

Study Report on Selective Disclosure

Overview and Classifications

**bg
in**

blockchain
governance
initiative
network

The Global Network for Blockchain Stakeholders™

This study is a work product of IAM, Key Management and Privacy Working Group of BGIN.

©2023 BGIN (Blockchain Governance Initiative Network) is registered as BN Association, a Japanese general association with its principal place of business at Shiba Building, 704, 4-7-6, Shiba, Minato-ku, Tokyo. All right reserved.

Introduction

IAM, Privacy and Key Management Study Group (IPWG) is chartered to provide guidance and good practice documents that describe Identity and Access management and its privacy considerations for access to crypto-currency exchange and its privacy considerations.

A user needs to provide its attributes when accessing a service, including crypto-currency exchange. The attribute needs to be certified as the service provider does not want a user to lie about its attribute. On the other hand, the user do not want to disclose all the attributes certified and want to selectively disclose its attribute according to the service provider's needs.

In this document, we provide guidance on various types of selective disclosure protocols and discuss its merits and demerits.

The technology described in this document was made available from contributions from various sources, including members of the BGIN and others. Although the BGIN has taken steps to help ensure that the technology is available for distribution, it takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any independent effort to identify any such rights. BGIN and the contributors to this document make no (and hereby expressly disclaim any) warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to this document, and the entire risk as to implementing this document is assumed by the implementer. The BGIN Intellectual Property Rights policy requires contributors to offer a patent promise not to assert certain patent claims against other contributors and against implementers. BGIN invites any interested party to bring to its attention any copyrights, patents, patent applications, or other proprietary rights that may cover technology that may be required to practice this document.

Table of contents

Introduction	2
Table of contents	3
1. Scope	3
2. Normative reference	3
3. Terms and definitions	3
4. Abbreviations and symbols	4
5. Purpose of Selective Disclosure	4
6. Online Authority	6
6.1 Alice talking to Online Authority	6
6.2 Bob talking to Online Authority	6
7. Offline Authority	7
7.1 Unlinkability in Selective Disclosure	7
7.2 Protocols without unlinkability	7
7.2.1 Naive Offline Protocol	7
7.2.2 Offline Protocol with hashed values	8
7.3 Offline Protocol with unlinkability	9
7.3.1 Naive Offline Protocol with unlinkability (one-show; one-time use)	9
7.3.2 Offline Protocol with unlinkability using zero-knowledge proofs	9
7.4 Predicate proofs using zero-knowledge proofs	10
8. Conclusions	11
Appendix A – Acknowledgement	11
A.1 Editor	11
A.2 Contributors	11
Appendix B – Informative reference	11

1. Scope

This report studies various ways to enable selective disclosure of attributes certified to an entity. The intended audience for this document includes developers, businesses, regulatory bodies, academic institutions, and any individual seeking to expand their understanding of selective disclosure and similar composability rights concepts.

2. Normative reference

This document has no normative reference.

3. Terms and definitions

This document uses the following terms as the shortcut for more complete wording provided as the definition. When the term appears within this document, it should be read as being replaced by the definition

3.1

selective disclosure

Process of revealing a subset of information where a full set of information has been certified by some authority, while ensuring the revealed information has been certified without disclosing the full set

3.2

online authority

authority who is available on demand when a certified information is needed

3.3

offline authority

authority who is not available once certified information has been issued

3.4

digital signature

cryptographic technique to ensure that an authority indeed certified the information

3.5

unlinkability

property of a protocol where it is impossible to relate two transactions originated from the same entity

3.6

zero-knowledge proof

ZKP

proof that asserts a proposition is true without revealing any other information

4. Abbreviations and symbols

In this document, the following abbreviations and symbols are used.

BGIN	Blockchain Governance Initiative Network
VC	Verifiable Credentials
DID	Decentralized Identifier
W3C	World Wide Web Consortium
IETF	Internet Engineering Task Force
SD-JWT	Selective Disclosure JWT
JWT	JSON Web Token
JSON	JavaScript Object Notation
ZKP	Zero-Knowledge Proof

5. Purpose of Selective Disclosure

This report studies various ways to enable selective disclosure of attributes certified to an entity.

