ZKPs & BBS-Signatures for Digital Identities
Overview and Perspectives

SPRIND Event: EUDI Wallet Prototypes | 3.9.2024 | Berlin

Anja Lehmann

Hasso Plattner Institute | University of Potsdam

Motivation | EUDI & Secure User Authentication

- EUDI aims to enable **strong user authentication** & attribute attestation

  Strong security (unforgeability) through public-key cryptography

  e.g., signed credentials, key-based authentication

- eIDAS regulation specifies several **privacy requirements**:
  - □ Selective disclosure
  - □ Unlinkability: RP ↔ RP and RP ↔ IdP (Untraceability)
  - □ Unobservability
  - □ Pseudonymous authentication

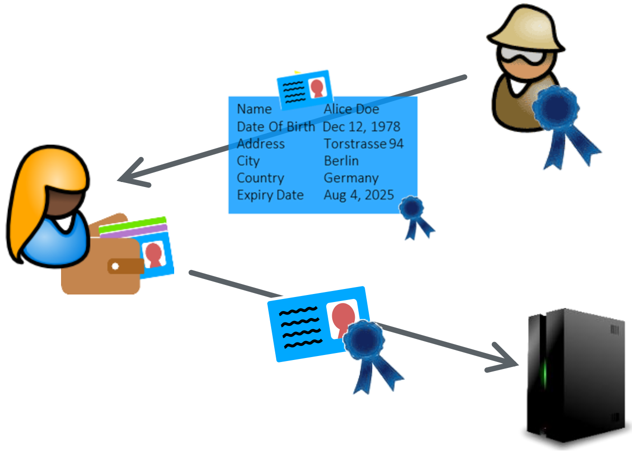

Name	Alice Doe
Date Of Birth	Dec 12, 1978
Address	Torstrasse 94
City	Berlin
Country	Germany
Expiry Date	Aug 4, 2025

- This talk:

  How to use modern cryptography to provide strongly secure & privacy-preserving EUDI

