Network Working Group
Internet-Draft
Intended status: Informational
Expires: 16 September 2026

A. Rosenberg
Veridigo

15 March 2026

Digital Emblems - Architectural Considerations
draft-rosenberg-diem-architecture-00

Abstract

This document explores architectural considerations for digital emblems as discussed in the DIEM working group. It is intended to complement the use cases and requirements in [DIEM-REQS] by sketching one way to think about how the pieces of a DIEM architecture might fit together. This document does not seek to propose a specific architecture. It instead captures the author's current mental model to help inform the requirements process and subsequent architectural work.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://ohmantics.github.io/diem-architecture/draft-rosenberg-diem-architecture.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-rosenberg-diem-architecture/.

Source for this draft and an issue tracker can be found at https://github.com/ohmantics/diem-architecture.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 16 September 2026.

Table of Contents

1. Introduction

The DIEM working group is developing standards for digital emblems: a means for associating metadata with physical or virtual assets, perhaps to indicate what treatments they should receive by reference to various normative frameworks [DIEM-REQS]. The requirements document identifies use cases ranging from International Humanitarian Law protections to trade and logistics labeling, and derives requirements for emblem format, discovery, validation, and extensibility.

This document sketches one way to think about the architecture. It is not a formal proposal of any kind, but an attempt to capture the author's mental model that may inform ongoing discussions. This is offered as points for conversation, not remotely as conclusions.

The mental model described here is primarily concerned with the shape of the data: what an emblem record looks like and how use cases configure it. Discovery follows naturally from DNS (given a FQDN, query for records), and trust follows from the existing DNSSEC chain of trust.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The following terms are used in this document to describe concepts in the mental model. These are not formal definitions and may evolve as the working group develops its architecture:

Profile: A description of how a particular use case shapes the data in an emblem record. Other possible terms are "schema" and "template" but this document uses "profile" for now and the final terminology is an open question.

Terminology for "(digital) emblem," "asset," "emblem issuer," "authorizing entity," "validator," and "validation" is defined in [DIEM-REQS] and used here with the same meaning.

3. Roles

From the perspective of this mental model, there are really only two actors: a domain holder who publishes emblem records, and a validator who queries for them. The FQDN identifies the asset, and the DNSSEC-based chain of trust provides attestation.

Everything else (who authorized the emblem in the real world, what legal framework applies, what visual symbol is associated with it) lives in the data as (possibly) optional fields defined by the profile. An emblem with specific legal protections would include a field referencing those protections. An emblem with a specific visual appearance (a Red Cross, a UN logo) would include a field carrying that representation, possibly as SVG or an image format. The architecture does not need to distinguish these concerns from any other metadata.

4. Data Model

At its simplest, an emblem record is a dictionary data structure: a collection of key-value pairs, including arrays and dictionaries as value types. Some fields are common across use cases (the asset identifier, emblem type, issuer identity, profile identifier). Others are defined by the profile for a specific use case, such as a UN Number for hazardous materials or a serial number for a consumer device.

Some value types are more complicated and require consideration such as the bounds of a protected location. Are such types constrained to the surface of Earth? Are they two-dimensional or three-dimensional? Must they be presented as pre-decomposed into convex shapes? Nearly every other value type the author has considered is simpler or has obvious precedent.

There are fields whose values are included in a cryptographic hash of the record (proving the record has not been tampered with) and fields that are not part of the hash and may change over time. Live GPS coordinates from an IoT-enabled shipping container are an example of the latter.

5. Hierarchical Structure

Emblem records are presented in a hierarchy. It is easiest to reason about this with physical assets, though the same structure applies to virtual ones. A pallet of smartphones, for example, has a record for the wood pallet itself (citing ISPM 15 compliance) and roughly 1400 records for each boxed smartphone. Each phone might carry an identifier record (serial number, model, color, branding), a hazardous materials record for its lithium battery (UN 3481), and recycling instruction data. The manufacturer attests to these device-level records. (Conceptually, some of this data could also live on the device itself and be read out via RFID, though that level of discovery is beyond this scope.) That may then be wrapped in shipper or customs facilitator records for logistics purposes (customs valuation, origin, carrier details).

