Proposal No. 830929
Call H2020-SU-ICT-03-2018

Project start: 1 February 2019
Project duration: 42 months

Cyber Security for Europe

D3.20

Final cybersecurity enablers and underlying technologies components

Document Identification	
Due date	30 June 2022
Submission date	30 June 2022
Revision	1.0

Related WP	WP3	Dissemination Level	Public
Lead Participant	UMU	Lead Author	Jorge Bernal
Contributing Beneficiaries	C3P, UMU, UMA, CNR, AIT, UNILU, DTU, UM, VTT, ATOS, CYBER, NEC, UPRC	Related Deliverables	D3.2, D3.11, D3.13

Abstract: This deliverable D3.20 - “Final cybersecurity enablers and underlying technologies components” reports about the assets devised, implemented, and evaluated in the scope of task T3.2. The document evolves former task’s T3.2 deliverable D3.13, which in turn, was an evolution of T3.2 deliverables D3.2 – “Cross Sectoral Cybersecurity Building Blocks” and D3.23.

This document is issued within the CyberSec4Europe project. This project has received funding from the European Union's Horizon 2020 Programme under grant agreement no. 830929.

Executive Summary

This document presents D3.20 – "Final cybersecurity enablers and underlying technologies components of task T3.2 – "Research and Integration on Cybersecurity Enablers and underlying Technologies". This deliverable includes the final and consolidated outcomes from T3.2 focused, mainly on privacy enablers and tools and as well as the impact and applicability that they can have in a general smart-campus demonstrator.

This deliverable describes the evolved version of some assets already described in D3.13, which have been improved with additional functionalities and have been analyzed and evaluated deeply. D3.2 gave a first overview of all CyberSec4Europe's enablers (also known as "assets" within the project), D3.13 described the first version of the demonstrator to show these assets working in a joint scenario and D3.20 is an updated version of the assets.

Moreover, this deliverable categorizes most of the assets in the different building blocks of the ATLAS in order to show the wide range of priority areas of research in which the assets belonging to T3.2 is being undertaken.

Finally, this deliverable reports research outcomes from other WP3 tasks that have been extended in the project and whose final associated deliverables were already delivered. For the sake of space, these additional assets and research outcomes beyond T3.2 are just outlined and referenced in the final sections of this document. The reader can find the completed documentation of these assets in their corresponding references to the official CS4E-WP3 Github page [cs4e-github].

Document information

Contributors

Name	Partner
Eda Marchetti	CNR
Said Daoudagh	CNR
Miryam Villegas Jimenez	Atos
Raquel Cortés Carreras	Atos
Juan Carlos Pérez Baún	Atos
Liina Kamm	CYBER
Alisa Pankova	CYBER
Baldur Kubo	CYBER
Hiroki Kaminaga	CYBER
Christos Xenakis	UPRC
Eleni Veroni	UPRC
Jorge Bernal	UMU
Antonio Skarmeta	UMU
Agustin Marín	UMU
Pablo Fernandez	UMU
Jesus Martinez	UMU
Juan F. Martinez	UMU

Ruben Rios	UMA
Rodrigo Roman	UMA
Javier Lopez	UMA
Alireza Esfahani	UNILU
Stephan Krenn	AIT
Thomas Lorünser	AIT
Florian Wohner	AIT
João Resende	C3P
Claudio Soriente	NEC
Alessandro Sforzin	NEC
Marko Kompara	UM

Reviewers

Name	**Partner**
Davy Preuveneers	KUL
Romain Laborde	IRIT

History

Version	Date	Authors	Comment
0.01	2022-03-28	Jorge Bernal Bernabe	1st Draft and TOC
0.02	2022-04-14	Eda Marchetti and Said Daoudagh	Start Drafting GENERAL_D Asset
0.03	2022-05-11	Juan Carlos Pérez Baún, Miryam Villegas Jimenez and Raquel Cortés Carreras	Section 5.7
0.04	2022-05-11	Liina Kamm, Alisa Pankova	PLEAK DP analysers
0.05	2022-05-12	Christos Xenakis, Eleni Veroni,	Section 5.4
0.06	2022-05-15	Eda Marchetti and Said Daoudagh	Completed GROOT description and application (GENERAL_D Asset, Section 4.12)
0.07	2022-19-05	Ruben Rios, Rodrigo Roman, Javier Lopez	Section 5.5
0.1	05/05/2022	UMU	Atlas priorities
0.2	20/04/2022	UMU	Updated SS-PP-IdM asset, pp-FL, pp-CTI assets
0.3	25/05/2022	Jorge Bernal	Editing process, Conclusions
0.4	31/05/2022	Jorge Bernal	Minor updates Document sent to internal review
0.5	17/06/2022	Davy Preuveneers (KUL) Romain Laborde (IRIT)	Internal review
0.6	27/06/2022	Jorge Bernal, Pablo Martinez, Said Daoudagh, Eda Marchetti, Stephan Krenn, Ruben Rios, Joao Resende, Marko	Revise the document and address the internal review comments

		Kompara, Alessandro Sforzin, Claudio Sonriente	
1.0	29.06.2022	Ahad Niknia	Final check, preparation and submission

Table of Contents

1 Introduction 1

1.1 Document Structure 1

2 Final T3.2 Architecture 2

3 Description of the Smart Campus Scenario 8

3.1 CCTV Surveillance in the Smart-Campus 10

3.1.1 Architecture 11
3.1.2 Administration Plane 13
3.1.3 Control and Management Plane 14
3.1.4 Control and Management Plane 14
3.1.5 User Domain 14
3.1.6 IoT domain 14

3.2 Identity management and service usage in Smart Campus 15

3.2.1 Architecture 16
3.2.2 Intelligence Plane 18
3.2.3 Control and Management Plane 18
3.2.4 User Domain 18
3.2.5 Blockchain Interfaces 19

3.3 Geolocation Service in the Smart Campus 19

3.3.1 Architecture 19
3.3.2 Control and Management Plane 21
3.3.3 User Domain 22
3.3.4 IoT Domain 22

4 Asset Demonstration 22

4.1 PTASC 22

4.1.1 Overview 22
4.1.2 Main asset improvements since D3.13 26
4.1.3 Research challenges addressed 26
4.1.4 Demonstrations Example 27
4.1.5 Future Work 27

4.2 ARGUS 27

4.2.1 Overview 27
4.2.2 Main asset improvements since D3.13 28
4.2.3 Research challenges addressed 28
4.2.4 Demonstrations Example 29
4.2.5 Future Work 33

4.3 Self-Sovereign Privacy-Preserving-IdM (SS-PP-IdM) 33

4.3.1 Overview 33
4.3.2 Main asset improvements since D3.13 36
4.3.3 Research challenges addressed 36
4.3.4 Demonstrations Example 38
4.3.5 Future Work 41

4.4 Password-less authentication 41

4.4.1 Overview 41
4.4.2 Main asset improvements since D3.13 44
4.4.3 Research challenges addressed 44
4.4.4 Demonstrations Example 45
4.4.5 Future Work 45

4.5 Edge-Privacy 45

4.5.1 Overview 45
4.5.2 Main asset improvements since D3.13 46

4.5.3 Research Challenges addressed 46
4.5.4 Demonstration Example 46
4.5.5 Future work 48
4.6 Privacy-Aware Aggregate Programming 48
4.6.1 Overview 48
4.6.2 Main asset improvements since D3.13 49
4.6.3 Future Work 49
4.7 DANS 49
4.7.1 Overview 49
4.7.2 Main asset improvements since D3.13 50
4.7.3 Research Challenges addressed 51
4.7.4 Demonstrations Example 51
4.7.5 Future Work 52
4.8 Cryptovault 52
4.8.1 Overview 52
4.8.2 Main asset improvements since D3.13 53
4.8.3 Research Challenges Addressed 53
4.8.4 Future Work 53
4.9 Elastic Deployment of TEE-based applications in the cloud 53
4.9.1 Overview 53
4.9.2 Main asset improvements since D3.13 54
4.9.3 Future Work 54
4.10 Backdoor-resistant TEEs 54
4.10.1 Overview 54
4.10.2 Main asset improvements since D3.13 54
4.10.3 Future Work 54
4.11 Privacy-Preserving for Genomic Data (PP4Genomic) 55
4.11.1 Overview 55
4.11.2 Main asset improvements since D3.13 55
4.11.3 Research Challenges Addressed 55
4.11.4 Future Work 55
4.12 GENERAL_D 56
4.12.1 Overview 56
4.12.2 Main asset improvements since D3.13 57
4.12.3 Research Challenges Addressed 57
4.12.4 Demonstration Example of GROOT in the context of CCTV Surveillance 57
4.12.4.1 Using GROOT 58
4.12.5 Future Work 59
4.13 Blockchain Platform 59
4.13.1 Overview 59
4.13.2 Main asset improvements since D3.13 61
4.13.3 Future Work 61
4.14 Sharemind 61
4.14.1 Overview 61
4.14.2 Main asset improvements since D3.13 63
4.14.3 Demonstration example 63
4.14.4 Future Work 65
4.15 Cloud-Based Credentials 65
4.15.1 Overview 66
4.15.2 Main asset improvements since D3.13 66
4.15.3 Future Work 66
4.16 Issuer-Hiding Anonymous Credentials 66
4.16.1 Overview 66

4.16.2 Main asset improvements since D3.13 67
4.16.3 Future Work 67
4.17 FlexProd and ArchiStar 67
4.17.1 Overview 67
4.17.2 Main asset improvements since D3.13 68
4.17.3 Research challenges addressed 68
4.17.4 Demonstrations Example 68
4.17.5 Future work 70
4.18 GDPR compliant user experience 70
4.18.1 Overview 70
4.18.2 Main asset improvements since D3.13 71
4.18.3 Future Work 71
4.19 Interoperability and cross-border compliance 72
4.19.1 Overview 72
4.19.2 Main asset improvements since D3.13 72
4.19.3 Research challenges addressed 73
4.19.4 Demonstration Example 73
4.19.5 Future work 75
4.20 Privacy-preserving Federated Learning 75
4.20.1 Overview 76
4.20.2 Main asset improvements since D3.13 76
4.20.3 Research challenges addressed 76
4.20.4 Demonstration Example 76
4.20.5 Future work 79
4.21 PLEAK Differential Privacy Analysers 79
4.21.1 Overview 79
4.21.2 Main asset improvements since D3.13 80
4.21.3 Research challenges addressed 80
4.21.4 Demonstration example 81
4.21.5 Future Work 85
5 Task 3.2 Outcomes 85
5.1.1 Overview 85
5.1.2 Cybersecurity Research and Areas Priority 85
6 Task 3.3 Outcomes 87
6.1.1 Overview 87
6.1.2 Cybersecurity Research and Areas Priority 87
7 Task 3.4 Outcomes 88
7.1.1 Overview 88
7.1.2 Cybersecurity Research and Areas Priority 88
7.1.3 Privacy Preserving Cyber Threat Information Sharing leveraging FL-based intrusion detection . 89
8 Task 3.5 Outcomes 92
8.1.1 Overview 92
8.1.2 Cybersecurity Research and Areas Priority 92
9 Task 3.6 Outcomes 93
9.1.1 Overview 93
9.1.2 Cybersecurity Research and Areas Priority 93
10 Task 3.7 Outcomes 94
10.1.1 Overview 94
10.1.2 Cybersecurity Research and Areas Priority 95
11 Task 3.8 Outcomes 96
11.1.1 Overview 96
11.1.2 Cybersecurity Research and Areas Priority 96
12 Task 3.10 Outcomes 96

12.1.1 Overview 96
12.1.2 Cybersecurity Research and Areas Priority 97

13 Conclusions 97

14 References 97

List of Figures

Figure 1: CyberSec4Europe Privacy-Preserving Functional high-level Architecture 3
Figure 2: Smart Campus 9
Figure 3: CCTV Scenario Overview 12
Figure 4: IdM scenario architecture overview and instantiation 17
Figure 5: Mapping of the Geolocation Service in the General Architecture. 20
Figure 6: Two deployment options for secure computing 21
Figure 7: Manager Setup Phase 23
Figure 8: Device authentication 24
Figure 9: Decentralized Secure End-to-End Communications 24
Figure 10: Merge Two Trusted Devices Pools 25
Figure 11: Privacy data controller 26
Figure 12 - Argus Architecture 28
Figure 13: Download scalability tests 32
Figure 14: Upload scalability tests 32
Figure 15 Generic overview of SS-PP-IdM asset 33
Figure 16. Smart Campus framework (Authentication and Authorization) 35
Figure 17: Instantiation of SS-PP-IdM in Smart Campus platform 38
Figure 18: Application Homepage of the Smart Campus 39
Figure 19: Using some attribute-based policy 39
Figure 20 – SS-PP-IdM asset Performance evaluation 40
Figure 21: Authentication and authorization processes time measurements 41
Figure 22: FIDO Authentication Concept 42
Figure 23: Password-less authentication. Registration process 42
Figure 24: Password-less Authentication. Authentication process 43
Figure 25: De-registration Process 43
Figure 26: Architecture of the Privacy Manager for IoT data. 45
Figure 27 : PMEC execution times for different database sizes 48
Figure 28: Incremental construction of a proximity field 48
Figure 29 - Components of the DANS as a service asset 50
Figure 30 - Flavours of DANS tool (a) Anonymisation as a Service, (b) Embedded library .. 52
Figure 31: The Authorization Policy Life Cycle. 56
Figure 32 - GROOT Methodology 57
Figure 33: Blockchain platform architecture showing three independent satellite chains in action 60
Figure 34 - Sharemind secure computing platform 61
Figure 35 - Sharemind HI security model 62
Figure 36 - Secret sharing two values (orange and blue) among three computation nodes and an MPC protocol that results in a secret-shared result (green). 63
Figure 37 - Smart Manufacturing scenario 69
Figure 38: The main steps in the DPIA Template. 71
Figure 39 - Proposed Federated Learning with Differential Privacy demonstration architecture 77
Figure 40 - FedAvg accuracy evolution for every perturbation mechanism. 78
Figure 41 - FedPlus Accuracy evolution for every perturbation mechanism. 78
Figure 42 - DP execution times per mechanism 78
Figure 43 - Histogram of tourists visits (data without differential privacy) 81
Figure 44 - Adding noise to a dataset 81
Figure 45 - Simple PLEAK model for counting seasonal travels 82
Figure 46 - Table schema for the trips data object 82

Figure 47 - Task “count seasonal travels” 82
Figure 48 - Attacker goal specified as a database query. 83
Figure 49 - Setting input parameters 83
Figure 50 - Changes in the recommended noise for the specified risk tolerance 84
Figure 51 - Example histogrm after sampling and adding the noise 84
Figure 52 - Task 3.3 tools 87
Figure 53 - CYTILIS architecture proposal 90
Figure 54 - Secure exchange of CTI events between organizations 91
Figure 55 - Producer workflow within the FL subsystem 91
Figure 56 - Consumer workflow within the FL subsystem 92

List of Acronyms

2	**2FA**	Two-Factor Authentication
A	**ABC**	Attribute-Based Credential
	AES	Advanced Encryption Standard
C	**CA**	Certification Authority
	CCTV	Closed-Circuit Television
	CS4E	CyberSec4Europe
D	**DLT**	Distributed Ledger Technology
	DPIA	Data Protection Impact Assessment
	DPO	Data Protection Officer
	DHE	Diffie-Hellman Ephemeral
	DIF	Decentralized Identity Foundation
	DID	Decentralized Identifier
E	**eID**	Electronic Identity
	ECDSA	Elliptic Curve Digital Signature Algorithm
	ECIES	Elliptic Curve Integrated Encryption Scheme
	ECDHE	Elliptic-Curve Diffie–Hellman
G	**GDPR**	General Data Protection Regulation
	GPS	Global Positioning System
I	**IdM**	Identity Manager
	IdP	Identity Provider

	IoT	Internet of Things
M	**MPC**	Multi-Party Computation
N	**NGS**	Next Generation Sequencing
P	**p-ABC**	Privacy-preserving Attribute-Based Credential
	PET	Privacy Enhancing Technologies
	PTASC	
S	**SAML**	Security Assertion Markup Language
	SSI	Self-Sovereign Identity
T	**TEE**	Trusted Execution Environment
U	**UAF**	Universal Authentication Framework
V	**vIdP**	Virtual Identity Provider
X	**XACML**	eXtensible Access Control Markup Language

1 Introduction

The main purpose of this document is to present the final version of the assets that have been designed, implemented and evaluated as part of task T3.2. Most of the assets define herein are an evolved version of the ones earlier described in CybserSec4Europe deliverable D3.13.

To validate the feasibility, effectiveness, novelty, soundness, accuracy and performance of the assets, some of them have been evaluated as part of a common scenario based on "Smart University Campus", where different CS4E partners (mainly in T3.2) have contextualized and in some cases integrated their assets for evaluation and demonstration. Thus, the smart-campus scenario describes the storyline, processes, and test-cases employed to validate the assets devised and implemented in task T3.2.

Note that the T3.2 demonstrator focuses on the evaluation of the research aspects of the T3.2 investigations and associated assets. The evaluation of the assets in a fully integrated software-implementation of a real pilot is conducted in WP5. Therefore, the integration of the implementations of all different T3.2 assets within the Smart Campus scenario is not mandatory; this is a demonstration on how the tools can work in this context. Thus, some assets have been evaluated and demonstrated in specific scenarios and testbeds outside the smart-campus scenario.

The smart-campus scenario was first defined in D3.13, and it is introduced here again for the sake of completeness. The scenario addresses challenges that are research topics in T3.2, including, above all, privacy-concerns in transactions and identity management, and also privacy-preservation involving IoT scenarios, TEE (Trusted Execution Environments), and blockchain systems.

As it is the final accumulative report with the final assets/enablers in T3.2. This document also includes the T3.2 privacy-preserving architecture as an instantiation of the general CS4E architecture. Nonetheless, those assets that have not evolved since the previous deliverable D3.13 are summarized here, but their demonstration is not repeated here again for the sake of space.

Additionally, for this deliverable we have categorized most of the assets in the cybersecurity research category to show the wide range of priority areas of research and problem solving in which this work is being undertaken.

Different assets have been analyzed and described during this deliberation and fitted into a table that, in a generalized way, shows the following priority research areas: Trust-Building Blocks, Trustworthy Ecosystems of Systems, Governance & Capacity Building, Disruptive & Emerging Developments.

1.1 Document Structure

This document is a follow-up of D3.13, with an updated version of the assets. Thus, it includes a use case scenario to make the demonstration of each asset.

The document is structured as follows:

- Section 2 describes the T3.2 general architecture.
- Section 3 gives an overview of the smart-campus scenario description, including the three sub-scenarios for demonstrations:
    - CCTV,
    - Identity Management, and
    - Geolocation service.
- Section 4 has a description of all assets. For each asset, the discussion is structured as follows:
    - An introduction to the asset and its modus operandi,
    - The research challenges addressed,
    - The demonstrations example.

- Section 5-12 shows achievements of each task in WP3, including task beyond T3.2:
    - A brief overview
    - Advancements of assets reached by the task
    - A table showing the relation between the assets and the cybersecurity research areas priorities.

Section 13 concludes the document.

2 Final T3.2 Architecture

The CyberSec4Europe Privacy-Preserving functional Architecture consists of several high-level building blocks that expand over several intertwined domains, including the user domain, the web domain, and the IoT domain, as shown in Figure 1. A detailed description can be found in D3.2 “Cross-Sectoral Cybersecurity Building Blocks.” We also map all the assets used in this deliverable with the correspondent block of the framework. Some assets of the D3.11 deliverable were replaced by new privacy tools or rebranded under a new name and this information is also available. The total number of assets in this task are 22, divided by “Services Plane”, “User domain”, “Administration Plane”, “Intelligence Plane”, “Control and management plane”, “Blockchain Plane”, “IoT domain” and “Web domain”. Table 1 categorizes the assets accordingly with the plane, and it also gives information regarding the difference regarding the previous deliverable and the collaboration with WP5.

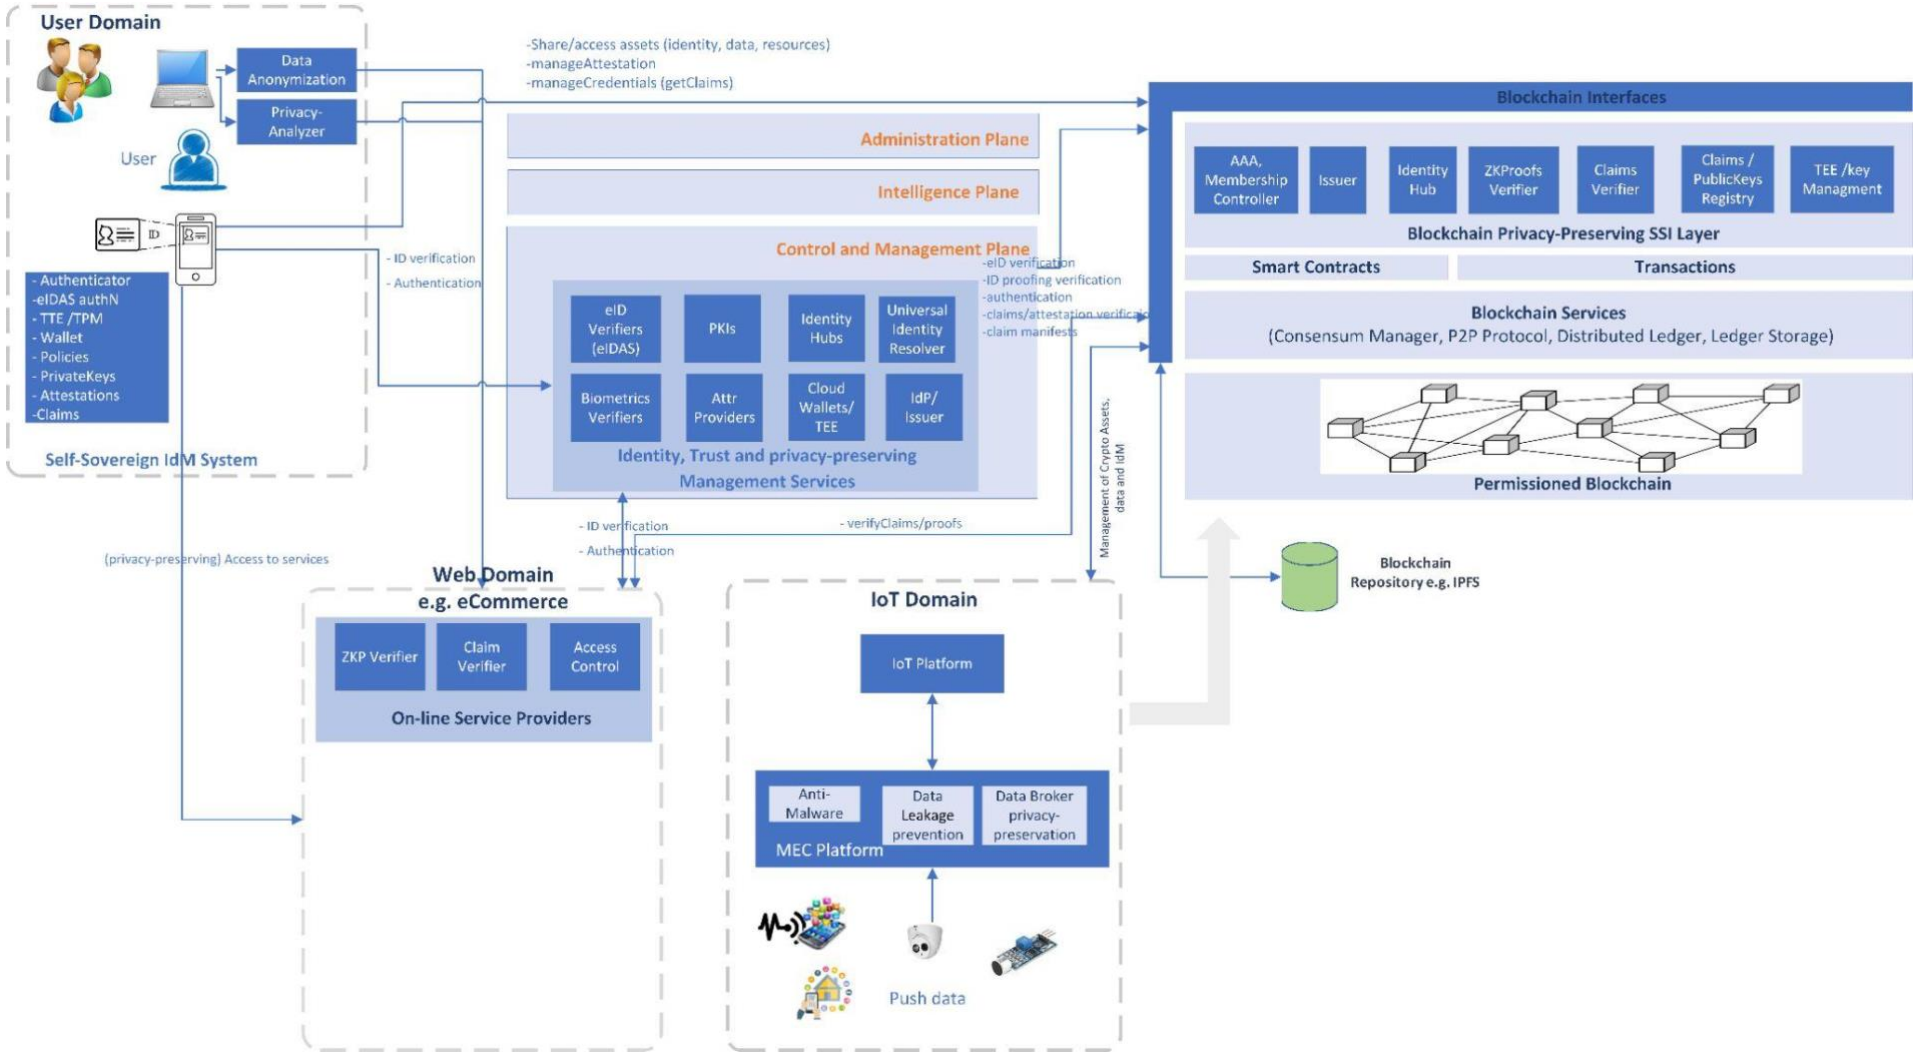

Figure 1: CyberSec4Europe Privacy-Preserving Functional high-level Architecture

The building blocks are defined for different purposes which range from compliance with current legal frameworks such as eIDAS and GDPR to mechanisms related to hardware-based solutions for managing keys and applications securely. Next, we give an overview of the different building blocks that are being proposed.

In the Control and Management plane of the CyberSec4Europe architecture, the Identity and Privacy-preservation Services plane includes the building blocks considered in the CyberSec4Europe Privacy-Preserving Architecture devoted to enabling privacy-respectful authentication based on the provision of anonymous credential systems and privacy-preserving identity management services, some of which rely on the use of secure distributed ledger technologies such as Blockchains to provide a self-sovereign identity (SSI) model. The Identity and privacy-preservation Services also include mechanisms for privacy-preserving computation technologies to reduce information leakage during the computations in the managed domain, thereby verifying that the systems comply with the users' privacy policies. Those privacy-preservation services can be run in the Cloud so that the architecture includes confidentiality-preserving and end-to-end secure sharing of sensitive data in the cloud among stakeholders using, for instance, secret sharing technologies. Besides, the architecture considers the privacy brokerage aiming at enhancing user trust in public cloud storage systems, guaranteeing data confidentiality and improving availability. The Privacy-preserving architecture includes functional building blocks for confidential and privacy-preserving storage that can employ techniques such as secret sharing to anonymize personal information during data analysis processes. Similarly, it also embraces privacy-preserving mechanisms for analyzing data from potentially different stakeholders in a way that gives high authenticity and integrity guarantees on the computation’s result, while protecting the confidentiality and privacy of the input data and ensuring data integrity.

On top of that, the Privacy-Preserving Architecture includes several mechanisms that use Trusted Execution Environments (TEE) for different purposes that range from securely storing and managing secret keys to remote anonymous attestation even in the presence of compromised hardware. The building blocks can be used on the virtualized applications in the Cloud or directly installed in the user domain.

In the User Domain, the privacy-preserving architecture encompasses the wallets and TEE needed to maintain securely protected credentials and manage key material obtained during the issuance and enrollment in diverse identity providers. The user domain is exemplified either with user mobiles, or software for desktop browsers. It contains the client-side software needed to perform authentication against service providers, eIDs-based authentication, and run protocols for proving privacy-Attribute Based credentials and claims (including zero-knowledge proofs).

Therefore, the user domain plays the role of Recipient and Prover in the privacy-ABC model. To this aim, the user domain interacts with diverse online identity services (including IdPs, Attribute providers, PKIs, biometric verifiers, eID verifiers) placed in the Control and Management Domain of the CyberSec4Europe architecture. In addition to credentials, the user domain needs to manage the attestations obtained from diverse attributes and identity providers, and short tokens obtained from IdPs (for single sign-on in Service Providers). The user-domain might also include ID-Proofing mechanisms, with client-side biometrics software needed to authenticate in biometric servers as a second authentication factor.

Furthermore, the user-domain considers the data anonymization building blocks to share in a privacy-preserving way data in transactions online and between organizations using diverse different privacy models (e.g., the k-anonymity, k-Map, Average risk model, among others). In addition, in the user-domain, the privacy-analyzer allows reducing the attack surface preventing privacy breaches when sensitive personal data are managed.

Decentralized authorization, privacy-preservation and distributed access control are also important features considered in this architecture. In the Blockchain privacy-preserving SSI Layer, this is achieved by means of building blocks that are aimed at making blockchain technologies and consensus mechanisms more scalable, efficient, guarantying on-chain transactional privacy. Besides, it includes building blocks for modifying transactions (fine-granular rewriting) already present in the blockchain in a limited and traceable manner, which may be important for legal reasons.

The architecture considers privacy-preservation of identities and personal data in blockchains. To that aim and following the Decentralized Identity Foundation (DIF)[1] standards and specifications, the architecture features the building blocks needed for the creation, resolution, and discovery of decentralized identifiers (DID identifiers[2]) and names in heterogeneous blockchains through resolvers. In addition, the Identity Hubs keep secure, encrypted, privacy-preserving personal data storage and computation of data. Where the resolver services link user's DID's employed in blockchain with Identity Hubs. The blockchain Identity services provide the means to create, exchange, and verify crypto credentials and claims in a decentralized identity ecosystem with the User, following a self-sovereign identity management model. Besides, the blockchain identity services might rely on authentication protocols open standards and cryptographic protocols, including DIDs and DID Documents.

Another group of solutions is intended to enable privacy preservation in Cloud computing environments as well as its extension towards the user side with Edge computing. The Privacy-Preserving architecture provides building blocks for secure data storage and processing in public clouds. In particular, it considers distributed data storage and privacy-preserving analytics as well as mechanisms for compliance with the provisions of GDPR regarding interoperability and cross-border data transfers.

The Edge is considered in this architecture as a security and privacy enabler especially for the IoT domain, where devices are typically extremely resource-constrained and may be subject to compromise or interference. In this respect, the proposed architecture includes a data broker for both handling sensitive data according to a set of privacy policies as well as tools for monitoring and sanitizing IoT devices for reducing the attack surface in this domain. Likewise, the privacy-preserving architecture considers the privacy-preserving middleware and software for the IoT domain aimed to ensure secure and authenticated communication channels between IoT devices. The managed domain in the global IoT architecture can be also instantiated through processes related to the Web domain (e.g., eCommerce) in the CyberSec4Europe privacy-preserving architecture. In this case, the Web domain is comprised of a set of functional components needed for the Service providers to authenticate their users, verify claims and privacy-preserving crypto-proofs (e.g., Zero-knowledge proofs). These service providers play the role of a Verifier in the privacy-ABC model.

Finally, our privacy architecture also considers the application of security and privacy by design mechanisms by introducing components for GDPR-compliant software development as well as analyzing the information leakage produced by some particular privacy solutions.

[1] DIF Identity Foundation. https://identity.foundation.org.
[2] Decentralized Identifiers (DIDs) v1.0. W3C. November 2019. https://w3c.github.io/did-core/.

Asset	Partner	Services Plane	User domain	Administration Plane	Intelligence Plane	Control and management plane	Blockchain Plane	IoT domain	Web domain	Covered
Issuer-Hiding Attribute-Based Credentials	AIT		x			x				In this deliverable and in WP5
Cloud-Based Anonymous Credential Systems (eABCs)	AIT		x			x				In this deliverable and in WP5
FlexProd	AIT					x				In this deliverable
ArchiStar	AIT		x							In this deliverable (as submodule of FlexProd)
SS-PP-IdM	UMU	x	x							In this deliverable and merged with Mobile p-ABC and eIDAS browser
Password-less authentication	UPRC		x			x			x	In this deliverable
Edge-Privacy	UMA							x		In this deliverable
Privacy-Aware Aggregate Programming	DTU			x						In this deliverable, replacing the effort with AntibIoTic
DANS	ATOS		x			x				In this deliverable and evaluated in WP5 – including the previous affords of SPeIDI
Cryptovault	VTT		x							In this deliverable

Elastic Deployment of TEE-based applications in the cloud	NEC		x							In this deliverable
Backdoor-resistant TEEs	NEC							x		In this deliverable
Blockchain Platform	NEC						x			In this deliverable and also in WP5.
Sharemind	CYBER		x							In this deliverable, including PLEAK differential privacy analyzers
Privacy-Preserving for Genomic data	UNILU					x				In this deliverable with a specific genomic use case
GENERAL_D	CNR		x			x		x	x	In this deliverable, in deliverable D3.13 and also in WP5
PTASC	C3P							x		In this deliverable and also in WP5
ARGUS	C3P					x				In this deliverable and also in WP5
GDPR compliant user experience	UM			x	x					In this deliverable and also in WP5.
Interoperability and cross-border compliance	UM			x						In this deliverable.
PLEAK Differential Privacy Analysers	Cyber		x			x				In this deliverable and also in WP5.
Privacy-preserving Federated Learning	UMU					x				In this deliverable and also in WP5

Table 1: Mapping of assets available and the CyberSec4Europe Privacy-Preserving Functional Architecture

3 Description of the Smart Campus Scenario

It should be noted that this description was already included in D3.13, and it is added here in D3.20 only for the sake of completeness of the whole document.
The Internet of Things (IoT) allows everyday objects (equipped with computational and communication capabilities) to connect to the Internet. The "things" can exchange data with each other and with the Internet, making decisions automatically, even without human interaction. IoT applications demand platforms that aim to facilitate their development process, which may involve integrating a diversity of heterogeneous devices with varying capacities, means of data transmission, and different communication protocols. The literature presents several middleware platforms that serve as the underlying infrastructure for the development of IoT applications [mineraud2016gap, ngu2016iot].

A smart campus uses technology solutions to manage services and users' life experiences. Sensors, networks, and applications are used to collect relevant data, such as the number of rooms used, energy use, and air quality. This data can be used to improve the smart campus services [sari2017study]. An example that improves the smart campus is the management of the identity of users when interacting with the smart campus where heterogeneous services will be available to potential users, with services for direct interaction with the University (e.g., enrolment in courses or activities). These services may have widely varying requirements for their usage. A public transport service may be available for any user that interacts with the platform, while some specific academic information should be restricted to users that are students from the university or even enrolled in a specific degree/course.

Cloud-based IoT applications receive, analyze, and manage data in real-time to help institutions in the smart-campus, businesses, and citizens make decisions that improve the quality of life. Students engage with smart campus ecosystems in a variety of ways, using smartphones and mobile devices, as well as connected transportation and homes. Pairing devices and data with the infrastructure and physical services can reduce costs and improve sustainability.

The smart campus allows a set of known interactions Services such as (an overview is represented in Figure 2):

- Public transport,
- Parking availability,
- Campus information,
- Academic information,
- Mobility data analysis (privacy-preserving).

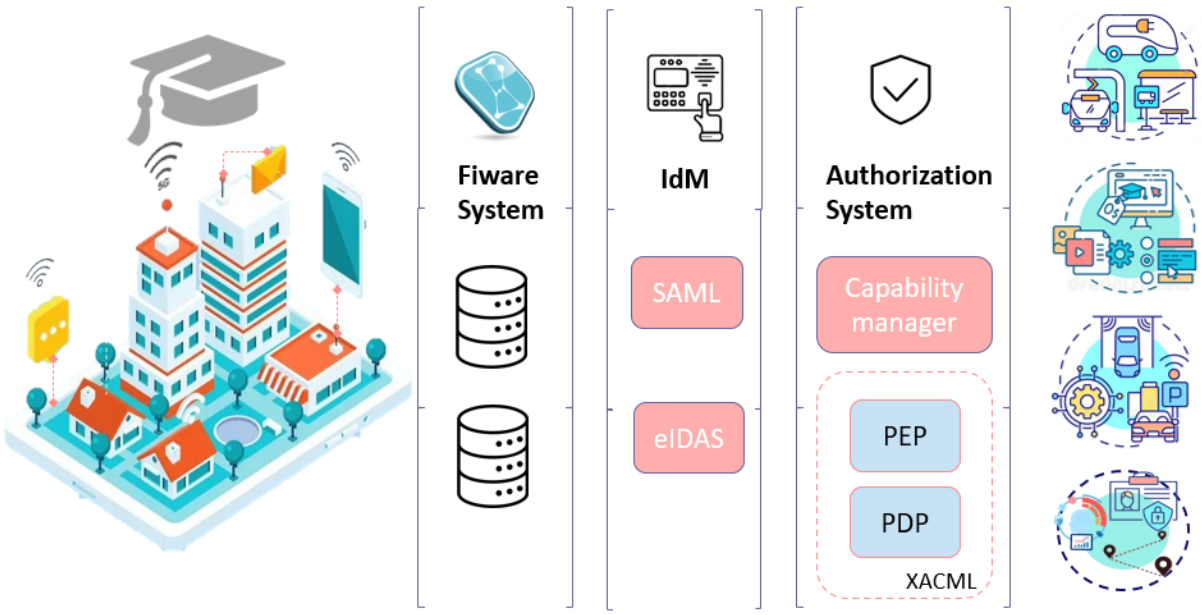

Figure 2: Smart Campus

A smart campus is composed of multiple infrastructure and components, such as:

- Smart-Building
- 5G - Base stations, WIFI Access Points (Edge Nodes)
- Smart-Campus IoT platform:
  - Fiware System (Orion Context Broker). Data Base that stores IoT data obtained from sensors (devices, users and services) in the smart campus.
  - IdM (Fiware Keyrock). Identity management system used for basic identity management and bridge to eIDAS (i.e., handles SAML communication flow with eIDAS node to obtain certified attributes).
  - Authorization System based on capability access control tokens. It includes (Policy Enforcement Point –PEP--, Capability Manager and Policy Decision Point --PDP)
    - **PEP:** Controls access to the services, checking that the request includes a valid capability token (i.e., the request is authorized).
    - **Capability Manager:** Generates capability tokens that bestow authorization to use specific services. Relies on the PDP for the decision (using XACML).
    - **PDP:** Checks if an authorization request should be conceded
  - User Mobile app: allow users to access Smart-campus services.

