Podda, Emanuela; Hölzmer, Pol; Amard, Alexandre; Sedlmeir, Johannes; Fridgen, Gilbert

Article

The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets

Internet Policy Review

Provided in Cooperation with:
Alexander von Humboldt Institute for Internet and Society (HIIG), Berlin

Suggested Citation: Podda, Emanuela; Hölzmer, Pol; Amard, Alexandre; Sedlmeir, Johannes; Fridgen, Gilbert (2025) : The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets, Internet Policy Review, ISSN 2197-6775, Alexander von Humboldt Institute for Internet and Society, Berlin, Vol. 14, Iss. 3, pp. 1-29, https://doi.org/10.14763/2025.3.2019

INTERNET POLICY REVIEW

Volume 14 **Issue** 3

RESEARCH ARTICLE

OPEN ACCESS

PEER REVIEWED

The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets

Emanuela Podda *Università degli Studi di Milano*
Pol Hölzmer *University of Luxembourg*
Alexandre Amard *University of Luxembourg*
Johannes Sedlmeir *University of Münster*
Gilbert Fridgen *University of Luxembourg*

DOI: https://doi.org/10.14763/2025.3.2019

Published: 30 July 2025
Received: 6 August 2024 **Accepted:** 1 April 2025

Funding: This research was supported in part by Luxembourg's Ministry for Digitalisation, PayPal and the Luxembourg National Research Fund (FNR), Luxembourg (P17/IS/13342933/ PayPalFNR/Chair in DFS/Gilbert Fridgen, as well as by the FNR, grant reference 16326754 (PABLO).
Competing Interests: The author has declared that no competing interests exist that have influenced the text.

Citation: Podda, E., Hölzmer, P., Amard, A., Sedlmeir, J., & Fridgen, G. (2025). The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets. *Internet Policy Review*, *14*(3). https://doi.org/10.14763/2025.3.2019

Keywords: Electronic attestation, Electronic identification, eIDAS, GDPR, Zero-knowledge proofs

Abstract: The recent amendment to the European eIDAS Regulation has established the European Digital Identity Framework, which introduces electronic attestations of attributes. Technically, these attestations involve auxiliary information to ensure their verifiability, leading to the generation, processing, and storage of more than just personal data. In particular, this auxiliary information contains globally unique information that can be misused as personal identifiers and poses risks to the privacy of individuals engaging in transactions using a European Digital Identity Wallet. As such, they create tension with the principle of data minimisation under the General Data Protection Regulation (GDPR). On the positive side, privacy-enhancing technologies, especially zero-knowledge proofs (ZKPs), are rapidly advancing and capable of addressing this tension. In this paper, we analyse the impact of the availability of these techniques on legal compatibility in the European electronic identification context and explore the tension field between the technical requirements of the digital identity wallet and the GDPR's data minimisation principle. We illustrate this dynamic through the specific examples of cryptographic data processed to ensure the authenticity and integrity of attributes' electronic attestations and shed light on how ZKPs can support legal compliance. This paper contributes to the privacy-oriented electronic identity management literature by providing policy and technical recommendations for achieving data minimisation compliance. We emphasise the necessity for regulatory bodies to enforce the use of advanced solutions like ZKPs to achieve unlinkability and unobservability. Accelerating the standardisation of these technologies is crucial for safeguarding user privacy and achieving seamless regulatory compliance in digital identity systems.

Introduction

Electronic identification is an essential process that empowers natural and legal persons to access a wide array of online services (Stevens et al., 2010). This dynamic further contributes to the increasing datafication of our society and economy (Van Dijck, 2014). Naturally, identification transactions involve the collection and processing of substantial amounts of personal data, raising important questions regarding privacy and data protection (Monteiro, 2023). In this paper, we address how privacy-by-design principles can be implemented by focusing on data minimisation in contexts that necessitate the verifiability of (legal) identity-related information (Tsakalakis, 2020).

Technical and management literature provide a spectrum of approaches and terminologies related to digital identity management (Pfitzmann et al., 2010), focusing on the growing trend in data-driven societies to digitally represent every facet of human life, while emphasising individual autonomy, self-empowerment, and data sovereignty. This reflects a broad expression towards the right to self-determination, meant as a form of control over personal information (McCorquodale, 1994). The design of corresponding privacy-oriented solutions is still being explored and debated in the context of modern information systems (IS) research (Giannopoulou & Wang, 2021; Mejias & Couldry, 2019). Furthermore, within the legal literature, Purtova recently highlighted that "*relatively little attention is paid in law and legal scholarship to what identification is*" (Purtova, 2022). Along with this line of re-

search, unfolding the limits of data protection law and personal re-identification, this article is intended for legal experts, policy-makers, and technologists alike. It thus adopts a dual perspective throughout the paper, grounded in legal reasoning but supported by technical analysis.

In the European Union (EU), the electronic Identification, Authentication and Trust Services Regulation (eIDAS) establishes norms for what and how citizens' identity information should be collected, shared, and verified with and by service providers (EU 910/2014). This allows qualified services to be provided to citizens, even cross-border, including identification and authentication for accessing governmental services and the issuance of qualified electronic signatures. By mandating specific identity processes to reach a high level of assurance (LoA), the regulation aims to establish the trustworthiness of these services (EU 1502/2015). Additionally, eIDAS 2.0 (EU 2024/1183) further introduced the concept of an EU Digital Identity Wallet (EUDIW) that citizens can use to store and present qualified and non-qualified electronic attestations of attributes.

However, the information processed in the context of eIDAS-compliant transactions does not consist exclusively of data subjects' identity attributes but also involves additional types of legally qualifiable personal information such as identifiers (Pfitzmann et al., 2010). This auxiliary data in identification transactions technically ensures identity assertions' integrity, authenticity, and validity (Grassi et al., 2017). Thus, electronic identification transactions may release personal data beyond what the holder intends to share, which is beyond the business case (Fuster, 2014). As such, in this paper, we argue that compliance with the data minimisation principle requires privacy-enhancing technologies (PETs), such as zero-knowledge proofs (ZKPs), to avoid the disclosure of auxiliary personal data without hampering verifiability.

While European jurisprudence and national and international advisory boards provide guidance on contextual application, their course struggles to keep pace with technological development. There is a lack of specific scientific evaluations on the impact of technological solutions on legal compliance in the context of identification transactions (EU Agency for Cybersecurity, 2022). While Recital 14 of the revised eIDAS regulation mentions ZKPs, their use is not mandatory. Accordingly, ZKPs are also not required in the latest version of the architecture and reference framework (ARF) (European Commission, 2025), a technical specification adopting the European Commission's (EC) Recommendation (EU) C(2021) 3968 on a common Union Toolbox for a coordinated approach towards a European Digital Identity Framework. This Toolbox is developed by the eIDAS Expert Group and the EC, in

collaboration with industry stakeholders (eIDAS Expert Group, 2021; EC, 2024). The ARF can be considered the technical counterpart of the eIDAS 2.0 regulation, outlining technical specifications, process guidelines, and system design recommendations for a robust, cross-border, and citizen-centric identity management model (eIDAS 2.0 (14, 15)).

Given this regulatory perimeter, we draw inspiration from legal literature (Purtova, 2022), building on a cross-disciplinary approach that integrates existing legal and technical scholarship on electronic identification, as well as the regulation of technological innovation. We provide both a legal analysis of data protection and electronic identification legislation, as well as a technical analysis of the latest technologies supporting data protection. We aim to explore the application of data minimisation in the context of electronic identification in the EU and the interplay between the technological evolution of identification transactions, the GDPR, and context-specific legislation, namely the eIDAS regulation. The research presented is rooted in doctrinal legal research (McConville & Chui, 2017), also known as traditional legal research or black-letter legal research. We provide a systematic exposition of the principles, rules, and concepts governing a particular area of law (Smits, 2015) and analyse their relationships to resolve ambiguities in existing law (Van Hoecke, 2011). Further, we encompass technical analyses and reflections, proposing a law and technology analysis (Kanevskaia & Pałka, 2023) on data protection and privacy (Monteiro, 2023).

This paper is structured as follows. Section 2 introduces the regulatory framework, specifically focusing on eIDAS and GDPR, and examines how data minimisation applies to electronic identification. Section 3 delves into the technical implementation of digital identity transactions and PETs, with an in-depth exploration of electronic attestations of attributes and ZKPs. In Section 4, we discuss the role of auxiliary data as identifiers in these transactions and how ZKPs can help address challenges related to data minimisation. Section 5 provides a broader discussion on how PETs can influence the application of the data minimisation principle. Finally, Section 6 summarises the contributions of the paper and discusses limitations, followed by the conclusion in Section 7.