This hierarchy is effectively a Merkle tree of chains of trust. Each level includes hashes of its children, so the integrity of the entire structure can be verified from the root. Different levels can be attested by different domain holders: the manufacturer attests to the device records, and the shipper wraps them in their own attestation at their own domain.

How deeply a validator walks the tree depends on the use case. A customs officer scanning a pallet might verify only the shipper's top-level attestation. A recycling facility processing an individual device might need to verify the hazmat data back to the manufacturer. A full audit would verify everything. The architecture supports the full tree; the use case determines how deep you go.

6. Serialization

Almost every realistic use case produces records larger than the roughly 1500 bytes of a typical UDP DNS response, so a mechanism for delivering larger payloads is needed. DNS could serve as a bootstrap (pointing to where the full record can be fetched), with the actual data retrieved over HTTPS, an extension to TCP DNS, or DoH.

An emblem record is a dictionary, and the choice of serialization format is an open question. Typical dictionary serialization formats are JSON, XML, Property Lists, and CBOR. Of these, only CBOR has direct support for selective disclosure through sub-dictionary encryption, allowing different fields to be readable by different validators. This is needed for use cases where some emblem data is confidential. Similar selective disclosure could be achieved in other formats through application-layer encryption of individual field values.

7. Discovery

Given an asset's FQDN, a validator queries DNS for emblem records. Emblem records are most likely to be published in a designated subdomain of the asset's domain (e.g., _diem.hospital.example.org). This would allow the zone to be delegated from the asset's regular DNS server to a DIEM-enabled authoritative server, which would use a different backing database laid out for emblem records and possibly integrated with an asset management or inventory ERP system.

For payloads that exceed DNS response sizes, the data could be retrieved via an extension to TCP DNS or DoH, or DANE could bridge from the DNS lookup to an HTTPS endpoint, maintaining the DNSSEC-based chain of trust through the transition.

8. Profiles

Profiles are expected to be defined in separate documents as extensions to the base standard, each specific to a use case.

Fields fall into two categories. Signed fields are included in a cryptographic hash of the record, proving it has not been tampered with. These are the immutable assertions like identifiers,

classifications, and legal protection citations. Informational fields are not part of the hash and may change over the lifetime of the emblem, such as live GPS coordinates or estimated arrival times.

Some signed fields may also be encrypted (for example, internal forensic data about who created the record, when, and on what equipment). This raises an open question: how is an encrypted subrecord included in the hash when the raw data is not accessible for validation? The ciphertext itself could be hashed, or a hash of the plaintext could be presented alongside the encrypted data, or some combination. This is a question for the cryptographic design work.

9. Security Considerations

This document sketches architectural considerations rather than proposing specific mechanisms, so detailed security analysis is left to the standards-track documents. A few observations:

The DNSSEC-based chain of trust proves integrity and authenticity for a given domain, but does not help a validator distinguish a legitimate domain from a deceptive lookalike. For example, a domain using Cyrillic "а" (U+0430) in place of Latin "a" (U+0061) is visually indistinguishable from the real domain but resolves to a completely different zone. How validators establish that they are querying the correct domain is an unsolved problem.

Some use cases (notably IHL) require that the act of querying for an emblem be undetectable. This involves making DNS queries anonymous, which is not the author’s area of expertise.

A full audit of a hierarchical emblem structure (validating every sub-record back to its origin) is a case of write amplification. For a pallet of small items, this could mean thousands of cryptographic verifications each including a DNS query. Use cases should consider whether full audit is actually required or whether top-level verification is sufficient.

Use cases employing selective disclosure (encrypted subrecords visible only to certain validators) introduce trade-offs between confidentiality and transparency that each profile will need to address.

10. IANA Considerations

This document has no IANA actions.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

11.2. Informative References

[CHARTER] "Digital Emblems", 27 May 2025, <https://datatracker.ietf.org/doc/charter-ietf-diem/01/>.

[DIEM-REQS] "Digital Emblems - Use Cases and Requirements", 2026, <https://datatracker.ietf.org/doc/draft-ietf-diem-requirements>.

Acknowledgments

AI was used for editorial review of this document. The author takes full responsibility.

Author’s Address

Alex Rosenberg
Veridigo
Email: alexr@veridigo.com