did:self A registry-less DID method

Nikos Fotiou*, George C. Polyzos and Vasilios A. Siris

ExcID, 113 62, Athens, Greece

Abstract

We introduce `did:self`, a Decentralized Identifier (DID) method that does not depend on any trusted registry for storing the corresponding DID documents. Information for authenticating a `did:self` *subject* can be disseminated using any means and without making any security assumption about the delivery method. `did:self` is lightweight, it allows controlled delegation, it offers increased security and privacy, and it can be used for identifying people, content, as well as IoT devices. Furthermore, DID documents in `did:self` can be implicit, allowing re-construction of DID documents based on other authentication material, such as JSON Web Tokens and X.509 certificates.

Keywords

Decentralized Identifiers, JSON Web Tokens, X.509 certificates

1. Introduction

In the age of "surveillance capitalism"[1] personal data is massively collected by large corporations that prioritize profit over privacy. At the same time, the alarming number of large scale data breaches, such as the Equifax data breach,[2] and incidents of massive disclosure of personal data, such as the Facebook-Cambridge Analytica scandal,[3] is a clear indication of the lack of proper security and privacy safeguards. To address these concerns, a significant number of efforts advocating decentralization has sprung up. W3C's *Decentralized Identifiers* (DIDs) is a notable example of such an effort in the domain of identification. W3C's DID specifications pursue a portable digital identity that does not depend on a centralized authority and can never be taken away by third-parties [1]. Decentralized identification systems are also beneficial to service providers, many of which collect personally identifying and private information to support their personalized services. A decentralized, user-centric approach for handling identity will alleviate service providers from the burden of collecting sensitive data [2] for which they are liable.

A DID is a new type of self-administered, globally unique identifier, which is *resolvable* and *cryptographically verifiable* [1]. In particular, a DID is resolved to a *DID document* that contains (among other things) cryptographic material that can be used for verifying DID ownership (e.g., a public key). DID documents are usually stored and maintained by a *registry* [3]. Registries may be administrated by a single entity or a consortium of trusted entities (e.g., a Web server or a permissioned blockchain), or they may be provided by public systems such as a public, permissionless blockchain (e.g., Ethereum). Registries of the former type require some level of trust and may introduce security and privacy risks. On the other hand, the interaction with blockchain registries may involve significant computational overhead (which may not be tolerable by an IoT device), or even monetary cost. Moreover, in general, existing registries lack interoperability and in many cases the information associated with a DID cannot be transferred from one registry to another. With these in mind, we propose `did:self`, a new DID *method*, which is compatible with W3C specifications and at the same time does not depend on any type of registry.

---

[1]https://en.wikipedia.org/wiki/Surveillance_capitalism
[2]https://en.wikipedia.org/wiki/2017_Equifax_data_breach
[3]https://en.wikipedia.org/wiki/Facebook–Cambridge_Analytica_data_scandal

Identifiers in `did:self` are created using the thumbprint of a public key. Accordingly, DID documents are protected by a *proof* generated by the DID controller; this proof can be validated with the public key that was used for creating the corresponding identifier. Therefore, given a `did:self` DID and the corresponding DID document and proof, any entity can verify their binding, as well as the correctness and the integrity of the DID document. Similarly to self-signed digital certificates, proof verification does not require any auxiliary information. A unique feature of `did:self` is that it supports secure DID sharing and controlled delegation. By leveraging this property, earlier versions of our DID method were used for implementing content-centric security for Information-Centric Networking architectures [4, 5], for securing IoT group communication [6], as well as for improving the security of the Inter-Planetary File System (IPFS) [7].

In this paper, we introduce an update to the `did:self` specification that brings the following improvements:

- DIDs can have a (human-readable) suffix that provides clearer semantics about the entities that share the same identifier.
- We generalize the concept of proof and introduce implicit DID documents. To this end, we propose two methods for generating DID documents using JSON Web Tokens and X.509 certificates. This enables interoperability with existing authentication and authorization systems.