Assume Alice has a set of attributes (a_1, a_2, \dots, a_n) that is certified by an authority I . Among those certified attributes, Alice wants to show to Bob say only (a_1, a_2, a_3). The goals of selective disclosure are:

Goal1) Bob can verify that indeed Alice has attributes (a_1, a_2, a_3) certified by the authority I .

Goal2) Bob does not learn other attributes (a_4, \dots, a_n).

This notion is important in the sense that while Alice has many attributes certified by an authority, she does not have to disclose all when the opponent, Bob, only wants to confirm a portion of her attributes. We assume here that Bob either "not need to know" or "don't want to know" the (a_4, \dots, a_n) attributes of Alice, respecting her privacy as her personal right.

As will be discussed in Section 6, if the authority is online, this property can be easily achieved when the authority only returns the subset of certified attributes. This is the situation where we have ID providers online, as in OpenID connect setting[1]. In the article[9], this case is described as 'just-in-time-issuance.'

The case becomes tricky when the authority is offline, that is, Alice has to do something with a previously signed certificate with all the attributes included. This technology is being discussed within the Verifiable Credential Working Group at the World Wide Web Consortium(W3C). The Working Group has published a W3C Recommendations on Verifiable Credentials Data Model v 1.1 (<https://www.w3.org/TR/vc-data-model-1.1/>) [2] where Verifiable Credentials are signed document where issuers certify a claim or attribute a subject has. The Data Model further defines the format for how to present the possession of Verifiable Credentials to a verifier. Such presentation considers the case where the subject only wants to disclose a portion of the attributes claimed in a Verifiable Credentials. The working group further discusses how to ensure correctness of such claims in the Verifiable Credential Data Integrity (<https://www.w3.org/TR/vc-data-integrity/>) [3] where a working draft is being published. In the draft, they discuss the importance of selective disclosure using the example of Drivers license as an example of Verifiable Credentials, and showing only a portion of attributes that appear on the license card. They also discuss the importance of unlinkability from the aspect of privacy. In IETF, there is a project on [Selective Disclosure for JWTs](#)[4]. This document classifies some approaches to achieve selective disclosure property and discusses its pros and cons.

While in this document we assume Alice is selectively disclosing her attribute to a single entity, Bob, it can be extended to the case where there are multiple recipients of her attributes and/ or the recipients are public like on blockchains[10]. There will be subtleties that need to be discussed in such cases, but it is out of scope of this document.

6. Online Authority

6.1 Alice talking to Online Authority

If Alice can contact the authority online, she can ask the authority I to issue a certificate with only attributes (a1, a2, a3). Namely, the certificate certifies 'Alice has (a1,a2,a3).' If Alice sends this certificate to Bob, (or the authority forwards the certificate to Bob on behalf of Alice) then it is clear that the above two goals are met.

Some disadvantages of this protocol are

Online-w-A-cons 1) Protocol is only possible when Authority is available. (Authority can deny issuance)

Online-w-A-cons 2) The authority learns which set of attributes Alice (and/or Bob) is interested in.

Online-w-A-cons 3) The authority learns when Alices discloses its attributes to Bob.

Some advantages of this protocol are

Online-w-A-pros 1) The authority can use any digital signature scheme to authenticate the attributes.

Online-w-A-pros 2) Alice can provide fresh attributes.

6.2 Bob talking to Online Authority

If Bob can contact the authority online, he can ask the authority if Alice has attributes (a1, a2, a3). Then it is clear that the above two goals are met.

Some disadvantages of this protocol are

Online-w-B-cons 1) Protocol is only possible when Authority is available. (Authority can deny responding to Bob)

Online-w-B-cons 2) The authority learns that Bob is interested in Alice's attribute, and in which set of attributes.

Online-w-B-cons 3) The authority learns when Alice discloses its attributes to Bob.

Online-w-B-cons 4) Alice does not know which attributes were asked.

Some advantages of this protocol are

Online-w-B-pros 1) The authority can use any digital signature scheme to authenticate the attributes.

Online-w-B-pros 2) Bob can confirm the freshness of the attributes

In order to avoid Bob learning too much of Alice's attributes, additional modification is possible.

- 1) Alice can tell the authority which attributes to give to Bob
- 2) Alice can provide a permission to Bob regarding which attributes Bob can ask to the authority, and authority checks the permission.
- 3) The authority can ask Alice if the authority can respond to Bob's request.

