Project number: 653454
H2020-DS-2014-1

Project start: October 1, 2015
Project duration: 3 years

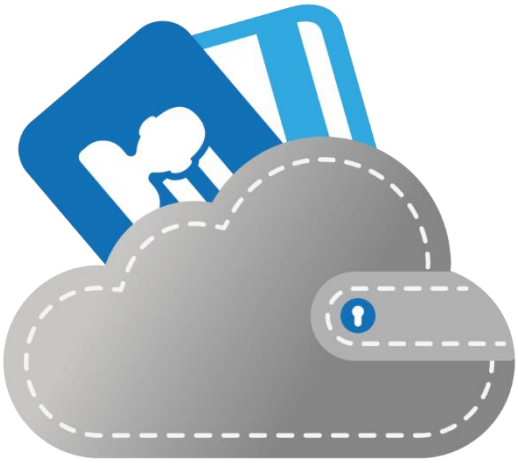

CREDENTIAL

D3.3

Recommendations on privacy enhancing mechanisms

Document Identification	
Due date	March 31, 2017
Submission date	March 31, 2017
Revision	1.0

Related WP	WP3	Dissemination Level	PU
Lead Participant	GUF	Lead Author	Fatbardh Veseli (GUF)
Contributing Beneficiaries	GUF, KAU, ATOS, AIT, TUG, LISPA	Related Deliverables	D2.3, D2.6, D6.1, D4.1, D5.1

Abstract: This deliverable reports on the privacy analysis performed on the core CREDENTIAL architecture. It relies on LINDDUN [1] to define a list of core privacy goals and their respective threats, and to provide the methodological basis for guiding the different steps in the privacy analysis. We model the system architecture using a DFD (Data Flow Diagram), which we subsequently use as a basis for identifying the privacy threats. Next, we categorize threats into high-, medium- and small-risk. For the first two categories, we propose individual mitigation strategies and a total of 29 mitigation mechanisms corresponding for the identified privacy threats. These mitigation mechanisms are meant to enable the design of (a more privacy-friendly) CREDENTIAL Wallet architecture. Finally, based on the experience gained in the project so far, we some high-level recommendations for CREDENTIAL especially targeting system architects and developers in following a Privacy-by-Design approach. Nevertheless, we consider that these high-level recommendations could also serve other projects that face a similar goal / challenge. Finally, we also provide concrete recommendations for regulatory authorities in order to help them enforce better privacy protection for the users beyond technology alone.

This document is issued within the CREDENTIAL project. This project has received funding from the European Union's Horizon 2020 Programme under grant agreement no. 653454.

Executive Summary

This deliverable performs a privacy analysis on the CREDENTIAL Wallet architecture, which is presented in CREDENTIAL Deliverable D5.1 [2]. We try to organise the work based on a sound methodology, which we considered to fulfil our needs. In our case, we considered LINDDUN [1], which is a methodology that helps consider privacy during the development lifecycle. However, we also acknowledge that there is a lack of a mature methodology, which we could fully rely on, therefore we made sure to adapt it to our needs rather than blindly follow it step by step. Therefore, we split the process of performing the privacy analysis in five main steps[1] corresponding to the steps performed in LINDDUN. Each step defines an own chapter in the deliverable.

The deliverable starts by introducing the project objectives and explain the organisational structure of the document. Next, we describe the methodology (LINDDUN), how we applied it, where we deviated from it, and why we did so. Following, we provide a system description of the CREDENTIAL Wallet architecture in the form of a DFD (Data Flow Diagram). This is then used as a basis for the next steps in the privacy analysis, most importantly for identifying the privacy threats in the system.

Next, following LINDDUN, we map privacy threats to the respective DFD elements in the system. Here, we end up with a large number of threats, which become unmanageable if we choose to follow LINDDUN, because of the complexity of the system. Therefore, we make a prioritization of threats based on interviews with experts into three main categories: high, medium, and small risk threats. We then focus on the first two categories, for which we then separately identify corresponding risk mitigation strategies. Here we also deviate partly from LINDDUN and rely on the mitigation strategies from ISO/IEC 27005 [3] instead, since LINDDUN does not provide a strategy that for instance accepts a certain risk, which we acknowledged is normally needed. We explain this in more detail in the corresponding chapter (Chapter 2 - Methodology).

[1] These steps are described in more detail in Chapter 2 - Methodology.

Finally, for the high and medium-risk threats, we propose concrete mitigation mechanisms, which is aimed to serve as a feedback to the ongoing work in the project in the design of the CREDENTIAL Wallet architecture. We categorize the mitigation mechanisms in three groups: first, the mitigation mechanisms that require integration of specific technological tools in the Wallet architecture; second, mitigation mechanisms that require adherence to particular development "guidelines and principles"; and third, mechanisms that rely on organizational measures, such as proper privacy policies or third party contracts.

Finally, based on the experience and the expertise in the project, and considering the requirements from Deliverable D4.1 [4] on the technologies, we make a number of high level recommendations for CREDENTIAL, but also for other projects that consider privacy. These recommendations are mostly targeting system architects and software developers that aim to consider privacy in their projects considering different requirements and limitations. However, we also give important recommendations, which target regulatory bodies, such as national data protection agencies, in monitoring companies and organisations in privacy-related matters.

Document information

Contributors

Name	Partner
Fatbardh Veseli	GUF
Tobias Pulls	KAU
Javier Presa Cordero	ATOS
Miryam Villegas Jimenez	ATOS
Christoph Striecks	AIT
Andrea Migliavacca	LISPA
Alberto Zanini	LISPA
Silvana Mura	LISPA
Felix Hörandner	TUG
Christof Rabensteiner	TUG
Pritam Dash	SIC

Reviewers

Name	Partner
Stephan Krenn	AIT
Miryam Villegas Jimenez	ATOS
Felix Hörandner	TUG

History

0.1	2016-01-14	Fatbardh Veseli	1st Draft
0.2	2016-11-02	Fatbardh Veseli	Initial table of contents (ToC)
0.3	2016-11-10	Fatbardh Veseli	Updated ToC and work plan, adjusted it to LINDDUN. Added contributors and reviewers.
0.4	2016-12-06	Fatbardh Veseli	Updating authors for different sections. Applied the new template.
0.5	2017-01-05	Fatbardh Veseli	Updated chapter 4 mapping privacy threats, added annex A, updated chapter 5 Privacy threat trees.
0.6	2017-01-15	Fatbardh Veseli	Updated the structure of the document, updated responsibilities
0.7	2017-01-20	Fatbardh Veseli	Updated the structure of Chapter 5, corrected the assignment of Chapter 7 & 8 between Christoph (AIT) and Christof (TUG). Updated the ToC, list of figures and list of tables.

0.8	2017-01-23	Tobias Pulls	Significantly refactored sections related to steps two to four in LINDDUN. Documented DFD generalization, interview structure, and preliminary privacy threats with ranking derived from partial interview results.
0.9	2017-01-25	Tobias Pulls	Updated method section, documenting the high-level changes and rational corresponding to steps two to four in LINDDUN.
0.10	2017-01-27	Javier Presa	Add generic definition of the elements of the DFD in section 3.
0.11	2017-01-30	Tobias Pulls	Updated threat tables after third and final interview completed. Expanded one threat (AC-NC) with an example. Replaced pruned DFD figure with the nicer version from Atos.
0.12	2017-02-01	Fatbardh Veseli	Updated Chapter 6 and proposed a table to address the threats.
0.13	2017-02-03	Tobias Pulls	Completed draft of interview process in Chapter 5 and detailed interview results in Annex B.
0.14	2017-02-07	Fatbardh Veseli	Reworked the table with high risks where different threats were grouped together. Split them into individual threats and proposed a mitigation strategy. Some comments added, to be resolved.
0.15	2017-02-13	Tobias Pulls, Christoph Striecks	In Chapter 5, added a short description of how each instantiated threat category was presenter to interviewees. In Chapter 6, included requirements from WP4 (D4.1).
0.16	2017-02-15	Fatbardh Veseli	Updated the table with mitigation strategies for high-risk and medium-risk threats. Updated ToC. Updated the list of bibliography items in Literature_D3.3.bib on Sharepoint and in the bib chapter here.
0.17	2017-02-16	Christof Rabensteiner, Tobias Pulls	Incorporated the user-centric DFDs and updated pruned DFD from Christof, removed the old mapping table and DFD, and replaced DFD elements from threat tables with figure refs.
0.18	2017-02-21	Tobias Pulls	Addressed most review comments from Felix and Stephan in Chapters 2, 4, and 5.
0.19	2017-02-21	Fatbardh Veseli	Integrated the input from Pritam restructured Chapter 7, provided some comments to some sections.
0.20	2017-02-24	Fatbardh Veseli	Integrated another input from Pritam, updated other parts of Chapter 7.
0.21	2017-02-27	Fatbardh Veseli	Added abstract and executive summary, updated the tables in Chapter 7, added a new table with a list of all

			mitigation mechanisms.
0.22	2017-02-28	Fatbardh Veseli, Christoph Striecks	Formatted some references, added some new ones, added small descriptions to Chapter 7 and 6. Addressed reviewer comments for Chapter 1-3.
0.23	2017-03-20	Fatbardh Veseli	Moved the tables from Section 6.3 to the annex, replaced the summary table with a new summary of the main mitigation mechanisms in Section 6.3. Other editorial changes, updated to Chapter 7. Added the names for the two new authors from LISPA.
0.24	2017-03-21	Fatbardh Veseli	Re-integrated manually the input from Tobias so that the track changes are still possible. Edited the captions of tables and figures.
0.25	2017-03-23	Silvana Mura	Integrated the updated version of Chapter 1.
0.26	2017-03-27	Miryam Villegas Jimenez	Provided 2nd review
0.27	2017-03-28	Fatbardh Veseli	Sorted the glossary terms alphabetically.
0.28	2017-03-29	Fatbardh Veseli	Integrated the new version of Chapters 4 and 5 from Tobias
0.29	2017-03-30	Fatbardh Veseli	Integrated updates in Chapter 1 from Silvana and Pritam. Updated text sections in Chapter 6 and 7, as well as Annex A & C.
0.30	2017-03-30	Fatbardh Veseli	Integrated the feedback from Christoph on Section 6.1.
0.31	2017-03-31	Fatbardh Veseli	Moved Annex C to Section 6.4, updated the content of Section 6.3, split it into three subsections.
0.32	2017-03-31	Fatbardh Veseli	Editorial corrections in the document, references, tables, etc. Updated the abstract, exec. summary.
0.33	2017-03-31	Fatbardh Veseli	Removed all comments, made a clean version. Updated the History versions.
1.0	2017-03-31	Stephan Krenn	Final copy-editing for submission

List of Contents

List of Figures

List of Tables

Glossary of Terms

The privacy-related terminology is based on [5], which is also adopted by the privacy analysis framework that we have based our work on [1].

Anonymity	Anonymity of a subject from an attacker's perspective means that the attacker cannot sufficiently identify the subject within a set of subjects, the anonymity set.
Confidentiality	Confidentiality means preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information.
Detectability	Detectability of an IOI means that the attacker can sufficiently distinguish whether such an item exists or not. If we consider messages as IOIs, it means that messages are sufficiently discernible from random noise.
Identifiability	Identifiability of a subject means that the attacker can sufficiently identify the subject associated to an IOI, for instance, the sender of a message.
Information Disclosure (Untintended)	Information Disclosure threats expose personal information to individuals who are not supposed to have access to it.
Linkability	Linkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, etc.) allows an attacker to sufficiently distinguish whether these IOIs are related or not within the system.
Non-compliance (Policy and consent)	Policy and consent Non-compliance means that even though the system shows its privacy policies to its users, there is no guarantee that the system actually complies to the advertised policies. Therefore, the user's personal data might still be revealed.
Non-repudiation	Non-repudiation, in contrast to security, this is a threat for privacy. Non-repudiation allows an attacker to gather evidence to counter the claims of the repudiating party, and to prove that a user knows, has done or has said something.
Plausible deniability	For privacy, plausible deniability refers to the ability to deny having performed an action that other parties can neither confirm nor contradict. Plausible deniability from an attacker's perspective means that an attacker cannot prove a user knows, has done or has said something.
Pseudonymity	A pseudonym is an identifier of a subject other than one of the subjects' real names. Pseudonymity is the use of pseudonyms as identifiers. A subject is pseudonymous if a pseudonym is used as

	identifier instead of one of its real names.
Undetectability	Undetectability of an item of interest (IOI) from an attacker's perspective means that the attacker cannot sufficiently distinguish whether it exists or not. If we consider messages as IOIs, this means that messages are not sufficiently discernible from, e.g., random noise.
Unlinkability	Unlinkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not.
Unobservability	Unobservability of an item of interest (IOI) means undetectability of the IOI against all subjects uninvolved in it and anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI.

1 Introduction

The aim of this report is to provide detailed analysis of privacy risks and to draft recommendations on how to mitigate the identified risks.

The risk analysis is based on the privacy analysis framework LINDDUN (**L**inkability, **I**dentifiability, **N**on-repudiation, **D**ectectability, **D**isclosure, **U**nawareness and **N**on-compliance threats), which foresees a six-step process for analysing a system, identifying privacy threats, and then selecting proper mitigations. This method was customized according to complexity of the CREDENTIAL Wallet Data Flow Diagram architecture (representing different services that attend to the Wallet architecture requirements), which covers 73 Data Flow Diagram (DFD) elements. A form of reduction of LINDDUN was therefore performed, identifying the following five steps:

- define a general DFD
- reduce the DFD
- identify and prioritize privacy threats
- elicit mitigation strategies, and
- provide recommendations.

Once defined the CREDENTIAL Wallet DFD, a pruned DFD model of manageable size, suitable for further analysis, was produced by removing elements not considered further for the first round of LINDDUN. Moreover, a user-centric model of the pruned DFD was created, identifying the four players of our model and eight user-centric actions that cover the vast majority of processes, entities, and data flows in the pruned DFD. A map of the user-centric model back to the pruned DFD is provided, showing how the different actions cover all relevant elements of the DFD.

Initially 56 combinations of potential threats were identified, by considering each privacy threat category of LINDDUN in combination with each action of the user-centric model (7 LINDDUN categories, 8 model actions = 56 combinations). Then, through the involvement of three external experts in the fields of identity management, privacy and/or cryptography, the *impact* and *likelihood* of each threat (combination) was assessed. 22 threats were therefore selected and classified in 2 priority levels: 12 high priority and ten medium priority (low priority was not further considered by the analysis as not relevant for the purpose of this deliverable).

By taking into consideration requirements deriving from technology evaluations and recommendations of D 4.1 [4]. mitigation strategies and mechanisms were identified both for high and medium risk threats. The identified strategy acknowledges the ISO/IEC 27005 risk treatment options, by defining a list of the appropriate mitigation strategy for each risk (risks *reduction*, risk *retention*, risk *avoidance* and risk *transfer)* and, where applicable, the mitigation mechanism(s).