Given the multiple domains addressed in the smart campus and according to the partners' feedback, we decided to set up a smart-campus high-level scenario from which we create 3 sub-scenarios or use cases. In this context, we will focus the next subsection on 3 main sub-scenarios from the general smart-campus scenario: CCTV Surveillance in the Smart-Campus; Identity management and service usage in Smart Campus; Geolocation Service in the Smart Campus. These sub-scenarios focus on the challenges addressed by the assets and available on the smart campus. Table 2 identifies the main partners/assets integrated into each sub-scenario of the smart-campus. The only partner who does not have assets directly integrated into the smart campus scenarios is UNILU since the asset focuses on privacy-preserving of Genomic data, which is unrelated to the smart campus. However, we still show this asset in Table 2.

Use case	Partners	Assets
CCTV Surveillance in the Smart-Campus	C3P, DTU, CNR, UPRC, UM	ARGUS, PTASC, Password-less authentication, GENERAL_D, Interoperability and cross-border compliance, Password-less authentication
Identity management and service usage in Smart Campus	UMU, AIT, UM, VTT, NEC, UPRC, CNR	Blockchain Platform, SS-PP-IdM, Cloud-Based Credentials, FlexProd, ArchiStar, Issuer-Hiding Anonymous Credentials, Cloud-Based Credentials, GDPR compliant user experience, Password-less authentication, Cryptovault, GENERAL_D
Geolocation Service in Smart Campus	CYBER, UMA, ATOS, NEC	Elastic-TEE, Sharemind, Edge-Privacy, Back-door resistant TEE
Genomic data	UNILU	Privacy-Preserving for Genomic data

Table 2: Assets/partners mapped by use case

3.1 CCTV Surveillance in the Smart-Campus

This section presents an application use case as a motivation for the demonstrator. CCTV is a common scenario in smart buildings, where a response team continuously monitors the campus for accidents, fires, to name a few. In case of emergency, police may be dispatched to the incident. However, the police must always have access to the video surveillance to gather information about incidents/unwanted situations, such as:

- A robbery is happening, and the robber is moving in a specific direction.
- A child is missing and the police needs to analyze the CCTV feeds for possible locations of the child.
- A bag is stolen, and the police officer needs to identify a possible suspect.

. However, depending on the authorization level, users can have different qualities of the CCTV recordings. A normal user can only be authorized to track traffic lights or traffic jams but not extra information such as person faces, clothes, license plates, or any other information that can re-identify a person. However, this can be done using ML algorithms and face defacement techniques. Then, suppose it is an operator (or police, or other members with special authorization level) from the university tracking the same information, in that case, they should be able to see the information with more quality, to allow tracking cars from the municipalities. However, in this case, the faces must always be anonymized[3]. In case of emergency, we have policies that enable to track in real-time a specific user or license plate of the car but the police may not have sufficient permissions to access the specific CCTV feed and the request for higher permissions can be slow and inhibit the viewing of crucial images. To solve this latency, we propose using a system that uses a break the glass, where the police will have access to the required information immediately. Break the glass is generally used to do something in an emergency, especially in a medical or fire context, and refers to a quick means for a person who does

[3] https://www.theverge.com/2020/6/11/21280293/anonymize-blur-faces-photos-videos-camera-app-ios

not have access privileges to certain information to gain access when necessary. In this case, a pop-up will be shown that notifies the user (either a police officer or a user with specific access to the information), informing that their action will be registered on the system logs and, later on, it will be verified by a police chief department[4] to validate the reason for accessing that specific information. This will allow a faster response in cases of emergency and allow users to access information available from the CCTV feeds, which otherwise will only be accessed by an operator in a control room from the university. Also, the operators from the university will be capable of accessing the CCTV feeds on their smartphones depending on the policies of the institution.

3.1.1 Architecture

Our architecture includes users, a video surveillance feed, an authentication and access control management, and an offline verifier. The general architecture describes a scenario where users access surveillance videos through an application. The users must have permission to access the videos, and this decision is made by the authentication management systems. The authentication depends on the type of authorization the person has, that is, users can have access to the videos in high quality, without quality, or with the faces blurred, creating the possibility of having different granularity levels. Still, there is also a registration mechanism for emergency access called Break the Glass. However, as it is an emergency case, it must be registered to be identified by an offline verifier if it is improper accessed. This register can be implemented as an immutable database, using blockchain technologies, in order to prevent the manipulation of data on the blockchain.

From the challenges from D3.11, we focus mainly on:

- DP-06 by providing mechanisms for control how the information is disseminated and control the private data on the communication.
- IDP-05 by providing a mechanism for guaranteeing proper identity management of things for authentication end-to-end.
- DP-07 by providing anonymization mechanism for control on how the information is stored and control the private data on the communication.
- DP-05 “Lack of mechanisms for controlling and limiting access to the data collected from numerous and geographically disperse IoT device
- Honest but curious Service Providers and Identity Providers that might be tracking users, IdP could also become a potential point of failure (identity/data leakage in the IdP) (IDP-04).

Figure 3 represents the architecture with the identification of the steps to implement the CCTV scenario. The explanation of how this scenario maps onto the general architecture shown in Figure 1 will be discussed later.

[4] This person in qualified to judge any misleading access to the information of the CCTV scenario.

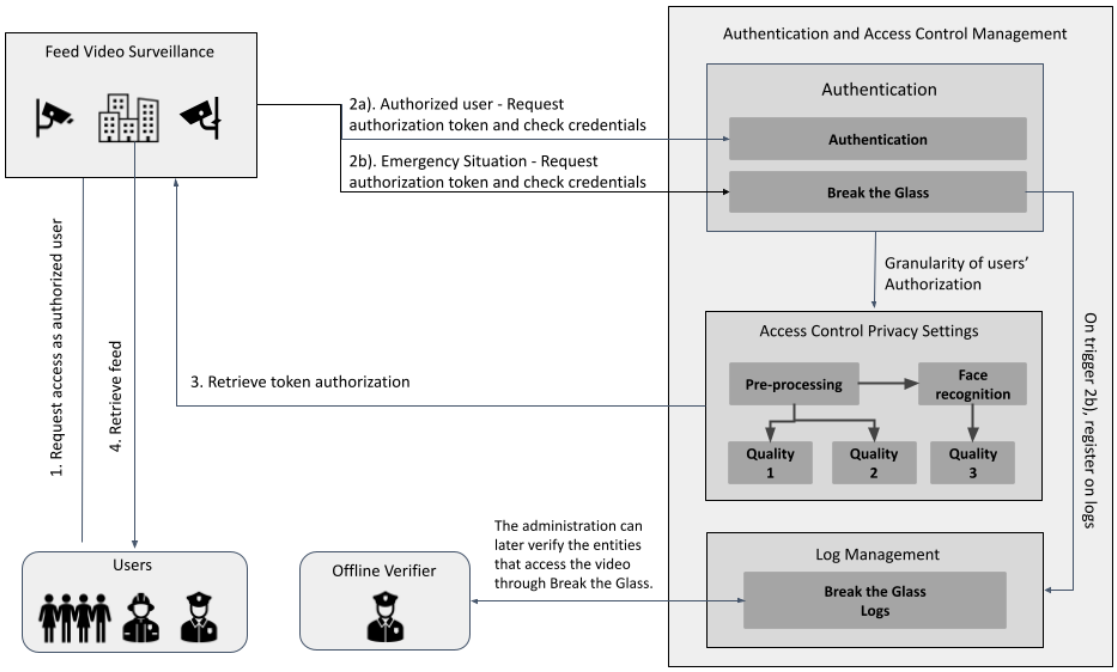

Figure 3: CCTV Scenario Overview

Video Surveillance Feed

The video surveillance feed is composed of a live stream of surveillance cameras, authenticated by login with username and password. The feeds are available depending on the logged users' authorization. For different users, it provides different video qualities and filters to comply with the data minimization principle.

This interface has an administration system that is accessible and protected by a username and password and which should only be accessed by the person in charge of the entities for managing video surveillance and security. In this interface, the admin can create or eliminate users and groups and set policies for them that define the different access controls for each entity.

Within the administration panel, it is possible to access records and logs (of those who accessed the system) for control and auditing purposes of the entity's internal platform, set the cameras that users and groups can access (e.g., by sectors), enable/disable filters for specific groups and users, remove non-accurate data and limit/define the period of conservation/retention of videos according to the GDPR.

Storage

Videos are stored for a limited time (defined by the DPO in the administration panel), but users who access these videos do not need to access all the information. Following the principle of data minimization, information should be reduced according to purpose and limited to what is strictly necessary. Thus, personal data must be hidden (for example, use filters to hide faces and license plates) to minimize exposure of personal data.

By default, the system automatically provides privacy filters and low-quality videos for users with general permissions.

The responsibility of the processing of information within the conditions defined in Article 9 from GDPR is from the administrator.

This storage is encrypted so that there is no access without permissions. Data access permissions are defined by a competent person (for example, DPO). Also, there are data retention policies because the lack of discipline around data lifecycle management means that organizations often hold onto data that add little value but lots of risks.

Authentication and Access Control

There are two types of authorization: Users and Groups. A group can have multiple users with the same policies. Individually, each user can have specific policies.

There are different authorization levels for different live feed streams. The first level presents the feed in high quality - the raw as the video camera produces it. The second level provides the image in low quality so that the image in high quality is not noticeable, for scenarios when it is not necessary to see all the details of the information. Lastly, the strictest level has privacy filters with face detection and concealment and detection of other types of personal data, such as license plates. For this type of personal data, it must be applied the appropriate filters, with the same type of policies also being applied to video storage. Our system only includes face detection, but the system is modular to include other filters.

Break the Glass

This system includes a recording mechanism for emergency access, called glass breakage. Breaking glass is often used to do something in an emergency, especially in a medical or fire context, and refers to a quick way for a person who does not have access privileges to information to gain access when needed.

This access breaks a security and privacy mechanism, especially concerning GDPR. The action of breaking the access control protection is a high-risk action, so it must be registered and identified and cannot be tampered with so that it can be audited later.

When there is an access attempt, the user must agree to record that intention and write a description to define the purpose of the access. Thus, it is possible to define the user's purpose limitation to the data.

An offline verifier - who must be registered and authenticated - must be promptly notified and verify the intent data. Furthermore, it can revoke user access (in real-time) to the data.

The access token given to the user in "break the glass" mode has a limited time and can be revoked during its use. Time is short (default 5 min) and can be adjusted by the responsible person. To access it again, users must accept the conditions, and the system records the identity and the timestamp.

The following sections describe the specific enablers based on Figure 1, which represents the CyberSec4Europe Privacy-Preserving Functional Architecture [D3.2] and address the challenges enumerated previously.

3.1.2 Administration Plane

Techniques for providing password-less authentication often use some form of biometric data for the purpose of identifying users (primarily because of its convenience). The use of biometric data is especially popular with mobile devices – e.g., smartphones as proposed in the Password-less Authentication demonstrator. Biometric data, which is considered as a special category of personal data in the GDPR, when it is used for the purpose of uniquely identifying a natural person. Processing of such data is typically prohibited unless one of the exceptions (e.g., explicit consent) defined in the second paragraph of the GDPR's Article 9 applies. Member States can also introduce their own conditions and/or limitations. Before implementing any such authentication solutions, it is therefore good to know whether local, national legislations extend the GDPR and when the use of such biometric data is permitted (this tool can be used generically for other types of data).

3.1.3 Control and Management Plane

3.1.4 Control and Management Plane

The storage of data is maintained on the public Cloud so that we keep the costs of this type of system very low.

The logs and video storage collected in this context must be kept privately and securely stored on a system that ensures privacy and data loss protection.

This can be achieved using ARGUS (Section 4.2), as a privacy brokerage system to enhance trust in public cloud storage systems, guarantee data confidentiality, and improve the information's availability, a comparative study is provided by João et al.[argusprivacy]. The focus of ARGUS in this scenario is to store the relevant information using a decentralized approach, where the public cloud providers don't have access to stored information (assuming that more than one cloud storage provider does not collude), this is accomplished by using erasure coding techniques. Only ARGUS can reconstruct all the bits of information to produce the original stream of video or logs. This also protects the privacy of information since all information is maintained in a secure public cloud.

It is also important to ensure the integrity of data because it is important not to tamper with the logs, modify the original identification of the person that accesses through the break the glass mechanism to the video surveillance. Thus, it is essential to have a tamper-resistant and time-stamped database.

Moreover, the security of the authentication and authorization processes is crucial in this scenario. It is known that passwords are targets of multiple attacks, as they can be leaked, key-logged, replayed, eavesdropped on, brute-force decoded and phished. Therefore, the password-less authentication asset is integrated with the Identity and Access Management system to enhance the authentication process. With the deployment of password-less authentication, the following aspects will be accomplished:

- Deployment of advanced authentication methods that combine strong cryptographic functions (e.g., FIDO protocol) with something the user knows (e.g., PIN), something the user has (e.g., USB key), or something the user is (e.g., biometrics).
- Utilizing a two factor authentication (2FA) mechanism. Depending on the authorization that each user has, one of the authentication methods described above will be implemented.

3.1.5 User Domain

As the number of accounts each user maintains has greatly increased in the last few years, users are having a hard time memorizing and managing all these passwords. To solve this password overload problem, users have come up with solutions that compromise the security of their accounts and the privacy of their data. For example, users either simplify their passwords to be easy to remember, or reuse the same password on different services, or store their passwords in a "secure" place, on paper, or use a password manager. Furthermore, the password overload problem significantly affects the usability of an application. Password-less authentication not only will increase security, but also will provide a more user-friendly authentication process. With its implementation users will not have to manipulate complex passwords to login into the application, instead, they will deploy the authentication method they prefer from a variety of available options (e.g., USB keys, biometrics, PIN, etc.).

In this process, particular attention will be devoted to developing a user-friendly mechanism for the management of personal data requirements (such as consent and purpose), and to automatically deriving GDPR-based access control policies useful for enforcing the user's preferences.

3.1.6 IoT domain

The Edge is considered a security and privacy enabler in this architecture, especially for the IoT domain, where devices are typically extremely resource-constrained and may be subject to compromise or interference. In this respect, the proposed architecture includes PTASC (pTrust Autonomous Secure Communication) an end-to-end identity provider (Section 4.1). This privacy component is a new

mechanism for authentication using a secure tokenhat allows extra control over the devices and has already been deployed in the smart-cities context.

In this case, the secure token will be used in the CCTV cameras' main categories and will allow two main privacy protection in the system: end-to-end communication with the central server; and privacy guarantees on the network to ensure that attackers do not try to connect to the system or connect to unauthorized remote servers.

Additionally, a GDPR-based Access Manager can manage data access in compliance with the GDPR and the data subject's consent.

3.2 Identity management and service usage in Smart Campus

This is another application case that justifies the selected demonstrator. In a Smart Campus, multiple heterogeneous services will be available to potential users, with services for direct interaction with the University (e.g., enrolment in courses or activities). These services may have widely varying requirements for their usage. A public transport service may be available for any user that interacts with the platform, while some specific academic information should be restricted to users that are students from the university or even enrolled in a specific degree/course.

Thus, the Smart Campus infrastructure needs to provide the means to perform authentication and authorization. Privacy concerns and related regulations like GDPR (which points to minimal disclosure, consent, etc.) must be taken into account. Indeed, users need to be empowered in their control over their identity, and over how much identifying information is shared in their interactions with the platform. In this context, we propose a capability-based authorization framework for the Smart Campus platform. To earn authorizations, users must authenticate using some properties of their identity, e.g., revealing an attribute or showing that it fulfils a condition (like being born in a specific period).

Within CyberSec4Europe, multiple related approaches to this challenge have been pursued, partially based on assets that were already developed before the start of the project, and also inspired by different requirements from the demonstration cases defined in WP5. Based on the concrete requirements, any of the assets below is used.

- SS-PP-IdM ensures privacy by design in authentication processes, leveraging p-ABC technologies. The usage of eIDAS for attribute population and DLTs for auditability and distributing public parameters gives strong trust assurances to the solution.
- While not focusing on auditability, eABCs are rather focusing on the scenario of highly resource-constrained devices such as smart cards on the end user side. eABCs allow one to outsource a large fraction of the computations to a largely untrusted cloud-provider without impacting the end-user's privacy.
- Finally, issuer-hiding ABCs consider the case where the precise issuer of a certificate is revealed. This might be interesting in the case of exchange students, where the issuer of the certificate might already uniquely identify the student within a university campus.

During student enrolment, the university collectes and processes an extensive collection of personal data of enrolled individuals. In the scenario of Smart Campus, Data Protection Impact Assessment (DPIA) is required on the basis that the processing involves the use of new technologies, processing affects a large number of data subjects, can include monitoring of publicly accessible areas on a large scale, and/or systematic monitoring. After the assessment shows that the use of personal data does not cause a high risk to the rights and freedoms of data owners (i.e., students), the University/Smart Campus can start using this data to provide their services. To help perform the assessment, we leverage the DPIA template we have designed in CyberSec4Europe.

Once students enroll in the university, the user's identity is augmented with extra information, like a university-specific identifier, an e-mail or other second factor authenticator. With the capabilities of our ABC-based assets, users can obtain credentials that attest with strong trust their personal and student

information, while ensuring that their privacy is kept when interacting with services. A typical example showing the advantages of this kind of system over existing approaches is the obtaining of student discounts in different services. Showing a student identification card (in “physical” scenarios, like buying a movie ticket) or sending university registration information or documents (requested in streaming services like Spotify) reveal more information than necessary, like full name, date of birth, etc. Instead, with a couple of clicks, users could leverage our privacy-preserving identity management solutions, so they learn, and control exactly which information is being shared. This would enable privacy principles like minimal disclosure (e.g., just proving you are a university student, but not revealing any extra information) or informed consent (i.e., the cryptographic mechanisms ensure that only the agreed data is revealed). The whole framework of the IdM ecosystem and the soundness of the involved proofs will lead to strong trust by service providers even in this privacy-preserving environment.

3.2.1 Architecture

From the general architecture, this scenario focuses on most of the identity management related components, as well as the service usage and the required authentication (highlighted in Figure 3). A more detailed vision of the internal architecture and interactions between the components can be seen in Figure 15. There, specifically, the SS-PP-IDM is shown to rely on the blockchain instance and a distributed identity provider, which issues p-ABC credentials. Users then take advantage of those credentials to carry out the authentication processes through zero-knowledge proofs. The high-level architecture for the other two assets is similar yet does not include the distribute-ledger components.

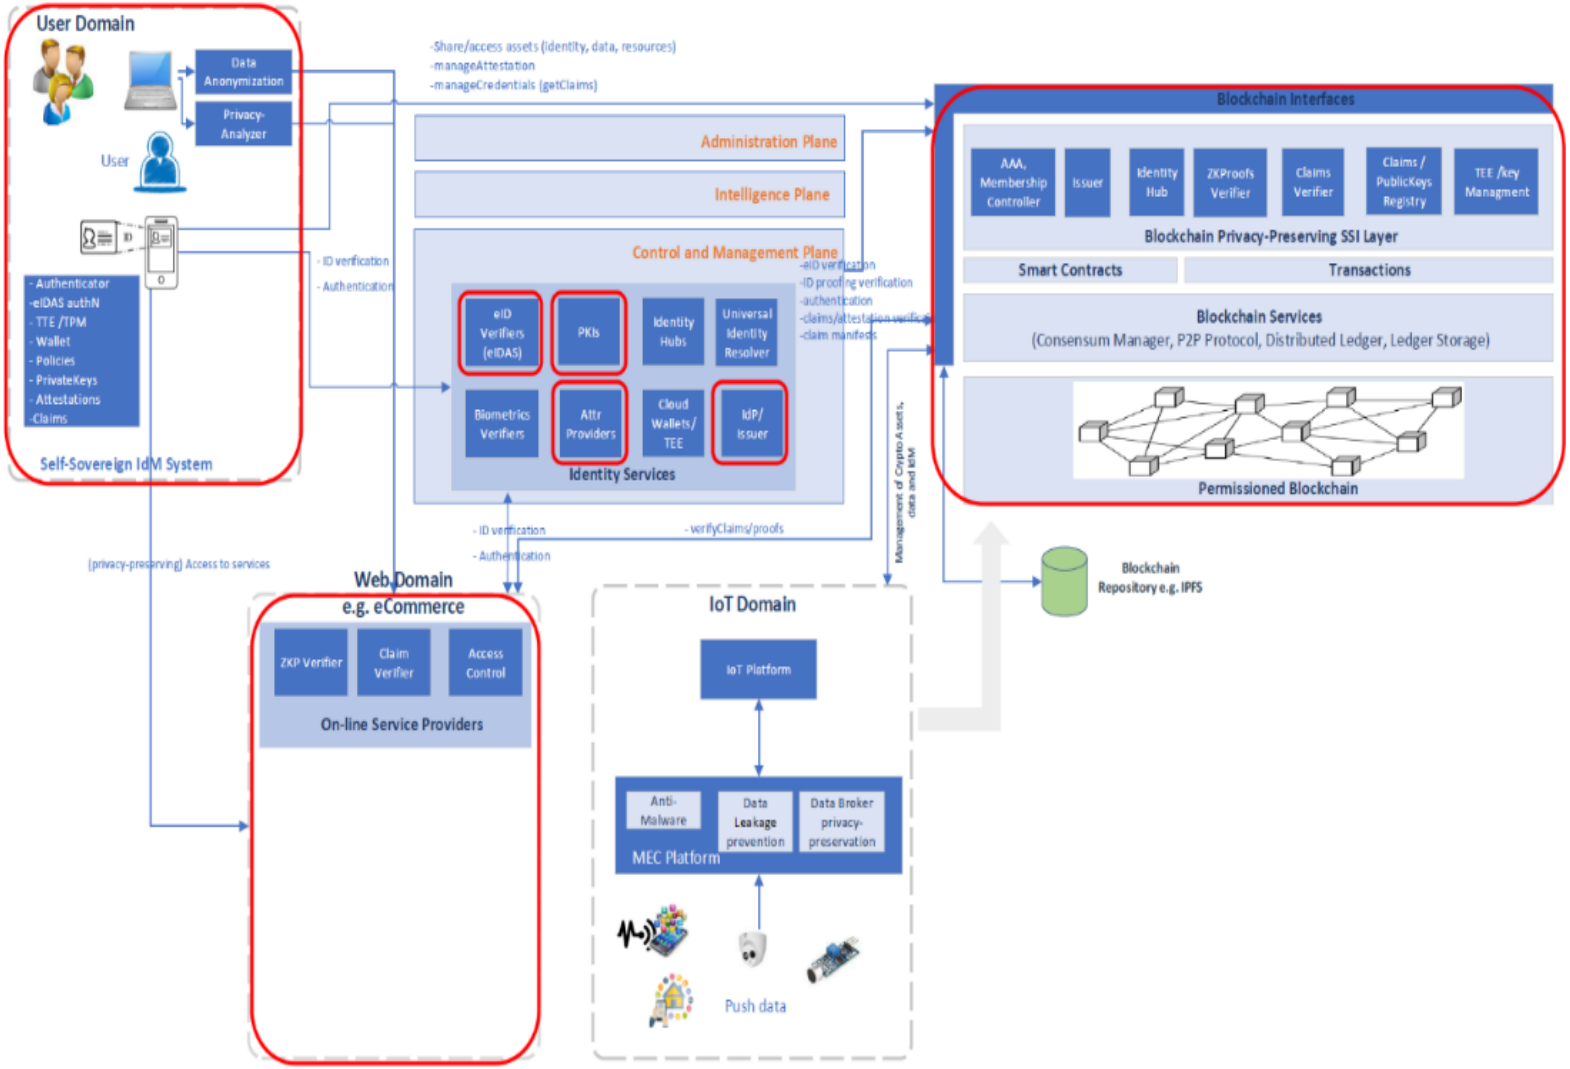

Figure 4: IdM scenario architecture overview and instantiation

From the challenges described in D3.11 this scenario addresses mainly the following:

- *IDP-02: Unnecessary over-identification and information disclosure due to a lack of awareness and usability drawbacks.*
- *IDP-03: User's privacy-preservation of transactions in distributed and immutable systems (e.g., blockchains).*
- *IDP-04: Honest but curious Service Providers and Identity Providers that might be tracking users, IdP could also become a potential point of failure (identity/data leakage in the IdP).*
- DP-08 *When uploading information to the cloud the user partially loses control over the data.*

3.2.2 Intelligence Plane

Before a service using personal data can be designed and implemented, it is important to comply with relevant regulations. Recent security and privacy relevant examples include eIDAS[5], GDPR and the upcoming ePrivacy[6]. One important step prescribed in the GDPR is the Data Protection Impact Assessment (DPIA) which should be done any time processing of personal data is likely to result in a high risk to violate the rights and freedoms of natural persons. The assessment serves as a University's/Campus' evaluation of compliance to the GDPR requirements and assessment of the impact the risks of the envisaged processing operations will have on the protection of personal data used in providing Smart Campus services.

3.2.3 Control and Management Plane

The Identity Provider must ensure that users are actually in possession of the attributes that they claim form their identity. To this end, it must collaborate/rely on external trustworthy attribute providers (which are supported by a PKI).

In this sense, eIDAS nodes are a great source of reputable identity information. Other attribute providers may be useful in this scenario. For example, the Smart Campus platform itself may provide an identifier for a user, while the university framework can provide other attributes like the user being enrolled in activities or courses.

The main function of the Identity Provider is issuing credentials to users so they can be used for authentication. In that task, the DLT infrastructure supports distributing the public parameters (e.g., verification keys) to all the entities that need them. However, it must assure that third-party elements cannot exploit the issuance process to obtain falsified/stolen credentials.

The Identity Provider will also need to fulfil all the account management operations contemplated in regulations like GDPR and in most use cases. Namely, it must allow users to manage the attributes associated to their account (add/remove) and delete the account altogether, among others (managing metadata like password if necessary…). Lastly, depending on the use case, it may be necessary for the management plane to have the means to suspend or remove accounts that misuse their credentials.

3.2.4 User Domain

Users access and interact with services using the available interfaces (e.g., a mobile application). To do that, they need to comply with the access requirements and policies. Identity management tools are necessary for the involved verifications, so they must be included in user applications. While most

[5] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv%3AOJ.L_.2014.257.01.0073.01.ENG
[6] https://digital-strategy.ec.europa.eu/en/policies/eprivacy-regulation

details about the underlying identity management solutions do not need to be apparent to users (e.g., how complex is the distributed identity provider, how the underlying cryptography works…), users will have some direct interactions related to identity management:

- During enrolment phases, users may need to do actions specific to the attribute providers. For example, using their eID for identification in eIDAS.
- For the issuance process, users have to identify against the identity provider, e.g., through traditional username and password and optionally two-factor authentication.
- For service usage, users will be prompted to consent to access policies, revealing information about their identity.

3.2.5 Blockchain Interfaces

When using any blockchain technology, the security of the private key is paramount. Typically, one would use a wallet application to interact with the blockchain: wallets are used to store the private keys, generate signatures, and encode the transactions on behalf of the user. Online wallets require a lot of trust in a third-party service provider, while mobile and desktop wallets place the onus on the user to choose how to keep their keys safe, balancing delicately between convenience and security. Hardware wallets are dedicated devices that host keys. They are very secure and the user is fully in control of the keys, but on the other hand they are not as convenient as the other wallet types and losing or breaking the device means that the keys are forever gone.

The technical know-how of the user is important to consider when choosing a wallet type; in the case of the Smart Campus, IT administrators can be considered technically savvy, and we can assume that in a professional context they need to prioritize security and control of the keys. Therefore, a hardware wallet would be a likely choice. The security of the private key is not only dependent on the secrecy (or confidentiality) of the key, but also on its availability and integrity: what if the key is somehow lost? It is important to have a reliable and secure backup mechanism. CryptoVault is used to demonstrate a secure hardware wallet and backup method for the private keys of blockchain applications.

3.3 Geolocation Service in the Smart Campus

As in smart cities, geolocation in the smart campus can be used to detect trends in people's movement. This in turn can help with planning public transport, urban planning, creating safer and more resident-friendly areas. However, geolocation data has a lot of sensitive information that can be used adversely. In this scenario, we look at outdoor geolocation based on GPS or cell tower signaling.. HHowever, the demonstrated assets are independent of the kinds of data collected and can be used for other data analysis as well.

In this scenario, we provide a means for analyzing the data and detecting patterns in a privacy-preserving way, so that the benefits of the data can be made use of without seeing any individual's records.

3.3.1 Architecture

From an architectural point of view, this scenario focuses on the User and IoT domains, as depicted in Figure 4.

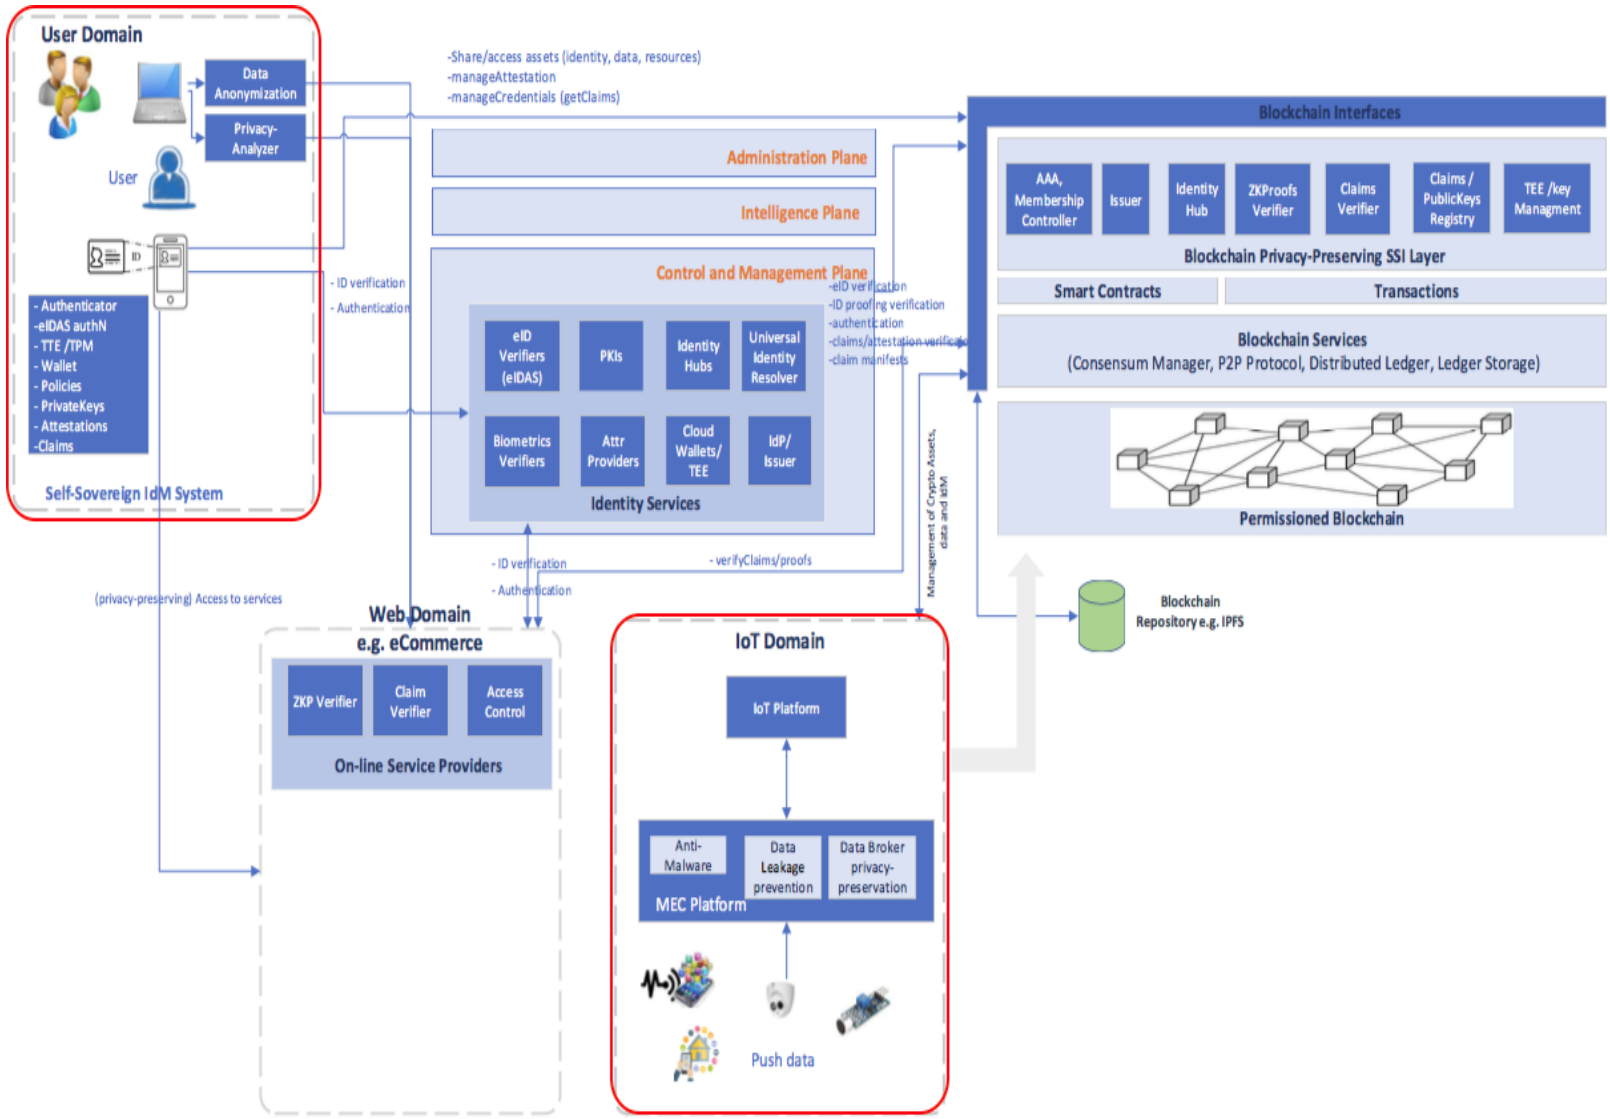

Figure 5: Mapping of the Geolocation Service in the General Architecture.

Data collection. We can have multiple sources for location data. First, sensor-based location can come from mobile phones and tablets with global positioning service (GPS) capability. These devices can determine their own location, combine it with a timestamp and encrypt the data. The data is uploaded then to the mobility analysis service, but without sharing the decryption key with the untrusted service provider.

Alternatively, location data can be collected from telecommunications operators who have bulk location data available from cell tower signaling (approximate triangulation of phone users' locations). This data is collected by the operators who can encrypt it in bulk and share it with the secure computing system. These two deployment options can be seen in Figure 5. Alternatively, the data set collected by the telecommunication operators could be anonymized to preserve the user's personal data privacy and facilitate the data analytics later on.

Data analytics. To calculate mobility statistics from the encrypted or anonymized data, the service uses a secure computing platform, e.g., a trusted execution environment (TEE). The exact method depends on the chosen technology.

Result presentation. Once the data preparation and analytics are completed, the results are stored in an encrypted form on the platform. Authorized visualization services can then download the resulting mobility statistics and visualize them for the user. Different users can be allowed to have access to data using different views, as long as privacy is guaranteed. Remote attestation can be implemented to ensure that the correct code is being executed.

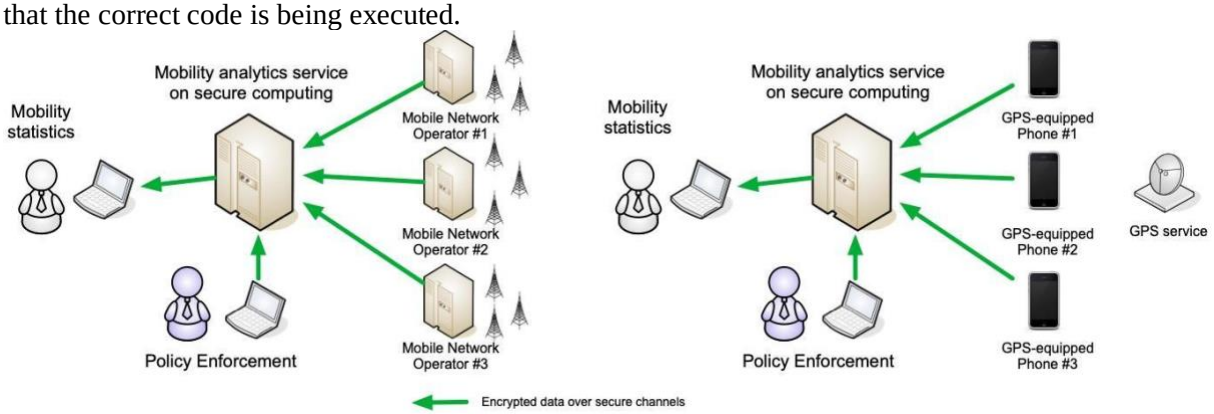

Figure 6: Two deployment options for secure computing

Research challenges. The main research challenge that the privacy-preserving data analysis of geolocation data addresses, is privacy by design. By using secure computing technologies like TEEs or secure multiparty computation, the privacy is built in at the start.

The more specific challenges are the following.

1. Data privacy challenge DP-02 (when using secure computing to analyze data, analysts are not able to see the individual data values).
2. Moving the internal state of an enclave to another platform is at odds with the main security provisions of SGX – the state of an enclave is private to the enclave itself.
3. TEEs need a subversion-resilient attestation mechanism that requires no secret (or secret share) on the hosting platform.
4. The privacy manager must be able to support the mobility of IoT devices without affecting data availability.

3.3.2 Control and Management Plane

This scenario uses privacy preservation tools to perform the data analysis to make the movement patterns of individuals invisible to the data analyst. We propose using a trusted execution environment (TEE) to store the encryption key inside a programmable hardware security module that loads the encrypted

location data from storage and processes it within the same environment. At all other times, data remains encrypted, and the encryption keys do not leave the hardware security module. This way, the data analysis can be carried out in a privacy-preserving manner so that even system administrators do not have access to individual values. In the case of anonymized data, the data analysis can be performed without compromising the user privacy as the anonymized data are no longer personal data.

Sharemind HI is used as the TEE whereas SR-EPID is used as a drop-in replacement of the attestation protocol in order to guarantee security even in case the trusted execution environment is subverted, whereas the Elastic TEE asset ensures that computing resources scale with the workloads.

3.3.3 User Domain

The application provides a PET client which is familiar to a data analyst (a language similar to R) that can easily be used to carry out data analysis or machine learning. The analyst will not be able to see individual records, just the aggregated results. The analyst does not need to know the details of how the underlying system works, the application provides them with a familiar workflow.

The Privacy Manager from the Edge-Privacy asset can be applied as a means for the controlled release of information collected from devices belonging to Smart Campus students. The students' smartphones can be configured to upload their data to a Privacy Manager under the control of the user, which has been configured to encrypt location information before storing them in the Edge. The Privacy Manager can then receive queries from third parties and after checking, they are entitled to get access to the location information of the user, which can then be entered into the central system running the TEE.

3.3.4 IoT Domain

The location information necessary for the realization of this scenario may be obtained from IoT devices belonging to members of the university, including students and staff. Wearables such as smartwatches or smartphones can provide location information since they are fitted with positioning technologies such as GPS. However, not all these devices will be able to provide data in encrypted form.

The data collected from user related IoT devices can be uploaded to the Edge Platform where the Privacy Manager resides. Upon the reception of location information, the Privacy Manager will process the data so that it is securely stored by means of encryption. The Privacy Manager also serves as an interface to control who has access to these data.

