An open-source platform for Resilient Secure Digital Identities: The RECITALS project

George Stamoulis[1,*,†], Dimitris Pavlou[1,†], Konstantinos Chousos[1,†], Manolis Koubarakis[1,*,†], George Papadakis[2,†], Christina Papapostolou[2,†], Georgios Smaragdakis[3,†], Zekeriya Erkin[3,†], Roland Kromes[3,†], Themis Palpanas[4,†], Salima Benbernou[4,†], Mourad Ouziri[4,†], Adel Boubetra[4,†], Paulo Correia[5,†], João Pedro[6,†], Ioan Constantin[7,†], Laurentiu Coica[7,†], Marco Fisichella[8,†], Apoorva Upadhyaya[8,†], Harshvardhan Pandit[9,†] and Kyriakos Dimitriou[10,†]

[1]*National and Kapodistrian University of Athens, Greece*
[2]*Dimosia Epicheirisi Ilektrismou Anonymi Etaireia, Greece*
[3]*Delft University of Technology, The Netherlands*
[4]*Université Paris Cité, France*
[5]*Projecto Desenvolvimento Manutencao Formacao e Consultadorialda, Portugal*
[6]*Hospital Do Espirito Santo de Evora EPE, Portugal*
[7]*Orange Romania SA, Romania*
[8]*Leibniz University of Hannover, Germany*
[9]*Trinity College Dublin, Ireland*
[10]*Center for Technology Research and Innovation, Cyprus*

Abstract

The RECITALS project is a three-year Horizon Europe Innovation Action that pioneers the development of a privacy-by-design, open-source platform for secure digital identity and data sharing. By integrating logic-based reasoning for compliance automation, explainable AI, and privacy-preserving technologies, RECITALS advances legal, secure, and transparent data exploitation across sectors. It leverages a Knowledge Graph approach combined with logic-based reasoning to ensure that decisions regarding identity lifecycle management and data governance are aligned with European regulations such as GDPR, eIDAS, NIS2, and the AI Act. Central to RECITALS is the implementation of self-sovereign identity solutions that place control in the hands of the user, enhanced by compliance and explainability mechanisms. The platform's deployment in real-world use cases -— from the energy, telecommunications, and healthcare sector – demonstrates its ability to reason over complex policies, explain automated decisions, and enforce compliance through interoperable infrastructures. RECITALS started on January, 2025 and will be concluded on December 2027, bringing together a consortium of 10 partners from 8 EU countries, with the aim of delivering a comprehensive open-source platform for resilient secure digital identities.

Keywords

Open-Source Platform, Privacy-Preserving Technologies, Self-Sovereign Identity, Identity Lifecycle Management, Data Protection Compliance, Secure Data Sharing

RuleML+RR'25: Companion Proceedings of the 9th International Joint Conference on Rules and Reasoning, September 22–24, 2025, Istanbul, Turkiye

[*] Corresponding author.

[†] These authors contributed equally.

gstam@di.uoa.gr (G. Stamoulis); sdi2100151@di.uoa.gr (D. Pavlou); sdi2000215@di.uoa.gr (K. Chousos); koubarak@di.uoa.gr (M. Koubarakis); g_a.papadakis@ppcgroup.com (G. Papadakis); c.papapostolou@ppcgroup.com (C. Papapostolou); g.smaragdakis@tudelft.nl (G. Smaragdakis); z.erkin@tudelft.nl (Z. Erkin); R.G.Kromes@tudelft.nl (R. Kromes); themis@mi.parisdescartes.fr (T. Palpanas); salima.benbernou@u-paris.fr (S. Benbernou); mourad.ouzri@u-paris.fr (M. Ouziri); adelarte.ab@gmail.com (A. Boubetra); paulo.correia@pdmfc.com (P. Correia); jassuncao@hevora.min-saude.pt (J. Pedro); ioan.constantin@orange.com (I. Constantin); laurentiu.coica@orange.com (L. Coica); mfisichella@l3s.de (M. Fisichella); upadhyaya@l3s.de (A. Upadhyaya); harshvardhan.pandit@adaptcentre.ie (H. Pandit); k.dimitriou@cetri.net (K. Dimitriou)

1. Introduction

The RECITALS project[1] stands at the forefront of advancing and implementing cutting-edge privacy-preserving technologies aimed at enabling secure and legally compliant data exploitation. It incorporates a suite of state-of-the-art tools, including cryptographic anonymous credentials, homomorphic encryption, secure multiparty computation, and differential privacy [1, 2]. Its innovation lies not only in its technical components, but also in its deliberate integration of logic-based reasoning techniques for compliance checks and decision explainability into every layer of the platform. One of the project's key differentiators is its *Compliance Manager*, which utilizes a Knowledge Graph approach combined with logic-based reasoning techniques [3, 4]. This component interprets and enforces compliance requirements—such as GDPR, NIS2, and the AI Act—using a framework that enables dynamic, transparent, and verifiable decision-making.