This document recommends 29 different mechanisms to mitigate the high- and medium-priority risks, consisting of technologies that should be integrated in the CREDENTIAL Wallet, development guidelines that should be followed, and organizational measures after the system is made available and used. Finally, we also provide general recommendations for the further development of the Wallet architecture following Privacy-by-Design, which could also be useful

for other projects beyond CREDENTIAL. Furthermore, we also provide a number of high-level recommendations for regulatory bodies, in order to better enforce the privacy protection of the users beyond technology alone.

1.1 Credential Project Highlight

The goal of CREDENTIAL is to develop, test and showcase innovative cloud based services for storing, managing, and sharing digital identity information and other critical personal data. The security of these services relies on the combination of strong hardware-based multi-factor authentication with end-to-end encryption representing a significant advantage over current password-based authentication schemes.

The use of sophisticated proxy cryptography schemes, such as proxy re-encryption and redactable signatures, will enable a secure and privacy preserving information sharing network for cloud-based identity information in which even the identity provider cannot access the data in plain-text and hence protect access to identity data.

Project goal is to extend the application of CREDENTIAL approach to a comprehensive cloud system and to existing solutions by using and exploiting recognized standards and protocols.

CREDENTIAL's basic architecture integrates cryptographic mechanisms into three key actors: user, Wallet, and data receiver, as shown in Figure 1.

- The **user** owns data that may be securely stored or shared with other members. In the user's domain is deployed a client application to handle operations such as signing or generating a re-encryption key.
- The **Wallet** is a cloud-based data storage and sharing service characterized by important advantages such as constant availability, scalability, and cost effectiveness. The Identity and Access Management system implements a multi-factor authentication and authorizes access to stored data. With proxy re-encryption, the confidentiality of the stored and shared data is ensured even when deploying the Wallet at a cloud provider, as no plain data is exposed. Moreover, once a re-encryption key is available, data can be shared with other participants even if the user or his/her client application are not available.
- The **data receiver,** that can be either another CREDENTIAL user or a service provider, relies on data stored in or authentication assertions issued by the CREDENTIAL Wallet. With this information, the data receiver reaches authorization decisions and performs arbitrary data processing.

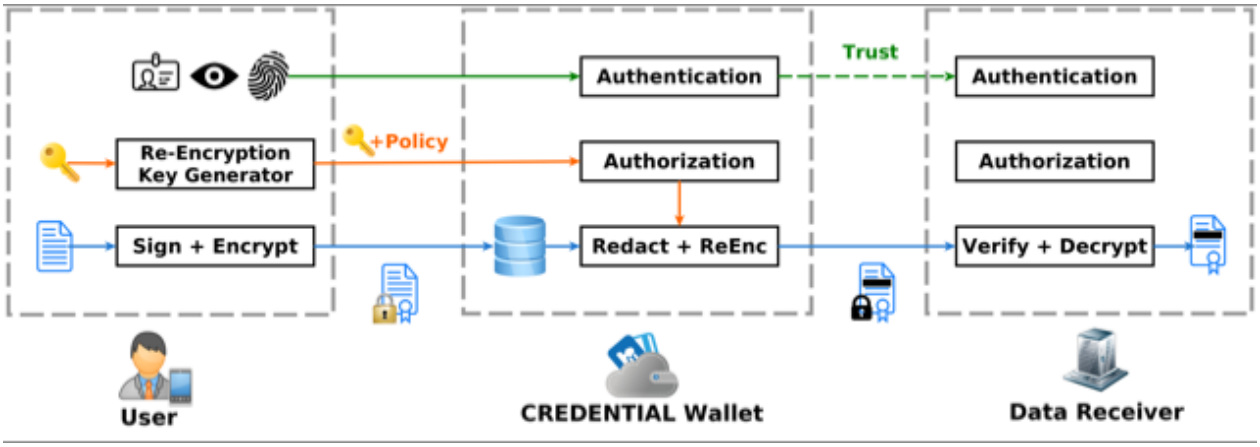

Figure 1: CREDENTIAL's core architecture components

The data sharing process is based on the following steps: the user authenticates at the Wallet to get permission to upload signed and encrypted data, data is encrypted for the user herself to keep maximum control, whilst a re-encryption key is generated by the Wallet towards a selected data receiver. The re-encryption key generation is performed on the user domain and a policy defining which data may be disclosed is transmitted to the Wallet.

Then, the user's ciphertext is transformed into encrypted data for the data receiver by using the re-encryption key. Finally, the data receiver decrypts the data and verifies the signature on the disclosed parts.

To showcase the functionality of the CREDENTIAL Wallet and to demonstrate how a higher security and privacy can be achieved by the means of the CREDENTIAL Wallet, three different pilots in the domains eGovernment, eHealth and eBusiness are performed by the project.

- **eGovernment**: the pilot focuses on identity management to authenticate citizens and assess their eligibility for a service, based on sensitive identity attributes.
  Standardized identity protocols such as SAML or OpenID Connect are used in CREDENTIAL's identity management data sharing process.
  Within this protocol, the service provider (i.e., the data receiver) triggers the process by requesting authentication and identity attributes from the identity provider (i.e., the CREDENTIAL Wallet). The pilot will concern not only enabling authentication for national eID solutions but also cross-border authentication according to the eIDAS regulation.
  The Wallet prompts the user for consent as well as a re-encryption key, and then selectively discloses re-encrypted attributes to the service provider.
- **eHealth**: The pilot focuses on secure data sharing between patients, doctors, and further parties, in the field of Type 2 diabetes.
  In particular, the pilot applies to patients, which can use mobile devices to record their health data, those data are collected by a CREDENTIAL eHealth mobile app which remotely stores

them in the CREDENTIAL Wallet. The user can decide who is allowed to access medical data and which specific parts, thus allowing authorized people to prepare and send advices back to the user. The fact of disclosing a minimal part of data brings several advantages to user such as saving time and money for personal visits and having a continuous control of health data.

- **eBusiness**: The pilot focuses on the integration of modular libraries implementing CREDENTIAL's technologies to increase the privacy offered by existing solutions.
  In particular, it tackles the issue of encrypted mails, which are nowadays increasingly used by companies to protect data and products but raise important challenges when employees are not at work.
  In fact, according the current legislation, workers have to provide their private keys to access company e-mail to give the possibility to other colleagues to still read and eventually take over incoming mail.
  Thanks to proxy re-encryption system, a worker can generate a re-encryption key and deliver it to an authorized colleague before leaving. The new person does not manage the key but the mail server is able to decode incoming mail during the worker absence.

1.2 Objectives

This document provides a detailed analysis of privacy risks and recommendations to mitigate the identified risks. The aim is to tackle CREDENTIAL Wallet architecture, as presented in CREDENTIAL Deliverable D5.1 [2], from a privacy perspective, and to set recommendations useful for CREDENTIAL project to design its information system.

To perform a privacy analysis on the CREDENTIAL Wallet architecture, we propose individual mitigation strategies, as well as concrete mitigation mechanisms corresponding to the mitigation strategy.
The overall objective is to provide recommendations also for other privacy-aware projects on how to mitigate certain risks, especially targeting system architects and developers, but also regulatory bodies in order to provide better protection of user's privacy.

1.3 Scope and Connection with Other Deliverables

The report, which provides detailed analysis of privacy risks and recommendations to mitigate the identified risks, is related to D2.3, D2.6, D6.1, D4.1 and D5.1.

In particular it takes into consideration results of the following CREDENTIAL deliverables:

- D6.1 "Pilot use case specification" as it provides detailed specification of use cases with consideration of pilot site backend infrastructure and client side integration;

- D4.1 “Assessment report on cryptographic technologies, protocols and mechanisms”, which addresses technology evaluations and recommendations used within CREDENTIAL, whose limitations have to be considered also in the CREDENTIAL privacy context;
- D2.3 “Cloud Identity Wallet requirements”, which makes a description of the functional requirements to be used as an input to design the CREDENTIAL architecture, identifies the privacy requirements to be addressed by privacy enhancing mechanisms.

So far, findings of this deliverable have to be considered as a reference for the following reports:

- D2.6 “User Centric privacy and usability requirements” that will provide a further analysis on requirements for centric privacy design aspects;
- D5.1 “Functional Design” that, on the basis of previously identified requirements such as recommendations on privacy-enhancing mechanisms, will define the functional design of the security protocols and of the IAM services.

1.4 Outline

The rest of this deliverable is structured as follows:

- *Chapter 1* – Introduction – outlines project activities focusing on the objectives of the analysis, describes the relation with other project deliverables and gives an overview of the document contents.
- *Chapter 2* – Methodology – describes the process followed to select the methodology and how the method was customized to apply to CREDENTIAL Wallet. The main steps of the method are therefore listed, making clear whether a deviation from the chosen methodology was made.
- *Chapter 3* – CREDENTIAL Wallet Architecture (DFD) – provides a visual representation of the architecture of the CREDENTIAL Wallet system and makes available a generic description of the data flows of the DFD.
- *Chapter 4* – Reducing the CREDENTIAL DFD – describes steps implemented to reduce the DFD to produce a generalized model of manageable size suitable for further analysis. A pruned DFD is represented, together with a user-centric model of the pruned DFD.
- *Chapter 5* – Threat Identification and Prioritization – presents the process followed to identify threats and results classified in 3 priority threats, either high, medium, or low priority. The chapter presents only high and medium priority threats, whilst low priority threats are not presented here, as they are not further considered in this deliverable.
- *Chapter 6* – Mitigation of Privacy Threats – identifies mitigation strategies and mechanisms for the high- and medium-priority privacy threats, which should serve as input to the development of the Wallet architecture in order to achieve better privacy protection.

- *Chapter 7* – Recommendations on Privacy-Enhancing Mechanisms – provides general recommendations for similar projects that integrate privacy-enhancing technologies, and provide examples of some technologies that could help reduce privacy threats.
- *Annex A* – Annex A – Mapping Privacy Threats to CREDENTIAL DFD Elements – contains the complete table mapping LINDDUN privacy threats to the elements of the general DFD.
- *Annex B* – Annex B – Detailed Interview Results and Rationale – contains eight tables with detailed interview results and the rationale behind them.

2 Methodology

Our method is based on the privacy analysis framework LINDDUN [1] that is intended for software engineers (not necessarily with a lot of privacy experience) to be able to account for privacy threats early in the software development lifecycle. LINDDUN is a short derived from seven high-level privacy threat categories **L**inkability, **I**dentifiability, **N**on-repudiation, **D**ectectability, **D**isclosure, **U**nawareness, and **N**on-compliance. The approach of LINDDUN is a six-step process for analysing a system, identifying privacy threats, and then selecting proper mitigations. Figure 2 shows the six main steps from the LINDDUN original paper [1] and they are:

1. Define a *Data Flow Diagram* (DFD) of the system to be analyzed.
2. Map the seven LINDDUN privacy threat categories to the elements of the DFD.
3. For each mapping, identify threat scenarios based on threat trees from LINDDUN.
4. Prioritize the threats using any appropriate method (not part of LINDDUN).
5. In order of importance, elicit mitigation strategies for the threats using identifiers from the threat trees (step 3) and a mapping to relevant mitigation strategies provided by LINDDUN.
6. Based on the mitigation strategies, select corresponding PETs from a table in LINDDUN.

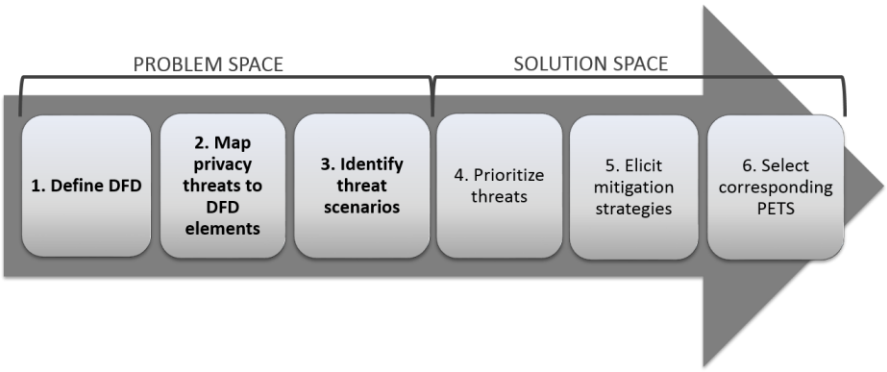

Figure 2: The main steps in the tutorial on applying LINDDUN [1]

A more recent tutorial on applying LINDDUN was later made available by the authors [6], which slightly differs in the final steps. We then decided to take the last version, which is also shown in Figure 3. We started early on step one of LINDDUN with defining a DFD. However, since requirements elicitation of the Wallet, including key functional requirements for our three pilots, was ongoing in parallel a significant time was spent on refining the DFD as the core functionality of the Wallet was determined. Finally, for this deliverable, we settled on the sixth iteration of the DFD containing in total 74 DFD elements consisting of 52 data flows, 14 processes, 3 external entities, and 4 data stores. Mapping each of the seven threat categories to 74 DFD elements (step two of LINDDUN) is impractical – especially if you consider that step three further potentially

multiplies the number of resulting threats – so we decided to deviate from strictly following LINDDUN and performing our own form of *reduction*.[2]

LINDDUN specifies that rows in the mapping table *after step two* can be reduced by combining and generalizing parts to reach a manageable size. Because the threat trees used in step three to identify threats are designed for software engineers with little privacy experience, and step four does not specify a method for prioritizing the threats, we decided tackle steps two to four as follows. First, we reduced the DFD and made a *user-centric model* of the system depicted by the DFD, describing the system using terminology familiar to non-consortium members with an expertise on privacy[3]. The model contains eight actions performed by end-users of the CREDENTIAL Wallet. With this generalized model and actions, we consider each combination of LINDDUN threat categories and actions (56 combinations) as potential threats. We then perform semi-structured interviews with privacy experts in the area of identity management to both identify and prioritize the most important privacy threats. By considering each threat category and interviewing privacy experts we tap into knowledge now assumed available in the LINDDUN method. Figure 3 shows the five steps of our selected method.

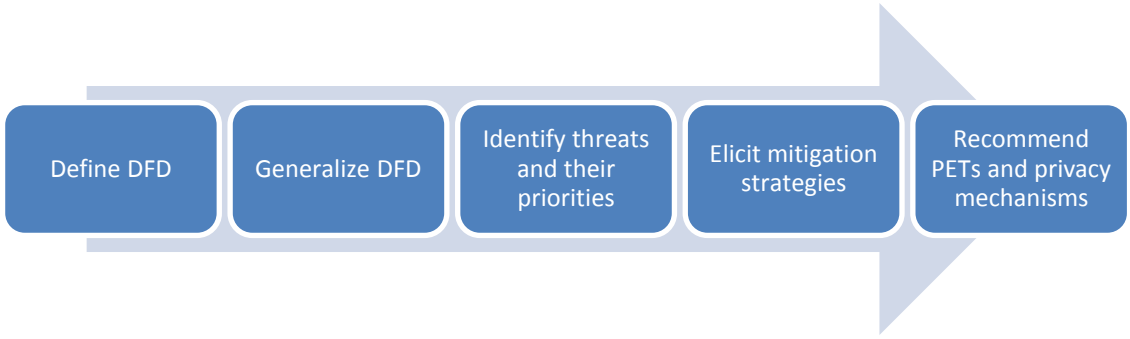

Figure 3: The main steps of our LINDDUN-based tutorial [6]

[2] However, we provide the complete list of threats that were considered in the mapping of each DFD element to a particular threat using the original proposal of LINDDUN in Annex A of this deliverable.

