Service Provider Accreditation: Enabling and Enforcing Privacy-by-Design in Credential-based Authentication Systems

Stefan More
smore@tugraz.at
Graz University of Technology
and Secure Information Technology Center Austria (A-SIT)
Graz, Austria

Jakob Heher
jakob.heher@iaik.tugraz.at
Graz University of Technology
and Secure Information Technology Center Austria (A-SIT)
Graz, Austria

Edona Fasllija
edona.fasllija@iaik.tugraz.at
Graz University of Technology
and Secure Information Technology Center Austria (A-SIT)
Graz, Austria

Maximilian Mathie
mathie@student.tugraz.at
Graz University of Technology
Graz, Austria

ABSTRACT

In credential-based authentication systems (wallets), users transmit personally identifiable and potentially sensitive data to Service Providers (SPs). Here, users must often trust that they are communicating with a legitimate SP and that the SP has a lawful reason for requesting the information that it does. In the event of data misuse, identifying and holding the SP accountable can be difficult.

In this paper, we first enumerate the privacy requirements of electronic wallet systems. For this, we explore applicable legal frameworks and user expectations. Based on this, we argue that forcing each user to evaluate each SP individually is not a tractable solution. Instead, we outline technical measures in the form of an SP accreditation system. We delegate trust decisions to an authorized Accreditation Body (AB), which equips each SP with a machine-readable set of data permissions. These permissions are checked and enforced by the user's wallet software, preventing over-sharing sensitive data. The accreditation body we propose is publicly auditable. By enabling the detection of misconduct, our accreditation system increases user trust and thereby fosters the proliferation of the system.

CCS CONCEPTS

• **Security and privacy** → **Authentication**; *Access control*; **Trust frameworks**;

KEYWORDS

E-ID, eIDAS, Privacy, Trust, Wallet, Self-sovereign Identity

ACM Reference Format:
Stefan More, Jakob Heher, Edona Fasllija, and Maximilian Mathie. 2024. Service Provider Accreditation: Enabling and Enforcing Privacy-by-Design in Credential-based Authentication Systems. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024), July 30–August 02, 2024, Vienna, Austria.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3664476.3669934

ARES 2024, July 30–August 02, 2024, Vienna, Austria

ACM ISBN 979-8-4007-1718-5/24/07
https://doi.org/10.1145/3664476.3669934

1 INTRODUCTION

In credential-based authentication systems, the exchange of personal data between users and Service Providers (SPs) is the basis for the SP's authentication decision [41]. However, this practice raises significant concerns regarding privacy, data security, and user trust. Users often find themselves in a position where they must entrust sensitive information to SPs [6, 39, 54]. Doing so, they solely rely on the assumption of the SP's legitimacy and the necessity of the data requested. The potential misuse of this data poses a considerable challenge, as identifying and holding SPs accountable for such breaches can be complex [62]. Mutual authentication is a response to these challenges [61]. By implementing mutual authentication protocols, users and SPs can verify each other's identities before engaging in data exchange. These identities allow them to establish a foundation of trust.

Need for Accreditation. However, more than mutual authentication is needed to address the broader data privacy issues. To establish trust in the SP's identity, it is essential to go beyond authentication and include the accreditation of SPs into the system [39, 61]. Accreditation involves the validation of SPs by a trusted party, the Accreditation Body (AB). This process ensures the SPs' legitimacy and adherence to specific standards and regulations (e.g., GDPR) [24, 25].

Enforceable Accreditation Constraints. After an SP is checked and accredited, it can request user data. However, users are often unable to judge whether the list of requested data is justified or lawful [18, 34]. For example, an SP could ask for more data than required to provide the requested service. To remove the burden of judging data requests from the user, we propose enforceable accreditation constraints. These constraints are attached to an SP's accreditation by the AB. Since the constraints are automatically verifiable, they can prevent users from sharing more data than needed. Additionally,

since the user's wallet enforces the constraints, the mechanism disincentivizes over-asking SPs.

Auditable Accreditations. The central AB is responsible for both checking SPs and issuing accreditation constraints. This power requires a lot of trust in the body. To improve user trust in the AB, we propose auditable accreditations. Our concept builds on the publication of accreditation records Further, it ensures that they are verifiable when showing a credential. By enabling the detectability of accreditation misconduct, we establish a reactive security model as an additional privacy layer.

Motivation & Examples. This paper's motivations stem from both users' privacy demands and legal considerations. Regulations such as the General Data Protection Regulation (DRPR) limit the collecting and processing of personal data to explicit and legitimate purposes [24]. Thus, there is a pressing need to develop robust mechanisms that uphold these principles.

To demonstrate the significance of this paper, let us consider a few examples. Imagine a scenario where a user accesses an online healthcare portal (the SP). In this context, the SP requires access to specific health-related attributes to fulfill its service. However, without proper accreditation and access constraints, there is a risk of unauthorized access to sensitive medical information. This risk can potentially violate the user's privacy rights. Similarly, in financial services, users often provide personal and financial data to SPs for transactions and account management. Without effective accreditation measures and access constraints, SPs may overreach in their data requests. This over-sharing could expose users to financial risks and privacy breaches. These examples underscore the critical importance of analyzing the accreditation of SPs and implementing effective access constraints.

Methodology and Contributions: We first collect the *privacy requirements* of credential-based authentication systems. For this, we perform an analysis of legal requirements and user expectations. We focus on data protection and electronic identity law for the legal analysis. Specifically, we explore the European GDPR [24] and the eIDAS regulation in its latest version from 2024 [25]. We filter both regulations for aspects concerning the transmission of user data to an SP. To extend the legal requirements and identify broader user expectations, we explore expectations on analog identity systems and transform them into the digital world. This empirical analysis was done by analyzing daily-life situations involving identity data (cf. Section 5.3).

We then systemize these requirements and present the *concept of an accreditation system* tailored to the identified privacy requirements.

In the course of this discussion, we present several privacy safeguards to fulfill the requirements and increase user privacy. Specifically, we propose automated constraints to prevent SPs from over-asking. Further, we introduce the concept of auditable accreditation registries to assist users with the assessment of accreditation and the stated purpose of processing.

Paper Outline: The rest of this paper is structured as follows: We start by introducing relevant concepts (Section 1.1) and discussing related work (Section 1.2). In Section 2, we then collect and systemize privacy requirements, which we use in Section 3 to show a compliant SP accreditation system. In Section 4, we extend this system with auditable accreditation registries. Finally, we discuss limitations and operational aspects of accreditation systems in Section 5, and propose future work in Section 5.3.

1.1 Background

In this section we introduce the entities and concepts relevant for the discussion of this paper [10, 12, 15]. The relationship between those entities is visualized in Figure 1.