Motivation | Privacy Requirements in eIDAS

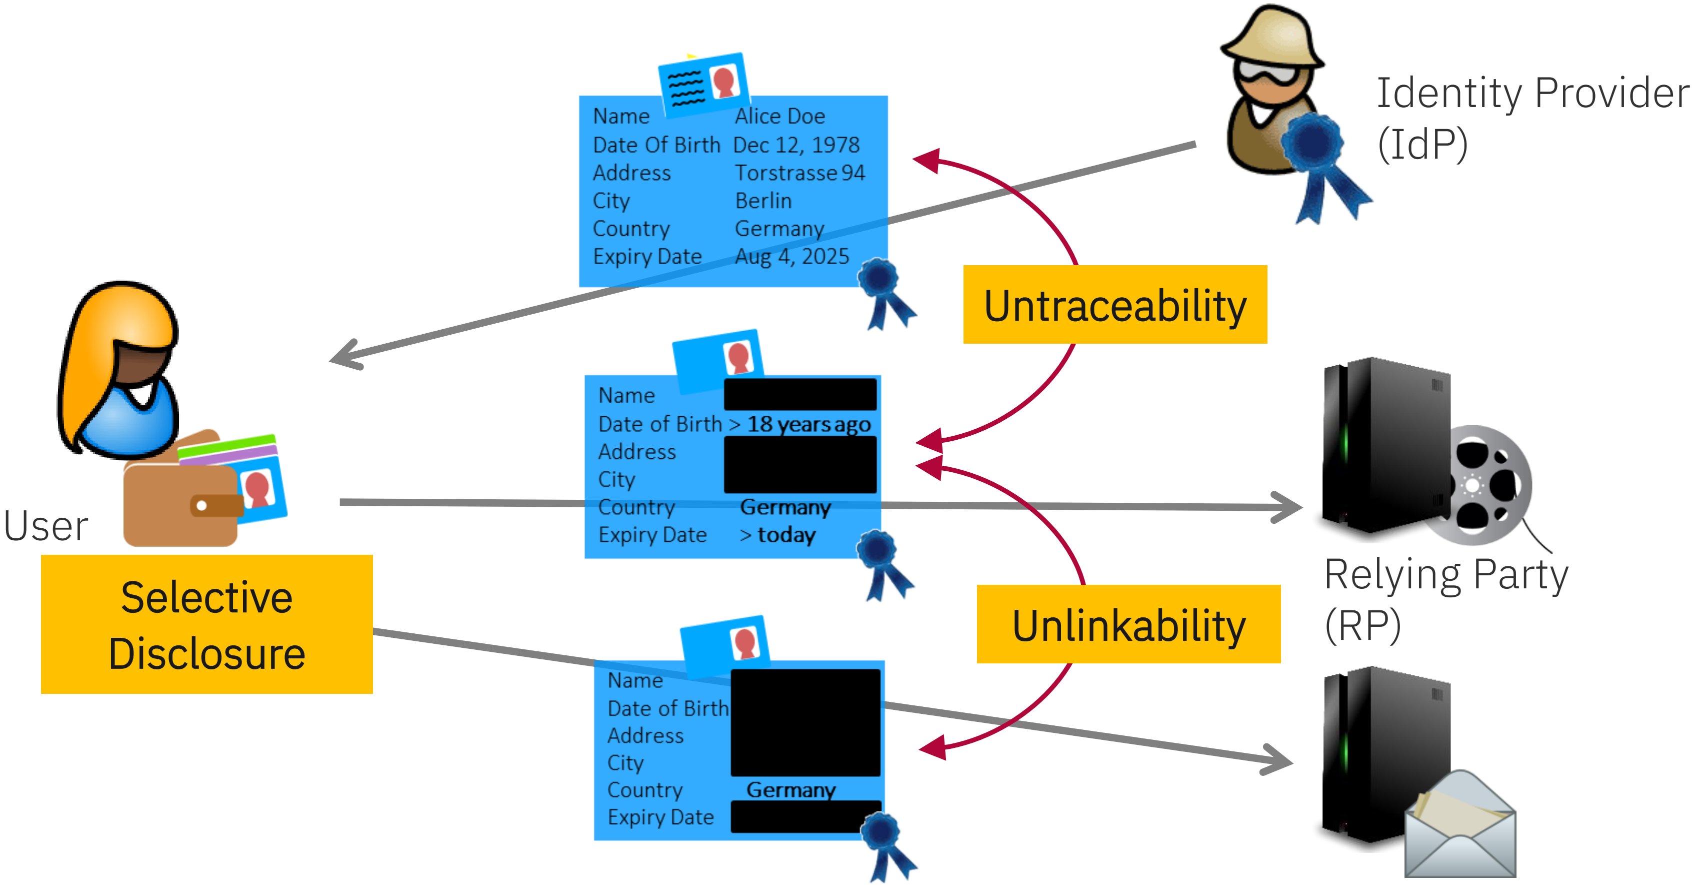

Identity Provider (IdP)

Name	Alice Doe
Date Of Birth	Dec 12, 1978
Address	Torstrasse 94
City	Berlin
Country	Germany
Expiry Date	Aug 4, 2025

User

Untraceability

Name	[redacted]
Date of Birth	> 18 years ago
Address	[redacted]
City	[redacted]
Country	Germany
Expiry Date	> today

Unlinkability

Relying Party (RP)

Selective Disclosure

Name	[redacted]
Date of Birth	[redacted]
Address	[redacted]
City	[redacted]
Country	Germany
Expiry Date	[redacted]