4 Asset Demonstration

In this section, we describe the tools for the demonstrators in more detail. For each tool, we first overview the components/architecture of the tool and include diagrams/screenshots to illustrate how the tool handles the challenges highlighted in the previous chapter. Then, for each asset we address a simple example on the integration on the smart campus integration demonstrators.

For each asset, we also provide a video to explain the main concept and integration in the demonstrator[7].

4.1 PTASC

4.1.1 Overview

PTASC presents a decentralized, secure device-to-device communications solution in which device provisioning is focused on improving usability while providing security by default. The solution focuses on using a PKI where the CA is represented by a manager device that can be switched on/off to reduce single point of failure (SPOF) problems. The solution combines public-key cryptography and symmetric keys with the One Time Password (OTP) concept using a secure token. Device identity is guaranteed by physical access to this physical token. In addition to generating an OTP, the physical token also stores

[7] https://www.youtube.com/channel/UCSAJ78frZjdUTooAC4t6Wuw

a public key to be transmitted to target devices only, eliminating attacks such as impersonation or man-in-the-middle. It also improves usability as we exclude configuration errors and difficulty choosing the right settings while provisioning the device. Although there is manual interaction to use the secure token, the process itself is as simple as finding the device to be provisioned and plugging in the security token. Along with the authentication and secure communications, PTASC encompasses a middleware layer solution for the IoT devices, which allows the control of the data generated by the device by its owner.

Manager Setup Phase

The manager device represents the CA system that plays an essential role in a certification system by signing public keys (or certificates) (Figure 7). This device should be assumed to be trusted and controlled only by trusted persons (such as the network owner). All certificates signed by the device will be implicitly trusted. Currently, systems that manage a PKI require a high degree of security and are installed on an isolated machine. In this proposed system, the PKI is installed on the manager device that is hybrid, meaning it may be offline from the network when not in use to prevent the possibility of the private key being stolen in a possible network intrusion. For added security, the manager device can use Intel SGX to secure all the cryptographic assets in a Trusted Execution Environment (TEE). The manager begins by setting a CA using 256-bit Elliptic Curve Digital Signature Algorithm (ECDSA) and a 256-bit Elliptic Curve Integrated Encryption Scheme (ECIES) key pair to generate a shared key without the need for Diffie-Hellman exchange.

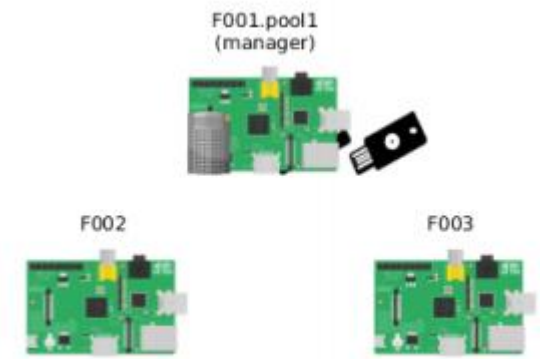

Figure 7: Manager Setup Phase

Device authentication

The authentication between a new device and the manager is essential for ensuring that it is added to the trusted device pool, Figure 8. To do this, the owner inserts the secure token into the target device, and then the new device is added to the pool (two cameras or one camera with another device).

As previously described, the secure token also has the manager's device public key. The new device starts by sending the Certificate Signing Request (CSR) of itself (which contains only the public key, not the private key, so the private key has not been compromised). When the new device sends the CSR to the manager, the latter will produce a signed x509 Certificate. Furthermore, it also sends the OTP to verify that the new device is in physical presence with the security token and is therefore the correct device to authenticate. All this information is sent encrypted with the shared key (ECDHE) generated for both parties (client and manager) to encrypt a message that the manager can only decrypt. After the authentication is successful, the manager device sends back the signed certificate to certify that the client is a secure device added to the trusted device pool. These certificates are used to establish trust between client devices and provide decentralized, secure communication between them without the manager device's intervention.

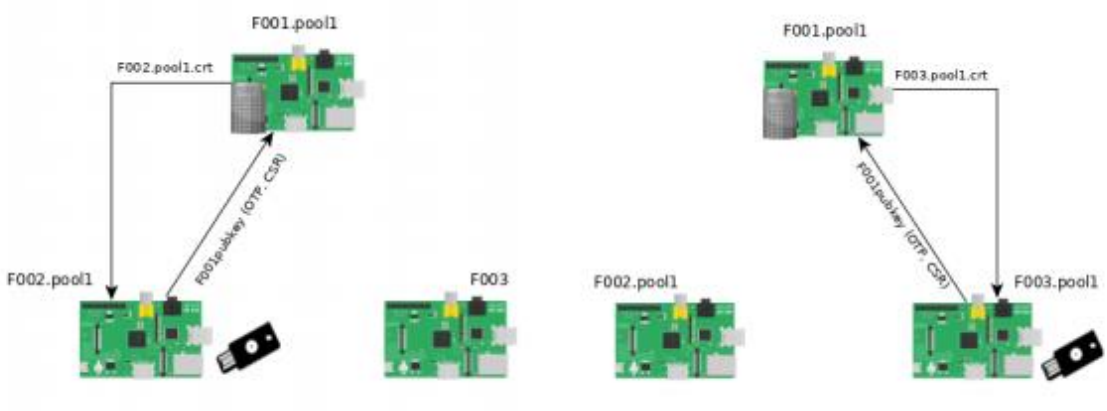

Figure 8: Device authentication

Decentralized Secure End-to-End Communications

After the discovery process, devices need to authenticate with each other. For mutual trust, both devices must exchange the manager's signed certificates. After verifying certificates' authenticity, a symmetric key is generated between both devices to establish secure communication. Symmetric key generation needs authentication so that the nodes know each other. ECDSA was chosen for signing and verification, and ECIES was chosen for encryption. Then, Diffie-Hellman Ephemeral (DHE) or Elliptic-curve Diffie-Hellman Ephemeral (ECDHE) is used for key exchange. Ephemeral mode is important because if the key pair is used for more than a few hours, it must be stored somewhere. After all, devices can be turned off. There is always some risk that a stored key pair may be compromised, although a wide variety of methods can be and are used to mitigate this issue. This mode avoids this type of attack by not storing key pairs and generating a new key pair every millisecond, thus ensuring Perfect Forward Secrecy. After establishing a shared secret using ECDHE, the devices can exchange data with symmetric encryption using the secure cipher AES256 to encrypt messages. When all devices are provisioned, the manager can be turned off, Figure 9, until a new device needs to be added to the pool or there is a change on the device pool, such as certificate revocation or renewal.

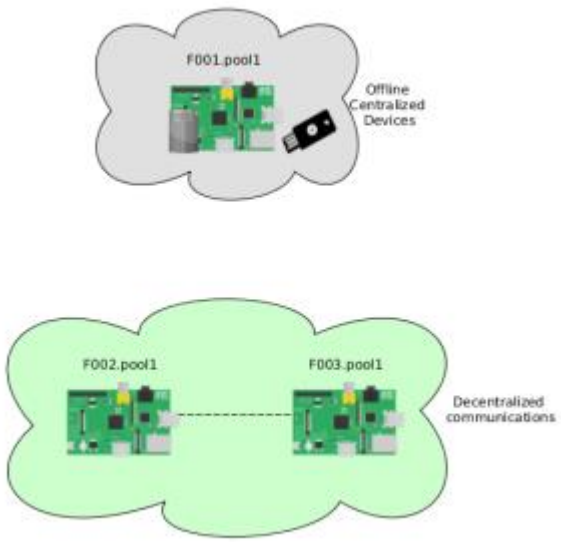

Figure 9: Decentralized Secure End-to-End Communications

Merge Two Trusted Devices Pools

An identity and authentication system must be flexible and highly scalable enough to handle billions of device infrastructures in multiple environments such as smart homes and smart cities in general. This system must support different environments, given the heterogeneity of applicability in IoT scenarios.

For greater scalability, there needs to be a useful way to integrate different device pools to make the system more practical as it would not be feasible to re-provision devices already provisioned with another manager so that devices from different pools can communicate with each other. To address this issue, the system replicates the traditional mechanisms of having multiple CAs supported by a client (bridge both CA). It is essential to ensure two points to deploy this in a real-world configuration: Use a secure token authentication scheme to enable enrollment and trust between different managers; and information dissemination on new pools among all new devices.
The authentication between managers (Figure 10 connection 1) uses the same mechanism to allow two pools to connect. After both managers perform mutual authentication, the next step is to spread the information across devices among different pools. To do this, the manager must send the signed and encrypted information to the devices. This allows any of the devices to read the information and verify the manager's signature. Figure 10 (connection 2) represents the agreement between managers and the corresponding spread of information from managers to their peers when they begin to trust each other and announce on the network that others should move to include these new trusted colleagues in their trusted network.

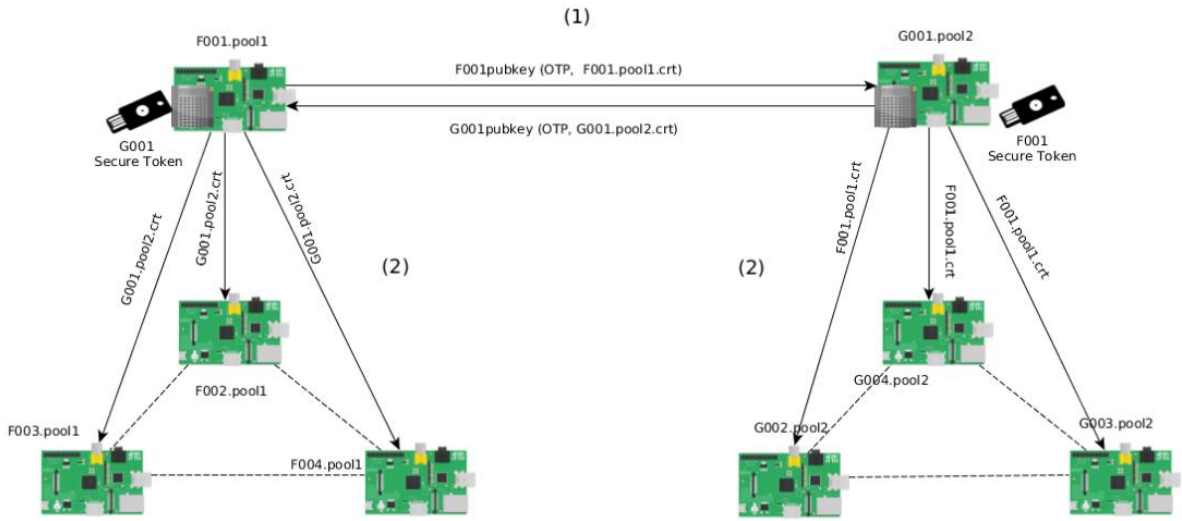

Figure 10: Merge Two Trusted Devices Pools

Privacy data controller

In PTASC, users can decide the data and traffic exchanged according to their preferences. Users have the option to block all traffic by default and to make exceptions for some specific domains. Therefore, users can block marketing/advertising sites and only communicate with the manufacturer's domains. Note that if users choose to block all communications from their devices to the Internet, some of the features may stop working, as some of these devices will not work in offline mode. Finally, users will have to choose between usability and privacy.
The middleware allows users to monitor incoming and outgoing connections to verify that the device is running an untrusted program and block or disable updates to specific resources, allowing them to block those connections "on the fly".
Along with this network traffic monitoring and, depending on the manufacturer's device firmware, users can store data offline on their home router for future reference. As users can connect with multiple routers (at home or work, for example), they can choose different permissions for their data depending on the device context, managing the data's life cycle.

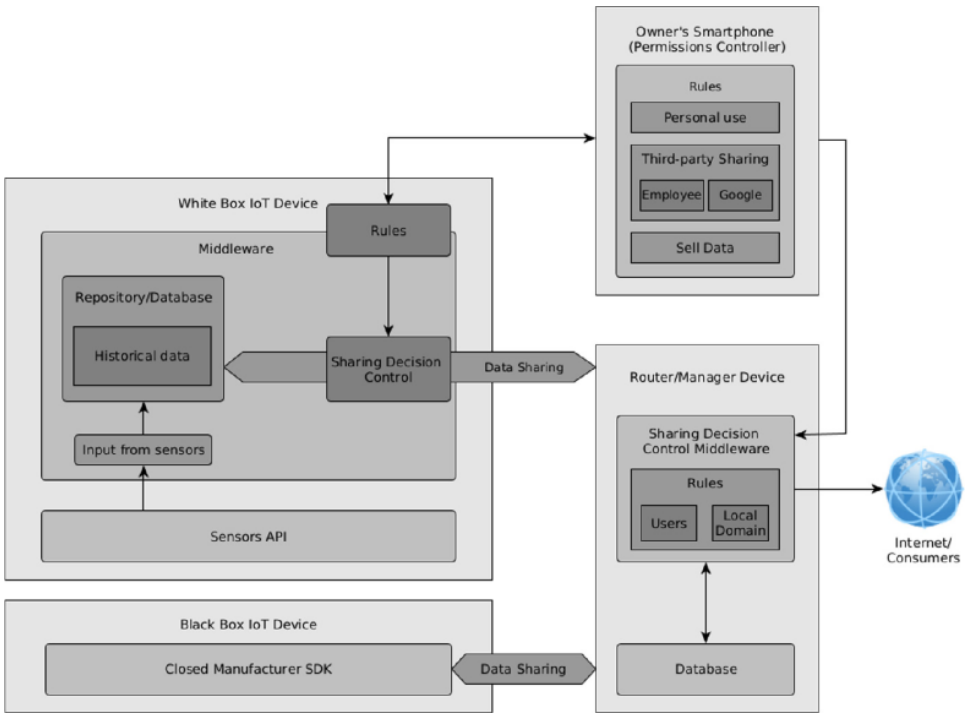

Figure 11: Privacy data controller

4.1.2 Main asset improvements since D3.13

From D3.13 PTASC has been improved mainly on the compliance with different IoT devices. The main motivation for the last work is the deployment in the WP5 since any device by default needs to connect directly with the manager without extra overhead the plug of the yubikey on the device must be sufficient to devices start to communicate end-to-end. Also we reviewed current certifications "IoT security certifications: Challenges and potential approaches" in an effort to define the requirements IoT devices must pass to be compliant and communicate end-to-end. For this the publication overviews four main questions: RQ1: What are the requirements that IoT environments expect from certifications? RQ2: What are the special requirements that different IoT application scenarios impose on certifications? RQ3: Are current certifications able to meet the requirements imposed by IoT? RQ4: Is ECSCF ready to integrate with the current certifications.

4.1.3 Research challenges addressed

Authentication and authorization of devices is a crucial aspect in the IoT ecosystem. PTASC addresses the current limitation of device pairing without using a PKI infrastructure or SSH keys. The main contribution of PTASC can be identified: Device Identity - A solution that relies on the combination of a security token (capable of generating OTP and storage of a PKI) with cryptographic algorithms to provide an identity to devices. Managers can authenticate the trusted devices in their pool, giving them an identity; Devices' Pools - After device provisioning, the system can provide a decentralized architecture where the trusted devices can communicate end-to-end between each other (if they are in the same or trusted pools). Different pools can be trusted between each other if both managers agree on that (for this reason, we consider that the scalability is better than other systems that need to authenticate all devices between each other).

We can map this general discussion about challenges to the ones defined in D3.11. Concretely, the asset addresses the following challenges:

- DP-06 by providing mechanisms for control how the information is disseminated and control the private data on the communication.
- IDP-05 by providing a mechanism for guaranteeing proper identity management of things for authentication end-to-end.

4.1.4 Demonstrations Example

From the development of D3.13 there is no update on the demonstration of the asset. For a full description of the demonstration refer to the original version on D3.13

4.1.5 Future Work

For future work, we plan to extend the integration of PTASC in the municipality of Porto and extent the current implementations with a revocation system.

4.2 ARGUS

Cloud storage allows users to remotely store their data, giving access anywhere and to anyone with an Internet connection. The accessibility, lack of local data maintenance and absence of local storage hardware are the main advantages of this type of storage. The adoption of this type of storage is being driven by its accessibility. However, one of the main barriers to its widespread adoption is the sovereignty issues originated by lack of trust in storing private and sensitive information in such a medium. Recent attacks to cloud-based storage show that current solutions do not provide adequate levels of security and subsequently fail to protect users' privacy. Usually, users rely solely on the security supplied by the storage providers, which in the presence of a security breach will ultimately lead to data leakage. We implemented a broker (ARGUS) that acts as a proxy to the existing public cloud infrastructures by performing all the necessary authentication, cryptography, and erasure coding. ARGUS uses erasure code to provide efficient redundancy (opposite to standard replication) while adding an extra layer to data protection in which data is broken into fragments, expanded, and encoded with redundant data pieces that are stored across a set of different storage providers (public or private). The key characteristics of ARGUS are confidentiality, integrity and availability of data stored In public cloud systems.

4.2.1 Overview

The main components of ARGUS are:

Privacy Broker:

We implement a broker that acts as a proxy to the public cloud infrastructures by performing all the necessary authentication, cryptography, and erasure coding. By doing so, it offloads the computational workloads from clients. Our approach ensures confidentiality, integrity, and availability of the data in the public cloud systems.

Dual Option File Encryption

In order to ensure the security and privacy of user data, the user can opt to encrypt everything locally and be responsible for the key management. In this way, the user does not need to rely on a third party. In addition, we have the option of delegating the encryption to a third party, which has the benefit of reducing the cost of execution on a limited device.

Confidentiality, Integrity and High-Availability

ARGUS maintains the integrity of the data as it stores an HMAC of all files. The confidentiality of the data is ensured as the data are encrypted, and the user can save their private key locally. ARGUS provides high availability through the redundancy that is assigned in the different cloud providers, that is, information is redundant on the three public clouds.

Intel SGX

The system uses Inte''s CPU SGX extensions to cipher user credentials (access tokens that give access to the use''s public cloud storage). This is an improvement over current implementations in systems that use the Google Drive API because the credentials are stored locally in the file system.

The general architecture of the ARGUS is represented in Figure 12 with all the components described in a detailed version.

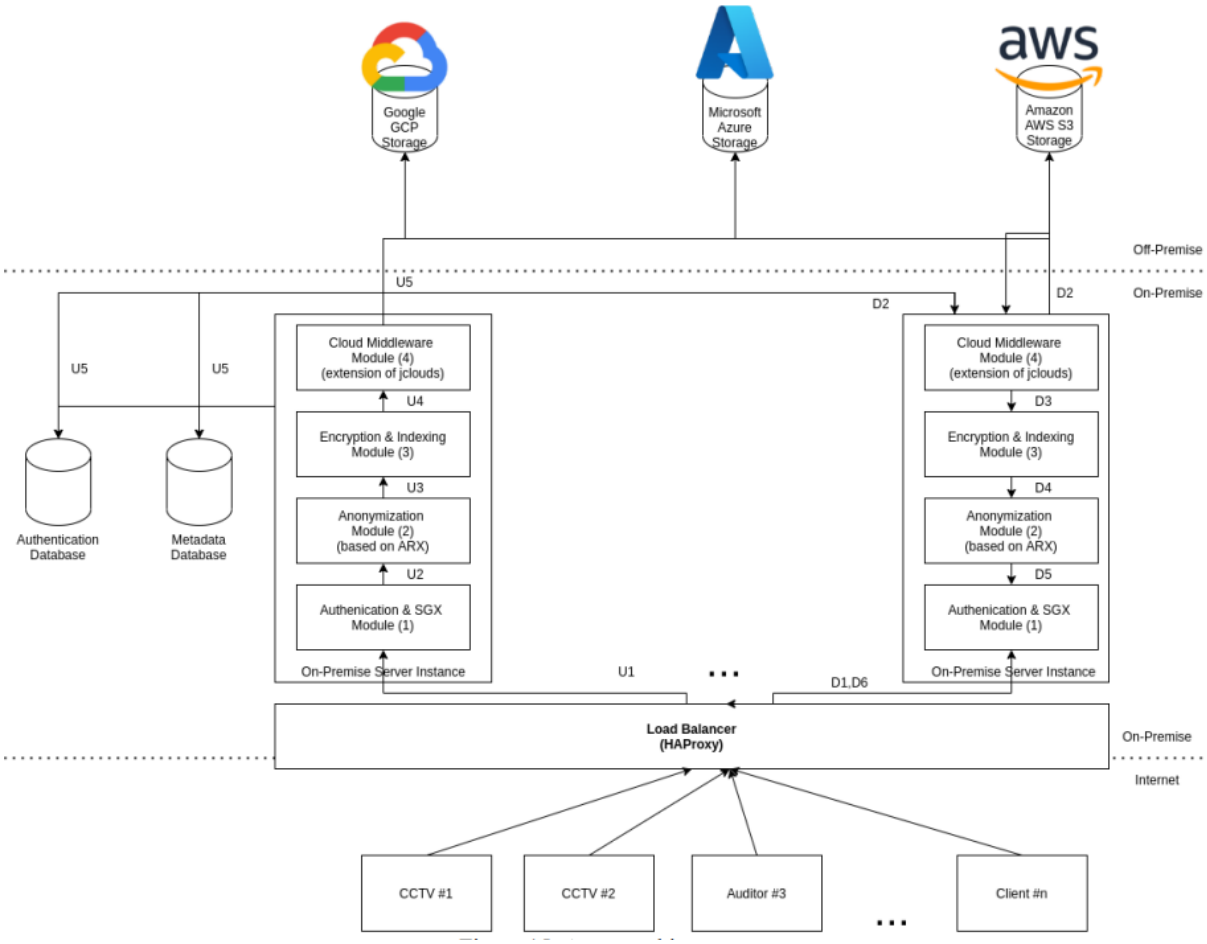

Figure 12 - Argus Architecture

4.2.2 Main asset improvements since D3.13

From D3.13 ARGUS has been improved mainly on the publication of "Towards a Modular On-Premise Approach for Data Sharin”" where we still required some minor implementation to automate some processing capabilities of ARGUS. This focus on a new caching mechanism based on ceph to locally be capable of storing the information in a distributed environment. Also we conducted a formal methods verification to Ceph, one of the most used open-source distributed storage platforms. Since the consensus algorithm is central to Ceph, it was formalized using the TLA+ language. After this formalization, properties of the algorithm were verified using model checkers. The specification was also evaluated in terms of efficiency by counting the number of generated states.

4.2.3 Research challenges addressed

Storage of information in the cloud privately is a topic with years of development, but it is crucial to use the public cloud, for this many solutions exist, such as implementing encryption on the client side, split over multiple files. But current solutions lack on the privacy when sharing information or managing private data. In this context, ARGUS in this is extended to include a Run-time Adjustable Privacy Schemes (RAPS) an adjustable privacy mechanism that enables us to tweak the anonymization, storage location, and persistence parameters, allowing them to have more control over the processing.

g that their data might suffer. This allows the system to detect anonymization patterns, using the parameters established on the RAPS, for possible privacy leakage or any other identifiable. This step is essential before sharing information with other parties, as it enhances anonymization and privacy. Also, we extended the work in the area of Secure sharing for Machine Learning using MPC.

We can map this general discussion about challenges to the ones defined in D3.11. Concretely, the asset addresses the following challenges:

- DP-07 by providing anonymization mechanism for control on how the information is stored and control the private data on the communication.
- DP-08 by introducing the possibility for users to store files locally or in a private cloud depending on their requirements.
- IDP-06 by ensuring that the encryption keys of the HTTPS protocol are manipulated in clear text only inside a trust zone (negotiating all the cryptographic material only inside the enclave)

4.2.4 Demonstrations Example

The demonstration of ARGUS in the smart campus is performed in D3.13. In this deliverable we demonstrate the results from ARGUS using two modules. In Test 1, we started by testing the performance overheads that SGX might cause. In the second and third tests, we followed a different approach, as we focused on a dedicated machine and removed the performance overhead of the communication between the on-premise server and the Android device, which allows us to compare the related work without having to include the overhead/latency of the client/server communications. Thus, in the second test, we compared with the related work, and finally, in the third test, we addressed the system's scalability.

Test 1—SGX Overhead
For this test, we defined three environments:

1. The on-premise server is used without SGX;
2. The SGX stores the private key and is only used on the public key operations of TLS;
3. Every cryptographic operation of TLS is handled on SGX, including a public key and symmetric key operations, such as the encryption and decryption of packets.

For this first test, we used a Xiaomi MI A2 lite (on-premise client) featuring Android 9, a Ubiquiti Uap-Ac-Pro Access Point (AP), and an Intel NUC6i7KYK for supporting our on-premise server. The smartphone was connected to the AP, using 802.11n AC with a standard 20 MHz channel width in the 5 GHz wireless frequency range. The network connectivity was a 1 GB upload and download internet. For these runs and each test, we used five repetitions of each scenario.
Table 3 shows that there is an increase in the performance overheads in the usage of SGX. However, using SGX to protect the HTTPS certificate private key has less performance overhead, providing usable upload and download performance. When SGX is responsible for all cryptographic operations over HTTPS, the performance overhead is too much to have a usable system. Because of this, the on-premise server implementation only has SGX responsible for public-key cryptography operations that require access to the certificate's private key. We also represent the standard deviation for the five repetitions of each sample.

	(1) On-premise Server without SGX		
File Size	100 MB	500 MB	1 GB
Upload MB/s	12.45 ± 0.43	11.90 ± 0.63	12.04 ± 0.84
Download MB/s	14.10 ± 2.76	13.66 ± 1.33	14.00 ± 1.44
	(2) On-premise Server SGX only PK		
File Size	100 MB	500 MB	1 GB
Upload MB/s	9.23 ± 0.33	10.90 ± 0.43	10.49 ± 0.14
Download MB/s	10.10 ± 1.05	12.45 ± 2.48	14.12 ± 1.85
	(3) On-premise Server SGX PK + SK		
File Size	100 MB	500 MB	1 GB
Upload MB/s	1.88 ± 0.02	1.54 ± 0.43	1.54 ± 0.33
Download MB/s	1.35 ± 0.47	1.37 ± 0.48	1.38 ± 0.08

Table 3: SGX overhead testing in MB/s

Test 2—Comparison with Other Cloud-of-Clouds

For this test, we used Google GCP storage as our cloud storage, writing in different buckets. For the setup, we used a network powered by a NOS Power Router 4.0 router with a 200 Mb/s internet down-link and a 150 Mb/s internet up-link. The machine used to host the services and run the tests used by the evaluation is a machine running in an Intel i7-8750H 6-core CPU with a 512 GB SSD connected to the router via a gigabit Ethernet cable.

Having an SGX mode selected, we started by comparing our solution with other state-of-the-art offloading technologies to see our performance overhead of encryption, erasure coding, and the cloud-of-clouds offload.

In our testing, we used three files with sizes of 100 MB, 500 MB, and 1 GB, uploading and downloading them ten times for each file size. To reduce the possibility of having noisy readings caused by network latency, we placed the client and the on-premise server on the same machine.

Contrarily to the previous test, we tested with the cloud-of-clouds and with local swift instances on the same machine, allowing us to test the scenario without the overhead associated with the public cloud provider. For this, we used four instances of the OpenStack Swift storage container deployed on the same machine as the client and the on-premise server that simulates the off-premise public cloud storage that our solution uses.

We tested our solution against Charon, configured without pre-fetching and compression, and against rClone. We compare our solutions' offload with the Charon, as it also uses multiple cloud providers for storing the information persistently. rClone tests should allow us to see the maximum throughput that the cloud backends allow, enabling us to verify how much offload performance is left on the table due to our overhead.

For sake of completeness, we ran the same tests using Google GCP storage as our cloud storage backends to see how the network latency affected our results. Table 4 shows the results of our testing in seconds with an average of ten runs for each scenario. We also have the standard deviations to understand the min and max values of each run.

		DOWNLOAD	
CLOUD	**100 MB**	**500 MB**	**1 GB**
Our Solution (s)	13.360 ± 1.427	45.718 ± 18.871	112.048 ± 2.269
Charon (s)	14.884 ± 0.752	77.998 ± 3.245	209.172 ± 7.658
Rclone (s)	18.486 ± 1.794	58.727 ± 5.022	84.364 ± 33.673
		UPLOAD	
CLOUD	**100 MB**	**500 MB**	**1 GB**
Our Solution (s)	18.657 ± 2.335	56.079 ± 0.451	79.031 ± 3.244
Charon (s)	8.904 ± 0.117	102.610 ± 0.781	152.324 ± 2.645
Rclone (s)	9.623 ± 0.524	52.657 ± 8.541	84.473 ± 5.368
		DOWNLOAD	
SWIFT	**100 MB**	**500 MB**	**1 GB**
Our Solution (s)	2.668 ± 0.236	9.870 ± 0.416	25.237 ± 3.365
Charon (s)	2.222 ± 0.117	10.333 ± 0.263	22.924 ± 0.293
Rclone (s)	1.230 ± 0.037	10.333 ± 0.437	12.838 ± 0.580
		UPLOAD	
SWIFT	**100 MB**	**500 MB**	**1 GB**
Our Solution (s)	5.537 ± 0.770	15.601 ± 2.162	34.354 ± 5.615
Charon (s)	8.904 ± 0.438	31.692 ± 0.800	65.567 ± 1.906
Rclone (s)	0.933 ± 0.150	6.126 ± 0.601	12.079 ± 0.426

Table 4. Off-premise data offload performance when compared with Charon and rClone, in seconds.

Regarding the tests with the cloud, we can achieve a better performance than Charon in the download process. However, in comparison with rClone, for 1 GB size files, we have a weaker performance. rClone only communicates with the cloud provider, so it has better performance. However, Charon has to wait for the cloud providers so that it can explain its results.
On the other hand, in the upload process, we only achieve the best results on the 1 GB file, as Charon seems optimized to small files because the difference between 100 MB and 500 MB is greater than 90 seconds.
Regarding the Swift tests, the download and upload times decrease for all solutions, as they are performed locally and there is no connection overhead with the cloud provider. rClone performs better in all the results, while our solution can achieve similar results as Charon, except for the 1 GB upload, where the performance of our solution is much better.

Swift tests also allow us to test the dependency of the public cloud providers, as rClone has better results than the other solutions. Compared with Charon, our solution has an overhead decrease of almost 50%.

Test-3—Scalability

In this test, we used the same setup provided in Test 2. We used a network powered by a NOS Power Router 4.0 router with a 200 Mb/s internet down-link and a 150 Mb/s internet up-link, and the machine used to host the services and run the tests used by the evaluation is a machine running in an Intel i7-8750H 6-core CPU with a 512 GB SSD connected to the router via a gigabit Ethernet cable.
In the architecture section, we stated that our architecture is built to be scalable. We achieved this by allowing multiple machines to run the same on-premise server code behind a load balancer. The load balancer allowed us to split the load equally among all the available machines running the on-premise server. In our implementation, we used HAProxy [46] as our load balancing technology with a round-robin configuration, meaning that the number of requests are divided equally among all the machines running the on-premise server. To verify whether our solution is indeed scalable, we devised a test where we have 1, 5, 10, and 20 concurrent users uploading files of 100 MB, 500 MB, and 1 GB, running ten times for each number of concurrent users and file size configuration. To simulate a more realistic use case, we used Google GCP as our storage backends. In our testing, we used three on-premise server instances running locally on the same machine where the requests were made. To help us compare our solution scalability with other state-of-the-art cloud offload technologies, we compared our solution with rClone. rClone is a standalone solution, meaning that it should have no scalability issues, as it connects directly with the cloud storage without any additional processing made locally. rClone

properties will help us to understand the maximum capacity that our testing environment has and to have a baseline on which to compare our solution. We did not use Charon on this test since it uses a caching service, making it only upload a file at a time instead of continuously uploading all the files, making it impractical for this performance comparison. Figure 12 and Figure 13 show a line graph of the meantime that took to, respectively, upload and download 100 MB, 500 MB and 1 GB files with 1, 5, 10 and 20 concurrent users.

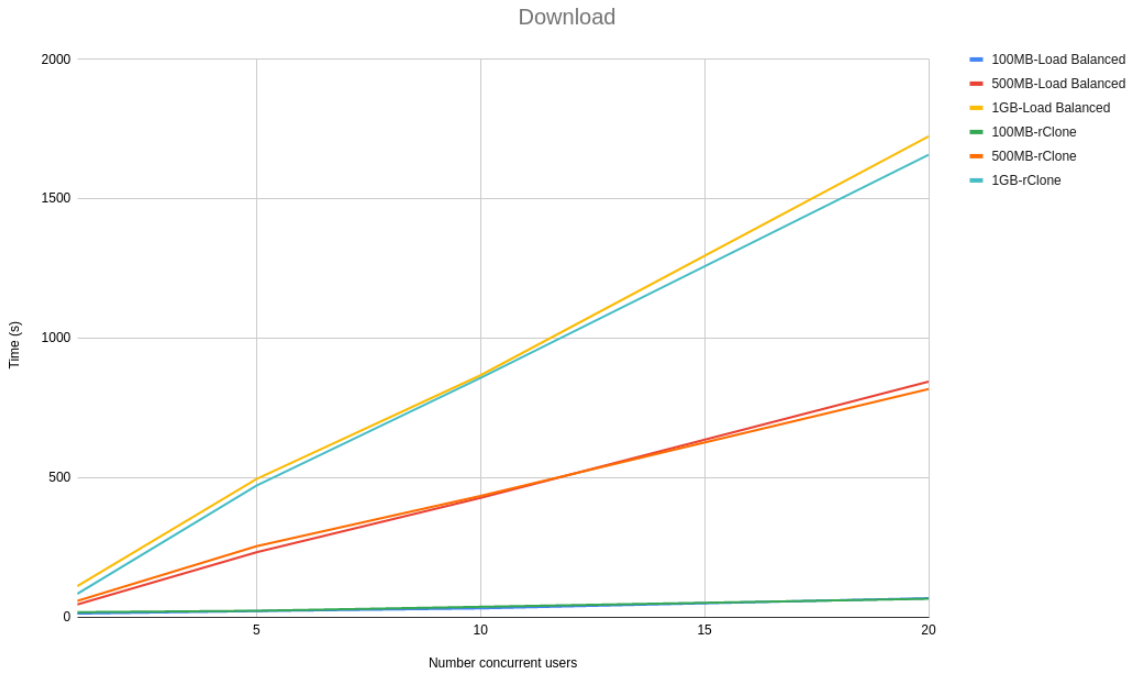

Figure 13: Download scalability tests

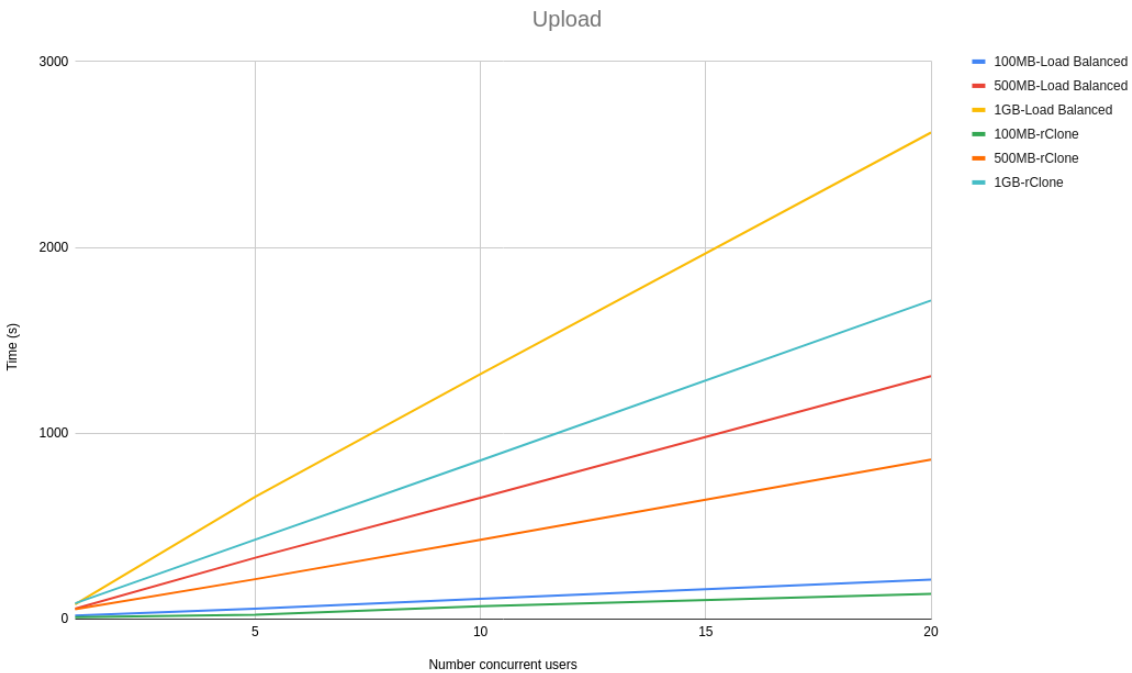

Figure 14: Upload scalability tests

From the line graphs, it is possible to see that the difference between our solution and *rClone* grows linearly with the size of the file and not with the number of users, meaning that our solution scalability is independent of the number of users but dependent on the file sizes. However, the growth is linear with the file size rather than being exponential, proving that our solution architecture is scalable

4.2.5 Future Work

Argus is currently being deployed in the WP5 and it is expected to be demonstrated in the final deliverable. As future work we expect to introduce a new mechanism to improve the upload speed to the cloud and tags to allow to identify personal information on the public cloud.

4.3 Self-Sovereign Privacy-Preserving-IdM (SS-PP-IdM)

This asset leverages the OLYMPUS virtual identity provider, which is comprised of multiple individual IdPs, to manage users' identities and authentication. It relies on distributed p-ABCs to offer privacy-preserving (minimal disclosure and unlinkability) and authentication (presentation of attributes) linked to eIDAS. Moreover, the asset proposes a trust framework based on Blockchain to complement the usage of credentials.

4.3.1 Overview

The SS-PP-IdM is an evolution of the OLYMPUS oblivious distributed Identity Provider, enhanced with Distributed Ledger Technologies for verifiable sharing of identity data and key materials. Figure 15 shows an overview of the asset, and the generic actors and interactions in an application of the asset.

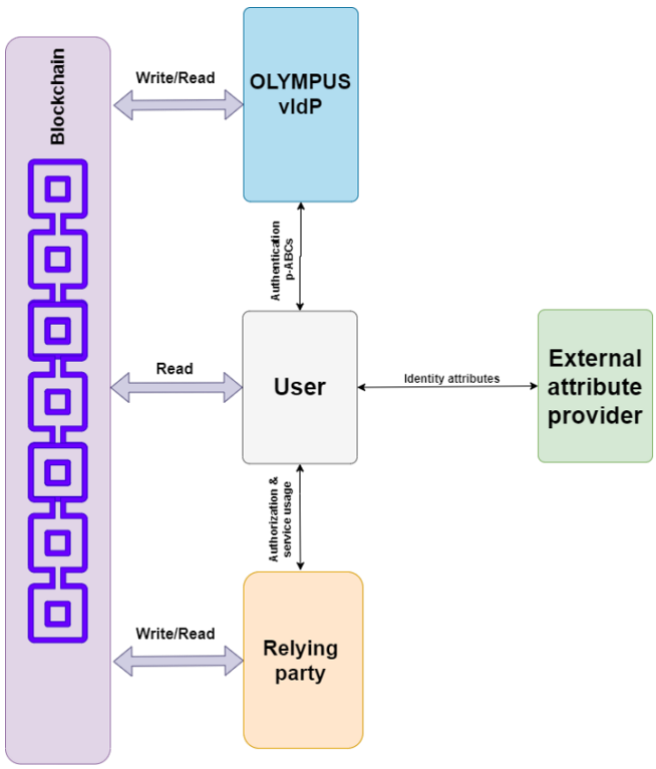

Figure 15 Generic overview of SS-PP-IdM asset

The main components of the SS-PP-IdM are:

- **OLYMPUS vIdP**: OLYMPUS virtual identity provider comprised of multiple individual IdPs. Leverages distributed p-ABCs to offer privacy-preserving (minimal disclosure and unlinkability) authentication (presentation of attributes).
- **Blockchain platform**: Stores public information about the OLYMPUS infrastructure (endpoints, vIdP and cryptographic parameters) in a trusted way. This information can be consulted at any time. In addition, the blockchain platform introduces a set of Smart Contracts

that enables a transparent way of operation between the Smart-Campus platform, OLYMPUS vIdP and the mobile application.