[3] The process of defining the DFDs identified that our setting is surprisingly symmetrical in the sense that users with a CREDENTIAL Wallet ("owners") and those interacting with the Wallet, such as *service providers or other end-users* ("participants"), can be modelled similarly. Commonly you find clear distinctions between human end-users and service providers ("relying parties").

Finally, the result is a prioritized list of privacy threats to be mitigated as specified by LINDDUN. Furthermore, we also deviate from LINDDUN in step 4 – Elicit Mitigation Strategies. In this respect, we chose from the mitigation strategies from ISO/IEC 27005 [3], since, unlike LINDDUN, this standard differentiates between mitigated, reduced, and remaining risks.

3 CREDENTIAL Wallet Architecture (DFD)

A data flow diagram (DFD) is a popular visual representation of how data will move within a software system, although theoretically could be applied to any other business models.

In order to represent internally the architecture of the CREDENTIAL Wallet system, there are some elements included in the DFD of Figure 4. The meaning of them is the following:

- Process **(P)**: it represents each of the services included in the CREDENTIAL Wallet.
- Entity **(E)**: it represents the actors involved in the system, namely Participant (using the Wallet services), Owner (of the Wallet account), and Identity/Attribute Provider (IdP).
- Data flow **(DF)**: it represents a communication between two components.
- Data storage **(DS)**: it represents the databases used in the system.

The core of the DFD consists of the different services that attend to the Wallet architecture requirements. Once these processes are represented, one-way arrows are included to draw the information communication, and the direction of it, between these processes or any two components. These data flows are labelled to indicate the type of data sent in that structure of the communication, but the accurate content of this structure could not be indicated, as each data structure should accommodate the requirements of each situation, e.g. assurance level of the service provider.

The trust boundaries are also represented in Figure 4; in particular the *CREDENTIAL Wallet boundary* is the border representation that defines the limits in the privacy threat analysis of the current document.

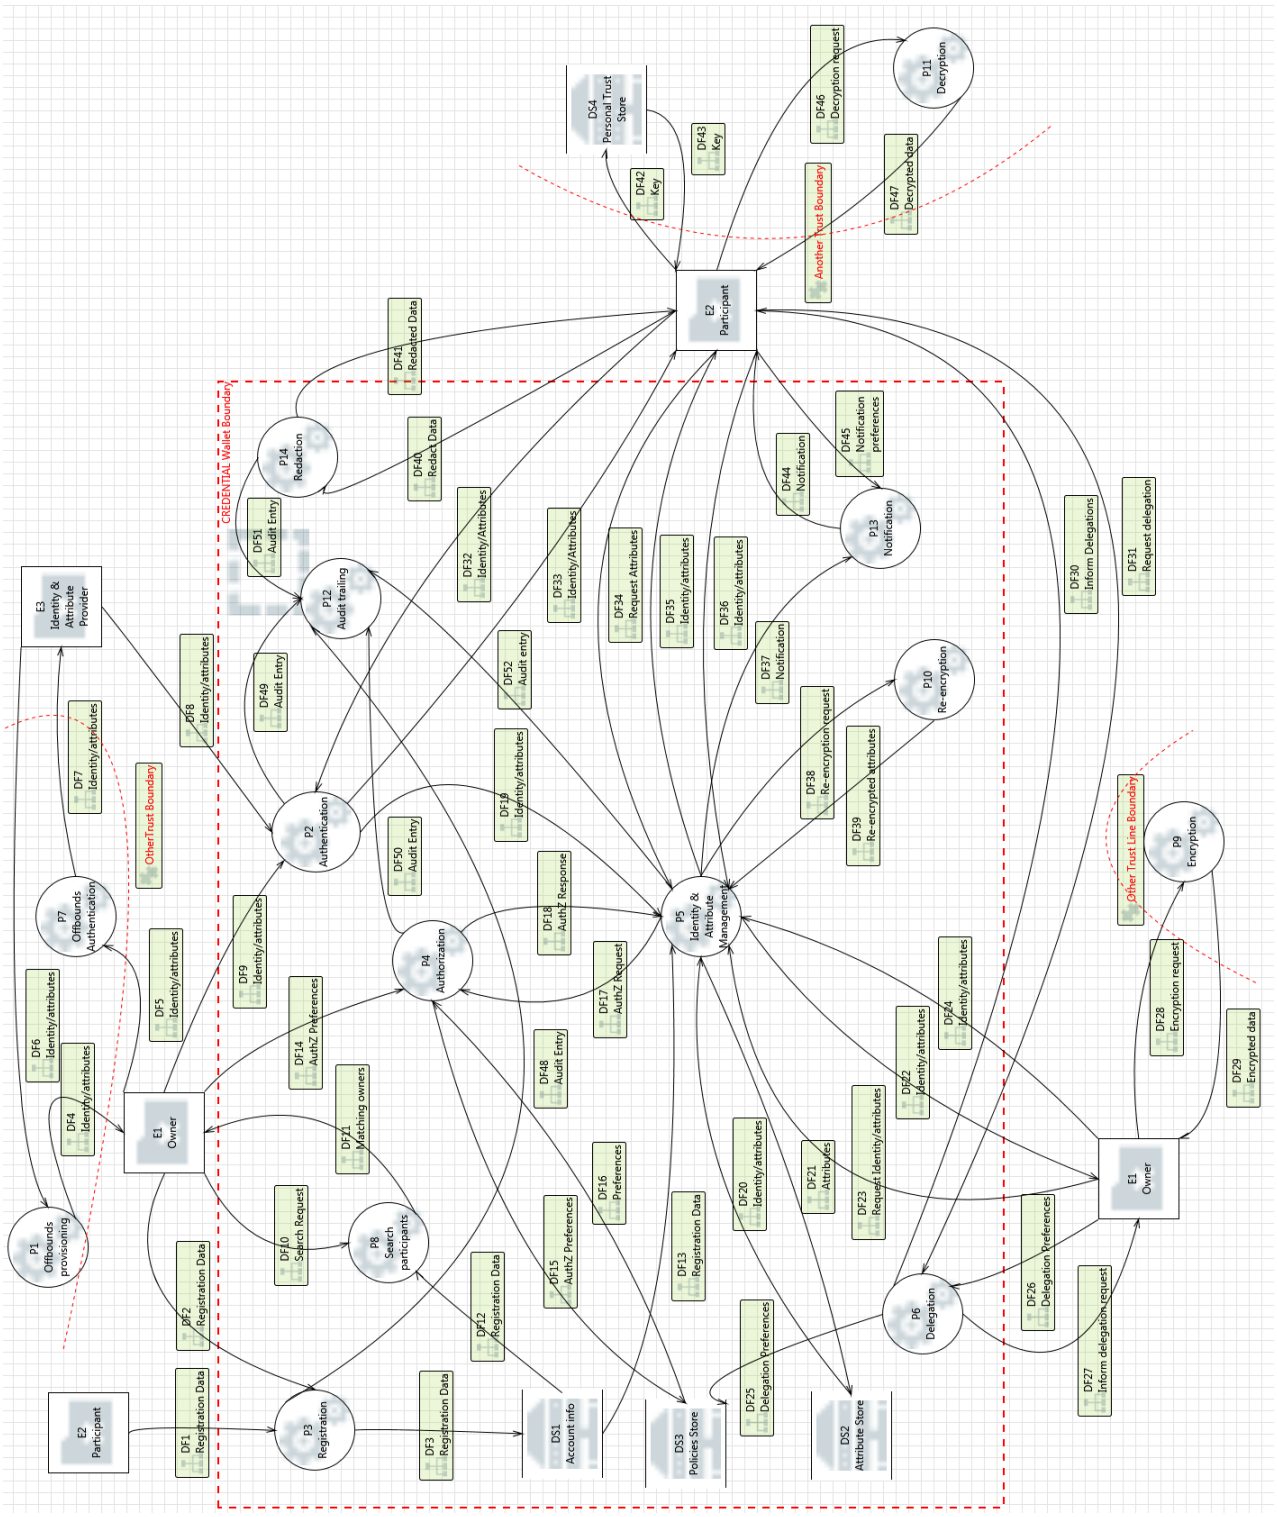

Figure 4: General Data Flow Diagram (DFD) of the CREDENTIAL Wallet architecture

As the processes, entities and data storages are uniquely named but not the data flows, a generic description of each is provided in Table 1.

Table 1: Definition of the data flows of the DFD

Component name	Definition / Description
DF1 Registration data	Participant information used to register or delete an existing account in the Wallet
DF2 Registration data	User information sent during registration or to be deleted in the Wallet
DF3 Registration data	Data of the account registered to be stored
DF4 Identity/attributes	Identity or attributes requested by user to an external IdP
DF5 Identity/attributes	User information sent to a trusted service to be authenticated in external IdP
DF6 Identity/attributes	Identity or attributes sent by external IdP to the external service in charge of provisioning
DF7 Identity/attributes	User information sent to an external IdP to be authenticated
DF8 Identity/attributes	External account data to be inbound authenticated by the Wallet
DF9 Identity/attributes	User information sent for authentication in the Wallet
DF10 Search Request	Query made by a user to search a participant
DF11 Matching owners	Information requested by the user query
DF12 Registration data	Stored account information requested by Wallet Search service
DF13 Registration data	Account information received by the data management service of the Wallet
DF14 AuthZ Preferences	Preferences of authorization sent by user to the service
DF15 AuthZ Preferences	Authorization preferences sent by the authorization service to be stored

Component name	Definition / Description
DF 16 Preferences	Response with the authorization preferences requested by the authorization service
DF17 AuthZ Preferences	Query of authorization preferences from data manager service
DF18 AuthZ Response	Response with the authorization preferences requested by data manager service
DF19 Identity/attributes	Account data sent by the authorization service to data manager
DF20 Identity/attributes	Information retrieved from database by the data manager
DF21 Attributes	Attributes to be stored in the database
DF22 Identity/attributes	Account information requested by user
DF23 Request Identity/attributes	Account information query made by user
DF24 Identity/attributes	Account information sent by user to the Wallet to be managed accordingly
DF25 Delegation preferences	Delegation preferences sent by the delegation service to be stored
DF26 Delegation preferences	Preferences of delegation sent by user to the service
DF27 Inform delegation request	Query of delegation preferences to the user from the delegation service
DF28 Encryption request	Information the user wants to encrypt
DF29 Encrypted data	Encrypted data sent by the service
DF30 Inform delegations	Response to delegation request made by participant
DF31 Request delegation	Participant query to ask for delegation
DF32 Identity/attributes	Participant information sent for authentication in the Wallet
DF33 Identity/attributes	Identity assertion issued by authentication service to

Component name	Definition / Description
	participant
DF34 Request attributes	Participant query of attributes to the Wallet
DF35 Identity/attributes	Response with information requested by the participant
DF36 Identity/attributes	Information sent by participant to the data manager
DF37 Notification	An event occurred in data manager to be notified
DF38 Re-encryption request	Information the Wallet needs to re-encrypt
DF39 Re-encrypted attributes	Re-encrypted data requested sent by the service
DF40 Redact data	Data where the redact operation shall be applied
DF41 Redacted data	Valid signed redacted document sent by service
DF42 key	Key to be stored
DF43 key	Retrieved key needed
DF44 Notification	Any event to be notified to the participant
DF45 Notification preferences	Preferences of notification chosen by participant
DF46 Decryption request	Information the participant needs to decrypt
DF47 Decrypted data	Decrypted data requested sent by the service
DF48 Audit entry	Event information of registration service to be logged
DF49 Audit entry	Event information of authentication service to be logged
DF50 Audit entry	Event information of authorization service to be logged
DF51 Audit entry	Event information of redaction service to be logged
DF52 Audit entry	Event information of data manager service to be logged

4 Reducing the CREDENTIAL DFD

With a DFD of the CREDENTIAL system, we now set out to reduce and thus simplify the CREDENTIAL DFD. We primarily do this because of scope: combining all DFD elements with the LINDDUN threat categories results in over 200 combinations to consider. 0contains a preliminary mapping of privacy threat categories to the architecture elements in the DFD that motivated our decision to generalise the DFD. We perform two steps of reductions of the DFD: first, we prune the DFD by removing elements, and then we make a user-centric model of the DFD suitable for use in the next step of our work. After presenting the two reductions (Sections 4.1 and 4.2), we thoroughly justify the design of our user-centric model and map it back to the reduced DFD elements (Section 4.3).

4.1 Pruning Elements

Figure 5 shows a pruned CREDENTIAL DFD, removing elements not considered further for the first round of LINDDUN. Rationale for the removal of elements (compare with the original DFD in Figure 4):

- We assume data stores are within the respective trust boundaries of where they will be located (in the *Wallet* and locally at the *participants*, respectively).
- We consider *off-bounds provisioning* (P1) out of scope due to being dependant on the *Identify & Attribute Provider* (E3).
- Without a clear picture of what information is available to the Wallet, discussing the implications of the *audit trail* (P14) is harder. We likely capture relevant privacy threats when going over the other parts related to Wallet functionality, since the audit trail cannot contain information not available to the Wallet.
- Flows presumably within the Wallet (between processes to stores) are removed.