The remainder of this paper is organized as follows. In Section 2 we introduce DIDs. In Section 3 we present the design of our solution. In Section 4 we discuss related work in this area. We conclude our paper in Section 5.

2. Background: W3C Decentralized Identifiers

Decentralized Identifiers (DIDs), defined by W3C, are a new type of globally unique identifier designed to enable individuals and organizations to generate their own identifiers using systems they trust. A DID architecture can be regarded as a key-value lookup system, where the key is the Decentralized Identifier (DID) and the value is a DID *document*. DID specifications allow multiple DID *methods* to co-exist; the main differentiating factor among existing DID methods is how the resolution from a DID to the corresponding DID document is implemented. A DID is expressed as a URI composed of three parts separated by ":"; the first part is the word "did", the second part is the DID method name, and the third part is a method specific identifier.

A DID document contains "properties", serialized according to a particular syntax, that may include (among other things) *verification methods* (e.g., public keys) and *verification relationships* that express a relationship between a verification method and the DID *subject*. Examples of verification relationships are *authentication*, which means that a verification method is used for authenticating the DID subject, and *assertion*, which means that a verification method is used for verifying assertions, e.g., digital signatures, made by the DID subject.

Listing 1 illustrates an example of a DID document. The DID of the subject is `did:self:iQ9PsBKOH1nLT9FyhsUGvXyKoW00yqm_-_rVa3W7Cl0/device1` (line 2). In lines 3-10 a verification method is defined, which is a public key. This public key has an identifier (i.e., "#key1"), and it is encoded using the type "JsonWebKey2020" [3], which is an encoding based on RFC 7517 JSON Web Keys. Next, two verification relationship are defined, i.e., "authentication" and "assertion" which include the identifier of the defined key. Therefore, in this example "#key1" can be used for authenticating the DID subject and for verifying digital signatures it generates.

Listing 1: A sample DID document

```
1   {
2     "id": "did:self:iQ9PsBKOH1nLT9FyhsUGvXyKoW00yqm_-_rVa3W7Cl0/device1",
3     "verificationMethod": [{
4       "id": "#key1",
5       "type": "JsonWebKey2020",
6       "publicKeyJwk": {
7         "kty": "EC",
8         "crv": "P-256",
9         "x": "YOGmYaMKzwTFytWHN2hGC-2VpPqGqj_sDSckB2IvCgI",
10        "y": "7iWuiXQlLXvROjdMA2WNHhGz0jxu6u41n83YupNteo"
11      }
12      }
13    ],
14    "authentication": ["#key1"],
15    "assertion": ["#key1"],
16  }
```

3. The did:self DID method

A `did:self` DID[4] is generated by a *controller* who creates a public/private key pair. The `did:self` identifier is the thumbprint of the JSON Web Key (JWK) representation of the controller's public key (as defined by RFC 7638 [8]), optionally followed by a suffix. Therefore a `did:self` identifier has the following form:

$$did{:}self{:}{<}thumbprint{>} \ / \ {<}suffix{>}$$

The corresponding DID document is protected by a *proof* generated using the controller's private key. A controller is not necessarily the DID *holder*. Our method supports multiple holders per DID, each of which may use a public/private key pair it controls for implementing a *verification relationship* (e.g., *authentication*). As a motivating example consider the case of an IoT system, where the IoT system owner is the controller and each IoT device can be a holder. IoT devices can use a key pair they control to implement a verification relationship. In this case, for each IoT device a DID document is created that defines a verification relationship based on the public key that the device controls (see for example Listing 1). The proof for this document is generated by the controller. In this use case, for each IoT device a suffix can be added to the DID. Additionally, two IoT devices may share the same DID, allowing, for example, device rotation in case of failure or representation of virtual devices composed of multiple physical ones (e.g., "smart home"). This use case is illustrated in Figure 1, which shows two drones that share the same DID, but use a different authentication key.