- **IdM client:** It performs client-side operations, like interaction with vIdPs, storage and management of credentials, and generation of zero-knowledge presentations**.** In this specific scenario, the client-side library has been used to develop a user mobile application.
- **Verifier-side elements:** The functionality needed by a relying party for verifying the presentations generated through the SS-PP-IdM.

In this case, we will rely on attributes strongly linked to user's identities through the use of eIDAS (and the eIDAS browser asset also included and developed in this project). Because of that, two extra components are involved in the identity management platform:

- **Keyrock**: Used as a bridge to eIDAS (i.e., handles SAML communication flow with eIDAS node to obtain certified attributes).
- **eIDAS node**: It handles authentication (of a natural person in the first pilot) with an electronic certificate or national eID following the eIDAS specification.

In any usage of our identity management solution there are three key processes:

- **User Enrolment (a)**: A user account is created and populated with the attributes that will comprise her identity in the scenario. This will only happen if the user correctly proves possession of said attributes, with assertions being verified by the vIdP. Note that the sub-process of managing (adding, removing…) attributes to the user account can also be executed at other times.
- **Credential Issuance (b):** After they have enrolled, users can authenticate against the vIdP to obtain a credential. The resulting p-ABCs will hold all the user attributes and can be (securely) stored for later use.
- **Privacy-preserving service access (c):** Accessing services can require that a user fulfils some attribute-based policy. With the credentials generated in process *b* the user can prove the necessary assertions while ensuring minimal disclosure and unlikability between requests.

The integration into the Smart Campus scenario means that there will be a close relationship between the identity management system and the smart platform's authorization framework. The generic scenario of the Smart Campus, focusing on the authentication/authorization interactions, can be found in Figure 16. The main components affecting the scenario of application of the SS-PP-IdM are:

- **Smart-campus platform**: Smart campus platform that offers services for Murcia city
  - **Services**: Public transport, parking availability…
  - **PEP**: Controls access to the services, checking that the request includes a valid capability token (i.e., the request is authorized).
  - **Capability Manager**: Generates capability tokens that bestow authorization to use specific services. Relies on the PDP for the decision (using XACML).
  - **PDP**: Checks if an authorization request should be conceded, using the OLYMPUS verification library to validate the presentation token against the policy.

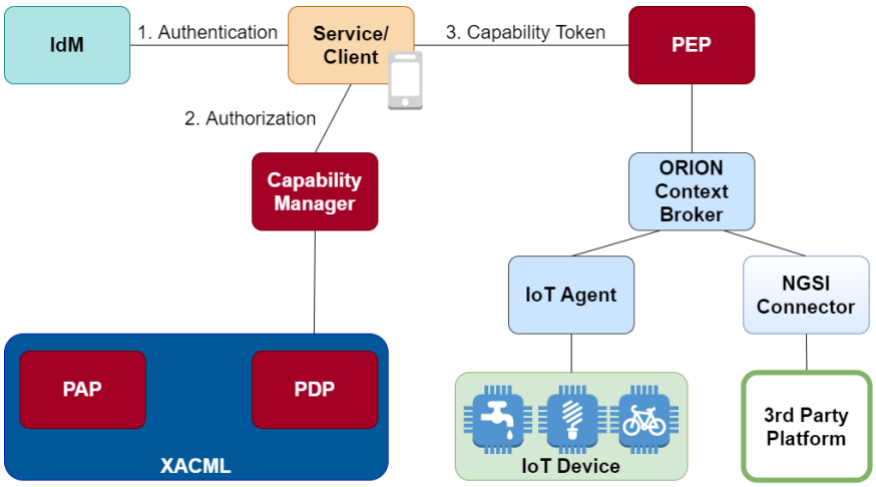

Figure 16. Smart Campus framework (Authentication and Authorization)

Smart Contract models

As mentioned, integration with the blockchain platform allows the use of smart contracts and a verifiable data registry. These contracts define a series of data models that represent different elements of the interaction. The most important models envisaged are:

- **Partial Identity Provider model**: A (partial) identity provider is characterized by several properties.
  - Status, active or inactive.
  - Last modification, that indicates when it was modified (RFC 3339 format).
  - Spawn date, that indicates when it was first deployed (RFC 3339 format).
  - Public key, associated to the IdP's p-ABC scheme. Includes the key, its id and type (determines how it was serialized, which EC was used...).
  - DID document that contains an identifier, context, and a service. The service establishes what type of service its endpoint.
- **Virtual Identity Provider model**: A vIdP is composed of 1 or more partial IdPs. In the same way as the partial IdP, this model collects the properties of status, spawn date and last modification. It also adds the following:
  - VIDP Document, is an extended version of the previous DID Document. In this case, instead of a single service, it includes a list of services.
  - Aggregated Public Key automatically computed using the public keys from the partial IdPs.
  - Schema ID contains the Id (DID) of the schema (pp and credSchema) associated to this vIdP.
  - Domain in which the vIdP is working. It can be used to separate vIdPs for different uses of the chain.
  - Alias of the vIdP for easy naming
- **Schema model**: Model for a scheme needed by an OLYMPUS vIdP, including public parameters and W3C's VC credential schema. Defined by:
  - DID Document, to assign concrete DID.
  - Public Parameters is a serialization of the public parameters for an IdP (i.e., p-ABC schema info, attribute definitions...). It is expected to be serialized in a JSON with two fields: type (e.g., OLYMPUS-pp) and data (e.g., public parameters serialized with OLYMPUS custom format).
  - Credential Schema is serialization of the credential schema as defined in W3C's Verifiable Credentials specification (JSON-LD).
- **External Service model**: This model refers to an external service willing to use the chain to offer itself and OLYMPUS for authentication of users. Defined by:

- DID Document, a DID identifier and service endpoint.
- Name refers to the service name.
- Status active o inactive.
- Domain in which the service is being offered.
- Policy requested for the usage of this service. It is expected to follow the serialized JSON format used in OLYMPUS.
- Spawn date that indicates when it was first deployed (RFC 3339 format).
- Last modification, that indicates when it was modified (RFC 3339 format).

Operating with the models described above, there are several smart contracts for both data query (*getschema*, *getvirtualidp*…) and ledger insertion (*addservice*, *addschema*...). Among them, the ***addpartialidp*** contract stands out, in charge of adding new partials IdP to the ledger. Moreover, this contract triggers internally the creation or update of the associated Virtual IdP, automatically calculating the aggregated public key associated to guarantee that it is correct. Therefore, whenever a new partial IdP is added, we can automatically and transparently update or create the related vIdP with its aggregated public key.

4.3.2 Main asset improvements since D3.13

The integration of features for inspection and revocation has been achieved in a theoretical way and a viable solution has been achieved that allows to provide these features to the cryptography applied in the solution. The implementation of SS-PP-IdM applications and functionalities has also been consolidated by improving serialization, cryptographic operations etc. This has resulted in better integration with standards including full integration with XACML including the design of an operational model for token verification. Integration with the W3C Verifiable credentials standard has also been improved by revising the data models and serialization methods.
The integration with DLT has been improved and finalized (though it remains extensible), with the definition of models that represent the necessary information for an environment with SS-PP-IdM (IdP, vIdP, keys…) and even other useful data like services/relying parties available, their policies… Additional work has been dedicated to smart contracts, with improvements on their code, the addition of useful queries (e.g., multiple ways for discovery of vIdPs, services…), and their integration with the other components. What is more, the "registration" contracts have been improved. The most relevant one in that respect is the smart contract for adding partial IdPs, which has been upgraded with logic for automatically adding/updating registries for vIdP, including the computation of the aggregated verification key which serves as a root for trust in the p-ABCs generated by the vIdP.

4.3.3 Research challenges addressed

Authentication and authorization of users is a crucial aspect in all online interactions, and an especially complex process in a large, heterogeneous, and dynamic environment like the smart campus scenario for this demonstrator. Protecting the smart platform's services and data is not a simple task, and further challenges arise when privacy and security are pivotal objectives.
In that respect, being able to perform access control by authenticating users while ensuring privacy principles are respected is a key challenge. Attribute-based access control has been proposed for authorization in smart platforms, and more specifically, solutions like XACML framework [OASIS13] offer capabilities for fine-grained access control. However, these solutions have been applied with traditional federated identity management systems like OAuth [Hardt12], which have glaring issues in terms of privacy and security, such as user tracking or breaches of minimal disclosure. The SS-PP-IdM asset is based on distributed p-ABCs used in the OLYMPUS project [MBGFSMSPS20], allowing users to control which information is revealed when interacting with the capability-based access control and avoiding the IdP as a single point of failure. What is more, it tackles the lack of homogeneity in representing p-ABCs, one of the issues that jeopardized adoption of previous p-ABC systems, by integrating with the emerging W3C Verifiable Credential standard for serialization [GMBS21].
However, the original OLYMPUS IdM, like other similar proposals about privacy-preserving identity management [BDMCBS20], does not address the challenge of establishing a complete ecosystem where the different components can interact and ensure trust. The SS-PP-IdM asset defines and implements a complete environment with the addition of blockchain for auditing, traceability, and trusted public

information (i.e., public keys, parameters, and support as verifiable data registry for W3C Verifiable Credential's elements). For instance, the public keys of the partial IdPs will be available and, what is more, the aggregated public key is also automatically added to the DLT's registries through a Smart Contract, ensuring its correctness. Thanks to this, users can safely discover and interact with identity and service providers [MGBS21].
Lastly, and also related to the challenge of security and trust in authentication, the asset addresses the challenge of establishing strong links to physical identities through the integration of the eIDAS [Dumortier17] system into the solution.

We can map this general discussion about challenges to the ones defined in D3.11. Concretely, the asset addresses the following challenges:

- *IDP-02: Unnecessary over-identification and information disclosure due to a lack of awareness and usability drawbacks*. The pp-IdM used in the scenario enables minimal disclosure through the application of p-ABCs. Also, it takes usability into account, offering simple authentication mechanisms to users easing their interaction with access policies (more specifics about the work on this matter within the SS-PP-IdM context can be seen in D3.16 and D3.17, from task 3.6).
- *IDP-03: User's privacy-preservation of transactions in distributed and immutable systems (e.g., blockchains).* This asset tackles the challenge by relying on zero-knowledge proofs and smart contracts (based on blockchain) that allow anonymity in interactions without storing sensitive data in the ledger.
- *DP-04: Honest but curious Service Providers and Identity Providers that might be tracking users, IdP could also become a potential point of failure (identity/data leakage in the IdP).* Our oblivious distributed pp-IdM avoids tracking by IdPs and the p-ABC scheme used provides unlinkability and minimal disclosure, avoiding tracking through the authentication mechanism by service providers (unless users consciously decide to reveal information that identifies them).

4.3.4 Demonstrations Example

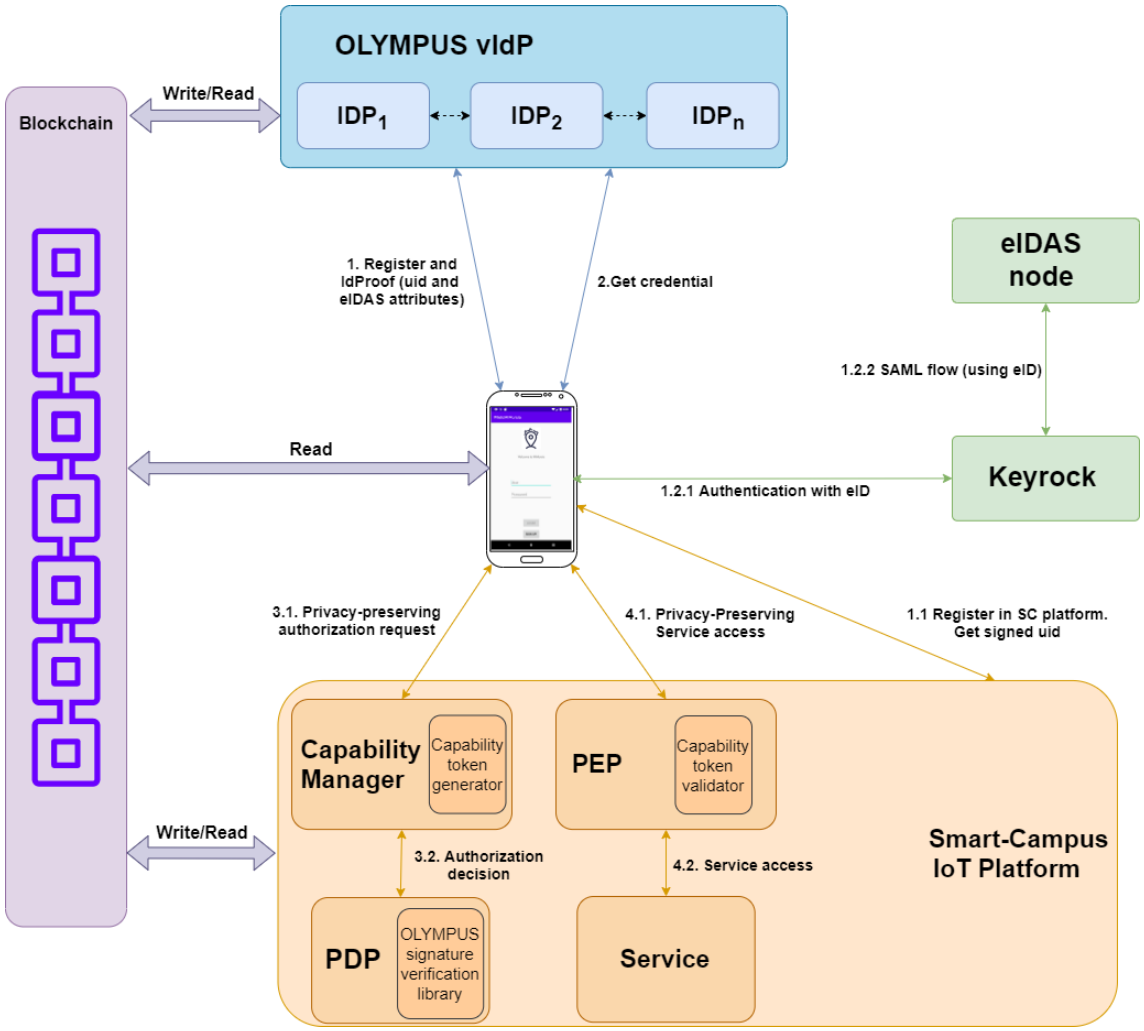

Figure 17: Instantiation of SS-PP-IdM in Smart Campus platform

The asset has been **validated** by using it to interact with a smart platform where the authentication and authorization processes completely rely on the SS-PP-IdM infrastructure, including the DLT as a public and verifiable data registry, not only for identity information but also for services and other information regarding the Smart Campus.

Figure 17 shows the architecture of the demonstrator, as well as an example flow involving the instantiation of the processes described in the previous section to this scenario. As we can see, the instantiation involves multiple steps specific to this use case, where the user relies on her mobile application to perform the different operations.

During **user enrolment**, the application redirects to the Keyrock platform so the user can obtain a set of attributes certified by an eIDAS node (using her eID). Also, the smart platform generates a user identifier (uID), so the user can be linked to an account in the platform. The application presents both assertions to the vIdP (eID Attributes, uID), and a successful verification will result in her OLYMPUS account having the information needed to generate credentials.

The **credential issuance** process, conversely, is equivalent to the one in any other application scenario. The user inputs her OLYMPUS account username and password to perform a login operation. If authentication is successful, the distributed issuance process will result in the user having a credential with her attributes certified by the IdM, which can be securely stored.

In this case, the **privacy-preserving service access** involves interaction with the smart platform. Figure 18 shows the services that appear on the application's homepage. It is conceptually divided into two

sub-processes, authorization, and actual service access. During the former, the user must prove that she fulfils some attribute-based policy (e.g., the one shown in Figure 19) to obtain the platform's authorization (in the form of a capability token issued by the capability manager). Later, the PEP checks that the user request includes a valid capability token to allow (or not if the verification fails) service usage.

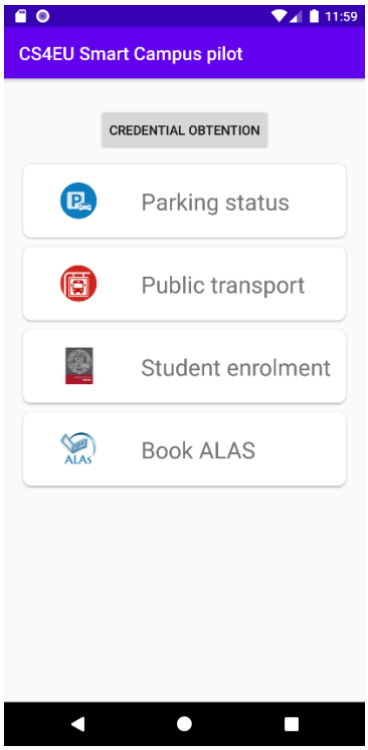

Figure 18: Application Homepage of the Smart Campus

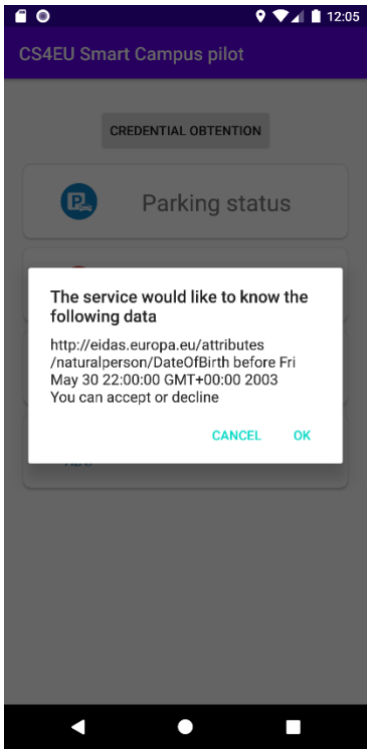

Figure 19: Using some attribute-based policy

Throughout the process, the involved actors will use the ledger to store and retrieve trusted public information. For example, vIdPs will register public information (endpoints, public keys…) in the ledger. The user application (and the verifier installed in the smart platform) can retrieve that information to perform the necessary configurations and ensure that the entity it is interacting with is trustworthy.

For this demonstration, we are deploying the necessary components (OLYMPUS partial IdPs, Keyrock, Smart platform's authorization components) in an instance (Ubuntu server 18.04, 8 GB RAM, 4 cores, 2 GHz) deployed in an OpenStack platform within the University of Murcia. The ledger deployment (Hyperledger Fabric v2) also relies on OpenStack virtualization. The Fabric infrastructure consists of 2 organizations with two peers each, a certification authority and an Orderer node. Every Hyperledger machine is running Docker v19, and Docker composes v6.14. Also, we implemented a user application in Android, controlling processes like enrolment, authentication and interaction with the smart platform. For the tests, we run various device, both emulated and physical (including a Poco X3 NFC for time measurements).

We have performed tests for the different needed functionalities with the current implementation, including (but not limited to): initialization with public information in the ledger, enrolment with Spanish eID, credential issuance, and authorization (and service usage) process. We generated Verifiable Credentials and Verifiable Presentations during a test execution of the application for

accessing a service. We have some preliminary performance evaluation, through execution time measurements (using a Poco X3 NFC device, Octa-core 2.3 GHz and 6 GB RAM) for different processes. For a more detailed benchmarking evaluation, we refer to the future evaluation of the asset within the project's pilots, specifically in T5.7 with a demonstration in a Smart City environment.

Figure 20 shows the results of the preliminary evaluations. Note that these measurements do not take into account user input (e.g., reading and deciding to accept the policy), but only computational procedures. It is also worth noting that users' identity credentials contained eleven attributes in these experiments, as this number impacts the cost of cryptographic operations.

The first graph within Figure 20 shows the results of the 3 main processes described in the previous section, though in this case service access (c) is divided in two subprocesses: getting authorization (obtaining a capability token) and accessing the service through the PEP (i.e., PEP verifying the capability token allows the operation and enabling service usage. This divide is specially interesting in this scenario, as capability tokens can be reused during their lifetime to access a service through the PEP directly. It is clear that the authorization process is the most expensive. This is a natural result, authorization is a complex operation, and it involves the bulkiest cryptographic procedures (zero-knowledge proof generation and verification). In any case, the execution times are not prohibitive, and the most expensive processes can be skipped in many of the user interactions.

The second graph of figure Figure 21 gives a detailed breakdown of the operations executed during the authorization process. It starts when the user accepts the policy and a Verifiable Presentation conforming to it is generated (GeneratePT). This operation is costly as the cryptographic method behind it is also one of the most expensive within the p-ABC scheme, and it is being executed on the (more constrained) mobile phone. Next, the token is sent to the capability manager, which controls the authorization flow in the smart platform. The capability manager contacts the PDP for a decision on the access request (following the XACML flows and models), the validity of the token is verified (SignVerify), and the capability token is generated (GenerateCT) and sent back to the user application.

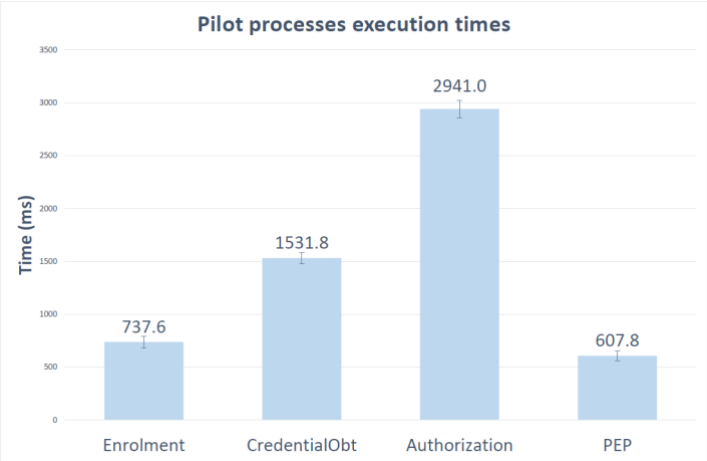

Figure 20 – SS-PP-IdM asset Performance evaluation

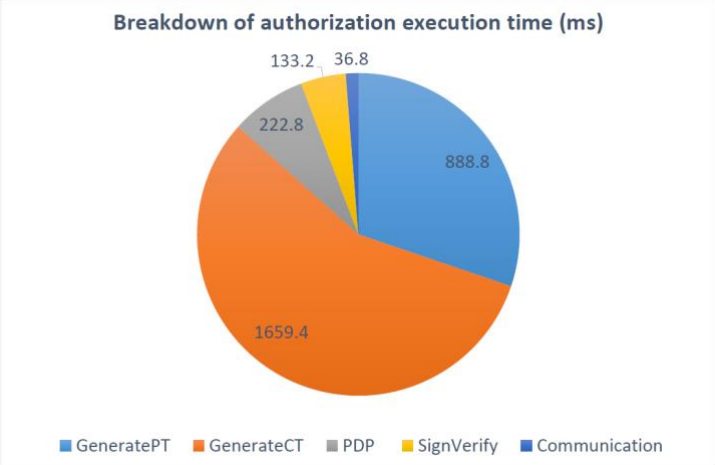

Figure 21: Authentication and authorization processes time measurements

4.3.5 Future Work

For future work that involves this asset, we plan to do a (feasibility-test) practical implementation of the inspection and revocation. Outside the project's lifetime, we will also explore the application of our p-ABC approach for identity management of constrained IoT devices. This might be accompanied by research on solutions for whole system integrations where identity management schemes coexist with other solutions to accommodate the different ranges of devices in terms of power (e.g., delegation for the most constrained, full capabilities for non IoT devices involved, reduced functionality for IoT devices "in the middle" etc). Lastly, we will explore architectures (and components that realize them) to provide pp-IdM applicable in distributed zero-trust scenarios, with dynamic participation, monitoring, dynamic trust, tackling the relationship between privacy and trust.

4.4 Password-less authentication

4.4.1 Overview

The password-less authentication asset is based on the FIDO standards[8]. It provides a device-centric authentication that implements a) a challenge-response scheme in which the user is authenticated locally (i.e., on the device that it is deployed to access the service) using alternative authentication methods, such as PIN, USB keys, and biometrics and b) public key cryptography to authenticate the device in the service. During the FIDO authentication, when a user (in our case a student) is authenticated in its device (for instance, using a USB key), it unlocks its private key, which subsequently is deployed to sign the challenge and the service deploys the user's public key, to decode the challenge. An overview of the FIDO authentication process is depicted in Figure 22

[8] https://fidoalliance.org/

Figure 22: FIDO Authentication Concept

The password-less authentication asset aims to replace or enhance the traditional username-password paradigm offering a more user-friendly and secure authentication procedure. Via this asset the students will deploy their smartphones, which are widely used in everyday life as well as in the campus premises, to be authenticated on numerous services. For this purpose, the FIDO protocol is used, which will be responsible for the registration, authentication, and de-registration of the students.

A) Registration

The registration process allows the server to verify the authenticity of the Authenticator and register it along with the user's account. The authenticator is represented by the device that the user will deploy to authenticate in the smart-campus platform, which in our case is the smartphone. Once an authenticator has been validated, the FIDO server can assign a unique identifier number (aaid) to the authenticator that can be used in future communication between the two parties. Specific policies are utilized in the FIDO server for the acceptance or not of the authenticator (for instance, some smartphones that have weak facial recognition can be omitted). The registration steps are described below:

1. The client application initiates the registration process.
2. The server sends a registration request.
3. The user enrolls in the client application and the key pairs (private and public) are created.
4. A registration response is sent to the server.
5. The server validates the response

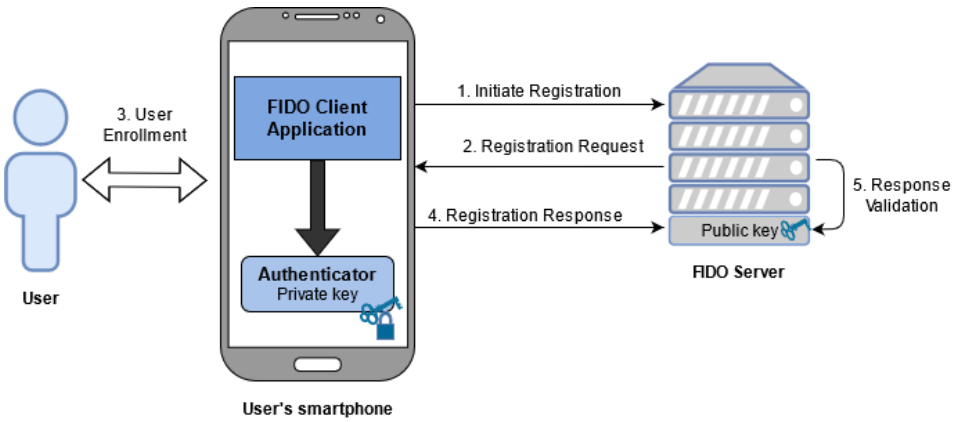

Figure 23: Password-less authentication. Registration process

B) User Authentication

The user authentication process is based on a cryptographic challenge-response scheme in which the user is prompted by the server to be verified by the authenticator that was used in the registration process.

The authentication is initiated by the client application, where the user chooses its preferred authentication method (i.e., fingerprint, PIN, pattern). The steps of the authentication process are:

1. The client application initiates the authentication process.
2. The server sends an authentication request along with a challenge.
3. The user is authenticated to the device using its preferred authentication method to unlock its private key.
4. The client application sends to the server the challenge signed by the user's private key.
5. The server validates the challenge using the user's public key.

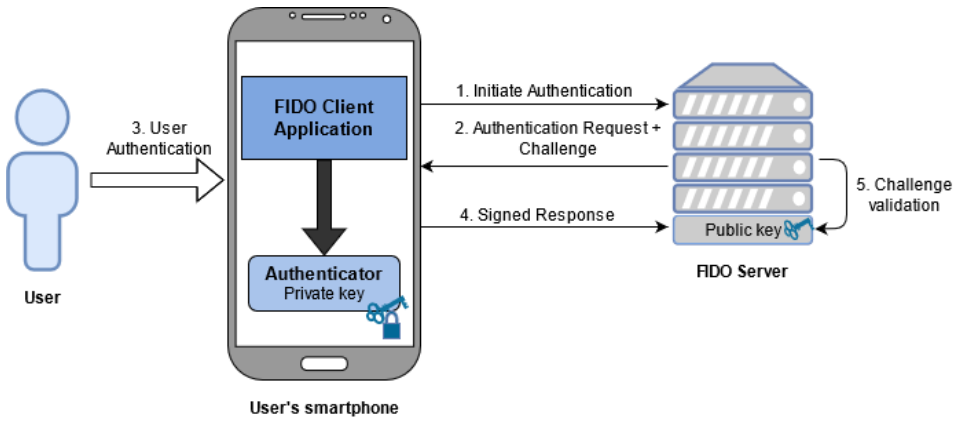

Figure 24: Password-less Authentication. Authentication process

C) De-registration

De-registration is required when the user account is removed from the server. The server can trigger the Deregistration by asking the Authenticator to delete the associated FIDO Credentials that are bound to the user's account. The de-registration steps are:

1. The client application initiates the de-registration process.
2. The server sends a de-registration request.
3. The server deletes the user's account and the public key from the DB.
4. The client application deletes the information from the device.

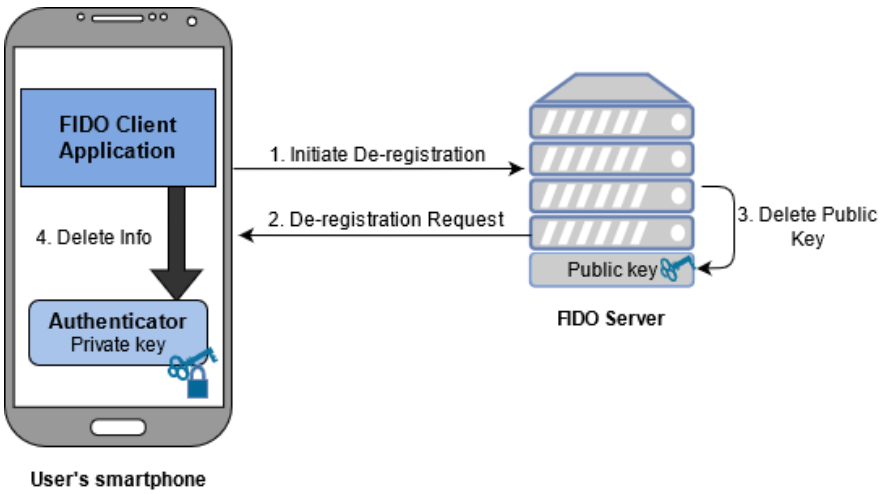

Figure 25: De-registration Process

In demonstration scenarios of CyberSec4Europe, the password-less authentication asset will provide an additional level of security in the authentication procedure, which will be based on the criticality of the service that the user wants to access.

4.4.2 Main asset improvements since D3.13

Since D3.13 the password-less authentication asset has been improved to support Single Sign On (SSO) capabilities. This is accomplished via the integration of the password-less asset with the OpenID connect protocol. Particularly for the implementation of the OpenID connect protocol, we deployed the Keycloak Identity Provider (IdP), which through a custom library is connected with the FIDO server that is shown in section 4.4.1. To this end, the user is authenticated on the Keycloak IdP using its smartphone via the FIDO client application. This improvement of the password-less asset can offer SSO capabilities, namely, a user can be authenticated on multiple services that implement the IdP using its smartphone replacing the traditional username-password paradigm that most SSO applications support.

4.4.3 Research challenges addressed

The authentication process of most web applications relies on the password paradigm. It is evident that a password can be considered secure when it contains 15 characters or more, is complex (is comprised of alphanumeric characters, symbols and non-dictionary words), is only stored in the brain of the user, is used only in one application and is changed frequently[9]. The massive number of online accounts has led to a password overload problem that directly impacts the security and privacy of users' data, since they try to deal with this problem by simplifying their passwords or reusing the same password on different accounts or keeping their passwords unprotected, on paper or password managers. At the same time, passwords are targets of multiple attacks, as they can be leaked, key-logged, replayed, eavesdropped, brute-force decoded and phished. All the aforementioned reasons have rendered the traditional username/password authentication solution unreliable and made user authentication an open research challenge.

The password-less authentication asset aims to contribute to the challenging task of authentication by providing a solution that is both secure and user-friendly. Particularly, it addresses the following research challenges [vasile2021web, Angelo2021how, panos2017security]:

- The mitigation of significant security gap that comes from the extended usage of the username/password scheme by providing a password-less authentication solution that deploys secure methods to authenticate a user that merge sophisticated cryptographic algorithms (e.g., elliptic curves) with *something the user knows* (e.g., PIN), *something the user has* (e.g., USB key), or *something the user is* (e.g., biometrics).
- The enhancement of the authentication process by utilizing a two-factor authentication (2FA) mechanism. With 2FA an additional layer of security is added on the authentication process, where the user after its initial authentication with username/password is requested to provide an additional piece of information (namely, *something the user knows*, *something the user has*, or *something the user is*).
- The limitation of password-related attacks, such as phishing, key-loggers, eavesdrop, and brute force.
- Unnecessary over-identification and information disclosure due to a lack of awareness and usability drawbacks (IDP-02).
- Honest but curious Service Providers and Identity Providers that might be tracking users, IdP could also become a potential point of failure (identity/data leakage in the IdP) (IDP-04).

[9] https://www.enisa.europa.eu/topics/csirts-in-europe/glossary/authentication-methods

4.4.4 Demonstrations Example

There have been no updates on the demonstration of the asset since D3.13. For a full description of the demonstration examples please refer to D3.13.

4.4.5 Future Work

Within the lifetime of the project the efforts regarding the password-less authentication asset will be focused on a thorough assessment and validation of its capabilities. After the end of the project the research will be focused on the integration of password-less authentication with the novel concepts of zero knowledge proof and zero trust to enhance the security and privacy of the authentication process.

4.5 Edge-Privacy

4.5.1 Overview

The Privacy Manager based on Edge Computing (PMEC) [rios2021personal] is aimed at helping users retain control of the sensitive data collected by their personal IoT devices. Users tend to own a number of IoT devices which might be spread across different locations, carried or even implanted. The data these devices collect is personally sensitive and must be handled with care.

This asset is devised as a decentralized set of virtualized services, called privacy manager instances, which can be seamlessly deployed in edge-ready devices in the vicinity of the personal IoT devices. The data collected by these devices are securely sent to a nearby privacy manager instance, which is in charge of interfacing third party users querying for these data. The data will be released according to the privacy preferences defined by the data owners, who can decide when, how and to which level of detail the data is shared with whom. See Figure 26 for the internal components of a privacy manager instance.

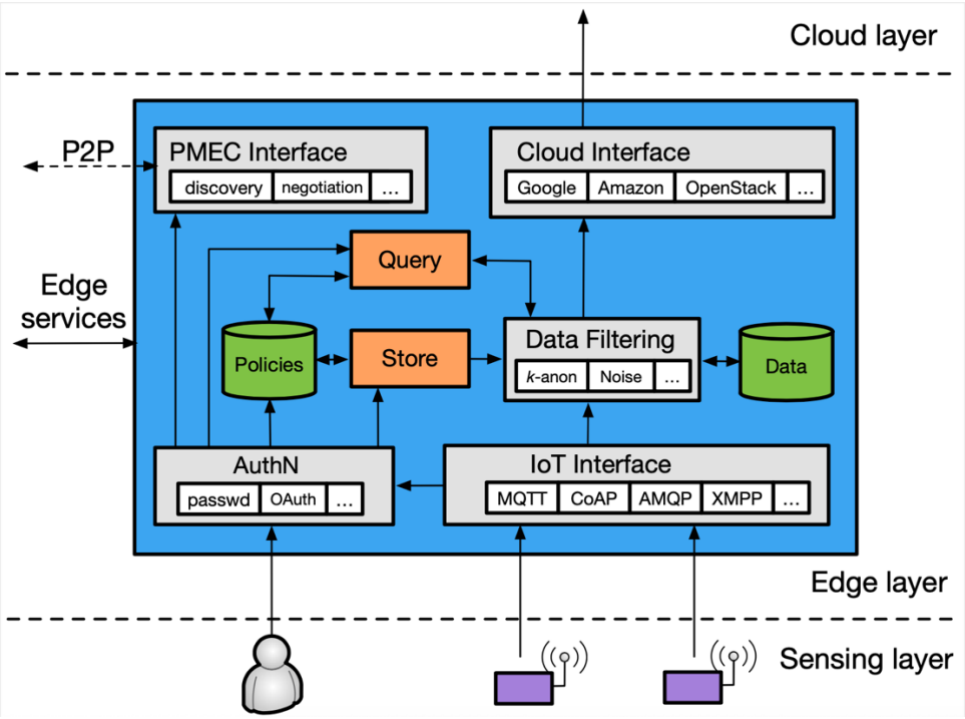

Figure 26: Architecture of the Privacy Manager for IoT data.

In essence, the PMEC asset operates as a privacy proxy between personal IoT devices and third-party consumers.

4.5.2 Main asset improvements since D3.13

Since D3.13, the PMEC asset has improved its stability and incorporated new privacy features. The previous version of this asset was capable of applying some ad-hoc anonymization techniques to the data obtained from IoT devices and later queried by the users. The evolved version of PMEC, which is still under development, supports new data transformation mechanisms and several privacy models, including *k*-anonymity, *l*-diversity and *t*-closeness.

In particular, the new version of the PMEC asset integrates an open-source library called ARX for the anonymization of datasets. More precisely, it uses the ARX as a Service (ARXaaS) [10] [i], which offers ARX features as a microservice, which is convenient for edge deployments.

4.5.3 Research Challenges addressed

The changes introduced in the evolved version of the Privacy Manager are not aimed to face new challenges but to improve the ability of the user to retain control of personal data collected by his/her IoT devices, which may be distributed across geographically disperse locations.

In summary, PMEC addresses three challenges defined in deliverable D3.11: the lack of mechanisms for controlling and limiting access to the data collected from numerous IoT devices (DP-05), the sanitization of data before it is sent to the cloud (DP-07), and the loss of control over personal data (DP-08).

4.5.4 Demonstration Example

The Privacy Manager can be used by different actors of the Smart Campus scenario to control the data collected by their IoT devices. In the previous deliverable (D3.13) we concentrated on the release of GPS information obtained from users' smartphones, which uploaded these data encrypted to the privacy manager. In turn, the privacy manager could serve these encrypted locations to authorized data requesters, for computing statistic by using tools like Sharemind, which is capable of performing computations over encrypted data.

