Security and Privacy Assessment Criteria for Decentralized Identifiers

July 20, 2021

Prepared for:
U.S. Department of Homeland Security
Science and Technology Directorate
Office of Industry Partnerships

Prepared by:
Deborah Shands
SRI International

CRATES

Cybersecurity Research and Technical Execution Support

This work was sponsored by the U.S. Department of Homeland Security (DHS) Science and Technology Directorate under Contract No. HSHQDC-16-C-00034. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DHS and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DHS.

Approved for Public Release

Contents

1. Introduction

Unique identifiers for people, organizations, things, and abstract entities are increasingly important in most contexts, but they are essential in the digital realm. The World Wide Web Consortium's (W3C) Decentralized Identifiers (DIDs) [4] define digital identities that are decoupled from centralized registries, identity providers, and certificate authorities. DID Methods are the mechanisms by which these DIDs and associated DID documents are created, resolved, updated, and deactivated. A DID Method is implemented in software that executes on hardware; both the software and hardware are operated (i.e., maintained and administered) by people who may be part of an organization (i.e., the DID Method operator).

DID Method implementations and DID Method operators may offer systems and services with very distinct security and privacy characteristics. This document recommends an approach to developing requirements for the security and privacy of DID Methods to enable assessment of different implementations and services.

In the following sections, we identify use cases of entities that require identifiers; discuss how various operational environments impact the security, functional, and assurance characteristics of DID Methods; discuss security threat models that focus on identity fraud; highlight DID Method-specific security and privacy objectives and requirements; and discuss approaches to producing a comprehensive set of requirements.

2. Use Cases for Decentralized Identifiers

DIDs are intended to identify a great range of entities including persons, organizations, conceptual things, and physical things. This section identifies some specific instances in each of these categories that inform the development of objectives and requirements. These use cases help illustrate important security or privacy issues that must be considered.

DIDs must be able to identify the following types of subjects:

2.1. Person

Legally capable adult	The most commonly discussed use case.
Child	This use case illustrates dependency and the need for a DID controller that is not the subject of the DID.
Dependent adult	This use case illustrates dependency and the need for a capability to address alternative DID controllers and potential disputes.

2.2. Organization

Business	This use case illustrates the need to establish the real-world analogue of the digital entity through some form of real-world investigation (e.g., D-U-N-S numbers [2]).

2.3. Conceptual Item

Bank account	This use case illustrates a business construction with regulatory obligations.
Postal address	(A mailable address [3]) This use case illustrates an item that is regulated by a governmental or commercial authority.

2.4. Physical Object

Traceable item	This use case illustrates an item with an associated history of properties. Examples include a food (farm or factory inspection of origination site); a chemical (purity or safety characteristics; manufacturing facility); imported items (bills of lading, customs documents).
Piece of equipment	This use case illustrates an object that is identical to many other items except for a unique sequence of characters attached to the item (e.g., serial number for an auto part).
Original painting [1]	This use case illustrates a unique object for which authenticity and provenance may be of concern.

3. DID Ecosystem and Operational Environment

DID Methods operate within an environment that includes infrastructure technologies (e.g., computing platforms, storage, and communication resources) that are configured, maintained, and administered by personnel, often in an organizational context. Implementations of DID Methods will use security-critical functionality provided by the operational environment, including, for example, persistent data storage and backups, logging and audit functions, time services, and network communications. Administrators of DID Methods will rely on system and network administration and provisioning tools, intrusion detection systems, network and host-based firewalls, and hardware and software mechanisms for generating and managing cryptographic keys.

The security and privacy functional and assurance characteristics of a DID Method implementation depend on the capabilities and quality of the operational environment in which it executes. The security and privacy protections provided by a DID Method operator depend on physical and environmental protections, contingency planning and incident response, personnel security and training, and maintenance and administration of systems.