1.1.1 *Architecture and Entities.*

- **Issuer CA/Trust Root**: The Issuer Certificate Authority (CA), or Trust Root, authorizes Issuers to issue certain Credentials.
- **Issuer/IDP**: Issuers, or Identity Providers (IdPs), issue Credentials to Users. These Credentials contain users' identity attributes, such as their name, email address, or other identity information. Each Credential is signed by the respective Issuer.
- **User/Holder**: Users, or Holders, are the recipients of Credentials. When they wish to authenticate to SPs, they do so by presenting credentials containing the requested attribute(s).
- **Wallet**: A User uses software, called a Wallet, that serves as a secure repository for storing and managing Credentials. One common form of Wallet would be an app on the user's phone, relying on the phone's secure hardware. The Wallet is responsible for technically facilitating the showing of Credentials to SPs on the User's behalf.
- **SP/RP/Verifier**: Service Providers (SPs), or Relying Parties (RPs), operate services and request authentication from Users. This process is technically facilitated by specialized software called a Verifier. Verifier software interacts with a User's Wallet and verifies the authenticity and validity of the Credentials presented by the User.
- **Accreditation Body/Registrar**: The Accreditation Body (AB) is an authoritative entity responsible for accrediting SPs based on predefined standards and regulations. The AB is trusted by the user to verify SPs' compliance to privacy an cyber security regulations. It issues accreditations to SPs, validating their legitimacy and adherence to established protocols. Additionally, it assesses the data processing purpose provided by the SP and issues constraints as part of the accreditation. The AB is in turn commonly accredited by a legal authority that enables some liability as part of a governance framework, e.g., by the government itself.
- **Credential**: A credential is a digital representation of a User's identity attributes, typically issued by an Issuer. It contains information such as the user's name, email address, role, affiliations, or other relevant details. Credentials are encoded in machine-readable form and serve as proof of identity when presented to SPs for authentication purposes. Cryptographic signatures are used to both ensure Credentials' authenticity (signature by the Issuer), and to link the Credential to a specific User (public key of the User in the Credential).

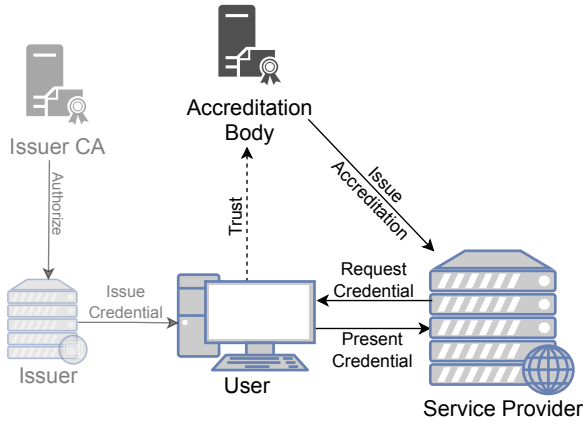

Figure 1: Architecture overview of a credential-based authentication system with SP accreditation

1.1.2 Pseudonyms. By allowing users to interact with SPs without revealing their true identities or a global persistent identifier, pseudonym schemes serve as a privacy-enhancing mechanism. [13, 40, 48]. One practical approach to pseudonymity are Pairwise Pseudonymous Identifiers.[1] Here, each pairwise relationship between a user and a SP is assigned at a unique identifier. This means that a user has a new pseudonym for each SP they interact with, and these pseudonyms are *unlinkable* between SPs. Other schemes may allow the user to have multiple independent pseudonyms towards the same SP, or allow them to prove correlation between pseudonyms at distinct SPs. Some pseudonym schemes (e.g., ABC4Trust Pseudonyms [51]) also allow for pseudonymous and unlinkable usage of a service, but enable a trusted entity (Inspector) to later unmask the true identity of a pseudonym.

1.1.3 Selective Disclosure. Selective disclosure (SD) is the ability to partially reveal individual credentials to SP during authentication processes [5, 8, 29]. For example, given a credential containing a user's date of birth, city of birth, and parents' names, it would be possible to only reveal the date of birth while keeping the remaining information hidden. This can be achieved using privacy-enhancing technologies such as attribute-based credentials or SD-JWTs.[2] Thus, only the necessary information is shared, minimizing the risk of privacy breaches and preventing unauthorized access to sensitive data. This fosters trust and confidence in the authentication system.

1.1.4 Zero-knowledge Proofs. Zero-knowledge proofs (ZKPs) represent a cryptographic technique utilized to demonstrate the validity of a statement without revealing any sensitive information [7, 27]. In the context of authentication systems, ZKPs enable users to authenticate themselves to SPs without disclosing their underlying credentials or personal data [44, 46]. By leveraging zero-knowledge protocols, users can prove possession of credentials meeting certain requirements (called ZK Predicates) without actually revealing the attributes themselves. This enhances user control over their data while ensuring the integrity and security of the authentication process. Incorporating ZKPs into authentication systems reinforces privacy protection measures and strengthens user trust in the system's security. In the context of this paper, zero-knowledge proofs and protocols can be used to take selective disclosure a step further and authorize a SP to only request predicates on user attributes. For example, while a user is in possession of a passport-credential containing the full date of birth, the SP can be authorized to only ask for the user's age, calculated from the date of birth. In that example, neither the date of birth itself nor any other attributes are revealed to the SP, but age verification is still possible.

1.1.5 Certificate Transparency. Transparency logs (or auditable registries) aim to provide visibility of authorities' actions [59]. These systems adopt a reactive security model and primarily focus on detecting misconduct by the authority rather than preventing it. The underlying assumption is that prominent, visible authorities are careful and hesitant to execute attacks that can be traced.

Certificate Transparency (CT) applies this concept to CAs for the Web Public Key Infrastructure [36–38]. It mandates that all valid certificates are recorded in a publicly accessible log, allowing legitimate domain owners to detect any wrongly issued certificates for their domains. The CT log is an append-only and tamper-evident list of all issued certificates that is based on a Merkle Tree structure. The structure allows for cryptographically efficient inclusion and consistency proofs, i.e., it is efficient to prove that a certificate is in the log, and that the current version of the log is an append-only successor of a previous version. In CT, individual users are responsible for monitoring their own entries. For example, domain owners can detect if a certificate was wrongly issued for their domain. Furthermore, these systems assume that global auditors act to prevent servers from carrying out split-view attacks; in other words, they ensure that all users see the same version of the log when accessing the server. This is typically achieved via a gossip protocols.

To ensure CT, web browsers check if the certificate presented by a web server is recorded in a log. This is commonly done by including a Signed Certificate Timestamp (SCT) issued by the log into the certificate or TLS handshake [38]. By including the SCT directly in the request, there is no need for a direct communication between the user and the infrastructure, thus mitigation observability [44, 45, 47].

1.2 Related Work

The *OpenID for Verifiable Presentations* (OID4VP) specification draft introduces a *Verifier Attestation JWT* token [58, Section 10]. This token is used by the user's wallet to authenticate a SP. While the trust framework of the token is out of the specification's scope, the token could serve as basis for accrediting SPs. This is similar to the accreditation certificate discussed in Section 3.

The recent European eIDAS 2 regulation updates the original eIDAS regulation from 2016 [25]. The revised regulation introduces the legal basis for establishing a European Digital Identity framework. More specifically, the provision of interoperable *European Digital Identity Wallets* by European Union member states. In addition to the existing privacy rules of the GDPR [24], the new eIDAS

[1] https://openid.net/specs/openid-connect-core-1_0.html#PairwiseAlg
[2] https://datatracker.ietf.org/doc/draft-ietf-oauth-selective-disclosure-jwt

regulation introduces further safeguards for wallet users. We analyze the regulation's impact in Section 2.