In order to show the new functionality added to the privacy manager, the data collected by personal IoT devices (e.g., smartphone) is stored in plaintext in the privacy manager. We consider the same workflow as in D3.13. Most changes in this new version are introduced in the Deployment and User Management phase since the data owner needs to define the type of protection offered to the data in the database.

Assuming the privacy manager instance has been deployed and all necessary credentials created, the user has to define the database structure for holding the data. On top of that, the user has to establish the level of sensitivity of each of the attributes in the database as well as the privacy requirements for these data depending on the credentials presented by the data requester.

Both categorical and numerical attributes can be used. The formers require an ontology in order to enable their generalization while for numerical attribute it is only necessary to define the range of possible values and minimum discretization jump. Some numerical attributes can be defined to be of type redaction, meaning that part of the attribute can be removed. This is typical for parameters such as ZIP codes or geodesic location information. An excerpt of the file used for the definition of these attributes is shown next. Note that the user might need update this file with his/her own criterion or in case a categorical attribute is not included or is not complete.

```
{
  "categorical": {
        "disease": {
```

[10] https://github.com/navikt/arxaas

```
          "respiratory": ["SARS", "pneumonia", "bronchitis"],
          "digestive": ["gastric flu", "gastric ulcer", "intestinal cancer"],
          ...
}
  },
  "numerical": {
      "age": {"min": 0, "max": 100, "jump": 10},
      "salary": {"min": 0,"max": 150000,"jump": 10000},
      ...
  },
  "redaction": ["lat", "lon"]
}
```

Next, the data owner can modify the privacy preferences configuration file to indicate the privacy model to use to protect the database. Different privacy models, such as k-anonymity or l-diversity, can be used for different data requesters. Also, a given data type may be considered to be sensitive for a given type of data requester (e.g., researcher) and non-sensitive for another type of data requester (e.g., friend). See the following example:

```
{
  "action_type": "GET",
  "privacy_method": "KAnonimity",
  "resource": "SELECT * FROM persons",
  "attributes": {
      "identifying": "name, id",
      "insensitive": "temperature",
      "sensitive": "salary",
      "level": "3"
  },
  "conditions": {
       "requester": {"role": "researcher"}
  }
}
```

In this example, the data owner defines that a data requester with the role researcher can query for attributes in the table *persons*, but responses will have a k-anonymity level equal to 3. For this role, attributes *name* and *id* are identifying and will this be completely obscured, the attribute *temperature* is considered insensitive and will be reported as it is, and *salary* is sensitive. Sensitive attributes are not altered but thanks to the required level of k-anonymity, the data requester is prevented from linking this attribute to a particular individual. Note that any attributes in the database table but not explicitly declared here are considered quasi-identifier attributes. As shown above, these must have been defined as categorical, numerical or redaction.

After receiving a query from an authorized data requester (e.g., a user presenting a JWT token with role researcher), the Privacy Manager checks whether it is entitled to access these data. Before delivering the data to the requester, it will transform the data according to the privacy preferences configuration file for that type of user.

We have conducted some experiments to check the overhead introduced by incorporating the new data transformation mechanisms and privacy models thanks to the integration of the ARXaaS microservice as part of our asset. The results of these experiments are shown Figure 27

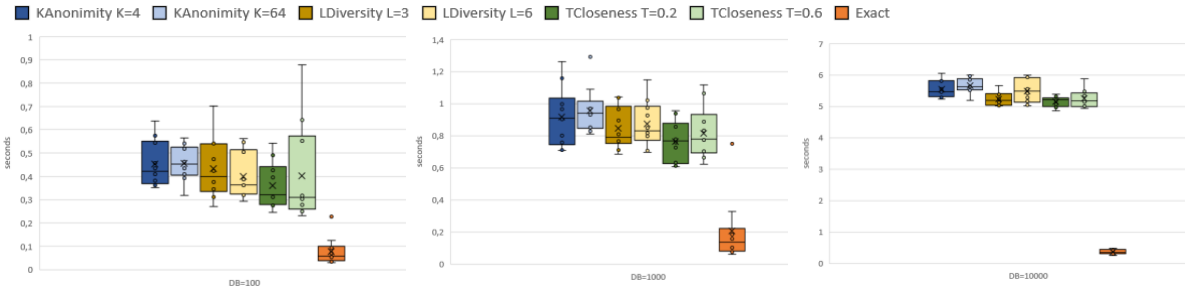

(a) Database size 100 (b) Database size 1.000 (c) Database size 10.000

Figure 27 : PMEC execution times for different database sizes

In particular, we performed queries to databases of three different sizes: 100, 1.000 and 10.000 entries. For each database we evaluated the performance of k-anonymity, l-diversity and t-closeness for different values of *k*, *l* and *t*. We also tested the performance of an exact query, which performs no privacy transformations to the dataset, to serve as a basis for comparison. As expected, the larger the database size, the longer it takes to process a query that requires data anonymization. When the database size is of the order or 10.000 entries, the response delay may be excessive for time-sensitive applications. Otherwise, the overhead introduced is reasonable.

4.5.5 Future work

Although there will be no further improvements of the PMEC asset within the project, we plan to continue researching on edge privacy and evolving the solution. We would like to provide user-friendly interfaces to facilitate the way users define their privacy preferences. Additionally, we aim to investigate mechanisms for holding sensitive information within the Privacy Manager with no trust assumptions on the edge infrastructure.

4.6 Privacy-Aware Aggregate Programming

4.6.1 Overview

Privacy-aware aggregate programming is a programming model centered around privacy protection and aggregate programming. The key concern is the trade-off between data utility and privacy protection. To illustrate this paradigm, consider one of the classical use-cases of aggregate programming for computing a so-called proximity field. The problem in this use-case is a 2D map where several agents, obstacles and locations of interest are spread. The agents want to collaborate in helping each other to find the locations of interest but they do not want to share their actual position with each other or with any central system. The aggregate programming solution to this problem is to build a proximity field: each agent in the map will indicate to closely located agents a notion of proximity. To calculate this field the agents, execute a simple aggregation-based program: iteratively aggregate the neighbors' proximity with the "min" operation, adding an estimate of the neighbor's distance.

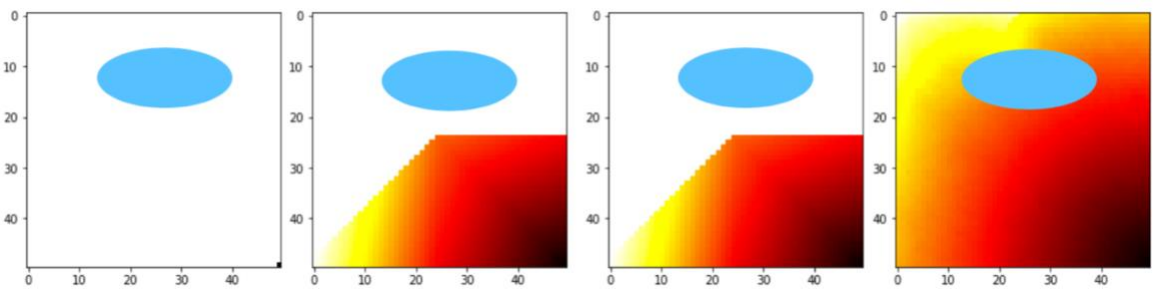

Figure 28: Incremental construction of a proximity field

In Figure 28 we see how the distributed computation of the proximity field evolves as the agents perform more and more rounds of the aggregation. In the Figure, there is only one point of interest (located at the bottom-right corner), one big oval (blue) obstacle, and agents in each of the rest of the coordinates in the map. The proximity field is indicated with a "heat" gradient that goes from "dark red" (very close) to "yellow" (far away), with a blank indicating unknown proximity. Such a field can then be used by the agents to navigate the map and reach the location of interest.

This asset is in a preliminary proof-of-concept stage and, as such, there is still not a proper tool supporting the approach. The current status is based on a prototype implementation that allows for simulation of proximity field constructions with distinct noise-based privacy protecting techniques (à la differential privacy).

4.6.2 Main asset improvements since D3.13

No improvements have been done for this asset since D3.13, so that the reader is referred to D3.13 for further information.

4.6.3 Future Work

A possible next step could be the development of an effective tool to support the approach described in this section. The tool could allow for assessing the impact of privacy-protecting techniques in the tradeoff between utility and privacy risks. This would require first a deeper investigation of privacy risks, e.g., related to inference of actual locations based on proximities, and on utility measures for proximity fields. Another long-term future work would be to extend the ideas to applications beyond proximity fields.

4.7 DANS

4.7.1 Overview

The Data Anonymization Service (DANS) is data protection tool, for preserving personal data privacy. "Considering regulatory aspects anonymized data are excluded from GDPR regulation because anonymized data is no longer "personal data"[11]. In this way, the DANS asset is an anonymization tool that avoids user tracking and user re-identification by the use of privacy and risk models which prevents privacy threats when data are managed. As perfect anonymization is not possible it is necessary to balance between privacy and data accuracy for analytics. The DANS is managing the user attributes in different ways, to be delivered to the data consumers:

- Identifying attributes which will be removed from the dataset.
- Quasi-identifying attributes which will be transformed accordingly with the specified transformation procedure, such as generalization or micro aggregation.
- Sensitive attributes will be kept as-is, but they can be protected using privacy models, such as t-closeness or l-diversity.
- Insensitive attributes which will be kept unmodified.

In order to preserve user data privacy, the DANS tool provides a k-anonymity privacy model for quasi-identifying attributes (as this model is the most used for protecting health data) and also provides an ℓ-diversity model for sensitive attributes" [D3.13].

DANS asset can be used in two ways:

[11] https://gdpr-info.eu/recitals/no-26/

- DANS as a service to be deployed on the data provider premises or a third party.
- As a java library to be embedded into the data provider system.

The DANS asset is based on the ARX library embedded in the core module. Figure 29shows the DANS main components [D3.2]:

- Core modules:
  - I/O Interfaces: services for input/export data
  - Data Encoding: transforms data
  - Data Management: orchestrator module
- Anonymization algorithms: algorithms to be used for k-anonymity or l-diversity
- Privacy criteria module: configuration module for setting the privacy criteria
- Public API: for accessing the anonymization service module
- DANS User Interface: Graphic interface allowing the user to interact with the anonymisation tool in a friendly way.

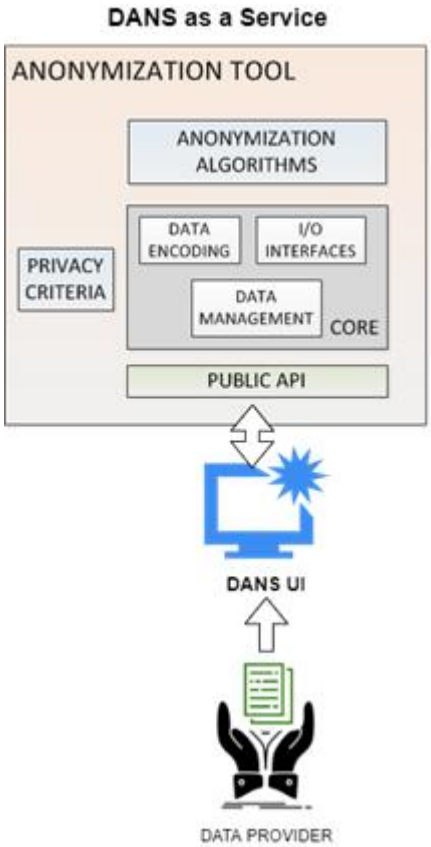

Figure 29 - Components of the DANS as a service asset

4.7.2 Main asset improvements since D3.13

Since D3.13, the DANS asset has been improved in two aspects as follows:

- Code improvement and polishing based on the feedback provided by end users during the testing and validation process performed in the context of the T5.6 Medical Data Exchange demonstrator. The public API offered has been polished for integrating the graphic user interface.
- Development of the graphic user interface (GUI) for accessing the cloud service and improve the user experience. This improvement also will facilitate the integration with any kind of data exchange platform requiring this cloud service.

4.7.3 Research Challenges addressed

The research challenges addressed by the updated DANS asset provided in Deliverable 3.13 [D3.13] are still valid. “Electronic health records gathered by hospitals and health organizations contain sensitive information. The aggregation of these data is an invaluable aid for preventing diseases, take accurate medical decisions and research purposes. These sensitive raw datasets cannot be directly shared and need to be protected in order to preserve the individual’s data privacy. With the aim to avoid a breach of data privacy when these data are delivered outside the medical institutions, anonymization techniques have been usually applied. De-identification of health records by applying the k-anonymity privacy model and using generalization techniques. With these techniques, identifiers such as name or unique identifiers are removed; quasi-identifiers such as sex, postcode or age are generalized; the sensitive data such as diagnosis are not modified for applying following analytics.

The application of these techniques provokes a loss of information, but it is necessary that the resulting anonymized dataset keeps the utility for subsequent analysis to be performed by the recipients of the data.

The anonymization approach developed by the DANS tool aims to provide a trade-off between privacy-preserving and utility.
We can map this general discussion about challenges to the ones defined in D3.11. Concretely, the asset addresses the following challenges:

- DP-07 by providing anonymization methods that detect and protect personal and sensitive identifiable information when data are managed out of the user’s control.
- IDP-01 by providing models that inform about the re-identification risks.
- IDP-04 by transforming medical data the ability to obtain information beyond the necessary is limited.
- LDP-01 is partially covered, as DANS is providing a statistical report regarding the anonymization process.”

4.7.4 Demonstrations Example

The use of the DANS asset could be helpful in the Geolocation Service in the Smart Campus scenario (section 3.3). “Once the geolocation data from the devices of university staff are collected and uploaded to the Edge Platform, the Privacy Manager could anonymize the geolocation data in two ways depending on how the DANS asset is used by the Edge platform either as a service or as an embedded library. The DANS asset can be deployed as a service providing the anonymization functionality on the service provider premises, in this case, the Edge Privacy Manager, leveraging the currently developed GUI, will communicate with the DANS service in order to anonymize the data location, as depicted in Figure 39 (a). The anonymization functionalities can also be embedded into the Privacy Manager component by using the DANS java library for anonymizing data location, as shown in Figure 30. (b). In both cases, the anonymized data can be used for analytics” [D3.13].

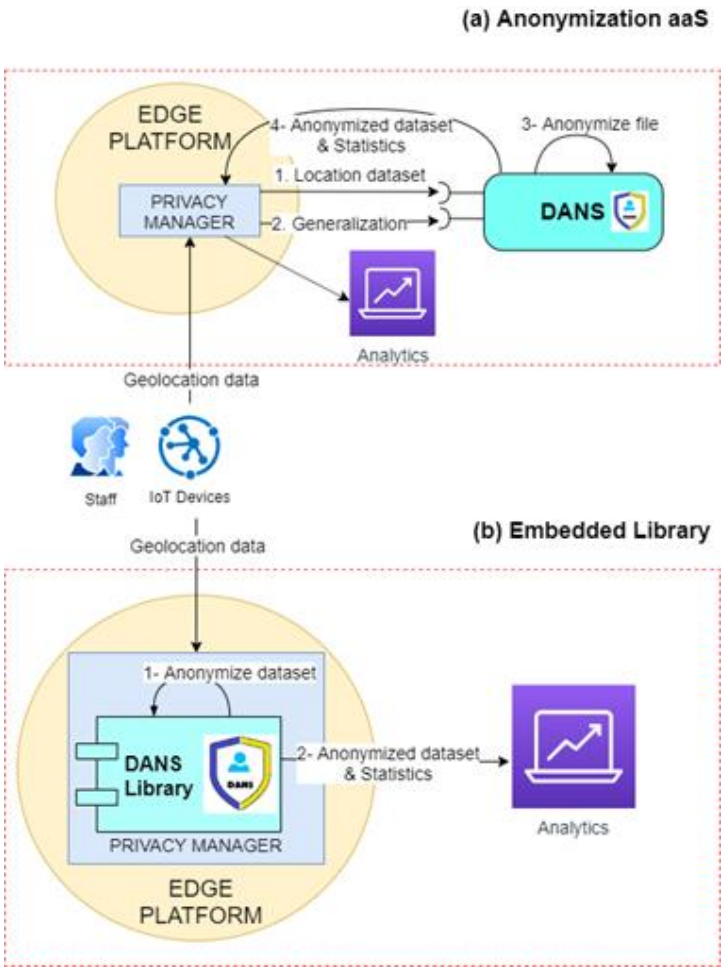

Figure 30 - Flavours of DANS tool (a) Anonymisation as a Service, (b) Embedded library

4.7.5 Future Work

Within the context of this project, there will be no further development of DANS. The focus will be the final integration, testing and validation of this asset in the WP5. Based on the feedback received from the end-users the DANS asset will be updated. Also, the core ARX library will be updated for including the latest functionalities and fixed bugs.

Additionally, is planned to study the use of standards such as HL7/FHIR for improving interoperability.

4.8 Cryptovault

4.8.1 Overview

CryptoVault is a system intended for users of different blockchain technologies. It comprises a hardware wallet that securely stores and manages the sensitive user keys, and a reliable method for backing up these keys as independent shares stored in multiple locations. The aim is to combine the best features of different wallet types while minimizing the risks related to these wallets. CryptoVault generates keys with high entropy, offers end-to-end protected key backup, signs transactions inside a trusted execution environment and can run key recovery without a single point of failure.

There are three security technologies used in the CryptoVault scheme: splitting keys into shares that are stored separately for secure backup purposes (on paper or on remote servers), isolating processing and storage from the main operating system by using the Intel SGX environment and using RSA to establish an end-to-end encrypted information channel between the SGX-enclave and remote servers when distributing the key shares.

In the identity management and service usage scenario, there are various options for using blockchain or smart contracts. CryptoVault can be integrated into this scenario as a tool for IT administrators to secure and backup their cryptographic keys.

4.8.2 Main asset improvements since D3.13

CryptoVault has not been developed further since deliverable D3.13, so that the reader is referred to D3.13 for demonstration results.

4.8.3 Research Challenges Addressed

One of the main selling points of blockchain technologies is their decentralized nature; there is no need for a trusted third party to check transactions. The drawback of the blockchain approach is the hassle of key management. There is no “blockchain support personnel” that can reset your password or create new credentials in case you lose the originals. Often the solution is to use a wallet application, which can be either software, hardware, or a combination of the two. Wallets store the private keys, generate signatures, and encode the transactions on behalf of the user.

There are several types of wallet applications [suratkar2020crypto]: Online wallets can be used with any web browser; hence they are easy to use, but require an internet connection, a secure browsing setup and placing trust on a third party once again. Mobile wallets are similarly easy to use and enable storing the keys on your personal device, but they require good overall security for the device as well as careful use of that device from the user. Desktop wallets that operate on a personal computer are slightly less convenient to use, and not all wallets are available for all operating systems. The security of the desktop environment is integral, and the keys are likely to be hosted locally. The user probably uses the wallet in a safe physical environment, reducing the risk of shoulder surfing attacks, thefts, and lost or broken devices. Finally, there are hardware wallets, which are dedicated devices for hosting and operating with user keys. These are significantly less convenient to use, but the security is increased as the device is not used for other purposes and the user is in full control.

There is one issue that is common to all the different wallet types: the loss of keys or the destruction of the device containing the keys. If that happens, the user will be unable to access their cryptocurrencies or other blockchain services. Again, we come to the original problem of trusting a third party, e.g., a wallet operator, to have a secure key recovery protocol in place, which is what the blockchain was supposed to mitigate. The CryptoVault [niko2021sec] solution provides a key backup and recovery method that is under the user’s control. The problem of a single point of failure with backing up a key to a different location, most probably controlled by a third party, is solved by using a secret sharing method and storing the shares separately. A highly secure but still accessible wallet is achieved by using a trusted execution environment on a regular, all-purpose computer.

By using the secret sharing method we can also address privacy challenge DP-08 from D3.11: When uploading information to the cloud the user partially loses control over the data. When using the CryptoVault system, the user is the only one who knows where all the key shares are stored, and the one share uploaded to a third-party cloud server is not enough to betray any information about the user's secret key.

4.8.4 Future Work

Within the context of this project, there will be no further development of CryptoVault. From a broader perspective, a new concept of social wallets has emerged as a way for users to back up their blockchain keys within their social circle. This is an interesting potential application for CryptoVault and for further research in general.

4.9 Elastic Deployment of TEE-based applications in the cloud

4.9.1 Overview

Processing a large amount of sensitive data requires secure and scalable solutions. For example, location data – as the one processed in the geolocation scenario - is considered as a privacy-sensitive data type so that the platform processing the data should be secured to minimize data leakage. At the same time,

the platform should allow for dynamic resource allocations so to cope with different amounts of data to be processed. ReplicaTEE is a solution that enables dynamic enclave replication and de-commissioning for TEE-based applications in the cloud. In particular, ReplicaTEE is designed to enable replication of applications that use Intel SGX – arguably the most popular TEE for workstations – as the undelaying TEE. Given the current deployment model of Intel SGX enclaves, replication of an enclave across machines requires the application owner to either be always online so to provision secret material to newly deployed enclaves, or to trust the cloud provider with managing the enclave secrets. An always-online application owner reduces the benefit of outsourcing to the cloud; trusting the cloud provider with managing enclave secrets avoids the advantages of using a TEE.

4.9.2 Main asset improvements since D3.13

There are no updates to this asset since the publication of D3.13.

4.9.3 Future Work

We expect to use the expertise collected while designing ReplicaTEE to design effective replication mechanisms for alternative TEEs. In particular, there is a growing trend of designing open-source TEE based on open architectures like Risc-V and we plan to enhance such new TEEs with replication frameworks so to foster the use of TEEs in the cloud.

4.10 Backdoor-resistant TEEs

4.10.1 Overview

SR-EPID is a subversion resilient version of Enhanced Privacy ID, the protocol used by Intel SGX for remote attestation. In the context of Trusted Execution Environments, remote attestation enables a party to establish trust in a TEE running on a remote platform. In the context of the demonstrator described earlier in this document, SR-EPID may be used in the geolocation scenario so that the integrity of the devices where geolocation data is going to be processed, is verified before such data is uploaded. All the popular remote attestation protocols balance authenticity and privacy by using group signature schemes. The latter is a digital signature scheme that allows a verifier to verify a signature as issued by a member of a trusted set while keeping the signer itself anonymous (within that set). In the context of Intel SGX, remote attestation uses a special system enclave, named Quoting Enclave, that certifies the application enclaves running on the same platform. Certification is realized by signing a report with the identity of the application enclave being certified. The group signature scheme being used – EPID – allows a party to verify that the report was issued by a genuine Intel SGX platform, without revealing any other information about the issuing platform, i.e., keeping the platform anonymous. Given the embedded nature of TEEs and the increasing concern on state-level adversaries, the scientific community is designing cryptographic protocols that can withstand subverted parties. In the context of remote attestation, a subverted signer may exfiltrate, through innocent-looking signatures, identifying information or even the signing key. Thus, an adversary can either identify the signer, thereby breaking anonymity, or obtain the signing key, thereby breaking the unforgeability of signatures.

4.10.2 Main asset improvements since D3.13

There are no updates to this asset since the publication of D3.13.

4.10.3 Future Work

We expect two main directions for future work. On the one side, we wish to improve the revocation mechanism of EPID to provide sub-linear complexity. As of today, the revocation mechanism of EPID has computational complexity that is linear in the number of revoked platforms and designing a solution

that reduces this overhead is key for TEE adoption, especially in embedded or edge devices. On the other side, we wish to design a post-quantum version of EPID that could withstand a quantum adversary.

4.11 Privacy-Preserving for Genomic Data (PP4Genomic)

This asset introduces a new genomic variation called Genome-wide Association Studies (GWASes). GWASes are statistically associated with a trait, such as a disease, in a group of individuals. Unfortunately, careless sharing of GWAS statistics might give rise to privacy attacks. Several works attempted to reconcile secure processing with privacy-preserving releases of GWASes. However, we highlight that these approaches remain vulnerable if GWASes utilize overlapping sets of individuals and genomic variations.

4.11.1 Overview

GWASes correlate differences in genomes to phenotypes, i.e., observable characteristics such as diseases, that they may cause. We proposed Interdependent Genome-Wide Association Studies (I-GWAS), a novel secure and privacy-preserving framework that supports interdependent GWASes. We show that unlike differentially private mechanisms, I-GWAS can afford a dynamic and unlimited number of safe releases by securely publishing GWAS results over selected batches of genomes and withholding only the genomic variations that would pose privacy risks.

4.11.2 Main asset improvements since D3.13

Since D3.13, the PP4Genomic asset has been improved in terms of efficiency and accuracy. More precisely, we show that even when relying on state-of-the-art techniques for protecting releases, an adversary could reconstruct the genomic variations of up to 28.6% of participants, and that the released statistics of up to 92.3% of the genomic variations would enable membership inference attacks.

4.11.3 Research Challenges Addressed

We highlighted that sharing genomes and Single Nucleotide Polymorphisms (SNP) positions in different GWASes requires the development of new privacy-preserving methods. In those settings, we showed that previous techniques that aim at preventing recovery and membership attacks might be ineffective. Using known statistical methods that enable safe and dynamic data releases for a single independent GWAS, we evaluated that up to 28.6% of the processed genomic information could be recovered by an adversary in a scenario that involved two interdependent GWASes. We also identified that between 80% and 92.3% of the SNPs might expose individuals to membership attacks. To fill this gap, we present novel frameworks that extend privacy-preserving algorithms to protect the genomic privacy of individuals participating in interdependent GWASes. In contrast, although Differential Privacy can be adapted to support dynamic releases of overlapping studies and presents a better release utility score in some scenarios, it imposes some restrictions, mainly on the number of releases that can be supported, and on the accuracy of the released statistics.

4.11.4 Future Work

Directions for future work include analyzing and improving the interplay between differential privacy and tolerable privacy budgets and combining differential privacy with genome-oriented statistical methods for protecting dynamic releases of interdependent GWASes.

4.12 GENERAL_D

4.12.1 Overview

As for any other requirement, a fundamental step for any organization (e.g., a Small and Medium-sized Enterprise (SME)) is to guarantee the compliant realization of the GDPR requirements by design. This means the integration of the data protection concepts into the overall software life cycle: from gathering the requirements to deployment and subsequent maintenance of the system. Research attention has been devoted to authorization systems because they are recognized, by scientific communities and private companies, as the successful elements for developing privacy-by-design solutions in compliance with the GDPR. However, to the best of our knowledge, most of the available proposals tend to target just a single aspect of authorization system development, and no integrated solutions for the GDPR-by-design compliant development through the entire life cycle are provided. Therefore, the GENERAL_D asset (or enabler) has the following objectives:

OBJ 1: defining a GDPR-based Life Cycle for authorization systems, i.e., defining a specific and integrated process development life cycle for the specification, deployment, and testing of adequate fine-grained authorization mechanisms, by considering legal requirements.

OBJ 2: providing an integrated environment for automatically enforcing the data protection or privacy regulations. Indeed, we define an integrated environment where some available solutions are combined: specifying the privacy requirements, controlling personal data, processing them, and demonstrating compliance with the GDPR in collecting, using, storing, disclosing, and/or disposing of the personal data. GENERAL_D Life cycle

We refer to Deliverable D3.2 - “Cross-Sectoral Cybersecurity Building Blocks” and D3.11 - “Definition of Privacy-by-Design and Privacy-Preserving Enablers” and the references from [generaDref1] to [generlaDref18] for a detailed description of the final release of GENERAL_D Life Cycle (Figure 31).

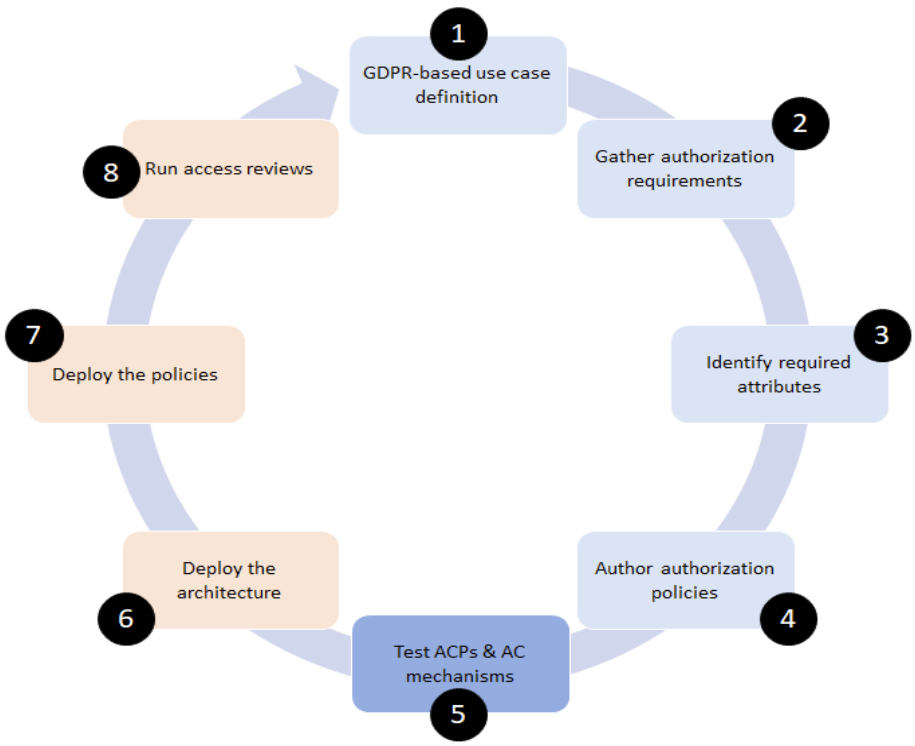

Figure 31: The Authorization Policy Life Cycle.

4.12.2 Main asset improvements since D3.13

GENERAL_D has been improved with test cases generation facilities starting from the developed GDPR-based Access Control Policies. The strategy we have integrated is GROOT, a general combinatorial testing approach for validating systems managing GDPR's concepts. A preliminary application of the GROOT strategy using the data derived in D3.13 (Demonstration Example 1: CCTV Surveillance) is also provided.

4.12.3 Research Challenges Addressed

Considering the challenges reported in D3.11 - "Definition of Privacy-by-Design and Privacy-Preserving Enablers, GENERAL_D targets the following challenges:

- **DP-05** by providing automatic facilities for assessing and testing access control systems that regulate/limit access to personal data [GDRef01, GDRef04, GDRef05, GDRef06, GDRef08, GDRef10, GDRef12, GDRef13, GDRef14, GDRef15].

4.12.4 Demonstration Example of GROOT in the context of CCTV Surveillance

GROOT is a general combinatorial testing approach for validating systems managing GDPR's concepts (e.g., Data Subject, Personal Data, or Controller). In the following, we first briefly illustrate the GROOT methodology, and then show its usage in the context of CCTV Surveillance access control. In describing the GROOT methodology, we use the following definitions:

Definition 1 (GDPR-based SUT Model) *A GDPR-based SUT Model is a tuple* $Model_{GDPR}(PAR, V)$*, where:*

- $PAR \subseteq \{DS, PD, DC, DP, C, P, PA, TP\}$ *is the set of parameters that affect the GDPR-based SUT, where DS = Data Subject, PD = Personal Data, DC = Controller, DP = Processor, C = Consent, P = Purpose, PA = Processing Activity, TP = Third Party, and*
- $V = \{V_i \mid i \in PAR$ *and* V_i *is the set of values for the parameter* $i\}$ *is the set of sets of the values that can be selected for each parameter.*

Definition 2 (GDPR-based Test Case) *Given a GDPR-based SUT Model* $Model_{GDPR}(PAR, V)$*, a GDPR-based Test Case is a tuple* $TC_{GDPR}(ATT)$ *where:* $ATT = \{ATT_i \mid ATT_i \subseteq V_i\ , i \in PAR$ *and* $V_i \in V\}$.

The GROOT methodology takes as an input a GDPR-based implementation, representing the GDPR in terms of a specification language. GROOT is composed of three main steps (see the following Figure 32): GDPR-based Model Derivation; Test Cases Generation; and Test Cases Translation

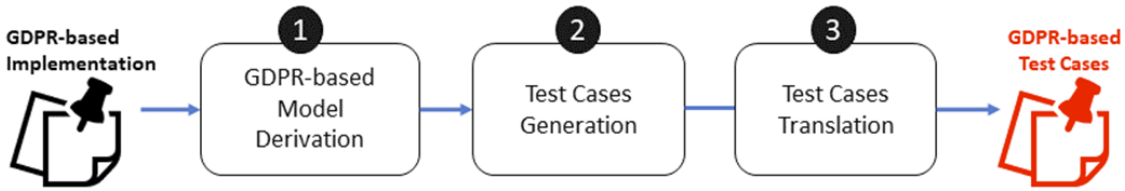

Figure 32 - GROOT Methodology

GDPR-based Model Derivation (Step1). In line with Definition 1, the GDPR-based SUT Model of the GDPR-based implementation is then derived. For this, the GDPR-based implementation is parsed in order to identify the set of parameters P, and the associated set of sets V. More precisely, for each parameter *i*, the subset Vi, containing the values used in the GDPR-based implementation, is derived.

Test Cases Generation (Step2). In this step, combinatorial testing is performed. Based on the derived parameters' values sets, different combinatorial strategies can be adopted such as : all-combinations, pairwise combinations, or t-wise combinations. For instance, in the all-combinations test strategy according to Definition 2, for each parameter *i* and its set of value V_i, the power set of $V_i(P(V_i))$ is derived, i.e., all possible subsets of V_i. Then, the obtained powersets $P(V_i)$ are combined so as to derive the test cases, i.e., the $TC_{GDPR}(ATT)$ tuples. Because combinatorial testing is costly, selecting the best combinatorial strategy that could be adopted may depend on different testing objectives such as coverage, effectiveness, reduction, or prioritization.

Test Cases Translation (Step3). According to the domain-specific language, each of the obtained $TC_{GDPR}(ATT)$ tuples in Step 2 is translated into a specific executable test case. In the context of access control, a test case is represented through an AC request that the access control mechanism can evaluate.

4.12.4.1 Using GROOT

GDPR-based Implementation. In this application example, the GDPR-based implementation refers to the policy associated with the consent management use case, GENERAL_D_CCTV_UC_1 reported in D3.13. In particular, we refer to UC_1_R1, with the corresponding access control policy (called UC_1_R1's policy) reported in the listing below. The policy allows a lawfulness of Personal Data processing related to Operator1 and Operator2, and it is composed of three rules (R1, R2, and R3):

UC_1_R1's Policy:

R1: permission(data_controller={Smart campus, CCTV team}, data_processor={Municipality Operator}, data_subject={Operator1}, personal_data{name, birth date, address}, purpose={Administrative}, action={Read, Write, Modify, Analyze}, consent={YES})

R2: permission(data_controller={Smart campus, CCTV team}, data_subject={Operator2}, personal_data{name, birth date, address}, purpose={Administrative}, action={Read, Write, Modify, Analyze}, consent={YES})

R3: permission(data_processor={Municipality Operator}, data_subject={Operator2}, personal_data{name, birth date, address}, purpose={Administrative}, action={Read, Write, Modify, Analyze}, consent={NO})

GDPR-based Model Derivation (Step1). According to the GROOT methodology presented in the previous section, the GDPR-based Model is parsed to derive the PAR, and the associated values of the parameters. In the case of UC_1_R1 Policy, the identified set of parameters derived from the policy elements is PAR ⊆ {DS, PD, DC, DP, C, P, P A}. For instance, the values associated with parameter DS is V_{DS} = {Operator1, Operator2}.

PE	PAR	V_{PAR}
data_subject	DS = Data Subject	Operator1, Operator2
personal_data	PD = Personal Data	name, birth date, address
data_controller	DS = Data Controller	Smart campus, CCTV team
data_processor	DP = Data Processor	Municipality Operator
consent	C = Consent	YES, NO
purpose	P = Purpose	Administrative
action	PA = Processing Activity	Read, Write, Modify, Analyze

Test Cases Generation (Step2). The combination of the parameters' values is computed in order to derive the set of test cases. Different strategies can be adopted in this step. By considering the all-combination, for each parameter j ∈ PAR, the power set of the associated values is derived. For instance, the power set associated with parameter DS (i.e., Data Subject) is P_{V_DS}

= {{}, {Operator1}, {Operator2},{Operator1, Operator2}}. Possible test cases are $TC_{GDPR}(ATT_1)$ and $TC_{GDPR}(AT\ T_2)$, where ATT_1 = {DS= {Operator1}, PD = {name, birth date}, DS = {Smart campus}, PD = {Municipality Operator}, C = {YES}, P = {Administrative}, PA ={Analyze}}, and ATT_2 = {DS= {Operator2}, DS = {Smart campus, CCTV team}, DP = {Municipality Operator}, C = {NO}, P = {Administrative}, PA = {Read, Write, Modify}}. For all-combination, the cardinality of test cases is 14.335, because the number of test cases follows exponential growth with the number of values' parameters. The number of generated test cases can be reduced by considering different approaches. For instance, by applying the pairwise technique the cardinality of test suite has been reduced to 160 covering the 14.335 variants.

Test Cases Translation (Step3). Finally, each of the obtained test cases is translated into an executable one. In the context of AC, possible AC requests associated with $TC_{GDPR}(ATT_1)$ and $TC_{GDPR}(ATT_1)$, respectively, are reported below. For instance, Req2 states that *Municipality Operator* (who) wants to process *name* and *birth date* (which resources) for *Administrative* purpose (under which circumstances).

Example of Access Control Requests Using GROOT:

Req1: request{DS= {Operator1}, PD = {name, birth date}, DS = {Smart campus}, PD = {Municipality Operator}, C = {YES}, P = {Administrative}, PA ={Analyze}},

Req2: request{DS= {Operator2}, PD = {name, birth date}, DP = {Municipality Operator}, C = {NO}, P = {Administrative}, PA = {Read, Write, Modify}}

4.12.5 Future Work

Future research activities will include full implementation of the GROOT methodology inside the GENERAL_D proposal and its application to additional case studies.

4.13 Blockchain Platform

4.13.1 Overview

This asset provides a blockchain-as-a-service platform but with improvements over state-of-the-art solutions (e.g., Hyperledger Fabric) that address several important issues the technology suitable for the fintech world [sforzin2017private]. Namely:

- Privacy: existing blockchain platforms assume work by broadcasting transactions to the entire network. However, both the corporate world and private citizens value their privacy, wanting to restrict data sharing to a few parties of their choice. Exchanging encrypted transactions is a good practice, but it still lets the network know when a particular transaction occurred.
- Scalability: In permissionless blockchains (e.g., Bitcoin, Ethereum) scalability is excellent, but sacrifices throughput. Permissioned blockchain can achieve higher throughput but sacrifice scalability. Today's standards require a scalable blockchain architecture that does not sacrifice throughput, making it useful for both organizations and private citizens.
- Lack of Governance: blockchains offer a distributed platform that does not need a trusted third-party. However, in real deployments, organizations and service providers want to oversee of their networks to enforce business logics and policies of their choice and to be able to grant or deny access to their services as they see fit.

The key feature of the platform is its architecture comprising “satellite chains”: small, independent blockchains with their own ledger, applications, consensus algorithm, and participants. Note that this model does not prevent the transfer of assets from one satellite chain to another, because all satellite chains are part of a larger “parent” network usually managed by the network’s owner (see “lack of governance” point above). The parent keeps track of the active satellite chains but does not interfere with their normal operations.