Assumptions about the security- and privacy-critical functionality provided by the operational environment depend on the implementation of a specific DID Method. For example, the `did:web` method, which relies on web hosts and Domain Name System (DNS) queries, and the `did:v1` method, which relies on the permissionless Veres One Blockchain public ledger, depend on very different operational environments; therefore, these DID Methods make different assumptions about many of the platforms and services provided by those environments. For each security and privacy requirement that a DID Method is obliged to meet, explicit assumptions about the operational environment should be presented. For example, a requirement for the DID Method to generate audit records for security-relevant events might assume that the operational environment provides mechanisms for generating portions of the audit data, storing audit records for the required period of time, and reviewing audit data. How these assumptions are met by organizations and systems within the DNS and web infrastructure will likely be quite different from how they are met by Veres One participants.

4. Security Threat Models

Historically, most security attacks against identification schemes have aimed at perpetrating fraud. Common examples include generating fake identity credentials for individuals (false passports, student, or state IDs) or unauthorized modifications to existing identity credentials (altered birth certificates). Often, falsified identity credentials establish a root that enables the acquisition of additional identifiers or credentials, such as a social security number or business tax identification number. Effective fraudulent credentials allow attackers to gain access to items with real-world value, such as bank accounts, deeds of trust to real estate, powers of attorney over individuals, and business assets.

A variety of entities may be involved in identity credential creation and management. Any one or more of these entities could participate in the commission of fraud, using either well-known or subtle weaknesses in the identity credential creation and management system. Common errors by well-meaning personnel or flawed processes and systems can sometimes be exploited repeatedly before attacks are detected.

The following sections identify some of the entities in the DID management ecosystem that could play a role in identity credential fraud.

4.1. Attacker is the Implementor of a DID Method

In this case, the attacker is the author of the DID Method software that expresses the logical semantics of the DID Method. By intentionally implementing software that undermines the security objectives of the DID management system, this attacker can fundamentally compromise its integrity, enabling a great variety of fraud or other harms to the DID subject or relying parties. For example, DID Method software could intentionally generate non-unique DIDs, allow DID controllers that are not identified in DID documents, or enable violation of DID integrity guarantees. Some of these could be accomplished through the use of flawed implementations of cryptographic algorithms, and others by incorrect implementation of access-control logic.

4.2. Attacker Operates the DID Method Service

In this case, the attacker operates a DID Method service, running and administering DID Method software (and, potentially, the hardware) to enable others to create and manage DIDs. Multiple DID Method service operators may run the same DID Method software, distinguishing their services based on price, reputation, or other factors. A DID Method service operator that administers the deployment and operation of DID Method software can host the service on unreliable or malicious infrastructure (e.g., storage systems that do not meet persistence obligations, weak entropy generators that undermine pseudorandom number generation for cryptographic operations), or establish weak or malicious deployment or operational configurations (e.g., configuring infrastructure integrity protection systems to mishandle DID Method operation metadata logging by dropping, modifying, or otherwise compromising system audit logs).

4.3. Attacker is the DID Subject

Often, the DID subject has strong incentives to commit fraud to gain inappropriate credentials. In the W3C model of DIDs, the DID subject is passive; however, a DID subject is often a DID controller for their own DID document. As an active participant in the DID creation and maintenance process for a DID document, a DID subject may manipulate the contents of the DID document, but this does not directly enable them to link inappropriate verifiable credentials to the DID. A compromised DID subject could misdirect relying parties to inappropriate service endpoints, which could later enable a broader attack.

4.4. Attacker is a DID Controller

In the W3C model, a DID controller has the authority to modify any element of the DID document. Most elements of a DID document are security-relevant, so modification by an attacker would have serious security consequences. For example, a DID controller might modify a DID document to add additional DID controllers, change or eliminate methods for verification or authentication, or redirect a service endpoint to trick relying parties into using an attacker-controlled service. In addition, a DID controller could deactivate (revoke) a DID, which could effectively disable any linked credentials. The real-world impact to a DID subject of malicious deactivation of a DID might be comparable to the impact of sudden passport revocation on a traveler.

