is not shared without explicit consent.
- **External Experts Advisory Board Members:** A formal collaboration agreement has been established from the outset of the project for members of the External Experts Advisory Board. This agreement outlines the expected contributions, defines the scope of collaboration, and, crucially, details the personal data protection measures to be applied.
- **REWIRE Partners' Personnel:** The personnel from REWIRE's partner organizations are subject to the data management and protection guidelines as outlined in the REWIRE Data Management Plan [1].

3.2 Relevant Legislation, Directives, and Guidelines

3.2.1 Cybersecurity Legislation at EU level

In the context of REWIRE, understanding and adhering to cybersecurity legislation and policies is vital for building a robust legal and regulatory foundation. This framework not only ensures compliance with existing laws but also fosters a secure environment for the protection of the platform and its associated data. Legislation and policies provide guidelines and requirements for various aspects of cybersecurity, including data protection, incident response, breach notification, and privacy. With the ever-increasing threat landscape, European regulations and legislations, such as the Network and Information Security Directive 2 (NIS2) and the EU Cybersecurity Act, play a crucial role in defining the security standards that must be met.

The NIS 2 Directive (EU) 2022/2555 aims to enhance cybersecurity within the European Union (EU) to improve the functioning of the internal market. It introduces several measures to achieve this goal. Member States are required to adopt national cybersecurity strategies and set up competent authorities, cyber crisis management bodies, designated points of contact for cybersecurity issues, and Computer Security Incident Response Teams (CSIRTs).

Additionally, the Directive establishes rules and responsibilities for the exchange of cybersecurity information. Lastly, it outlines supervisory and enforcement obligations for Member States. More specifically, in this context, the MUD is an IETF standard aimed to define the intended behaviour of the device through Access Control Lists (ACLs), to restrict communication to/from a certain device. While MUD was recently standardised (March 2019), it has received strong interest from the research community and standardisation entities worldwide. REWIRE is not only fully aligned with the MUD but also extends the use MUDS through the REWIRE's zero-touch onboarding protocol. The NIST has also recommended MUD files to complement security credentials to reduce the attack surface2. MUD is focused on the definition of network access control policies. Therefore, these restrictions can be straightforwardly enforced through the Software-Defined Networking (SDN) paradigm.

The EU Cybersecurity Act is a legislative framework established by the European Union to enhance cybersecurity measures and foster a more secure digital environment within its Member States. It aims to strengthen the EU's ability to prevent, detect, and respond to cyber threats by establishing a European cybersecurity certification framework for products, services, and processes. The Act also establishes a permanent mandate for the European Union Agency for Cybersecurity (ENISA). It provides increased

responsibilities and resources to support Member States in their cybersecurity efforts. The support is delivered through a range of activities spanning five key areas:

- Expertise: provision of information and expertise on key network and information security issues.
- Policy: support to policy making and implementation in the Union.
- Capacity: support for capacity building across the Union (e.g. through trainings, recommendations, awareness raising activities).
- Community: foster the network and information security community (e.g. support to the Computer Emergency Response Teams (CERTs), coordination of pan-European cyber exercises).
- Enabling (e.g. engagement with the stakeholders and international relations).

The EU Cybersecurity Act is crucial in protecting critical infrastructure, businesses, and individuals against cyber threats. As the EU Cybersecurity Act aims to ensure the proper functioning of the internal market while achieving a high level of cybersecurity, cyber resilience, and trust within the European Union, it achieves this by addressing two main aspects:

- establishing objectives, tasks, and organisational matters related to ENISA, the European Union Agency for Cybersecurity, and
- providing a framework for establishing European cybersecurity certification schemes.

It is important to note that the Act respects the competences of Member States in areas such as public security, defense, national security, and activities related to criminal law.

3.2.2 AI Legislation at EU level

REWIRE will not function predominantly upon the basis of Artificial Intelligence (AI). Yet, specific REWIRE AI-based mechanisms will be developed, which will be based only on threat intelligence data, attestation-related data, and software updates. The mechanisms will not operate on any PII that may breach a user's privacy. Therefore, the REWIRE consortium deemed as relevant and important aspect to investigate the AI policy landscape at EU level. More specifically, REWIRE examines the EU AI Act, which outlines 7 key areas for trustworthy AI and employs them as a basis for determining the trustworthiness principles in the Compute Continuum field. This section sets the scene for the detailed analysis in Chapter 4.

The EU AI Act is the EU's flagship new AI regulation which entered into force on August 1st, 2024. The Act imposes risk and technology-based obligations on organizations that develop, use, distribute or import AI systems in the EU, with significant fines for non-compliance. Its goal is to promote the ethical and trustworthy development, deployment, and use of AI systems while safeguarding fundamental rights and values. The AI Act outlines requirements related to the classification of AI systems, risk assessment, transparency, data governance, and accountability. It also addresses specific areas such as biometric identification, high-risk AI systems, and remote biometric identification systems.

Definition of an artificial intelligence system (AI system)

> *"A machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments."*

According to the Act, high-risk AI systems[1] should indeed be designed and developed in a manner that allows effective oversight by natural persons while the AI system is in use. The objective of this oversight is to prevent or minimize risks to health, safety, or fundamental rights from both intended uses and reasonably foreseeable misuse. This requirement is crucial for ensuring accountability in critical applications such as biometric identification, employment decision-making, and law enforcement.

[1] High-risk systems refer to those involved in critical infrastructure, education, employment, or law enforcement.

To ensure human oversight, the AI Act suggests implementing one or more of the following measures to enhance accountability and mitigate potential risks associated with high-risk AI systems:

- Incorporating oversight features into high-risk AI systems from the design phase, if technically feasible. These could include the ability to intervene or override decisions made by the AI.
- Defining appropriate human oversight measures to be implemented by the users of the system, allowing them to maintain control over the AI's functioning and decisions. This could involve providing users with necessary training or creating processes to ensure human judgment is part of critical decision-making.

The AI Act introduces new obligations for entities both within and outside the EU. An interactive tool, the "Compliance Checker" [3] is available to help determine whether an AI system falls under these requirements.

3.2.3 General Product Safety Regulation

The REWIRE project's objectives are also closely aligned with the European Union's General Product Safety Regulation (GPSR) [4], particularly regarding the continuous security assessment and lifecycle management of IoT devices. The GPSR mandates that products placed on the market, including IoT devices, must maintain a high standard of safety throughout their lifecycle, a principle mirrored by REWIRE's focus on security-by-design and continuous security assessment. By incorporating real-time monitoring and oversight through cryptographically verifiable security proofs, REWIRE ensures that IoT devices remain secure and compliant with safety regulations, reducing risks associated with device misuse or cyberattacks.