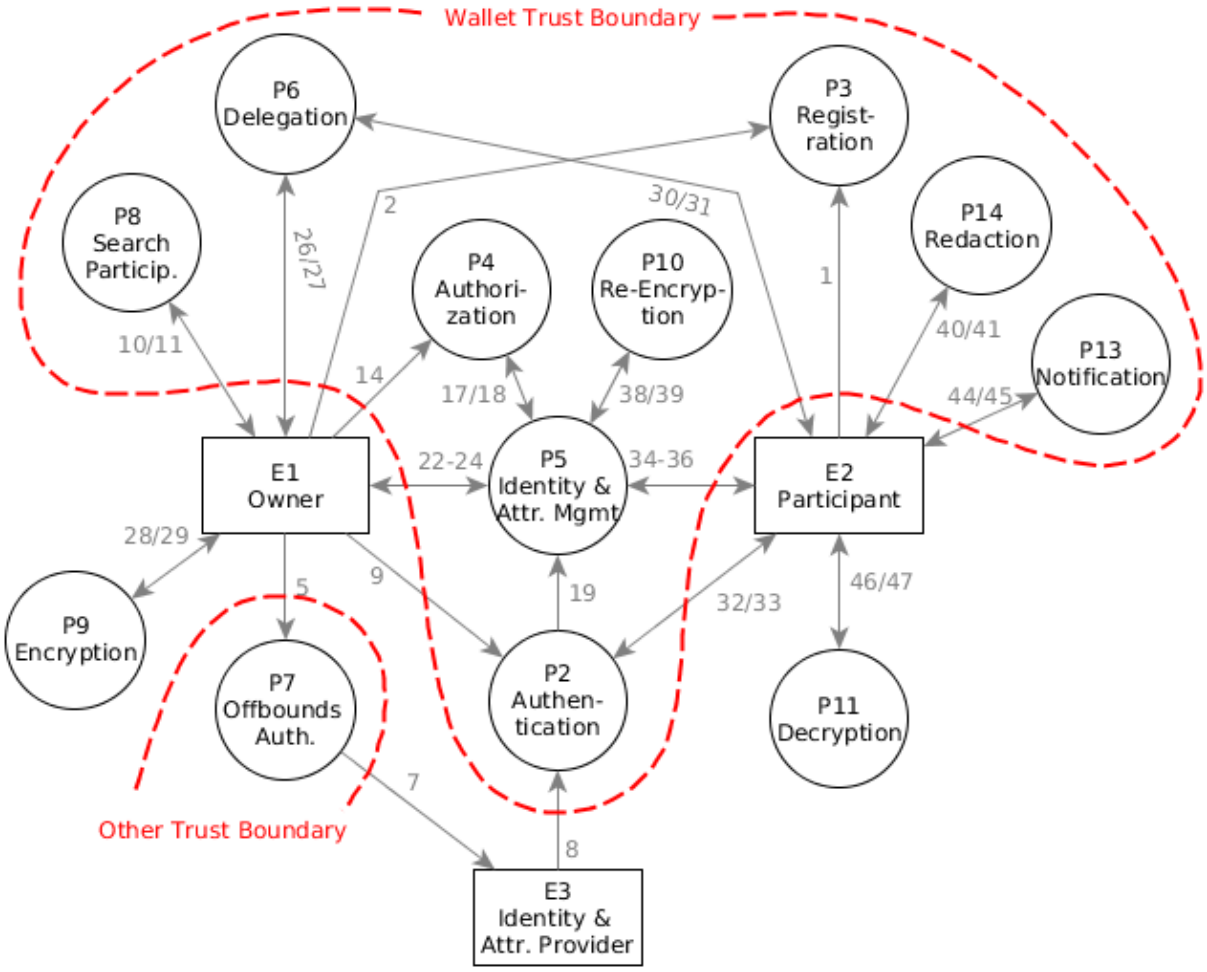

Figure 5: The pruned CREDENTIAL DFD, removing elements not considered for the first round of LINDDUN

4.2 A User-Centric Model

For the purpose of more easily being able to perform semi-structured interviews with people outside of the consortium (next step), we create a user-centric model of the pruned DFD. The model uses terminology and generalizations that are easy to relate to for experts with experience of currently deployed identity management systems. Figure 6 shows the four players in our model together with examples (a hospital and Facebook) used for our analysis:

- **User:** the individual end-user (data subject) that uses the CREDENTIAL Wallet.
- **Wallet:** the CREDENTIAL Wallet that stores data in an encrypted format using a proxy re-encryption scheme.
- **Relying Party:** a hospital that *uses* the Wallet for *managing* users' health data. Usage involves, e.g., retrieving data from and storing data in the Wallet.

- **External IdP:** Facebook, an external identity provider that users already today can use to log-in to services online.

Figure 6: The four players in our user-centric model and the examples (hospital and Facebook) used in the interviews

Based on the players in Figure 6, we define eight user-centric actions that cover the vast majority of processes, entities, and data flows in the pruned DFD. We stress that Facebook as an external identity provider (IdP) and a hospital as a relying party are only examples and does not necessarily reflect planned work within the project. We selected Facebook and a healthcare-related relying party to emphasize privacy-related threats, since Facebook's impact on privacy is debated and most people consider health and healthcare as private affairs. The eight user-centric actions, depicted in Figure 7 together with the relevant players, are:

- **Register:** the user creates an account at the Wallet.
- **Authenticate:** the user authenticates to the relying party with the help of the Wallet.
- **Get/set data:** the user or relying party gets or sets data in the Wallet.
- **Authorize:** the user (or relying party) authorizes the relying party (or user) to access parts of the Wallet.
- **Notify:** the user or relying party are notified by the Wallet (e.g., on authorization request or data modified in the Wallet).
- **Search:** the user (or relying party) searches for other users or relying parties that are registered at the Wallet.
- **External IdP authenticate:** the user authenticates to the relying party with the help of the Wallet and the external IdP Facebook.
- **External IdP data import:** the user imports authenticated data from the external IdP Facebook into its Wallet for use by the relying party.

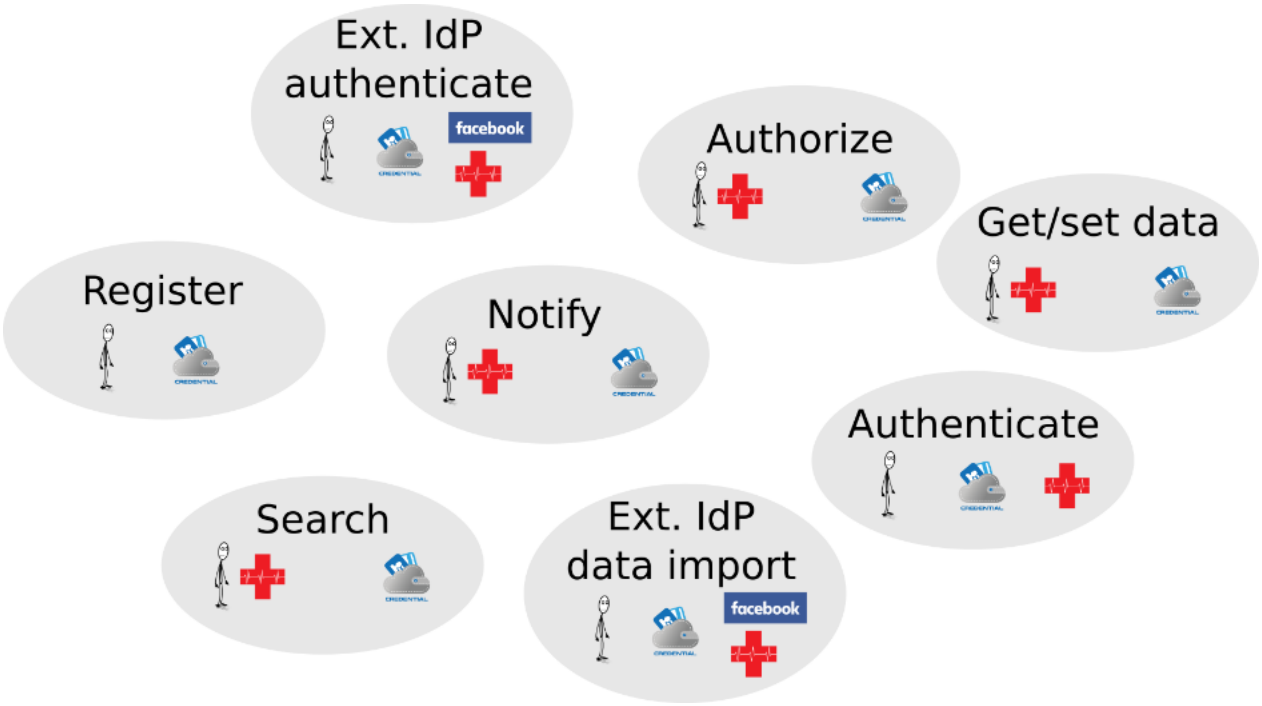

Figure 7: The eight actions of a user

4.3 Mapping the User-Centric Model to the DFD

Next we map the user-centric model (Figure 7) back to the pruned DFD (Figure 5) and show how the different actions cover all relevant elements of the DFD. In general, each user action in the user-centric model depicts the relevant entities/actors, and leaves out data flows. The vast majority of data flows are nothing more than response-reply pairs, as shown in the pruned DFD.

First out are the DFD elements of the *registration* action, shown in Figure 8. A participant (E2) is turned into an owner (E1), supplying data (DF 1-2) for the registration (P3) process in the Wallet. The registration action depicts the user (E1-E2) and the Wallet (P3).

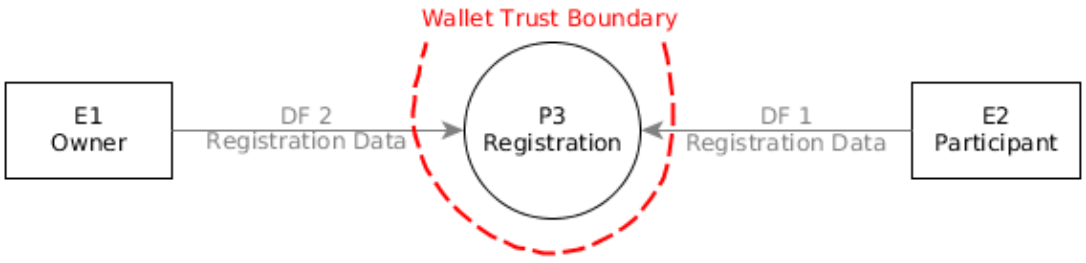

Figure 8: The DFD elements for the *register* action

Figure 9 shows the *authentication* action. The user (E1) is authenticating towards a relying party (E2) through the registration process (P2) at the Wallet with associated data flows (DF 9 and DF 32-33).

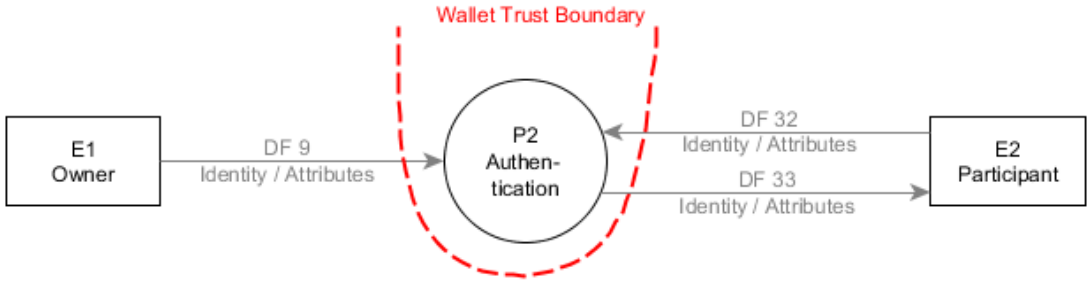

Figure 9: The DFD elements for the *authenticate* action

Figure 10 shows the DFD elements for the *get/set data* action. Either the user (E1) or a relying party (E2) are involved in getting/setting data (P5) at the Wallet through associated data flows (DF 22-24 and DF 34-36).

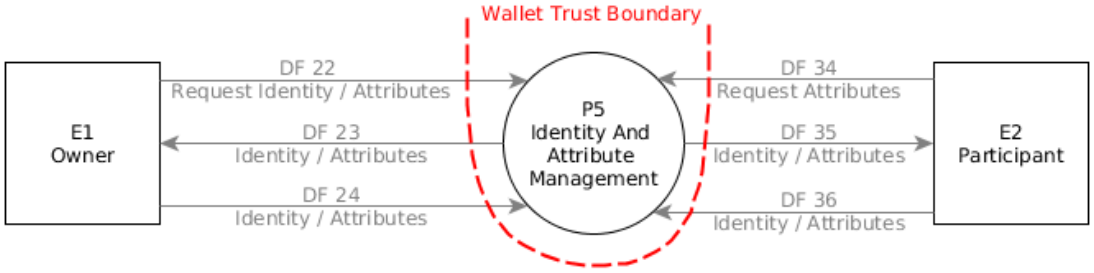

Figure 10: The DFD elements for the *get/set data* action

Figure 11 shows the DFD elements for the *authorize* action. Authorizations are authorized/set by the user (E1) and concern the rights for a relying party (E2). In the pruned DFD, the authorizations are expressed either directly as authorization (P4) or delegations (P6), with associated data flows from the user to the Wallet (DF26-27 and DF 14) or relying party to Wallet (DF 30-31).

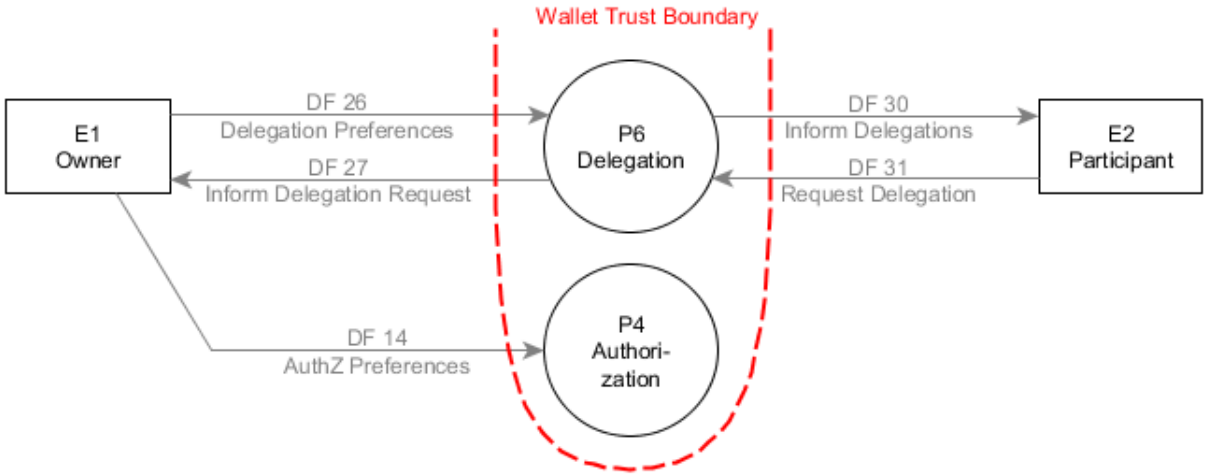

Figure 11: The DFD elements for the *authorize* action

Figure 12 shows the DFD elements for the *notify* action. A user and/or relying party (E2) sets preferences and receives notifications (DF 44-45) from the Wallet that sends the notifications (P13).

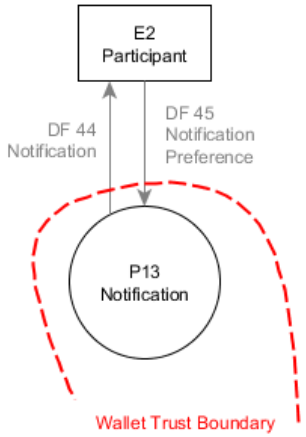

Figure 12: The DFD elements for the *notify* action

Figure 13 shows the DFD elements for the *search* action. The user (E1) searches (D 10-11) for other users (Wallet owners) at the Wallet (P8).

