A Policy-Aware Data Sharing Application Using Solid, Semantic Provenance, and IDS-Inspired Data Governance

Muhammad Raza Naqvi*[1], Antoine Zimmermann[1] and Christoph Braun[2]

[1]Mines Saint-Etienne, CNRS, UMR 6158 LIMOS, 42023 Saint-Etienne, France.
[2]Karlsruhe Institute of Technology, Institute AIFB, Web Science, Karlsruhe, Germany

Abstract

Secure and governed data sharing remains a persistent challenge for organizations that need to exchange resources while retaining control over access conditions, permitted use, provenance, and accountability. Existing approaches often address decentralized storage, access control, policy representation, or semantic interoperability separately, but provide limited support for connecting these concerns within a single operational workflow. This work presents a Solid-based, IDS-inspired framework for secure, policy-aware, and semantically auditable data sharing. Solid pods, WebID-based identities, and Solid authentication provide the decentralized storage and identity layer, while IDS-inspired provider–consumer interactions, transfer requests, agreements, policies, and transfer decisions form the governance and exchange-coordination layer. Machine-readable policies are evaluated before data movement, allowing the transfer component to permit or block an exchange according to conditions such as consumer identity, declared purpose, redistribution permission, temporal constraints, and agreement reference.

The framework is implemented as a modular prototype comprising resource publication, contract negotiation, policy definition and evaluation, governed PUSH and PULL transfer, semantic provenance generation, and SPARQL-based auditing. Rather than implementing a complete IDS connector stack, the prototype retains Solid pods as the participant-controlled resource layer and represents requests, agreements, policy decisions, transfers, and provenance as lightweight RDF resources. The prototype demonstrates how IDS-inspired governance can be integrated with Solid-based decentralized storage to support pre-transfer policy enforcement and semantic inspection of recorded transfer activities.

Keywords

Linked Data, Solid, International Data Spaces, Semantic Auditing

1. Introduction

Data sharing across independent organizations involves more than transferring data between systems. It also requires governance, accountability, and control over how data are accessed and used [1]. These challenges are more pronounced in cross-organizational settings, where providers and consumers are administratively independent and may rely on different infrastructures, identities, authorization mechanisms, and governance procedures [2].

A data-sharing system must therefore support not only storage and access, but also usage conditions, agreements, provenance, and auditability [3]. Solid is a set of Web specifications that

DMKG'26: 2nd International Workshop on Data Management for Knowledge Graphs, October 2026, Bari, Italy
syed.naqvi@emse.fr (M. R. Naqvi*); antoine.zimmermann@emse.fr (A. Zimmermann); braun@kit.edu (C. Braun)
0000-0002-9652-5164 (M. R. Naqvi*); 0000-0003-1502-6986 (A. Zimmermann); 0000-0002-5843-0316 (C. Braun)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-4322/paper_14.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-09

separates applications from data storage and enables resources to be maintained in participant-controlled storage spaces called pods [4]. A pod is an HTTP-accessible data store that may be self-hosted or operated by a selected provider. Participants are identified by WebIDs, which are URI-based identifiers associated with Web-accessible profile documents [5].These capabilities provide a decentralized storage and identity foundation. However, these mechanisms do not by themselves define a complete governed exchange process involving transfer requests, negotiated agreements, policy evaluation, transfer decisions, and semantic audit records.

Data-space architectures address these broader requirements through the concept of data sovereignty, according to which participants should retain control over the conditions under which their data are shared and used [2] [6]. International Data Spaces (IDS) structure exchange as a controlled interaction between providers and consumers, through negotiation, usage-control, and transfer-coordination mechanisms [6].

A remaining challenge is to connect machine-readable policies to the component that performs the transfer. Policies must be evaluated against contextual information before data movement, and the resulting decision must be able to permit or block the transfer. The workflow should also record its decisions and transfer activities for later audit.

To address this need, this paper proposes a Solid-based, IDS-inspired framework for secure, policy-aware, and semantically auditable data sharing. Solid pods, Solid authentication, and WebID identities provide the storage and identity layer, while IDS-inspired concepts structure transfer requests, agreements, and provider–consumer interactions. Before a transfer is executed, the framework evaluates policy conditions such as consumer identity, declared purpose, redistribution permission, temporal constraints, and agreement reference. Compliant requests proceed to transfer and provenance generation, whereas non-compliant requests are blocked before the resource is copied.