The General Product Safety Regulation (GPSR) is indeed a pivotal component of the European Union's product safety legal framework, aimed at ensuring that all products placed on the EU market, including emerging digital products like Internet of Things (IoT) devices, meet stringent safety requirements. The regulation mandates that manufacturers, distributors, and importers ensure their products are safe for consumers throughout the product's lifecycle, from design to disposal. It emphasizes proactive risk management, requiring that any foreseeable risks, including those posed by new technologies, be identified and mitigated early in the development process. Moreover, the GPSR incorporates provisions for market surveillance and rapid recall mechanisms for unsafe products, and it holds economic operators accountable for ensuring compliance with safety standards. By harmonizing product safety laws across the EU, the GPSR ensures a high level of consumer protection and enhances the smooth functioning of the internal market, particularly as technology continues to evolve.

3.3 Alignment of REWIRE with relevant regulations

In a nutshell, the core regulations that govern research data stems from the GDPR. REWIRE, complies to GDPR since no sensitive and personal information is collected throughout its research and evaluation activities. As already mentioned, all the performed experiments are based on simulated data that emulate realistic use case scenarios. This is an intentional design choice since has more flexibility to be able to stress test the various aspects of the system. It has also to be noted here that all the research artifacts are stored on the project's repository and will be released as open source either via publications or via technical libraries.

On top of that, REWIRE examines the newly introduced standards and regulations and more specifically the EU AI Act, which outlines 7 key areas for trustworthy AI and employs them as a basis for determining the trustworthiness principles in the Compute Continuum field. More details of this roadmap can be found in Chapter 4.

3.4 Engagement with Compute Continuum stakeholders

As part of REWIRE's social and ethical analysis and the on-going discussions with set of expertise for the technical advancements of the project, REWIRE engages with several stakeholders from both industry and academia. More specifically, REWIRE not only leverages the expertise and the already established network of its partners and its advisory board, but also experts beyond the consortium for a better and more complete feedback. Table 3 summarizes the deferent types of involved stakeholders and its relevance to REWIRE.

Table 3 Stakeholders and their relevance to REWIRE

Type of stakeholder	Relevance
Open source HW (e.g., RISC-V) related stakeholders	Discussions beyond the consortium are in progress. More specifically with Markku-Juhani Saarinen from University of Tampere who is also the chair of the Crypto Extensions WG at RISC-V International which is tasked with developing official RISC-V ISA extensions that improve the efficiency of modern public-key crypto.
Attestation related stakeholders	Discussions and collaboration with Chris Fenner who is the head of TPM WG from TCG and Ahmad Atamli from NVDIA for TDI SP are in progress.
HW-based RoT and Virtualization related stakeholders	Discussions and collaboration with consortium members of CROSSCON H2020 project. More precisely with Bruno Crispo from University of Trento regarding Virtualization aspects in RISC-V environments.

4. REWIRE Roadmap Towards Analyzing Ethical & Legal Dimensions of Trust across the Compute Continuum

This Chapter focuses on the **interplay analysis between trust and ethics**, focusing on the concept of trust and its explanation to the users. Figure 1 below summarizes this interplay and can be divided in two parts separated by the trust evidence (e.g., Integrity, Safety, Consistency, Usability, Verifiability, Transparency, Accuracy, Privacy, Security, Reliability, etc.). The above-part of the figure depicts the collection of trustworthiness evidence and their interpretation from a technical standpoint – how they can be leveraged as part of the trust assessment process for the trust relationship of interest between the Trustor and Trustee. The below-part depicts how these decisions may affect the users. Essentially, this does not only consider how the use of trustworthy systems can affect the digital life of users but it considers the important question of how users understand trust: *Would the provision of such strong trust extensions lead the enhanced adoption of the system by the users?* REWIRE's roadmap sets out the context in which we examine the dimensions and determinants that affect the trust that humans can understand, so that a system like REWIRE can offer ways of conceptualizing trust beyond the technical dimension.