Legal background

The EU's General Data Protection Regulation (GDPR) defines "*personal data*" as "*any information related to an identified or identifiable natural person*," underscoring the implications of individual identity in the digital realm. This definition was transposed from the Data Protection Directive (95/45/EC) to the GDPR and echoes the

traditions of the legally binding international data protection instruments: the 1980 OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data and the 1981 Council of Europe Convention 108 for the Protection of Individuals with regard to Automatic Processing of Personal Data (Fuster, 2014). These documents intertwine data protection with privacy, mirroring the approach pioneered by the United States (US) for privacy that is anchored to the protection of any information related to an individual, defined as Personally Identifiable Information (PII) (NIST SP 800-122).

Despite the varied and nuanced lexical differences in EU and US legislative terminology (*personal data* in EU versus *PII* in the US), it is unequivocal that any electronic or digital representation of PII qualifies as personal data as it permits personal re-identification (Figure 1). Nevertheless, the conceptualisation of these electronic and digital representations may vary depending on the technical and computational perspectives for investigating personal re-identification and privacy. In this regard, according to Torra (2017), three main scientific communities are working on privacy and personal re-identification, addressing the protection and minimisation of personal data. One with a statistical background and working on statistical disclosure control (e.g., Hundepool et al. 2012), one with a background on databases and data mining, working on privacy-preserving data mining (e.g., Aggarwal et al. 2008, Samarati 2001), and one with a communication and security background, working on secure and verifiable computations (e.g., Heurix et al. 2015). We focus on the perspective of PETs, unfolding and addressing the limits of data protection law with a computational or cryptographic approach that is reasonably valuable in the context of electronic identification, for several considerations that will be discussed further.

However, before delving into the technical perspective, plainly speaking, electronic identification pertains to the process of using personal data, in electronic form, that uniquely represents a natural person. Within this context, the terminology used by the European legislator seems to differ from the general umbrella term of "*personal data*" used in the GDPR. Such terminology seems to be rather enriched. To this purpose, with the aim of connecting legal and technical nuances, it appears imperative to unfold and determine the applicable legislation and the correlated terminology for the multitude of data generated by electronic identification. Some of it falls literally within the scope of application of the context-specific legislation (eIDAS), while some other, although generated in the process of electronic identification, falls outside of it as it is not specifically regulated. Nevertheless, given that it could permit personal re-identification (as auxiliary information not serving the

purpose of electronic identification), it must necessarily fall within the scope of application of the GDPR and, legally qualifying as personal data, it must be subjected to GDPR principles, among which data minimisation. Such considerations appear to be reasonably necessary even in case the context-specific legislation is without prejudice to the GDPR.

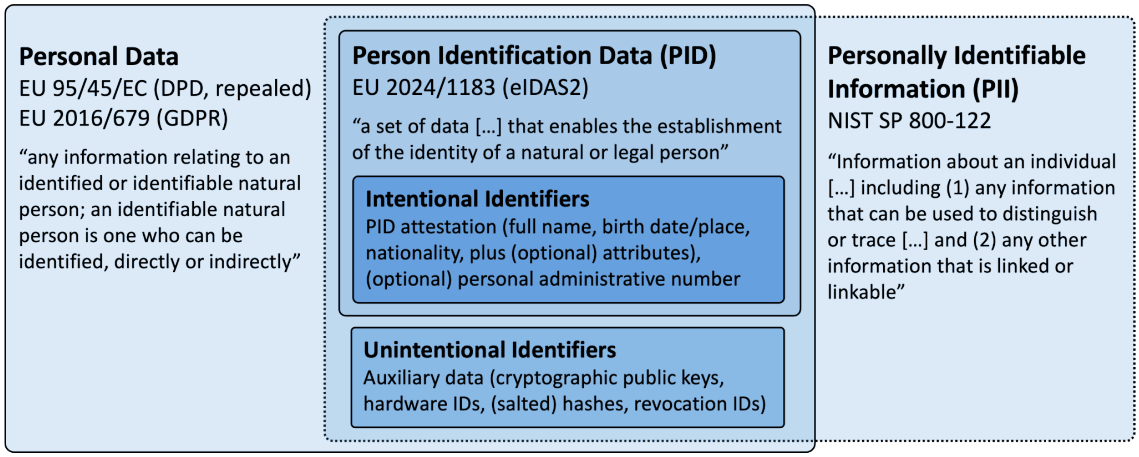

FIGURE 1: Identification data in different contexts and by overlapping definitions relevant to the technical implementation

2.1 eIDAS regulation

The first regulatory framework on electronic identification in the EU was the regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market, which repealed Directive 1999/93/EC. As literally stated in Recital 2, the regulation aimed to "*enhance trust in electronic transactions in the internal market by providing a common foundation for secure electronic interaction between citizens, businesses and public authorities, thereby increasing the effectiveness of public and private online services, electronic business and electronic commerce in the Union.*" After carrying out its evaluation (*ex* Art. 49) in 2021, the EC presented the Evaluation Report (European Commission, 2021a) and a Working Document (European Commission, 2021b) providing an overview of the implementation and application of the regulation. The EC acknowledged benefits and limitations, highlighting five shortcomings as *key issues*. Among them is the lack of available digital identity solutions to sufficiently address the evolving data control and security concerns.

To address these key issues, the EC proposed a revision of the eIDAS regulation (European Parliament and Council, 2021) (eIDAS 2.0) to re-set rules and principles of electronic identification in the EU. The EC agreed-on text of eIDAS 2.0 (Euro-

pean Parliament and Council, 2024) listed, among its objectives, the need to give citizens a higher degree of control over the release of their identity data and strengthen security in the digital market. This was essentially a response to non-EU private sector companies increasingly acting as gatekeepers to access services through "*single sign-on*" features and using their position to collect vast amounts of identity and service usage data (Weigl et al., 2022).

The new regulation (eIDAS 2.0) shifts the focus to a citizen-centred approach, placing the concept of decentralised digital identity at the core of its strategy. The eIDAS 2.0 entered into force in May 2024, introducing the EUDIW, a digital identity wallet that allows EU citizens, residents and businesses to store and manage their personal identification data and other information defined as *attributes*[1] in a bid to provide control over their digital identities and offer a level playing field for service providers in the single digital market (Rieger et al., 2022). In addition, it streamlines verifiable data exchange, which today is often insecure (e.g., video-based identification), not machine-readable (e.g., scanned signature), or inflexible (e.g., limited to national electronic identification) (Rieger et al., 2024). Digital identity wallets are no new concept and exist in various forms, such as mobile apps, browser extensions, or dedicated hardware (Podgorelec et al., 2022). However, eIDAS 2.0 marks the first time for deployment at a continental scale, reshaping the boundaries and interactions in identity and access management as a data-centric domain (Glöckler et al., 2023).

2.2 Engineering data minimisation in electronic identification transactions

Electronic identification under eIDAS is the process of using person identification data, *inter alia*, to authenticate individuals in online services (EU/910/2014, Art. 3(1) and (5)). Electronic identification, like any transaction involving personal information of EU citizens, is subject to data minimisation requirements. Article 5.1(c) of the GDPR introduces specific requirements for data collection and processing, demanding that data should be "*adequate, relevant and limited to what is necessary to achieve the purpose that justified the initial collection*." Recital 39 of the GDPR further clarifies this principle, stating that, on the one hand, personal data should only be processed if the purpose cannot be "*reasonably achieved by other means*" and, on the other hand, data should not be kept longer than necessary to achieve the purpose. This principle also appears in Article 25 as a requirement in engineering

1. According to the legal definitions provided in eIDAS regulation, '*attribute*' means a characteristic, quality, right or permission of a natural or legal person or of an object; '*electronic attestation of attributes*' means an attestation in electronic form that allows the authentication of attributes.

data protection "*by-design and by-default*" (European Data Protection Board, 2020).

Based on the literal interpretation of the law, data minimisation has two primary purposes. Firstly, to reduce personal data collection, storage, and processing, to mitigate risks to privacy (Deng et al., 2010). Secondly, to reduce other risks related to profiling, automated decision-making, incompatible uses of data, and lack of control and transparency (Tosoni, 2020). In addition, the data minimisation principle imposes further considerations on determining the purpose of processing (Finck & Biega 2021) and the period during which personal data may be stored, the so-called "*retention period*" (European Data Protection Board, 2022). Despite the many specifications and legal nuances in the GDPR, this requirement is often considered vague, difficult to interpret, and not aligned with technical approaches. In fact, the European regulator favours technological neutrality by not prescribing specific solutions and regulating information systems through legal principles whose operationalisation and integration have already been considered difficult (Finck & Biega, 2021).

In the context of digital identity management, technical approaches generally seek to reconcile the legal nuances between identification and anonymity. Pfitzmann, Dresden, and Hansen (2010) were pioneers in this line of research. They acknowledged the different approaches used to describe anonymity, highlighting the need for precise technical terminology. Their work clarifies that privacy can be engineered to minimise information disclosure through different anonymity measures: *unlinkability* (the inability to correlate separate items with the same subject), *undetectability* (the inability to verify whether an action relates to a subject), *unobservability* (the inability to detect a subject's actions), and *pseudonymity* (the use of alternative identifiers to mask an identity).