The European Commission is publishing a technical *Architecture and Reference Framework (ARF)* [15] for the EUDIW, in an attempt to avoid fragmentation and diverging standards [25, Recital 70]. The ARF discusses users' trust in the SP [15, Section 6.3.2.1] and introduces a certificate to accredit SPs (Relying Party Instance certificate) [15, Section 7.5]. The ARF also covers the control of presented data by means of "[a] user approval mechanism [...], technical measures in the Wallet Instance, and legal and organizational measures on Member State or EU level" [15, Section 7.6]. In Section 3, we will argue that "user approval" cannot be the sole barrier of defense against excessive data queries. On this point, the ARF agrees with our position that a manual assessment places an undue burden on the user [15, Section 7.6.1]: "For example, if a Relying Party presents the release of all requested attributes as a precondition for the use case that is going on, a User may not have the resources to determine whether this is indeed the case, or to judge the impact of sharing attributes that the Relying Party does not actually need." Thus, automated enforcement is required; but the ARF does not currently describe any means for it. Conceptualizing such a system is a core contribution of the present work.

While the eIDAS 2 regulation and the ARF only lay out the framework for European Identities, the actual creation and design of the wallet software will be left up to each member state. For example, Germany's wallet plans are described in the *Architecture Proposal for the German eIDAS Implementation* working document [20]. The document covers privacy requirements and explicitly mentions data minimization and the prevention of overidentification. To achieve this requirement, the document proposes a *Enforced Disclosure Limitation*, but in its current version does not implement it. This requirement is similar to the attribute-constraint mechanisms proposed in Section 3.

The Dutch EU wallet prototype involves a *DeviceRequest* containing the identity of the SP [60, Disclosure flow]. That information covers the accreditation of SPs, but does not involve over-sharing-prevention measures apart from manual user consent.

The Italian EU wallet implementation profile describes the accreditation of SPs by a Trust Anchor or its intermediates [21]. The SP accreditation builds on OpenID Federation (OIDF) Trust Marks (signed JWTs) [28]. Over-sharing prevention constraints (cf. Section 3.4) are supported using OIDF's metadata policies [28, Sec. 6.1].

The European Blockchain Services Infrastructure (EBSI) provides a Trusted Issuer Registry (TIR) [30] to allow Verifiers (SPs) to check the identity and the accreditations of Trusted Issuers. According to their Issuer Trust Model, Trusted Accreditation Organizations (TAOs) authorize Trused Issuers to issue particular Verifiable Credentials (VCs). However, EBSI does not provide a Trusted Verifier Register; verifiers are not onboarded or accredited as issuers are.

The civil society watchdog epicenter.works provides an analysis of privacy safeguards for digital public infrastructure systems, with a focus on recent EU law [23].

Extended Access Control (EAC) is a security control for restricting access to sensitive biometric data in electronic passports [31]. In EAC, certificates (i.e., accreditations) are used to authorize terminal readers to access sensitive data on the passport. These certificates are separately issued by the respective country that also issues the passport. The certificates are then verified by the chip in the passport, and can grant access to different types of data (e.g., just to the fingerprint, or also to the iris scan). This represents a simple version of the attribute constraints described in this paper.

2 REQUIREMENTS

In this section, we explore both legal aspects and user expectations, to derive minimal requirements for any electronic identity system. On the legal side, we focus on the EU's GDPR and eIDAS regulations, and derive legal requirements (LREQs). On the user side, we analyze analog "plastic-card based" authentication scenarios [15, Section 7.7.2]; users' privacy in any credential-based authentication system should be at least as good. From this, we derive user requirements (UREQs).

2.1 Legal Requirements

2.1.1 GDPR compliance. Any SP's authentication system must comply with data protection regulations like the GDPR. GDPR Articles 5 and 6 emphasize the purpose and lawfulness of processing personal data [24]. Any data processing must have a purpose in alignment with GDPR principles [24, Article 5]. Additionally, not every purpose necessitates accessing qualified data, such as a government-issued legal name; accessing non-qualified data, such as a user's self-provided name with no verifying information, is often sufficient. More generally, this underscores the importance of adhering to the principle of "privacy by design", disclosing information on a "need to know" basis, and implementing the principle of least privilege [24, Article 25]. A well-designed wallet system based on these principles can enhance GDPR compliance, mitigate privacy risks, and uphold user rights to data protection.

LREQ1 **Purpose**: Personal data shall be collected for specified, explicit and legitimate purposes [24, Article 5].

LREQ2 **Data Minimization**: Personal data collected shall be limited to what is necessary in relation to the purposes for which they are processed [24, Articles 5 and 25].

2.1.2 eIDAS 2 privacy rules. The European eIDAS 2 regulation amends the original eIDAS regulation from 2016, introducing a framework for European digital identity wallets [25]. The new eIDAS regulation introduces further privacy rules, in addition to the existing privacy rules of the GDPR [23]. Those rules are tailored to the wallet context. For this paper, the most relevant articles are Article 5a – introducing the identity wallet – and Article 5b – covering SPs. Furthermore, Article 5 establishes a right to the use of pseudonyms.

LREQ3 **SP Registration**: A SP shall register in the Member State where it is established. SPs shall identify themselves to the user [25, Article 5b].

LREQ4 **Purpose Registration**: During registration, a SP shall provide indication of the data to be requested from users, and the SP shall not request any other data than indicated [25, Article 5b].

LREQ5 **Purpose Information**: Wallets shall inform the user whether the SP has the permission to access a credential [25, Article 5a].

LREQ6 **Auditability**: The list of registered SPs and their indicated data processsing shall be public in a form suitable for automated processing [25, Article 5b].

LREQ7 **Unlinkability**: The technical framework shall ensure unlinkability [25, Article 5a].

LREQ8 **Selective Disclosure**: The technical framework shall enable selective disclosure of data [25, Article 5a].

LREQ9 **Unobservability**: The technical framework shall not allow Issuers or any other party to track, link or correlate user behavior [25, Article 5a].

LREQ10 **Pseudonyms**: The use of pseudonyms chosen and managed by the user shall not be prohibited [25, Article 5]. Wallets shall enable the user to generate pseudonyms and store them encrypted and locally [25, Article 5a]. SPs shall not refuse the use of pseudonyms, except where the identification of the user is required by law [25, Article 5b].

The GDPR requires that a data subject has given consent to the processing of his or her personal data (for one or more specific purposes) [24, Article 6]. However, this is only one of the six legal bases stated by the GDPR. In contrast, the eIDAS revision highlights the user's full control over their data [25, Article 5a]. Additionally, the eIDAS ARF states that "[a Wallet] SHALL always ask the User to give approval for any attribute released. This goes for any use case. [...] That is, a user SHALL always be able to refuse presenting an attribute that is requested by a [SP], even when knowing that the consequence of that refusal may have negative consequences for the user" [15, Section 7.7.2].

LREQ11 **Control**: Users shall have full control of the use of the wallet and of the data in their wallet [25, Article 5a].

2.1.3 Other regulations. While we focus on the European sphere in this paper, we note that many jurisdictions recognize similar rights. For example, the U.S. Supreme Court recognizes a constitutional right to anonymity derived from the First Amendment's freedom-of-speech protections; and that the "decision in favor of anonymity may be motivated by fear of economic or official retaliation, by concern about social ostracism, or merely by a desire to preserve as much of one's privacy as possible" [16, p. 341].