A straightforward approach for creating a proof is by generating a JSON Web Signature (JWS) [9] using the private key of the controller's key pair. In that case, the JWS header must include the `jwk` field, which is used as follows:

- `jwk` The JWK representation of the controller's public key, which can be used for verifying the proof. The thumbprint of this key must match the thumbprint included in the `did:self` identifier.

The payload of the JWS is the DID document. Therefore, validating such a proof is a simple three-steps process: (i) extract the public key from the JWS header, (ii) verify that the thumbprint of the extracted

---

[4]In the following we use the terms `did:self` DID and `did:self` identifier interchangeably

public key is equal to the thumbprint included in the `did:self` identifier, and (iii) validate the JWS signature using as input the extracted public key and the DID document.

In the next sections we discuss how `did:self` can be used with JSON Web Tokens (JWT) and X.509 certificates.

3.1. Representing DID documents as JSON Web Tokens

A `did:self` DID document that includes a single verification method can be represented as a JSON Web Token (JWT) [10]. The following table specifies how the payload claims of such a JWT are created:

Claim	Value
iss	The `did:self` identifier without any suffix
sub	The suffix of the `did:self` identifier
cnf	A public key controlled by the DID holder represented as a JWK

Table 1
The claims of a JWT representing the DID document of a did:self DID.

Listing 2 includes the payload of a JWT representing the DID document of Listing 1. Such a JWT is signed by the controller using its private key. The JWS header must include either the `jwk` claim whose value will be the JWK representation of the controller's public key or the `x5c` field representing a certificate chain (see next section). Again, the JWT verification is a three-steps process: (i) extract the public key from the JWS header, (ii) verify that the thumbprint of the extracted public key is equal to the thumbprint included in the `iss` claim, and (iii) validate the JWS signature using as input the extracted public key and the JWT.
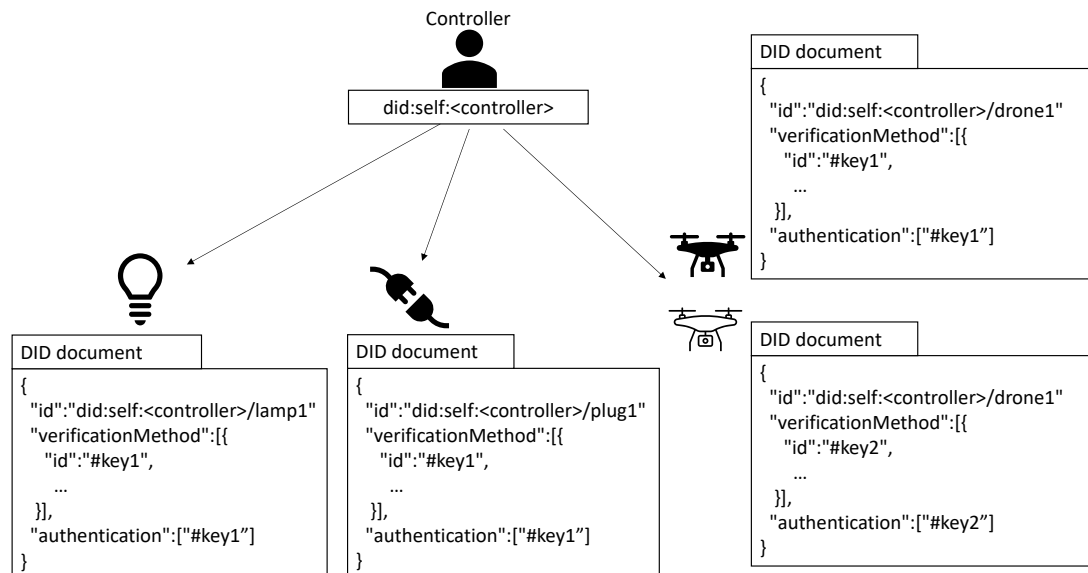
Listing 2: A did:self DID document represented as a JWT

```
1    {
2      "iss": "did:self:iQ9PsBKOH1nLT9FyhsUGvXyKoW00yqm_-_rVa3W7Cl0",
3      "sub": "device1",
4      "cnf": {
5        "jwk": {
6          "kty": "EC",
7          "crv": "P-256",
8          "x": "YOGmYaMKzwTFytWHN2hGC-2VpPqGqj_sDSckB2IvCgI",
9          "y": "7iWuiXQlLXvROjdMA2WNHhGz0jxu6u41n83YupNteo"
10        }
11      }
12   }
```

Using this approach `did:self` can be straightforwardly used with protocols such as Demonstrating Proof-of-Possession at the Application Layer (DPoP) of OAuth 2.0 [11].