The remainder of the paper is structured as follows. Section 2 reviews the related work and positions the proposed framework. Section 3 presents the methodology and implementation, while Section 4 discusses the conclusion, and challenges.

2. Related Work

2.1. Centralized and Decentralized Data Sharing

Existing centralized data-sharing approaches include cloud-based data-sharing platforms [7, 8, 9, 10, 11]. These approaches commonly organize data storage, integration, governance, and access control through centrally managed infrastructures. Although they can simplify access, centrally controlled storage models may couple data to particular applications or service providers. Such coupling can limit portability, application independence, and the ability to move between providers [12].

The Solid literature identifies this platform dependency as part of a broader centralization problem in which applications and service providers control access to user or organizational data [13]. Solid addresses this problem by separating applications from data storage and enabling individuals and organizations to maintain resources in participant-controlled pods. A pod may be self-hosted or hosted by a selected pod provider, while applications access its resources through standard Web mechanisms without controlling the underlying storage [14]. These

resource interactions also follow Linked Data Platform principles for HTTP-based creation and management of Linked Data resources [15].

Early Solid work demonstrated its use for decentralized Web applications and participant-controlled data storage [16, 17]. Solid relies on WebID-based identification, Solid-OIDC authentication, and resource-level authorization mechanisms such as WAC and ACP [5, 18, 19, 20].

Solid also supports Linked Data and RDF representations, allowing structured metadata to be represented and exchanged in a machine-readable form. This makes it possible for applications to interact with shared resources without requiring the data to be stored in an application-specific database.

However, these mechanisms do not by themselves provide a complete governance process for inter-organizational data exchange. Solid does not define a full provider–consumer workflow. In particular, resource-level authorization does not completely govern the conditions under which transferred data may subsequently be processed, redistributed, or reused.

2.2. Data Spaces and Data Sovereignty

The limitations of resource-level access control become particularly important when data are exchanged across organizational boundaries. From a data-space perspective, the broader requirement is described as data sovereignty, meaning that participants should be able to determine the conditions under which their data are shared and used [2].

International Data Spaces frames data exchange as a controlled and sovereign interaction between providers and consumers rather than as a simple file-access or resource-copying operation [6]. The IDS Reference Architecture Model defines the main roles and sovereignty principles of provider–consumer data exchange. The Dataspace Protocol further specifies contract negotiation, agreement handling, and transfer coordination between dataspace participants [21].

Solid and data-space architectures are complementary in this respect. IDS concepts can consequently extend Solid-based resource sharing into a governed provider–consumer workflow.

Previous work has already explored Solid as a foundation for data spaces and federated multi-stakeholder environments. Werbrouck et al. addresses complementary aspects of Solid-based collaboration and data-space integration [22]. Meckler et al. [23] examine Solid and Linked Data as a foundation for dataspaces, while Braun et al. [24] study Solid as a transport mechanism in IDS environments. Our proposed work does not claim that Solid and data-space concepts have never been combined.

2.3. Access Control, Usage Conditions, and ODRL

A key limitation in many data-sharing workflows is the separation between access control and usage control [1]. Access control determines whether a user, consumer, or application can read or write a resource. Usage control concerns the conditions attached to the use of data after access has been granted.

A system may authorize access to a resource while remaining unable to determine whether the consumer's declared use satisfies the provider's conditions. In practice, usage policies are often expressed in human-readable contracts or agreements but are not directly evaluated by the software component responsible for performing the transfer [1]. As a result, data movement

may occur even when the declared purpose or transfer context does not satisfy the provider's policy. Policy evaluation should therefore be connected to the transfer workflow before data are copied or otherwise made available to the consumer. This provides an enforcement point at which the software can allow or block the transfer according to the applicable policy.

The Open Digital Rights Language (ODRL[1]) provides a machine-readable model for representing permissions, prohibitions, duties, constraints, and the parties involved in a policy [25]. ODRL therefore addresses the representation of usage conditions. However, representing a policy in ODRL does not by itself ensure that the policy can be operationally evaluated or enforced.