4.5. Attacker Operates a Verifiable Data Registry

A DID Method relies on a verifiable data registry to record DIDs and the data necessary to produce DID documents. Like other essential system infrastructure required by most applications (operating system, network, etc.), a verifiable data registry must function correctly to support fundamental security requirements for DID management. Failing to record DIDs or DID documents, returning garbled files that prevent reconstruction of DID documents, deleting DID Method metadata (e.g., activity logs, version stamps) are a few ways by which problems in a verifiable data registry could impact DID Method operations. Many of these problems could occur due to benign system failures, but a malicious operator of a verifiable data registry or author of verifiable data registry software could undermine DID security for a DID Method on which others rely. For example, the selective mishandling of DID document metadata, like the `deactivated` property (which indicates revocation of the DID), could trick relying parties into using revoked identifiers.

4.6. Attacker Operates a DID Resolver

A DID resolver implements a “lookup” or “read” function for a relying party; given a DID, the resolver finds and returns the corresponding DID document and, potentially, some DID document metadata and DID resolution metadata. This process could be compromised via DID resolver software, an operator of a DID resolution service, or client-side software that transforms legitimate resolution results to present false information to the requestor. The `id` property of a DID document must match its corresponding DID. Assuming that integrity protection was applied to the DID document (e.g., a digital signature), the relying party could check that a matching and unadulterated DID document was provided by the resolver. Realistically, however, relying parties (especially human users) will depend on layers of software to correctly perform system, network, and cryptographic operations. The deeper the layers of convenience software that mediate between the DID resolver and human user, the more opportunities for exploitable flaws or malicious intervention.

5. Security Objectives and Requirements for DID Methods

While a complete enumeration of the security objectives and requirements for DID Methods is beyond the scope of this document, this section highlights some of the security objectives and requirements that are important to address identity fraud threats (discussed in Section 4) against DID Methods.

One model for developing and expressing a comprehensive list of security objectives and requirements for a security-critical information-technology product is the protection profile used in security assessments under the internationally recognized Common Criteria for Information Technology Security Evaluation [5, 7]. The technology product evaluation criteria are also licensed to the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) and are available as ISO/IEC Standard 15408:1:2009 [8]. Protection profiles identify the security functional and assurance requirements for products that implement the technology and state explicit assumptions about the security characteristics of the operational environment of the product by referencing requirements for technology infrastructure systems such as operating systems, databases, and network devices. Industry consortia and standards bodies have developed protection profiles that capture the common security requirements of their technologies, which enable standards-based product security evaluations.

Protection profiles have been developed for technologies that have security functional and assurance requirements similar to those of DID Methods, including certification authorities [6], enterprise identity and credential management systems, and electronic signature creation modules. A DID Method protection profile should provide a detailed exposition of the security functional and assurance requirements for DID Methods and careful linkage to assumptions about elements of the DID ecosystem. A protection profile could be used as a basis for technical product evaluations of DID Method software.

A second model for expressing common security objectives and requirements is the U.S. National Institute of Standards and Technology’s (NIST) Security and Privacy Controls for Information Systems and Organizations, which is available as Special Publication 800-53 [9]. The Common Criteria applies to products, while NIST’s catalogue of controls applies to systems. The NIST catalogue identifies requirements on the technical components of the system and on the organization that operates the system. For example, the Common Criteria identifies a requirement that a product provides an access control mechanism for restricting access to system services according to an access control policy, while the NIST requirements also identify the obligation of the organization to establish access policies according to a principle of least privilege.

DID Methods have security objectives and requirements in common with most security-critical applications and systems that execute within a complex computing, storage, and communication ecosystem (e.g., requirements for various system administrator roles, secure updates, secure transmission of audit data to a remote repository, self-test capabilities, re-authentication after password change). They also have some unique security objectives and requirements. This section focuses on the relatively unique security objectives and requirements of DID Methods. A complete DID Method protection profile would detail the specialized objectives and requirements and would either directly include the more generic objectives and requirements or reference the source that detailed them.

This section highlights DID Method-specific security objectives and requirements that may be challenging to meet, given the DID v1.0 core architecture, data model, and representations. The goal is to raise potential issues, encourage discussion among DID Method designers and standards authors, and further community understanding of DID system security obligations and dependencies. Requirements for cryptographic algorithms and implementations are addressed in a separate document.