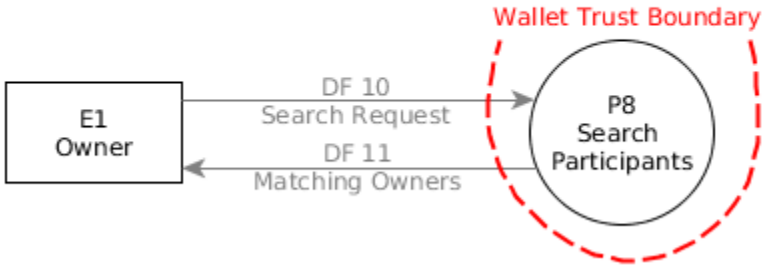

Figure 13: The DFD elements for the *search* action

Figure 14 shows the DFD elements for *authentication with an external IdP*. The external IdP (E3) provides attributes to the Wallet (DF 8) together with the user (E1 and DF 9), that in turn authenticates the user towards the relying party (E2 and DF 32-33).

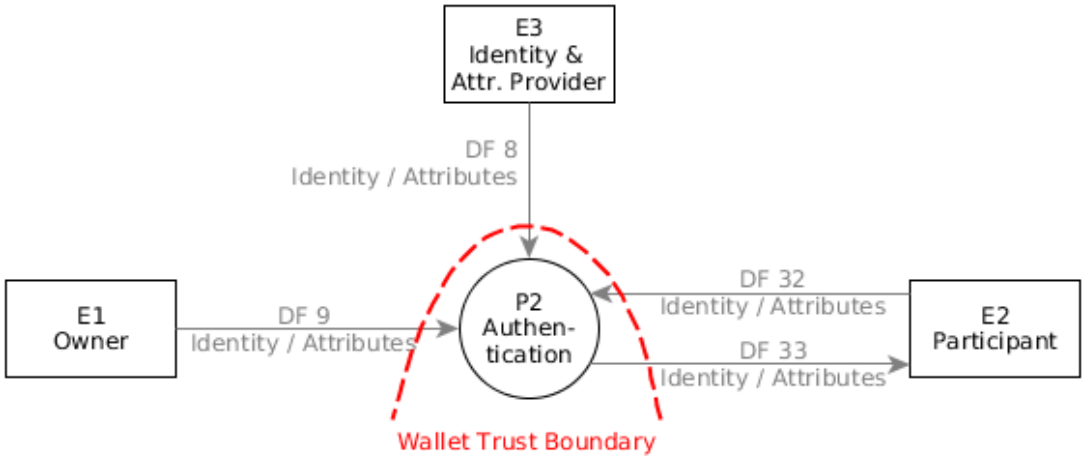

Figure 14: The DFD elements for *the external IdP authenticate* action

Figure 15 shows the DFD elements for *the importing data from an external IdP* action. The external IdP (E3) provides authenticated data (DF 8) concerning the user (E1 and DF 9) to the Wallet (P2). The Wallet stores the data locally (DF 19 and P5) and later provides it for use by a relying party (E2 and DF 32-33).

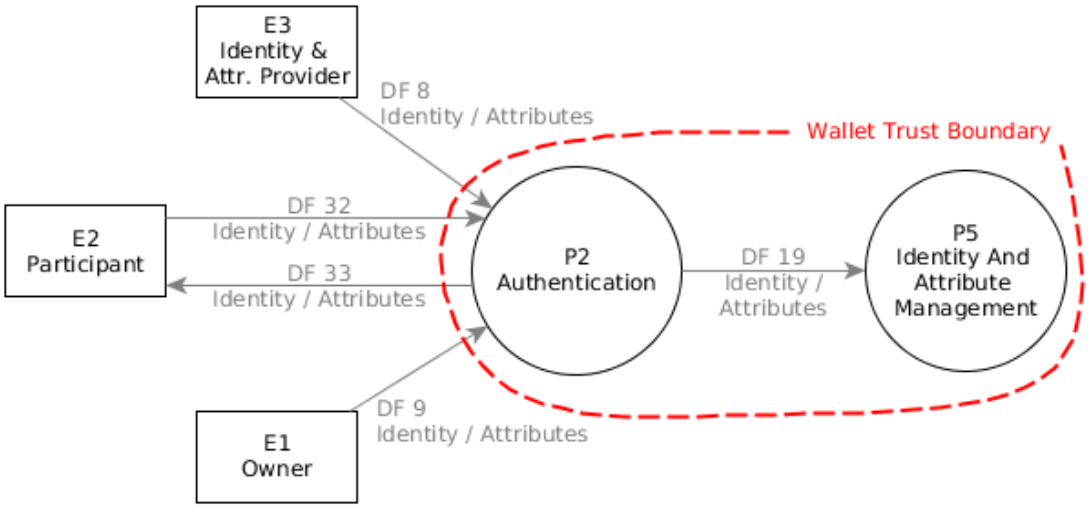

Figure 15: The DFD elements for the *external IdP data import* action

The DFD elements left to consider are the specific proxy re-encryption and malleable signatures processes (P9-11) and data flows (DF 28-29, DF 38-39, and DF 46-47), which are tightly coupled to the redaction process (P14) and associated data flows (DF 40-41). We address this omission in part by emphasising the use of these cryptographic building blocks in our semi-structured interviews (see the following chapter) and by assuming that most relevant privacy threats that have implications on these building blocks will be captured by the actions above. For example, by identifying that non-repudiation is a privacy threat when importing data from an external IdP, one possible mitigation strategy is to ensure that our cryptographic building blocks provide *deniability*, thus indirectly the actions above have the potential to identify relevant privacy threats for the building blocks.

5 Threat Identification and Prioritization

We first give an account of our process of eliciting and prioritizing threats using semi-structured interviews, then present the resulting relevant threats and their priorities. The mitigation mechanisms for these threats are further addressed in Chapter 6 and the basis for our recommendations in Chapter 7.

5.1 Process

Instead of using the threat trees of LINDDUN to identify threats, we consider each privacy threat category of LINDDUN in combination with each action in our user-centric model (seven from LINDDUN and eight actions from the model, in total 56 possible combinations) as potential threats. Using semi-structured interviews with three experts (*not part of the CREDENTIAL consortium*) in the area of identity management, privacy and/or crypto, we asked the experts to assess the *impact* and *likelihood* of each threat (combination). Each interviewee was interviewed separately and the interviews lasted about an hour. The interviews were structured as follows:

1. Explain that notes are only kept on paper, no names or raw estimates will be shared
2. Introduction to the CREDENTIAL Wallet
3. High level idea behind proxy re-encryption if requested (passive/semi-trusted proxy)
4. Explain the user-centric setting from Section 4.2
5. For each action:
    a. Describe the action and the relevant players
    b. For each LINDDUN category:
        i. Briefly describe the threat (threat category and action in setting → threat)
        ii. Have interviewee reason about the threat, discussion if requested
        iii. Note observations of interviewee
        iv. Ask for impact and likelihood estimates
6. Thank interviewee

The introduction to the CREDENTIAL Wallet started as follows:

"In the CREDENTIAL project we're building a secure cloud identity Wallet. The Wallet will use proxy re-encryption and malleable signatures to enable selective disclosure of encrypted and signed data. The encryption protects the contents of the data from the Wallet.

We're currently designing the Wallet, and as part of this process we're working on identifying and ranking privacy threats. This is where you---the interviewee---come in. We use a modified version of the LINDDUN method."

Note that an emphasis was put on the fact that the encryption hides the content of the data stored in the Wallet from the Wallet: this is the only information that was provided to the interviewees with regard to our trust model. After the introduction we discussed proxy re-encryption (one interviewee with prior knowledge) and LINDDUN (one interviewee encountered it before, had

no comment on our modifications to the methodology), depending on interviewee preference. The LINDDUN threat categories were presented as follows:

- **L**inkability: the lack of unlinkability of items of interest, e.g., if multiple actions (register, authenticate, etc.) can be associated (not necessarily directly to a user, see next category).
- **I**dentifiability: linking an item of interest, such as an action, to the subject (user).
- **N**on-repudiation: the user cannot deny having performed the action.
- **D**etectability: that an attacker can detect that an action has taken place.
- **D**isclosure: threat related to the content of the action (registration data, authorization data, etc.).
- **U**nawareness: that the user's unawareness of what happens during an action is a threat.
- **N**on-compliance: that an action or part of it is non-compliant to, e.g., the GDPR, privacy policies, or legislation/regulation related to the domain (healthcare).

When explaining the user-centric setting from Section 4.2, the external IdP was set to Facebook and the relying party to a healthcare provider like a hospital. As intended, due to the selection of Facebook as the external IdP, interviewees were likely to view this entity as a potential privacy threat. For the first action, *registration*, we slowly went over each LINDDUN threat category in more detail to familiarise the interviewee. The order of actions was: *register* → *authenticate* → *get/set data* → *authorize* → *notify* → *search* → *ext. IdP authenticate* → *ext. IdP data import*.

5.2 Results

We consider in total 22 threats, where twelve are high priority (Table 2) and ten medium priority (Table 3). The remaining low priority threats are not presented here and will not be considered further in this deliverable. The priorities were derived by aggregating the results from the three interviews, first constructing relative rankings of threats for each action by each interviewee, normalizing the relative ranks, then merging the resulting ranks from all interviewees. We set a conservative breakpoint for the low-to-medium ranks, erring on the side of caution to not exclude relevant threats. In general, threats that all three interviewees ranked highly are high threats, while low threats were consistently ranked low by all three interviewees. Annex B contains further details and rationale for how each threat below was constructed. Note that several of the threats cover multiple LINDDUN categories due to how the threats were linked by the interviewees. For example, one common theme is the merger of linkability and identifiability, where the threat of profiling users' actions (linking together multiple actions such as registrations) is likely to lead to identification of the subject with widely available side-information, and such profiling is clearly possible if users are identifiable.

Table 2 and Table 3 present the high and medium priority threats, respectively. The tables have four columns: ID, Action, Categories, and Threat description. The identifier is constructed by combining the action (from the user-centric model) with one of the threat categories. For example, the *R-ID* threat is for the registration action and the LINDDUN threat category identifiability. The action column lists the relevant action from the user-centric model together

with a reference to the figure depicting the action in the previous chapter, while the categories column lists one or more LINDDUN threat categories. If more than one category is listed, Annex B explains why. Finally, each threat has a description that summarizes the threat. The description is based on how the threats where described by the interviewees in the interviews. Note that health-related threats are prominent in the descriptions due to the relying party in the user-centric model being described as a healthcare provider.

Table 2: High priority threats

ID	Action	Categories	Threat description
R-ID	Register, see Figure 8	Linkability and identifiability	When a user registers multiple accounts, presumably this is a form of identity management, so if accounts can be linked and the owner identified is a threat.
AC-ID	Authenticate, see Figure 9	Linkability, identifiability, and disclosure of information	When authenticating, entities that can identify the user and link multiple authentications have the capability to profile the user. Given examples are the Wallet, multiple relying parties (e.g., the same user is authenticating to Facebook and the anonymous counselling service), and other users (such as, for the health related scenario, spouses or family in domestic abuse cases). This also includes the case that the contents of the authentication messages enable this profiling.
AC-NC	Authenticate, see Figure 9	Non-compliance	The authentication process is likely non-compliant (at least in spirit with the GDPR) since many business models rely on being a third-party (like the Wallet) that tracks and profiles users for advertisement or related business models.
GS-ID	Get/set data, see Figure 10	Linkability and identifiability	Usage patterns when performing read/write operations on storage, if they are linkable to an identifiable user, are likely to reveal a lot information about users. For example, in the healthcare scenario, the magnitude of data (MRI scan compared to text entry) and frequency of operations leak information about health status and potential treatments/illnesses.

ID	Action	Categories	Threat description
AZ-DI	Authorize, see Figure 11	Disclosure of information	For authorizations, the set of permissions to a user's Wallet may reveal highly sensitive information about the user, depending on the granularity of permissions. For example, if the permission is given to a hospital relying party, this may reveal that the user is a patient at the hospital. If permissions are given to individual care-givers, this is significantly more invasive. Even allowing the Wallet to have knowledge of the authorizations should ideally be avoided.
N-ID	Notify, see Figure 12	Linkability and identifiability	For notifications, any party in the position to link notifications and identify the recipient/sender can profile and track users/relying parties.
S-ID	Search, see Figure 13	Linkability, identifiability, and detectability	Being able to link search queries (and their replies) together enables profiling, especially if linked to identity. Detecting that a search is going on can be enough in some settings.
S-UA	Search, see Figure 13	Unawareness	Depending on how the search feature is implemented, users are likely to be unaware of the details and implications of searching and making their identity searchable (if configurable).
EA-UA	External IdP authenticate, see Figure 14	Unawareness	Given the complex setting, users are likely unware of the implications of the action (external IdP profiling, Wallet profiling, what content is transmitted from external IdP to Wallet).
EA-NC	External IdP authenticate, see Figure 14	Non-compliance	Given the complexities of using an external IdP to authenticate and the business models of popular external IdPs today (profiling), that some party is not legally compliant is likely (e.g., if a user involves Facebook as part of authenticating to a healthcare related service, are there any legal obligations under (inter)national law for Facebook or the Wallet in processing this?)

ID	Action	Categories	Threat description
EI-UA	External IdP data import, see Figure 15	Unawareness	When importing data into the Wallet from an external IdP, the setting is complex and users are likely not aware of data flows or what exactly is being imported into the Wallet.
EI-DI	External IdP data import, see Figure 15	Disclosure of information	The external IdP might not be very clear with what exactly is being imported by the user into his/her Wallet. For example, the external IdP might leave around seemingly (to the user) arbitrary identifiers that might later be disclosed to relying parties and leak information about the user.

Table 3: Medium priority threats

ID	Action	Categories	Threat description
R-UA	Register, see Figure 8	Unawareness	When registering an account, users are likely to be unaware of the protections provided by the Wallet against different threats (active vs passive attacker, Wallet capabilities, etc.).
GS-DI	Get/set data, see Figure 10	Disclosure of information	When data is read or written, despite being encrypted the lack of adequate padding or other features of the content may disclose information.
GS-NC	Get/set data, see Figure 10	Non-compliance	The detailed usage patterns, potentially linked to a user, are likely (sensitive) personal data yet also likely to not be covered by privacy policies or attempts made to make users aware of this data leak and usage.
N-DI	Notify, see Figure 12	Disclosure of information	Any content of notifications, even if details are removed (e.g., “you have a new authorization request”), could enable profiling by entities with access to the content.