The challenge is consequently not limited to policy representation. It also concerns connecting ODRL-based policies to executable policy evaluation and to an enforcement point that can permit or block data movement.

In the proposed framework, policy enforcement occurs before the transfer. The framework does not claim to provide continuous technical control over the consumer's actions after the data have been transferred.

2.4. Semantic Provenance and Auditability

Semantic Web technologies provide a suitable foundation for representing the elements of governed data sharing in a machine-readable form. RDF represents information as graph-based statements, while SPARQL[2] provides a standardized query language for RDF[3] graphs [26].

Provenance is important for accountability and the later inspection of data exchanges. Previous work has demonstrated the value of recording the origin, derivation, processing, and flow of data [27, 3]. PROV-O provides a standard semantic model for representing entities, activities, agents, and the relationships among them [28].

Such semantic records support the reconstruction and inspection of recorded transfer activities. They do not by themselves prove that all real-world actions were recorded, but they provide a structured audit trail for the operations captured by the system.

A SPARQL-based audit layer can be used to inspect these records. SPARQL supports graph-pattern queries over RDF information sources [26], while Comunica provides a modular framework for querying heterogeneous Linked Data interfaces, including federated querying scenarios [29]. Auditing can consequently become part of the data-sharing workflow rather than remaining an external and predominantly manual activity [3, 29].

2.5. Positioning of the Proposed Framework

Existing work highlights complementary requirements for decentralized and provider-controlled storage [12], sovereign governance of provider–consumer exchanges [2, 6], machine-readable policy representation and evaluation [1, 25], and semantic provenance for accountability and auditability [3, 28]. The reviewed literature addresses these requirements through partly separate lines of work. Solid provides decentralized storage, identity, and resource-level authorization.

[1] https://www.w3.org/TR/odrl-model/
[2] https://www.w3.org/TR/sparql11-query/
[3] https://www.w3.org/RDF/

IDS and related data-space approaches provide concepts for sovereignty, provider–consumer agreements, and governed exchange. ODRL provides a model for expressing machine-readable policies, while provenance and Semantic Web standards support structured transfer records and semantic inspection.

The proposed framework builds on these complementary contributions by connecting them within one operational workflow. Its specific focus is the integration of Solid-based storage and identity, IDS-inspired transfer requests and agreements, pre-transfer policy evaluation, transfer-point enforcement, RDF-based provenance generation, and SPARQL-based auditing. This integration is intended to support decentralized data transfers that are not only technically possible, but also associated with explicit participants, purposes, agreements, policy decisions, and queryable audit records.

3. Proposed System

The proposed system defines a governed data-sharing workflow. Solid provides the decentralized storage and identity foundation, IDS-inspired concepts structure the provider–consumer exchange, and Semantic Web technologies represent requests, agreements, policies, transfer activities, provenance, and audit metadata.

Figure 1 presents both the conceptual workflow and its implementation in the prototype. Figure 1a shows the eight methodological phases, while Fig. 1b shows the numbered interface modules through which these phases are demonstrated. Some closely related phases are grouped within the same prototype tab. In particular, transfer-request creation and agreement generation are implemented through the *Contract Negotiation* module, policy definition and evaluation are implemented through the *Policy Enforcement* module, and the two governed transfer alternatives are represented by the PUSH and PULL modules.

The workflow consists of eight phases.

3.1. Participant Identity and Identity Setup

Phase 1 establishes the identities of the data provider and data consumer. The input is the Solid WebID of each participant. Both participants authenticate through the prototype using Solid identity mechanisms, after which the application establishes an authenticated Solid session and retrieves the corresponding WebID.