We note that this type selective disclosure is implemented in OpenIDconnect specifications[1].

7. Offline Authority

7.1 Unlinkability in Selective Disclosure

In this section, we the case where authority does not need to be online once it has issued Alice a certificate with all of her attributes. Again, the purpose of the selective disclosure remains the same:

Goal1) Bob can verify that indeed Alice has attributes (a1, a2, a3) certified by the authority I.

Goal2) Bob does not learn other attributes (a4,...,an).

We will consider an additional scenario, where Chris wants to verify that Alice has attributes (a4,a5,a6). The question is, if Bob and Chris collude, would they learn that Alice has attributes (a1, a2, a3, a4,a5, a6)? From a privacy point of view, this is not preferable. Therefore, if the protocol does not allow this, we say the protocol has ‘unlinkability’. However, as we see below, there are some costs to achieve unlinkability. We begin by looking into simpler protocols without unlinkability property.

7.2 Protocols without unlinkability

7.2.1 Naive Offline Protocol

If Alice wants to perform selective disclosure without help of online authority, she can obtain certificates from the authority in advance. One naive solution is she is going to obtain n certificates, saying ‘Alice has attribute ai’ for $i=1,..n$.

When disclosing to Bob, she can choose the set of certificates that she wants to disclose to Bob.

It is clear that the above two goals are met.

Some advantages of this protocol are

NaiveOffline-pros 1) Alice can selectively disclose without help of Authority

NaiveOffline-pros 2) The authority does not learn which set of attributes Alice (and/or Bob) is interested in.

NaiveOffline-pros 3) The authority does not learn when Alice discloses its attributes to Bob.

NaiveOffline-pros 4) The authority can use any digital signature scheme to authenticate the attributes.

Some disadvantages of this protocol are

NaiveOffline-cons 1) Alice needs to manage multiple certificates.

NaiveOffline-cons 2) Bob cannot confirm the freshness of the attributes.

NaiveOffline-cons 3) Alice cannot be anonymized, because a common identifier to link multiple certificates is necessary.

The third cons prevents the protocol from achieving unlinkability. The reason why we need a common identifier in each certificate is as follows: Assume Alice has only (a1, a2) but there is Dave who has attribute (a3). We want to avoid Alice and Dave to collude (perhaps by sharing certificates) and deceive Bob that Alice has attributes (a1, a2, a3). In order to do this, we need a mechanism to show that indeed certificates for a1, a2, and a3 were issued to the same person. However, this mechanism, if not carefully designed, tends to serve as a common identifier and thus the protocol fails to achieve unlinkability. An example of a carefully designed mechanism avoiding linkability will be discussed in 7.3.

We note that with online authority, the authority can certify the attributes of Alice by using a suitable pseudonym for the session. Thus achieving unlinkability is possible, assuming we have a trustworthy authority.

7.2.2 Offline Protocol with hashed values

In contrast to Naive Offline Protocol, the protocols in this category issue only one certificate to each person. The authority is going to sign on $(\text{hash}(a_1), \text{hash}(a_2), \dots, \text{hash}(a_m))$ where hash represents a cryptographically secure hash function, and (a_1, \dots, a_m) are Alice's attributes. That is, one can easily compute $\text{hash}(a_1)$ from a_1 , but it is difficult to guess what a_1 is from $\text{hash}(a_1)$. The certificate is composed of a signature on $(\text{hash}(a_1), \text{hash}(a_2), \dots, \text{hash}(a_m))$. Note that the certificate itself does not reveal a_1, \dots, a_m directly.

When Alice is revealing her attributes (a_1, a_2, a_3) , she is going to send (a_1, a_2, a_3) and signed $(\text{hash}(a_1), \text{hash}(a_2), \dots, \text{hash}(a_m))$, which is the certificate. Bob is going to verify the signature of the authority, and that $\text{hash}(a_1)$, $\text{hash}(a_2)$, and $\text{hash}(a_3)$ are equal to hash of a_1 , a_2 , and a_3 respectively.

The protocol cannot achieve unlinkability because the same message $(\text{hash}(a_1), \text{hash}(a_2), \dots, \text{hash}(a_m))$ and the signature of the authority will be shown to both Bob and Chris. If they collude, they will know that they are talking with the same person and are linkable.