RECITALS underscores the paramount significance of trusted digital identities, aligned with the European eID and forthcoming eIDAS regulations [5]. The project promotes the development of self-sovereign identity solutions that grant users full autonomy over their personal data and its usage. These user-centric digital identities are key to fostering trust across digital ecosystems [6]. Moreover, the project prioritizes usability, scalability, and the seamless integration of privacy-preserving technologies into existing infrastructures, particularly within supply chain environments.

RECITALS addresses the complexity of privacy-preserving data management across three use cases with diverse organizational models in the energy, telecommunications, and healthcare sectors. The proposed solutions are not only innovative but also practical, undergoing rigorous validation and pilot testing in real-world, federated data infrastructures to ensure their effectiveness and adaptability. To enhance user trust in AI-driven processes, RECITALS provides the *explAIner* component, a library of state-of-the-art xAI methods [7] that enhance the transparency of its automatic processes, while ensuring that they are not only compliant with EU regulations but also understandable, and user-centric.

In this context, RECITALS aims to deliver a highly efficient and scalable open-source platform tailored for both industrial and scientific applications that require privacy-preserving data sharing and identity management. With a robust privacy-by-design architecture, the platform is engineered to resist advanced and AI-driven cyber threats, comply with EU regulations, and provide value-added services for a broad spectrum of use cases.

In the following, we delve into the main objectives of RECITALS, the architecture of its open-source platform and the partners comprising its consortium.

2. Main Objectives and Goals

The RECITALS project is driven by a vision to enable secure, privacy-preserving, and regulation-compliant data sharing and identity management across critical sectors. The primary goal of RECITALS is to design an open-source, privacy-by-design platform that supports next-generation data sharing and identity management. This platform is tailored for resilience and full compliance with evolving European regulations, such as GDPR, NIS2, and the AI Act. RECITALS integrates fundamental but common operations for privacy-preserving data sharing and identity management into modules that facilitate the development of diverse applications through a library of state-of-the-art techniques. On top of these modules, the project aims to develop advanced services that support industrial-strength applications with automated compliance capabilities. These services are specifically designed to meet the requirements of domain-specific EU regulations and ensure seamless deployment across varied operational contexts. Special care is taken to identify and mitigate common and AI-powered threats against these services through the cybersecurity component. This component protects the platform's core infrastructure as well as its value-added services, guaranteeing security across the data lifecycle.

To validate these developments, RECITALS demonstrates the platform's potential in real-world business scenarios within the energy, telecommunications, and healthcare sectors. These use cases

[1] https://recitals-project.eu/

are not only critical from a cybersecurity perspective but also serve to test the platform's usability, interoperability, and effectiveness in diverse data-sharing environments.

RECITALS aligns with the European Commission's Increased Cybersecurity Destination of the Strategic Plan 2021–2024, contributing to enhanced data protection, network security, and technological sovereignty across the EU. The project aims to achieve tangible impacts, including improved software, hardware, and supply chain security, and the creation of a cybersecurity culture through interdisciplinary collaboration and stakeholder engagement. RECITALS envisions a strengthened EU cybersecurity posture and sovereignty in digital technologies. It promotes privacy-by-design architectures and interoperable standards to ensure that infrastructures and processes remain resilient against emerging threats. By aligning its goals with the European cybersecurity initiatives and engaging with a wide range of stakeholders, RECITALS lays the foundation for smart and qualified security assurance and certification that is shared across the EU.

3. The RECITALS Platform

Figure 1 illustrates the high-level architecture of the RECITALS platform, which is composed of three main components: the Core, a suite of Value-Added Services, and a dedicated Security Manager. Each component plays a pivotal role in enabling a secure, compliant, and efficient identity management and data sharing ecosystem across organizational boundaries. Together, they ensure that RECITALS delivers a secure, privacy-aware, and regulation-compliant platform for identity and data management across multiple domains and stakeholders. Below, we elaborate on their role and modules.