2.2 User Expectations

Next, we focus on users' expectations from a credential system. The underlying theory is simple: by invoking the term "Wallet", we evoke a particular mental framework from users. We therefore describe the privacy posture of revealing the contents of one's physical wallet, and map it to the digital sphere.

In any credential system, users' privacy should be no worse than this. Such an assurance, if communicated properly, allows users to feel comfortable using the system, leading to increased adoption.

User perception: *When I submit a claim to my insurance company, I need to present bills and medical documents. Only the insurance clerks can see my documents. They are not visible to, e.g., the postal office or bystanders.*

UREQ1 **Confidentiality**: When presenting a credential to some SP, the credential is only accessible by that specific SP. The system must be able to authenticate the SP, verify its identity, and ensure that the data is directly and securely transmitted to that SP.

User perception: *When I want to enter a bar, I need to prove I'm of the legal drinking age. I do this by showing my government-issued ID card to the bouncer. The ID card also has my photograph, name, date of birth, and city of residence; but while the bouncer can see this data, I am not expecting them to memorize the ID cards of any of the hundreds of partygoers passing by each night. If I found out they were doing so, I would be disturbed and concerned.*

This scenario introduces an important distinction that only exists in the analog environment: humans generally have limited and imperfect memory. This informs the expectations we have when presenting information to other humans, and does not translate directly to a digital context. Computers, with their perfect recall, would have no trouble recording all information that we make available to them in passing. Therefore, in the digital context, we cannot rely on them only picking up the necessary data and forgetting the rest, but need to prevent disclosing that superfluous data in the first place.

UREQ2 **Selective Disclosure**: When presenting a credential to a SP, the SP can only see the information relevant to that particular transaction.

User perception: *When I shop at a grocery store chain, I pay in cash. The cashier sees what I bought. But, similarly to the bouncer in our previous scenario, there's no way they'll remember most of the hundreds of customers coming through, or what they bought.* These interactions are *essentially anonymous* [48].

The perception is the case for most daily life interactions in the physical world where no identification is required. However, this experience again does not translate perfectly to the digital sphere; computers have perfect recall, after all. The digital equivalent is, again, to not disclose the identity information instead of relying on human memory's fallibility.

UREQ3 **Anonymity and Pseudonymity**: Using services should not reveal the user's "true" identity, or any long-lived identifier, unless desired and required.
Depending on the context, this can be realized either by disclosing a disposable reusable pseudonym, or no identifier at all.

A similar thing happens when a user visits the same store twice in a week, or when they visit multiple stores. In that case users do not expect that the show remembers their visit, or that different shops exchange data about a user. However, in the digital world, both remembering behavior of customers and sharing between different entities is easily automated [48]. It is thus important that a wallet system offers a functionality to use services without being tracked.

UREQ4 **Unlinkability**: Users expect that they can use a credential or a set of credentials with different SPs without reducing their privacy. Specifically, the involved SPs should not be able to link the interactions with each other, which would allow for user profiling. In the case of anonymous/pseudonymous interactions, users would like to use the same credential twice (i.e., for age verification) in two interactions. The expectation is that the service cannot

link the two interactions with each other, i.e., does not know that the same user interacted with it twice (multi-show unlinkability).

Let us return to the party-going example from earlier.

User perception: *To enter the bar, I show my government-issued ID to the bouncer. There's no reason for me to suspect that the government could tell that I used my ID card to get into this bar, or even that I took it out of my wallet at all.*

However, in poorly designed digital systems it is common that the issuer or infrastructure takes part in the showing of a credential [48]; e.g., for a revocation check [3].

UREQ5 **Unobservability**: Only the user and the SP learn that a showing takes place.
Neither the issuer nor any other authority (e.g., the government) should learn where the user is using their wallet, or even *that* the user is using their wallet.

User perception: *If I carry an ID card in my wallet, nobody can see it unless I take it out, or either the card or the wallet is taken from me.*

The GDPR states consent as only one of the six legal bases for the processing of personal data [24, Article 6]. However, when it comes to identity documents, the offline-world experience relies on the user's cooperation when accessing an ID document (with the exception of physical force, see Section 5.1) [15, Section 7.7.2]. In the same way, wallet software must rely on the user's approval, and must never act on its own.

UREQ6 **Consent**: No credentials should be disclosed or shown unless a user explicitly operates the wallet software for that purpose.

3 SP ACCREDITATION AND CONSTRAINTS

While user consent is a necessary requirement for the disclosure of data [24, 53], we argue that it cannot be the *sole* requirement. Intuitively, our reasoning is simple: users will typically encounter credential requests in contexts in which they wish to complete some other task. Their goal in the ongoing context is to achieve that other task, taking whatever steps are necessary.

Let's say I want to order pizza for a group of five. We've identified a suitable restaurant, and I've passed my phone around, having everyone choose from the menu. I now go to place the order, and we're all ready to have some great pizza. But wait – it says it needs to verify my identity to prevent hoax orders. I get a notification from my wallet: the pizza place wants to query my complete passport information. At this point, I can either grumble and accept, or I'll need to explain to the group why we need to re-do the entire order somewhere else. It should be intuitively obvious that only the most privacy-minded among us will choose the latter option.

This intuition is also backed up by real-world research into comparable privacy consent prompts [34, 35]. In a study by Laine, 25% of respondents state that they accept GDPR cookie banners automatically, without making a conscious decision [35]. Some also expressed frustration at "distracting" or "annoying" prompts. Simultaneously, among the respondents that accepted these cookie prompts, 80% state that they would have chosen to reject cookies if they could do so through a global browser setting. Even though the study setup uses a banner that places equal weight on "accept" and "reject" options, 50% of participants both accepted cookies on the study site and expressed a general desire to reject cookies globally. Laine concludes that these users were conditioned to accept all such prompts, since this is the only option guaranteed to let them view the website's contents [35].

In light of these arguments, it seems apparent to us that while user consent is required, it should not be *sufficient*. Instead, some other entity needs to assess the basis for SPs' data queries without being rushed, pressured, or otherwise influenced towards allowing the transfer. We will refer to this entity as an *Accreditation Body (AB)*. The AB checks the SP and its data processing, and issues a accreditation to the SP. This accreditation then allows the SP to query data from users' wallets. The rest of this chapter will be dedicated to conceptualizing an accreditation system for SPs.

3.1 Accreditation Process

Before a SP can request credential showings from users, they must first undergo an accreditation process at an AB. This involves, at minimum, them stating what information they're looking to request from users, and what basis they need this data on. The AB can then evaluate whether the stated justifications are sufficient to obtain the data in question.

Depending on the nature of the data, additional organizational scrutiny of the SP's processes and systems may also be conducted. For example, if a SP seeks to operate on users' health information, an independent evaluation of its information security posture may be required.

Once the process is complete to the AB's satisfaction, it issues an *accreditation certificate* to the SP in question. This certificate contains both the SP's identity, as well as the maximum extent of data that the SP may request [18]. The latter information is encoded in machine-readable form and called *accreditation constraints*.

In our example from earlier, a pizza store would not be accredited to request my full passport contents, as there is insufficient justification to access this data. Instead, after a balancing test by the AB, they may only be allowed to verify my country of residence. Since this constraint is recorded in the store's accreditation certificate, them asking customers for passport contents would be an exercise in futility. Therefore, they will be forced not to overreach in their requests, and I will never have to weigh my desire for privacy against my desire for pizza.