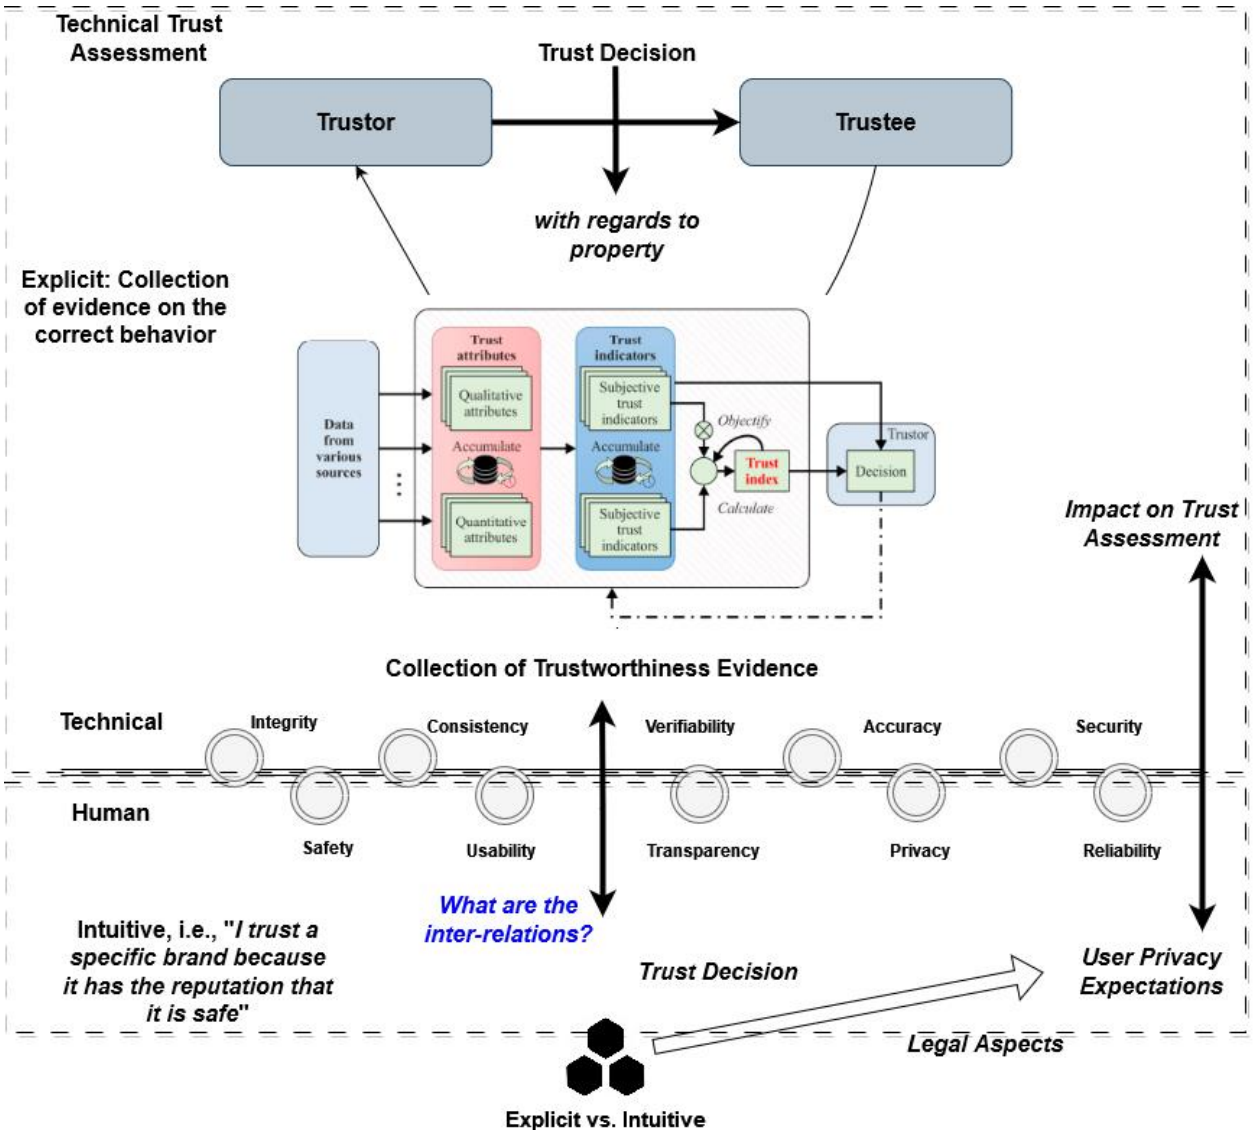

Figure 1 – Interplay Analysis between Trust and Ethics

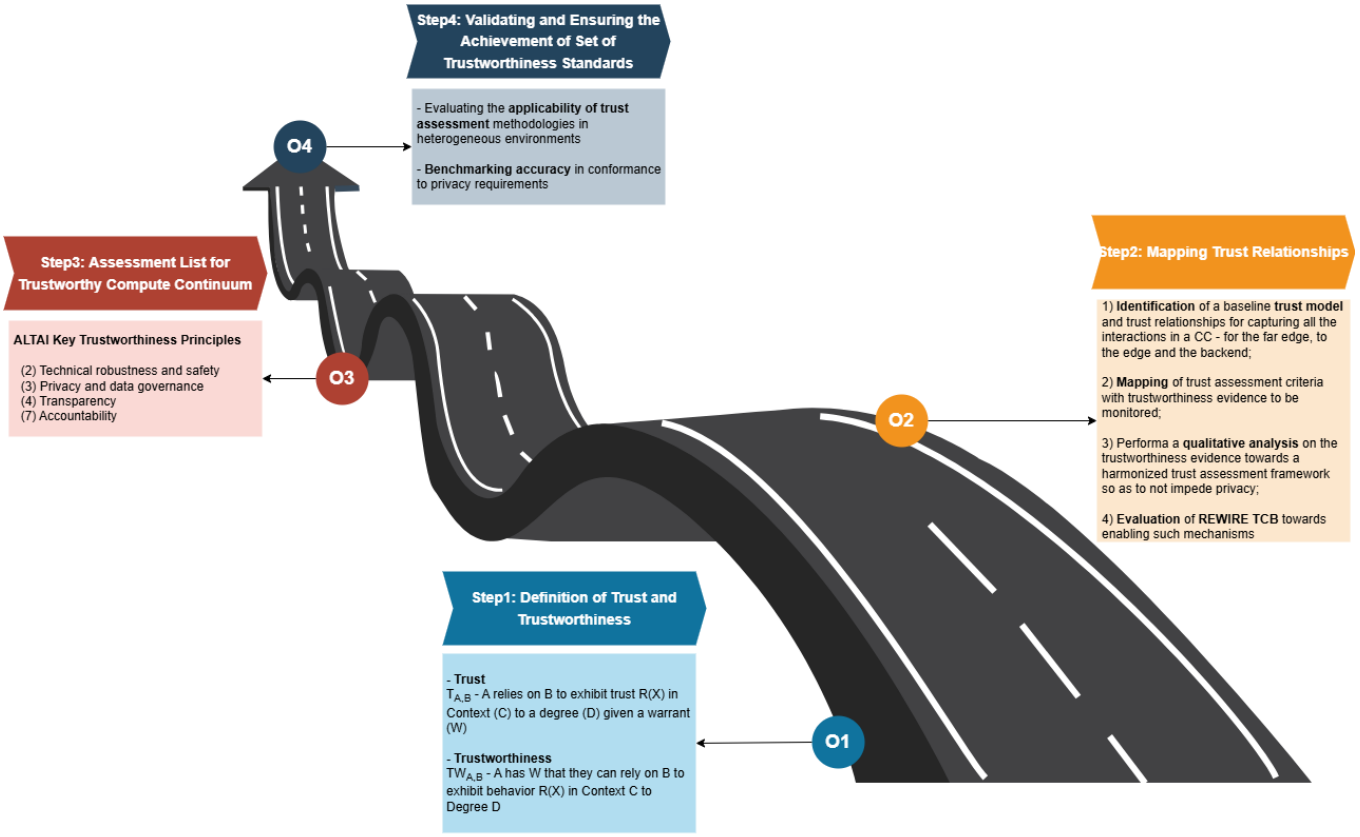

Figure 2 Roadmap to Analyze Ethical and Legal Dimensions

In order to answer this question (on whether the expansion of the trust concept as an integral part of any system) can enhance user adoption, based on a set of trust extensions and a cybersecurity framework, we look in detail in this interplay between trust and users' privacy expectations. It has to be noted that there is an implicit interplay between what a system can achieve technically (in a trustworthy manner) and how this relates to what is valued by users. In what follows, we proceed with the underpinnings between the technical and ethical/legal considerations of trust across the Compute Continuum. On top of that, we identify the key areas that capture the social/ethical dimensions of trust as they apply to such safety-critical applications. These conceptual and ethical discussions set the foundations to enable key users to identify what is valuable within the REWIRE system

The first aspect we need to consider in this context, is that the **trustworthiness evidence** generated by the device for a trust assessment, **does not impede the privacy of the device** compared without the trust assessment. For instance, identifying a vehicle through, vehicle fingerprinting, based on information from trustworthiness evidence, should not be possible. Hence, there needs to be an interplay between trust and privacy related to the information extracted by the vehicles. In other words, the generated trustworthiness evidence does not provide to the adversary any advantage in compromising the devices' or end users' privacy profile. This leads to the adoption of **abstraction layers** in order to harmonize what type of information is need for trustworthiness to satisfy this requirement. Meaning that, there is no need to provide the evidence, based on which the quantification of trust was performed, as is but the trust characterization as an abstraction layer. For instance, users care that a system has integrity based on verifiable and observable evidence but not on the actual evidence themselves.