The following requirements are color-coded according to how (and if) the requirement may be assessed.

Red: May be impossible to verify through assessment of the DID Method alone.

Blue: Verifiable during onboarding (inspect software, operational configurations, administrative processes). A profile may restrict the set of accepted DID Methods to those that implement optional portions of the specification.

Green: May be verifiable automatically by a resolver (table lookup, check a hash value).

Brown: Not supported by the current specification.

Black: Informational.

5.1. DID Uniqueness

A DID must be globally unique to prevent identifier collisions that would lead to identity ambiguity for relying parties. The operator of a DID Method is responsible for ensuring that its DID Method (software implementation, configurations, and operation) generates unique identifiers. Two different DID Method operations may generate identifiers that are unique to the method and operational instance but are not globally unique. Ensuring global uniqueness of DIDs, therefore, requires that relying parties be able to distinguish a DID generated by one DID Method operator from an otherwise identical DID generated by a different DID Method operator. If DID Method names are not globally unique, this is not possible.

Requirements:

- A DID Method name SHALL be globally unique.
- A DID Method SHALL generate unique method-specific identifiers.

- A DID Method SHALL resolve to a unique DID document, or the resolver SHALL return an error.

5.2. Integrity of DID-Related Data

A DID Method shall protect the integrity of DID-related data. Protecting the integrity of DID-related data is essential to defend against identity fraud. Typically, the supporting infrastructure systems in the operational environment provide some integrity protection mechanisms, and a DID Method must use those correctly to satisfy its obligation for protecting the integrity of DID-related data.

Requirements:

- A DID Method SHALL apply tamper protections to DID-related data.
- A DID Method SHALL check DID-related data for evidence of tampering before completing operations. If evidence of tampering is found during a DID Method operation, the DID Method SHALL return an informative error message in response to the operation request.
- DID-related data include DIDs, DID documents, and DID metadata.

5.3. Availability of DID Method Operations and Persistence of DID-related Data

For relying parties, the availability of DID Method operations may be essential to security-critical operations. Some DID Method operations may require higher availability and performance than others. For example, a relying party may need to resolve a DID to provide a time-critical service for its subject (e.g., verifying investor instructions for an online securities trade). While DID resolution sometimes requires very high availability and performance, obligations to confirm real-world subject and controller characteristics may result in DID creation and update availability only during regular business hours. While DID resolvers must immediately identify deactivated DIDs, sometimes deactivating a DID involves legal processes, so the performance of DID deactivation operations can be tied to the pace of legal processes.

The DID ecosystem relies on the persistence of DID-related data. A DID Method must correctly use the data persistence mechanisms provided by the supporting technology infrastructure systems in the operational environment.

Requirements:

- A DID Method operator SHALL specify the availability characteristics of its DID Method operations (Create, Resolve, Update, and Deactivate). This specification MAY be expressed through a service level agreement (SLA).
- A DID Method operator SHALL specify the persistence characteristics of DID-related data. This specification MAY be expressed through SLA.
- DID-related data include DIDs, DID documents, and DID metadata.

5.4. Access Control and Access Policy

The contents of a DID document provide security-critical information to relying parties. Knowing which entities are permitted to update a DID document enables a relying party to determine the trustworthiness of the DID document contents. To make this possible, DID documents must list all of the DID controllers.

Not all DID controllers should be authorized to modify all the contents of a DID document. For example, the standard practice of rotating the cryptographic keys, used for encrypting messages to be sent to the DID subject, requires periodic updates to the `keyAgreement` property of the DID document. While authorizing an automated process to make such periodic updates to batches of managed DID documents is likely good security practice, authorizing that same process to modify other elements of those DID documents could have negative security consequences.

DID documents should include authorization policy statements that define each DID controller's authority over each element of the document. DID Methods must then enforce those policies.