The advantage is that information is exchanged only between the intended parties, thus limiting the sharing of knowledge to who matters. The only requirement is instantiating a satellite chain that they all join. Transactions exchanged within that satellite chain are visible only to its members. See also Figure 33 for an illustration. In the example, the “circles” satellite chain does not know that the “triangles” satellite chain exists because they are independent. Similarly, the “squares” satellite chain and the “circles” satellite chain are not aware of each other, even if one entity is a member of both. Only that member knows the existence of both. There can be as many satellite chains as the use case needs, there is no upper bound.

This design gives two benefits: efficiency, because participants do not have to store transactions that are not relevant to them; scalability, because by allowing an unbounded number of satellite chain to work in parallel, the platform achieves a higher throughput than existing deployments.

Note that smart contracts are available as well: organizations can thus implement their business logics and policies via a program uploaded to the blockchain. In particular, the organization managing the network can assume the role of a “regulator” that implements policies that apply to all satellite chains active in the network. For the second cycle of the project, the asset named “Consensus Research” [d312] has been integrated into this one.

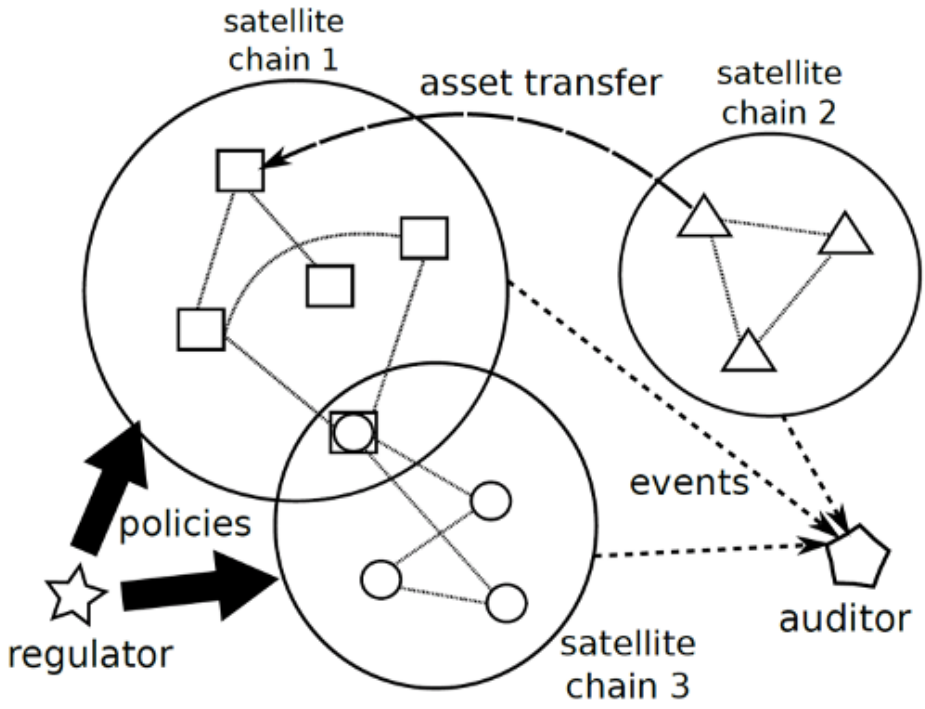

Figure 33: Blockchain platform architecture showing three independent satellite chains in action

Supporting these features is a Byzantine Fault Tolerant (BFT) consensus protocol - named FastBFT - NEC has perfectioned over the course of the project. [liu2018scalable]. The recent researchers’ interest in BFT protocols comes after years of the blockchain being in the spotlight. For example, Bitcoin relies on proof-of-work (PoW) to agree on the order and correctness of transactions. However, PoW has been proven to be inefficient and too slow to be useful in the fintech world. Therefore, research institutions and private organizations are investing in designing new BFT protocols that could finally allow

organizations, (e.g., financial institutions and supply chains) to leverage the blockchain for their business.

4.13.2 Main asset improvements since D3.13

There are no updates to this asset since the publication of D3.13.

4.13.3 Future Work

Presently, there are two research and development paths we plan to follow for future work on this asset:

- Blockchain interoperability: with the advent of web3, the blockchain has only become more heterogeneous. We plan to investigate protocols allowing for the transfer of assets and transaction between our platform and other prominent blockchain ecosystems.
- Smart contracts verification: smart contracts are a key element of modern blockchains, yet they are neglected. Indeed, they are a major vector of attacks; bugs in smart contracts code caused losses in the tens of million-dollar range. Finding and fixing smart contracts bugs is not trivial, and often costly (e.g., reuploading a smart contract in Ethereum is subject to fees). For these reasons, we plan to investigate static analysis methods that would spot and fix smart contract bugs, and integrate them in our platform.

4.14 Sharemind

4.14.1 Overview

Sharemind (Figure 34) is a secure computing platform that consists of Sharemind MPC (based on secret sharing) and Sharemind HI (based on trusted execution environments).

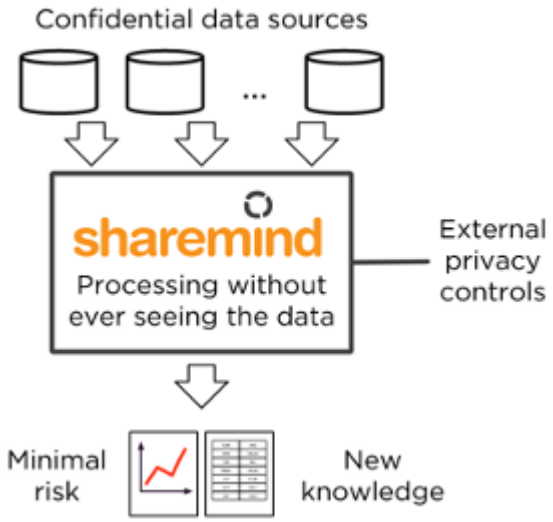

Figure 34 - Sharemind secure computing platform

Sharemind HI is a development platform for the confidential analysis of data from multiple parties on a centralized server with full control over exposing the data and results to others. It was created to reduce the risk of a privacy breach when processing confidential data. The data is encrypted at the source, by the data owner, and only then sent to the Sharemind HI service. The host of the service will not have access to the unencrypted data nor the encryption keys. Sharemind HI does not remove the data protections even while processing it, the data will remain protected by cryptographic means during the whole analysis.

Sharemind HI relies on a trusted execution environment (TEE) technology to provide security guarantees. A TEE isolates the security sensitive parts of an application from the rest of the system with the help of some trusted hardware. The TEE technology used in Sharemind HI to implement the privacy-

preserving data processing is Intel® Software Guard Extensions (SGX) which is available in modern Intel® processors.
Sharemind HI is built as a client-server service. The solution is based on tasks that run inside SGX. Each task resides in a separate SGX enclave. The client is an application that calls operations on the server, encrypts data and performs remote attestation on the server. The Sharemind HI server does the bulk of the work and is responsible for the following:

- checking if a user has the right to access the system (authentication),
- checking if a user has proper roles and data access permissions to perform an operation (authorization),
- managing the encryption keys of the data,
- managing the secure data transport in the solution between tasks and external stakeholders including data upload and download,
- storing a log of the operations performed in the server,
- scheduling the solution tasks to run.

The security model of Sharemind HI relies on the security guarantees provided by SGX. The data encryption model of Sharemind HI is illustrated in Figure 35. The input data, shown in light blue, is encrypted at the client side and sent to the server. The input data encryption keys of the data are securely transferred to the SGX protected enclaves. Likewise, the output data, shown in green, is encrypted inside the enclave and stored on the server. When requested, the enclave securely transfers the output data encryption keys to the authorized clients.

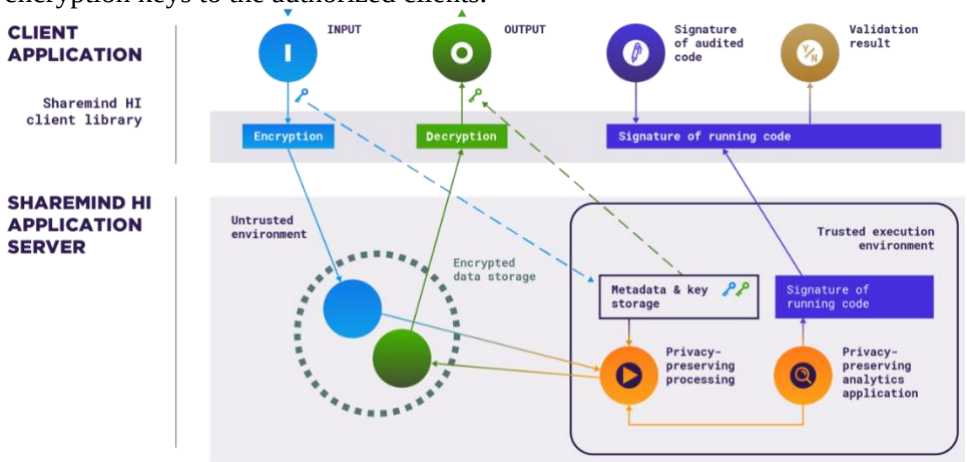

Figure 35 - Sharemind HI security model

It is the obligation of the enforcers to verify that a task is configured as expected. Input providers and output consumers specify which enforcers they trust with this task. Sharemind HI ensures that they can only upload data to and download data from tasks which have been approved by their trusted enforcers. This link of trust prevents clients from sending data to or receiving data from a wrong task.
At any point during the deployment, a client can request cryptographic proof of what analysis code is running in the server, shown in dark blue in Figure 35. This proof can be compared against a previously generated proof by an auditor who has validated the code to be secure.

For each Sharemind HI deployment, a deployment coordinator has to generate a deployment specific private key and public key certificate. This private key is used to sign all the client keys that want to communicate with the Sharemind HI server. The signed deployment certificate is loaded into the server at startup and is used to authenticate clients in the remote attestation. The root CA certificate is embedded into the server and verifies the validity of the deployment certificate. This process is needed to establish a root of trust for the server.

Sharemind MPC [sharemind] uses secret sharing and secure multi-party computation to protect confidential data at rest, in transit and in use. Data is secret-shared by the data provider at the source. All calculations on secret-shared data are done using secure multiparty computation, thus protecting data during the whole analysis lifecycle. Even the Sharemind MPC hosts providing the service will not have access to unencrypted data.

Sharemind MPC is built as a client-server service. The solution is based on tasks that run on Sharemind virtual machines that compute on secret-shared data using hundreds of implemented protocols for different operations and data types. Sharemind MPC servers carry out calculations and jointly manage:
• user rights to access the system (authentication),
• user access permissions to perform an operation and read or write a particular data table (authorization),
• authentication keys of system users and client applications,
• the secure transport of secret-shared data including data upload (secret-sharing) and download (reconstruction of secret-shared results),
• storing logs of operations performed in the servers,
• running applications invoked by users in real time or, scheduled to run periodically.

The model secret-sharing data and computing on secret-shared data in Sharemind MPC is illustrated in Figure 36. Confidential data of two data owners, shown in orange and blue, is secret-shared at the client side and shares are sent to the three Sharemind servers (computation nodes). All computations are conducted on secret-shared data by the Sharemind servers using secure multiparty computation protocols. Allowed final result shares (in green) are returned to the result party and reconstructed on-site.

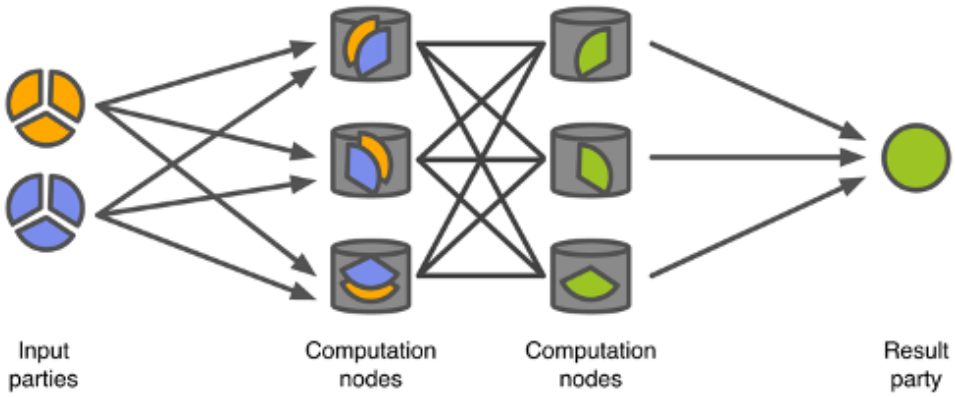

Figure 36 - Secret sharing two values (orange and blue) among three computation nodes and an MPC protocol that results in a secret-shared result (green).

4.14.2 Main asset improvements since D3.13

The Sharemind HI part of the Sharemind asset has not been changed, the Sharemind MPC part has been added. The research challenges addressed remain the same as described in D3.13. However, we have added the demonstrator description for the prototype Sharemind MPC deployment we completed in collaboration with Task 5.1 (Open Banking). The following subsection describes this new demonstrator.

4.14.3 Demonstration example

Banks often need to answer the simple YES/NO question “if any bank in the network has seen a specific IBAN in the context of financial fraud”. However, the banks need to complete this comparison in a privacy-preserving manner. In the present OBSIDIAN (Open Banking Sensitive Data Sharing Network for Europe) implementation pseudonymised (hashed) IBAN data was shared in an encrypted format.

If secure data exchange between banks would also include the date and transfer amount to this IBAN, i.e.
Date |Amount |hash
20/04/2020|2987|135febf76bfe3646b708be92e56fcc5c9063d328328d984f

and we would use secret-sharing and secure multi-party computation, we could also compute reputation data, share KPIs and global insights to fuel transaction risk scoring.

Cybersec4Europe partners I-BP and CYBER designed and implemented a Sharemind MPC demonstrator that allows the banks to secret-share details about potentially fraudulent transactions and compute the following KPI-s using secure multiparty computation. For system analysis and design we used the open-source software PLEAK (Privacy LEAKage analysis tool). The PLEAK model and disclosure analysis can be accessed from the URL: https://pleak.io/pe-bpmn-editor/viewer/y2uwdy5aZ3sENRZBsLdC/. For the demonstrator we generated 1000 records of artificial data for each of sample participating banks, designed and tested the demonstrator against the defined security goals for the system:

Confidentiality

- A participant should be able to choose, whether to secret-share hashed or IBAN data.
- A bank should be able to query only those hashes/IBANs they have secret-shared themselves.
- Bank's own reputation data should be provided only to the bank.
- Global KPI-s should be only provided to participating banks.
- Reputation data for each bank should contain only those hashes/IBANs, that the bank has provided.
- Network quality KPI-s should be visible to all configured data suppliers.
- No confidential data should be written into server logs. (hashes, ibans and local and global view data)

Integrity

- Each bank could verify that the data they have secret-shared has not been tampered with.

Availability

- Updated calculations should be available at least for the last day of provided data.
- There should be a configurable (cron job) parameter to periodically trigger calculations of global KPIs.
- The last successful calculation time of KPI-s should be available for all participants. (when the last successful calculation was performed)
- Every participant should be able to trigger the calculation of their own KPI-s.

The system allows each bank to secret-share basic transaction data (IBAN/hash, data and amount) and computes the following reputation information and KPI-s from secret-shared data using secure multiparty computation.

Reputation information

- Number of transfers realized to this account
- Average transfer amount for this account
- How frequently the account is used for transfers
- First time IBAN has been added as beneficiary
- Last time a transfer has been realized to this account

Example reputation information query result:

```
IBAN: 135febf76bfe3646b708be92e56fcc5c9063d328328d984f
Number-of-transfers: 5
Average-amount: 1542.0
Average-interval(in days): 268
First-date-added: 24/3/2021
Last-date-added: 22/12/2021
```

The results of the query are calculated from secret-shared data using secure multiparty computation. Result shares are returned to the allowed bank (result party) and reconstructed in the client computer.

The other participating banks and the Sharemind server hosts do not see the data provided by each bank or the results of other banks.

KPI-s and global insights
Min/Max/Average number of entries shared by participants to the network
Number of entries common to 1,2,…N participants
Amount of time over which a fraudulent IBAN is active Min/Max/ Average

Example query result (for generated test data of 6 sample banks):
Note that in a real deployment scenario the min, max and average number of entries contributed to the OBSIDIAN MPC extensions will be highly dependent on the size of the participating bank and the rigour of anti-fraud initiatives in the particular bank.

```
Min number of entries: 1000
Max number of entries: 1000
Avg number of entries: 1000
Number of entries common to participants: 68

Amount of time over which IBAN is active (in days)::
->      Horizontal Histogram (avg: 590.0, avg(non 0): 592.0)
        50      | oo (29)
        100     | o (15)
        150     | oooo (46)
        200     | ooo (40)
        250     | ooo (44)
        300     | ooooo (58)
        350     | oooooooo (98)
        400     | oooooooooooooo (175)
        450     | ooooooooooooooooo (205)
        500     | oooooooooooooooooooooo (271)
        550     | ooooooooooooooooooooooooooo (343)
        600     | oooooooooooooooooooooooooooooo (368)
        650     | ooooooooooooooooooooooooooooooooo (423)
        700       | oooooooooooooooooooooooooooooooooooooooooooo
(530)
        750      | ooooooooooooooooooooooooooooooooooooooooooooo
(556)
        800     | ooooooooooooooooooooooooooooooooo (433)
        850     | oooooooooooo (155)
```

4.14.4 Future Work

For Sharemind, we plan to implement new machine learning algorithms, so the data analysis that can be done with secure computing will be even more extensive. We are investigating ways of making the implementation more developer friendly, so more people would be able to implement applications for Sharemind, thus making the adoption more accessible.

4.15 Cloud-Based Credentials

Cloud-based credential systems provide a mechanism for privacy-friendly identity management. They enhance the state of the art by not only offering means for selective disclosure and data minimization (cf. also asset SS-PP-IdM (Section 4.3)) but by additionally focusing on resource-constrained devices.

This is achieved by outsourcing all computationally heavy operations to the cloud without putting a user's privacy at risk.

4.15.1 Overview

Cloud-based anonymous credentials were first introduced within the H2020 CREDENTIAL project[12]. The main components of cloud-based anonymous credential systems are as follows:

- **Identity provider**: This is the identity provider (or issuer) in a cloud-based credential environment. In our implementation, this is given by a web service where users can register and receive certificates on their sensitive data.
- **User application**: This is the user's local application required for performing privacy-preserving authentications. In our implementation, this is realized through an Android mobile app. However, the corresponding cryptographic libraries require only a minimum computational capacity and could easily be carried out also, e.g., on a student ID card.
- **CREDENTIAL Wallet**: This is a cloud-based service where users can upload encrypted versions of their credentials. To authenticate, the user application triggers the necessary computations in the Wallet, which, depending on the concrete implementation, directly computes with the relying service, or routes all communications through the user's device.
- **Relying party**: Due to the cryptographic specificalities of cloud-based credentials, relying parties (e.g., cloud services) need to integrate this end point into their system.

4.15.2 Main asset improvements since D3.13

This asset has not been further developed since its presentation in the previous iteration of this deliverable, so that the reader is referred to D3.13 for evaluation and demonstration results.

4.15.3 Future Work

An interesting open challenge for this asset would be to support presentation policies with the same expressiveness as traditional attribute-based credential systems. In particular, proving attributes like "older than 18" or "not yet expired" in a cloud-based scenario remains challenging, as already the fact that such a presentation policy is satisfied would leak information to the Wallet, and would thus require to carry out parts of the computation on the user's device at minimal costs.

4.16 Issuer-Hiding Anonymous Credentials

Issuer-hiding anonymous credential systems provide a mechanism for privacy-friendly identity management. They enhance over the state of the art by not only offering means for selective disclosure and data minimization (cf. also asset SS-PP-IdM (Section 4.3) and cloud-based credentials (Section 4.15)), but also allow for hiding the issuer of a certain credential.

4.16.1 Overview

An issuer-hiding anonymous credential system consists of the three main actors:

- **Identity provider**: The identity provider certifies users' attributes and issues anonymous credentials on them.
- **User application**: This is the user's local application required for performing privacy-preserving authentications.

[12] https://credential.eu/

- **Relying party**: Due to the cryptographic specificalities of issuer-hiding credentials, relying parties (e.g., cloud services) need to integrate this end point into their system.

The main difference to other attribute-based credential systems (E.g., SS-PP-IdM or cloud-based credentials) is that all of them reveal the issuer of the credential. This may seem like a natural property upon first glance. After all, the relying party must be able to decide whether it is willing to trust attributes certified by this issuer. However, this is too restrictive in many interesting scenarios. For instance, a national electronic identity might be used to participate in a European wide opinion poll: a priori, there is no need to reveal the citizenship of a participant. Similarly, students might want to use their student identities in different contexts: when requesting access to the university campus, it might be necessary to prove that they are currently enrolled at this university; however, when authenticating towards a cloud service, it might be sufficient to prove that they are enrolled at *some* university in order to get a student discount, but there is no need to reveal the precise university. In issuer-hiding anonymous credentials, the relying party can, in an ad-hoc fashion, decide which issuers it is willing to accept, and the user then shows that she owns a credential that was issued by one of these accepted issuers, without revealing which one.

By following this approach, issuer-hiding credentials immediately also address the problem of revocation of issuers upon corruption. In traditional settings, it would be required to invalidate all certificates that have ever been issued by the given issuer, leading to significant scalability issues in case of, e.g., millions of potentially affected users. With issuer-hiding ABCs, this challenge is solved by allowing for federated systems, where only a limited number of credentials is issued under each specific key, without compromising the users' privacy.

4.16.2 Main asset improvements since D3.13

Compared to D3.13, no further demonstration of this asset has been carried out. However, further research directions – e.g., combination with self-sovereign identity management solutions, have been identified and will be further analyzed beyond the project's duration.

4.16.3 Future Work

Future work for this asset might include, e.g., transferring the ideas into a post-quantum setting. Additionally, optimizing the functionality in terms of efficiency for specific application scenarios might further increase the change of a real-world uptake of the solution.

4.17 FlexProd and ArchiStar

The goal of the FlexProd platform is to enable computations on sensitive data without compromising the privacy of the data sources, while at the same time giving high integrity guarantees. Besides working on the functionality and efficiency of the existing platform prototype, research performed within CyberSec4Europe during the last months mainly focused on possible extensions to achieve end-to-end integrity and authenticity. The envisioned features are still under analysis and no implementation supporting these features is available for demonstration purposes yet.

4.17.1 Overview

FlexProd is an integrity- and privacy-preserving platform for distributed computations on potentially sensitive data, using ArchiStar libraries as subcomponents. Like Sharemind (cf. Section 4.14), FlexProd is based on secure multi-party computation, which allows multiple computation nodes to jointly perform computations on their respective inputs, without the other nodes learning any information about the other nodes' input. Users can now share their data using the ArchiStar secret sharing libraries, and store one share for each compute node, which can then perform, e.g., statistical analytics on data from potentially numerous users.
The following actors are involved in such a computation:

- **Data sources**: These are individuals or devices contributing data.

- **Compute nodes**: These are the entities performing the secure multiparty computations.
- **Data consumers**: These are the receivers of the computations of the result.

In an ongoing development and research effort, the asset is enhanced to also give the end-to-end integrity guarantees, by letting users sign their data before splitting it for the different nodes. The nodes then jointly perform the computation and furthermore will generate a zero-knowledge proof (potentially a so-called zk-SNARK for efficiency reasons) that all computations were performed correctly, and that the input data was equipped with valid signatures.
For a more detailed description of these assets, we also refer to the respective sections in D3.2 and D3.11.

4.17.2 Main asset improvements since D3.13

Since the release of D3.13, we improved this asset in different directions. In particular, we proved its efficiency and practicability in two additional application domains (smart manufacturing and air traffic management). These new use cases are described in more detail in Section 4.17.4.

4.17.3 Research challenges addressed

While a variety of secure multi-party computations exists (cf. Hastings et al. [HHNZ19] for an overview), existing frameworks do not directly cover the end-to-end authenticity and integrity of the results provided by the MPC framework. That is, the privacy and integrity of the computation are always based on the assumption that a sufficiently high fraction of the MPC nodes behaves honestly. With this asset, we aim at overcoming this limitation in the case that authentic (i.e., signed) data is processed by the MPC nodes. In this case, we aim at jointly generating a cryptographic proof (a zk-SNARK) that will allow the receiver to verify that all computations were performed correctly, where even a fully malicious MPC network would not be able to forge a proof. By doing so, all integrity guarantees can be based, e.g., on a physical trust anchor located in the users' devices collecting and generating the data. This approach could thus be seen as one possible approach towards addressing the integrity and privacy of supply chain information assets, cf. D4.4 (Section 4.5.10).

Regarding the challenges identified in D3.11, FlexProd also addresses the following challenges:

- *DP-02: When using secure multi-party computation (MPC) to analyze data, analysts are not able to see the individual data values.* By design and definition, MPC aims at not revealing the input data to any entity, including the receiver. However, one ambition of this asset is to remove any trust assumptions that need to be put into the MPC network by the data analyst. We believe that by doing so, and by aiming for formal integrity and authenticity guarantees, data analysts will be more willing to rely on the results provided by the MPC network.
- *DP-08: When uploading information to the cloud the user partially loses control over the data.* In collaboration with other projects (e.g., H2020 KRAKEN), the goal is to let the user, on a fine granular basis, specify for which computations her data may be used. Before triggering any computation, the MPC network would then verify that these policies are satisfied, thereby minimizing the risk of data abuse by malicious parties.

4.17.4 Demonstrations Example

In deliverable (D3.13) we concentrated on a possible integration of this asset in the context of a hypothetical scenario within the smart campus scenario.
In the last months, we have continued our work and adopted our framework for two specific application scenarios in complementary and independent application domains, in close collaboration with liaison projects such as KRAKEN[13], FlexProd[14], and SlotMachine[15], which we will briefly describe in the following.

[13] https://www.krakenh2020.eu/
[14] https://flexprod.at/
[15] https://www.frequentis.com/en/research/projects/slotmachine

Smart Manufacturing. The digitization trend in the manufacturing industry is gaining pace and novel cloud based market places will play an important role in the transformation, as it will support the establishment of a sharing economy. This is expected to have a significant impact in multiple dimensions, ranging from reduced costs, over increased innovativeness, and competitiveness to considerable environmental benefits.
However, mutual distrust and data sovereignty is of special concern in the manufacturing industry, e.g., related to corporate secrets and customer data, and thus centrally organized platforms cannot provide the required level of data privacy and trustworthiness needed for the manufacturing industry.

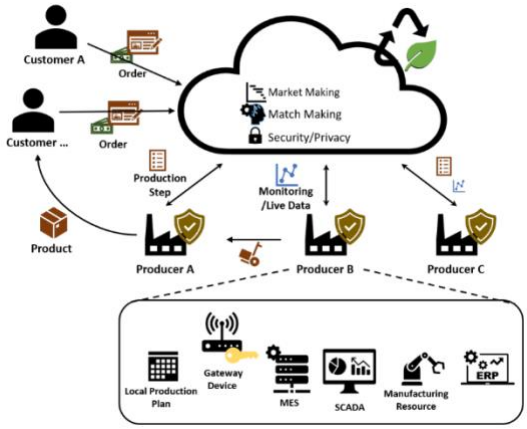

Figure 37 - Smart Manufacturing scenario

In [LWK22a] we studied the security and privacy aspects for the case of a market platform for outsourcing in manufacturing. Specifically, we considered the scenario where producers offer production resources to customers, cf. Figure 37**Error! Reference source not found.**. We showed that the requirements identified together with relevant stakeholder are challenging and sometimes also contradicting on the first sight. To address this challenge we combined different cryptographic building blocks into a novel framework for more secure but transparent decentralized data markets. In particular the framework combines secure multiparty computation with zero-knowledge proof of knowledge methods and blockchain to enable flexible sealed-bid auctions which are also publicly verifiable.
A proof-of-concept was developed, and benchmarking results show that the framework can efficiently address all requirements established.

Air Traffic Management.
An airport slot allows an airline to operate at an airport at a specific point in time. Such slots are necessary, e.g., due to runway throughput limitations or available parking space. However, due to external events such as weather conditions, deviations from the original fight plan (and thus assigned slots) are part of the everyday business at any airport. To minimize delays and costs at large costs, it would thus be beneficial to optimize the assignment of starting and landing rights, also across competing airports, based on continuous monitoring of the current situation at the airport, and taking into consideration airlines' preferences and requirements.
With User-Driven Prioritization Process (UDPP), a first such system has already been developed and deployed, which goes beyond the current slot swapping process implemented in EUROCONTROL's Network Manager, and shows large benefits in practical evaluations [PGB+21]. However, reaching a global optimum across multiple airlines is currently not possible, as airlines are reluctant to share their preferences with other airlines, thereby not achieving potentially significant further improvements.
Leveraging multi-party computation, we thus developed a decentralized platform that also enables the collaboration and optimization across airlines for optimal flight sequencing. From a modeling point of view, it turns out that the problem to solve corresponds to a linear sum assignment problem (LSAP): in such a problem, a number of tasks (e.g., takeoffs or landings) is assigned to a number of agents (e.g., airplanes), where each agent-task assignment comes with certain costs (e.g., priority of the slot for the airline). The goal of LSAP is then to find an assignment between tasks and agents optimizing the overall costs.
Given the financial and economic impact of slot assignments, it is furthermore required to prove the authenticity of any slot assignment, to minimize the risk of biased MPC nodes, e.g., giving preference to a specific competitor. Furthermore, due to the high frequency of slot assignments (multiple times per hour on large airports), the frequency of delays, unforeseen changing weather conditions, etc., the efficiency of such a system is of utmost importance.
In [LWK22b] we present the detailed evaluation results including benchmarks and additional findings.

4.17.5 Future work

The development of FlexProd is still ongoing, and many further improvements and applications are envisioned. For instance, the approach could be used in the context of data markets, where analytics of sensitive information are shared, but authenticity is important to the data buyer. Also the inclusion of heuristic algorithms to find high-quality solutions in complex optimization tasks is an interesting open question.

4.18 GDPR compliant user experience

GDPR compliant user experience is combined from two sections. The first is the Guidelines for General Data Protection Regulation (GDPR) which present the regulation's requirements through the GDPR principles. The second part of the enabler is the Data Protection Impact Assessment (DPIA) template. As the name suggests, this part of the enabler can be used to help guide the user through the process of doing a DPIA and also serve as documentation for the performed analysis.

A DPIA could be considered in many areas of the Smart Campus scenario as many of the services will use the personal data of the attending students and personnel. Here are the main properties of processing personal data in Smart Campus scenario can cause the DPIA to be required:

- is likely to pose a high risk,
- involves the use of new technologies,
- involves systematic monitoring,
- involves sensitive personal information (e.g., biometric data),
- refers to a significant amount of personal data at the regional level,
- can affect a large number of data subjects,
- prevents individuals from using the service or contract,
- includes monitoring of publicly accessible areas on a large scale (e.g., public university).

A DPIA must be performed before any type of processing is carried out and is an ongoing process that has to be regularly reviewed and brought up to date. A finished and properly performed DPIA will also help an organization evaluate, document, and later show how they comply with all the personal data protection requirements. A DPIA template could therefore be implemented as a forerunner to many different scenarios within Smart Campus.

The purpose of the DPIA is to systematically analyze, identify and minimize the impact the identified risks could have on the privacy of the data subjects. We will therefore prepare a DPIA example; however, in it, we will limit ourselves to the data collected and later processed in a typical student enrolment process.

4.18.1 Overview

Regulation (EU) 2016/679 of the European Parliament and of the Council or more commonly known as General Data Protection Regulation (GDPR), is a legal framework that sets guidelines for the collection and processing of personal information.

Guidelines for GDPR Compliant User Experience is a deliverable that was produced as D3.6 in the CyberSec4Europe project. As its name implies, it is a collection of guidelines, best practices and recommendations for achieving GDPR compliance. However, here we will focus on a specific section of the deliverable, which was designed to serve as a template for the process of performing a Data Protection Impact Assessment (DPIA). We will refer to it as the DPIA template. The template is like a to-do list with guidelines on how to perform specific tasks and some pre-prepared structures to support the user. DPIA template is a combination of a guide and pre-prepared content in the form of table templates that personal data controllers can use to perform the DPIA. This solution aims to be primarily

of use to the smaller organizations having problems performing or having questions about the assessment's specific steps by giving them a starting point on which they can build.

DPIA is meant to identify and minimize personal data protection risks by systematically analyzing the processing of personal data. Unlike most other risk analyses , DPIA is concentrated on the prevention of harm to data subjects, individuals, and overall society rather than the risk to the organization itself. A DPIA is a legal requirement under the GDPR when the processing is likely to result in a high risk to natural persons' rights and freedoms. This is an excellent example of a condition set by the GDPR for which it is difficult to instinctively know whether it applies or not because there is no definition for "likely to result in a high risk" and the type of issue the enabler is meant to resolve.

The major elements of the DPIA template are presented in Figure 38. The DPIA template aids with the initial decision on the necessity of performing a DPIA. If the circumstances demand the organization to perform the assessment, then the template describes and provides guidelines for the DPIA steps. The "Conduct the self-assessment" (bottom left former in the Figure) is optional and the last step in the DPIA. Before the solution/process can be implemented in the organization, it is important to also make sure all other GDPR requirements are met, which is the purpose of the more broad GDPR compliant user experience enabler. DPIA template contains all the basic information about the assessment as well as many recommendations and good practices on how to perform it. For further detail please refer to D3.13.

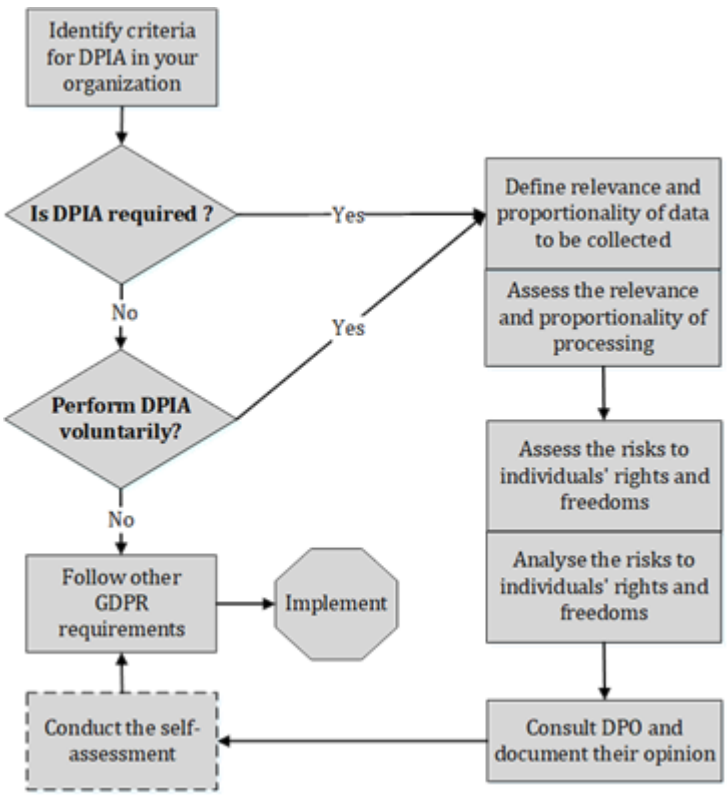

Figure 38: The main steps in the DPIA Template.

4.18.2 Main asset improvements since D3.13

This asset was presented in the previous demonstration deliverable (D3.13), and it has not changed since. If you want to know more, refer to D3.13, the original deliverable introducing the asset D3.6, or you can see the overview of the asset, relevant resources and any future changes or additions at https://github.com/cs4ewp3/wp3/tree/main/3.7.

4.18.3 Future Work

In the future, some effort should be put into keeping the guidelines up to date as and if new recommendations or rulings on achieving GDPR compliance become available. An important goal for

the future is to apply the guidelines and especially the DPIA template in real-world organizations and use their feedback to improve the guidelines and the template. An interesting question is also what potentially additional privacy and personal data protection aspects brought in by local legislation should organizations from the individual EU Member States pay attention to. This is, at least partially, addressed in the following Interoperability and cross-border compliance asset.

4.19 Interoperability and cross-border compliance

The Interoperability and cross-border compliance enabler address issues related to different eIDAS (electronic IDentification, Authentication and trust Services) implementations and legislation differences in EU member states, ultimately hampering the idea of a Single European Market. The eIDAS regulation provides a common foundation for secure electronic interaction between citizens, businesses and public authorities. EU is currently working on amending and improving the regulation to create the so-called eIDAS 2. In this asset, which was published in D3.18, we focused on analysing a sample of current eDIAS implementations and selected eIDAS network cross-border use-cases with the intent to identify any deficiencies in interoperability or cross-border compliance.

In the Smart Campus scenario, identity management and service usage will provide means to perform authentication and authorization. This brings university services to the students after they have been physically authenticated and granted credentials. However, this approach is limiting to students wanting to enroll remotely. By connecting to the eIDAS network, students could use their eIDs to remotely enroll and receive the campus credentials that they would then use when they actually go to the university or use them remotely where applicable in the case where the lectures are held remotely – for example during the Covid-19 pandemic foreign students would not have to travel to the university at all.

4.19.1 Overview

The eIDAS regulation aims at increasing the effectiveness of public and private online services, electronic business and electronic commerce in the Union. To this end, it includes provisions for electronic identification and trust services. However, after more than 5 years, only 14 Member States (representing 59% of the EU's population) have notified at least one eID scheme. This and other issues have caused the EU to review the eIDAS (electronic IDentification and Trust Services) regulation and change it to reduce the issues and improve the adoption of eIDs and trust services across the Member States. On June 3rd 2021, the European Parliament and European Commission published a Proposal for amending Regulation (EU) No 910/2014 to establish a framework for a European Digital Identity, which will someday become what is being called eIDAS 2. In our research as part of the CyberSec4Europe's deliverable D3.18 Analysis of interoperability and cross-border compliance issues, we considered some problems with the current situation, namely how eIDAS compliant implementations across the EU Member States are different and how that could be a problem for interoperability and cross-border cooperation of service providers and users.

In the research, we have focused on a specific sample of common eIDAS network cross-border use-cases to identify additional shortcomings of the current framework with the intent to suggest further additions to the current regulation and in this way achieve EU's long-term cybersecurity goals as well as to promote the use of eIDs in the commercial sector of the European Single Market. By analysing real-world implementations (from Italy, Slovenia, Spain, and Switzerland) in the chosen use-cases, we found shortcomings of the current regulation in the areas of organizational independence, remote access to the banking services, remote video identification, use of electronic signatures in public administration and remote access to the EU Digital COVID Certificate, commercial access to the eIDAS network, biometric authentication mechanisms, and finally some technical issues with mechanisms used to provide security and authentication.

4.19.2 Main asset improvements since D3.13

The previous demonstration deliverable of this task (D3.13) presented the GDPR section of the Interoperability and cross-border compliance. Since then, we have finalized our research into the interoperability and cross-border compliance of eIDAS. Therefore, the research described here is completely new. As such content in this section is completely different from the results presented in D3.13. If you are interested in the previous GDPR research refer to D3.13, the original deliverable D3.18

addressing both GDPR and eIDAS aspects, or you can see the overview of the asset, relevant resources and any future changes or additions at https://github.com/cs4ewp3/wp3/tree/main/3.7. Video describing the assets and results from this chapter is available at https://www.youtube.com/watch?v=k8ZPbeXkTM4.