Classic Signatures & Limitations

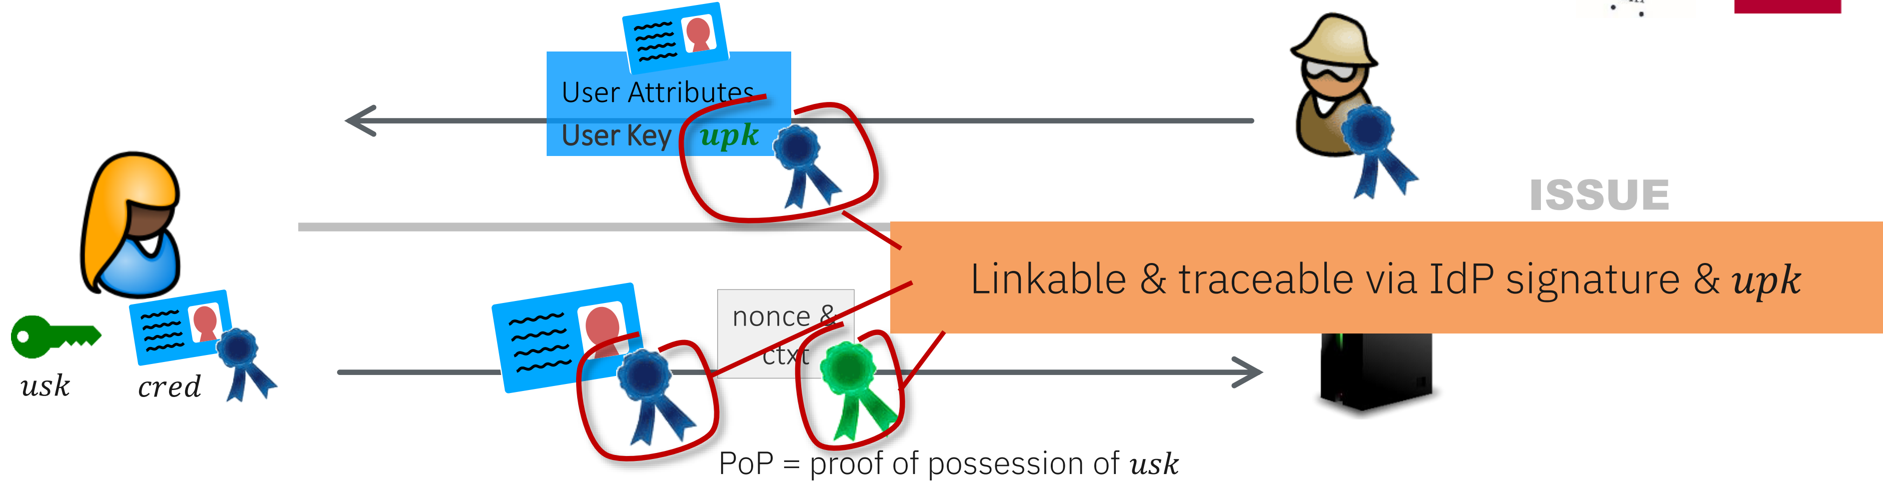

Linkable & traceable via IdP signature & upk

PoP = proof of possession of usk

Properties	Plain Signatures	„Patched“ Signatures	AnonCreds
Selective Disclosure	✗	Salted hashes ✓	✓
Unlinkability (RP ↔ RP)	✗	Batch issuance ✓	✓
Untraceability (RP ↔ IdP)	✗	Impossible ✗	✓

Anonymous Credentials | Privacy through ZKPs

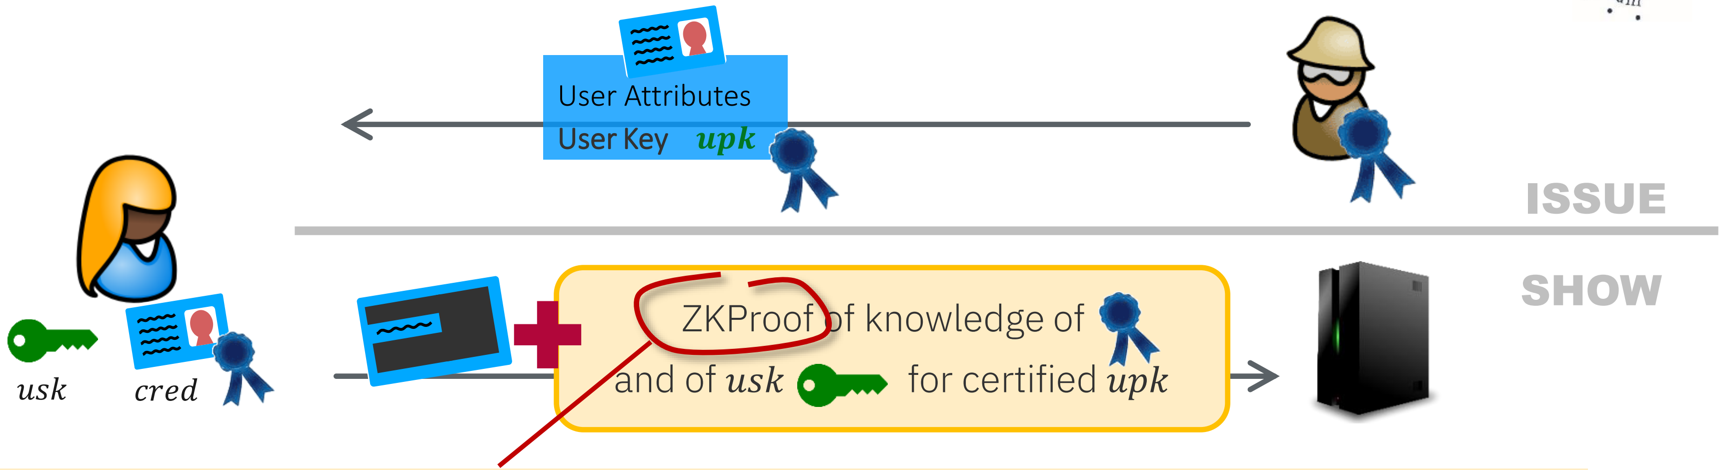

User Attributes

User Key upk

ISSUE

SHOW

usk $cred$

ZKProof of knowledge of

and of usk for certified upk

Zero-Knowledge Proof (ZKP) (idea &first schemes invented in 1985!)

Proof of a statement that reveals *nothing* beyond validity

Here: user proves she owns $cred$ from IdP on the revealed attributes & knows usk but reveals nothing about IdP‘s signature, her usk or upk (!)

User generates fresh ZKP from the <u>same</u> $cred$ and usk for every presentation