The legislative action attempted to reflect on these nuances, stimulating the integration of technological development in the rollout and evolution of the eIDAS regulation. It should be noted that data minimisation can also be relevant when full unlinkability cannot be achieved, e.g., because the relying party requires the disclosure of some identity attributes for their business processes. In this case, reducing the amount of linkable information shared to the minimum, under the restrictions posed by the verifier, is still a legally desirable outcome (Babel & Sedlmeir, 2023).

In fact, in its original setting, the eIDAS (2014) regulation had been criticised for its insufficient implementation of the data minimisation principle and specifically for lacking selective disclosure for identifiers (Tsakalakis et al., 2019). Selective

disclosure is the process of revealing only specific pieces of information necessary for a particular purpose. The eIDAS regulation originally required that identification transactions disclose the full set of mandatory personal identification attributes, even when a relying party could fulfil its service with only a subset of that information (eIDAS eID Technical Subgroup, 2019). The EC acknowledged this criticism in the 2019 eIDAS Impact Assessment (European Commission, 2021). It introduced digital wallets as a means to enable citizens to have a higher degree of control over their identity data and its usage. In this regard, Recital 29 of the agreed eIDAS 2.0 text clarifies that "*selective disclosure is a concept empowering the owner of data to disclose only certain parts of a larger data set, for the receiving entity to obtain only such information that is necessary for the provision of a service requested by a user.*" It further specifies that selective disclosure as a basic design feature of the EUDIW. This requirement is introduced by Article 6a (3) (a) in combination with recognising the legal effect of the electronic attestation under Article 45a (2).

Technical implementation

3.1 Electronic attestations and verifiable presentations

Electronic attestations, in the form of digital certificates, have been heavily used since the 1990s, when the secure socket layer (SSL) protocol was developed, to prove a server's identity and establish authenticated and confidential channels on the Web (Housley et al., 2002). They have since become a critical component of online security, establishing trust and confidence in online communication and transactions. Subsequently, they have been further adapted to attest the identities of other entities, including natural persons (W3C, 2022).

The EUDIW is a modern approach to user-centric identity management that enables natural and legal persons to manage their personal identification data (PID) alongside other generic electronic attestations of attributes (EAA). PID, in the context of eIDAS 2.0, is a special type of EAA, which constitutes a quasi-identifier with a well-defined set of attributes that shall be unique in the context of the issuing country (See Figure 1). Individuals ("*holders*" as data subjects) can share (parts of) their EAAs with relying parties ("*verifiers*"). More specifically, verifiers can request proof of specific attributes and restrict which issuers they consider trustworthy to provide this proof. In response to this "*proof request*" – and contingent on the user's consent – the holder's digital wallet then selects suitable EAAs and creates a "*verifiable presentation*" that is subsequently sent to the verifier.

Verifiability is the prerequisite of trust in digital interactions and is essential to

mitigate fraud (Menezes et al., 1996b). In this sense, a verifiable presentation is a technical means to package identity data derived from electronic attestations and enrich it with auxiliary information that can be used to verify its integrity, authenticity, and validity (Grassi, 2017; W3C, 2022). A verifiable presentation is issued by a digital identity wallet with the user's consent in response to a verifier's "*proof request*" (Glöckler et al., 2023).

Figure 2 illustrates the technical representation of an EAA, as per the "*verifiable credential*" data model (W3C, 2022), which is one of two mandated data formats as per EU/2024/2977. This data model extends beyond mere identity attributes; it embodies a comprehensive construct that integrates attributes and auxiliary information, consisting of attestation metadata and cryptographic proofs. For any relying party to whom an electronic attestation is presented, identity attributes and attestation metadata are mere claims, which require proofs to ensure integrity, authenticity, and validity. Cryptographic proofs are, therefore, additional (non-human readable) attributes, such as cryptographic (binding) keys and signatures. Most prominently, cryptographic signatures ensure the integrity and authenticity of all other attestation attributes. Given those premises assured, the attestation metadata can be used to verify validity (e.g., expiry data, issuing authority, and issuing country as mandatory for PID as per EU/2024/2977), as well as identity attestation semantics (e.g., data types, namespace, and schema as per the verifiable credentials data model (W3C, 2022)). Figure 2 simplifies these technical concepts by comparing them to their analogue ID counterparts and (some of) their security features (Council of the European Union, 2022) to illustrate how this type of credentials ensures trust and verifiability in offline presentations.

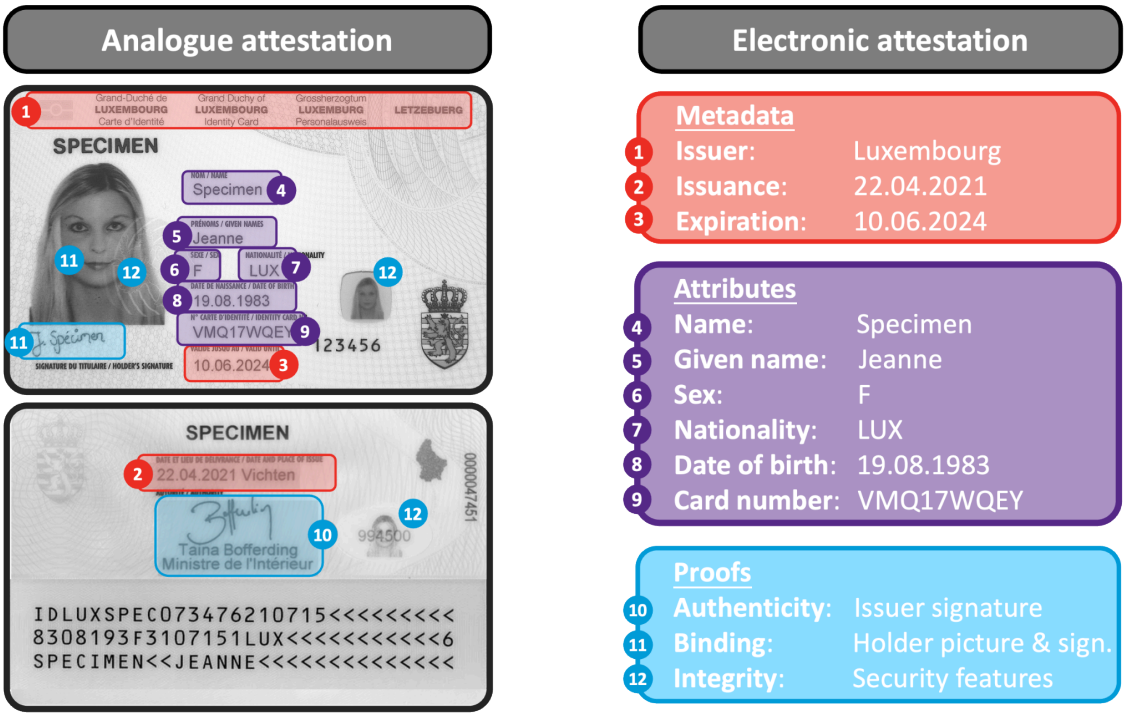

FIGURE 2: Electronic attestations illustrated as the natural analogue of a physical identity card

As illustrated in Figure 2, electronic attestations and, in turn, verifiable presentations thereof rest on three core security guarantees—integrity, authenticity, and validity—that together establish trust in digital identity transactions. Integrity ensures that an attestation remains unaltered from issuance through presentation: the issuer computes a collision-resistant checksum ("hash") of the full attestation and encrypts it with its private key; during verification, the verifier retrieves the issuer's public key from a trust registry, decrypts the signature, and compares the result with the result of hash against the received data to detect any tampering (Menezes et al., 1996a). Authenticity has two facets: issuer authenticity guarantees that the attestation originates from a trusted authority, achieved using cryptographic signatures, which implicitly also guarantees integrity, and holder authenticity (or "holder binding") confirms that the presenter is the legitimate owner of the attestation. In the case of an analogue attestation, this can be achieved by comparing the picture on the credential with the face of the presenter. In contrast, digital holder binding in remote interactions is achieved by including the owner's public key in the electronic attestation in the issuance process. In the verifiable presentation, the user can then prove they are the legitimate holder through a challenge–response protocol in which the wallet's secure element signs a verifier-issued nonce with the holder's private key; the verifier then checks this signature against the public binding key embedded in the attestation (ARF, 2025). Finally, validity assures that the credential is current and has not been revoked by em-

bedding "not valid before" and "not valid after" timestamps, as well as revocation IDs, in the metadata and consulting revocation registries or status lists at verification time. Together, these mechanisms enable digital wallets to produce verifiable presentations that mirror the integrity, authenticity, and validity properties of physical identity credentials.