Moreover, this scheme as it is has further issues. Because of the nature of deterministic properties of hash functions (that is, the same a_1 always gives the same $\text{hash}(a_1)$) and that the hash function is public, one may be able to guess several candidates of attributes and see if they are included in Alice's certificate. In order to avoid this, instead of ordinary deterministic hash function, we can use what is called salted hash function. This hash function takes two inputs, a salt and a message and returns a hashed value.

So the authority will issue a certificate on $(\text{hash}(s_1, a_1), \text{hash}(s_2, a_2), \dots, \text{hash}(s_m, a_m))$ where s_1, s_2, \dots, s_m are independently chosen salt values for each attribute. The authority will give the values of salts to Alice. When Alice is revealing her attributes (a_1, a_2, a_3) , she is going to send (a_1, a_2, a_3) (s_1, s_2, s_3) together with signed $(\text{hash}(s_1, a_1), \text{hash}(s_2, a_2), \dots, \text{hash}(s_m, a_m))$. Bob is going to verify the signature of the issuer, and that $\text{hash}(s_1, a_1)$, $\text{hash}(s_2, a_2)$, and $\text{hash}(s_3, a_3)$ are correctly computed from (a_1, a_2, a_3) and (s_1, s_2, s_3) .

This protocol is now being standardized at IETF as Selective Disclosure JWT(SD-JWT)[4].

Details can be found at:

<https://datatracker.ietf.org/doc/draft-fett-oauth-selective-disclosure-jwt/>

Some advantages of this protocol are

SD-JWT-pros 1) Alice can selectively disclose without help of Authority

SD-JWT-pros 2) The authority does not learn which set of attributes Alice (and/or Bob) is interested in.

SD-JWT-pros 3) The authority does not learn when Alice discloses its attributes to Bob.

SD-JWT-pros 4) The authority can use any digital signature scheme to authenticate the attributes.

Some disadvantages of this protocol are

SD-JWT-cons 1) Bob cannot confirm the freshness of the attributes.

SD-JWT-cons 2) The protocol is linkable, because it is showing the same credential to all verifiers.

7.3 Offline Protocol with unlinkability

7.3.1 Naive Offline Protocol with unlinkability (one-show; one-time use)

Alice has all certificates issued for all possible subsets of attributes. This is similar to Naive Offline Protocol described in 7.2.1 except that Alice will be issued certificates for all possible sets of attributes, in order to avoid unlinkability. As can be easily guessed, cons for this protocol is Alice has to maintain a large number of certificates, with all possible combination of attributes.

7.3.2 Offline Protocol with unlinkability using zero-knowledge proofs

In this protocol, Alice will be issued one certificate including all n attributes of Alice, digitally signed by the authority beforehand.

In an ordinary digital signature scheme, when a signature is performed over a message containing (a_1, \dots, a_n) , the exact message is necessary to verify the signature. Therefore, even if Alice wants to show a portion of her attributes, she needs to disclose all other attributes in order for Bob to verify the signature of the authority.

However, in some special signatures, Alice can convert the signature of the authority into a 'proof' so that Bob can verify that there is authority's signature on (a_1, a_2, a_3) issued to Alice without learning other attributes, meeting the aforesaid goals. Examples of such special signatures are CL signatures[5] and BBS+ Signatures[6][7][8][9].

More specifically, given a signature of (a_1, \dots, a_n) , public key of the authority, and (a_i, a_j, a_k) , Alice generate a proof that the i -th, j -th and k -th message in the signed message $(a_1, a_2, a_3, \dots, a_n)$ are a_i, a_j and a_k respectively.

Bob can verify the proof using the public key of the authority and (a_i, a_j, a_k) , that indeed Alice knows a signature to a message where its the i -th, j -th message are a_i, a_j and a_k . The proof is called zero-knowledge proof if it does not contain any other information except that the statement is true, including the other a_i 's and even the original signature.

This protocol using zero-knowledge proofs achieves unlinkability, as it does not disclose any information to link two proofs.

Some advantages of this protocol are

ZKIP-pros 1) Alice can selectively disclose without help of Authority

ZKIP-pros 2) The authority does not learn which set of attributes Alice (and/or Bob) is interested in.