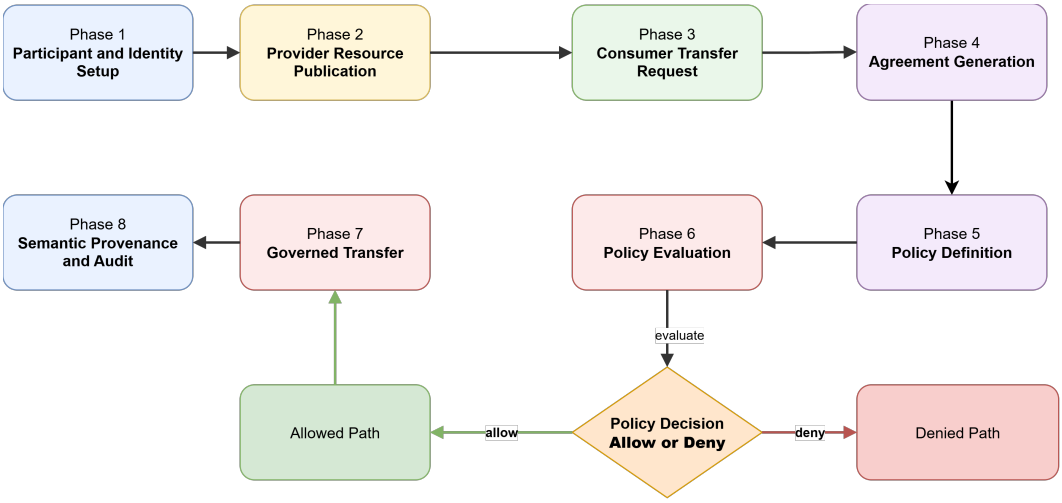
In the demonstration, this phase is implemented through 1 *Authentication*, shown in Fig. 1b.

The output of this phase is an authenticated Solid session and a verified participant WebID.

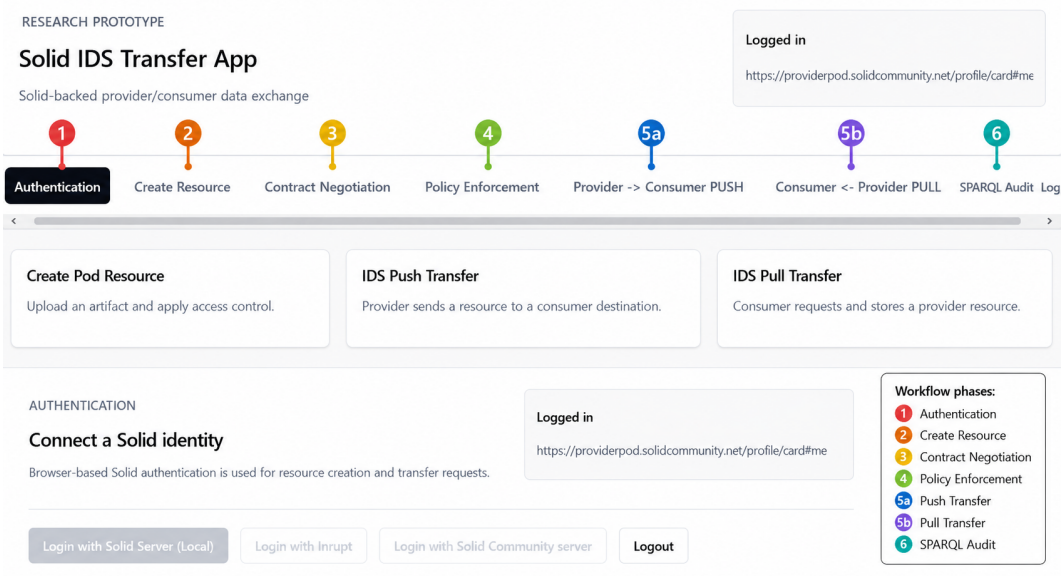
3.2. Provider Resource Publication

Phase 2 allows the authenticated provider to create the resource that will later be exchanged. As shown in Fig.[4], the provider specifies the target pod URL, content type, resource content, and access mode.

[4] https://tinyurl.com/38fa8drw



(a) Conceptual workflow of the proposed methodology.



(b) Numbered prototype modules used to demonstrate the workflow.

Figure 1: Proposed governed data-sharing methodology

The resource is stored in the provider's Solid pod using the active authenticated session. The output is a provider-controlled resource URL.

3.3. Consumer Transfer Request

Phase 3 initiates the contract-negotiation process and does not itself trigger data transfer. As shown in Fig.[5], the consumer specifies the provider resource, its WebID and destination pod, the requested transfer mode, the declared purpose, and any proposed usage conditions. The resulting RDF/Turtle request is stored at the specified request URL and made accessible to the provider for review in Phase 4.