Another aspect is, to identify what is the **appropriate set of evidence** that technically can allow a high accuracy for a trust assessment framework but at the same time can also have a direct impact on the safety of the system in an explainable way to the end user, so that is can increase their adoption. In the context of REWIRE, we investigate the alignment conformance of REWIRE with the technical accuracy and robustness reequipments as highlighted in the EU AI Act. High accuracy directly affects the type of evidence dives the trust extension that needs to be deployed, which directly affects the trust extensions and security controls that are designed as part of the device's TCB.

In the context of REWIRE, we envision to **perform a qualitative analysis of the type of evidence that a trust assessment framework needs to have for achieving high accuracy and at the same time be ethically sound** and to **evaluate if the REWIRE TCB can enable such mechanisms**. Figure 2 depicts these four steps of REWIRE's roadmap towards analyzing ethical/legal dimensions of trust. Steps 1 and 2 will be carried out through an iterative process. Initially, broad definitions of trust and trustworthiness are established, followed by the identification of stakeholders. As more specific definitions of trust and trustworthiness are developed (in the context of the envisioned use cases), this may lead to the identification of additional stakeholders, creating a cyclical refinement of the process. The ultimate goal is to produce a detailed map of trust relationships among all relevant stakeholders in the various layers of the CC – starting from the far-edge devices and leading to the edge and cloud-based backend where additional service-graph-chains may be deployed. ***This map will clarify who is trusting whom, for what purposes, and outline the trust expectations, contexts, degrees, and evidence required to establish and maintain those trust relationships***. Below we provide an overview of these four steps.

Step 1: This step focuses on **defining the broad requirements** for a system's trustworthiness. It involves applying the **general concepts of trust and trustworthiness** to the specific system being evaluated. It requires clearly outlining the conditions of the trust relationship between the trustor (A) and the trustee (B). Towards this direction, (a) A and B must be identified within the trust relationship; (b) A's goal that B is expected to fulfill within the particular context of the expectation must be defined; (c) the level of trust A seeks and requires from B, including the reasons, evidence, or data necessary for A to make an informed trust assessment must by clarified. Similarly, to evaluate whether B is worthy of A's trust, it is essential to (a) specify the goal and context in which B is expected to perform; (b) determine the degree of trust B must achieve to successfully fulfill the goal and (c) identify the reasons, evidence, or data B can provide to A to demonstrate its trustworthiness or lack thereof.

Step 2: This step focuses on **mapping the trust relationships within the system**, **identifying all trustors and trustees** involved. In the context of REWIRE, while we have considered a generic methodology, this instantiation has taken place in the context of the envisioned use cases. In this step, it is important to recognize that A and B may share multiple trust relationships. For example, B might be expected to perform X, Y, and Z, or carry out Z in various contexts such as Context 1, Context 2, and Context 3. Additionally, given the complexity of the system, there are likely to be multiple stakeholders forming a network of trust. For instance, (a) A may need to trust B, C, and D; (b) B might need to trust E, F, and G; (c) C could require trust in H, I, and J, and so forth. For each trust relationship, it is essential to clearly define (a) what the trustor expects from the trustee; (b) the specific context(s) of the expectation, (c) the required degree of trust and (d) the reasons, evidence, or data necessary to support the trust relationship. This detailed mapping ensures a comprehensive understanding of all stakeholder interactions within the trust network.

Step 3: This step focuses on **defining the standards that determine what constitutes a trustworthy relationship**. For example, if the robustness of a system within the Compute Continuum is questioned, there must be clear and specific standards that define and clarify the concept of robustness. These standards are typically external and independent of both the trustor and trustee, emphasizing the critical distinction between trust and trustworthiness. In addition, this step bridges trust and ethics by establishing ethical benchmarks that define what should be considered worthy of trust. Within the REWIRE framework, this involves a **qualitative analysis of the type of evidence** required by a trust assessment framework to **ensure high accuracy while remaining ethically sound**. These ethical considerations are informed by the seven key principles identified by the EU High-Level Expert Group (HLEG) on Artificial Intelligence and the Assessment List for Trustworthy Artificial Intelligence (ALTAI). Even though, REWIRE is not strictly bound to AI, the foundational ethical values outlined in these seven key principles are directly relevant to REWIRE and to the Compute Continuum as a whole.

Step 4: This last step focuses on the **validation and testing of the previous three steps**. In general, this step is a way to validate the process, and test the specific questions and responses with a range of stakeholders, to tighten and improve the trust method in practice. At the end, we will have a specified REWIRE trust methodology which can be applied equally to REWIRE or any Compute Continuum system.

4.1 Step 1: Definition of Trust and Trustworthiness

As aforementioned, the first step is the definition of trust and trustworthiness as a general concept (based on the standards) and its concretization in the context of Compute Continuum. The main question here revolves around the ethical and social dimensions of trust, and their interpretation by the stakeholders. When conceptualizing trust, it is essential to first define the specific aspects of trust being addressed and then distinguish it as a concept separate from, yet closely connected to, trustworthiness. The key difference lies in perspective that trust pertains to the judgment a trustor makes about whether to place trust in a subject, whereas trustworthiness focuses on the qualities of the subject that determine whether a subject of trust is worthy of that trust.