ZKIP-pros 3) The authority does not learn when Alice discloses its attributes to Bob.

ZKIP-pros 4) The protocol is unlinkable

Some disadvantages of this protocol are

ZKIP-cons 1) Bob cannot confirm the freshness of the attributes.

ZKIP-cons 2) The authority should use a special digital signature scheme to authenticate the attributes.

This type of proofs is considered in W3C Working Draft on Verifiable Credentials Data Integrity[3].

7.4 Predicate proofs using zero-knowledge proofs

A selective disclosure protocol discloses a subset of attributes. Using zero-knowledge proofs, we can disclose not a subset, but only a property of an attribute without disclosing the attribute itself.

Assume Alice has a certificate signed by the authority saying her age is 24. By designing a suitable zero-knowledge proof, Alice can prove that she has a certificate signed by the authority of her age, and that the age is over 20 without disclosing her age itself. That is, she can prove that her age attribute a_1 satisfies the following predicate ' $a_1 > 20$ ' without disclosing a_1 .

This type of proofs is also considered in W3C Working Draft on Verifiable Credentials Data Integrity[3].

8. Conclusions

We have discussed various ways to perform selective disclosure of attributes given to a subject. The protocol can be easily implemented when there is an online authority who is disclosing the attributes on behalf of the subject. If there is no such online authority, or there is a risk to have such authority to intermediate the disclosure, we can consider having an offline authority who signs on a credential that confirms the attributes of the subject. We would need special tricks or special digital signature schemes to enable selective disclosure with an offline authority.

In our model, we assumed that Bob asks Alice to reveal a subset of her attributes written on credentials. This should be under

In either case, selective disclosure is an important feature to enhance privacy of the subject.

Appendix A – Acknowledgement

(Informative)

A.1 Editor

- Kazue Sako (Waseda University)

A.2 Contributors

- Joseph Beverley (Soulbis)
- Michi Kakebayashi (University of California, Berkeley)
- Ken Katayama (Nomura Research Institute, Ltd.)
- Shin'ichiro Matsuo (Georgetown University)
- Nat Sakimura (OpenID Foundation)
- Yuji Suga (Internet Initiative of Japan)
- Dan Yamamoto (Internet Initiative of Japan)
- Tomonori Yuyama (Georgetown University)

Appendix B – Informative reference

(Informative)

[1] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros and C. Mortimore, "**OpenID Connect Core 1.0 including errata set 1**," November 2014.

[2] W3C Recommendations: Verifiable Credentials Data Model v 1.1
<https://www.w3.org/TR/vc-data-model-1.1/>

[3] W3C Working Draft: Verifiable Credentials Data Integrity v1
<https://www.w3.org/TR/vc-data-integrity/>

[4] D. Fett and K. Yasuda, "Selective Disclosure JWT (Draft x)," <https://datatracker.ietf.org/doc/draft-ietf-oauth-selective-disclosure-jwt/>

[5] J. Camenisch and A. Lysyanskaya, "A signature scheme with efficient protocols," in SCN 02, ser. LNCS, S. Cimato, C. Galdi, and G. Persiano, Eds., vol. 2576. Springer, Heidelberg, Sep. 2003, pp. 268–289.

[6] M. H. Au, W. Susilo and Y. Mu, "Constant-size dynamic k-TAA," in SCN 06, ser. LNCS, R. D. Prisco and M. Yung, Eds., vol. 4116. 61 Springer, Heidelberg, Sep. 2006, pp. 111–125.

[7] J. Camenisch, M. Drijvers and A. Lehmann, "Anonymous attestation using the strong diffie hellman assumption revisited," in International Conference on Trust and Trustworthy Computing. Springer, 2016, pp. 1–20.

[8] T. Looker, V. Kalos, A. Whitehead and M. Lodder, "The BBS Signature Scheme"
<https://identity.foundation/bbs-signature/draft-irtf-cfrg-bbs-signatures.html>

[9] MATTR Concepts 'Selective Disclosure'
<https://learn.mattr.global/docs/concepts/selective-disclosure>

[10] Polygon ID
<https://polygon.technology/blog/introducing-polygon-id-zero-knowledge-own-your-identity-for-web3>

