Towards Queryable Verifiable Credentials

Gertjan De Mulder*,†, Ruben Dedecker*,†, Ben De Meester and Pieter Colpaert

IDLab, Department of Electronics and Information Systems,
Ghent University – imec, Technologiepark-Zwijnaarde 122, 9052 Ghent, Belgium

Abstract

With initiatives such as the European Digital Identity Wallet, the exchange of verifiable data via digital credentials using digital wallets is reaching mainstream adoption. While digital wallets gain adoption and the ecosystem grows, the flexibility and interoperability of querying verifiable data remains limited: digital credentials can only be directly queried based on JSON keys, thus hampering support for internationalization and alternative semantics. With this work, we present a practical approach to leverage Semantic Web technologies to combine the expressiveness of RDF with the semantic nature of W3C's recommended Verifiable Credentials and existing wallet protocols like OID4VP. We leverage SPARQL queries to request specific claims within the OID4VP protocol, which are evaluated over the combined RDF representation of the wallet credentials, to return Verifiable Presentations that contain the requested claims using selective disclosure. We applied named graphs with blank node graph names to ensure a uniform and globally unique connection between credentials and their corresponding claims and proof graphs when storing credentials in a wallet's Knowledge Graph. To retrieve which claims need to be selectively disclosed from their original credentials, we applied an initial conversion from RDF triple predicates to JSONPath pointers, thus currently supporting only a subset of SPARQL expressivity. Through the addition of a SPARQL query in the OID4VP authorization flow, we enable semantically enriched query evaluation over the stored credentials, opening the way to semantic alignment of multilingual vocabularies and similar ontologies used in different online ecosystems. Future work is needed to improve SPARQL support to enable querying over complex claim requirements, in terms of the query to JSONPath pointers, as well as addressing metadata requirements for the requested claims.

Keywords

Verifiable Credentials, SPARQL

1. Introduction

Digital credentials – i.e., tamper-evident digital assertions issued by a trusted authority to represent claims about a subject – are gaining adoption: in government by the European Parliament [1] and implemented in the European Digital Identity Wallet [2], in industry [3], and in international organizations through standardization, showcased by, e.g., the recently W3C recommended Verifiable Credentials (VC) 2.0 specifications [4]. Due to its gaining adoption, the legal obligations for preserving the privacy of the individual and ensuring data minimization [5] have to be taken into account. This is typically achieved through cryptographic methods such as selective disclosure [6].

Selective disclosure for Verifiable Credentials is made available by deriving Verifiable Presentations (VPs) [4]: creating newly combined verifiable subsets of claims derived from credentials stored in a personal wallet. For example, at the job office, instead of requesting individuals to provide a complete set of documents containing more information than needed (potentially introducing bias), a selection of relevant claims can be requested, such as (i) a valid driver's license; (ii) a valid diploma degree; and (iii) the job title of the applicant's current job (if available). Minimizing the information shared reduces the likelihood of discrimination against potential candidates.

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan
*Corresponding author.
†These authors contributed equally.
gertjan.demulder@ugent.be (G. De Mulder); ruben.dedecker@ugent.be (R. Dedecker); ben.demeester@ugent.be (B. De Meester); pieter.colpaert@ugent.be (P. Colpaert)
https://ben.de-meester.org/#me (B. De Meester)
0000-0001-7445-1881 (G. De Mulder); 0000-0002-3257-3394 (R. Dedecker); 0000-0003-0248-0987 (B. De Meester); 0000-0001-6917-2167 (P. Colpaert)

As credentials become increasingly integrated in digital data flows, we expect an increasing need to request and share flexible combinations of credentials.

At the heart of this flexibility requirement, the semantics of the exchanged data need to be well understood. Here, the Verifiable Credential Data Model [4] makes use of the JSON-LD format [7] to represent the data stored in the credentials using the Resource Description Framework [8] in JSON format. This gives the double advantage of (i) managing the syntactic representation of the credential using (widely adopted) JSON tooling, whilst (ii) handling the semantic interpretation of the data using RDF. However, while existing works aim to formalize selective disclosure and ZKP algorithms for RDF triples and RDF terms [9, 10], the current protocols for Verifiable Credentials and the resulting Verifiable Presentations tend to disregard this semantic information throughout their lifespan. Instead, they rely on JSONPath pointers through protocols such as the Presentation Exchange [11] or the Digital Credential Query Language [12] to specify the selection of claims from specific credentials that are being requested for disclosure.