→ All presentations are unlinkable & untraceable due to ZK property

Anonymous Credentials | Signatures with ZKPs

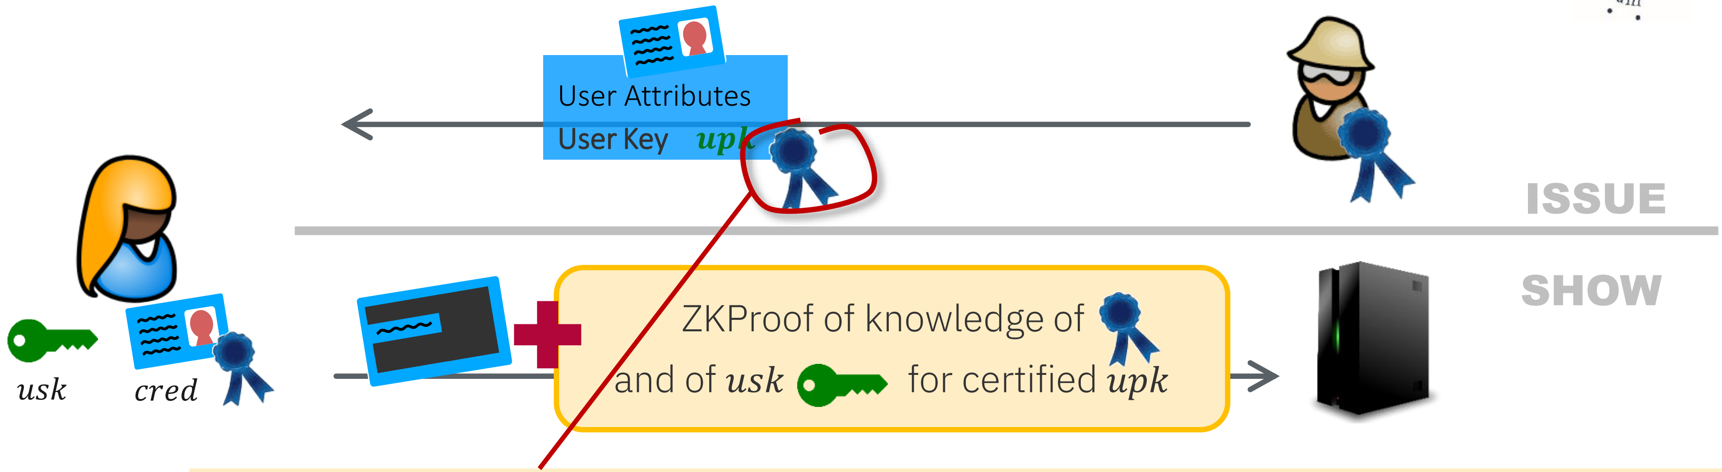

User Attributes

User Key upk

ISSUE

SHOW

usk

$cred$

ZKProof of knowledge of

and of usk for certified upk

Needs **signature scheme** (for IdP) **that allows for efficient ZKP of a signature**

Option 1 | Dedicated signature scheme with„build-in“ ZKP-capabilities

E.g., CL/BBS/PS-signatures → this talk

Option 2 | Use any signature scheme (e.g., ECDSA) & generic (circuit-based) ZKP

Generic, but less efficient & more complex → abhi‘s talk

BBS Signatures | Overview

- Core scheme proposed by Boneh, Boyen, and Shacham [BBS04]
  Extended & improved through series of works [CL04, ASM06, CDL16,TZ23]

- Mature (20 years!) and provably secure scheme (DL-related → unforgeability not quantum-safe
  privacy can hold perfectly)
  Requires pairing-friendly curve, e.g., BLS12-381

- Very simple, compact & with efficient ZKPs of signatures
  Signatures: 80bytes, ZKPs: 272 bytes
  Issuance: ~6ms, ZKProof: ~9ms, ZKVerify: ~20ms

- **Real-world adoption:** e.g., ISO, IETF Draft, W3C VC,
  Implemented in TPM2.0 DAA (2014), SGX EPID (2008)

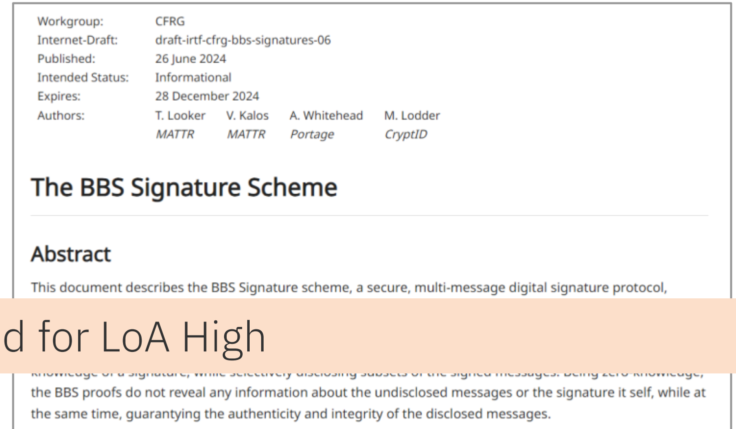

Workgroup:	CFRG
Internet-Draft:	draft-irtf-cfrg-bbs-signatures-06
Published:	26 June 2024
Intended Status:	Informational
Expires:	28 December 2024
Authors:	T. Looker *MATTR* V. Kalos *MATTR* A. Whitehead *Portage* M. Lodder *CryptID*

The BBS Signature Scheme

Abstract

This document describes the BBS Signature scheme, a secure, multi-message digital signature protocol,

knowledge of a signature, while selectively disclosing subsets of the signed messages. Being zero-knowledge, the BBS proofs do not reveal any information about the undisclosed messages or the signature it self, while at the same time, guarantying the authenticity and integrity of the disclosed messages.