3.2 Showing Process

With the SP having obtained an accreditation certificate, they can now query users for credentials. We first provide a high-level overview of the process steps before discussing the details in the next section.

(1) Service Access. The user wants to access some service or resource. Since the service is protected by some access control system, the SP asks the user to present certain credentials so that it can authenticate the user and grant (or deny) access.

(2) Presentation Request. The user receives a request from the SP to present credentials for authentication. In this request, the SP tells the user what credentials (or attributes) it requires. For example, if the SP needs to perform age verification, the SP asks the user

for a government-issued credential containing the user's date of birth. The SP also provides its SP accreditation certificate, which authorizes them to ask for the information in question.

(3) Authenticate SP. The user's wallet verifies the identity of the SP against the accreditation certificates, checking that the provided SP accreditation represents the entity that sent the presentation request. The wallet also validates the SP's certificate for trustworthiness; for example by verifying if the accreditation certificate was issued by a trusted AB.

(4) Check Constraints. The wallet evaluates the access constraints imposed on the SP by the AB in the accreditation certificate. This ensures that the requested attributes align with regulations and that the SP has a legally justifiable purpose to process these data. For example, if the SP asks the user for a date of birth, the wallet checks if this specific SP is authorized to ask for that date.

(5) User Consent. The user is prompted to consent to the sharing of their credentials and attributes based on the SP's request. This ensures that the requested attributes align with the user's expectations and privacy preferences. Note that this step is only reached if the request aligns with the SP's accreditation.

(6) Build Presentation. Upon receiving user consent, a presentation containing the requested credentials and attributes is constructed by the wallet. Only the requested attributes are included in the presentation, i.e., the user's date of birth.

(7) Sign and Send Presentation. The wallet digitally signs the presentation and sends it securely to the SP. The wallet must ensure that only the authenticated SP can receive the data.

3.3 Additional Steps & Pitfalls

Now that we have established the general process, we highlight some additional considerations.

On the observability of revocation checks. When issuing certificates, the question of revocation is always relevant. If naively implemented, revocation checks for accreditation certificates involve the certificate's issuer, i.e., the AB. Doing so violates the *Unobservability* goal, allowing the AB to observe which SPs a user interacts with [55].

This is a well-known issue, and thus also has well-known solutions. For instance, in the Web PKI, bundling a certificate's status to the authentication in *OCSP stapling* using widespread [45, 47].

This technique, or equivalent, needs to be employed to ensure that accreditation certificates do not inadvertently compromise users' privacy expectations.

On accountability for service providers. In any system, actors may at times misbehave. In such a case, it is desirable to allow the misbehavior to be reliable documented. This allows organization or legal penalties to be imposed on the bad actor.

In the context of SPs, we are concerned with SPs attempting to over-ask additional constraints not strictly enforced by the wallet, such as making supposedly-optional attributes mandatory. Thus, it would be desirable to ensure *non-repudiation* of the SP's presentation request. Depending on the authentication method used, this may be easy or hard; for instance, a TLS connection handshake can be non-repudiable when using the right cipher suite, but the actual data sent over the TLS connection can be forged by either side of the connection [11]. A non-repudiable wallet interaction can be achieved by signing the SP's presentation request directly with their key.

Alternative accreditation model. Accreditation registries are an alternative to credential-based accreditation. In this more centralized approach, the result of the accreditation process is not a certificate, but an entry into some public registry, e.g., list or database. While this entry is linked to the SP by means of a (self-signed) certificate, the accreditation information itself is directly stored in the registry. The advantage of this approach is that competent entities can monitor and audit the list of accredited SPs, and that the information can more easily be updated. When naively implemented, a disadvantage of this approach is that the registry potentially learns when a wallet interacts with a specific SP, resulting in observability. This could be mitigated by continuously downloading a snapshot of the registry or other more privacy-preserving architectures (see Section 4).

3.4 Types of Accreditation Constraints

The constraints attached to the accreditation can take different forms, constraining different levels of information:

- **Credentials:** Accreditation certificates contain a list of all the types of credentials the SP can access. Credential types can be for example identified using Uniform Resource Names (URNs), Digital Object Identifiers (DOIs), and Credential Schemas [32, 42, 56]. These constraints could also encompass a group of credentials, e.g., granting access to different health-data credentials.
- **Attributes:** To enable more flexible control, accreditation certificates can also contain a list of attributes as constraints. In the same way than for credentials, attributes can be identified using various schemes, e.g., URNs or other identification schemas.
- **Predicates:** To support privacy-preserving and -enhancing technologies like attribute-based credentials and zero-knowledge proofs (cf. Section 1.1), accreditation certificates can restrict the access to specific predicates on that data. For example, if the only legitimate purpose stated by a SP is an age-check of its users, the AB could accredit this SP only to `age-check` or `date-difference` predicates [46]. By doing so, the SP has no access to other data, not even the user's date of birth.
- **Pseudonyms:** Another type of accreditation constraint is concerned with the handling of the user's identifiable information, i.e., identity identifier or pseudonyms (cf. Section 1.1). For example, a social network might be accredited to use the wallet system merely to (re-)authenticate users (e.g., as a more convenient or secure replacement for a password manager or second factor system). In that case, the social network acting as SP has no business in learning any other information about the user. While the access to the user's credentials and attributes is easily restricted using the constraint types discussed above, the SP still has access to the user's identifier. Thus, the SP can link multiple visits by the same user and apply other profiling techniques. To support

users in their choice of identities, ABs can restrict SPs to only access user's pseudonyms, and to prevent that only a single pseudonym is possible for each SP. This is the wallet counterpart to a user that creates multiple accounts with an SP, without creating any link between those accounts. For example, OIDC's pairwise pseudonymous identifiers prevent system-wide profiling by colluding SPs, but do not stop a SP from linking and profiling multiple visits by the user. In contrast, pseudonyms that are freely generated by the users (somehow linked to their wallet identity or not) preserve the user's privacy and freedom of choice. While in general it's up to the users how many pseudonyms they use for each service, there might be a limit to sybil restrictions (see discussion in Section 5).

4 AUDITABLE ACCREDITATION REGISTRY

ABs check the compliance of SPs with data protection and cyber security regulations. Further, they accredit SPs and use accreditation constraints to control SPs' access to user data. It is thus of paramount importance that ABs comply with all regulations and act in the users' interest. The AB's behavior is the basis for the users' trust in the identity system and is thus crucial for a system's adoption.

To increase the users' trust in the system even further, we propose that ABs publish their accreditation decisions in *auditable accreditation registries*. This allows third parties to audit the AB's decisions and check whether the issued accreditation constraints match the SPs' stated purpose for lawful processing. Doing so allows the detection of misconduct and serves as an additional incentive for ABs to comply with regulations. For example, competent entities like data protection authorities and NGOs could be enabled to monitor the AB system. By doing so, the burden of verifying whether an AB issued an accreditation that is too permissive is transferred from the user to these competent entities.

Auditable Accreditation: Combining Certificate and Registry. On a technical level, auditable accreditation registries combine both Accreditation Certificates and Accreditation Registries. In doing so, the SP is authenticated using its accreditation certificate. That certificate also contains the SP's accreditation constraints. The trust in this certificate is then established by querying the registry. This combination enables both privacy during the authentication process, and auditability of the system.