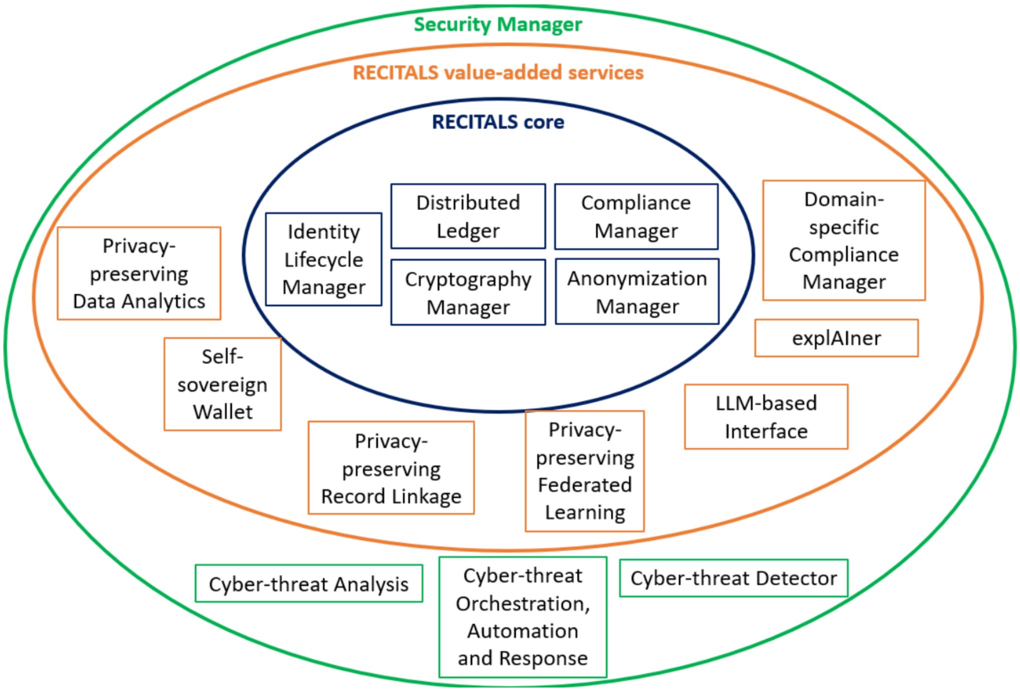
Security Manager
RECITALS value-added services
RECITALS core
Identity Lifecycle Manager
Distributed Ledger
Compliance Manager
Cryptography Manager
Anonymization Manager
Privacy-preserving Data Analytics
Domain-specific Compliance Manager
explAIner
Self-sovereign Wallet
LLM-based Interface
Privacy-preserving Record Linkage
Privacy-preserving Federated Learning
Cyber-threat Analysis
Cyber-threat Orchestration, Automation and Response
Cyber-threat Detector

Figure 1: Overview of the RECITALS platform. The main components correspond to ovals, and their modules to rectangles.

3.1. RECITALS Core

This is the backbone of the RECITALS platform, with its modules providing the fundamental operations that lay the ground for the more complex services. At its heart lies the *Distributed Ledger*, which establishes a decentralized and immutable foundation for all transactions and records [8]. This component fosters trust among participants by providing transparency, ensuring data integrity through tamper-proof immutability, and guaranteeing system resilience by avoiding a single point of failure. Special attention is paid to supporting private ledgers for intra-organizational identity management and hybrid ledgers for secure inter-organizational data exchanges.

Complementing the ledger is the *Identity Lifecycle Manager*, which governs digital identities throughout their entire lifecycle, from creation to deactivation. It ensures regulatory and organizational compliance via audit and reporting functionalities and enforces policy alignment through automated checks. Self-service capabilities empower users to manage their identities autonomously, while administrative tools centralize control. Notifications and approval workflows enhance user engagement and oversight, whereas built-in password and compliance management mechanisms mitigate security risks and promote adherence to EU standards.

Privacy and security are further strengthened through the *Cryptography Manager*, which incorporates state-of-the-art techniques such as differential privacy, homomorphic encryption, secure multi-party computation, zero-knowledge proofs, and verifiable credentials. These tools enable secure computations on encrypted data, privacy-preserving authentication, and trusted identity assertions without compromising sensitive information.