The reliance on JSONPath pointers – which are one-to-one coupled with the Verifiable Credential's schema – imposes an implicit reliance on the issuer's chosen context, which is typically influenced by the issuer's local context such as language and culture. This limits the reusability of the pointer outside the issuer's local context. Taking the example of a diploma credential: the European Learning Model[1] specifies its model using English keys (e.g., `accreditation` linking to http://data.europa.eu/snb/model/elm/accreditation); where the Flemish derived application profile[2] specifies the same terms, but using Flemish keys (e.g., `Bewijs.accreditatie`, also linking to http://data.europa.eu/snb/model/elm/accreditation). We assume such discrepancies in JSON keys to occur more and more frequently, making interoperability difficult even when the underlying managed data models are correctly reused. Secondly, as JSONPath pointers use tree-based traversal, they do not support graph-based selections over multiple credentials. Even simple cases could require graph-based selections. Taking the example of proving ownership of a car: A person holding the credentials of both a valid license and insurance would require claim retrieval over multiple interlinked credentials, which would benefit from graph-based selection.

With this work, we present a novel method to leverage mature Semantic Web technologies to more flexibly but still effectively combine claims represented in Verifiable Credentials, leveraging a SPARQL subset as a selective disclosure mechanism to automatically combine claims from multiple Verifiable Credentials into a Verifiable Presentation. We apply our method to the maturing OpenID for Verifiable Presentations protocol (OID4VP) [12].

2. Method and Implementation

In this section, we show how we can retrieve and combine arbitrary claims from multiple Verifiable Credentials *semantically* – i.e., without needing to rely on the specific credential schemas used – by extending the OID4VP authorization request [12] with SPARQL query functionality. A proof of concept implementation of the SPARQL query evaluation over Verifiable Credentials can be found on Github[3]. Our method consists of five steps (Figure 1): (i) integrate all wallet's credentials in a single graph store without losing the credentials' internal connections; (ii) request a selectively disclosed set of claims using a SPARQL query instead of a list of JSONPath pointers; (iii) transform the SPARQL query to match the credentials' internal representation to retrieve which credentials contain which claims; (iv) transform the semantic relations back to the credentials' JSONPaths to construct the Verifiable Presentation; and (v) return the compliant Verifiable Presentation to the original requester.

Step 1: Integrate all wallet's credentials in a single graph store To enable the evaluation of a SPARQL query over stored Verifiable Credentials, these credentials are stored in a graph store on the

[1]Retrieved via https://op.europa.eu/en/web/eu-vocabularies/dataset/-/resource?uri=http://publications.europa.eu/resource/dataset/snb-model&version=3.2 at 10/07/2025.
[2]Retrieved via https://data.vlaanderen.be/doc/applicatieprofiel/leerprestatiecredential/ at 04/07/2025.
[3]Github repository for the POC is located at: https://github.com/KNowledgeOnWebScale/queryable-vcs