3.2. Representing DID documents as X.509 certificates

Similarly, a `did:self` DID document that includes a single verification method can be represented as an X.509 certificate chain. The first certificate of the chain is a self-signed certificate that includes the controller's public key and the `did:self` identifier (without any suffix) in the `Subject Alternative Name` field. This certificate is signed using the controller's private key. The last certificate of the chain includes the holder's public key and the `did:self` identifier (with a suffix if necessary) in the `Subject Alternative Name` field. The chain may include intermediate certificates issued by the DID controller. The intermediate certificates give greater flexibility to the controller, which can have multiple public

Figure 1: An IoT use case where did:self is applied. It can be observed that all IoT devices share a DID with the same prefix. Furthermore, the two drones have the same DID.

keys associated with its DID with varying level of security. The validation of the certificate chain follows the standard X.509 validation rules. Additionally, a verifier should verify that the thumbprint included in the `Subject Alternative Name` is the same as the thumbprint of the JWK representation of the public key included in the first certificate of the chain.

4. Related work

Although DID document registries management, they create lock-in conditions, decrease DID self-sovereignty, and add communication overhead, which in some scenarios is intolerable. For this reason, various efforts propose registry-less DID systems. NaCL:did [12] and did:key [13] are two such systems that use (the hash of) a public key as the DID. Similarly, ethr-did [14] uses Ethereum addresses as DIDs. The main drawback of these systems is that they support a single verification method, which is the public key from which the DID has been derived. For this reason, they cannot not support real DID documents; instead, in these methods DID documents are implied. In contrast, did:self supports DID documents, which may include multiple verification methods.

did:peer [15] is a registry-less DID system, which allows DID documents. DID document modifications are encoded in authenticatable *deltas*. A DID document can be modified by multiple entities, and for this reason did:peer relies on a *conflict-free replicated data type* for consolidating all deltas and eventually creating the final DID document. This approach has some side-effects: it makes DID document reconstruction complex and it poses restrictions on the format of the document. Our approach is much simpler, since the only additional operation required for retrieving a DID document is the validation of at most two digital signatures. Furthermore, our approach does not pose any restriction to the fields a DID document may contain, therefore, our solution is sustainable and future proof; this is very important considering that DID specifications have not yet been finalized.

did:web [16] is a registry-based DID system that allows users to use any Web server as a registry. With did:web, DIDs are bound to the URL of the Web server and DID document resolution is trivially implemented using HTTPS. did:web offers a higher degree of sovereignty compared to many existing systems, but still relies on the security of HTTPS for securely distributing DID documents. On contrary, did:self does not need any secure communication channel for distributing DID documents.

ION [17] is also a registry-based DID system that has been adopted by Microsoft.[5] In ION, DID

---

[5]https://techcommunity.microsoft.com/t5/identity-standards-blog/ion-we-have-liftoff/ba-p/1441555

documents are stored in a "Content-addressable storage network" (e.g. the InterPlanetary File System (IPFS) [18]). Modifications to the DID documents are stored in "delta" files, which are eventually "anchored" in a decentralized sequencing oracle (ION uses Bitcoin for this purpose). Therefore, DID documents are reconstructed by retrieving the appropriate "delta" files, following the "pointers" stored in the sequence oracle. Compared to did:web, ION relies on a more distributed trust model, but at the same time adds complexity and in some cases monetary cost.

