The Need for Usage Control in Decentralized and Federated Ecosystems

Wout Slabbinck[1]

[1]IDLab, Departement of Electronics and Information Systems, Ghent University - imec, Belgium

Abstract

In recent years, Personal Data Stores and dataspaces have emerged with the aim to provide data control and data sovereignty for their users. A method often adopted is to allow users to declare their own access control policies, containing rules for controlling who can perform which action upon what resource. Unfortunately, this does not take into account the full usage of the data, i.e. what happens after given action is performed. Another method, adopted by few, is to allow users to declare usage control policies (UCPs) which, when enforced correctly, allow to govern the full usage lifecycle of data. However, solutions supporting UCPs lack formalisation, thus reliance on specific implementations is required. We aim to provide said formalism to an interoperable usage control policy standard through a practical implementation of the UCP evaluator. Furthermore, we aim to integrate this evaluator into a usage control framework with the objective of working with any kind of resource server. Ultimately, the hope is that this research provides a foundation for future research on UCP enforcement alignment with legislation and improving user understanding of UCP implications.

Keywords

Decentralization, Federation, Usage Control, Personal Data Stores, Dataspaces

1. Introduction

In the mid-2000s, Web 2.0 emerged, which shifted the focus of the Web to centralized platforms [1]. Compared to the original World Wide Web, which was more consumer-oriented due to the technical expertise required and inherently decentralized by design, it then became easier for general end-users to write on the Web. The centralized model came at the cost of delegating the responsibility of storing, sharing and governance of this data to these new platforms. This shift necessitated users placing trust in upholding these responsibilities by these platforms. However, recent years have shown that this trust was misplaced as there have been multiple cases of data breaches and data misuse in centralized platforms [2, 3]. To combat challenges created by centralization, federated and decentralized systems have emerged that empower data sovereignty and control over data [4, 5, 6]. The latter entails that the producer of data has control over how the data is used and who has access while adhering to the law[1]. In the decentralized spectrum, a plethora of Personal Data Store (PDS) technologies have been created such that end-users achieve control over their data [7]. On the federated side, the idea of dataspaces as a solution that aims to gain data sovereignty.

The common ground of these initiatives allows us to extract several requirements that are important to reach the aim of data control and sovereignty. (i) First, there is a need for enforcing mechanisms pertaining to the use of the data, rather than the access [8, 9]. With existing systems users can give access to a certain resource, but they should be able to say that is only for the purpose of verification, only possible for one week and that the data should be removed from their local machine after one week. (ii) A language is required to capture all those constraints in one policy. However, since it is under the control of the user, it means that they must be able to express it. It cannot be expected that everybody is technically gifted, so there is a need for inclusive mechanisms to express usage control

Proceedings of the Doctoral Consortium at ISWC 2024, co-located with the 23rd International Semantic Web Conference (ISWC 2024)

wout.slabbinck@ugent.be (W. Slabbinck)
https://woutslabbinck.com (W. Slabbinck)
0000-0002-3287-7312 (W. Slabbinck)

[1]definition of data sovereignty: https://www.techtarget.com/whatis/definition/data-sovereignty

policies [10]. (iii) To achieve full data sovereignty, it is not enough that end-users govern the usage of data. They might declare policies that are actually illegal. There is thus a need to verify that the expressed policies conform to the laws of the country [11, 12]. (iv) Finally, we are designing this to be used in a decentralized or federated environment. As a result, every node in the environment must be able to interpret and enforce policies the same. Achieving this requires formalism and interoperability regarding both the policy language and the evaluation of the policy in monitoring scenarios and during access requests [13, 14]. (v) While all the previously mentioned arguments show the necessity to fundamentally achieve full control over data and data sovereignty, they lack expressing the scalability of such an ideal system. It does not allow to answer questions about how fast should a revocation of consent be handled. In other words how long can end-users be sure that they are still conforming to the policies?

These requirements, correspondingly summarized as usage control over access control enforcement, understandable policies, legal alignment, formalisation and interoperability, and scalability, are the necessary (but not automatically sufficient) conditions identified through the research on data sovereignty and control in decentralized and federated systems. In Section 2, this research on the state-of-the-art is further elaborated. In my Ph.D. I will focus on tackling challenges (i), (iv), and (v) and follow up on the developments of (ii) and (iii) so that they can be integrated with my work. To further narrow down what I will contribute to these challenges, I have formulated two research questions in Section 3. The practical execution plan to provide answers is detailed in Sections 4 and 5. My progress and insights thus far are summarized in Section 6. Finally, I conclude by giving an overview and future prospects regarding data sovereignty and control through decentralized or federated usage control.

2. Related Work

Achieving data sovereignty and data control requires usage control, which extends traditional access control by incorporating deontic concepts [8]. Next to permissions and prohibitions, deontic concepts include obligations, a party must do something before or after a certain event, and dispensations, a party is excluded from an obligation. As a consequence, it is not only important to verify compliance at an access request, but also during the whole lifecycle of data use. Indicating the need for mechanisms to facilitate the continuous monitoring and auditing of data use [9]. In their survey on usage control, Akaichi and Kirrane [15] concluded the importance of semantic technologies for encoding usage control policies regarding interoperability. Furthermore, they highlight the need for research on verification and testing tools and the importance of benchmarks with regard to usage control solutions. There exist multiple semantic usage control languages, though, to the best of my knowledge, there is only one which is also a W3C Recommendation, more specifically the Open Digital Rights Language (ODRL) [16]. In an ODRL policy, permission, prohibitions and obligations can be defined as rules. In each rule, the involved parties, possible actions and target resources can be encoded. Constraints facilitate further refinements such as intended purpose, cardinality and temporal limitations. A point of critique stated by multiple researchers is the lack of formalism to this date [13, 14, 15, 17, 18, 19, 20]. However, a W3C Community Group (CG) is formed to create formal semantics for ODRL and create an ODRL evaluator for an access control and policy monitoring scenario[2].