3.4. Provider Review and Agreement Generation

Using the same 3 *Contract Negotiation* module, the provider loads the request and verifies its resource, consumer, purpose, transfer mode, and proposed conditions. The provider may approve or reject the request and, when appropriate, grant the consumer read access. An approval produces a machine-readable agreement linking the request, provider, consumer, resource, and accepted conditions, while a rejection is stored as a decision record. The resulting agreement or rejection record is referenced by the later policy evaluation and transfer phases.

3.5. Policy Evaluation

Using the same 4 *Policy Enforcement* module, the application evaluates the transfer context before accessing or copying the provider resource. The requester WebID, declared purpose, transfer mode, redistribution status, resource URLs, and agreement reference are compared with the policy conditions. The output is an explicit *allow* or *deny* decision. An allow decision permits the transfer phase to proceed, whereas a deny decision blocks data movement. The interface therefore connects machine-readable policy representation directly to pre-transfer decision making rather than treating the policy as passive documentation.

3.6. Governed Transfer

Phase 7 executes the data movement only after a positive policy decision. Its inputs are the provider resource URL, consumer destination URL, authenticated Solid session, transfer mode, agreement reference, and allow decision produced during Phase 6. The prototype supports two governed transfer patterns: 5a *Provider* $\rightarrow$ *Consumer PUSH* and

5b *Consumer* $\leftarrow$ *Provider PULL*, as shown in Fig. 1b. In the PUSH workflow, the provider initiates the transfer and sends the approved resource to a destination controlled by the consumer. In the PULL workflow, the consumer retrieves the provider resource and stores a copy in its own pod. Both patterns are executed through authenticated Solid operations. The output is a copied resource stored at the approved destination URL. The transfer is validated through the HTTP responses returned during the provider-side read and consumer-side write operations. For example, a successful PULL transfer produces a successful GET response for the provider resource and a successful PUT response for the consumer destination. When Phase 6 returns a deny decision, neither the PUSH nor the PULL operation is executed.

[5] https://tinyurl.com/3b2576sd

3.7. Phase 8: Semantic Provenance and Audit

After a transfer, the application generates RDF/Turtle metadata describing the source and destination resources, participants, transfer mode, agreement, policy decision, status, and timestamp. This record represents the transfer as a traceable activity and provides the input for later inspection. The 6 *SPARQL Audit* module shown in Fig.[6] loads one or more metadata resources and supports preset or custom SPARQL queries. Transfers can be filtered by participant, resource, policy keyword, transfer mode, or date. The resulting table enables users to inspect recorded transfers.

Compared with standard Solid, the framework reuses pods, WebID-based identity, authentication, and resource access, while adding negotiation, policy evaluation, transfer provenance, and audit functions. Compared with IDS, it reuses provider–consumer governance, agreement, policy, and transfer concepts, but differs by retaining Solid pods and representing workflow records as lightweight RDF resources rather than implementing a complete connector stack.

The complete source code, local execution instructions, workflow sequence diagram, and an end-to-end video demonstration of the prototype are publicly available in the project repository.[7]

4. Conclusion & Future work

This paper presented a Solid-based, IDS-inspired framework for secure, policy-aware, and semantically auditable data sharing. The framework combines Solid pods for decentralized storage, IDS-inspired provider–consumer governance, machine-readable policy evaluation, RDF-based provenance, and SPARQL-based auditing. Its main contribution is an operational workflow that connects data requests, agreements, policy decisions, controlled transfer, provenance generation, and audit inspection within a single prototype.

The current implementation remains a proof of concept. Policy enforcement is performed only before transfer and does not provide continuous control over how the consumer uses the data afterward. The policy model supports a limited set of conditions and does not yet include advanced conflict resolution or richer reasoning. The prototype has also been evaluated only in a controlled demonstration and still requires larger-scale performance testing, stronger security analysis, and validation in real inter-organizational settings. Future work will address richer policy reasoning, post-transfer usage control, stronger security guarantees, IDS connector interoperability, performance evaluation, and validation in real organizational environments. Overall, the framework provides a practical foundation for governed, policy-aware, and auditable decentralized data exchange.