The Key Event Receipt Infrastructure (KERI) proposal [19] considers as the primary root-of-trust self-certifying identifiers that are strongly bound at issuance to a cryptographic signing (public, private) key-pair, similar to our proposal. Key rotation is performed with a signed transfer statement. The KERI proposal considers an indirect mode where a set of witnesses (replicas) store and forward key events to any requester/validator. Although the existence of witnesses enhances fault tolerance and availability, it entails additional requirements and complexity related to witness designation and policy, witness consensus, and handling of out-of-order events. did:self avoids such complexity by focusing on use cases where it is sufficient to disseminate the corresponding DID document through any method, even between untrusted parties.

Compared to all these DID methods, did:self is the only DID method that allows multiple DID documents per DID, as well as controlled DID document delegation. As we discuss in Section 4, these properties enable intriguing security and privacy features.

Due to the security and privacy characteristics of the DID paradigm, many research efforts investigate the potential of using DIDs to improve the security and privacy of emerging technologies. Chadwick et al. [20] propose the integration of DIDs and Verifiable Credentials with the FIDO Universal Authentication Framework (UAF) in order to provide safer and more private online account management. DIDs are also considered for improving the security and privacy of IoT systems (e.g., [21],[22],[23],[24]). Davie et al. [25] propose a four-layer architectural stack, based on DIDs, to establish trust between peers over the Internet and other digital networks. Lagutin et al. [26] investigate the application of DIDs and Verifiable Credentials in the OAuth 2.0 authorization process. Finally, Munoz [27] discusses the advantages of integrating eIDAS and DIDs. Our work is complementary to these approaches: being standard compliant, did:self could be the DID technology that these proposed systems may eventually use.

5. Conclusion

In this paper we presented the `did:self` DID method and introduced two new extensions: identifier suffixes and implicit DID documents. Our DID method gives controllers absolute control over their identifiers and removes any dependency on middlemen, hence increasing self-sovereignty and privacy. At the same time, `did:self` allows multiple DID documents per DID enabling delegation and sharing. Our solution is not only compatible with W3C specifications, but is also easier to integrate in existing systems, by removing the complexity of interacting with a DID registry.

By leveraging implicit DID documents, we presented two constructions for generating a `did:self` DID document: one using JSON Web Tokens and another based on X.509 certificates. This enables interoperability with existing authentication and authorization systems. At the same time however, it prevents taking advantage of the full potential of DIDs, since legacy DID documents can carry more information beyond a single public key.

Although the lack of a registry and the support for multiple DID documents per DID enable novel applications and offer significant security and privacy advantages, they create challenges related to revocation. Currently, `did:self` replies mainly on the expiration time to handle revoked verification material, but other more efficient mechanisms are also being explored.

References

[1] W3C Credentials Community Group, A primer for decentralized identifiers, 2019. Https://w3c-

ccg.github.io/did-primer/.

[2] K. C. Toth, A. Anderson-Priddy, Self-sovereign digital identity: A paradigm shift for identity, IEEE Security Privacy 17 (2019) 17–27.

[3] O. Steel, M. Sporny, DID specification registries, 2020. Https://w3c.github.io/did-spec-registries/.

[4] N. Fotiou, Y. Thomas, V. Siris, G. Xylomenos, G. Polyzos, Securing Named Data Networking routing using decentralized identifiers, in: SARNET-21: Semantic Addressing and Routing for Future Networks, IEEE International Conference on High Performance Switching and Routing (HPSR) 2021 workshop, IEEE, 2021.

[5] N. Fotiou, Y. Thomas, V. A. Siris, G. Xylomenos, G. C. Polyzos, Self-verifiable content using decentralized identifiers, Computer Networks 230 (2023) 109799.

[6] N. Fotiou, V. A. Siris, G. Xylomenos, G. C. Polyzos, IoT group membership management using decentralized identifiers and verifiable credentials, Future Internet 14 (2022).

[7] N. Fotiou, V. A. Siris, G. C. Polyzos, Enabling self-verifiable mutable content items in IPFS using decentralized identifiers, in: 2021 IFIP Networking Conference (IFIP Networking), 2021, pp. 1–6.