Auditability: Link between Certificate and Registry. To ensure auditability, it is important that the accreditation registry records contain the same constraints as the accreditation certificate. To enable a wallet to check for this match, we propose the use of transparency logs. One suitable mechanism is the Web PKI's Certificate Transparency (see Section 1.1.5). Since accreditation certificates are encoded as X.509 certificates, they can be combined with CT's signed SCT proof of log inclusion. This mechanism provides the basis for the verification of both the accreditation certificate and its auditability. To ensure the unobservability of the showing, this verification of auditability can take place offline.

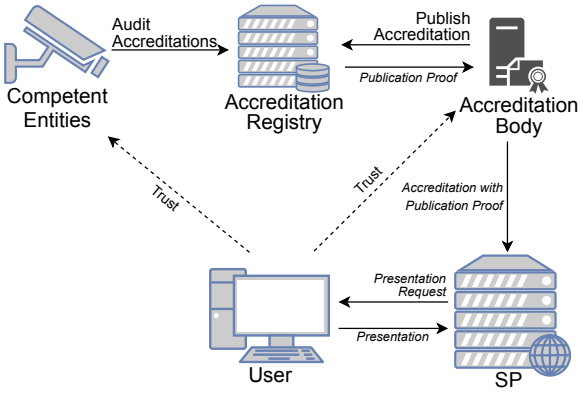

Figure 2: Accreditation and Showing process

We now outline the extension of the accreditation and showing process with auditable registries. We visualize the overall dataflow in Figure 2.

Extended Accreditation Process. During the accreditation process, the AB adds the accreditation to the public registry. This results in a proof of publication of the accreditation. The proof cryptographically links the accreditation certificate and constraints to the registry record, e.g., in the form of a SCT. The AB then adds the proof to the accreditation certificate. After the SP obtained the auditable accreditation certificate, they can now query users for credentials.

Extended Showing Process. From the SP's side, the showing process takes place unaltered. The user's wallet receives the accreditation certificate alongside the SP's presentation request. During the authentication of the SP (Step 3 in Section 3.2), the wallet checks the SCT inclusion proof.

This guarantees that the provided certificate and stated constraints are auditable, and that a misconduct of the AB would be observable. The wallet then continues evaluating the access constraints imposed on the SP by the AB in the accreditation certificate.

5 DISCUSSION

We now state the limitations of accreditation systems, discuss the constraining of pseudonym generation, and propose further research directions.

5.1 Limitations

While accreditations play a crucial role in ensuring the trustworthiness of SPs, it is essential to recognize their limitations.

5.1.1 Trust in AB. Accreditations, while valuable, cannot resolve all issues, particularly those rooted in legal or governance frameworks.

For instance, certain jurisdictions may mandate the disclosure of real names or identifiers. If this is enforced by law, a government-controlled system would not prevent SPs from executing this regulation. Specifically, accreditation bodies would then accredit all SPs to retrieve names or other identifiable data. However, users can still

refuse to share credentials – with the possible effect of no access to services [15, Section 7.7.2].

Additionally, accreditations are susceptible to exploitation. For example, accreditation bodies may covertly issue accreditations with overly permissive constraints. Consequently, trust in the AB is paramount, necessitating transparency measures such as public logs of accreditations and constraints (cf. Section 4). These logs enable competent entities to monitor accreditation activities and ensure compliance with established standards. While this cannot prevent intrusive accreditation constraints, it at least allows for their detection by the public, potentially holding the authority accountable.

5.1.2 Availability and Voluntariness. Availability is a core security goal of every system. This is especially the case for critical systems like electronic identity frameworks. It is thus crucial that wallet systems are designed in a way that ensures the highest availability, e.g., by introducing redundancies and limiting the use of centralized systems in everyday interactions.

In the same way, privacy is all about the user's control of how their personal information is collected, processed, and shared [50]. A consequence of that is that it remains a free choice to use an electronic identity system. However, it is impossible to prevent SPs from mandating wallet systems on a technical level. It is thus important that legal frameworks ensure useful and practical alternatives. For example, the updated eIDAS regulation mandates that "the use of European Digital Identity Wallets shall be voluntary. [...] It shall remain possible to access public and private services by other existing identification and authentication means" [25, Article 5a (15)].

5.1.3 Purpose Enforcement. Another limitation lies in the nature of digital data. Once data is *lawfully* shared, enforcing usage and storage limitations becomes challenging: The user has no control about what the SP *really* does with the data, and with whom it shares the data. While liability mechanisms exist, such as presentation/access logs stored by the user, they do not inherently prevent misuse or unauthorized access. However, they serve as evidence of data sharing and can be instrumental in holding parties accountable for breaches.

5.1.4 Physical Threats. Additionally, accreditation primarily covers the showing process and does not extend to protecting against unauthorized access to the wallet itself. For instance, in scenarios where a border authority confiscates a user's phone, accreditation mechanisms may not offer direct protection against data exposure.[3]

5.2 Constraining the Number of Pseudonyms

Often, SPs want to restrict the creation of fresh user accounts; this may be done to combat spam, to ensure the authenticity of reviews or comments, or the integrity of other reputation systems. Fraudulently presenting multiple identities is commonly termed a *Sybil attack*, and systems preventing this are referred to as *Sybil-resistant* [22].

[3]While mechanisms like (verifier) app attestation could make direct access to the data harder, it cannot fully prevent unauthorized access by entities with physical access to the user's device.

Many real-world SPs try to obtain Sybil resistance. Common means include phone number verification, scans of identity documents, or in-person account registration. These interactions are, by their very nature, highly invasive to a user's privacy.

Government issuers are uniquely positioned to enable Sybil resistance; however, it is unclear to what extent we should want them to.

On the one hand, its implications and nuances are a complex subject: the ability to present multiple independent pseudonyms is undoubtedly a core tool in the privacy toolbox of today's Internet citizen [43, 49]. A single stable identifier – even if scoped – would undoubtedly be a significant windfall for the many actors that make up today's panoptical Internet. As privacy advocates, it seems self-evident that we should not want our governments to give it to them.

On the other hand, current Sybil resistance efforts are also deeply invasive to users' privacy. By offering a more privacy-friendly pseudonymized method that still provides *some* Sybil resistance, such privacy invasions could be discouraged – or even regulated.

Furthermore, SPs are concerned with individuals presenting hundreds, or even thousands, of independent, disposable, identities; meanwhile, a single individual requiring hundreds or thousands of legitimate online personas seems far-fetched to us.

Might it be possible to find some number n – perhaps for a given SP, recorded in its accreditation – and to allow each user to obtain up to n independent, stable, unlinkable pseudonyms for this SP, but no more? This still places *some* limit on users' privacy; but is there a reasonable trade-off to be found? We posit this as a starting point for further discussion.

5.3 Future Work

In this section we propose future endeavours and next steps for SP accreditation and constraints

Explore users' privacy expectations: In addition to regulatory compliance, user concerns and expectations are a cornerstone for a system's privacy requirements [26]. This paper builds on a empirical analysis of user expectations, but a more systematic view is needed to fine tune design and implementation. While research into users' expectations in a system's privacy exists [19, 50, 52, 57], we propose a more focused evaluation and user-study in the context of digital identities, specifically wallet-bases systems.