Acknowledgments

This work was supported by the Alliance of Universities in Europe – European Universities Linking Society and Technology (EULiST), a European University Alliance officially recognized and funded by the European Union under the European Universities Initiative.

[6] https://tinyurl.com/y5k4j8uc

[7] https://github.com/MRNaqvi/Solid-and-International-dataspaces-protocol-solid-ids-transfer-app-

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT & CodeX (OpenAI) for grammar and spelling checks, literature search, and rephrasing to improve the clarity and readability of the manuscript. After using this tool, the authors reviewed and edited all suggested content as needed and take full responsibility for the content of the publication.

References

[1] T. Dam, A. Krimbacher, S. Neumaier, Policy patterns for usage control in data spaces, in: Proceedings of the 19th International Conference on Semantic Systems, SEMANTiCS 2023 Workshops, volume 3510 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.

[2] B. Otto, M. Jarke, Designing a multi-sided data platform: Findings from the international data spaces case, Electronic Markets 29 (2019) 561–580. doi:10.1007/s12525-019-00362-x.

[3] T. Pasquier, J. Singh, J. Powles, D. Eyers, M. Seltzer, J. Bacon, Data provenance to audit compliance with privacy policy in the Internet of Things, Personal and Ubiquitous Computing 22 (2018) 333–344. doi:10.1007/s00779-017-1067-4.

[4] S. Capadisli, T. Berners-Lee, K. Kjernsmo, Solid Protocol, Draft Community Group Report Version 0.11.0, W3C Solid Community Group, 2024. URL: https://solidproject.org/TR/2024/protocol-20240512.

[5] A. Sambra, H. Story, T. Berners-Lee, WebID 1.0: Web Identity and Discovery, W3C Editor's Draft, World Wide Web Consortium, 2014. URL: https://www.w3.org/2005/Incubator/webid/spec/identity/.

[6] B. Otto, S. Steinbuss, A. Teuscher, S. Lohmann, S. Bader, et al., IDS Reference Architecture Model: Version 3.0, Technical Report, International Data Spaces Association, Berlin, 2019. doi:10.24406/publica-fhg-299836.

[7] F. Hörandner, S. Krenn, A. Migliavacca, F. Thiemer, B. Zwattendorfer, CREDENTIAL: A framework for privacy-preserving cloud-based data sharing, in: 2016 11th International Conference on Availability, Reliability and Security (ARES), IEEE, 2016, pp. 742–749. doi:10.1109/ARES.2016.79.

[8] M. Crosas, The dataverse network: An open-source application for sharing, discovering and preserving data, D-Lib Magazine 17 (2011). doi:10.1045/january2011-crosas.

[9] D. Hong, Y. Zhang, J. Luo, C. Liu, M. Xu, Z. Suo, Research on centralized data-sharing model based on master data management, MATEC Web of Conferences 139 (2017) 00195. doi:10.1051/matecconf/201713900195.

[10] P. Sawadogo, J. Darmont, On data lake architectures and metadata management, Journal of Intelligent Information Systems 56 (2021) 97–120. doi:10.1007/s10844-020-00608-7.

[11] A. Nambiar, D. Mundra, An overview of data warehouse and data lake in modern enterprise data management, Big Data and Cognitive Computing 6 (2022) 132. doi:10.3390/bdcc6040132.

[12] A. V. Sambra, E. Mansour, S. Hawke, M. Zereba, N. Greco, A. Ghanem, D. Zagidulin, A. Aboulnaga, T. Berners-Lee, Solid: A platform for decentralized social applications based

on linked data, in: Proceedings of the 25th International Conference Companion on World Wide Web, International World Wide Web Conferences Steering Committee, 2016, pp. 223–226. doi:10.1145/2872518.2890529.

[13] J. Werbrouck, P. Pauwels, J. Beetz, L. Van Berlo, Towards a decentralised common data environment using linked building data and the solid ecosystem, in: 36th CIB W78 Conference, 2019, pp. 113–123.

[14] R. Dedecker, W. Slabbinck, J. Wright, P. Hochstenbach, P. Colpaert, R. Verborgh, What's in a pod? a knowledge graph interpretation for the solid ecosystem, in: 6th Workshop on Storing, Querying and Benchmarking Knowledge Graphs (QuWeDa) at ISWC 2022, volume 3279, CEUR, 2022, pp. 81–96.