ID	Action	Categories	Threat description
S-DI	Search, see Figure 13	Disclosure of information	When searching, the contents of your search queries and their replies may reveal sensitive information. For example, if a user is searching for competing relying parties (like a new healthcare provider), other users (like doctors with particular specialities), or particular types of relying parties (like lawyers specialising in divorce cases).
S-NR	Search, see Figure 13	Non-repudiation	Depending on if other threats related to search are adequately mitigated or not, non-repudiation of having performed particular searches (or received a particular response to a query) can be used against the user.
S-NC	Search, see Figure 13	Non-compliance	Depending on how the search functionality is implemented and the varying degrees of information leakage to different parties, the use of the search in the context of health is likely swamped with complications.
EA-LI	External IdP authenticate, see Figure 14	Linkability and identifiability	Since many external IdPs have as a business model to profile its users, given the capability to, they would likely profile the relying parties the user authenticates to.
EI-NC	External IdP data import, see Figure 15	Non-compliance	Importing data from an external party into the Wallet for future use towards relying parties is a complex operation, if anything, because we don't know the intended use. Is, e.g., Facebook in any way liable if data from it is used in medical contexts or for insurance?
EI-LI	External IdP data import, see Figure 15	Linkability and identifiability	Since many external IdPs have as a business model to profile its users, given the capability to, they would likely profile the relying parties the user shares data with and in any way interacts with.

6 Mitigation of Privacy Threats

In this chapter, we first introduce the requirements that came from the project in terms of technologies that were evaluated and assessed to be integrated. Next, we provide mitigation strategies for the identified high- and medium-risk privacy threats from the previous chapter, and considering both, we also identify corresponding mitigation mechanisms.

6.1 Technology Requirements from Deliverable D4.1

In the means of a practical and interoperable CREDENTIAL system, Deliverable D4.1 [4] provides technology evaluations and recommendations that are used within CREDENTIAL for a later development stage.

Core cryptographic technologies that are addressed and evaluated in D4.1 are:

- **Secure data sharing** – In particular, classical Proxy Re-Encryption (PRE) and conditional PRE (c-PRE) is recommended by D4.1 for the use within CREDENTIAL.
- **Authenticated data disclosure** – In particular, Redactable Signatures (RS) are recommended by D4.1 for the use within CREDENTIAL and further research is needed on Privacy-enhancing Attribute-Based Credentials (Privacy-ABCs).

Additional cryptographic technologies that are addressed and evaluated in D4.1 are:

- **Authentication** – D4.1 does not give any recommendations on authentication technologies such as TPASS since this is not at the core of CREDENTIAL.
- **Access to encrypted data** – In particular, Searchable Encryption (SE) is recommended by D4.1 for the use within CREDENTIAL.
- **Other technologies** – D4.1 considers Secret Sharing and Verifiable Computation for potential future revisions of the CREDENTIAL Wallet. For the time being, D4.1 does not explicitly recommend Attribute-Based Encryption (ABE), Fully Homomorphic Encryption (FHE), Privacy-ABCs, Private Information Retrieval (PIR), Oblivious RAM (ORAM) to be used within CREDENTIAL, mostly because of efficiency and compatibility reasons. For a detailed evaluation of the technologies, we refer to the corresponding sections of D4.1 [4].

Furthermore, the use of Provable Data Possession (PDP) and Proofs of Retrievability (PoR) techniques are not at the core interest of CREDENTIAL as they focus on auditing features rather than authentication or data sharing.

Authentication-to-the-cloud technologies that are addressed and evaluated in D4.1 are:

- **Authentication technologies** – FIDO UAF and OAuth are recommended by D4.1.
- **Underlying Technologies for authentication** – D4.1 suggests that the technology of Trusted Platform Module (TPM) should (if at all) only be used to encrypt the secret

key of the PRE scheme and needs further evaluations. Furthermore, D4.1 recommends Trusted Execution Environment (TEE) to use within CREDENTIAL.

Remartk: even though D4.1 recommends FIDO UAF and TEE, both technologies need further investigations.

Identity and access management protocols that are addressed and evaluated in D4.1 are:

- **Identity protocols** – SAML and OpenID Connect are recommended by D4.1 for use within CREDENTIAL.
- **Authorization protocols** – UMA and OAuth2 are recommended by D4.1 for use within CREDENTIAL.
- **Policies** – For access control standard, D4.1 recommends XACML.

6.2 Mitigation Strategies

LINDDUN identifies strategies for the threats, but assumes that all the threats are mitigated, namely avoided. LINDDUN defines the following mitigation strategies:

1. Warn the user
2. Remove or turn off a feature
3. Counter threats with preventive or reactive privacy enhancing technologies

However, in a real life setting, similar to a security assessment, there are limitations that inhibit risk avoidance, or some risks cannot be completely mitigated. In this regard, industry established standards, such as ISO/IEC 27005 [3] define mitigation strategies for the risks instead of the threats. Furthermore, there is a difference between the categories of mitigation strategies in LINDDUN compared to the industry standards for general IT Security.

In this regard, we found especially the mitigation strategies from ISO/IEC 27005 useful, which are comparable to those of BSI 100-3, and are split into the following strategies, also shown in Figure 16:

1. Risk reduction
2. Risk retention
3. Risk avoidance
4. Risk transfer

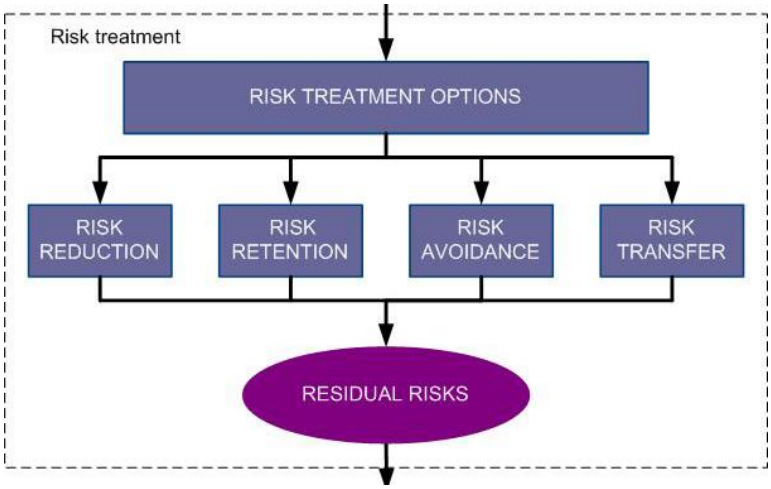

Figure 16: Risk treatment options (strategies) according to ISO/IEC 27005 [3]

6.3 Mitigation Mechanisms (MM)

Finally, besides the decision on the risk mitigation strategies, tha last step in the privacy analysis process is the identification of mechanisms to mitigate the identified privacy threats. These mechanisms are then to be integrated in the CREDENTIAL Wallet architecture, followed during the development of the architecture, as well as after the architecture is being used, e.g. in the pilots.

As a result of the privacy threat analysis, we came up with a total of propose 29 different mitigation mechanisms. A summary of all the recommended mitigation mechanisms is shown in Table 4. However, we can categorize these mitigation mechanisms in the following three main groups, namely into *technology-based mitigation mechanisms* that rely on the integration of privacy-enhancing technologies; *development principles and guidelines*, requiring particular features to be developed considering privacy; and *mechanisms that require the application of organizational and contractual measures*. These are explained in the following sections.

Table 4: An overview of all proposed mitigation mechanisms

Mitigation Mechanism ID	Mitigation Mechanism (MM)
MM1	Tor [7] or comparable network anonymization tool
MM2	"1-hop" onion service
MM3	Define a privacy policy that *limits the purpose ofr the use of the data*
MM4	Define a policy that *requires minimal information to be disclosed (data

	minimization)*.
MM5	Enforce transaction-based pseudonyms
MM6	Avoid persistent identifiers from the wallet to the RP
MM7	UnlimitID [9]
MM8	Control access to the authentication logs, e.g. require the app to be "unlocked" using owner credentials again in order to access this feature.
MM9	Make sure the privacy policy is compliant with the regulation
MM10	Make sure to double check that the privacy policy is enforced by the system through adequate, documented processes.
MM11	Offer Private Information Retrieval (PIR)
MM12	Searchable Encryption
MM13	Apply need-to-know principle. Limit the use of third-party notification service to only the minimal needed scenarios.
MM14	AnoNotify [8] or comparable service, which hides the relation between the notification server and the notified user AND send no contents over the notification server.
MM15	UI to offer transparency over what the wallet may learn about the user based on the search terms / functionality, and how this is used in a system
MM16	Define a privacy policy that documents the MM15
MM17	Provide IdP-risk aware User Interface 1
MM18	Define a privacy policy that documents the MM17
MM19	Define a policy / contract (SLA-like) with the IdP, making sure they comply with the legal / data processing requirements of the project, with the privacy policy, and any other policy of regulation that the project complies with, which applies to the IdP.
MM20	Provide IdP-risk aware User Interface 1
MM21	Define a privacy policy that documents MM20

MM22	Encrypt the data using PRE (Proxy Re-Encryption) for the User, avoiding the wallet to see the values, unless needed.
MM23	Make available a clear, understandable, simple policy where the users get informed about the guarantees and potential threats from using the CREDENTIAL wallet.
MM24	Split messages into same size packets, apply padding.
MM25	Define a privacy policy that includes clear statements about what inference can be made about the users based on their usage behaviour.
MM26	Make searchs terms digitally signed (plausible deniability), provide the user the chance to delete/disable search history.
MM27	Define in the privacy policy what information is leaked to whom when searching AND limit it on a need-to-know basis (minimal disclosure).
MM28	Define a privacy policy through a clear statement that limits the purpose for which the data import functionality can be used AND add a clausel in the SLA with the SPs to limit the purpose of the data import.
MM29	Redactable digital signatures

6.3.1 Technology-based Mitigation Mechanisms

Considering the goals of CREDENTIAL as a project and the privacy goals of LINNDUN, our preferred method for mitigating identified risks is by proposing privacy-enhancing technologies, wherever this is feasible. In the following, we list the main technologies that we recommend to be integrated in the architecture of the CREDENTIAL Wallet in order to mitigate the identified risks:

1. "The onion router" (Tor) [7] or comparable network anonymization tool – Tor uses "a group of volunteer-operated servers that allows people to improve their privacy and security on the Internet" [7]. It allows CREDENTIAL users to connect "through a series of virtual tunnels rather than making a direct connection, thus allowing both organizations and individuals to share information over public networks without compromising their privacy." [7] We consider this or similar tool to be adequate for partially mitigating linkability and identifiability threats. Tor anonymizes both the identity of the user and of the service being accessed. This corresponds to MM1 in Table 4.

2. A „1-hop" onion service – It is a modified version of Tor, but enables more efficient performance, since it only hides the identity of the user, but not of the network. Hence, referred to as "1-hop". This tool can therefore be used instead of Tor for partially mitigating identifiability threats. This corresponds to MM2 in Table 4.
3. UnlimitID [9] – it is an enhancement (extension) for OpenID Connect which enables unlinkable pseudonyms for users. In OpenID Connect, which is integrated in our project (see Section 6.1), an honest but curious identity provider can link users with service providers and collect data on the user's behavior, which allows the IDP to build profiles. UnlimitID prevents the IdP from learning this information by allowing the user to create and use service specific pseudonyms; these pseudonyms cannot be traced back to the original identity. UnlimitID extends the OpenID Connect protocol with additional steps that rely on anonymous attribute-based credentials and zero knowledge proofs.[4] This tool corresponds to MM7 in Table 4.
4. Private Information Retrieval (PIR) – a PIR scheme enables the user to retrieve some item from the Wallet without the Wallet learning which item was retrieved. This tool corresponds to MM11 in Table 4.
5. Searchable Encryption – It enables the user to search through a list of encrypted documents without the Wallet having to decrypt those documents. This requires specific cryptographic schemes for searchable encryption. This tool corresponds to MM12 in Table 4.
6. AnoNotify [8] – is a private and timely notification service that prevents linking notification subscribers to notification publishers. In conventional notification services, the notification server can observe the content and the routing of notifications. AnoNotify eliminates this privacy threat by building an infrastructure based on anonymity networks, bloom filters, and database shards. However, in our project, we discussed also additional options, and AnoNotify is just one recommendation. As described in D5.1 [2], CREDENTIAL will most likely go for the MQTT [9], which is an OASIS-standard, and

[4]We are also aware that it not that easy to integrate UnlimitID with our current CREDENTIAL architecture, because there is a major difference between CREDENTIAL and UnlimitID. While in CREDENTIAL architecture the user data are stored in the Wallet (cloud), UnlimitID stores them locally at the user's device. Nevertheless, we recommend to investigate this or some other possibility further in the implementation of the Wallet.

which runs integrated in the Wallet, thereby avoiding third party risks. Nevertheless, we find it important to consider the privacy goals that the notification service should aim, in this case minimal information disclosure. This tool corresponds to MM14 in Table 4.

7. Proxy Re-Encryption (PRE) – to enable the User to share encrypted messages / documents stored on the wallet to be re-encrypted by the Wallet for a Service Provider (to which the User agrees sharing) without the Wallet learning the contents of the encrypted messages / documents. This tool corresponds to MM22 in Table 4.
8. Redactable digital signatures – to enable minimal disclosure of information when using identity services for digitally signed credentials. Similar to "traditional" digital signatures, these schemes allow signing and verification. However, unlike "traditional" digital signatures, these schemes allow the User to be able to get a signature verified even if parts of the original signed messages are hidden. This tool corresponds to MM29 in Table 4.

6.3.2 Development Principles and Guidelines

Besides integration of special technologies that enable the mitigation of certain privacy threats, some threats can be mitigated by following certain system development principles and guidelines. We recognize the following in this category:

1. Enforce transaction-based pseudonyms – so that each transaction is assigned a new pseudonym, making it unlinked to a previous one. This way we do not need to link different transactions of the user to a single persistent pseudonym. However, it should be made possible to choose otherwise, if the user wishes to. This corresponds to MM5 in Table 4.
2. Avoid persistent identifiers from the Wallet to the Service Provider – Besides enforcing transaction-based pseudonyms, the Wallet should also avoid other types of identifiers that could potentially link different transactions of the user. This corresponds to MM6 in Table 4.
3. Enforce access control to the authentication logs – e.g. require the app to be "unlocked" using owner credentials again in order to access this feature. This would prevent unauthorized users who may accidentally get access to the app / phone of the user to see those logs, which may reveal some private transactions. This corresponds to MM8 in Table 4.
4. Document the enforcement of privacy policy – Make sure to double check that the privacy policy is enforced by the system through adequate, documented business processes. Some guarantees should be in place that document how the promised privacy policy is enforced in a way that it can be checked for consistency and completeness. This corresponds to MM10 in Table 4.
5. "Pull" instead of "push" notifications – A user can manually check on his client and retrieve any notifications directly from the Wallet, so this does not require an explicit, separate notification mechanism. This does not mean no push notifications to be implemented, but rather to only do that on a "need to push" basis considering the urgency of the notification. Hence, the concrete application scenario must be analysed, and a notification service must not be present for all new events that are in the system for a user. This corresponds to MM13 in Table 4.