In the context of REWIRE, the **general definition of trust** is intended to be general, and generalizable and is developed in such a way as to be adapted and specified to specific applications and particular trust relationships. In other words the general definition of trust that describes a trust relation between A (trustor) and B (trustee), given at $T_{A,B}$ is the following:

$T_{A,B}$ – A relies on B to exhibit trust R(X) in Context (C) to a degree (D) given a warrant (W)

In addition, we have to recognize that trust comes in degrees. This means that trust judgments will be affected by both uncertainty and disbelief. For instance, when considering the way that people trust each other, A will trust B more or less, according to his/her experiences. However, in other cases, trust might be considered in a binary fashion and not in degrees. This is also the case in REWIRE, where the core objective is to explore a way of providing reasons, evidence, or data, that give a component, user or other stakeholder that warrant. Next, we note how these trust relations depend on W, that the trust judgments being formed by A about B are accurate, leading to the notion of trustworthiness. Towards this direction we need to answer: What is sufficient reason, evidence or data to justify A trusting B? When building trust into complex systems in the Compute Continuum, it is not enough to simply say A trusts B, we must also ask if B is worthy of A's trust. To this end, we must also reformulate the general definition of trust to be a **general definition of trustworthiness**.

$TW_{A,B}$ – A has W that they can rely on B to exhibit behaviour R(X) in Context C to Degree D

This updated definition is concerned with a trust relation between A and B, in which A has warrant to rely on B to exhibit particular behaviour in a specific context to a particular degree. This definition is now focused on whether B is worthy of that trust. It is crucial in the complex systems of Compute Continuum to consider both trust relations and trustworthiness relations.

4.2 Step 2: Mapping Trust Relationships

Having provided these general definitions of trust and trustworthiness, this step establishes the **link between trust and ethics**, whereby ethics is concerned with setting some standards about what ought to be considered worthy of trust. Thus, each identified stakeholder relationships need to be specified, such that it is clear what each trustor is expecting of each trustee, in which context(s), to what degree, and what reasons, evidence, and data are needed. However, first is crucial to identify who is the trustor, the trustee, the expected behaviour, the context of the trust relationship, the degree and the reasons of trust. This step includes a comprehensive **mapping of the trust relationships among all the relevant stakeholders**, who they are trusting to do what, and what those trust expectations, contexts, degrees, and warrants require. In REWIRE, we provide the mechanisms that provide the trustworthiness evidence in a verifiable manner (e.g., attestation mechanisms). Further, we consider that there are criteria that a Compute Continuum device must meet in order to be worthy of trust:

> *Is Device B worthy of the trust of other devices in a Service Graph Chain (SGC), with regard to robustness or safety of the overall service?*
> *Can Device B produce/communicate information accurately?*

Is the information produced and/or communicated by Device B accurate, free from external interference?

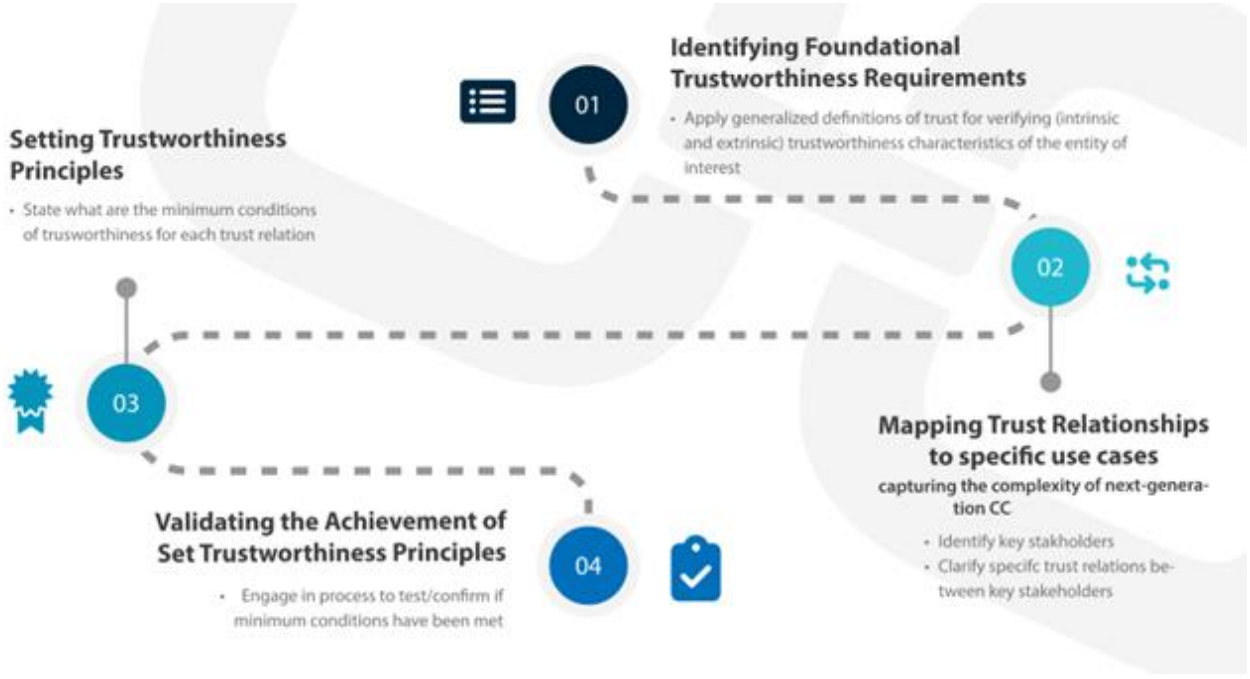

Figure 3: Roadmap for understanding, mapping, setting, and ensuring trust relations in Compute Continuum

In order to answer all these questions and determine if a device can be trusted or not, we need to qualitative assess CC devices for trustworthiness (see step 3 in next sub-section). Recall, also the four-step roadmap (see Figure 3).

4.1 Step 3: Assessment List for Trustworthy Compute Continuum

In the third step of the roadmap, REWIRE will perform the qualitative analysis of the type of evidence that a trust assessment framework needs to have for achieving high accuracy and at the same time be ethically sound based on the seven key principles identified from the EU High-Level Expert Group (HLEG) on Artificial Intelligence, the Assessment List for Trustworthy Artificial Intelligence (ALTAI)[2]. It has to be noted here that the HLEG group and ALTAI were specifically for AI, and REWIRE is not strictly bound to AI considerations. However, the seven core values are highly relevant to REWIRE and generally to the Compute Continuum. Even though REWIRE is not fundamentally connected to AI regarding the notion trust, still the key values for a system operation are commonly envisioned.