Requirements:

- A DID Method SHALL ensure that only entities identified as controllers within a DID document can modify that DID document.
- A DID document SHALL explicitly define the extent of a DID controller's authority over specific elements of the DID document.
    - For example, some DID controllers are permitted to add a `verification method` and `verification relationship` and then modify or delete those same elements; other DID controllers are permitted to modify a `Service`.
    - This requires the development of a method for specifying DID controller authorizations within a DID document and a method for policy enforcement through DID Method operations.
- A DID Method SHALL enforce the security policies expressed in a DID document.

5.5. Accountability

To enable security audits of DID Method technology components and the individuals and organizations that operate them, systems must be able to log security- and privacy-relevant events. Because DID metadata captures many security-relevant events at the application level of DID Method operation, DID metadata will be an important component of the DID Method audit trail. The technology infrastructure supporting DID Method execution will also contribute to security logs.

Requirements:

- DID metadata SHALL include:
    - DID of the DID controller responsible for a DID update.
    - Timestamps applicable to the DID document, identified through the `created` or `updated` properties.

- DID Methods SHALL generate audit records for security-relevant events.
    - Audit records SHALL include at least the date and time of the event; type of event; identity of the subject (e.g., user identifier of a method user, platform-specific identifier of a site administrator, IP address associated with an external request); and outcome of the event.
- Integrity protection SHALL be applied to DID metadata and activity logs and their association with DID Method operations.

5.6. Non-repudiation of DIDs and DID document updates

Linking digital identities to their real-world counterparts is essential and challenging. While DID Methods focus exclusively on managing digital representations of identity, they must interact with other systems and external processes that address real-world subjects. Typically, this will involve communication with an individual who is either the subject of a DID or has custody of an item that is the subject of a DID.

Many forms of identity fraud target weak links between digital identities and their real-world counterparts. To prevent or detect attempts at fraudulent creation or manipulation of digital identities, DID Method operators must substantiate links to real-world subjects and custodians before creating or updating DID documents. Evidence of these interactions should be preserved in case individuals later attempt to repudiate DID operations.

Requirements:

- A request for DID creation SHALL NOT be processed until the DID subject (or legal custodian or guardian) provides external evidence of authorization.
    - Before a new DID/DID document is recorded, the system SHALL verify that an individual who is legally associated with the DID subject has taken an approval action external to the DID Method (e.g., the individual must browse to a web site and enter a PIN code that was delivered by postal mail).
- Any requested DID document updates that are not routine (e.g., key rotation) SHALL NOT be processed until a designated DID controller provides external evidence of authorization.
    - A designated DID controller SHALL be alerted to the requested update.
    - The system SHALL verify that the designated DID controller has taken an approval action external to the DID Method.
- The DID controller designation SHALL be made through a DID document security policy assertion.

6. Privacy Objectives and Requirements for DID Methods

Government regulations such as the California Privacy Rights Act (CPRA) of 2020 [14], which will take effect in January 2023, and the European Union's General Data Protection Regulation (GDPR) of 2016 [12] define data privacy rights of individuals and data protection obligations of entities that collect data about individuals. Variations of the Fair Information Practice Principles

(FIPPs) [13] have been tailored to guide the data-collection and protection practices of organizations such as U.S. federal agencies.

Data privacy regulations like GDPR and CPRA are essential to defining rights and obligations, but current engineering standards fail to characterize systems that can protect personal data. Privacy engineers are systems engineers who translate high-level goals for personal data protection into concrete system requirements. In [10], NIST introduces privacy engineering and risk management in the context of federal systems, defining three privacy engineering objectives of predictability, manageability, and disassociability, and clustering nine FIPPs principles under these three objectives. MITRE's Privacy Requirements Definition and Testing [11] sketches an approach to decomposing high-level privacy objectives, first to the level of the general system, then to more system-specific requirements.

The DID core architecture was defined with an explicit goal of enabling individual control over revelations of personal information. While the standard strongly encourages DIDs and DID documents without Personally Identifiable Information (PII), DID Methods run within an operational environment that links to other elements of the DID ecosystem, the Internet, and the physical world. Many DID privacy challenges are likely to emerge from this larger context. Privacy objectives apply to the technical components of the system and to the organization that operates the system.