[8] M. Jones, N. Sakimura, JSON Web Key (JWK) Thumbprint, RFC 7638, IETF, 2015. URL: https://www.rfc-editor.org/rfc/rfc7638.txt.

[9] M. Jones, J. Bradley, N. Sakimura, JSON Web Signature (JWS), RFC 7515, IETF, 2015. URL: https://tools.ietf.org/html/rfc7515.

[10] M. Jones, J. Bradley, N. Sakimura, JSON Web Token (JWT), RFC 7519, IETF, 2015. URL: https://tools.ietf.org/html/rfc7515.

[11] D. Fett, et al., OAuth 2.0 Demonstrating Proof-of-Possession at the Application Layer (DPoP), RFC 9449, IETF, 2023. URL: https://datatracker.ietf.org/doc/rfc9449/.

[12] uPort Project, NaCL DID resolver and manager, 2020. Https://github.com/uport-project/nacl-did.

[13] D. Longley, D. Zagidulin, M. Sporny, The did:key method, 2020. Https://w3c-ccg.github.io/did-method-key/.

[14] uPort Project, Ethr-DID Library, 2021. Https://github.com/uport-project/ethr-did.

[15] D. Deventer, et al., Peer DID method specification, 2020. Https://openssi.github.io/peer-did-method-spec.

[16] D. Zagidulin, O. Terbu, A. Guy, did:web Method Specification, 2020. Https://w3c-ccg.github.io/did-method-web/.

[17] Identity Foundation, did:ion Home Page, 2022. Https://identity.foundation/ion/.

[18] J. Benet, IPFS - content addressed, versioned, P2P file system, 2014. URL: http://arxiv.org/abs/1407.3561.

[19] S. Smith, Key Event Receipt Infrastructure (KERI), 2019. URL: https://arxiv.org/abs/1907.02143.

[20] D. W. Chadwick, R. Laborde, A. Oglaza, R. Venant, S. Wazan, M. Nijjar, Improved identity management with verifiable credentials and fido, IEEE Communications Standards Magazine 3 (2019) 14–20.

[21] R. Ansey, J. Kempf, O. Berzin, C. Xi, I. Sheikh, Gnomon: Decentralized identifiers for securing 5G IoT device registration and software update, in: 2019 IEEE Globecom Workshops (GC Wkshps), 2019, pp. 1–6.

[22] G. Fedrecheski, J. M. Rabaey, L. C. P. Costa, P. C. Calcina Ccori, W. T. Pereira, M. K. Zuffo, Self-sovereign identity for IoT environments: A perspective, in: 2020 Global Internet of Things Summit (GIoTS), 2020, pp. 1–6.

[23] Y. Kortesniemi, D. Lagutin, T. Elo, N. Fotiou, Improving the privacy of IoT with decentralised identifiers (dids), Journal of Computer Networks and Communications 2019 (????).

[24] S. Terzi, C. Savvaidis, K. Votis, D. Tzovaras, I. Stamelos, Securing emission data of smart vehicles with blockchain and self-sovereign identities, in: 2020 IEEE International Conference on Blockchain (Blockchain), 2020, pp. 462–469.

[25] M. Davie, D. Gisolfi, D. Hardman, J. Jordan, D. O'Donnell, D. Reed, The trust over ip stack, IEEE Communications Standards Magazine 3 (2019) 46–51.

[26] D. Lagutin, Y. Kortesniemi, N. Fotiou, V. A. Siris, Enabling decentralised identifiers and verifiable credentials for constrained IoT devices using oauth-based delegation, in: Proceedings of the

Workshop on Decentralized IoT Systems and Security (DISS 2019), in Conjunction with the NDSS Symposium, San Diego, CA, USA, volume 24, 2019.

[27] C. Munoz, SSI and eIDAS: a vision on how they are connected, 2019. Https://ec.europa.eu/futurium/en/system/files/ged/eidas_supported_ssi_may_2019_0.pdf.