In parallel, the *Anonymization Manager* applies rigorous data anonymization algorithms (e.g., k-anonymity, l-diversity, t-closeness) to manage controlled data publishing, limiting disclosure through selective access and attribute generalization.

Ensuring that all operations conform to legal and regulatory expectations, the *Compliance Manager* leverages knowledge graphs and logic-based reasoning, drawing on prior EU projects. It extends beyond GDPR to encompass emerging frameworks like the DGA and NIS2, supported by extensible vocabularies such as the Data Privacy Vocabulary [9]. It also automates key compliance checks, enforces security principles such as confidentiality, integrity, and availability, and promotes trustworthy services aligned with evolving regulatory landscapes across the EU Data Spaces.

3.2. RECITALS Value-Added Services

To operationalize and extend RECITALS into real-world applications, a range of Value-Added Services will be developed on top of RECITALS Core. The main service is the intuitive user interface, which supports natural language interactions through LLMs, i.e., ChatBots. The *LLM-based Interface* is powered by a Retrieval-Augmented Generation framework, which ensures that outputs are grounded in factual and updated documentation. Thus, it bypasses the limitations of static LLM training by dynamically incorporating evolving RECITALS platform knowledge.

The *Self-Sovereign Wallet* gives users complete control over their identities and data through decentralized identifiers and verifiable credentials. This wallet enhances privacy and user autonomy while ensuring interoperability and ease of use with the LLM-based Interface.

The *Privacy-Preserving Record Linkage* enables secure identification of duplicate records across datasets without revealing underlying sensitive data [10]. This is achieved by transforming identifiers into encrypted or obfuscated formats and by leveraging cryptographic methods such as secure multi-party computation to perform record linkage in a privacy-preserving manner.

The *Privacy-Preserving Federated Learning* facilitates collaborative model training without centralizing data [11]. It supports techniques like federated averaging, knowledge distillation, and federated reinforcement learning, allowing users to jointly improve models while safeguarding data privacy.

Transparency in automated decisions is achieved through the *explAIner component*, which integrates explainable AI techniques such as Local Interpretable Model-agnostic Explanations, SHapley Additive exPlanations, Partial Dependence Plots, and Individual Conditional Expectation. These tools provide post-hoc interpretations of model behavior, supporting compliance with transparency mandates and enhancing user trust.

Domain-specific compliance is addressed through an extension of the *Compliance Manager*, targeting the energy, telecommunications, and healthcare sectors. This extension adapts the core compliance infrastructure with tailored vocabularies and validation rules for domain-specific use cases, facilitating operational deployment and reducing compliance overhead.

Finally, *Privacy-Preserving Data Analytics* enables the extraction of insights from sensitive data using methods that combine encryption, aggregation, and quasi-identifier handling [12]. These analytics techniques maintain data confidentiality throughout processing, thereby empowering users to derive value from sensitive datasets without compromising privacy.

3.3. RECITALS Security Manager

The responsibility for safeguarding the RECITALS platform against cyber threats lies with the Security Manager. This component performs two interlinked functions that are based on *Cyber-Threat Analysis*, i.e., an systematic overview of the most likely attacks for the RECITALS platform (e.g., AI-powered attacks, threats to identity). First, the *Cyber-Threat Detector* identifies threats through anomaly detection, signature-based recognition, behavioral analysis, endpoint monitoring, and centralized event correlation. These mechanisms can operate independently or in combination, often enhanced by machine learning for real-time threat identification.

Second, once threats are detected, the *Cyber-Threat Orchestration, Automation, and Response* module coordinates the necessary countermeasures. It leverages threat intelligence from standards and open-source sources, such as SIGMA rules and the MITRE framework, to automate response workflows and dynamically adapt defenses. This enables RECITALS to maintain operational integrity and proactively manage evolving security risks.

4. The RECITALS Consortium

The RECITALS project brings together a carefully selected consortium of 10 partners from 8 EU countries: Greece, The Netherlands, France, Portugal, Romania, Germany, Ireland, and Cyprus. The composition of the consortium was strategically designed to address the multifaceted challenges associated with the development of the RECITALS platform, combining expertise, diversity, and cohesion.