In this context, in what follows we provide an extended analysis of each of the directly applicable key areas and adapt them for the embedded systems and IoT elements across the Compute Continuum, with particular attention to the REWIRE use cases. ALTAI's key EU values are: (**1) Human agency and oversight; (2) Technical robustness and safety; (3) Privacy and data governance; (4) Transparency; (5) Diversity, non-discrimination, and fairness; (6) Environmental and societal well-being; (7) Accountability**. In REWIRE the directly applicable key values from the technical standpoint are: the (2) Technical robustness and safety; (3) Privacy and data governance; (4) Transparency and (7) Accountability. These seven key areas must be adapted to REWIRE to ensure applicability. However, as we move from key area 1 out to 7, the values become more generalised, and hence the adaption needs to become more general. That is, the first two key areas are much more easily translatable to specific engineering requirements and existing standards, whereas the other five areas are applicable to the Compute Continuum more generally, and must also be understood by reference to existing law, legislation, policy, as well as emerging and evolving social norms. Table 4 summarize how REWIRE adheres to these principles.

[2] https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html

Table 4 Mapping of ALTAI principles to REWIRE

ALTAI-based principles	How REWIRE achieves it
(2) Technical robustness and safety	Compute Continuum systems must rely on technical robustness to both ensure and assure users of their reliability and safety. This approach is essential to mitigate and prevent even unintentional harm effectively. REWIRE provides state of the art attestation mechanism towards establishing robustness and safety for Compute Continuum systems
(3) Privacy and data governance	To ensure full adherence to privacy and data protection standards, robust data governance mechanisms must be established. These should focus on maintaining data quality and integrity while providing legitimized and controlled access to the data. REWIRE's approach supports zero-touch onboarding aligned with the SSI concept.
(4) Transparency	Data ad systems should operate with transparency, supported by traceability mechanisms to ensure accountability. To do so, REWIRE capitalizes on Blockchain technology, where every transaction is linked through cryptographic hashes, allowing detailed and transparent audits.
(7) Accountability	To ensure accountability, robust mechanisms must be established. REWIRE's Blockchain technology through its immutability feature records all transactions and changes in a secure and accountable manner.

Below we perform a first investigation for the two requirements (e.g., #2 Technical Robustness and Safety and #3 Privacy and Data Governance) that are core in the context of REWIRE. D6.2 will include more details for the rest of the requirements.

REQUIREMENT #2 Technical Robustness and Safety

Compute Continuum rely on technical robustness to both ensure, and assure, users of their reliability and safety. A trustworthy Compute Continuum device and system is one in which the components work as expected. A crucial requirement for achieving trustworthy systems is their dependability (the ability of the overall system services in which trust is justified), robustness - produces stable results when facing changes such that it doesn't break or fail easily, and resilience means recovering quickly from challenges. Technical robustness requires that systems are developed with a preventative approach to risks and that they behave reliably and as intended while minimizing unintentional and unexpected harm as well as preventing it where possible. The questions in this section address four main issues: a) security; b) safety; c) accuracy; and d) reliability, fallback plans and reproducibility.

Table 5 Security and Robustness

Security and Robustness
Did you assess potential forms of attacks to which the system could be vulnerable?
Did you consider different types of vulnerabilities and potential entry points for attacks?
Did you put measures in place to ensure the integrity, robustness, and overall security of the system against potential attacks throughout its lifecycle?
Did you perform penetration testing on the system?
Did you inform end-users of the duration of security coverage and potential updates?

Table 6 Safety

Safety
Did you define risks, risk metrics and risk levels of the system per use case?
Did you identify the possible threats to the system and the corresponding consequences?
Did you assess the dependency of a critical system's decisions on its stable and reliable behaviour?

Did you plan fault tolerance via a duplicated system or another parallel system?
Did you develop a mechanism to evaluate when the system has been changed to merit a new review of its technical robustness and safety?

Table 7 Accuracy

Accuracy
Could a low level of accuracy of the system result in critical or even damaging consequences?
Did you put in place measures to ensure that the data used to develop the system is up-to-date, of high quality, complete and representative of the environment the system will be deployed in?
Did you put in place a series of steps to monitor, and document the system's accuracy?
Did you consider whether the system's operation might lead to adversarial attacks?
Did you put processes in place to ensure that the level of accuracy of the system to be expected by end-users and/or subjects is properly communicated?

Table 8 Reliability, Fall-back plans and Reproducibility

Reliability, Fallback plans and Reproducibility
Could the system cause critical, adversarial, or damaging consequences in case of low reliability and/or reproducibility?
Did you put in place verification and validation methods and documentation such as logging to evaluate and ensure different aspects of the system's reliability and reproducibility?
Did you define tested failsafe fallback plans to address system errors of whatever origin and put governance procedures in place to trigger them?
Did you put in place a proper procedure for handling the cases where the system yields results with a low confidence score?
Is your system using online continual learning?

REQUIREMENT #3 Privacy and Data Governance

Privacy, a fundamental right, is closely tied to the principle of harm prevention, especially in the context of Compute Continuum. Safeguarding privacy requires robust data governance measures that ensure the quality and integrity of the data. Privacy concerns typically arise when information is produced, used, or linked to individuals. Table 9 provides how to self-assess the impact of the CC systems on privacy and data protection. These are interconnected fundamental rights that are also intrinsically linked to the broader fundamental right to personal integrity, encompassing both mental and physical well-being. Table 10 provides how to self-assess the adherence of the CC systems to various elements concerning data protection.

Table 9 Privacy

Privacy
Did you consider the impact of the system on the rights to privacy and data protection?
Depending on the use case, did you establish mechanisms that allow flagging issues related to privacy concerning the system?
Have you considered if your product/system utilizes PIIs and sensitive information, and do you take the requisite steps to ensure that such data is processed in a lawful and ethical manner?
Have you considered legal and ethical implications of any non-sensitive information your system might be using?

Table 10 Data Governance

Data Governance
Is your system developed, by using or processing PIIs and sensitive information?
Did you implement the right to withdraw consent, the right to object and the right to be forgotten into the development of the system?
Did you consider the privacy and data protection implications of data collected, generated or processed over the course of the system's life cycle?
Did you consider the privacy and data protection implications of the system's processed non-sensitive data?
Did you align the system with relevant standards (e.g. ISO25, IEEE26) or widely adopted protocols for (daily) data management and governance?