- But: Secure Elements don‘t support BBS (yet) → required for LoA High

BBS Signatures | Device Binding for LoA High

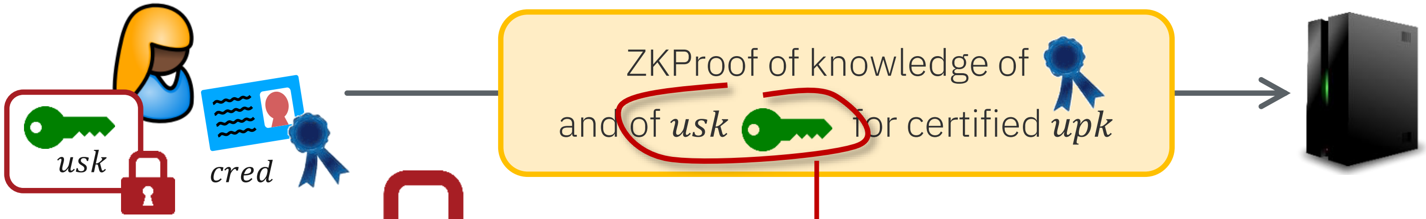

ZKProof of knowledge of and of usk for certified upk

```
Create()
  draw $sk \in \mathbb{Z}_q$, store $sk$
  output $pk \leftarrow g^{sk}$

Commit($P^r$)
  choose $r \in \mathbb{Z}_q$, store $(ctr, r)$
  $t \leftarrow P^r$
  output $(ctr, t)$

Hash($t, m$)
  output $c \leftarrow H(t, m)$

Sign($c, ctr$)
  get $(ctr, r)$
  output $s \leftarrow r + c \cdot sk$
```

High-level idea only

- Secure Element only needs to create proof of usk !

  For BBS: single exponentiation, no pairings

- E.g., generic interfaces in the spirit of TPM2.0 DAA-APIs
  - □ Support for BBS, other AnonCred signatures & extensions

    The original TPM APIs have security shortcomings

    See [CCD+17] for a revised version (only simple modification)

- Generic arithmetic operations might be supported by deployed Secure Elements already

→ Only curve needs update (via Secure Applet ?)

Comparison & Further Features

Properties	Salted Hashes + Batch Issuance	BBS + ZKPs
Selective Disclosure	✓	✓
Unlinkability	✓ security trade-off	✓
Untraceability	✗ impossible	✓
Device Binding	✓	not yet* / ✓ in 2 years (?)
Pseudonyms	?	✓
Deniability	✗	✓

Many extensions exist, e.g.:

- Privacy-preserving revocation
- Pseudonyms
- Designated verifier proofs
- Conditional disclosure
- Multi-credential proofs
- Threshold signing
- Blind signing

Helpful in Cloud HSM setting!

* BBS/ECDSA-bridge can be done now, if device-binding/LoA High only needed when user identifiable

Further Features | Pseudonyms and Deniability

BBS signatures support pseudonymous authentication

- Unlinkable pseudonyms derived from single usk & re-authentication requires usk
- Pseudonyms can additionally be RP-specific:

  Unique nym per RP, but unlinkable across RPs

  → Cloning detection & prevention of sybil attacks

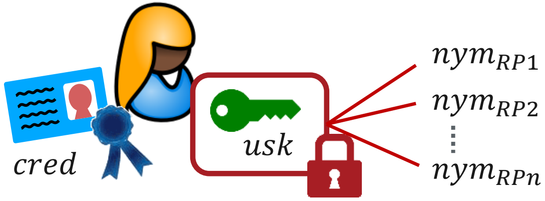

(also covered by APIs from previous slide)

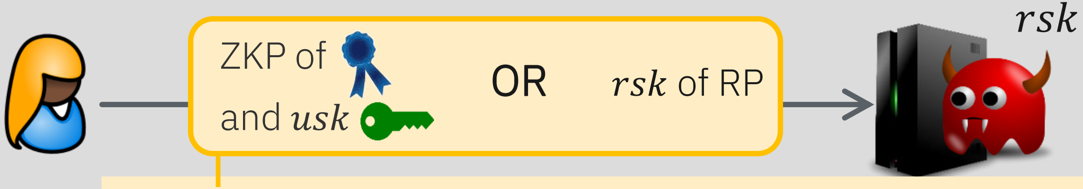

bound to ctx=„Only for RP-XY, must be deleted by 19.9.24, any further processing fined by 1Mio €“

ZKP-based presentation can be RP-bound & deniable:

- Validity of user attributes w.r.t. IdPs' key can only be verified with session-specific ctx

  Sticky context can disincentive malicious RPs from data sharing, makes leak traceable
- Designated verifier proof → can be generated by either user or RP → deniable yet signed

Summary

- Anonymous Credentials & ZKPs yield EUDI with privacy by design

  Fully satisfies all privacy requirements in eIDAS regulation

  Provide better security than current solutions → user has single credential & key

  Most efficient & mature instantiation: BBS

- Want to know more?

https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/issues/200

- Open questions:
  - □ What are exact functional & security requirements for full EUDI system?
  - □ Timeline for BBS-support on hardware?

    How to get certification for pairing-friendly curves, e.g., BLS12-381?
  - □ Quantum-safe or hybrid constructions?