The verification process (depicted in Figure 3) relies on a challenge–response protocol: the (trusted) verifier sends a (random) value, which the wallet cryptographically signs using the private binding key stored in the secure element as proof of possession (holder-binding), and the verifier confirms possession by checking the embedded public binding key (Babel & Sedlmeir, 2023). This process, however, exposes the full attestation, including auxiliary information such as the holder's public key, the issuer's digital signature, and the attestation revocation identifier. Modern systems commonly employ "selective disclosure" to limit exposure of identity attributes: each attribute (e.g., date of birth) is first combined with a secret (random) value (the "salt") and then hashed, yielding a unique "salted hash". During verification, only the specific salted hash and its salt are disclosed; the verifier recomputes the hash, confirms that it matches the corresponding part of the signed record, and then validates the digital signature.

This approach has been standardised in different data formats such as JSON-SD by the Internet Engineering Task Force (IETF) and as MDOC by the International Organisation for Standardisation (ISO) (ETSI, 2023), and is, therefore, also included in the current version of the ARF (European Commission, 20254b). While this approach effectively conceals any non-requested attributes—advancing data-minimisation goals—every salted hash and salt still acts as a persistent fingerprint of the electronic attestation. Furthermore, selective disclosure cannot be applied to hide the holder-binding public key and the issuer's digital signature. By correlating these fingerprints across multiple transactions, an observer can link separate presentations, demonstrating that metadata alone can undermine anonymity even when direct attribute leakage is prevented. Indeed, even presentations with selective disclosure are vulnerable to various types of linkability (ISO/IEC 27551:2021; Pfitzmann et al., 2010). Consequently, the ARF potentially violates eIDAS 2.0 Article 5a(16). Linkability is especially problematic because untrusted entities, such as attribute providers and relying parties acting together, can correlate and link auxiliary information to the same user, thereby breaching privacy and enabling tracking, profiling, or de-anonymisation.

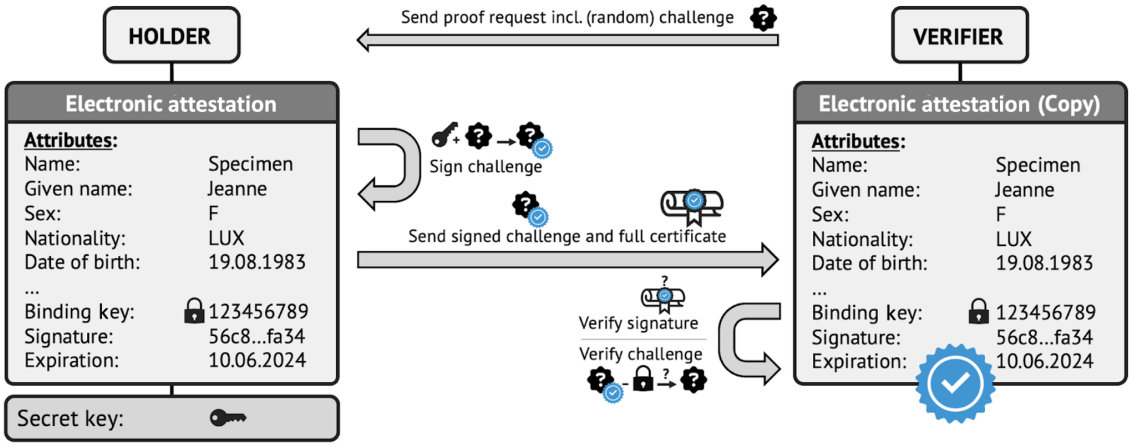

FIGURE 3: Presenting and verifying an electronic attestation.

3.2 Privacy enhancing technologies

In the context of electronic identification, achieving data minimisation in the form of selective disclosure is not trivial because sharing only a subset of an electronic attestation may compromise its verifiability. More specifically, selective disclosure focuses on an interaction between a user and a service that depends on the user providing certain identity attributes or credentials in a verifiable way. While a minimal set of identity attributes must be transferred to deliver the service (Sedlmeir et al., 2021), *ex* GDPR requirement, data minimisation entails limiting personal data exchange to what is strictly necessary for the provision of the service. For instance, unlinkability may be desirable when considering repeated interactions of a user with the same service or different interactions with different services in which no identity attributes need to be disclosed. On the other hand, as a corollary to the mandatory disclosure of attributes, it becomes possible to use the linkability of this data to derive a sophisticated user profile. This is particularly true when auxiliary information is collected beyond the identity attributes explicitly required for the identification transaction. In this context, PETs, such as ZKPs, "*allow a relying party to validate whether a given statement based on the person's identification data and attestation of attributes is true, without revealing any data on which that statement is based, thereby preserving the privacy of the user*" (EU/2024/1184, Recital 14). In the well-established (Wairimu et al., 2024) LINDDUN framework (Deng et al., 2010), privacy is defined as the protection of individuals' rights to control how their personal data is collected, processed, and shared, going beyond mere confidentiality (unauthorised-access protection) to encompass the broader notion of user control over information disclosure. To realise this, LINDDUN augments the classic security goals (confidentiality, integrity, availability) with three privacy-specific protection goals: unlinkability (preventing personal data linkage across con-

texts), intervenability (enabling individuals to influence or halt data processing), and transparency (ensuring individuals are informed about how their data is handled)—thereby countering threats such as linkability, identifiability, detectability, data disclosure, unawareness, non-repudiation, and non-compliance. PETs guarantee a reasonable accuracy-privacy trade-off (Ah-Fat & Huth, 2019; Garrido et al., 2022).

PETs encompass a range of solutions from software-based to hardware-based components. Software-based PETs rely on applications and algorithms. In contrast, hardware-based PETs leverage tamper-resistant secure elements (SEs) featuring dedicated processors, memory, and cryptographic co-processors to process and store sensitive data while protecting it from extraction and manipulation, even by parties with physical access to the hardware. These chips are available in various formats, such as (embedded) SIM cards (eSIM) or (remote) trusted execution environments (TEEs), with varying degrees of adoption and security. (Bastian et al., 2023). These technologies ensure secure data processing, storage, and management, often surpassing the security capabilities of software-based approaches alone, working as support or enablers for software-based PETs for storing and generating cryptographic material.

Among the software-based techniques, ZKPs seem to be particularly suitable for identification and authentication transactions (Sedlmeir et al., 2021). ZKPs are characterised by allowing "*for the validation of a claim, fact, capacity, or identity, without requiring the 'prover' to reveal to the 'verifier' any underlying information beyond the validity of the assertion in itself*" (Bamberger et al., 2022). As such, the zero-knowledge property can be considered a precise formulation of data minimisation in verifiable bilateral interactions (Babel & Sedlmeir, 2023). This feature becomes extremely valuable in identity management, as data subjects may be empowered to prove the strictly necessary subset of identity attributes and authorisations or statements derived from them. In other words, ZKPs can grant unobservability and unlinkability beyond the identity attributes that were explicitly requested – unless the verifier specifically asks for connections with previous identification processes. Other technologies that rely on hardware can provide comparable privacy guarantees, provided the hardware manufacturer is trusted (Garrido et al., 2022). The hardware-based German eID, for instance, is implemented based on this paradigm, thus achieving effective data minimisation in the form of selective disclosure and unlinkability (Bender et al., 2010). Yet, this approach is not easily portable to mobile phones and is challenging to scale to a variety of electronic attestations with heterogeneous security requirements (Babel & Sedlmeir, 2023).

Therefore, it is reasonable to consider that, after a potential transition period, the eIDAS 2.0 has explicitly opted for an approach based on electronic attestations stored in digital wallet apps running on mobile phones.

While our analysis focuses on ZKPs, the core premise of this paper holds for any (combination of) software- or hardware-based PETs that can eliminate the exchange of unnecessary personal data in identity interactions, such as any type of secure, multi-party computation, cloud-based TEEs, or homomorphic encryption, to name a few alternative approaches (Baum et al., 2023; Spensky et al., 2016). Openness with regard to the used PETs also resonates with the principle of technological neutrality in the EU regulatory framework. We rely on ZKPs due to their comparably high degree of practical maturity with relatively straightforward integration into today's wallet architectures (Schwarz et al., 2022; Frigo & Shelat, 2024), which strengthens the case for their mandatory adoption under eIDAS as further highlighted in the explicit mention in the (non-binding) Rectial 14, which states that: "*Member States should integrate different [PETs], such as [ZKP], into the [EUDIW]. Those cryptographic methods should allow a relying party to validate whether a given statement based on the person's identification data and attestation of attributes is true, without revealing any data on which that statement is based, thereby preserving the privacy of the user.*"

Data minimisation: the promise of zero knowledge proofs

4.1 Issuer digital signature

As shown in the previous section, the issuer's digital signature represents an essential part of an electronic attestation for verifying integrity and issuer authenticity. Because each signature is generated by the issuer's private key over the exact attestation data, it serves as a unique cryptographic identifier, backed by the collision resistance of the hash function. This uniqueness causes concerns regarding linkability and observability. (Colluding) Verifiers can use this unique identifier to combine and track (cross-system) usages of the same attestation. Therefore, even when a holder consents to share or present electronic attestations with a verifier through their digital wallets, they may not realise that an attestation-specific unique identifier is also shared as embedded metadata (Glöckler et al., 2023). This issue even persists when subsets of digital identity attributes are selectively disclosed in different verifiable presentations using salted hashes. Consequently, if there are other ways to verify the integrity and authenticity of electronic attestations, the sharing of digital signatures may conflict with the data minimisation

principle, as the verifier can use digital signatures beyond integrity checks to link a user's actions that involve this attestation.