4.19.3 Research challenges addressed

The eIDAS regulation is currently under review by the EU. A revised proposal of the regulation was proposed by the EU in June of 2021. The main challenge of the contribution is to identify additional current eIDAS properties that should be considered in the future alterations of the regulation.

In different analyses of the current regulation, very little notice has been given to the differences in eIDAS implementations between the EU Member States, which could lead to reduced interoperability or even unusable (non-compliant) eIDs, when used in different countries. This contribution looks at how eIDAS solutions are implemented in a selected EU Member States and identifies situations and areas where there are significant differences between them. The identified issues and their consequences include privacy/security/trust vulnerabilities, unbalanced markets, incomparable authentication levels of trust, etc. The identified issues do not include situations that would be improved on by the recently proposed regulation update (i.e. eIDAS 2) and should, therefore, be considered in improving the current eIDAS regulation.

4.19.4 Demonstration Example

Here we summarize the findings of the research into the interoperability and cross-border compliance of eIDAS implementations. They are mostly specific to an area (e.g. banking), or they are really only a problem from the perspective of the entire eIDAS network (because of the differences between countries). For this reason, some of them do not impact the Smart Campus scenario directly. However, if the identity management and authentication mechanism of the Smart Campus would be linked to the Spanish (for the purpose of this exercise, we assume the campus is in Spain) eIDAS node, the identified issues could affect how, mostly foreign students, interact with the service.

The identified problem areas in ensuring interoperability and cross-border compliance of the eIDAS are:

1. Organisational independence between supervisory authorities and service providers

   The results show a pattern where supervisory authorities are providing trust services at the same time. This is contradictory to the families of standards (International standard of Auditing 200, IESBA Code, ISO 19011) that define the independence of auditors or other types of professional reviewers. Since cybersecurity incurs costs, providing the function of supervision and provision of services in the same entity could affect the independence, at least in appearance. This arrangement could also affect the competition in the market as the main entity that is offering trust services and defining legislation or even access to the eIDAS network is also competing with other service providers in a closed market.

   Spanish supervisory authority for trust services provides non-qualified trust services at the same time.

2. Remote access to the banking services across borders

   When setting up eIDAS, one of the use-cases was to enable the remote opening of bank accounts across borders. However, the area of banking is heavily regulated, and there is much local legislation that could hinder the envisioned seamless connection with foreign banks. The examples from the included countries showed that each of the countries has a different system in place. Every Member State has its own solution and its own requirements for the remote opening of a banking account.

3. Remote video identification

   With the digital interoperability of the eIDAS network, borders are fading away. A citizen obtaining an identity in one Member State could use acquired identity in any other Member State if eIDAS is followed strictly. In marginal scenarios, a citizen of one Member State could even obtain digital identity using remote video identification in another Member State and later use that identification in its own state. This could already be a case, for example, if the bank in Spain decided to provide

banking services and trust services according to eIDAS to be freely used by their customers. Consequently, the questions of regulating requirements for remote video identification are bleeding into the regulation area of the eIDAS network.

Again, the results from the selected countries show that, where remote video identification is allowed, each state has their own requirements regarding the security of remote identification. Spain does allow for video identification with a level of security equal to physical personation. The difference between how remote video identification is performed will potentially lead to different levels of trust in the obtained title and difficulties in their cross-border recognition. In the case of Smart campus, this could mean that foreign students that have obtained their eID over a video identification with different levels of standards than are required in Spain might have problems using the campus services or will be able to use them with a different level of trust.

4. The use of electronic signatures in public administration and remote access to EU digital COVID certificate

   eIDAS provides different levels of assurance and different kinds of electronic signatures. Every level of assurance and every kind of electronic signature brings additional costs and complexity into the information system and user experience. Businesses are therefore reluctant to use higher levels of assurance than necessary. Since eIDAS was primarily targeted at the public sector, we were interested to understand whether assurance levels and the kind of signatures used in the public sector could be comparable across the Member States.

   Even though different levels of assurance and different kinds of signatures are defined in the eIDAS, there is little convergence in understanding what levels should be used in specific use-cases. The surveyed countries support different authentication methods, with no clear indication that they prefer or promote the highest assurance level electronic signatures. When governments are not promoting the use of the highest assurance levels and qualified electronic signatures with their services, the technology support and acceptance in the population is lower, and this spills over to the commercial market. Commercial services that require higher levels of assurance (e.g. banking, insurance) either by the law or because they are not prepared to take the risk of lower assurance levels are left alone to promote the use of higher assurance level technologies with the citizens.

   In the case of the EU digital COVID certificates, where the implementation was more consistent across the Member States, the minimum assurance level to access the EU Digital COVID certificate was the same across all surveyed Member States.

   In Spain, the use of electronic signatures is not required, but it is well supported in the public administration. This would give foreign students an additional bonus of easily and remotely dealing with any requirements for moving and living in Spain (e.g. report their temporary residency etc.).

5. Commercial access to the eIDAS network

   Since eIDAS does not have a condition that the eIDAS network has to be accessible to private entities, this is left to the regulation of the Member States. The results from the selected Member States vary. Some Member States do not allow access for the public sector (Spain), access is envisioned but not implemented (Slovenia), or access is allowed even for foreign businesses (Italy).

   With some governments providing access to foreign entities, the competition between local regulations will also start to build. Companies will have the option to choose an eIDAS network entry point and consequently control their costs. That will put pressure on the public providers to stay competitive or local nodes may start to lose interest. This may have a negative impact. If there is too much open competition between trusted service providers in different Member States, that may have a direct effect on the security of the network in the different Member States. Security incurs costs. If security requirements in some Member State are lower than in other Member State, this will give local providers a competitive advantage. Consequently, care has to be put into requiring the exact basic security requirements in all Member States.

   Spain does not support private entities connecting to the Spanish eIDAS node (and as a rule, foreign eIDAS nodes do not accept connections from private organizations unless the organization’s home

country also supports connections from private entities), however, if Smart Campus is a part of a public university, then this problem would be avoided, and the connection would be allowed.

6. Biometrics as Authentication Mechanism

   In the private sector, especially banking and the general public identity providers (e.g. Google, Microsoft), we see the rise of the use of biometrics.

   The use of biometrics is currently a "grey area". The use cases are primarily based on the biometric capabilities of current mobile devices. That means that the service providers (e.g. banks) are not processing biometric data and are consequently not under the GDPR requirements. There is no certification scheme in place, and there are no specific requirements for the use of such devices. Even though these devices have a direct impact on the security of the service for the end-user, the service providers don't have contracts with "biometric security device providers", e.g. Apple, Samsung etc. even though that was the case when the banks were buying authentication solutions on the market to meet their needs and the needs of their customers. Consequently, the user is left to their own selection of the mobile device and the final security of the service will vary depending on the device selected. For this kind of authentication, we are proposing the term "Bring Your Own Authentication Device – BYOAD".

   The overview of the selected countries showed that Switzerland (member of the European Economic Area, not the EU) is the only one that supports biometric authentication with eIDAS services.

7. Technical authentication and onboarding security mechanisms/requirements

   The authentication mechanisms in the eIDAS network across the Member States are using technical solutions with questionable security attributes, like SMS One Time Password (OTP) codes. Some Member States are even turning to security questions that are deemed obsolete. Considering the stakes at play, it is better not to trade ease of use (which, as far as we can tell, is the main reason for their use) for actual security. This is especially true because alternatives do exist, and steering away from using weak authentication would also improve the trust users have in the system. Therefore, current standards may not be detailed or current enough to support the latest findings in cybersecurity. This finding is another indication that a call to increase the speed of security standards development is justified.

   When analysing security requirements of electronic signatures, it also became evident that eIDAS uses very strict wording regarding capabilities of the qualified electronic signature, namely that "any subsequent change in the data is detectable". However, using the current state-of-the-art technologies for the creation of electronic signatures, an absolute guarantee of no changes in a message is not possible (there is an extremely tiny chance of different messages resulting in identical signatures). To avoid any differences in how courts in different Member States may explain and understand this term, we suggest that the wording be amended in a way that will refer to how existing technology works.

4.19.5 Future work

The work done in this research has some limitations, like the language barrier for which reason we have to rely on cooperation with partners from other countries that can supply information on their own eIDAS nodes and their legislation surrounding it. Considering this limitation, we would want to extend the study by including more EU Member States in the comparison. Additionally, there are other possible point of contention we have not considered in this study worth analyzing in the future as well as going into more detail of the already identified issues.

4.20 Privacy-preserving Federated Learning

4.20.1 Overview

Over the last years, Federated Learning (FL) [mcmahan2017communication] has attracted a lot of interest due to its inherent advantages and applicability in many use cases, including intrusion detection. Compared to centralized machine learning approaches, FL enables training global machine learning models from decentralized data. Based on local model updates and fusion algorithms normally executed at a central server, generally called aggregator, it offers clients the possibility to participate in the training process of a model without sharing their local data to a third-entity, as it was necessary in centralized approaches, thus preserving the privacy.

However, it has been demonstrated recently that this federated training process is still exposed to some privacy threats. Specifically, a man-in-the-middle between a client and the aggregator could infer sensitive information about the data that the client used to generate a model update, since some techniques can be applied to the exchanged model weights [mothukuri2021survey].

Given these circumstances, this asset proposes the application of Differential Privacy (DP) throughout the FL process to ensure that model updates shared by clients in the scenario cannot be used by an attacker to infer anything from the clients' private data, thus providing a more robust FL mechanism compared to the classical approach.

4.20.2 Main asset improvements since D3.13

The first version of this asset is presented in this deliverable.

4.20.3 Research challenges addressed

Below, the main research challenges addressed are listed:

1. Evaluate the feasibility of applying DP to FL in the context of Intrusion Detection Systems (IDS).
2. Adapt the well-known ToN-IoT dataset to provide different partitions for different FL clients considering non-iid data distributions.
3. Test different aggregation algorithms and the impact of different DP mechanisms in them.
4. Analyze and compare quantitatively diverse DP mechanisms, based on some aspects such as computational performance, using different privacy-factor values and FL settings.

4.20.4 Demonstration Example

The architecture of the demonstrator is depicted in Figure 39. The two main entities in the scenario are described below:

- Client: represents the end device where data is held, local training is performed, and DP mechanisms are applied. Since the FL process starts, it is able to re-train the global model sent by the aggregator with its local data and provide a private model update. This task is done for every FL round.
- Aggregator: represents the central entity where the aggregation of model updates is performed, based on some aggregation algorithm such as FedAvg or FedPlus. Once all updates are merged into a single model, it is sent back to clients to continue the process.

Before launching the system, the aggregator is configured with the following parameters:

- Number of FL rounds that it will execute. Once this number is reached, it will terminate the process and send the final model to clients.
- Number of local epochs to be executed by clients.

- Learning rate for the ML algorithm. This can be adjusted to tune the speed at which the model learns.

In order to evaluate the impact of DP to the process, the following DP perturbation mechanisms have been used:

- Truncated Laplacian Mechanism [geng2018private]
- Uniform. This is a special case of Truncated Laplacian Mechanism.
- Gaussian [balle2018mechanism]
- Analytic Gaussian Mechanism [balle2018mechanism]
- Bounded Domain Laplace Mechanism [nolohan2018bounded]
- Bounded Laplace Noise Mechanism [dagan2020noise]

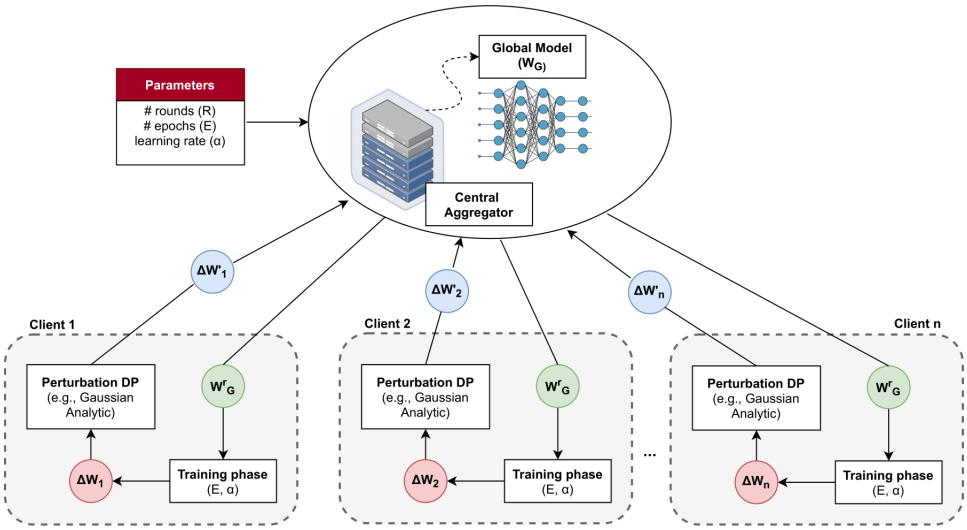

Figure 39 - Proposed Federated Learning with Differential Privacy demonstration architecture

Regarding the ML model, a simple logistic regression classifier has been used along with four partitions of the ToN-IoT dataset [alsaedi2020iot], resulting in four FL clients. Each partition has samples for different attacks or labels including XSS, Injection, Password or Scanning, as well as benign traffic samples.

For three different DP perturbation mechanism and aggregation algorithm, the accuracy of the model is measured in every FL round. In Figure 40, the results using FedAvg algorithm are shown. On the other side, in Figure 41, the same results but using FedPlus instead of FedAvg can be consulted. As a reminder, a smaller epsilon value means that the output of the DP algorithm is more obfuscated or private, then, accuracy when using lower epsilon values should be lower since processed data is farer from the original. However, at it can be seen in the graphs, for almost all DP perturbation mechanisms, there is no big difference between the accuracy reached by a mechanism configured to use high epsilon values and the ones configured to use a smaller one.

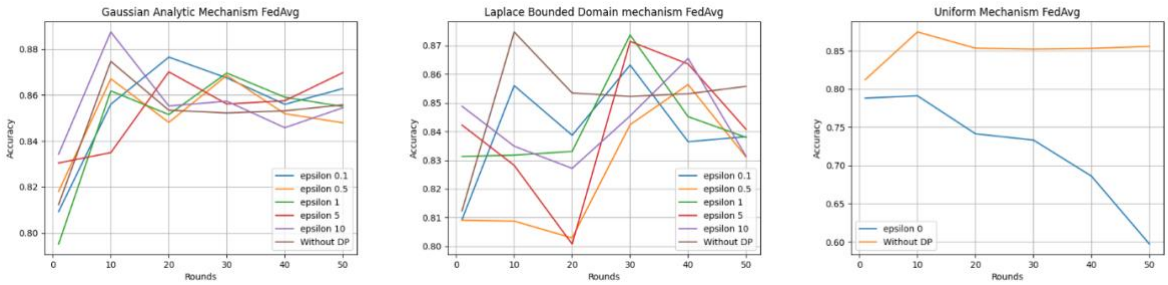

Figure 40 - FedAvg accuracy evolution for every perturbation mechanism.

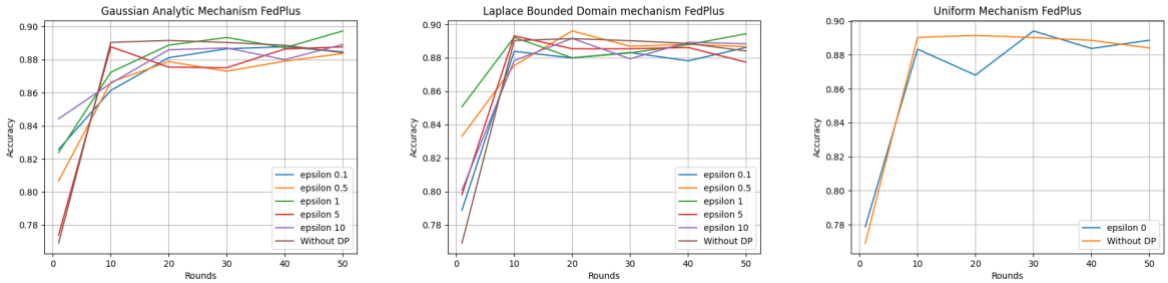

Figure 41 - FedPlus Accuracy evolution for every perturbation mechanism.

In addition, in Figure 42, a box-plot graph which summarizes the maximum, minimum and average execution times for every DP perturbation mechanism can also be consulted. The graph is the result of 10 iterations per mechanism with the same epsilon (epsilon=1), excluding Uniform mechanism since it can only be configured with epsilon=0.

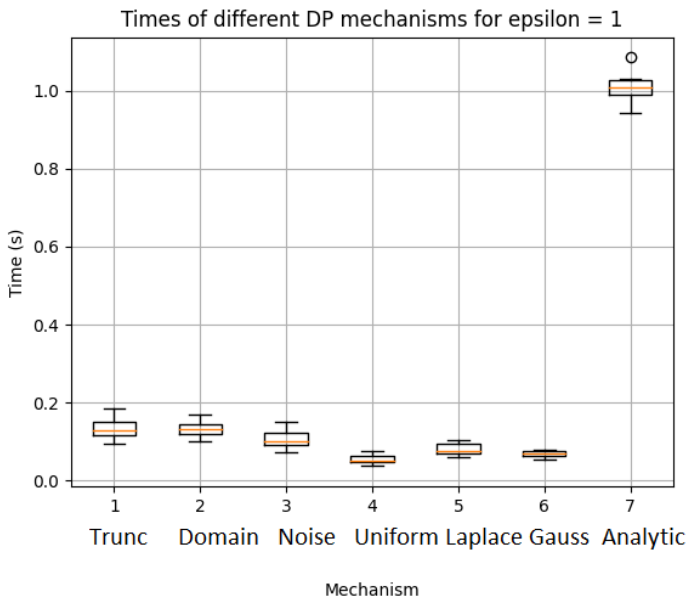

Figure 42 - DP execution times per mechanism

4.20.5 Future work

As future work, we will focus on designing an enhanced system that allows different clients to have different privacy requirements, as well as the use of gradient compression techniques to reduce the latency in scenarios with network limitations, such as IIoT environments.

4.21 PLEAK Differential Privacy Analysers

4.21.1 Overview

The web-based tool Pleak (pleak.io) has been previously described in Deliverable 3.15. Pleak allows to model and analyse business processes specified in privacy-enhanced Business Process Model and Notation (PE-BPMN). It supports several different kinds of privacy leakage analysis of PE-BPMN models, some of which have been described in Deliverable 3.15. In this section, we give an overview of the differential privacy analysis of Pleak.

Using Pleak starts with a BPMN model that can either be created in Pleak or imported if it already exists. Let us assume that a model represents computation of some output data from some input data. Such a model can be analysed using differential privacy analysers. The goal of these analysers is to recommend a certain way of adding noise to the output so that a certain level of input data privacy would be preserved after disclosing the output.

Let us first provide the main intuition behind differential privacy. There is a dataset with private data, which is used as a part of the BPMN model. The BPMN model describes a process that computes some output from this input data, and releases the output to a party whom we consider as an attacker who is trying to learn something about the input data. Even if the observed output is something very general like an aggregated statistic (min, max, average), it may still leak something interesting about the input data if the attacker already has some additional knowledge a priori. For example, if the attacker already knows the salaries of all users except Alice, then the average salary will leak precisely the salary of Alice. Such leakage can be prevented by adding noise to the output.

The definition of differential privacy is based on comparing two settings: when the output is computed from the input *with* the private data of Alice, and *without* private data of Alice. It assumes a probabilistic output. If the probability of getting the same output in both cases is “nearly” the same (quantified by a privacy parameter ε), then Alice can feel safe since the output does not depend “much” on her private data. This property must hold not just for Alice, but for any individual whose data is in the dataset. Since the output of a BPMN process is typically deterministic, differential privacy can be achieved by adding a certain amount of random noise to the released output, making it probabilistic.

Pleak provides two analysers related to differential privacy.

- *Sensitivity analyser* quantifies how much the output of the process depends on the input of the process. Sensitivity can in turn be used to suggest the way(s) of adding noise to the output to obtain a certain level of differential privacy (for the given ε).
- *Guessing advantage analyser* is a user-friendly extension of the sensitivity analyser. The problem of differential privacy is that the parameter ε that determines the level of privacy is difficult to interpret. It may range from 0 to infinity, and its goodness depends on the query and the data. The guessing advantage analyser allows the user to state a certain attacker goal, i.e., which data the attacker would like to guess and with which precision (in the case of numeric data). Instead of ε, the user comes up with a number in the range 0...1, representing the desired upper bound on the probability of attacker’s success in achieving the goal. Pleak then suggests the way(s) of adding noise to the output to achieve that bound.

In practice, it may happen that, e.g., the user has chosen 1% as an acceptable bound on the probability of attacker's success, but relaxing it to 1.1% would result in a sudden decrease in noise magnitude. If the user knew it, he would probably agree to accept 1.1% instead of 1%. Ideally, the user would like to see an interactive plot representing the relation between the privacy level and the noise level. This visual interface has been developed as a complementary component to Pleak.

4.21.2 Main asset improvements since D3.13

The first description of this asset is presented in this deliverable.

4.21.3 Research challenges addressed

The standard definition of differential privacy quantifies indistinguishability of the cases "with an individual's data" and "without an individual's data". If we are interested in privacy of some particular attribute (or a certain combination of attributes, such as geolocation data), it can be more reasonable to consider a certain change in a particular attribute instead of considering the case where an individual's record has been removed entirely. For example, we can agree to leak the city that the tourist has visited, but conceal the particular district. [dersens2020] proposes some methods to achieve such kind of differential privacy (also called d-privacy in the literature).

One problem of differential privacy is that the parameter ε that determines the level of privacy is difficult to interpret. It may range from 0 to infinity, and its goodness depends on the query and the data. [ga2022] proposes some methods to relate differential privacy (and also d-privacy) to attacker's success in guessing or approximating sensitive attributes.

In order to make differential privacy accessible to non-expert users who do not have deep understanding of the mathematical definitions behind it, we need a user interface that would hide the most complicated part under the hood. In the following publication, we assume that the potential user is a policy authority who is able to assess the following:

- How trusted is the party that receives the output?
- How sensitive is the input data?
- How high risk of data leakage can be tolerated?

These three parameters can be provided in a human-friendly format ('low', 'medium', 'high'), which can be calibrated in advance by a knowledgeable expert. These parameters are more generic than the ε of differential privacy whose interpretation is application-specific and cannot be set up once for all datasets and all queries in advance. The tool provides a visual interface that allows to examine the trade-off between the risk of data release, and the expected amount of noise that should be added.

[prisms2021] Publication: Pankova, Alisa; St. John, Mark F.; Denker, Grit; Laud, Peeter; Martiny, Karsten; Pavlovic, Dusko (2021). Decision Support for Sharing Data Using Differential Privacy. 18th IEEE Symposium on Visualization for Cyber Security, VizSec 2021, online, 27.10.2021. IEEE, 1-10. https://doi.org/10.1109/VizSec53666.2021.00008

It is relatively easy to come up with a differential privacy mechanism for a process with a single aggregated output. What if the analyst prefers to work with the input directly and does not know in advance exactly which statistics he is going to compute? In the following publication, differential privacy was applied in the context of process mining, where the noise has been added directly to the input. In the considered scenario, not only the attributes, but also the number of records in the dataset were considered private, so noise addition mechanisms had to be combined with under- and oversampling of records.

[pm2021] Elkoumy, Gamal; Pankova, Alisa; Dumas, Marlon (2021). Mine Me but Don't Single Me Out: Differentially Private Event Logs for Process Mining. 3rd International Conference on Process Mining (ICPM). IEEE Computer Society, 80−87. https://doi.org/10.1109/ICPM53251.2021.9576852.

4.21.4 Demonstration example

Our demo extends the demonstration example of Sharemind HI considered in Section 4.14.4 of Deliverable 3.13. In this scenario, an analyst observes a counting histogram of tourists visiting some foreign countries in a particular time interval, grouped by month. The counts are computed from private geolocation data using a trusted execution environment of Sharemind HI. The analyst does not see the actual input data, and may only observe some aggregated data, i.e., the counts. The technology used by Sharemind HI guarantees that the input data is not just hidden from the analyst visually, but it does not even exist (in plaintext) in the machine that computes the counting histogram (Figure 43).

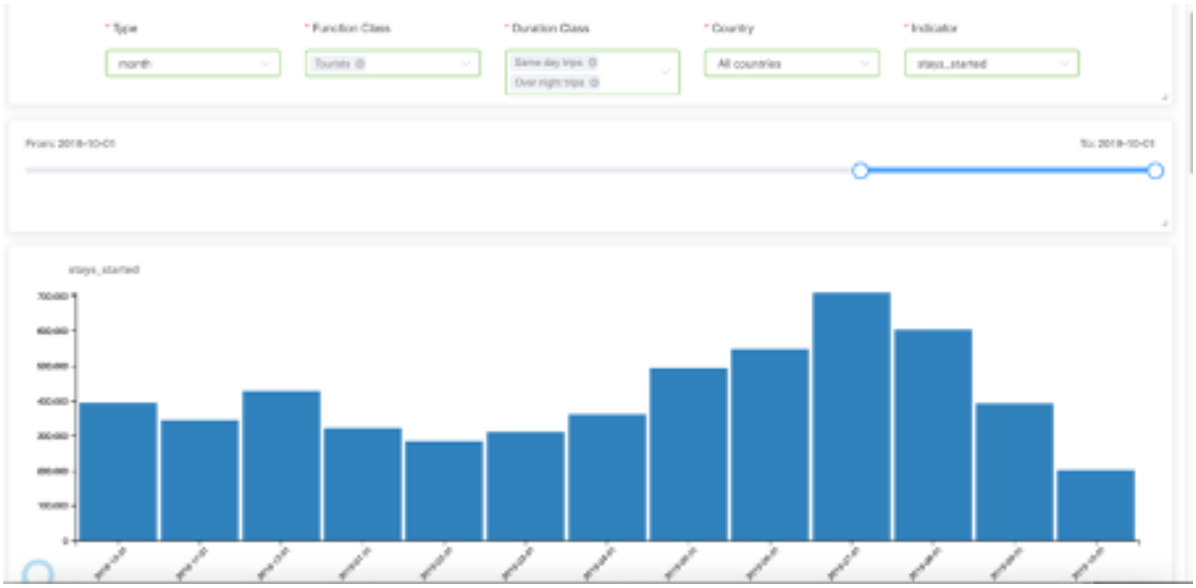

Figure 43 - Histogram of tourists visits (data without differential privacy)

If Alice is one of the tourists who has contributed her data, can she be sure that her privacy is preserved? Of course, aggregated data is much less informative than the input data, but it may still leak some sensitive information. For example, in an extreme case where no tourists have visited a country in April, the analyst learns that Alice has not visited this country in April a well. To prevent leakage of data via the final output, we can add some random noise to it (Figure 44).

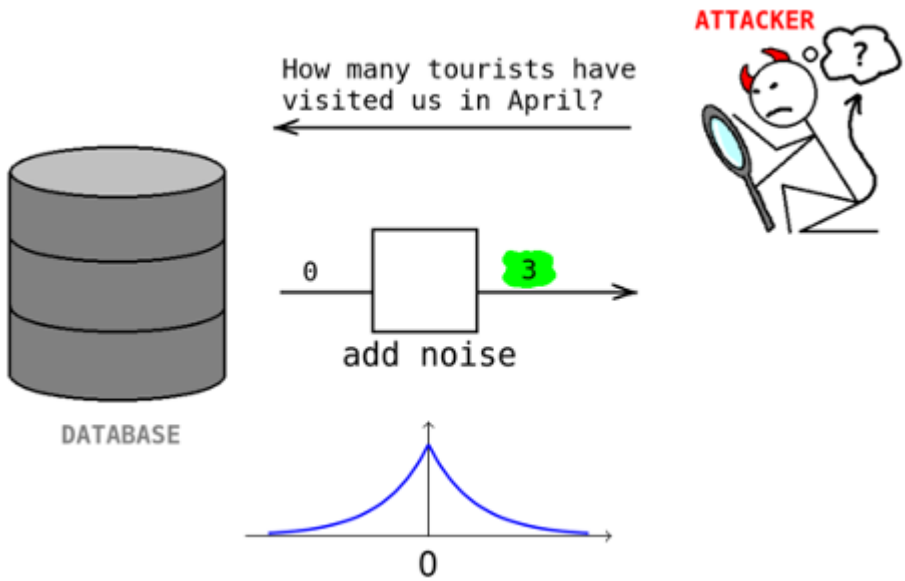

Figure 44 - Adding noise to a dataset

While more noise means more privacy, a too large noise may lead to poor utility of the released data. How much noise would be enough for our particular use case? We can ask Pleak differential privacy analysis tool.

Let us consider a very simple model, consisting of two data objects, "trips" and "counts", and a single task "count seasonal travels" (Figure 45).

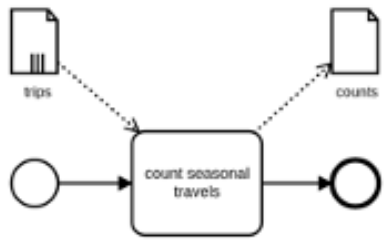

Figure 45 - Simple PLEAK model for counting seasonal travels

The data object "trips" contains the input data (Figure 46). It is a simple table, which says where and when a certain trip was made. It also contains some more attributes that are used for filtering, and whose exact meaning is not relevant for this demo and is explained in the Sharemind HI demo of Section 4.14.4 of Deliverable 3.13.

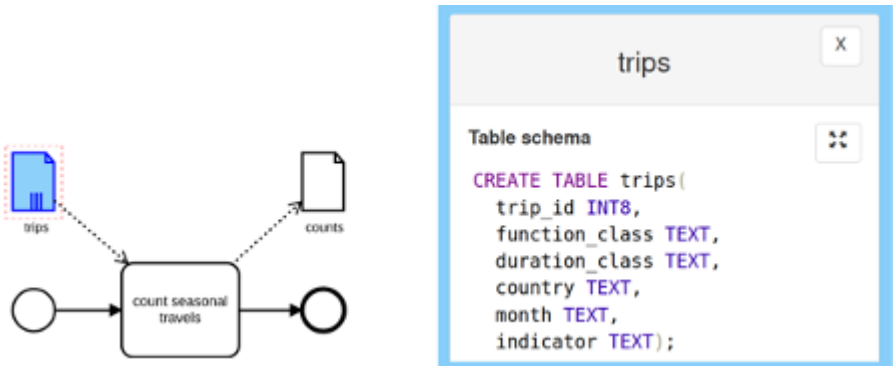

Figure 46 - Table schema for the trips data object

The task "count seasonal travels" describes the query whose output the analyst is going to see (Figure 47). Shortly speaking, the query computes a counting histogram of trips grouped by month. It also applies some additional filters to the input data. In Pleak, the data object “task” is highlighted in green since it is an input. The data object “counts” is highlighted in red since it is an output.

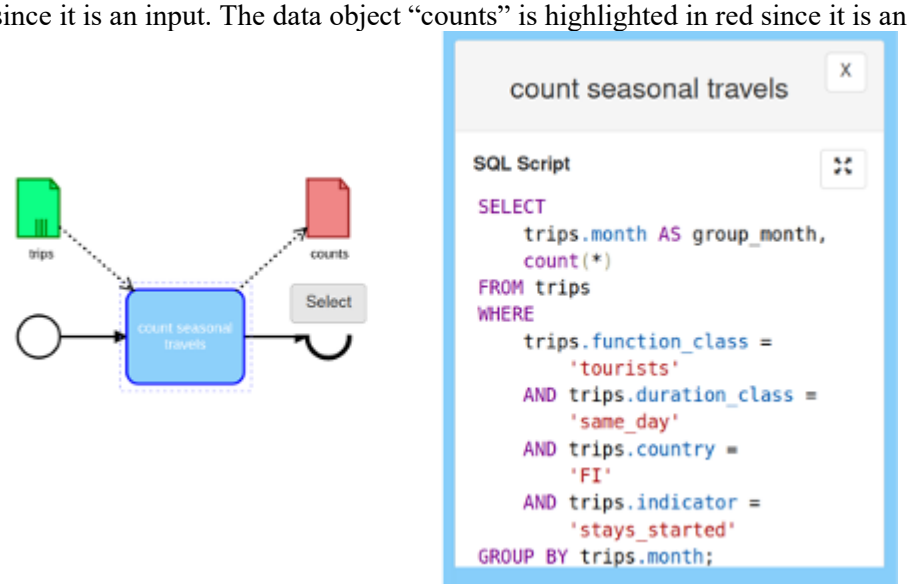

Figure 47 - Task “count seasonal travels”

Let us estimate how much the output of this task may leak to an analyst. In order to analyse the leakage, we need to specify the attacker's goal. Let us assume that the attacker wants to learn where and when the target user has been. The goal can be specified in form of a database query (Figure 48).

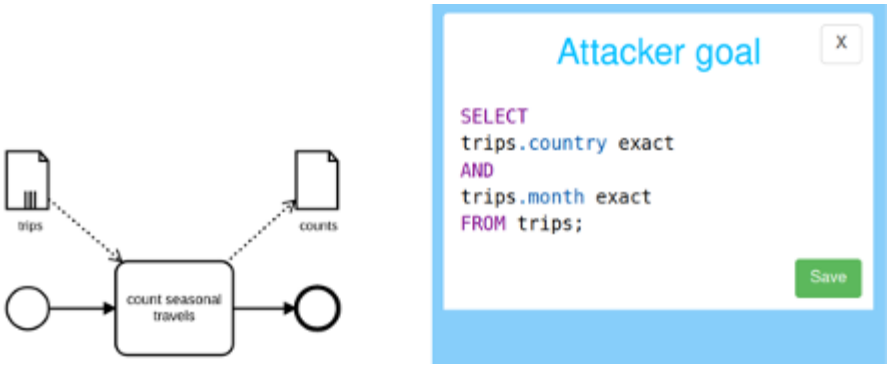

Figure 48 - Attacker goal specified as a database query.

After the attacker goal have been set, the tool can propose the amount of noise that needs to be added to achieve a certain bound on attacker's success in achieving this goal. The user is interested in balancing the degree of attacker's success and the amount of noise.

In differentially private mechanisms, the level of privacy is quantified by a special parameter ε. Pleak analysis tools allow to provide the ε directly as an input. The problem is that ε is a positive value ranging from 0 to infinity. While the user knows that smaller ε means more privacy, it is difficult to tell in advance how small the ε should be to provide enough privacy. For that reason, Pleak also allows to measure privacy in terms of risk of data leakage, which depends on the particular dataset and the attacker goal. The risk is easier to interpret than the ε. The user can interact with the tool to balance the risk and the noise.

The user starts from fixing several input parameters. Each parameter has some quantitative interpretation, ranging from 0 to 1. Instead of choosing a numerical value that can be hard to interpret, the user will only see a list of pre-defined choices, which have been set by an expert once in advance (Figure 49).

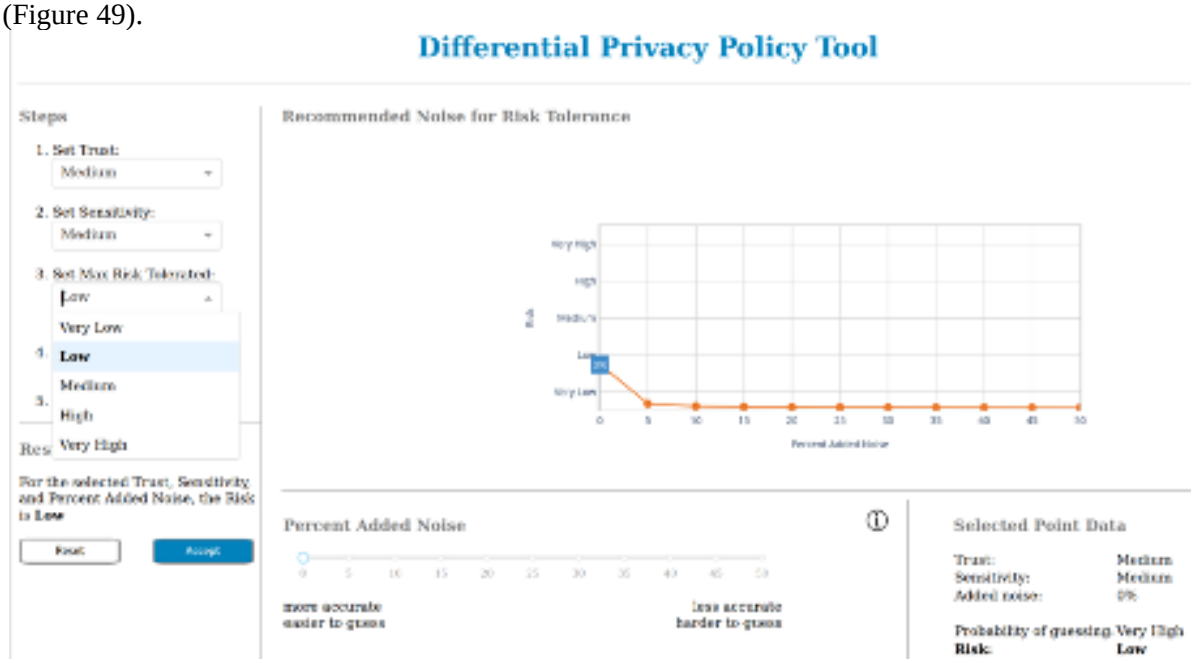

Figure 49 - Setting input parameters

First of all, the user chooses how much the analyst can be trusted. Let us try to be safe and assume that the trust is *very low*. The next thing to set up is the data sensitivity. Let it be v*ery high*. Intuitively, if the trust was high, or the data sensitivity was low, we would not need any noise. The next thing that the user can fix is the maximum tolerated risk. Let it be *very low*. We can see how the plot on the right-hand side has changed in Figure 50. The gray line shows that, in order to achieve very low risk, it is recommended to add 5.6% of noise, i.e., if the true count was 100, then 5.6% would mean that the noisy count stays somewhere between 94.4 and 105.6.