4.2 Step 4: Validating and Ensuring the Achievement of Set of Trustworthiness Standards

In this last step of the REWIRE trust methodology, we focus on involving key stakeholders of the consortium and gathering their insights. This process comprises three sub-steps: **identifying** those key stakeholders of interest (we are considering the definition of "users" from the perspective of service providers that wish to expose services with inherent trust), **validating** trustworthiness standards through stakeholder input, and **ensuring** those standards are met (or addressing gaps if they are not). Identifying key stakeholders involves selecting representatives from specific groups closely associated with the Compute Continuum (CC). Four broad categories are considered as potential contributors: (a) design experts, (a) development and production specialists, (c) oversight and accountability professionals, and the broader group of (d) non-expert users and community members. Each group brings unique perspectives and expertise to the trust assessment process, ensuring a holistic and robust evaluation of CC systems. These are listed below:

(a) Design Experts: These are individuals with the technical expertise required to comprehend and contribute to the foundational theoretical and technical aspects of Compute Continuum components, processes, and systems. Examples include software engineers developing the trust assessment framework, cybersecurity experts evaluating risks of malicious intrusions, and other technical specialists, such as those within the REWIRE project team and the Advisory Board members.

(b) Development and Production Experts: This category includes professionals working for organizations that utilize CC components, processes, and systems in production, such as original equipment manufacturers (OEMs), vendors, and security service providers. While they can be seen as users, they are distinct from end users or consumers. Example of this group is the REWIRE use case partners.

(c) Oversight and Accountability Experts: This group comprises institutions and individuals responsible for creating and enforcing relevant standards and policies. They assess the trustworthiness of CC systems based on established or emerging standards. Examples include standards bodies focusing on safety or IoT cybersecurity and experts in areas like privacy, data management, and legislation (e.g., GDPR, the European AI Act). While their evaluations of CC trustworthiness extend beyond REWIRE's scope, their input would be vital in broader trust assessments.

(d) Non-Expert Users and Community Members: This set involves the general public, including non-experts who either use or are affected by the final CC product. Essentially, they are the consumers or customers. As this group could encompass entire populations, its inclusion in trust assessments lies outside REWIRE's scope but remains crucial for comprehensive evaluations.

In the context of REWIRE, for validation purposes, the first two stakeholder groups will participate in brief testing sessions and surveys. They will address the general questions outlined in Steps 1 and 2, followed

by the specific inquiries in Step 3. These validation activities aim to ensure the questions are relevant and easy to understand. The REWIRE trustworthiness framework is designed to be adaptable, evolving in response to stakeholder input and recommendations. Simultaneously, stakeholders will refine and evaluate the trustworthiness standards for specific CC components, processes, and systems. This effort aligns with the goal of identifying and translating the ethical principles outlined in the HLEG's seven key areas for AI into a defined Required Trust Level (RTL).

5 Conclusions and Future Work

The present deliverable corresponds to the T1.3 “Legal and Ethics Issues and Guidelines” and aims to outline the overall holistic plan of REWIRE to address the potential legal and ethical risks. To this end, the present document represents a blueprint for the identification, definition, description and mitigation of the associated risks within REWIRE project.

More specifically, the deliverable begins with a brief and comprehensive overview of the legal and ethical aspects of the REWIRE project. Furthermore, provides the final version of the data management plan, demonstrating the REWIRE’s approach to managing and processing research data, either regarding the artefacts or in terms of the evaluation activities. Overall, the present deliverable manages to address and deal with all potential legal and ethics related aspects and activities of the REWIRE project focusing also on the interplay analysis of trust and ethics. More specifically, the present document evaluates the trust and trustworthiness from an ethical and legal standpoint, providing a detailed four-step roadmap.

List of Abbreviations

Abbreviation	Translation
AI	Artificial Intelligence
CERTs	Computer Emergency Response Teams
CSIRTs	Computer Security Incident Response Teams
CRF	Charter of Fundamental Rights of the European Union
EB	Executive Board
EEAB	External Experts Advisory Board
ENISA	European Union Agency for Cybersecurity
EU	European Union
FW	Firmware
GA	General Assembly
GDPR	General Data Protection Regulation
GPSR	General Product Safety Regulation
HLEG	High-Level Expert Group
IoT	Internet of Things
MSA	Multi-factor Authentication
OSD	Open-Research Development Plan
PII	Personally Identifiable Information
RoT	Root-of-Trust
SW/HW	Software/ Hardware

References

[1] REWIRE. Data Management Plan. Deliverable D1.2, The REWIRE Consortium, 03 2023.
[2] United Nations Publication. 2016. Data protection regulations and international data flows: Implications for trade and development. United Nations Conference of Trade and Development. Available at: https://unctad.org/system/files/official-document/dtlstict2016d1_en.pdf
[3] EU AI Act Compliance Checker. Available at: https://artificialintelligenceact.eu/assessment/eu-ai-act-compliance-checker/
[4] Regulation (EU) 2023/988 on general product safety, also known as the General Product Safety Regulation (GPSR) will replace the General Product Safety Directive (GPSD) 2001/95/EC after a transitional period that will end in December 2024.
[5] REWIRE. Project Handbook. Deliverable D1.1, The REWIRE Consortium, 12 2022.
[6] REWIRE. Rewire operational landscape, requirements, and reference architecture - initial version. Deliverable D2.1, The REWIRE Consortium, 12 2023.
[7] Mckenna, D. and Automotive, B. (n.d.). Available at: https://www.nxp.com/docs/en/whitepaper/Making-Full-Vehicle-OTA-Updates-Reality-WP.pdf.
[8] ISO. ISO 31000, Risk management. Available at: https://www.iso.org/iso-31000-risk-management.html
[9] REWIRE. Dissemination, Communication, Clustering and Exploitation Activities – Initial Version. Deliverable D7.2, The REWIRE Consortium, 01 2024.
[10] High Level Expert Group on Artificial Intelligence. Ethics Guidelines for Trustworthy AI (hereinafter, “the Ethics Guidelines”), 2019. Available at: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai.