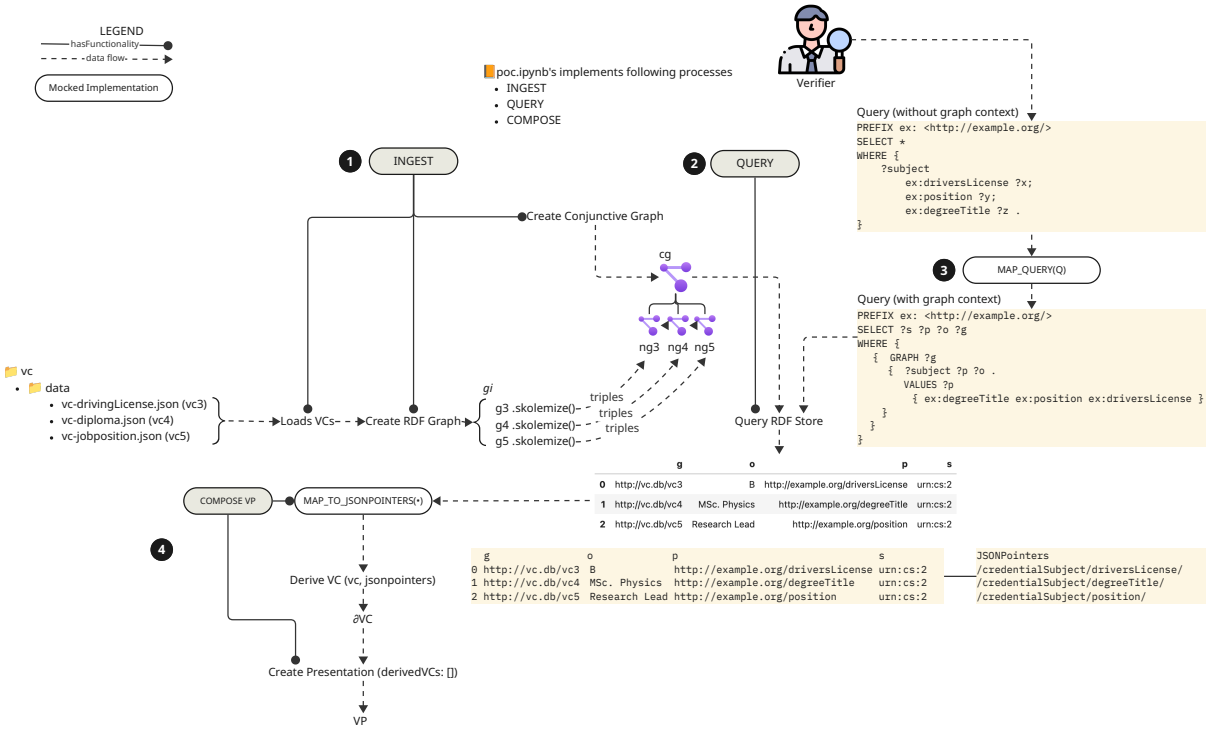

Figure 1: Our proposed method to support SPARQL queries as a selective disclosure pointing method: (i) integrate all wallet's credentials in a single graph store; (ii) request a selectively disclosed set of claims using a SPARQL query; (iii) transform the SPARQL query to match the credentials' internal representation; (iv) construct the Verifiable Presentation; and (v) return that to the original requester.

wallet instance. Verifiable Credentials can be interpreted in RDF. However, as indicated by Braun et al. [10], the Verifiable Credential Data Model mandates linking the proof graph from a node within the default graph. This is problematic when storing multiple credentials in their RDF representation: all claims of all credentials, being part of the default graph, are subject to automated merging when inserted into a knowledge graph. This makes it impossible to distinguish which claims originate from which credential in their RDF representation.

To ensure we keep the link between claims and their originating credential, the ingestion of these credentials into the graph store requires the contents of their default graph to be embedded in a named graph [13]. A subsequent skolemization step ensures that all graph names used are converted to blank nodes that are relabeled on ingestion in the graph store, preventing coreferencing between the content graphs of different credentials.

Step 2: Request a selectively disclosed set of claims using a SPARQL query To query the wallet's credentials, we envision the reuse of the OID4VP Authorization Request [12], where we introduce a `sparql_query` parameter, similar to how the `dcql_query` parameter is used to pass DCQL queries to the wallet [12].

We assume that the SPARQL query is formed based on the native RDF representation of the credentials over which the query is evaluated (Listing 1). Querying the credential metadata currently means that the requester must be made aware of how the wallet links the claims with the credential metadata (i.e., via the minted credential graphs). Given this was introduced in Step 1 (combining multiple credentials in one Knowledge Graph), this is no longer aligned with the VC standard. Extending the VC standard to support these types of metadata queries is out of the scope of this paper, and we focus on querying the claims themselves.

```
PREFIX ex: <http://example.org/>
SELECT *
WHERE {
    ?s ex:degreeTitle ?z ;
       ex:driversLicense ?x ;
       ex:jobPosition ?y
}
```