Notably, advanced verification methods that circumvent the sharing of linkable digital signatures have existed for years. Electronic attestations that provide the required features to ensure verifiability are commonly termed anonymous credentials or attribute-based credentials (Camenisch & Lysyanskaya, 2001; ETSI, 2023; Kaaniche et al., 2020). In fact, the current version 1.9 of the ARF acknowledges the suitability of mechanisms underlying anonymous credentials for data minimisation compliance. The arguably most common implementations of anonymous credentials rely on Camenisch-Lysyanskaya (CL) (Camenisch & Lysyanskaya, 2004) and Boneh-Boyen-Shacham (BBS) (Boneh & Boyen, 2004) signatures. While CL signatures are (to the best of the authors' knowledge) not standardised, BBS signatures have been specified by IETF (Looker et al., 2023). In both approaches, the holder no longer needs to send the electronic attestation directly to the verifier. Instead, the holder sends the requested identity attributes to the verifier and generates a cryptographic ZKP to convince the verifier of integrity and issuer authenticity. This proof convinces the verifier that the holder indeed possesses an electronic attestation that includes these identity attributes that have been signed by the issuer (Glöckler et al., 2023). Moreover, this proof will always look different, even if the requested attributes coincide. A notable application of these concepts in practice is the British Columbia Government's use of anonymous credentials based on the Hyperledger AnonCreds specification, demonstrating a commitment to user-centric and privacy-focused digital identity solutions (BCGov, 2024).

4.2 Holder binding

Similarly to the case of the issuer's digital signature, the public holder binding key is used for verifying holder authenticity via a signed challenge, which also constitutes a unique cryptographic identifier (Paquin et al., 2024; Frigo & Shelat, 2024). Holder binding creates a non-negligible linkability risk since the holder's public key needs to be shared with every attribute presentation. As elaborated in Section 4.1, this enables a verifier to potentially link separate transactions originating from the same holder when it uses an electronic attestation repeatedly, or even different verifiers to track the use of an attestation if they share data from their identification transactions. As in the case of the issuer's digital signature, the user/data subject is arguably unaware of this unique identifier, and it is unclear if the processing is strictly necessary for checking the authenticity of the verifiable presentation.

Indeed, CL and BBS-based anonymous credentials can also hide the holder-binding public key from the verifier (Kakvi et al., 2023). Instead of sharing the attestation that includes the signed holder public key and the challenge signed with the secret key, the holder creates a cryptographic ZKP that they know the secret key corresponding to the holder binding public key specified in the attestation, using the random challenge specified by the verifier, without revealing the holder binding public key. As such, CL and BBS-based anonymous credentials also allow us to avoid this unintentional identifier. However, because CL and BBS signatures are not widely used, the corresponding algorithms for creating a cryptographic proof of holder binding are not implemented in dedicated cryptographic hardware and, in particular, not those available in mobile phones on the market in the form of embedded SEs (Bastian et al., 2023). Because the SE must protect the private key needed to generate the ZKP of possession of the corresponding private key in high LoA settings, the aforementioned anonymous credential schemes are thus currently not compatible with hardware binding in consumer devices commonly used to host digital wallets.

However, there are novel constructions of anonymous credentials that avoid this shortcoming, which are based on general-purpose ZKPs (Frigo & Shelat, 2024). They allow the holder to sign the random challenge from the proof request using their secure element and then to use this signed challenge as an input to creating the ZKP that convinces the verifier of the integrity of the digital signature and the knowledge of the private holder binding key by means of knowledge of the signed challenge. This approach was first proposed by Delignat-Lavaud et al. (2016) for turning X.509 certificates into anonymous credentials and is further discussed, e.g., by Rosenberg et al.(2023). Plastically speaking, these general-purpose ZKPs change the paradigm of verifiable presentations in the following way: Instead of sending the electronic attestations and auxiliary data to the verifier to run the verification computations, the user performs the verification locally, thereby obviating the need to share the auxiliary data with the relying party. To convince the relying party, even though they lack access to the user's device, the user generates a proof about the correct execution of the verification, i.e., that they indeed can create a verifiable presentation satisfying all the requirements (Babel & Sedlmeir, 2023). This (zero-knowledge) proof is sent to the relying party, and by definition of the zero-knowledge property, reveals no auxiliary information beyond the identity attributes requested by the relying party. However, general-purpose ZKPs arguably have a higher degree of technical complexity and a lower degree of technical maturity, which introduces potential risks when implemented in EUDIWs (Fernández, 2024).

Discussion

The preceding sections have highlighted a fundamental tension at the heart of modern electronic identification: the need for auxiliary information to verify claims over identity attributes versus the GDPR's necessity to minimise data collection. We first pointed out that selective disclosure paired with salted hashes already provides a technically mature means to reduce the set of attributes disclosed to the relying party without harming verifiability. Selective disclosure is a requirement in the ARF and mandated by eIDAS 2.0, which states that "*citizens [...] should be empowered to securely request, select, combine, store, delete, share and present data related to their identity [...], under the sole control of the user, while enabling selective disclosure of personal data*" (Art. 5a(4)(b)).

However, as we argue, even with selective disclosure, unintended identifiers remain. Specifically, we showcased that digital signatures used in almost any digital interaction to perform integrity and issuer authenticity checks, as well as holders' public keys used to perform proofs of holder binding, may be (mis)used as unique identifiers. As such, they enable the linking of transactions back to the same individual (i.e., re-identification), even if only attributes with weak linkability are selectively disclosed. We also argue that there are technical means to avoid this release of auxiliary data in verifiable presentations. Furthermore, expiration and revocation metadata (to ensure the validity of EAAs) may also lead to re-identification. These resulting linkability risks can, however, all be avoided using general-purpose ZKPs (Babel & Sedlmeir, 2023; Frigo & Shelat, 2024).

The described risks of linkability through auxiliary information should shed light on the necessity to comply with the data minimisation principle, especially in the context of identity management. This assessment will eventually improve the alignment between data minimisation and the original purpose of processing *ex* Art. 5 of the GDPR, namely electronic identification. Constraints within the design of current technical systems (e.g., ARF) make the implementation of data minimisation remarkably challenging (to impossible) without using advanced PETs, like ZKPs.

Acknowledging that the technological neutrality principle permits a margin of discretion in designing systems that process personal data, this acknowledged latitude necessitates ongoing assessment and elucidation of technological advancements, paired with the evolution of the regulatory frameworks. Such scrutiny aims to facilitate interdisciplinary discourse, thereby enhancing the understanding of the complexities associated with safeguarding personal data within a perpetually

advancing technological *milieu*.

De facto, according to the letter of the law, the eIDAS regulation's original setting allowed selective disclosure only for the primary mandatory attributes (i.e., name, surname, date of birth). However, the regulation did not consider or address the protection of information qualifiable as personal data, in the guise of auxiliary information allowing for linkability, and thus re-identification. Recalling the given examples, in both instances, the processing of the identifier had a clear and valid purpose: to prove the integrity, authenticity, and validity of EAAs. With the use of PETs – in particular, ZKPs – disclosing linkable auxiliary cryptographic information to compute such proofs becomes unnecessary. According to these premises, we derived two essential considerations to address: at first, concerning data minimisation in the context of electronic identification, and second, the legal interplay of the two main relevant bodies of legislation: the eIDAS and GDPR.

Regarding the first consideration, data minimisation compliance in the context of electronic identification depends on the current state of technology and its capabilities. This necessitates context-based assessments of data-driven systems where data quality and verifiability are evaluated. In this respect, PETs, such as ZKPs, can be implemented in verifiable presentations to provide the same integrity, authenticity, and validity assurances without disclosing linkable auxiliary information.

Regarding the second consideration, it is of paramount importance to distinguish between, on the one hand, identification and identifiability made possible by the massive amount of personal data (released, collected, shared, processed, and reused in the digital environment) and on the other hand, identification and identifiability in the context of electronic identification (aimed at authentication for access to and use of online services in bilateral interactions, where only data from one individual is required for decision-making).