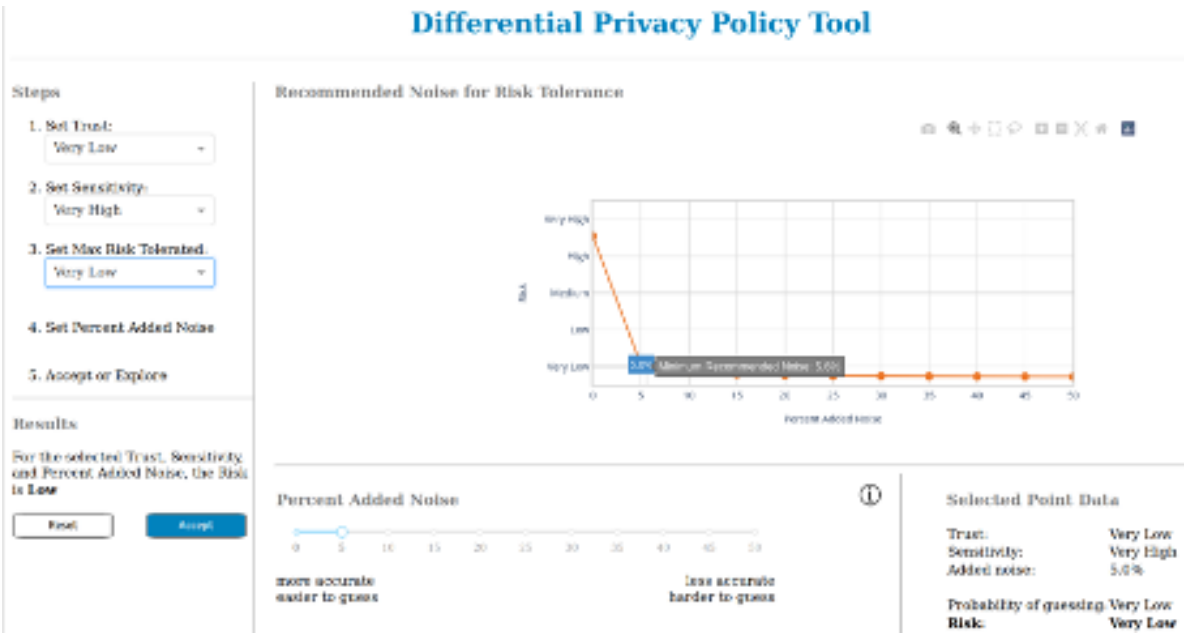

Figure 50 - Changes in the recommended noise for the specified risk tolerance

From the orange plot, we can see that reducing the noise level slightly does not increase the risk significantly. If we add 5% of noise, the risk is still closer to "very low". If we add 3%, it would already grow too much and be closer to "low". Let us decide to allow 5% of noise and accept. Pleak can now internally compute the ε that corresponds to this noise level, and come up with a description of noise distribution that it sufficient for ε-differential privacy. Figure 51 shows what our example histogram (Figure 51) looks like after sampling the noise according to the tool's suggestions.

Figure 51 - Example histogrm after sampling and adding the noise

Since the noise is sampled randomly, and it can be both positive and negative, some histogram bars may increase, while some may decrease. Each such difference contributes to the error. Since we have chosen noise magnitude 5%, the difference between the true bar height and the noisy bar height does not exceed 5% of the true bar height.

4.21.5 Future Work

The next logical step in Pleak development would be integration of the visual decision support tool of [prisms2021] with the existing differential privacy analysers of Pleak, which currently compute the noise magnitude for a single numerical level of the attacker's success (without demonstrating a plot) and provide a quite technical output. This addition would make adoption of differential privacy more accessible for non-technical people.

The differential privacy analysers of Pleak are currently primarily based on [dersens2020]. A useful thing to consider is integrating into them the latest results of [ga2022], which provide more tight bounds on the required noise magnitude for the case where an attacker is trying to approximate a numeric attribute. While this would not make any difference from the user perspective, it would provide better trade-off between privacy and utility.

Currently, Pleak can only recommend to add noise to numeric outputs, which is primarily designed for statistical aggregation queries. In a BPMN model, it may happen that one party sends to another party a non-aggregated piece of raw private data. It would be interesting to enhance Pleak with mechanisms recommending addition of a certain amount of noise to non-aggregated data, possibly applying some ideas from [pm2021].

Currently, Sharemind and PLEAK are two completely independent tools. Ideally, Pleak could be connected directly to Sharemind, and instead of just providing a recommendation of noise distribution, it could actually execute noise sampling inside the trusted execution environment of Sharemind, and apply it to the output before releasing it to the result parties.

5 Task 3.2 Outcomes

5.1.1 Overview

Task 3.2 "Research and Integration on Cybersecurity Enablers and underlying Technologies" of CyberSec4Europe had as one of its main goals to identify research challenges, requirements and approaches in privacy preserving tools. In the last update the assets are deployed using the smart-campus ecosystem to leverage three main scenarios: CCTV surveillance in the smart-campus, Identity Management and Service usage in Smart Campus, and Geolocation Service in Smart Campus. In these scenarios we integrated a total of 20 assets, showing how the enablers can help to improve privacy protection on general systems. More information in the official task T3.2 github page https://github.com/cs4ewp3/wp3/tree/main/3.2

5.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
PTASC				✓	✓
ARGUS					
GDPR compliant user experience					

Interoperability and cross-border compliance	✓				
Edge Privacy (UMA)				✓	
Password-less AuthN				✓	✓
SS-PP-IdM					
CryptoVault					
GENERAL_D				✓	✓
pp-FL					
Sharemind				✓	
PLEAK DP analysers					

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
PTASC							
ARGUS	✓		✓				
GDPR compliant user experience	✓						✓
Interoperability and cross-border compliance							
Edge Privacy (UMA)							✓
Password-less AuthN			✓				
SS-PP-IdM	✓						✓
CryptoVault				✓			
GENERAL_D	✓		✓				✓
pp-FL	✓	✓				✓	
Sharemind					✓	✓	✓
PLEAK DP analysers	✓			✓	✓		✓

6 Task 3.3 Outcomes

6.1.1 Overview

Task 3.3 “Software Development Lifecycle” of CyberSec4Europe had as one of its main goals to identify research challenges, requirements and approaches in all stages of the lifecycle of software. The analysis of the research challenges lead to the identification of promising tool-supported approaches to be developed in order to address the challenges. A total of 13 tools to support different activities in the lifecycle of software was developed and demonstrated on a common smart-cities scenario as illustrated in Figure 52. More information in the official task T3.3 github page https://github.com/cs4ewp3/wp3/tree/main/3.3

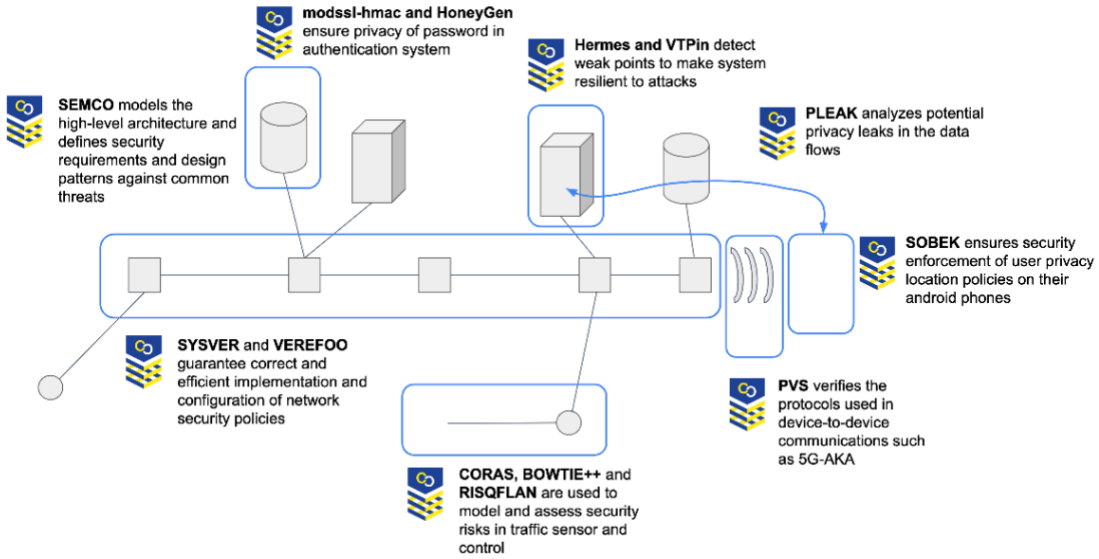

Figure 52 - Task 3.3 tools

6.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
SOBEK					✓
HERMES					✓
RisQFlan					
Pleak					
Verefoo					
CORAS					
PVS			✓		
BowTiePlus					
HoneyGen					
VTPin					
Modssl-hmac					
SEMCO					

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
SOBEK			✓				✓
HERMES			✓				
RisQFlan			✓				
Pleak			✓				✓
Verefoo			✓				
CORAS			✓				
PVS			✓				
BowTiePlus			✓				
HoneyGen			✓				
VTPin			✓				
Modssl-hmac			✓				
SEMCO			✓				

7 Task 3.4 Outcomes

7.1.1 Overview

Task 3.4 “Security Intelligence” of CyberSec4Europe is focused on defining the requirements and mechanisms to share digital evidence between the different expert systems, providing solutions to allow interoperability, either through the unification of languages, formats and interfaces, or through trusted intermediate translators systems respecting the privacy, business requirements and the regulations of the different countries. We propose a general framework of interaction and cooperation with threat intelligence services and provided three specific use cases where these services can cooperate. The partners contributed by creating tangible software assets, by integrating them into joint proof-of-concepts and illustrating the practical feasibility of a modular cybersecurity platform able to provide key information about the status of a system to monitor. More information in the official task T3.3 github page https://github.com/cs4ewp3/wp3/tree/main/3.4

7.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
Briareos (C3P)					✓
EBIDS (CNR)					✓
ENIDS (FBK)					✓
HADES (UMA)					✓

IntelFrame (DTU)					✓
JUDAS (UMA)					✓
NetGen (Polito)					✓
PP-CTI (UMU)					
Reliable-CTI (UMU)					
RoCe (UNITN)					✓
TATIS (KUL)					✓
TIE (ATOS)	✓				✓
UASD (CNR)					

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
Briareos (C3P)							
EBIDS (CNR)		✓				✓	
ENIDS (FBK)		✓					
HADES (UMA)							
IntelFrame (DTU)		✓				✓	
JUDAS (UMA)							
NetGen (Polito)		✓				✓	
PP-CTI (UMU)	✓						✓
Reliable-CTI (UMU)							
RoCe (UNITN)							
TATIS (KUL)	✓	✓					
TIE (ATOS)	✓						✓
UASD (CNR)		✓					

7.1.3 Privacy Preserving Cyber Threat Information Sharing leveraging FL-based intrusion detection

This asset has been added to task T3.4 recently, and it is the first time that it is reported as part of T3.4, that is the reason why it is summarized in this deliverable.

As a consequence of the briskly digital transformation and technological developments that are happening in recent years in many sectors such as financial, industrial or healthcare, organizations are starting to collaborate closely to build effective countermeasures to detect and mitigate new evolving cyber threats and risks. In order to address this, CYTILIS asset provides an architectural proposal to deal with private Cyber Threat and Information (CTI) Sharing and enhanced anomaly detection models based on Federated Learning (FL). In this line, a trusted privacy-preserving ecosystem of threat intelligence platforms based on MISP is proposed, aimed to exchange and process information automatically and in an auditable manner. Alongside, a Federated Learning scenario is deployed to enrich machine learning models used for anomaly detection. Once obtained, these models will be shared through the proposed ecosystem as CTI events. This whole brings the opportunity to build more trusted, resilient and secured organizations and institutions.

For the demonstration of the privacy-preserving subsystem of this asset, a dataset simulating sensitive data has been created to simulate a CTI interchange where values can disclosure a personal identity. The inputs which this subsystem will receive are events whose sensitive values are visible together with policy files that specify which attributes to hide and with which techniques to do so. These policies are designed by the data holder and must be in accordance with the input dataset and the level of privacy to be achieved. The result of this process will be a privatized event published on the organization's MISP instance, which has been transformed according to the policies created by the user. To obfuscate the selected attributes, several techniques related to data mining field have been configured to be applied, such as suppression, generalization, K-anonymity, L-diversity or T-closeness. These give an automated way of suppressing or generalize certain values that could be detrimental to individuals in data if they were unveiled.

On the other side, for the Federated Learning subsystem, a basic scenario with one aggregator and a pair of clients has been set up. To test this subsystem, the ToN-IoT dataset is leveraged, as well as producer and consumer modules have been created. On one hand, producer will create and update events based on samples of the dataset to the MISP instance. Then, consumer will receive these in real-time and, once an instance threshold is reached, will create a client and a FL process along with the aggregator and the other client will be started. Once the final model is obtained, consumer will upload it back to the MISP instance as an event with an attachment.

The overall architecture for the asset, differentiating both subsystems, can be consulted in Figure 53.

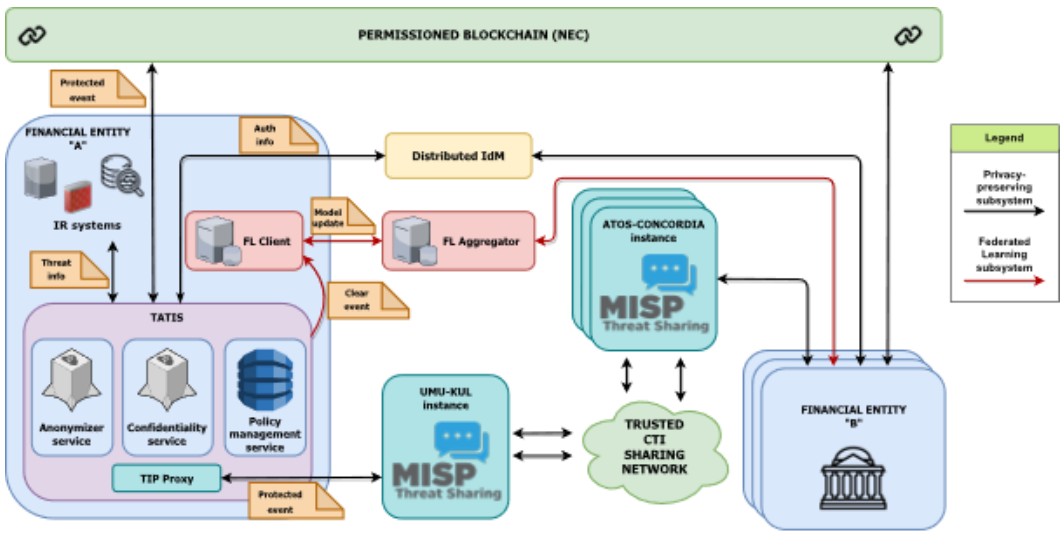

Figure 53 - CYTILIS architecture proposal

Below, in Figure 54, the sequence diagram which describes the secure exchange of CTI events enabled by the privacy-preserving subsystem can be also consulted.

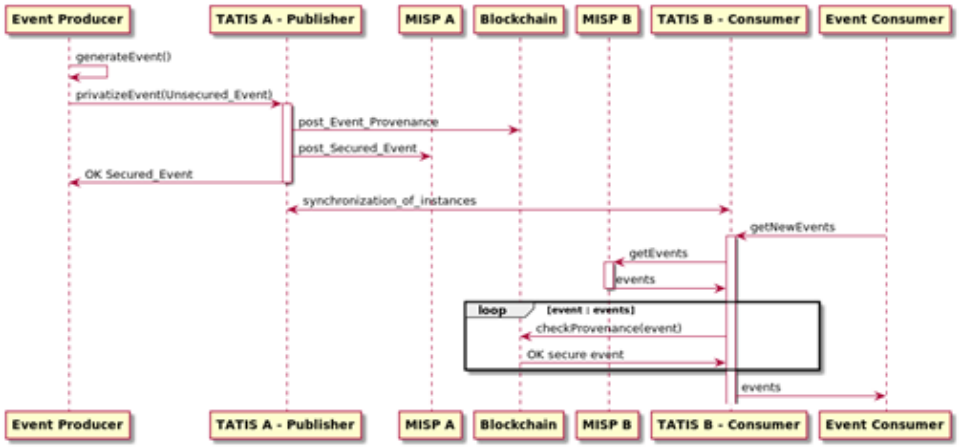

Figure 54 - Secure exchange of CTI events between organizations

In the same way, diagrams contained in Figure 55 and Figure 56 depict the workflow of both producer and consumer modules in the FL subsystem.

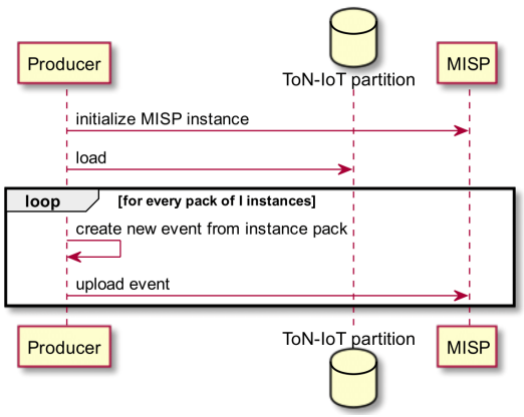

Figure 55 - Producer workflow within the FL subsystem

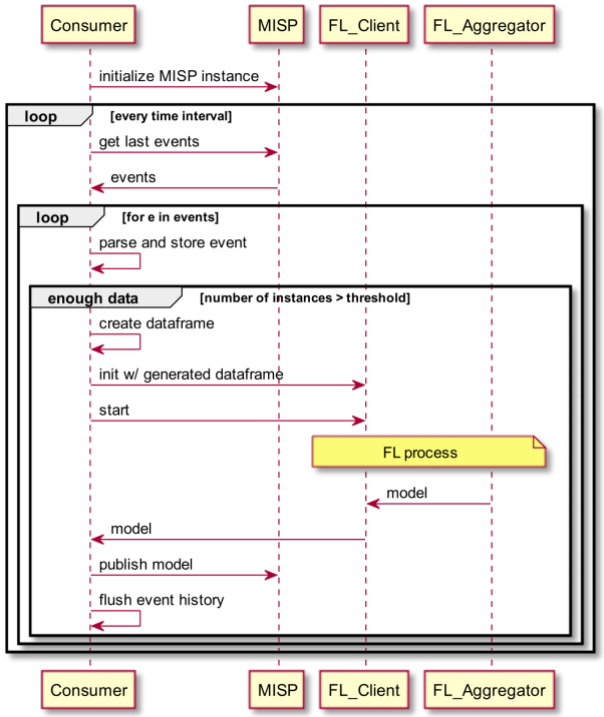

Figure 56 - Consumer workflow within the FL subsystem

The asset is implemented and it is now in the process of preparing a paper targeting a high-impact research journal.

8 Task 3.5 Outcomes

8.1.1 Overview

Task 3.5 “Adaptive Security” of CyberSec4Europe explores the development of flexible security solutions that can quickly adapt security controls in response to security changes such as new attacks or changes in security requirements, improving the modelling and analysis of dynamic systems and providing provide scalable architectures supporting security situation computation and risk assessment, and also selection and deployment of security controls that could satisfy security requirements and policies, also enabling awareness of the current system status. More information in the official task T3.5 github page https://github.com/cs4ewp3/wp3/tree/main/3.5

8.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
SPARTA (KUL)					
Adaptive Authentication (UCD)					
Situation-driven risk assessment and security enforcement framework (UPRC + IRIT)					✓
SYSVER (CNR)					

AIRE (ATOS)					

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
SPARTA (KUL)			✓				✓
Adaptive Authentication (UCD)			✓				✓
Situation-driven risk assessment and security enforcement framework (UPRC + IRIT)							✓
SYSVER (CNR)							✓
AIRE (ATOS)			✓				✓

9 Task 3.6 Outcomes

9.1.1 Overview

Task 3.6 “Usable Security (Human-centred Cybersecurity)” of CyberSec4Europe has the motivation of empowering users to make sensible security choices. We have researched methods on how to advise or convince users on different security solutions such as authentication methods or privacy settings, and how to make visible the underlying structures such as security policies or cryptographic protocols. More information in the official task T3.6 github page https://github.com/cs4ewp3/wp3/tree/main/3.6

9.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
Privacy-preserving IdM					
DPIA template					
HATCH		✓			

CyberSecurity Awareness Quiz		✓			
LiSRA					
Adaptive Authentication					
EEVEHAC					
HAMSTERS		✓			
SYSVER					
AuthGuide					
LEECH		✓			

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
Privacy-preserving IdM	✓						✓
DPIA template	✓						✓
HATCH	✓		✓				
CyberSecurity Awareness Quiz	✓						
LiSRA							
Adaptive Authentication							
EEVEHAC	✓						
HAMSTERS	✓			✓			
SYSVER			✓				
AuthGuide			✓				
LEECH							✓

10 Task 3.7 Outcomes

10.1.1 Overview

Task 3.7 focused on the design of best practices for innovative and GDPR compliant user experience and the investigation of the interoperability issues in identity technologies (eIDAS and GDPR). In the task two assets were developed.

Guidelines for GDPR Compliant User Experience is an asset/deliverable that was produced as D3.6 in the CyberSec4Europe project. As its name implies, it is a collection of guidelines, best practices and recommendations for achieving GDPR compliance. It is combined from two parts. The first is the

Guidelines for General Data Protection Regulation (GDPR) which present the regulation's requirements through the GDPR principles. The second part of the enabler is the Data Protection Impact Assessment (DPIA) template. As the name suggests, this part of the enabler can be used to help guide the user through the process of doing a DPIA and serve as documentation for the performed analysis. This asset was already a part of the demosntration in D3.13.

Analysis of interoperability and cross-border compliance issues builds on the work already done by the European Commission and ENISA as a basis for monitoring the current eIDAS network and proposing eIDAS 2. Additionally, it is an extension to the research described in the previous asset, where, in the process of creating the guidelines for GDPR, we have noticed differences between the EU Member States that could potentially cause problems in situations involving more than one Member States. The asset tries to identify areas of difference in how eIDAS and GDPR are applied in the EU Member States, that could couse potential interoperability and cross-border compliance issues.
This asset was a part of the demonstration in D3.13 and this D3.20 document.

Both assets and their byproducts are described in the project's T3.7 GitHub page: https://github.com/cs4ewp3/wp3/tree/main/3.7.

10.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
Guidelines for GDPR compliant user experience					
Analysis of interoperability and cross-border compliance issues	✓				

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
Guidelines for GDPR compliant user experience	✓						✓
Analysis of interoperability and cross-border compliance issues							

11 Task 3.8 Outcomes

11.1.1 Overview

The main result of Task 3.8 - Conformity, Validation and Certification is SURFACE (Support Framework for Certification) – an integrated approach that can be used to support certification and recertification. Present SURFACE – a support framework for certification which integrates and combines steps and processes from the ARMOUR methodology, the ECSO meta-scheme, the European Cybersecurity Candidate Scheme and the continuous monitoring process from NIST 800-137. The process in SURFACE is divided into phases inspired by the ARMOUR methodology.

CSA is a web-based asset management and certification assistant tool. The certification workflow has been built based on SURFACE. The certification dimension adds another layer to asset management. It allows the manufacturer or vendor to carry out incremental certification based on sub-assets. More information in the official task T3.8 github page https://github.com/cs4ewp3/wp3/tree/main/3.8

11.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
CSA			✓		✓

Asset\Research Challenge	Trust-Building Blocks				Disruptive Emerging Development		
-	Holistic Data Protection	AI-based Security	Systems Security & Security Lifetime Management	Secure Architectures for Next Generation Communication	Secure Quantum Technologies	Secure AI Systems	Personalized Privacy Protection
CSA			✓				

12 Task 3.10 Outcomes

12.1.1 Overview

In task 3.10 – “Impact on Society” one of the main goals is to propose a conceptual framework for the monitoring and evaluation of a cybersecurity awareness program. The framework provides guidelines and practical advice on “what to do in each phase” of a cybersecurity awareness program. Moreover, it also answers “what to expect in each phase” i.e., expected outputs and outcomes, after those guidelines and advice are followed. The earlier information can be helpful for designing an effective awareness program, whereas the latter information will more specifically facilitate monitoring and evaluation of the program. In addition, it provides evaluation criteria of two cybersecurity awareness mechanisms, which are posters (simplest and one of the least interactive media that is known to almost everyone), and serious games (highly interactive media that is gaining popularity). More information in the official task T3.10 github page https://github.com/cs4ewp3/wp3/tree/main/3.10

12.1.2 Cybersecurity Research and Areas Priority

Asset\Research Challenge	Governance and Capacity Building			Trustworthy Ecosystems of Systems	
-	Collaborative Networks	Education & Training	Certification	Secure Platforms of Platforms	Infrastructure Protection
Conceptual Framework & Guidelines		✓			
Criteria for Poster Evaluation		✓			
Criteria for Serious Game Evaluation		✓			

13 Conclusions

This document has described the results and experiments carried out in the scope of T3.2 to evaluate the benefits, feasibility and performance of the assets. Several T3.2 assets have been demonstrated in the same context of a smart-campus framework, whereby each asset aims at strengthening the overall privacy and security of the campus in from different angles. The smart campus global demonstrator has been divided in three main different scenarios to demonstrate a wide range of security and privacy capabilities, namely:

(i) Video surveillances use case to manage the feeds, as it is a common scenario in smart buildings that allows a response team to continuously monitor the campus, such as monitoring fires or accidents.

(ii) Services with different access controls - Identity management for user authentication in various heterogeneous smart-campus services, including services for direct interaction with the University (e.g., enrollment in courses or activities).

(iii) A smart-campus geolocation service is used to detect trends in the movement of people, deal with public transport planning, urban planning, creating safer and friendlier areas for residents.

The document has described the smart-campus scenario, and how each asset is mapped or contextualized in the scenario. Each T3.2 asset is described and evaluated in detail, showing their benefits to improve security and privacy in broad range of cases.

Furthermore, the document has outlined the assets reported by other WP3 tasks, linking them to the official github page where the reader can find additional information.

Finally, this deliverable has shown a categorization of the assets in the cybersecurity research categories, showing the priority areas of the research that has been addressed in WP3.

14 References

[cs4e-github] CS4E-WP3 Github page, https://github.com/cs4ewp3

[sari2017study] Sari, Marti Widya, Prahenusa Wahyu Ciptadi, and R. Hafid Hardyanto. "Study of smart campus development using internet of things technology." IOP Conference Series: Materials Science and Engineering. Vol. 190. No. 1. IOP Publishing, 2017.

[mineraud2016gap] Mineraud, Julien, et al. "A gap analysis of Internet-of-Things platforms." Computer Communications 89 (2016): 5-16.

[ngu2016iot]Ngu, Anne H., et al. "IoT middleware: A survey on issues and enabling technologies." IEEE Internet of Things Journal 4.1 (2016): 1-20.

[D3.2] Stephan Krenn, D3.2 –"Cross Sectoral Cybersecurity Building Blocks"

[D.13] Joao Resende, D3.13-" Updated version of enablers and components "

[sousa2021provisioning] Sousa, P. R., Magalhães, L., Resende, J. S., Martins, R., & Antunes, L. (2021). Provisioning, Authentication and Secure Communications for IoT Devices on FIWARE. Sensors, 21(17), 5898.

[OASIS13] OASIS. eXtensible Access Control Markup Language (XACML) Version 3.0. http://docs.oasis-open.org/645xacml/3.0/xacml-3.0-core-spec-os-en.html, 2013

[Hardt12] D. Hardt et al., "The oauth 2.0 authorization framework," tech.rep., RFC 6749, October 2012

[MBGFSMSPS20] R. T. Moreno, J. Bernal Bernabe, J. García Rodríguez, T. K. Frederiksen, M. Stausholm, N. Martínez, E. Sakkopoulos, N. Ponte, and A. Skarmeta, "The olympus architecture—oblivious identity management for private user-friendly services," Sensors, vol. 20, no. 3, p. 945, 2020

[BDMCBS20] J. B. Bernabe, M. David, R. T. Moreno, J. P. Cordero, S. Bahloul, and A. Skarmeta, "Aries: Evaluation of a reliable and privacy-preserving european identity management framework," Future Generation Computer Systems, vol. 102, pp. 409 425, 2020.

[MGBS21] R. T. Moreno, J. García-Rodríguez, J. B. Bernabé and A. Skarmeta, "A Trusted Approach for Decentralised and Privacy-Preserving Identity Management," in IEEE Access, vol. 9, pp. 105788-105804, 2021, doi: 10.1109/ACCESS.2021.3099837.

[Dumortier17] Dumortier, J. (2017). Regulation (EU) No 910/2014 on Electronic Identification and Trust Services for Electronic Transactions in the Internal Market (eIDAS Regulation). In EU Regulation of E-Commerce. Edward Elgar Publishing.

[rios2021personal] R. Rios, J. A. Onieva, R. Roman, and J. Lopez, "Personal IoT Privacy Control at the Edge", IEEE Security & Privacy, vol. 20, issue 1, Early Access. Doi: 10.1109/MSEC.2021.3101865

[roberto2021Eng] Roberto Casadei, Mirko Viroli, Giorgio Audrito, Danilo Pianini, Ferruccio Damiani: Engineering collective intelligence at the edge with aggregate processes. Eng. Appl. Artif. Intell. 97: 104081 (2021)

[viroli2013advance] Viroli M., Damiani F., Beal J. (2013) A Calculus of Computational Fields. In: Canal C., Villari M. (eds) Advances in Service-Oriented and Cloud Computing. ESOCC 2013. Communications in Computer and Information Science, vol 393. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-45364-9_11

[dwork2013alg] Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci. 9(3–4), 211–487 (2013)

[suratkar2020crypto] Suratkar, S., Shirole, M., & Bhirud, S. (2020, September). Cryptocurrency Wallet: A Review. In 2020 4th International Conference on Computer, Communication and Signal Processing (ICCCSP) (pp. 1-7). IEEE.

[niko2021sec] N. Lehto, K. Halunen, O. -M. Latvala, A. Karinsalo and J. Salonen, "CryptoVault - A Secure Hardware Wallet for Decentralized Key Management," 2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS), 2021, pp. 1-4, doi: 10.1109/COINS51742.2021.9524133.

[gu2017secure] J. Gu, Z. Hua, Y. Xia, H. Chen, B. Zang, H. Guan, and J. Li, "Secure live migration of SGX enclaves on untrusted cloud," in International Conference on Dependable Systems and Networks, DSN, 2017, pp. 225–236.

[guerreiro2020Tee] J. Guerreiro, R. Moura, and J. N. Silva, "TEEnder: Sgx enclave migration using HSMs", Computers & Security, 2020

[jan2017ano] Jan Camenisch, Manu Drijvers, Anja Lehmann: Anonymous Attestation with Subverted TPMs. CRYPTO (3) 2017: 427-461

[j2018accurate] J. Decouchant, M. Fernandes, M. Volp ¨ et al., "Accurate filtering of privacy-sensitive information in raw genomic data," Journal of Biomedical Informatics, vol. 82, pp. 1–12, 2018.

[HK19] Ulrich Haböck, Stephan Krenn: Breaking and Fixing Anonymous Credentials for the Cloud. CANS 2019: 249-269

[K18] Dominik Koehle: ABC for PrivacyImplementing a Demonstrator Model for Privacy-Preserving Authentication in the Cloud. MSc Thesis, 2018

[KLSS17] Stephan Krenn, Thomas Lorünser, Anja Salzer, Christoph Striecks: Towards Attribute-Based Credentials in the Cloud. CANS 2017: 179-202

[LWK22a] Thomas Lorünser, Florian Wohner, Stephan Krenn: A Privacy-Preserving Auction Platform with Public Verifiability for Smart Manufacturing. ICISSP 2022: 637-647

[LWK22b] Thomas Lorünser, Florian Wohner, Stephan Krenn: A Verifiable Multiparty Computation Solver for the Assignment Problem and Applications to Air Traffic Management. 2022. Available at: arXiv:2205.03048

[PGB+21] Nadine Pilon, Laurent Guichard, Zoltan Bazso, Giuseppe Murgese, Marie Carré: User-Driven Prioritisation Process (UDPP) from advanced experimental to pre-operational validation environment: Journal of Air Transport Management, Volume 97, 2021.

[BEK+21] Jan Bobolz, Fabian Eidens, Stephan Krenn, Sebastian Ramacher, Kai Samelin: Issuer-Hiding Attribute-Based Credentials. Technical report, currently under submission. 2021

[HHNZ19] Marcella Hastings, Brett Hemenway, Daniel Noble and Steve Zdancewic, "SoK: General Purpose Compilers for Secure Multi-Party Computation," 2019 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 2019, pages 1220-1237

[REF4.19.1] Data protection comparison," activeMind.legal. https://www.activemind.legal/law/

[REF4.19.2] "GDPR Tracker," Bird & Bird. https://www.twobirds.com/en/in-focus/general-data-protection-regulation/gdpr-tracker

[REF4.19.3] GDPR Derogations Tracker," Latham & Watkins, Apr. 2018. https://gdpr.lw.com/Home/Derogations

[GDRef1] Said Daoudagh: The GDPR Compliance Through Access Control Systems. PhD Thesis, 2021

[GDRef2] Paolo Barsocchi, Antonello Calabrò, Antonino Crivello, Said Daoudagh, Francesco Furfari, Michele Girolami and Eda Marchetti: COVID-19 & Privacy: Enhancing of Indoor Localization Architectures towards Effective Social Distancing. Array. https://doi.org/10.1016/j.array.2020.100051

[GDRef3] Said Daoudagh, Francesca Lonetti and Eda Marchetti: An automated framework for continuous development and testing of access control systems. J Softw EvolProc. 2020; e2306. https://doi.org/10.1002/smr.23063

[GDRef4] Said Daoudagh, Francesca Lonetti and Eda Marchetti: XACMET: XACML Testing& Modeling. Softw. Qual. J. 28(1): 249-282 (2020)

[GDRref5] Said Daoudagh and Eda Marchetti: GROOT: A GDPR-based Combinatorial Testing Approach. ICTSS 2021.

[GDRef6] Said Daoudagh and Eda Marchetti: GRADUATION: A GDPR-based Mutation Methodology. QUATIC 2021.

[GDRef7] Said Daoudagh, Eda Marchetti, Vincenzo Savarino, Roberto Di Bernardo and Marco Alessi: How to Improve the GDPR Compliance Through Consent Management and Access Control. ICISSP 2021.

[GDRef8] Said Daoudagh, Francesca Lonetti and Eda Marchetti: Continuous Development and Testing of Access and Usage Control: A Systematic Literature Review. ESSE2020

[GDRef9] Paolo Barsocchi, Antonello Calabrò, Antonino Crivello, Said Daoudagh, Francesco Furfari, Michele Girolami and Eda Marchetti: A Privacy-By-Design Architecture for Indoor Localization Systems. QUATIC 2020: 358-366

[GDRef10] Said Daoudagh, Francesca Lonetti and Eda Marchetti: Assessing Testing Strategies for Access Control Systems: A Controlled Experiment. ICISSP 2020: 107-11

[GDRef11] Said Daoudagh and Eda Marchetti: A Life Cycle for Authorization Systems Development in the GDPR Perspective. ITASEC 2020: 128-140
[GDRef12] Said Daoudagh and Eda Marchetti: Defining Controlled Experiments Inside the Access Control Environment. MODELSWARD 2020: 167-176
[GDRef13] Said Daoudagh, Francesca Lonetti and Eda Marchetti: A Framework for the Validation of Access Control Systems. ETAA@ESORICS 2019: 35-51
[GDRef14] Said Daoudagh, Francesca Lonetti and Eda Marchetti: A Decentralized Solution for Combinatorial Testing of Access Control Engine. ICISSP 2019: 126-135
[GDRef15] Said Daoudagh, Francesca Lonetti and Eda Marchetti: A General Framework for De-centralized Combinatorial Testing of Access Control Engine: Examples of Application. ICISSP (Revised Selected Papers) 2019: 207-229
[GDRef16] Cesare Bartolini, Said Daoudagh, Gabriele Lenzini and Eda Marchetti: Towards a Lawful Authorized Access: A Preliminary GDPR-based Authorized Access. IC-SOFT 2019: 331-338
[GDRef17] Antonello Calabrò, Said Daoudagh and Eda Marchetti: Integrating Access Control and Business Process for GDPR Compliance: A Preliminary Study. ITASEC2019
[GDRef18] Cesare Bartolini, Said Daoudagh, Gabriele Lenzini and Eda Marchetti: GDPR-Based User Stories in the Access Control Perspective. QUATIC 2019: 3-17
[GMBS21] Jesús García-Rodríguez, Rafael Torres Moreno, Jorge Bernal Bernabé, and Antonio Skarmeta. 2021. Towards a standardized model for privacy-preserving Verifiable Credentials. In The 16th International Conference on Availability, Reliability and Security (ARES 2021). Association for Computing Machinery, New York, NY, USA, Article 126, 1–6. DOI: https://doi.org/10.1145/3465481.3469204

[kaminskas2018agg] Kaminskas L., Lluch Lafuente A. (2018) Aggregation Policies for Tuple Spaces. In: Di Marzo Serugendo G., Loreti M. (eds) Coordination Models and Languages. COORDINATION 2018. Lecture Notes in Computer Science, vol 10852. Springer, Cham. https://doi.org/10.1007/978-3-319-92408-3_8
[arguspriv acy]Resende, J. S., Magalhães, L., Brandão, A., Martins, R., & Antunes, L. (2021). Towards a Modular On-Premises Approach for Data Sharing. Sensors, 21(17), 5805.
[panos2017security]Panos, C., Malliaros, S., Ntantogian, C., Panou, A., & Xenakis, C. (2017, September). A security evaluation of FIDO's UAF protocol in mobile and embedded devices. In *International Tyrrhenian Workshop on Digital Communication* (pp. 127-142). Springer, Cham.
[Angelo2021how]Angelogianni, A., Politis, I., & Xenakis, C. (2021). How many FIDO protocols are needed? Surveying the design, security and market perspectives. *arXiv preprint arXiv:2107.00577*.

[vasile2021web]Vasileios Grammatopoulos, A., Politis, I., & Xenakis, C. (2021, August). A web tool for analyzing FIDO2/WebAuthn Requests and Responses. In *The 16th International Conference on Availability, Reliability and Security* (pp. 1-10).

[mcmahan2017communication] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics* (pp. 1273–1282).

[mothukuri2021survey] Mothukuri, V., Parizi, R., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems, 115, 619–640*.

[geng2018private] Geng, Q., Ding, W., Guo, R., & Kumar, S. (2018). Privacy and utility tradeoff in approximate differential privacy. *arXiv preprint arXiv:1810.00877*.

[balle2018mechanism] Balle, B., & Wang, Y.X. (2018). Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning* (pp. 394–403).

[nolohan2018bounded] Holohan, N., Antonatos, S., Braghin, S., & Mac Aonghusa, P. (2018). The bounded Laplace mechanism in differential privacy. arXiv preprint arXiv:1808.10410.

[dagan2020noise] Dagan, Y., & Kur, G. (2020). A bounded-noise mechanism for differential privacy. *arXiv preprint arXiv:2012.03817*.

[alsaedi2020iot] Alsaedi, A., Moustafa, N., Tari, Z., Mahmood, A., & Anwar, A. (2020). TON_IoT telemetry dataset: a new generation dataset of IoT and IIoT for data-driven Intrusion Detection Systems. *IEEE Access, 8, 165130–165150.*

[sharemind] Sharemind MPC overview and documentation https://docs.sharemind.cyber.ee/ May 2022

[sforzin2017private] Li, W., Sforzin, A., Fedorov, S. and Karame, G.O., 2017, April. Towards scalable and private industrial blockchains. In *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts* (pp. 9-14).

[liu2018scalable] Liu, J., Li, W., Karame, G.O. and Asokan, N., 2018. Scalable byzantine consensus via hardware-assisted secret sharing. IEEE Transactions on Computers, 68(1), pp.139-151.

[d312] D3.2 - *Cross Sectoral Cybersecurity Building Blocks, Cybersec4Europe deliverable.*

[dersens2020] Laud, Peeter; Pankova, Alisa; Pettai, Martin (2020). A Framework of Metrics for Differential Privacy from Local Sensitivity. Proceedings on Privacy Enhancing Technologies, 2020 (2), 175−208**.**

[ga2022] Publication: Pankova, Alisa; Laud, Peeter (2022). Interpreting Epsilon of Differential Privacy in Terms of Advantage in Guessing or Approximating Sensitive Attributes. Submitted to CSF 2022.

[i] https://navikt.github.io/arxaas/