Appendix A: Examples of REWIRE research datasets

5.1 Trustworthiness Claims JSON file Example


```
{
  "tchReport": {
    "trusteeReports": [
      {
        "attestationReport": [
          {
            "appraisal": 1,
            "claim": "secure-boot",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "control-flow-integrity",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "access-control",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "secure-communication",
            "timestamp": "2024-07-01T14:05:15Z"
          }
        ],
        "trusteeID": "Service 1"
      },
      {
        "attestationReport": [
          {
            "appraisal": 1,
            "claim": "runtime-integrity",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "application-isolation",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "secure-communication",
            "timestamp": "2024-07-01T14:05:15Z"
          },
          {
            "appraisal": 1,
            "claim": "keystore-integrity",
            "timestamp": "2024-07-01T14:05:15Z"
          },
```

```
          {
            "appraisal": 1,
            "claim": "network-intrusion-detection-report",
            "timestamp": "2024-07-01T14:05:15Z"
          }
        ],
        "trusteeID": "Infrastructure Node 1"
      }
    ]
  },
  "evidence": {
    "timestamp": "2024-07-01T14:05:15Z",
    "signatureAlgorithmType": "ECDSA-SHA256",
    "signature":
"30440220676790c8092f9830afd0141100dcf364f08a742a10ef0c4580bb761395121625022046a
3b05f282dda822c9b88b0c0732adab319074b0fc3a84137ee2f7edf21f2ba",
    "keyRef": "tch_public_key"
  }
}
```


5.2 Risk Assessment JSON file Example


```
{
  "content": [
    {
      "id": "663368075a7f965400954f55",
      "taraRiskassessment": {
        "id": 102,
        "name": "TARA R2"
      },
      "businessPartner": {
        "id": 1,
        "name": "Ubitech"
      },
      "attackPathProfile": {
        "id": 2,
        "name": "Information Disclosure on Camera",
        "businessPartner": {
          "id": 1,
          "name": "Ubitech"
        },
        "assets": [
          {
            "id": 502,
            "name": "Camera",
            "tags": [],
            "relationships": [
              {},
              {},
              {}
            ],
            "attributes": [],
            "taraAttributes": [
              {
                "key": {
                  "id": "CONFIDENTIALITY"
                },
                "value": {
                  "id": 1,
                  "name": "The image from the camera and location of the vehicle can be accessed by externals"
```

```
            }
          },
          {
            "key": {
              "id": "INTEGRITY"
            },
            "value": {
              "id": 2,
              "name": "The location of the vehicle can be manipulated by externals"
            }
          },
          {
            "key": {
              "id": "INTEGRITY"
            },
            "value": {
              "id": 3,
              "name": "The pictures from the camera are manipulated and it causes collision"
            }
          }
        ],
        "tagsToString": "N/A"
      }
    ],
    "targetAsset": {
      "id": 502,
      "name": "Camera"
    },
    "targetProperty": {
      "id": "CONFIDENTIALITY",
      "description": "Confidentiality"
    },
    "elapsedTime": {
      "id": "LESS_THAN_1_WEEK",
      "description": "LESS_THAN_1_WEEK"
    },
    "expertise": {
      "id": "EXPERT",
      "description": "EXPERT"
    },
    "knowledge": {
      "id": "PUBLIC",
      "description": "PUBLIC"
    },
    "windowsOfOpportunity": {
      "id": "EASY",
      "description": "EASY"
    },
    "equipment": {
      "id": "BESPOKE",
      "description": "BESPOKE"
    },
    "attackFeasibilityRating": {
      "id": "M",
      "description": "M"
    }
  },
  "damageScenarioProfile": {
    "id": 1,
    "name": "The image from the camera and location of the vehicle can be accessed by externals",
    "businessPartner": {
      "id": 1,
```

```
        "name": "Ubitech"
      },
      "safetyImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "financialImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "operationalImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "privacyImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "overallImpact": {
        "id": "SEVERE",
        "description": "Severe"
      }
    },
    "riskLevel": {
      "id": "H"
    }
  },
  {
    "id": "663368075a7f965400954f56",
    "taraRiskassessment": {
      "id": 102,
      "name": "TARA R2"
    },
    "businessPartner": {
      "id": 1,
      "name": "Ubitech"
    },
    "attackPathProfile": {
      "id": 5,
      "name": "Information Disclosure on GNSS",
      "businessPartner": {
        "id": 1,
        "name": "Ubitech"
      },
      "assets": [
        {
          "id": 502,
          "name": "Camera",
          "tags": [],
          "relationships": [
            {},
            {},
            {}
          ],
          "attributes": [],
          "taraAttributes": [
            {
              "key": {
                "id": "CONFIDENTIALITY"
              },
              "value": {
                "id": 1,
```

```
            "name": "The image from the camera and location of the vehicle can be accessed by externals"
          }
        },
        {
          "key": {
            "id": "INTEGRITY"
          },
          "value": {
            "id": 2,
            "name": "The location of the vehicle can be manipulated by externals"
          }
        },
        {
          "key": {
            "id": "INTEGRITY"
          },
          "value": {
            "id": 3,
            "name": "The pictures from the camera are manipulated and it causes collision"
          }
        }
      ],
      "tagsToString": "N/A"
    },
    {
      "id": 501,
      "name": "Global Navigation Satellite System (GNSS)",
      "tags": [],
      "relationships": [
        {},
        {}
      ],
      "attributes": [],
      "taraAttributes": [
        {
          "key": {
            "id": "CONFIDENTIALITY"
          },
          "value": {
            "id": 1,
            "name": "The image from the camera and location of the vehicle can be accessed by externals"
          }
        },
        {
          "key": {
            "id": "INTEGRITY"
          },
          "value": {
            "id": 2,
            "name": "The location of the vehicle can be manipulated by externals"
          }
        }
      ],
      "tagsToString": "N/A"
    }
  ],
  "targetAsset": {
    "id": 501,
    "name": "Global Navigation Satellite System (GNSS)"
  },
  "targetProperty": {
    "id": "CONFIDENTIALITY",
```

```
        "description": "Confidentiality"
      },
      "elapsedTime": {
        "id": "LESS_OR_EQUAL_THAN_3_YEARS",
        "description": "LESS_OR_EQUAL_THAN_3_YEARS"
      },
      "expertise": {
        "id": "LAYMAN",
        "description": "LAYMAN"
      },
      "knowledge": {
        "id": "PUBLIC",
        "description": "PUBLIC"
      },
      "windowsOfOpportunity": {
        "id": "MODERATE",
        "description": "MODERATE"
      },
      "equipment": {
        "id": "SPECIALIZED",
        "description": "SPECIALIZED"
      },
      "attackFeasibilityRating": {
        "id": "M",
        "description": "M"
      }
    },
    "damageScenarioProfile": {
      "id": 1,
      "name": "The image from the camera and location of the vehicle can be accessed by externals",
      "businessPartner": {
        "id": 1,
        "name": "Ubitech"
      },
      "safetyImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "financialImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "operationalImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "privacyImpact": {
        "id": "MODERATE",
        "description": "Moderate"
      },
      "overallImpact": {
        "id": "SEVERE",
        "description": "Severe"
      }
    },
    "riskLevel": {
      "id": "H"
    }
  }
 ],
 "total": 1,
 "index": 0,
```

```
  "size": 10,
  "contentSize": 1
}
```