While the first consideration is generally regulated by the GDPR, the second enjoys a context-specific regulation given by the eIDAS on electronic identification. This means that PID falls under the scope of the eIDAS. Conversely, auxiliary (personal) information, such as unintentional identifiers released in the context of identification and authentication, should be minimised as they fall under the scope of application of the GDPR, as long as they are necessary for the corresponding verifiability checks in electronic identification, and not expressively recalled by eIDAS regulation.

The eIDAS regulation does not reference the processing of identifiers used as tech-

nical means to ensure an electronic attestation's integrity, authenticity, and validity of an electronic attestation. Nevertheless, eIDAS is without prejudice to the GDPR (*ex* Art. 2.4 eIDAS 2.0). Moreover, as already highlighted by legal scholars (Ortalda et al., 2021) regarding the relationship between the two regulatory frameworks, neither eIDAS nor the GDPR clarifies which piece of legislation takes precedence over the other. In the event of such a contrast, given that the eIDAS regulation regulates a *specific data processing* – the processing of personal data for the purpose of authentication – would imply that the eIDAS regulation would be considered a *lex specialis* and thus prevailing over the GDPR as *lex generalis*. However, in the event of a regulatory gap in the *lex specialis*, the *lex generalis* would apply. The latter would be the case of unintentional identifiers, which do not find a normative recognition in the eIDAS regulation. Consequently, it should be subject to the principles and rules encompassed in the GDPR, therefore being compliant with *privacy by-design and by-default* settings.

Moreover, on a general note, the implementation of data minimisation sometimes generates a wrong perception of the strict correlation between data minimisation requirements and those imposed by purpose limitations. Due to system constraints, such correlation often imposes a trade-off between the requirements of data minimisation and the requirements of purpose limitation. As discussed above, electronic identification data transactions require auxiliary information; this is rarely necessary and proportional to the required information for the consented purpose of collection and processing.

Contributions and limitations

In summary, in this paper, we discuss three main considerations. First, in the context of electronic identification transactions, the protection of the inherent release of auxiliary personal data is not addressed by the eIDAS regulation (1.0). Despite the lack of specific protection in sector-specific legislation (eIDAS) and given the importance of protecting such information for reducing the auxiliary knowledge needed for re-identification (by linkability and observability), we consider that the protection of such identifiers should be granted in light of the general requirements imposed by the GDPR. We first unfold the different computational perspectives on privacy and re-identification, along with data protection, and pinpoint how PETs address the apparent conflict between achieving data minimisation and verifiability. Second, we argue why the reliance of digital identity wallets on electronic attestations, as foreseen in the ARF by the EU governments and relying parties, introduces a tension field with the GDPR's data minimisation principle that goes far

beyond a lack of support for anonymity in authentication. In particular, we discuss that capabilities for selective disclosure and the avoidance of unintended (cryptographic) identifiers as part of auxiliary information exchanged in verifiable presentations must be respected in assessing compliance with the GDPR principle of data minimisation, and that a consent-based approach to resolving this tension is insufficient. Third, the revised eIDAS text leading to eIDAS 2.0 may be considered a door opening towards more compliant minimisation solutions (e.g., selective disclosure). In this regard, we sustain that, among the different technical solutions that can be tailored to minimise the unintentional release of data, ZKPs seem to positively impact the minimisation of (auxiliary) information needed for verifiability. We discuss that the cryptographic techniques facilitating selective disclosure and eliminating the need for unintended (cryptographic) identifiers have different degrees of maturity. Consequently, to decide whether the GDPR mandates the use of these PETs and thus assess the GDPR compliance of the current ARF, a sophisticated maturity assessment of ZKPs is indispensable to weight data minimisation opportunities with technical complexity. We urge standardisation bodies such as ETSI to intensify their corresponding activities (e.g., ETSI 2023).

However, in this paper, we focus on the scenarios described in Sections 4.1 and 4.2 to support our conclusions and do not elaborate on the full scope of auxiliary information types in verifiable presentations that impose risks to privacy. In particular, we did not discuss linkability aspects of validity metadata (e.g., evidence for non-revocation, which is usually based on a unique credential identifier, and expiration, where the linkability risk increases with time resolution). Furthermore, we did not cover all scenarios where the data minimisation principle is violated if, e.g., only a range is relevant (e.g., status level on a loyalty card at least silver, date of birth at least 18 years in the past, salary at least 3000€ a month). More generally, if the verifier is required to run a program on one or multiple identity attributes to check a certain statement ("*predicate*"), it would be possible for the holder to create a cryptographic proof that this statement is valid ("*predicate proof*"), with a range proof being the most straightforward example but also complex statements possible (Rosenberg et al. 2023). Moreover, if the verifiable presentation is supposed to convince only the intended verifier, cryptographic evidence that would convince any verifier may also violate the data minimisation requirement. Both arbitrary predicate proofs and "*designated verifier proofs*" that are not convincing to any entity but the intended relying party can be implemented with today's ZKP tooling (Frigo & Shelat, 2024).

In this paper, we focus on the legal necessity for minimising disclosure of auxiliary

information in electronic attestations' proofs (see Figure 2). At the same time, we acknowledge that the technical implementation of the EUDIW entails wider privacy and security risks, which fall outside the scope of this paper. In Abellán et al. (2025), we provide a formal privacy evaluation of the EUDIW's architecture and reference framework (ARF). Using the LINDDUN methodology, we evaluate surveillance, secondary-use, and other privacy threats across a suite of five carefully selected threat cases. Note, however, that Van Dijck (2014) highlights that risks such as surveillance must be justified on a case-by-case basis against Europe's purpose limitation, proportionality standards, and only exceptionally for lawful interception. The Regulatory framework foresees the broadest technical implementation, hence ensuring that each Member State tailors the best solution that is in line with the history and values of each European country. This will also minimise the risk of eroding public trust in otherwise compliant identity infrastructures, as recent empirical work on the spread of technology conspiracy beliefs demonstrates (Trang et al., 2024).

Conclusion

The changing landscape of electronic identification requires a careful balance between verifiability and data minimisation. The interaction of GDPR and eIDAS regulation emphasises the challenge of implementing strong privacy protection in identity transactions. Transitioning from traditional cryptographic methods to advanced PETs like ZKPs is essential for enhancing privacy on a large scale. These innovative solutions provide greater control over data disclosure without compromising on security, ensuring compliance with data minimisation principles and better alignment with purpose limitations. This shift represents a significant advancement in privacy and identity management.

In this paper, we highlight three considerations concerning the usage of personal data in electronic identification, authentication, and trust services. First, their protection is critical to reduce the risk of re-identification by linkability and observability. Second, they are not specifically protected in eIDAS as context-specific legislation, and it remains unclear which regulatory framework – eIDAS or the GDPR as the gold standard for data protection – should apply in this context. We consider that the protection of such personal data should be granted in the light of the general requirements imposed by the GDPR, regardless of whether sector-specific legislation takes it into account. Third, we point out, using the example of ZKPs, that today's PETs may prevent the disclosure of additional personal data without hampering verifiability, and argue for their application being mandated by law-

makers.

Our analysis emphasises that the current reliance on conventional (cryptographic) methods and protocols conflicts with the GDPR's data minimisation requirements. The adoption of advanced technologies like ZKPs can mitigate these challenges by improving current approaches to selective disclosure and removing unnecessary data linkability. However, achieving widespread implementation of these technologies requires overcoming technical and regulatory hurdles. It is imperative that standardisation bodies, legal frameworks, and technological solutions evolve in tandem to ensure that data protection principles are upheld without compromising the functionality and trustworthiness of electronic identification systems. The interpretation of data protection principles and requirements should be tailored to the scenario and continuously re-evaluated, integrating the rapid pace of technical developments and feasibility. Technical and legal designers, as well as GDPR compliance auditors, must continually re-assess the interpretation of data protection requirements and integrate technological developments, as also mandated in Art. 25 of the GDPR.

References

Abellán Álvarez, I., Hölzmer, P., & Sedlmeir, J. (2025). *Privacy evaluation of the European Digital Identity Wallet's architecture and reference framework*. https://ssrn.com/abstract=5247521

Aggarwal, C. C., & Yu, P. S. (2008). *A general survey of privacy-preserving data mining models and algorithms, 11-52*. Springer US.