6. Design a UI that is transparent to the user – It should enable the user to see what the Wallet may learn about the user based on the search terms / functionality, and how that information is used in the system. This corresponds to MM15 in Table 4.
7. Design a authentication aware UI – offer the possibility to the user to see what the Wallet and the Identity Provider (IdP) can learn from their authentication(s) using the external IdP. This corresponds to MM17 in Table 4.
8. Provide IdP-risk aware UI – that offers them to see what the Wallet and the IdP can learn about them by observing the data import feature from the external IdP. This corresponds to MM20 in Table 4.
9. Split messages into same size packets, apply padding – This should prevent a malicious attacker (even the Wallet) to infer information about the contents of the messages by simple observing their size. Splitting the messages into same size packets and applying padding is one way to mitigate this threat. This corresponds to MM24 in Table 4.
10. Implement the search using no digitally signed searches – search terms should not be digitally signed (provide plausible deniability) and enable the user the chance to effectively delete / disable search history. This corresponds to MM26 in Table 4.

6.3.3 Organizational- and Process-based Mitigation Mechanisms

Finally, the third category of the mitigation mechanisms requires special organizational efforts when the system is up and running. Mostly this includes special considerations in the privacy policy, or special contracts with other entities. A summary of the main mitigation mechanisms that resulted from our analysis consists of the following:

1. Define a privacy policy that clearly limits the purpose of data usage. This can be used to partially mitigate identifiability threats and disclosure of information. This corresponds to MM3 in Table 4.
2. Define a policy that requires minimal information to be disclosed (data minimization). This corresponds to MM4 in Table 4.
3. Make sure the privacy policy is compliant with the applicable regulation(s). A special focus should be made to the tracking of the user activities by third parties and using such information for advertisement or related business processes. This corresponds to MM9 in Table 4.
4. Define a privacy policy that correctly documents the what the Waller or some other service can learn about the User from the use of Search functionality. This corresponds to MM16 in Table 4.
5. Define a policy/ contract (SLA-like) with the IdP, making sure they comply with the legal / data processing requirements of the project, with the privacy policy, and any other policy of regulation that the project complies with, which applies to the IdP. This corresponds to MM18 in Table 4.
6. Define a privacy policy that correctly documents potential inferences that an external IdP could make about the User from using the authentication through an external IdP. This corresponds to MM19 in Table 4.
7. Define a privacy policy that correctly documents what the Wallet can learn about the User by observing the “Data import from external IdP”. This corresponds to MM21 in Table 4.

8. Define and a clear, understandable, simple policy where the users get informed about the privacy guarantees and potential privacy threats from using the CREDENTIAL Wallet in terms of the information that can be learnt, shared, or saved about the User. This corresponds to MM23 in Table 4.
9. Define a privacy policy that includes clear statements about what inference can be made about the users based on their usage behaviour. This corresponds to MM25 in Table 4.
10. Define in the privacy policy what information is leaked to whom when searching and limit it on a need-to-know basis (minimal disclosure). This corresponds to MM27 in Table 4.
11. Define a privacy policy through a clear statement that limits the purpose for which the data import functionality can be used. In addition, define clearly in a contract with the Service Provider(s) to limit the purpose for which the data import feature can be used. This corresponds to MM28 in Table 4.

6.4 Mapping Privacy Threats to Mitigation Strategies and Mechanisms

Previously we provided a list of the mitigation mechanisms, which were subsequently categorized and described in Section 6.3. In the following, we will show explicitly how the list of identified threats maps to the a mitigation strategy and to (one or more) corresponding mitigation mechanisms. Table 5 maps the mitigation strategies and mitigation mechanisms for the high-risk threats from Table 2. The fourth column defines an identifier of the mitigation mechanism proposed, whereas concrete mitigation mechanisms corresponding to this identifier is presented earlier in this document in the overview list of mitigation mechanisms from Table 4.

Table 5: Mitigation strategies for high-risk threats

ID	Threat category	Mitigation strategy	Mitigation Mechanism (MM)	Limitations / Comments
R-LI	Linkability and identifiability	Risk avoidance	MM1, MM2, MM3	MM2 enables more efficient performance and better security, since it only hides the identity of the user, but not of the network. Hence, referred to as "1-hop".
AC-LI1	Linkability and identifiability	Risk reduction	MM4, MM5, MM6	We limit the analysis only to the architecture level. We omit the application specific / browser specific threats. Privacy-ABC technologies are excluded from the project (see Requirements in Section 6.1). Hence, risk reduction.

ID	Threat category	Mitigation strategy	Mitigation Mechanism (MM)	Limitations / Comments
AC-LI2	Linkability and identifiability	Risk reduction	MM7	Considering that we will use OpenID and OAuth, which by default link and observe user authentications, we would at least protect from this threat by integrating UnlimitID. We could investigate the proposal by some project partners on implementing encrypted attribute-based credentials if feasible.
AC-LI3	Linkability and identifiability	Risk avoidance	MM8	We have to assume that the users will not share their credentials that are used to authenticate to the mobile app / logs. Assess the usability impacts.
AC-DI1	Disclosure of information	Risk reduction	MM2, MM20, MM29	Privacy-ABC technologies are excluded from the project (see project requirements in Section 6.1). Hence, risk reduction.
AC-DI2	Disclosure of information	Risk reduction	MM3, MM29	
AC-DI3	Disclosure of information	Risk reduction	MM3, MM8	A user may nevertheless want to disclose some information to other users by their own will. This is beyond our requirements.
AC-NC	Non-compliance	Risk reduction	MM9, MM10	
GS-LI	Linkability and identifiability	Risk reduction	MM11, MM12	ORAM and PIR impractical (see project requirements in Section 6.1).
AZ-DI	Disclosure of information	Risk reduction	MM3/MM9, MM20	The project has already decided to use the XACML protocol (see project requirements in Section 6.1). Hence, we have to provide soft protection, namely by policy.
N-LI	Linkability and identifiability	Risk avoidance	MM13, MM14	

ID	Threat category	Mitigation strategy	Mitigation Mechanism (MM)	Limitations / Comments
S-LI	Linkability	Risk reduction	MM1, MM11, MM12	
S-ID	Identifiability	Risk reduction		
S-DE	Detectability	Risk retention		
S-UA	Unawarenes	Risk reduction	MM15, MM16	Risk can only be reduced if the user wants to be informed by clicking on the respective UI components. Hence, risk reduction is selected.
EA-UA	Unawarenes	Risk reduction	MM17, MM18	Risk can only be reduced if the user wants to be informed by clicking on the respective UI components. Hence, risk reduction is selected.
EA-NC	Non-compliance	Risk avoidance	MM19	
EI-UA	Unawarenes	Risk avoidance	MM20, MM21	
EI-DI	Disclosure of information	Risk avoidance	MM22	

Similarly, for the medium-risk threats, we have chosen for each risk the appropriate mitigation strategy, and where applicable, the corresponding mitigation mechanism(s). This is shown in Table 6.

Table 6: Mitigation strategies and mitigation mechanisms for medium-risk threats

ID	Threat category	Mitigation strategy	Mitigation Mechanism (MM)	Limitations / Comments
R-UA	Unawareness	Risk reduction	MM23	Risk can only be reduced, as it requires the user to want to become aware, not feasible to force the user to read the privacy policy.
GS-DI	Disclosure of information	Risk reduction	MM24	Project requirements (see Section 6.1) rule out more privacy-friendly solutions, such as PIR or ORAM. Hence, risk is only reduced by the mechanism that we suggest, but not avoided.
GS-NC	Non-compliance	Risk reduction	MM25	We can only reduce this risk by identifying it as such in the privacy policy.
N-DI	Disclosure of information	Risk avoidance	MM13, MM14	
S-DI	Disclosure of information	Risk reduction	MM12	This does not protect from searching for other participants though, only to searched tags, which is a category search more or less.
S-NR	Non-repudiation	Risk reduction	MM26	This reduces the risk that it is used in a court case.
S-NC	Non-compliance	Risk avoidance	MM27	
EA-ID	Identifiability and linkability	Risk avoidance	MM4, MM5, MM6	Privacy-ABC technologies are excluded from the project (see Section 6.1). Hence, risk reduction.

ID	Threat category	Mitigation strategy	Mitigation Mechanism (MM)	Limitations / Comments
EI-NC	Non-compliance	Risk transfer	MM28	
EI-ID	Identifiability and Linkability	Risk reduction	MM4, MM5, MM6	

7 Recommendations on Privacy-Enhancing Mechanisms

In this chapter, we summarize the main lessons that the CREDENTIAL project has learnt and should keep implementing in order to further fulfill the aimed privacy-by-design approach. However, these lessons should also be useful for other projects that follow a similar goal. We split our recommendations in two main parts. In the first part, we target system architects and developers, whereas in the second part we provide some concrete recommendations for the regulatory organisations, which should help enforce some of the mitigation mechanisms that improve the privacy protection of the consumers in and potentially beyond the European Union.

7.1 Recommendations For System Architects And Developers

Privacy-enhancing Technologies (PETs) helps to achieve compliance with protection of personal data [10]. PETs aim to help users to make informed choices and also, provide control over personal data sent to or used by service providers. Use of PETs makes it technologically more difficult to carry out certain breaches resulting in invasions of privacy [11].

Among the main recommendations for system architects, we can list the need for data minimization by design, and making information flows transparent using system diagrams.

7.1.1 Adapt Methodologies to Your Scenario

In our experience, it was not straightforward to select a methodology for doing the privacy analysis. It was clear to us already in the beginning that none of the existing methodologies are mature enough to be directly applied in an engineering project, such as ours. However, we still consider that agreeing on a certain methodology to guide the privacy analysis is useful, and instead of applying it blindly, you have to learn to adapt the methodology to your needs. We have reported on the methodology and the needed deviations in Chapter 2. However, we had a strong privacy expertise in the project, which made it possible, and which may not always be the case in every project. Nevertheless, the conclusion is that whoever is responsible for the privacy analysis should keep in mind that any method is a tool aimed at achieving its goal, and adapting whenever necessary is the key to an effective analysis.

7.1.2 Apply “Need To Know” as a Privacy Design Principle

Data minimization is a fundamental step towards data protection. The “need to know” principle that is applied to security systems for access control can also be adapted for privacy. A typical measure to prevent misuse of personal data is by eliminating or reducing collection of personal data and by preventing undesired processing of personal data [12]. This can be achieved by separating authorization from authentication. Design decisions on inbound and outbound system interfaces and use of cryptographic schemes are a few choices that can support privacy expectations. Furthermore, this can also be applied, as we saw, to other services, such as the notification service in our analysis. Such privacy threats can be mitigated by limiting data flows through third parties, e.g. only providing notifications when strictly required, such as a doctor urgently informing a patient about some test result, or requiring the user to share some document.

7.1.3 Make Information Flows and Storage Points Transparent Using System Diagrams

Our analysis showed that it is useful to use system diagrams to perform privacy analysis. However, at times, these diagrams can get complex. Nevertheless, this is not an excuse in order to omit important information flows in the system, e.g. in a DFD. Important for privacy are the data flows outside of the trust boundaries of the system, and data flows from one entity to another one, especially the not-so-obvious data flows to third parties, such as brokered-notification service that we reported on.

7.1.4 Document Mitigation Strategies and Limitations

Not all privacy threats will be mitigated. However, it is important to identify as many as possible. A privacy analyst needs feedback from the system designers and developers on many aspects of the project, which may not be straightforward. A clean design requires to list the potential privacy implications of a certain design choice, technical limitation, third party app, etc. Furthermore, what measures were taken to address the privacy threats (implications) (if any), what could be alternatives (for future projects or updates to existing systems).

7.2 Recommendations for the Regulatory Bodies

Regulatory bodies, such as data protection authorities, which are responsible for supervising institutions and companies on the compliance with the applicable regulation, can learn from our experiences. The following shows the main lessons, which we assume to be most relevant for the regulatory bodies.

7.2.1 Check For the Completeness of Information on Privacy Policies

A privacy policy should define all the ways the personal data of the users are treated within the system of a given company, the parties that have access to them, data retention and processing policies, purpose of use, etc. All of these are listed in the regulation. However, since privacy policies are usually long and use legal terms, it is challenging to verify the completeness. However, the goal should be to be able to clearly identify the entities that have any access and the types of data they process, relevant contracts with third parties, etc.

7.2.2 Require Informative Privacy Policies

Besides completeness, users should also be able to have access to a simple version of the privacy policies. It is clear that completeness can contradict clarity, and it is not easy to achieve both, but that should be the final goal. A privacy policy should not only make sure to comply with the regulation, but also for the users to be able to get informed about how their information is treated.

Users should be informed whenever personal data is processed [10]. They should be notified about which type of data is processed (name, date of birth, social security number, health records, credit score, etc.), for what purpose, and by whom (service provider). The technical complexity of identity management systems makes it difficult to assume that an average user will be able to understand the system and make informed decisions. Hence, it is important to provide easy and understandable information to achieve informed consent about sharing user data.

Furthermore, user interfaces should provide the possibility for users to see whenever some access is requested to their data, if they wish so. At best, system settings should be configurable to enable this (as not all users would like to constantly see each data flow, because that would kill usability in many scenarios).

As an example, let us consider the following: Instead of a push notification saying "*You have a new authorization request from serviceX*", it is more informative to say "*Do you want to share data with service?*". If the user agrees, allow the user to choose the items to share with serviceX. This prevents user from providing too much information (which they may ignore) and also prevents causing wrong decision or actions.

7.2.3 Empower Users with Control over Their Personal Data

Even if users consent to the initial collection of their personal information, they must also be given a mechanism to specify whether or not to consent to the future use of that information. To provide users with higher transparency and control over their data, a dashboard could be used to summarize data used by service providers and provide links to personal setting.

7.2.4 Make Explicit Contracts with Third Parties

As we saw, not all risks could be mitigated through technical means. Some threats need to be mitigated via contracts or other legally binding documents instead. As we saw, we had a number of proposal for mitigation mechanisms in our document that required special considerations to be documented in privacy policies, or contracts to be made with third parties. This is especially important for risks that we transfer to third parties. Therefore, these need not only be properly documented, but also carefully analyzed when doing audits, in order to check conformity to the privacy policies, how this is enforced in the contracts with third parties, and what would happen in case of one party not complying.

References

[1] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements," *Requir. Eng.*, vol. 16, no. 1, pp. 3–32, 2011.

[2] N. Notario (ed.), "D5.1 Functional Design", March 2017 (CREDENTIAL Deliverable).

[3] "ISO/IEC 27005 Information technology -- Security techniques -- Information security risk management," ISO copyright office Case postale 56 • CH-1211 Geneva 20, 2008.

[4] F. Hoerandner (ed.), "D4.1 Assessment report on cryptographic technologies, protocols and mechanisms", March 2017 (CREDENTIAL Deliverable)..

[5] A. Pfitzmann and M. Hansen, "A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management." 2009.