[15] S. Speicher, J. Arwe, A. Malhotra, Linked Data Platform 1.0, W3C Recommendation, World Wide Web Consortium, 2015. URL: https://www.w3.org/TR/ldp/.

[16] A. V. Sambra, A. Guy, S. Capadisli, N. Greco, Building decentralized applications for the social web, in: Companion Proceedings of the 25th International Conference on World Wide Web, WWW '16 Companion, Association for Computing Machinery, New York, NY, USA, 2016, pp. 1033–1034. URL: https://doi.org/10.1145/2872518.2891060. doi:10.1145/2872518.2891060.

[17] E. Mansour, A. V. Sambra, S. Hawke, M. Zereba, S. Capadisli, A. Ghanem, A. Aboulnaga, T. Berners-Lee, A demonstration of the solid platform for social web applications, in: Companion Proceedings of the 25th International Conference on World Wide Web, WWW '16 Companion, Association for Computing Machinery, New York, NY, USA, 2016, pp. 223–226. URL: https://doi.org/10.1145/2872518.2890529. doi:10.1145/2872518.2890529.

[18] A. Coburn, E. Pavlik, D. Zagidulin, Solid-OIDC, Community Group Specification Version 0.1.0, W3C Solid Community Group, 2022. URL: https://solidproject.org/TR/2022/oidc-20220328.

[19] T. Berners-Lee, H. Story, S. Capadisli, C. Braun, Web Access Control, Community Group Report Version 1.0.0, W3C Solid Community Group, 2024. URL: https://solidproject.org/TR/2024/wac-20240512.

[20] M. Bosquet, Access Control Policy, Editor's Draft, W3C Solid Community Group, 2022. URL: https://solid.github.io/authorization-panel/acp-specification/, work in progress.

[21] International Data Spaces Association, Dataspace Protocol 2024-1, Technical Specification, International Data Spaces Association, 2024. URL: https://docs.internationaldataspaces.org/ids-knowledgebase/dataspace-protocol, release candidate 2024-1.

[22] J. Werbrouck, P. Pauwels, J. Beetz, R. Verborgh, E. Mannens, Consolid: A federated ecosystem for heterogeneous multi-stakeholder projects, Semantic Web 15 (2024) 429–460.

[23] S. Meckler, R. Dorsch, D. Henselmann, A. Harth, The web and linked data as a solid foundation for dataspaces, in: Companion Proceedings of the ACM Web Conference 2023, WWW '23 Companion, Association for Computing Machinery, New York, NY, USA, 2023, pp. 1440–1446. URL: https://doi.org/10.1145/3543873.3587616. doi:10.1145/3543873.3587616.

[24] C. H.-J. Braun, Y. Wuwang, Z. Wang, X. Hou, T. Käfer, Towards using the solid protocol for data transport in international data spaces (IDS), in: Proceedings of the Third International Workshop on Semantics in Dataspaces, volume 4007 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2025. URL: https://publikationen.bibliothek.kit.edu/1000186449. doi:10.5445/IR

/1000186449.
[25] R. Iannella, S. Villata, ODRL Information Model 2.2, W3C Recommendation, World Wide Web Consortium, 2018. URL: https://www.w3.org/TR/odrl-model/.
[26] S. Harris, A. Seaborne, SPARQL 1.1 Query Language, W3C Recommendation, World Wide Web Consortium, 2013. URL: https://www.w3.org/TR/sparql11-query/.
[27] P. Buneman, S. Khanna, W.-C. Tan, Why and where: A characterization of data provenance, in: Database Theory – ICDT 2001, volume 1973 of *Lecture Notes in Computer Science*, Springer, 2001, pp. 316–330. doi:10.1007/3-540-44503-X_20.
[28] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, W3C Recommendation, World Wide Web Consortium, 2013. URL: https://www.w3.org/TR/prov-o/.
[29] J. Van Herwegen, R. Taelman, M. Vander Sande, R. Verborgh, Demonstration of Comunica, a web framework for querying heterogeneous Linked Data interfaces, in: 17th ISWC, 2018.