Extensive survey of existing identity systems: In this paper we focus on electronic identity systems in the scope of the EU's eIDAS wallet regulation [15, 25] as well as related Self-sovereign Identity (SSI) systems [1, 2]. To broaden the understanding of the issue of overidentification and potential mitigation strategies, we propose a survey of existing identity systems worldwide. One interesting aspect is how the systems differ depending on their instantiation context, e.g., political system of the respective country. We consider it especially worthwhile to explore the level of resilience of privacy safeguards and potential extended safeguards, e.g., by combining technical and legal measures.

Instantiation and implementation: This paper does not provide a concrete instantiation of an accreditation system. Instead, we identify privacy requirements and propose conceptual building blocks to comply to them. To better understand the implications of

our proposal on specific identity systems, we propose a more concrete instantiation, e.g., in the context of the EU's eIDAS regulation and its ARF [15, 25]. For the encoding of the SP's attribute requests and their purpose, we propose to look into privacy negotiation literature [9, 14, 33]. For example, a promising body of literature are the successors of the Platform for Privacy Preferences Project (P3P) standard [4, 17, 18].

5.4 Conclusions

Accreditations and accreditation constraints represent foundational pillars in the architecture of credential-based authentication systems. Those techniques play an important role in establishing trust, ensuring accountability, and safeguarding user privacy. By subjecting Service Providers (SPs) to accreditation processes, users are enabled to assess the SPs' legitimacy and trustworthiness. Moreover, accreditation constraints add a layer of protection, preventing overasking and unnecessary data disclosure while empowering users with greater control over their personal information. The effectiveness of accreditations and constraints depends on Accreditation Bodies' (AB) rigor to check SPs and their processing needs. Striking a balance between facilitating seamless authentication experiences, enabling innovation, and upholding privacy principles remains the goal. Auditable accreditation registries can help to increase trust in electronic identity systems. As authentication technologies continue to evolve, ongoing efforts to refine accreditation processes and integrate constraint mechanisms will be crucial in ensuring the integrity, security, and user-centricity of authentication systems in the digital age.

ACKNOWLEDGMENTS

This work was supported by the European Union's Horizon 2020 research and innovation programme under grant agreement № 101020416 (ERATOSTHENES). We thank our reviewers for their well-considered suggestions.

REFERENCES

[1] Andreas Abraham. 2017. *Whitepaper: Self-Sovereign Identity.* Technical Report. https://technology.a-sit.at/en/whitepaper-self-sovereign-identity/ online, accessed on 27 April 2022.
[2] Andreas Abraham. 2022. *Qualified Self-Sovereign Identity: Addressing the gaps between Self-Sovereign Identity and traditional Identity Systems.* Ph. D. Dissertation. https://doi.org/10.13140/RG.2.2.29266.22728
[3] Andreas Abraham, Stefan More, Christof Rabensteiner, and Felix Hörandner. 2020. Revocable and Offline-Verifiable Self-Sovereign Identities. In *TrustCom*. IEEE, 1020–1027.
[4] Paul Ashley, Satoshi Hada, Günter Karjoth, and Matthias Schunter. 2002. E-P3P privacy policies and privacy authorization. In *WPES*. ACM, 103–109.
[5] David Bauer, Douglas M. Blough, and David Cash. 2008. Minimal information disclosure with efficiently verifiable credentials. In *Digital Identity Management*. ACM, 15–24.
[6] Eleanor Birrell and Fred B. Schneider. 2013. Federated Identity Management Systems: A Privacy-Based Characterization. *IEEE Secur. Priv.* 11, 5 (2013), 36–48.
[7] Manuel Blum, Paul Feldman, and Silvio Micali. 1988. Non-Interactive Zero-Knowledge and Its Applications (Extended Abstract). In *STOC*. ACM, 103–112.
[8] Stefan A. Brands. 2000. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy.* MIT Press, Cambridge, MA, USA.
[9] Jan Camenisch, Sebastian Mödersheim, Gregory Neven, Franz-Stefan Preiss, and Dieter Sommer. 2010. A card requirements language enabling privacy-preserving access control. In *SACMAT*. ACM, 119–128.
[10] Yuan Cao and Lin Yang. 2010. A survey of Identity Management technology. In *2010 IEEE International Conference on Information Theory and Information Security*. 287–293. https://doi.org/10.1109/ICITIS.2010.5689468
[11] Srdjan Capkun, Ercan Ozturk, Gene Tsudik, and Karl Wüst. 2021. ROSEN: RObust and SElective Non-repudiation (for TLS). In *CCSW*. ACM, 97–109.
[12] David W. Chadwick, Michael Kubach, Ioram Schechtman Sette, and Isaac Henderson Johnson Jeyakumar. 2023. Establishing Trust in SSI Verifiers. In *Open Identity Summit (LNI, Vol. P-335)*. Gesellschaft für Informatik e.V.
[13] David Chaum. 1985. Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Commun. ACM* 28, 10 (1985), 1030–1044.
[14] Juri Luca De Coi and Daniel Olmedilla. 2008. A Review of Trust Management, Security and Privacy Policy Languages. In *SECRYPT*. INSTICC Press, 483–490.
[15] European Commission. 2024. *Architecture and Reference Framework (Version 1.3).* https://eu-digital-identity-wallet.github.io/eudi-doc-architecture-and-reference-framework/1.3.0/arf
[16] U.S. Supreme Court. 1995. *McIntyre v. Ohio Elections Comm'n, 514 U.S. 334 (1995).* https://supreme.justia.com/cases/federal/us/514/334/case.pdf
[17] Lorrie Faith Cranor. 2003. P3P: Making Privacy Policies More Useful. *IEEE Secur. Priv.* 1, 6 (2003), 50–55.
[18] Lorrie Faith Cranor. 2019. *Platform for Privacy Preferences (P3P).* Springer Berlin Heidelberg, Berlin, Heidelberg, 1–2. https://doi.org/10.1007/978-3-642-27739-9_759-2
[19] Adéle da Veiga. 2022. A study on information privacy concerns and expectations of demographic groups in South Africa. *Comput. Law Secur. Rev.* 47 (2022), 105769.
[20] Bundesministerium des Innern und für Heimat. 2024. *Germany: eIDAS 2.0 Architecture Concept (Version 2).* https://gitlab.opencode.de/bmi/eudi-wallet/eidas-2.0-architekturkonzept
[21] AGID + Team Digitale. 2024. *Italian EUDI Wallet implementation profile (Version 0.6.1).* https://italia.github.io/eudi-wallet-it-docs
[22] John R. Douceur. 2002. The Sybil Attack. In *IPTPS (LNCS, Vol. 2429)*. Springer, 251–260.
[23] Epicenter.works. 2024. *Analysis of Privacy-by-Design EU Legislation on Digital Public Infrastructures.* https://epicenter.works/en/content/analysis-of-privacy-by-design-eu-legislation-on-digital-public-infrastructures
[24] European Parliament and Council of the European Union. 2016. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation, GDPR).* https://data.europa.eu/eli/reg/2016/679/oj
[25] European Parliament and Council of the European Union. 2024. *Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework (eIDAS 2).* https://eur-lex.europa.eu/eli/reg/2024/1183/oj
[26] Mohamad Gharib, Paolo Giorgini, and John Mylopoulos. 2020. An Ontology for Privacy Requirements via a Systematic Literature Review. *J. Data Semant.* 9, 4 (2020), 123–149.
[27] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1985. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *STOC*. ACM, 291–304.
[28] R. Hedberg, M.B. Jones, et al. 2024. OpenID Federation 1.0 - draft 36. tps://openid.net/specs/openid-federation-1_0.html
[29] Jason E. Holt and Kent E. Seamons. 2002. Selective disclosure credential sets. *IACR Cryptol. ePrint Arch.* (2002), 151.
[30] European Blockchain Services Infrastructure. 2024. *Issuer Trust Model.* https://hub.ebsi.eu/vc-framework/trust-model/issuer-trust-model-v3
[31] International Civil Aviation Organization. 2021. *Doc 9303: Machine Readable Travel Documents.* https://www.icao.int/publications/pages/publication.aspx?docnum=9303
[32] ISO 26324:2022 2022. *Information and documentation – Digital object identifier system.* Standard. International Organization for Standardization, Geneva, CH.
[33] Saffija Kasem-Madani and Michael Meier. 2015. Security and Privacy Policy Languages: A Survey, Categorization and Gap Identification. *CoRR* abs/1512.00201 (2015).
[34] Lin Kyi, Abraham Mhaidli, Cristiana Teixeira Santos, Franziska Roesner, and Asia J. Biega. 2024. "It doesn't tell me anything about how my data is used": User Perceptions of Data Collection Purposes. In *CHI*. ACM, 984:1–984:12.
[35] Jere Laine. 2021. *There is no decision: design of cookie consent banner and its effect on user consent.* Master's thesis. Tampere University. https://trepo.tuni.fi/handle/10024/135598
[36] Ben Laurie. 2014. Certificate transparency. *Commun. ACM* 57, 10 (2014), 40–46.
[37] Ben Laurie, Adam Langley, and Emilia Käsper. 2013. Certificate Transparency. *RFC* 6962 (2013), 1–27.
[38] Ben Laurie, Eran Messeri, and Rob Stradling. 2021. Certificate Transparency Version 2.0. *RFC* 9162 (2021), 1–53.
[39] Hosub Lee and Alfred Kobsa. 2017. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *PerCom*. IEEE, 276–285.
[40] Anna Lysyanskaya, Ronald L. Rivest, Amit Sahai, and Stefan Wolf. 1999. Pseudonym Systems. In *SAC (LNCS, Vol. 1758)*. Springer, 184–199.
[41] Tewfiq El Maliki and Jean-Marc Seigneur. 2007. A Survey of User-centric Identity Management Technologies. In *SECURWARE*. IEEE, 12–17.
[42] Ryan Moats. 1997. URN Syntax. *RFC* 2141 (1997), 1–8.