Ah-Fat, P., & Huth, M. (2019). Optimal accuracy-privacy trade-off for secure computations. *IEEE Transactions on Information Theory*, *65*(5), 3165–3182. https://doi.org/10.1109/tit.2018.2886458

Babel, M., & Sedlmeir, J. (2023). *Bringing data minimization to digital wallets at scale with general-purpose zero-knowledge proofs* (No. arXiv:2301.00823). arXiv. https://doi.org/10.48550/arXiv.2301.00823

Bamberger, Kenneth A.; Canetti, Ran; Goldwasser, Shafi; Wexler, Rebecca; Zimmerman, Evan J.; (2022). *Verification dilemmas in law and the promise of zero-knowledge proofs*. https://doi.org/10.15779/Z38P55DH8N

Bastian, P., Kraus, M., & Fischer, J. (2023). Konzepte für sichere wallets in dezentralen Identitätsökosystemen. *HMD Praxis der Wirtschaftsinformatik*, *60*(2), 381–404. https://doi.org/10.1365/s40702-023-00954-4

Baum, C., Chiang, J. H., David, B., & Frederiksen, T. K. (2023). SoK: Privacy-enhancing technologies in finance. *5th Conference on Advances in Finanical Technologies (AFT 2023)*, *282*, 12:1-12:30. https://doi.org/10.4230/LIPICS.AFT.2023.12

Bender, J., Kügler, D., Margraf, M., & Naumann, I. (2010). Privacy-friendly revocation management

without unique chip identifiers for the German national ID card. *Computer Fraud & Security*, *2010*(9), 14–17. https://doi.org/10.1016/s1361-3723(10)70122-6

Boneh, D., & Boyen, X. (2004). Short signatures without random oracles. In Cachin, C. & Camenisch, J.L. (Eds), *Advances in cryptology – EUROCRYPT 2004* (Vol. 3027, pp. 56–73). Springer. https://doi.org/10.1007/978-3-540-24676-3_4

Camenisch, J., & Lysyanskaya, A. (2001). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Pfitzmann, B. (Ed.), *Advances in cryptology – EUROCRYPT 2001* (Vol. 2045, pp. 93–118). Springer. https://doi.org/10.1007/3-540-44987-6_7

Camenisch, J., & Lysyanskaya, A. (2004). Signature schemes and anonymous credentials from bilinear maps. In Franklin, M. (Ed.), *Advances in cryptology—CRYPTO 2004* (Vol. 3152, pp. 56–72). Springer. https://doi.org/10.1007/978-3-540-28628-8_4

Council of the European Union & General Secretariat. (2022). *Glossary: Technical terms related to security features and to security documents in general (in alphabetical order)* (No. 12344/22; Public Register of Authentic Travel and Identity Documents Online (PRADO)). https://www.consilium.europa.eu/prado/en/prado-glossary/prado-glossary.pdf

Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., & Parno, B. (2016). Cinderella: Turning shabby X.509 certificates into elegant anonymous credentials with the magic of verifiable computation. *2016 IEEE Symposium on Security and Privacy (SP)*, 235–254. https://doi.org/10.1109/sp.2016.22

Deng, M., Wuyts, K., Scandariato, R., Preneel, B., & Joosen, W. (2011). A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, *16*(1), 3–32. https://doi.org/10.1007/s00766-010-0115-7

eIDAS eID Technical Subgroup. (2019). *eIDAS SAML AttributeProfile v1.2*. EU Login. https://ec.europa.eu/digital-building-blocks/wikis/digital-building-blocks/wikis/display/DIGITAL/eIDAS+eID+Profile

eIDAS Expert Group. (2021). *eIDAS expert group meeting: The toolbox process*. European Commission. https://webgate.ec.europa.eu/regdel/web/meetings/2409/documents/6690

ETSI. (2023). *Electronic signatures and infrastructures (ESI); Analysis of selective disclosure and zero-knowledge proofs applied to Electronic Attestation of Attributes* (Technical Report No. ETSI TR 119 476). https://www.etsi.org/deliver/etsi_tr/119400_119499/119476/01.01.01_60/tr_119476v010101p.pdf

European Commission. (2014). Regulation (EU)) No 910/2014 of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC. *Official Journal of the European Union*, *910/2014*.

European Commission. (2015). *Commission implementing Regulation (EU) 2015/ 1502—Of 8 September 2015—On setting out minimum technical specifications and procedures for assurance levels for electronic identification means pursuant to Article 8(3) of Regulation (EU) No 910/ 2014 of the European Parliament and of the Council on electronic identification and trust services for electronic transactions in the internal market*.

European Commission. (2021a). *Commission Staff Working Document accompanying the document report from the Commission to the European Parliament and the Council on the evaluation of Regulation (EU) No 910/2014 on electronic identification and trust services for electronic transactions in the internal market (eIDAS)*. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52021SC0130

European Commission. (2021b). *Recommendation—C(2021) 3968—A trusted and secure European e-ID*.

https://digital-strategy.ec.europa.eu/en/library/trusted-and-secure-european-e-id-recommendation

European Commission. (2021c). Report from the Commission to the European Parliament and the Council on the evaluation of Regulation (EU) No 910/2014 on electronic identification and trust services for electronic transactions in the internal market (eIDAS). *European Commission*. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52021DC0290

European Commission. (2024, January 10). *EU Digital Identity Wallet Toolbox process*. European Commission. https://digital-strategy.ec.europa.eu/en/policies/eudi-wallet-toolbox

European Commission. (2025, April). *The European Union Digital Identity (EUDI) Architecture Reference Model (ARF) v1.9.0*. https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/blob/85756d5b958c77053e4086f94c62c9155e07bf5f/docs/architecture-and-reference-framework-main.md

European Commission: Directorate General for Communications Networks, Content and Technology., PwC., DLA Piper., Garbasso, G., & Bianchini, D. et al. (2021). *Study to support the impact assessment for the revision of the eIDAS regulation: Final report*. Publications Office. https://data.europa.eu/doi/10.2759/671740

European Data Protection Board. (2020). *Guidelines 4/2019 on Article 25 Data Protection by Design and by Default (Version 2.0 10/2022)*. European Data Protection Board, EDPB. https://www.edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_201904_dataprotection_by_design_and_by_default_v2.0_en.pdf

European Data Protection Board. (2022). *Guidelines 01/2022 on Data Subject Rights—Right of Access (Version 2.1 03/)*. European Data Protection Board, EDPB. https://www.edpb.europa.eu/system/files/2023-04/edpb_guidelines_202201_data_subject_rights_access_v2_en.pdf

European Parliament & Council. (2021). *Proposal for a Regulation of the European Parliament and of the Council amending Regulation (EU) No 910/2014 as regards establishing a framework for a European Digital Identity*.

European Parliament & Council. (2024). *Agreed text for a Regulation of the European Parliament and of the Council amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework* (No. 2021/0136 (COD)). https://data.consilium.europa.eu/doc/document/PE-68-2023-COR-1/en/pdf

European Union Agency for Cybersecurity. (2022). *Data protection engineering: From theory to practice*. Publications Office. https://data.europa.eu/doi/10.2824/09079

Fernández, R. R. (2024). Evaluation of trust service and software product regimes for zero-knowledge proof development under eIDAS 2.0. *Computer Law & Security Review*, *53*, 105968. https://doi.org/10.1016/j.clsr.2024.105968

Finck, M., & Biega, A. J. (2021). Reviving purpose limitation and data minimisation in data-driven systems. *Technology and Regulation*, *2021*, 44–61. https://doi.org/10.71265/z7r0t122

Frigo, M. & Shelat. (2024). *Anonymous credentials from ECDSA*. *Cryptology ePrint Archive*. https://ia.cr/2024/2010

Fuster, G. G. (2014). *The emergence of personal data protection as a fundamental right of the EU, 16*. Springer Science & Business.

Garrido, G. M., Sedlmeir, J., Uludağ, Ö., Alaoui, I. S., Luckow, A., & Matthes, F. (2022). Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A

systematic literature review. *Journal of Network and Computer Applications*, *207*, 103465. https://doi.org/10.1016/j.jnca.2022.103465

Giannopoulou, A., & Wang, F. (2021). Self-sovereign identity. *Internet Policy Review*, *10*(2). https://doi.org/10.14763/2021.2.1550

Glöckler, J., Sedlmeir, J., Frank, M., & Fridgen, G. (2024). A systematic review of identity and access management requirements in enterprises and potential contributions of self-sovereign identity. *Business & Information Systems Engineering*, *66*(4), 421–440. https://doi.org/10.1007/s12599-023-00830-x

Government of British Columbia & Canada. (2023). *BC Digital Trust: Leveraging Hyperledger tools for digital trust*. LF Decentralized Trust. https://www.lfdecentralizedtrust.org/blog/bc-digital-trust-leveraging-hyperledger-tools-for-digital-trust

Grassi, P. A., Garcia, M. E., & Fenton, J. L. (2017). *Digital identity guidelines: Revision 3*. National Institute of Standards and Technology. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf

Heurix, J., Zimmermann, P., Neubauer, T., & Fenz, S. (2015). A taxonomy for privacy enhancing technologies. *Computers & Security*, *53*, 1–17. https://doi.org/10.1016/j.cose.2015.05.002

Housley, R., Polk, T., Ford, D. W. S., & Solo, D. (2002). *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL)* (No. RFC 3280). https://www.rfc-editor.org/info/rfc3280

Hundepool, A., Domingo-Ferrer, J., Franconi, L., Giessing, S., Nordholt, E. S., Spicer, K., & Wolf, P. P. (2012). *Statistical disclosure control*. John Wiley & Sons.