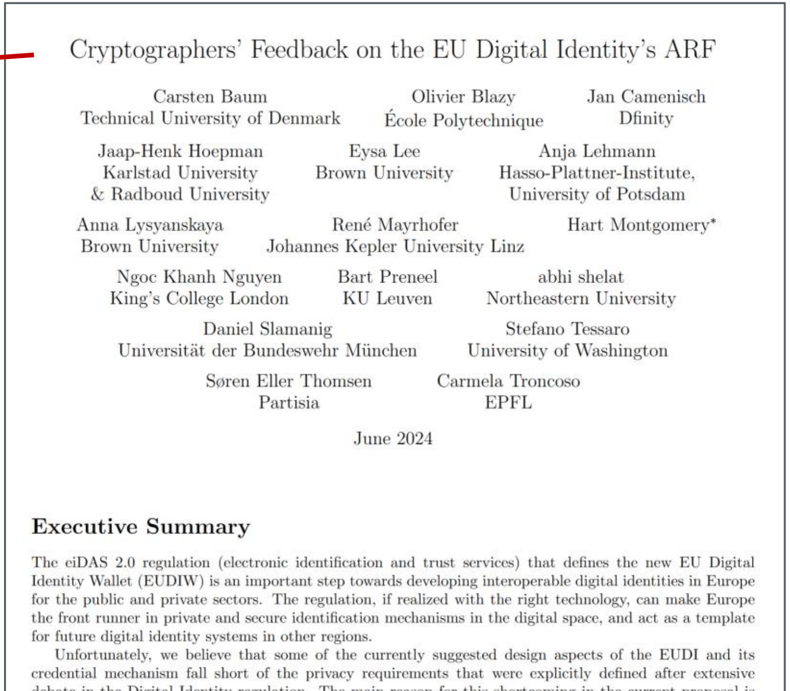

Cryptographers' Feedback on the EU Digital Identity's ARF

Carsten Baum
Technical University of Denmark

Olivier Blazy
École Polytechnique

Jan Camenisch
Dfinity

Jaap-Henk Hoepman
Karlstad University
& Radboud University

Eysa Lee
Brown University

Anja Lehmann
Hasso-Plattner-Institute,
University of Potsdam

Anna Lysyanskaya
Brown University

René Mayrhofer
Johannes Kepler University Linz

Hart Montgomery*

Ngoc Khanh Nguyen
King's College London

Bart Preneel
KU Leuven

abhi shelat
Northeastern University

Daniel Slamanig
Universität der Bundeswehr München

Stefano Tessaro
University of Washington

Søren Eller Thomsen
Partisia

Carmela Troncoso
EPFL

June 2024

Executive Summary

The eiDAS 2.0 regulation (electronic identification and trust services) that defines the new EU Digital Identity Wallet (EUDIW) is an important step towards developing interoperable digital identities in Europe for the public and private sectors. The regulation, if realized with the right technology, can make Europe the front runner in private and secure identification mechanisms in the digital space, and act as a template for future digital identity systems in other regions.

Unfortunately, we believe that some of the currently suggested design aspects of the EUDI and its credential mechanism fall short of the privacy requirements that were explicitly defined after extensive debate in the Digital Identity regulation. The main reason for this shortcoming in the current proposal is

References & Resources

References:

[JSI96] *Designated Verifier Proofs and Their Applications*. Jakobsson, Sako and Impagliazzo

[BBS04] *Short Group Signatures*. Boneh, Boyen, and Shacham

[CL04] *Signature Schemes and Anonymous Credentials from Bilinear Maps.* Camenisch and Lysyanskaya

[ASM06] *Constant-size Dynamic k-TAA*. Au, Susilo, and Mu

[CDL16] *Anonymous Attestation using the strong Diffie-Hellman Assumption Revisited.* Camenisch, Drijvers, and Lehmann

[CCD+17] *One TPM to Bind Them All: Fixing TPM 2.0 for Provably Secure Anonymous Attestation*. Camenisch, Chen, Drijvers, Lehmann, Novick, and Urian

[TZ23] *Revisiting BBS Signatures.* Tessaro and Zhu

[DG23] *RETRACT: Expressive Designated Verifier Anonymous Credentials*. Debes and Giannetsos

Further Resources

- □ The BBS Signature Scheme. IRTF Draft. Looker, Kalos, Whitehead, and Lodder. https://datatracker.ietf.org/doc/draft-irtf-cfrg-bbs-signatures/06/
- □ BBS per Verifier Linkability. Kalos and Bernstein https://www.ietf.org/archive/id/draft-vasilis-bbs-per-verifier-linkability-01.html
- □ W3C Data Integrity BBS Cryptosuites v1.0. Bernstein and Sporny. https://www.w3.org/TR/vc-di-bbs/
- □ ISO/IEC 20008-2:2013 Information technology — Security techniques — Anonymous digital signatures
- □ BBS+ Applications, Standardization, and a Bit of Theory. Bernstein and Kalos. https://csrc.nist.gov/csrc/media/presentations/2023/crclub-2023-10-18/images-media/20231018-crypto-club--greg-and-vasilis--slides--BBS.pdf
- □ Benchmark of the BBS+ signature scheme (2024) https://news.dyne.org/benchmark-of-the-bbs-signature-scheme-v06/