[43] Alfred Moore. 2018. Anonymity, Pseudonymity, and Deliberation: Why Not Everything Should Be Connected. *Journal of Political Philosophy* 26, 2 (2018), 169–192. https://doi.org/10.1111/jopp.12149 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/jopp.12149

[44] Stefan More. 2023. *Trust and Privacy in a Heterogeneous World.* PhD thesis. Graz University of Technology. Available at https://doi.org/10.3217/p65m8-j3q66.

[45] Stefan More, Jakob Heher, and Clemens Walluschek. 2022. Offline-verifiable Data from Distributed Ledger-based Registries. In *SECRYPT*. SCITEPRESS, 687–693.

[46] Stefan More, Sebastian Ramacher, Lukas Alber, and Marco Herzl. 2022. Extending Expressive Access Policies with Privacy Features. In *TrustCom*. IEEE, 574–581.

[47] Yngve N. Pettersen. 2013. The Transport Layer Security (TLS) Multiple Certificate Status Request Extension. *RFC* 6961 (2013), 1–10.

[48] Andreas Pfitzmann and Marit Hansen. 2010. A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management. v0.34.

[49] David J. Phillips. 2002. Negotiating the Digital Closet: Online Pseudonymity and the Politics of Sexual Identity. *Information, Communication & Society* 5, 3 (2002), 406–424. https://doi.org/10.1080/13691180210159337 arXiv:https://doi.org/10.1080/13691180210159337

[50] Nils Quermann and Martin Degeling. 2020. Data Sharing in Mobile Apps - User Privacy Expectations in Europe. In *EuroS&P Workshops*. IEEE, 107–119.

[51] Kai Rannenberg, Jan Camenisch, and Ahmad Sabouri (Eds.). 2015. *Attribute-based Credentials for Trust: Identity in the Information Society.* Springer.

[52] Ashwini Rao and Jürgen Pfeffer. 2020. Types of Privacy Expectations. *Frontiers Big Data* 3 (2020), 7.

[53] Marco Robol, Elda Paja, Mattia Salnitri, and Paolo Giorgini. 2018. Modeling and Reasoning About Privacy-Consent Requirements. In *PoEM (Lecture Notes in Business Information Processing, Vol. 335)*. Springer, 238–254.

[54] Yashothara Shanmugarasa, Hye-Young Paik, Salil S. Kanhere, and Liming Zhu. 2021. Towards Automated Data Sharing in Personal Data Stores. In *PerCom Workshops*. IEEE, 328–331.

[55] John Solis and Gene Tsudik. 2006. Simple and Flexible Revocation Checking with Privacy. In *Privacy Enhancing Technologies (LNCS, Vol. 4258)*. Springer, 351–367.

[56] Manu Sporny, Kyle Den Hartog, Grant Noble, Daniel Burnett, Dave Longley, and Brent Zundel. 2022. *Verifiable Credentials Data Model v1.1.* W3C Recommendation. W3C. https://www.w3.org/TR/2022/REC-vc-data-model-20220303/.

[57] Jonah Stegman, Patrick J. Trottier, Caroline Hillier, Hassan Khan, and Mohammad Mannan. 2023. "My Privacy for their Security": Employees' Privacy Perspectives and Expectations when using Enterprise Security Software. In *USENIX*. USENIX Association, 3583–3600.

[58] O. Terbu, T. Lodderstedt, K. Yasuda, and T. Looker. 2023. OpenID for Verifiable Presentations - draft 20. https://openid.net/specs/openid-4-verifiable-presentations-1_0.html

[59] Alin Tomescu, Vivek Bhupatiraju, Dimitrios Papadopoulos, Charalampos Papamanthou, Nikos Triandopoulos, and Srinivas Devadas. 2019. Transparency Logs via Append-Only Authenticated Dictionaries. In *CCS*. ACM, 1299–1316.

[60] Ministerie van Binnenlandse Zaken en Koninkrijksrelaties. 2024. *Netherlands: Public Reference Wallet (Version 0.1.26).* https://github.com/MinBZK/nl-wallet

[61] Jan Vossaert, Jorn Lapon, Bart De Decker, and Vincent Naessens. 2013. User-centric identity management using trusted modules. *Math. Comput. Model.* 57, 7-8 (2013), 1592–1605.

[62] Surong Yan, Xiaolin Zheng, Deren Chen, and Wen-Yu Zhang. 2011. User-centric trust and reputation model for personal and trusted service selection. *Int. J. Intell. Syst.* 26, 8 (2011), 687–717.