A critical design consideration was the balance between academic and industrial perspectives. The consortium comprises five well-established academic institutions and five industrial organizations, creating an ideal synergy between theoretical innovation and real-world applicability. The academic partners include the National and Kapodistrian University of Athens (NKUA), the Delft University of Technology (TUD), the Université Paris Cité (UPC), the Leibniz University of Hannover (LUH), and the Dublin City University (DCU). They are renowned universities with significant contributions to research in fields such as AI, cybersecurity, data management, and privacy preservation On the industrial side, the consortium includes two large companies, Public Power Corporation SA (PPC) and Orange Romania SA (ORO), the SMEs Projecto Desenvolvimento Manutencao Formacao e Consultadorialda (PDM) and the Center for Technology Research and Innovation (CETRI), as well as a major public-sector organization, the Hospital Do Espirito Santo de Evora EPE (HES).

This diversity ensures that the RECITALS platform is validated across a range of sectors, making the project highly adaptable and future-proof. In fact, each partner brings a distinct area of specialization, as shown in Table1, contributing to a comprehensive skill set that spans both research excellence and market-oriented implementation. The complementary backgrounds of the consortium members are a key strength of the project, ensuring that collectively, the consortium covers the full spectrum of competencies required for building a secure, privacy-respecting, and AI-enhanced identity management and data-sharing platform.

Acknowledgments

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No.101168490. The European Commision authority managing RECITALS project is the European Cybersecurity Compentence Center[2].

[2] https://cybersecurity-centre.europa.eu/index_en

Table 1
Expertise areas covered by the RECITALS consortium partners.

Expertise	NKUA	PPC	TUD	UPC	PDM	HES	ORO	LUH	DCU
Distributed Ledger			X		X				
Identity Management	X	X			X				
Cryptography	X				X				
Anonymization	X				X				
Compliance								X	X
Record Linkage	X	X		X					
Federated Learning				X				X	
LLMs	X	X							
Explainable AI				X				X	
Data Analytics				X				X	
Cyber-threat Detection		X	X		X	X	X		
Cyber-threat Mitigation		X	X		X	X	X		

Declaration on Generative AI

During the preparation of this work, the authors used OpenAI GPT-4o for *grammar and spelling check*. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] J. Katz, Y. Lindell, Introduction to Modern Cryptography, Second Edition, CRC Press, 2014.

[2] C. Paar, J. Pelzl, T. Güneysu, Understanding Cryptography - From Established Symmetric and Asymmetric Ciphers to Post-Quantum Algorithms, Second Edition, Springer, 2024.

[3] D. Golpayegani, I. Hupont, C. Panigutti, H. J. Pandit, S. Schade, D. O'Sullivan, D. Lewis, AI cards: Towards an applied framework for machine-readable AI and risk documentation inspired by the EU AI act, CoRR abs/2406.18211 (2024).

[4] D. Golpayegani, H. J. Pandit, D. Lewis, Aicat: An AI cataloguing approach to support the EU AI act, CoRR abs/2501.04014 (2025).

[5] A. Sharif, M. Ranzi, R. Carbone, G. Sciarretta, F. A. Marino, S. Ranise, The eidas regulation: a survey of technological trends for european electronic identity schemes, Applied Sciences 12 (2022) 12679.

[6] A. I. Blasco, Digital identity in a european user-centric ecosystem and its similarities with the digital euro proposal, in: Governance and Control of Data and Digital Economy in the European Single Market, 2025, pp. 453–471.

[7] G. Schwalbe, B. Finzel, A comprehensive taxonomy for explainable artificial intelligence: a systematic survey of surveys on methods and concepts, Data Mining and Knowledge Discovery 38 (2024) 3043–3101.

[8] S. G. Savadatti, S. Krishnamoorthy, R. Delhibabu, Survey of distributed ledger technology (DLT) for secure and scalable computing, IEEE Access 13 (2025) 8393–8415.

[9] H. J. Pandit, et al., Creating a vocabulary for data privacy - the first-year report of data privacy vocabularies and controls community group (DPVCG), in: OTM, volume 11877, 2019, pp. 714–730.

[10] A. Gkoulalas-Divanis, D. Vatsalan, D. Karapiperis, M. Kantarcioglu, Modern privacy-preserving record linkage techniques: An overview, IEEE Transactions on Information Forensics and Security 16 (2021) 4966–4987.

[11] T. H. Rafi, F. A. Noor, T. Hussain, D.-K. Chae, Fairness and privacy-preserving in federated learning: A survey, 2023. URL: https://arxiv.org/abs/2306.08402. arXiv:2306.08402.

[12] H.-Y. Tran, J. Hu, Privacy-preserving big data analytics a comprehensive survey, Journal of Parallel and Distributed Computing 134 (2019) 207–218.