Listing 1: Querying claims from different credentials

```
PREFIX ex: <http://example.org/>
PREFIX vc: <https://www.w3.org/2018/credentials#>
SELECT
    ?credential ?subject ?p ?o
WHERE {
    GRAPH ?credentialGraph {
        ?credential a vc:VerifiableCredential ;
            vc:credentialSubject ?subject ;
        ?subject ?p ?o .
        VALUES ?p { ex:degreeTitle ex:position ex:driversLicense }
    }
}
```

Listing 2: Rewritten query to target named graph with credential contents

Step 3: Transform the SPARQL query to match the credentials' internal representation At the credential wallet, an incoming SPARQL query must be converted to take into account the internal RDF representation of the stored credentials. To achieve this, the query statements are embedded in the GRAPH keyword to direct the queried statements to the named graphs storing the credential claims (Listing 2). This aligns the queries with the conversion performed in Step 1. Within the context of this proof of concept, our considered query examples are chosen with simplicity in mind, allowing us to discuss the system as a whole. Hence, future research is needed to support more complex query mapping scenarios (e.g., in the case of VCs that have a more complex structure).

Step 4: Construct the Verifiable Presentation Having found the relevant credentials for the requested claims, we can rely on mature JSON-LD compaction algorithms [14] for the wallet to transform the semantic relations used in the SPARQL query back to JSONPaths conforming to the credentials' original JSON representation. These JSONPath pointers are needed to construct the Verifiable Presentation using existing Verifiable Credential libraries (e.g., https://github.com/digitalbazaar/vc).

Step 5: Return the compliant Verifiable Presentation to the original requester At the requesting side, the Verifiable Presentation returned from the authorization request can now be validated based on its signature and the identity of its issuer. The original query, as shown in Listing 1, can now be evaluated over the Verifiable Presentation's native RDF representation to receive the verified results for the query.

3. Conclusions

With this initial work, we demonstrate a practical solution for the selective disclosure of claims via Verifiable Presentations based on SPARQL queries applied to the OID4VP protocol, by taking into account

the underlying RDF semantics of Verifiable Credentials. The result of the query is a presentation of claims from the original credentials, which is then, at the requesting side, re-processed with the same mappings to provide a query response. As such, we ensure that claims are verified correctly.

The addition of a `sparql_query` as an alternative to the `dcql_query` parameter is a mild intrusion in the OID4VP authorization request flow. We provide a straightforward adoption path for integrating SPARQL queries without impacting the overall flow or functionality. Our method allows for more flexible querying of claims, but does not influence authorization decisions, as the authorization mechanisms are based on the claims that are being returned, not on the internal query evaluation process. The evaluation of the returned presentations for their metadata and query results occurs outside the existing protocol stack.

Integrating the semantics of credentials into the request flow enables the use of existing Web standards such as OWL for schema alignment, facilitating alignment of credentials over multiple languages and ecosystems. Where initial wallet implementations can rely on ecosystem enforcement of standardized schemas (e.g. ISO standards), further adoption is likely to cause divergence between data models. Standardized mappings can be provided at an ecosystem level, and ad hoc mappings can be passed by a requesting party to allow wider matching of requested claims to stored credentials.

Although the provided implementation validates its feasibility, the avenues for future work are manifold.

A JSONPath-to-SPARQL transformation could be devised to serve as a proxy for existing OID4VP clients (i.e., reinterpreting the `dcql_query` parameter as a SPARQL query), resulting in a completely OID4VP-compatible solution.

Introducing SPARQL support strains the current Verifiable Credentials Data Model, as the native RDF representation requires additional transformations to allow for unambiguous query evaluation matching claims with the credential they originate from, supporting only a subset of the current RDF and SPARQL capabilities. Further investigation is needed to assess the possibility for the Verifiable Credential model to support native querying over both claims and metadata in its RDF representation.

We currently need to rely on JSON-based Verifiable Credential libraries to create the Verifiable Presentation (i.e., the reverse transformation into JSONPaths in Step 4). The work of Braun et al. [10] provides a basis for an RDF-native representation of signatures over RDF graphs with support for selective disclosure of individual terms, opening avenues to extending OID4VP to an RDF-native protocol to further increase its flexibility.