We recommend that stakeholders in the DID technology and user community work together to develop common objectives that address personal data protection in the larger context of the DID ecosystem (e.g., acknowledging that verifiable credentials will often be attached to DIDs) and addressing realistic threat models.

The following sections summarize NIST's three privacy engineering objectives [10] of predictability, manageability, and disassociability and discuss how these objectives apply to DID Methods. These objectives use the definition of PII given by the U.S. Office of Management and Budget [13]: "information that can be used to distinguish or trace an individual's identity, either alone or when combined with other information that is linked or linkable to a specific individual."

This section highlights DID Method-specific privacy objectives and requirements that may be challenging to meet given the DID v1.0 core architecture, data model, and representations. The goal is to raise potential issues, encourage discussion among DID Method designers and standards authors, and further community understanding of DID system privacy obligations and dependencies.

6.1. Predictability

Predictability is defined as enabling reliable assumptions by individuals, owners, and operators about PII and its processing by an information system. This objective maps to the accountability, authority, purpose specification and use limitation, and transparency FIPPs.

The predictability objectives and associated FIPPs may focus on how closely the contents of the public-facing element of a DID (the DID document) matches the information processed and stored by the DID Method and the operational environment in which it executes. A DID Method or associated system may collect information that is appropriate for bootstrapping trust in the identity of the subject and performing the service of generating the DID and managing the DID document (e.g., payment information such as a credit card holder name and billing address, external information confirming a real-world identity), but users would reasonably expect that this additional information would be limited and would not be retained longer than necessary.

Requirement:

- A DID Method operator SHALL provide the capability for a DID controller to request a list of all types of PII maintained by the system and the entities that can view and/or modify each type.

6.2. Manageability

Manageability is defined as providing the capability for granular administration of PII, including alteration, deletion, and selective disclosure. This objective maps to access and amendment, accountability, minimization, quality, integrity, and individual participation FIPPs.

It is important to identify the PII about the DID subject (or the human custodian of a non-human DID subject) that is stored by the DID Method and its operational environment and various users' (e.g., DID controllers) authority to alter, delete, or disclose elements of the PII. For example, the DID controller (possibly the DID subject), who requested the creation of the DID and supplied a postal address to enable real-world confirmation, should be permitted to view and update the postal address. However, it may be inappropriate for other DID controllers to have access to the postal address. Can other users access PII about the DID subject? Can administrators of the systems in the operational environment access PII?

Some technology infrastructures that underlie DID Method implementations may pose additional challenges for the manageability of PII. For example, while it is important to enable the correction or deletion of PII that was inadvertently included in a DID document, persistent, distributed storage systems may have limited support for full erasure. The correction or deletion of information is difficult for many current technologies. Because DIDs are specifically intended to provide strong privacy protection for individuals, DID Method implementors should take special care to address manageability limitations of the technology infrastructure on which their DID Methods depend.

Requirement:

- A DID Method operator SHALL provide the capability for a DID controller to request and confirm successful modification and/or deletion of PII from the system.

6.3. Disassociability

Disassociability is defined as enabling the processing of PII or events without association to individuals or devices beyond the operational requirements of the system. This objective is most closely linked to the minimization FIPP (which limits the PII collected/processed to only that necessary), but the objective also maps to the accountability FIPP.

One of the fundamental goals of the DID core architecture is to enable separation of identities, personas, and interactions. By intentionally defining DIDs (and DID documents) distinct from verifiable credentials, the conceptual architecture of the DID ecosystem strives to meet the disassociability objective.

While the conceptual architecture of the DID ecosystem supports disassociability, the technology infrastructure underlying DID Method implementations is likely to present challenges to disassociability that are similar to those faced by other security- and privacy-critical systems. For example, Common Criteria protection profiles often express a requirement that no general-purpose computing capabilities be available on the system. This prevents a user from executing a spreadsheet application on the server that runs the certification authority service. The intent of such requirements is to prevent security attacks that target general-purpose applications from impacting security-critical services. Similar requirements are needed to guard against intentional or unintentional association of PII within the system.