Further Features | Pseudonyms & User Binding

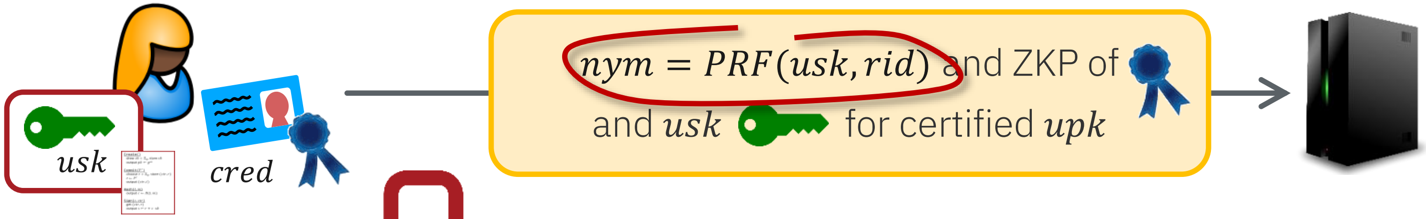

$nym = PRF(usk, rid)$ and ZKP of
and usk for certified upk

```
Create()
  draw $sk \in \mathbb{Z}_q$, store $sk$
  output $pk \leftarrow g^{sk}$

Commit($P^r$)
  choose $r \in \mathbb{Z}_q$, store $(ctr, r)$
  $t \leftarrow P^r$
  output $(ctr, t)$

Hash($t, m$)
  output $c \leftarrow H(t, m)$

Sign($c, ctr$)
  get $(ctr, r)$
  output $s \leftarrow r + c \cdot sk$
```

High-level idea only

- BBS Signatures support pseudonymous authentication
  - Pseudonym ~ privacy-preserving version of public key
  - Pseudonym is derived from certified usk & re-authentication requires usk – but user can derive many unlinkable nym
  - Pseudonyms can additionally be RP-specific:

    Unique nym per RP, but unlinkable pseudonyms across RPs

    E.g., ensuring that users can only have single account per RP

    → Cloning detection & prevention of sybil attacks

Further Features | Deniability

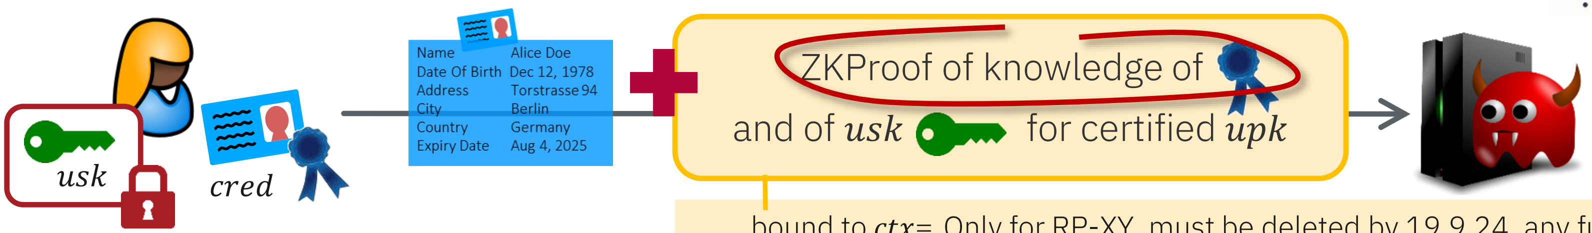

Name	Alice Doe
Date Of Birth	Dec 12, 1978
Address	Torstrasse 94
City	Berlin
Country	Germany
Expiry Date	Aug 4, 2025

ZKProof of knowledge of and of usk for certified upk

bound to ctx=„Only for RP-XY, must be deleted by 19.9.24, any further processing fined by 1Mio €“

RP-bound Presentation:

- User never sends the original IdP signature, only a ZKP of it
- User can bind every ZKP to a session-specific ctx

  Validity of user attributes w.r.t. IdPs’ key cannot be verified w/o ctx
- Sticky context can disincentive malicious RPs from data sharing, data leak is traceable

Deniable Presentation:

- ZKP-based presentation can be done as designated verifier proof [JSI96, DG23]
- ZKP proves that sender is either the user *or the designated RP*

  → will convince the targeted RP, but no one else → deniable yet signed data

Revised TPM2.0 Interfaces

TPM

```
TPM.Create()
    draw $sk \in Zq$, store sk
    output $pk \leftarrow g^{sk}$

TPM.Hash(t,m)
    output $c \leftarrow H(t,m)$

TPM.Commit(bsn)
    random nT, $hT \leftarrow H(nT)$
    choose $r \in Zq$, store (ctr, r, nT)
    $P \leftarrow H(bsn)$, $t \leftarrow P^r$
    output (ctr, t, hT)

TPM.Sign(c, ctr, nH)
    get (ctr, r, nT)
    $c' \leftarrow H(nH \oplus nT, c)$
    output nT, $s \leftarrow r + c' \cdot sk$
```