In terms of functionality requirements, the proposed approach imposes an overhead on the wallet instance: for managing Verifiable Credentials as a knowledge graph, and for matching a user query over that knowledge graph to individual claim requirements over the stored credentials. Research is needed towards the added benefit of semantic cross-credential querying and ontology alignment outweighs this added cost.

Acknowledgments

The described research activities were supported by SolidLab Vlaanderen (Flemish Government, EWI and RRF project VV023/10), and Interreg project SecuWeb (0100085).

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-4o to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] Consolidated text: Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC, Official Journal of the European Union (2024).

[2] EU Digital Identity Wallet – Architecture and Reference Framework, Technical Report, European Commission, 2025. URL: https://eu-digital-identity-wallet.github.io/eudi-doc-architecture-and-reference-framework/2.4.0/, version 2.0.0.

[3] D. Deimel, et al., Organisations-identitäten, Technical Report, Bitkom, 2025. URL: https://www.bitkom.org/sites/main/files/2025-01/bitkom-whitepaper-organisationsidentitaeten.pdf.

[4] M. Sporny, T. Thibodeau Jr, I. Herman, G. Cohen, M. B. Jones, Verifiable Credentials Data Model v2.0, W3C Recommendation, Verifiable Credentials Working Group, 2025. URL: https://www.w3.org/TR/vc-data-model-2.0/.

[5] I. Lella, M. Theocharidou, E. Magonara, A. Malatras, R. Svetozarov Naydenov, C. Ciobanu, C. Georgios, ENISA Threat Landscape 2024, Technical Report, The European Union Agency for Cybersecurity (ENISA), 2024. URL: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024.

[6] Š. B. Ramić, E. Cogo, I. Prazina, E. Cogo, M. Turkanović, R. T. Mulahasanović, S. Mrdović, Selective disclosure in digital credentials: A review, ICT Express 10 (2024) 916–934. doi:10.1016/j.icte.2024.05.011.

[7] M. Sporny, G. Kellogg, M. Lanthaler, JSON-LD 1.1 – A JSON-based Serialization for Linked Data, Recommendation, World Wide Web Consortium (W3C), 2020. URL: http://www.w3.org/TR/json-ld/.

[8] R. Cyganiak, D. Wood, M. Lanthaler, RDF 1.1 Concepts and Abstract Syntax, Recommendation, World Wide Web Consortium (W3C), 2014. URL: https://www.w3.org/TR/rdf11-concepts/, accessed: 2023-09-29.

[9] D. Yamamoto, Y. Suga, K. Sako, Formalising linked-data based verifiable credentials for selective disclosure, in: 2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), 2022, pp. 52–65. doi:10.1109/EuroSPW55150.2022.00013.

[10] C. H.-J. Braun, T. Käfer, RDF-Based Semantics for Selective Disclosure and Zero-Knowledge Proofs on Verifiable Credentials, in: The Semantic Web, volume 15718, 2025, pp. 383–402. doi:10.1007/978-3-031-94575-5_21.

[11] D. Buchner, B. Zundel, M. Riedel, K. H. Duffy, Presentation Exchange 2.1.1, DIF Ratified Specification, Decentralized Identity Foundation, 2024. URL: https://identity.foundation/presentation-exchange/spec/v2.1.1/.

[12] O. Terbu, T. Lodderstedt, K. Yasuda, D. Fett, J. Heenan, OpenID for Verifiable Presentations 1.0, Technical Report, OpenID Digital Credentials Protocols, 2025. URL: https://openid.net/specs/openid-4-verifiable-presentations-1_0.html.

[13] R. Dedecker, J. De Roo, B. Esteves, P. Colpaert, Demonstrating a pragmatic solution to context associations in RDF using Blank Node Graphs, in: The Semantic Web: ESWC 2025 Posters & Demos, 2025.

[14] D. Longley, G. Kellogg, P.-A. Champin, JSON-LD 1.1 Processing Algorithms and API, W3C Recommendation, JSON-LD Working Group, 2020. URL: https://www.w3.org/TR/json-ld11-api/, https://www.w3.org/TR/2020/REC-json-ld11-api-20200716/.