[6] K. Wuyts and W. Joosen, "LINDDUN privacy threat modeling: a tutorial," 2015.

[7] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," *SSYM'04 Proc. 13th Conf. USENIX Secur. Symp.*, vol. 13, p. 21, 2004.

[8] A. Piotrowska, J. Hayes, N. Gelernter, G. Danezis, and A. Herzberg, "AnoNotify: A Private Notification Service.," *IACR Cryptol. ePrint Arch.*, vol. 2016, p. 466, 2016.

[9] OASIS, "MQTT Version 3.1.1," *OASIS Stand.*, no. December, p. 81, 2015.

[10] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. Le Metayer, R. Tirtea, and S. Schiffner, *Privacy and Data Protection by Design - from policy to engineering*, no. December. 2015.

[11] "Privacy Enhancing Technologies (PETs)," 2007.

[12] M. Hansen, J.-H. Hoepman, and M. Jensen, *Readiness Analysis for the Adoption and Evolution of Privacy Enhancing Technologies*, December. 2015.

Annex A – Mapping Privacy Threats to CREDENTIAL DFD Elements

Table 7 shows the complete list of the mapping of privacy threats to the general Wallet DFD elements from Figure 4 following instructions from LINDDUN [1]. The threats outside of the CREDENTIAL trust boundary are highlighted and considered in the subsequent steps of LINDDUN.

Table 7: Mapping privacy threats to general Wallet DFD elements

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	Threat target	**L**	**I**	**N**	**D**	**D**	**U**	**N**
Data store	DS1 Account info	x	NC	x	x	x		
	DS2 Attribute Store	x	x	x		x	x	x
	DS3 Policies Store	x						
	DS4 Personal Trust Store	x	x	x		x	x	x
Data flow	DF1 Registration data (E2 Participant)	x	x	x	x	x	x	x
	DF2 Registration data (E1 Owner)	x	x	x	x	x	x	x
	DF3	x	x	x	x	x	x	x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	Registration data (E1 Owner or E2 Participant)							
	DF4 Identity/attributes (E1 Owner)	NC	x	x	x	x	x	x
	DF5 Identity/attributes (E1 Owner)	x	x	x	x	x	x	x
	DF6 Identity/attributes (E3 Identity & Attribute Provider)	x	x	x	x	x	x	x
	DF7 Identity/attributes (E1 Owner to E3 Identity & Attribute Provider)	x	x	x	x	x	x	x
	DF8 Identity/attributes (E3 Identity & Attribute Provider)	x	x	x	x	x	x	x
	DF9 Identity/attri	x	x	x	x	x	x	x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	butes (E1 Owner)							
	DF10 Search Request (E1 Owner)			x		x		
	DF11 Matching owners (E1 Owner)			x	x			
	DF12 Registration data		x	x	x	x	x	x
	DF13 Registration data		x	x	x	x	x	x
	DF14 AuthZ Preferences (E1 Owner)			x				
	DF15 Authz Preferences			x				
	DF 16 Preferences			x				
	DF17 AuthZ Preferences			x				
	DF18 AuthZ Response			x				

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	DF19 Identity/attributes	x	x	x	x	x	x	x
	DF20 Identity/attributes	x	x	x	x	x	x	x
	DF21 Attributes	x	x	x			x	x
	DF22 Identity/attributes	x	x	x	x	x	x	x
	DF23 Request Identity/attributes	x	x	x	x	x	x	
	DF24 Identity/attributes	x	x	x	x	x	x	x
	DF25 Delegation preferences	x		x		x	x	x
	DF26 Delegation preferences	x		x		x	x	x
	DF27 Inform delegation request	x		x		x	x	x
	DF28 Encryption request	x	x	x		x		x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	DF29 Encrypted data	x	x	x		x		x
	DF30 Inform delegations	x		x		x		x
	DF31 Request delegation	x		x	x	x	x	
	DF32 Identity/attributes	x	x	x	x	x	x	x
	DF33 Identity/attributes	x	x	x	x	x	x	
	DF34 Identity/attributes	x	x	x	x	x	x	
	DF35 Identity/attributes	x	x	x	x	x	x	
	DF36 Identity/attributes	x	x	x	x	x	x	
	DF37 Notification	x	x	x				x
	DF38 Re-encryption request			x		x	x	x
	DF39 Re-			x		x	x	x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	encrypted attributes							
	DF40 Redact data	x	x	x		x	x	x
	DF41 Redacted data	x	x	x		x	x	x
	DF42 key	x	x	x	x	x	x	x
	DF43 key	x	x	x	x	x	x	x
	DF44 Notification	x	x	x				x
	DF45 Notification preferences	x	x	x				
	DF46 Decryption request	x	x	x		x	x	x
	DF47 Decrypted data	x	x	x		x	x	x
Process	P1 Offbounds provisioning	x	x	x	x	x	x	
	P2 Authentication	x	x	x	x	x	x	x
	P3 Registration	x	x	x	x	x	x	x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
	P4 Authorization	x	x	x	x	x	x	x
	P5 Identity & Attribute Management	x	x	x	x	x	x	x
	P6 Delegation	x	x	x	x	x	x	x
	P7 Offbounds Authentication	x	x	x	x	x		x
	P8 Search participants			x		x		
	P9 Encryption			x		x	x	x
	P10 Re-encryption			x		x	x	NC
	P11 Decryption			x		x	x	x
	P12 Audit Trailing	x	x	x	x	x	x	x
	P13 Notification	x		x	x	x		x
	P14 Redaction	x		x		x	x	NC
Entit	E1 Owner		x		x	x	x	x

		Linkability	*Identifiability*	*Non-repudiation*	*Detectability*	*Disclosure of information*	*Unawareness*	*Non-compliance*
y	E2 Participant		x		x	x	x	x
	E3 Identity & Attribute Provider		x					x

Annex B – Detailed Interview Results and Rationale

Below are eight tables with the aggregate results of our semi-structured interviews. There is one table for each action, and for each threat category in LINDDUN related to the particular action, the table contains the aggregated relative risk from the interviewees, and the resulting relative risk and priority. For each combination of action and threat category, the interviewees estimated the *impact* and *likelihood* of the threat on a scale from 1 (low impact / unlikely) to 5 (high impact / likely). This results in the *risk* associated with each threat (*impact*likelihood*). Because interviewees may use different scales and find some actions riskier than others (without this being clear in the data: the interviewees did not get an overview of all actions in the introduction), we first transform each risk assessment into a *relative risk per action*, e.g, "linkability during registration is a significantly higher threat than detectability".

With the relative risks of each threat per interviewee calculated for each action, we aggregate the relative risks. The aggregated relative risks are then normalized to values between 0-1 and the priority based on the relative risk. The median relative risk, after aggregation, is 0.77. Since our goal is to prioritize threats, we set the breaking point for low priority threats at 0.77: this deprioritizes half of the potential threats. Further, we set the breaking point between medium and high risk at 0.9, which upon casual inspection roughly splits the remaining threats into two equal halves. Our analysis in this deliverable considers both high and medium risk threats, so the selection of where we make a distinction between high and medium threats is less important.

The upside of our approach is that we are likely to capture high/low priority threats *for each action*, where each interviewee more or less consistently ranked a threat relatively high/low. The downside of our approach is that we lack insights into the relative priority *between actions*: e.g., there might be more privacy threats around authentication and authorization actions compared to the get/set data action. We believe that such a comparison between actions is not possible given how small our data set is (three expert interviews) and how we conducted the interviews. This needs to be addressed by us as part of our analysis. Further, another downside is the limited size of our data set: we cannot draw many conclusions more than for the threats where there is consensus between all experts. By only excluding about half of the potential threats (56 possible combinations) we err on the side of caution.

Next, we go over each action, its associate table with interview results, and motivate why we created the high and medium priority threats found in Table 2 and Table 3. Section 5.2 explains how the identifiers of the threats are constructed.

Table 8 contains the results for the register action. We merged the high priority threat of linkability of registrations with the medium threat identifiability into the **R-ID** threat because two out three interviewees linked the linkability threat with the identifiability threat during the interviews (paraphrasing: "users probably register multiple accounts for a reason, and if registrations are linkable, identifying the user is probably easy"). The unawareness threat was turned into the medium **R-UA** threat, focused on users not understanding or overestimating the privacy protections provided by the Wallet.

Table 8: Aggregate results for the *register* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2.59	1	High
Identifiability	2.2	0.85	Medium
Non-repudiation	1.32	0.51	Low
Detectability	0.8	0.31	Low
Disclosure of information	1.28	0.49	Low
Unawareness	2.15	0.83	Medium
Non-compliance	1.13	0.44	Low

Table 9 contains the results for the authenticate action. The linkability, identifiability and disclosure of information threats during authentication were merged into the high **AC-ID** threat. Once again, identifiability and linkability were linked by the interviewees. The disclosure of information threat is also merged because two interviewees associated the threat to disclosing that the user is authorizing to a spouse or significant other, and not necessarily the Wallet or relying party. In other words, different threat actors got associated to different threat categories. Finally, the threat of non-compliance got turned into another high **AC-NC** threat, primarily due to the business models of popular IdPs today (profiling) and the setting of the interview (both GDPR and the healthcare domain apply).

Table 9: Aggregate results for the *authenticate* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2.2	0.95	High
Identifiability	2.31	1	High
Non-repudiation	1.49	0.64	Low
Detectability	1.29	0.56	Low
Disclosure of information	2.11	0.91	High
Unawareness	1.26	0.55	Low

Non-compliance	2.13	0.92	High

Table 10 shows the results for the get/set data action. As for the prior actions, the linkability and identifiability threats were merged into the **GS-ID** high threat. One interviewee mentioned that usage patterns (linkability) are likely easily tied to a user (identifiability) with little side information. **GS-DI** is the resulting medium threat from disclosure of information during the get/set data action. Even if data is encrypted (a premise in the interviews due to the use of proxy re-encryption), adequate padding or usage patterns are likely not hidden, and the processing will likely be on sensitive medical data (from the example relying party in the interviews). Finally, the non-compliance threat was turned into the medium **GS-NC** threat with the motivation that the kind of data that leaks from this action is likely not covered by privacy policies or considered important during operations.

Table 10: Aggregate results for the *get/set data* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2.75	1	High
Identifiability	2.46	0.89	Medium
Non-repudiation	0.92	0.34	Low
Detectability	0.8	0.29	Low
Disclosure of information	2.46	0.89	Medium
Unawareness	1.71	0.62	Low
Non-compliance	2.43	0.88	Medium

Table 11 shows the results for the authorize action. Here, most interviewees found most threats to be “a feature, not a bug”, e.g., non-repudiation is likely to be in the user’s interest. The only high threat identified is disclosure of information, **AZ-DI**, focusing on the actual “permissions”, “authorization table”, “who can access what” data generated in the process. Knowledge of the relying parties that can access a user’s Wallet is likely sensitive personal data. The non-compliance threat is close to being a medium threat priority. We decided not to merge it into the disclosure of information threat because of the significant difference in risk, and the fact that the **GS-NC** threat is related.

Table 11: Aggregate results for the *authorize* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	1.57	0.55	Low
Identifiability	1.56	0.55	Low
Non-repudiation	0.75	0.26	Low
Detectability	0.71	0.25	Low
Disclosure of information	2.86	1	High
Unawareness	1.71	0.6	Low
Non-compliance	2	0.71	Low

Table 12 shows the results for the notify action. For the same reason as before, linkability and identifiability are merged into the high priority threat **N-ID**, focusing in the ability of a party (potentially a third party, like Google's GCM) to profile user's usage of the Wallet. Further, the **N-DI** medium threat is on the disclosure of information within the contents of notifications.

Table 12: Aggregate results for the *notify* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2.58	1	High
Identifiability	2.3	0.89	Medium
Non-repudiation	1	0.4	Low
Detectability	0.69	0.27	Low
Disclosure of information	2.1	0.81	Medium
Unawareness	1	0.38	Low
Non-compliance	1.69	0.66	Low

Table 13 shows the results for the search action. These results are the least helpful: each interviewee estimated all threats as relatively risky (the small difference in relative risk is not an artefact of all interviewees estimating low risk across the board). One resulting high threat **S-ID**,

is a merger of the detectability, linkability, and identifiability threats because the main concern of two of the interviewees were profiling of users' behaviour leaking information to different parties (including potentially relying parties). The second high-priority threat is **S-UA**, since depending on how the search mechanism works users are likely unware of how it works which may be used against them. The first medium threat is **S-DI** on disclosure of information, covering the contents of search queries and replies. The second medium threat is **S-NC**, since search queries, their results, and metadata are unlikely covered by privacy policies and the Wallet will inadvertently deal with sensitive medical data (paraphrasing: "a user searches for a healthcare provider specialising on a particular type of treatment"). Finally, we can see that the non-repudiation threat is borderline of being included (0.76). We opt for including it as a conditional medium threat **S-NR**, since if the other related threats to search is not adequately mitigated, then non-repudiation becomes a more serious threat. As a side note from one interviewee, this transparency might actually prevent misuse of the search feature if transparency is made adequately clear to those searching.

Table 13: Aggregate results for the *search* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2	0.86	Medium
Identifiability	1.79	0.77	Medium
Non-repudiation	1.78	0.76	Low
Detectability	2.31	1	High
Disclosure of information	2	0.87	Medium
Unawareness	2.19	0.95	High
Non-compliance	1.84	0.8	Medium

Table 14 shows the results for the external IdP authenticate action. We define two high priority threats, **EA-UA** and **EA-NC**, from the unawareness and non-compliance categories, respectively. They both relate to the complexities of multiple third-parties (the Wallet and external IdP) and the potential of health data being processed/leaked. **EA-LI** is a medium threat that merges linkability and identifiability threats due to the prevalence of IdPs with a business model built on profiling.

Table 14: Aggregate results for the *external IdP authenticate* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	2.55	0.85	Medium
Identifiability	2.36	0.79	Medium
Non-repudiation	1.29	0.43	Low
Detectability	1.41	0.47	Low
Disclosure of information	1.5	0.5	Low
Unawareness	3	1	High
Non-compliance	3	1	High

Table 15 shows the results for the external IdP data import action. The first high priority threat is **EI-DI** on inadvertently disclosing information from the external IdP to relying parties. The second high-priority threat is **EU-UA**, that similar to **EA-UA**, stems from the complex setting. The non-compliance threat **EI-NC**, unlike **EA-NC**, is a medium priority threat. Finally, the linkability and identifiability threats are just below the medium threshold. To err on the side of caution, we include the medium priority threat **EI-LI** with a similar rational as **EA-LI**.

Table 15: Aggregate results for *the external IdP data import* action

Threat	Aggregated relative risk	Relative risk	Priority
Linkability	1.95	0.74	Low
Identifiability	1.95	0.74	Low
Non-repudiation	1.7	0.65	Low
Detectability	1.05	0.4	Low
Disclosure of information	2.63	1	High
Unawareness	2.6	1	High
Non-compliance	2.05	0.78	Medium