2017 IEEE Symposium on Security and Privacy

One TPM to Bind Them All: Fixing TPM 2.0 for Provably Secure Anonymous Attestation

Jan Camenisch*, Liqun Chen†, Manu Drijvers*‡, Anja Lehmann*, David Novick§, and Rainer Urian¶
*IBM Research – Zurich, †University of Surrey, ‡ETH Zurich, §Intel, ¶Infineon

***Abstract*—The Trusted Platform Module (TPM) is an international standard for a security chip that can be used for the management of cryptographic keys and for remote attestation. The specification of the most recent TPM 2.0 interfaces for direct anonymous attestation unfortunately has a number of severe shortcomings. First of all, they do not allow for security proofs (indeed, the published proofs are incorrect). Second, they provide a Diffie-Hellman oracle w.r.t. the secret key of the TPM, weakening the security and preventing forward anonymity of attestations. Fixes to these problems have been proposed, but they create new issues: they enable a fraudulent TPM to encode information into an attestation signature, which could be used to break anonymity or to leak the secret key. Furthermore, all proposed ways to remove the Diffie-Hellman oracle either strongly limit the functionality of the TPM or would require significant changes to the TPM 2.0 interfaces. In this paper we provide a better specification of the TPM 2.0 interfaces that addresses these problems and requires only minimal changes to the current TPM 2.0 commands. We then show how to use the revised interfaces to build *q*-SDH- and LRSW-based anonymous attestation schemes, and prove their security. We finally discuss how to obtain other schemes addressing different use cases such as key-binding for U-Prove and e-cash.**

1. Introduction

The amount of devices connected to the Internet grows rapidly and securing these devices and our electronic infrastructure becomes increasingly difficult, in particular because a large fraction of devices cannot be managed by security professional nor can they be protected by firewalls. One approach to achieve better security is to equip these devices with a root of trust, such as a Trusted Platform Module (TPM), a Trusted Execution Environment (TEE), and Software Guard Extensions (SGX), and then have that root of trust attest to the state of the device or to computations made. When doing such attestations, it is crucial that they be privacy-protecting. On the one hand, to protect the privacy of users of such devices, and on the other hand, to minimize the information available to attackers. Realizing this, the Trusted Computing Group (TCG) has developed a protocol called direct anonymous attestation (DAA) [1] and included it in their TPM 1.2 specification [2]. The protocol allows a device to authenticate as a genuine device (i.e., that it is certified by the manufacturer) and attest to messages without the different attestations being linkable to each other and has since been implemented in millions of chips.

Later, Brickell and Li [3] proposed a scheme called Enhanced-privacy ID (EPID) that is based on elliptic curves and adds *signature-based* revocation which is a revocation capability based on a previous signature of a platform. This scheme has become Intel's recommendation for attestation of a trusted system, has been incorporated in Intel chipsets and processors, and is recommended by Intel to serve as the industry standard for authentication in the Internet of Things. Being based on elliptic curves, EPID is much more efficient than the original RSA-based DAA scheme. Therefore, the TCG has revised the specification of the TPM and switched to elliptic curve-based attestation schemes [4], [5]. The design idea of this new specification is rather beautiful: the TPM only executes a simple core protocol that can be extended to build different attestation schemes. Essentially, the core protocol is a Schnorr proof of knowledge of a discrete logarithm [6], the discrete logarithm being the secret key stored and protected inside the TPM. Chen and Li [5] describe how to extend this proof of knowledge to DAA schemes, one based on the *q*-SDH assumption [14] and one based on the LRSW assumption [15]. The idea here is that the host in which the TPM is embedded extends the protocol messages output by the TPM into messages of the DAA protocol. They further show how to extend it to realize device-bound U-Prove [7], so that the U-Prove user secret key is the one stored inside the TPM.

Unfortunately, the core protocol as specified has severe shortcomings. First, the random oracle based security proof for attestation unforgeability by Chen and Li is flawed [8] and indeed it seems impossible to prove that a host cannot attest to a message without involving the TPM. Second, the core protocol can be abused as a Diffie-Hellman oracle w.r.t. the secret key *tsk* inside the TPM. It was shown that such an oracle weakens the security, as it leaks a lot of information about *tsk* [26]. Further, the presence of the oracle prevents forward anonymity, as an attacker compromising a host can identify the attestations stemming from this host.

These issues were all pointed out in the literature before and fixes have been proposed [8]–[10]. However, the proposed fixes either introduce new problems or are hard to realize. Xi et al. [8] propose a change to the TPM specification that allows one to prove the unforgeability of TPM-based attestations. This change introduces a subliminal channel though, i.e., a subverted TPM could now embed information into the values it produces and thereby into the final attestation. This covert channel could be used to break anonymity of the platform and its user, or to leak the secret key held in the TPM. The proposed fixes to remove the static Diffie-Hellman oracle [8]–[10] either require substantial changes to the TPM to the extend that they are not implementable, or restrict the functionality of the TPM too much, excluding some major DAA schemes from