The technology infrastructure used in the implementation of DID Methods includes the usual servers and networking equipment, and special-purpose persistent storage facilities, such as verifiable data registries. To prevent intentional or unintentional associations with PII, the technology infrastructure used by DID Methods to process and store DIDs and DID documents should not be used to process and store other materials. In particular, using the same verifiable data registry to host DIDs and verifiable credentials increases the risk that privacy-damaging associations will occur between DIDs and the verifiable credentials.

Requirement:

- A DID Method implementation, including storage systems for DIDs and DID documents SHALL NOT provide capabilities for general-purpose computing, and SHALL NOT provide capabilities for other security-critical functions.
    - Platforms, storage systems, verifiable data registries, and other technologies used by DID Method implementations SHALL NOT enable processing or storage of verifiable credentials.

7. Conclusion

W3C's DIDs have the potential to become a valuable tool for representing subject identity in the digital realm. Establishing the trustworthiness of specific DID Method technologies and operators will be essential to enable individuals and organizations to rely on the digital identities they manage.

This report outlined an approach to developing requirements for the security and privacy of DID Methods that will enable the assessment of different implementations and services. We recommend that members of the W3C DID Working Group collaborate to draft a set of security and privacy requirements for DID Methods. Publicly available Common Criteria protection profiles [5] provide examples of security requirements tailored to various information technologies that could be used as models for a DID Protection Profile.

By drafting a set of DID Method security and privacy requirements before the DID core architecture, data model, and representations are finalized, the DID community can identify elements of the standard that should be modified to better support security and privacy.

8. References

[1] Patricia Harpring, “Cataloguing Art and Architecture: Introduction and Application of CDWA and CCO,” Getty, May 2020. https://www.getty.edu/research/tools/vocabularies/intro_to_cco_cdwa.pdf

[2] Duns and Bradstreet, “What is a D-U-N-S Number?” https://www.dnb.com/duns-number.html

[3] Delivery Point Validation against the Address Management System (AMS) Database, Comprehensive Statement on Postal Operations, Chapter 2 Postal Operations, 2004. https://about.usps.com/strategic-planning/cs04/contents.htm

[4] W3C Decentralized Identifiers (DIDs) v1.0 Core Architecture, Data Model, and Representations, Candidate Recommendation Draft, 30 April 2021. https://www.w3.org/TR/did-core/

[5] Common Criteria Portal. https://commoncriteriaportal.org

[6] National Information Assurance Partnership Approved Protection Profiles. https://www.niap-ccevs.org/Profile/PP.cfm

[7] “Common Criteria for Information Technology Security Evaluation, Version 3.1, Revision 5,” April 2017. https://commoncriteriaportal.org/cc/

[8] “Information technology—Security techniques—Evaluation criteria for IT security, Third Edition,” International Standard ISO/IEC 15408-1, December 15, 2009. https://www.iso.org/standard/50341.html

[9] National Institute of Standards and Technology (NIST), “Security and Privacy Controls for Information Systems and Organizations, Revision 5, Special Publication 800-53,” September 2020. https://doi.org/10.6028/NIST.SP.800-53r5

[10] National Institute of Standards and Technology (NIST), “An Introduction to Privacy Engineering and Risk Management in Federal Systems, Internal Report 8062,” January 2017. https://doi.org/10.6028/NIST.IR.8062

[11] MITRE, “Systems Engineering Guide: Privacy Requirements Definition and Testing,” https://www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/engineering-informationintensive

[12] Regulation (EU) 2016/679 General Data Protection Regulation, adopted April 27, 2016. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679

[13] Office of Management and Budget (OMB) Circular A-130, *Managing Information as a Strategic Resource*, July 2016. https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/OMB/circulars/a130/a130revised.pdf

[14] The California Privacy Rights Act of 2020 (CPRA). https://oag.ca.gov/system/files/initiatives/pdfs/19-0021A1%20%28Consumer%20Privacy%20-%20Version%203%29_1.pdf