International Organization for Standardization & International Electrotechnical Commission. (2021). *Information security, cybersecurity and privacy protection—Requirements for attribute-based unlinkable entity authentication (ISO/IEC Standard No. 27551:2021)*.

Kaaniche, N., Laurent, M., & Belguith, S. (2020). Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey. *Journal of Network and Computer Applications*, *171*, 102807. https://doi.org/10.1016/j.jnca.2020.102807

Kakvi, S. A., Martin, K. M., Putman, C., & Quaglia, E. A. (2023). SoK: Anonymous credentials. In Günther, F. & Hesse, J. (Eds), *Security Standardisation Research* (Vol. 13895, pp. 129–151). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-30731-7_6

Kanevskaia, O., & Pałka, P. (2023). Introduction to the research handbook on law and technology. In *Research handbook on law and technology* (pp. 1–10). Edward Elgar Publishing. https://www.elgaronline.com/view/book/9781803921327/chapter1.xml

Looker, T., Kalos, V., Whitehead, A., & Lodder, M. (2023). *The BBS signature scheme (Internet Draft draft-irtf-cfrg-bbs-signatures-05*. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-irtf-cfrg-bbs-signatures

McConville, M., & Chui, W. H. (Eds). (2017). *Research methods for law* (Second). Edinburgh University Press.

McCorquodale, R. (1994). Self-determination: A human rights approach. *International & Comparative Law Quarterly*, *43*(4), 857–885.

Mejias, U. A., & Couldry, N. (2019). Datafication. *Internet Policy Review*, *8*(4). https://doi.org/10.1476

3/2019.4.1428

Menezes, A. J., Vanstone, S. A., & Oorschot, P. C. V. (1996a). Chapter 9—Hash functions and data integrity. In *Handbook of applied cryptography* (1st edn). CRC Press, Inc. https://cacr.uwaterloo.ca/hac/about/chap9.pdf

Menezes, A. J., Vanstone, S. A., & Oorschot, P. C. V. (1996b). Chapter 10—Identification and entity authentication. In *Handbook of applied cryptography* (1st edn). CRC Press, Inc. https://cacr.uwaterloo.ca/hac/about/chap10.pdf

Menezes, A. J., Vanstone, S. A., & Oorschot, P. C. V. (1996c). Chapter 11—Digital signatures. In *Handbook of applied cryptography* (1st edn). CRC Press, Inc. https://cacr.uwaterloo.ca/hac/about/chap11.pdf

Monteiro, A. P. L. (2023). Privacy at a crossroads. In *Research Handbook on Law and Technology* (pp. 214–221). Edward Elgar Publishing. https://www.elgaronline.com/view/book/9781803921327/chapter13.xml

Paquin, C., Policharla, G. V., & Zaverucha, G. (2024). *Crescent: Stronger privacy for existing credentials*. Cryptology ePrint Archive. https://ia.cr/2024/2013

Pfitzmann, A., Dresden, T., & Hansen, M. (2010). *A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management*.

Podgorelec, B., Alber, L., & Zefferer, T. (2022). What is a (Digital) Identity Wallet? A systematic literature review. *2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC)*, 809–818. https://doi.org/10.1109/compsac54236.2022.00131

Purtova, N. (2022). From knowing by name to targeting: The meaning of identification under the GDPR. *International Data Privacy Law*, *12*(3), 163–183. https://doi.org/10.1093/idpl/ipac013

Rieger, A., Roth, T., Sedlmeir, J., Weigl, L., & Fridgen, G. (2021). Not yet another digital identity. *Nature Human Behaviour*, *6*(1), 3–3. https://doi.org/10.1038/s41562-021-01243-0

Rieger, A., University of Arkansas, Roth, T., University of Arkansas, Sedlmeir, J., University of Luxembourg, Fridgen, G., University of Luxembourg, Young, A., & University of Arkansas. (2024). Organizational identity management policies. *Journal of the Association for Information Systems*, *25*(3), 522–527. https://doi.org/10.17705/1jais.00887

Rivest, R. L., Shamir, A., & Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, *21*(2), 120–126. https://doi.org/10.1145/359340.359342

Rosenberg, M., White, J., Garman, C., & Miers, I. (2023, May). Zk-creds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure. *2023 IEEE Symposium on Security and Privacy (SP)*. 2023 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA. https://doi.org/10.1109/sp46215.2023.10179430

Samarati, P. (2001). Protecting respondents identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, *13*(6), 1010–1027. https://doi.org/10.1109/69.971193

Schwarz, F., Do, K., Heide, G., Hanzlik, L., & Rossow, C. (2022). Feido: Recoverable FIDO2 tokens using electronic IDs. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2581–2594.

Sedlmeir, J., Huber, J., Barbereau, T., Weigl, L., & Roth, T. (2022). *Transition pathways towards design principles of Self-Sovereign Identity*. Forty-Third International Conference on Information Systems.

Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering*, *63*(5), 603–613. https://doi.org/10.1007/s12599-021-00722-y

Smits, J. M. (2015). What is legal doctrine? On the aims and methods of legal-dogmatic research. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2644088

Spensky, C., Stewart, J., Yerukhimovich, A., Shay, R., Trachtenberg, A., Housley, R., & Cunningham, R. K. (2016). SoK: Privacy on mobile devices - It's complicated. *Proceedings on Privacy Enhancing Technologies*, *2016*(3), 96–116. https://doi.org/10.1515/popets-2016-0018

Stevens, T., Elliott, J., Hoikkanen, A., Maghiros, I., & Lusoli, W. (2010). The State of the electronic identity market: Technologies, infrastructure, services and policies. *JRC Scientific and Technical Reports*. https://papers.ssrn.com/abstract=1708884

Torra, V. (2017). *Data privacy: Foundations, new developments and the big data challenge*. Springer International Publishing. https://doi.org/10.1007/978-3-319-57358-8

Tosoni, L. (2020). Article 4(25). Information society service. In *The EU General Data Protection Regulation (GDPR)* (pp. 292–302). Oxford University PressNew York. https://doi.org/10.1093/oso/9780198826491.003.0031

Trang, S., Kraemer, T., Trenz, M., & Weiger, W. H. (2025). Deeper down the rabbit hole: How technology conspiracy beliefse emerge and foster a conspiracy mindset. *Information Systems Research*, *36*(2), 709–735. https://doi.org/10.1287/isre.2022.0494

Tsakalakis, N. (2020). *Analysing the impact of the GDPR on eIDAS: Supporting effective data protection by design for cross-border electronic identification through unlinkability measures*.

Tsakalakis, N., Stalla-Bourdillon, S., & O'Hara, K. (2019). Data protection by design for cross-border electronic identification: Does the eIDAS interoperability framework need to be modernised? In Kosta, E., Pierson, J., Slamanig, D., Fischer-Hübner, S., & Krenn, S. (Eds), *Privacy and identity management. Fairness, accountability, and transparency in the age of big data* (Vol. 547, pp. 255–274). Springer International Publishing. https://link.springer.com/10.1007/978-3-030-16744-8_17

Van Dijck, J. (2014). Datafication, dataism and dataveillance: Big Data between scientific paradigm and ideology. *Surveillance & Society*, *12*(2), 197–208. https://doi.org/10.24908/ss.v12i2.4776

Van Hoecke, M. (2011). Legal doctrine: Which method(s) for what kind of discipline? In *Methodologies of legal research: Which kind of method for what kind of discipline?* (pp. 1–18). Hart Publishing. http://hdl.handle.net/1854/LU-1091776

W3C. (2022). *Verifiable credentials data model v1.1*. W3C. https://www.w3.org/TR/vc-data-model/

Wairimu, S., Iwaya, L. H., Fritsch, L., & Lindskog, S. (2024). On the evaluation of privacy impact assessment and privacy risk assessment methodologies: A systematic literature review. *IEEE Access*, *12*, 19625–19650. https://doi.org/10.1109/access.2024.3360864

Weigl, L., Amard, A., Codagnone, C., & Fridgen, G. (2022). The EU's digital identity policy: Tracing policy punctuations. *Proceedings of the 15th International Conference on Theory and Practice of Electronic Governance*, 74–81. https://doi.org/10.1145/3560107.3560121

Published by

Alexander von Humboldt
INSTITUTE FOR INTERNET
AND SOCIETY

in cooperation with

UNIVERSITY OF TARTU
Johan Skytte Institute of
Political Studies