2.1. Usage Control in Decentralized Systems

Personal Data Stores (PDSs), a decentralized solution, allow users to store, manage and control their data [7, 21]. In the Semantic Web, the most widely adopted technology is the Solid Protocol[3] [22]. Solid is built on existing W3C Recommendations and has as its goal, next to control over data, to promote data reuse which is achieved through the decoupling of data and applications. Data is stored in resources on a *pod*. On this level of granularity, access control rules can be configured using either the Web Access

[2]ODRL Formal Semantics W3C CG: https://w3c.github.io/odrl/formal-semantics/
[3]The Solid protocol: https://solidproject.org/TR/protocol

Control (WAC) or the Access Control Protocol (ACP). These specifications only allow to specify access control and not usage control. Thus there is a need in Solid to define proper usage control policies such that the data sovereignty goal can be achieved. Different papers expressed this need, though to this day, no consensus in the Solid CG has been reached on how to fulfill this need [19, 23, 24, 25, 26].

2.2. Usage Control in Federated Systems

In recent years, dataspace initiatives have gained significant attention. One of the core principles of these federated systems is data sovereignty [4]. The International Data Space Association (IDSA) and Gaia-X are the two leading initiatives which have been proposing architectural solutions to standardize and build dataspaces. A recurring concept is the *dataspaces connector* as an entry point within dataspaces. They facilitate data exchange which is safeguarded by data usage contracts, ensuring data sovereignty. This agreement defines the usage of the data and is the result of negotiation between data consumers and providers prior to the exchange. To the best of our knowledge, they are either expressed or have been heavily influenced by ODRL policies[4][5][6] [27, 28, 29]. Current dataspace connectors have been implementing their interpretation of ODRL[7][8] due to the lack of formalisation of how to evaluate ODRL.

Recently, technical attempts were undertaken to combine both PDS and dataspace technologies [30, 31].

3. Problem statement and Research Questions

Achieving data sovereignty and control includes the **enforcement** of the use of the data and the necessity of decentralized or federated systems that can **interoperate** as one network. As such, the objective of my research is centred around enforceable and interoperable usage control systems. The following research questions aim to contribute to the usage control domain for decentralized and federated systems.

RQ-I **Accountability** This entails transparency by proving that the data use was correct, is currently correct and will be correct in the future. The first aspect of correctness is the fact that the usage of data can be expressed as a policy in a formal and interoperable language. Interoperability in the sense that every actor within a network must interpret the policies in the same manner and formality in the sense that the rules and constraints of those policies prove that certain action was allowed at a given time. Furthermore, tracking data usage over time allows for the verification and auditing of past actions, and the dynamicity of future actions can be reasoned upon. How would such a proof look like? What are the necessary and sufficient conditions for conformance to a policy which entails both the request, the policy itself but also the state of the world?

RQ-II **Scalability** Optimisation of the theoretical foundation that allows estimating the needs on computation and storage, and the limits of how long policies can be trusted. To the best of my knowledge, there are no benchmarks on usage control policy evaluations. Therefore it is hard to estimate what the requirements and behaviour are for usage control enforcement systems. Performance metrics for usage control systems are vital for ensuring real-time data usage guarantees. This would allow us to determine bounds for the duration of a policy revocation or confirmation, which in turn ensures that every actor uses data under the constraints of the policies at any given time based on their available knowledge. What are the key performance metrics for a benchmark that evaluates a real-time usage control policy enforcement system?

[4] IDS Usage Control Policies (v6): https://international-data-spaces-association.github.io/DataspaceConnector/Documentation/v6/UsageControl
[5] IDS USage Control Contracts https://github.com/International-Data-Spaces-Association/IDS-G/tree/main/UsageControl/Contract
[6] Gaia-X and contracts: https://gaia-x.eu/news-press/gaia-x-and-contracts/
[7] Prometheus-X Data Space Connector (Gaia-X): https://github.com/Prometheus-X-association/dataspace-connector
[8] MYDATA technologies (enforcing ODRL policies, used by IDSA): https://www.mydata-control.de/

How can we devise scalable simulations for usage control enforcement scenarios that accurately reflect real-world scenarios? How do these simulations compare with usage control on the Web itself, that is do the results reliably translate to Web-scale environments?

4. Methodology

The starting point for describing my methodology depends on revisiting the state-of-the-art as they both research questions rely on an interoperable usage control policy model. To this end, a decision has been made to continue with ODRL as a policy model.

Answering RQ-I partially aligns with the work completed over the years by the ODRL Formal Semantics CG: the initiative of building an *ODRL Evaluator*. The idea is that based on a set of policies, performed and requested actions it provides proofs regarding whether usage of data was and is permitted and whether all obligations have been met. My approach to contribute consists of the creation of a test suite for validating the correctness of an ODRL Evaluator. Each test case is a tuple consisting of the input and the expected output of such evaluator. To this day, there exists no formal representation for certain input elements and the output. This indicates that the first task for creating a test suite is to establish models for the performed and requested actions, core attributes relating to constraints such as cardinality, and the conformance report.

To tackle the benchmark aspect for RQ-II, further analysis of usage control systems in the decentralized and federated systems is required. The evaluator could then either be included in an existing modular continuous usage control enforcement system or, based on the knowledge gained, such a system must be designed and created to integrate the evaluator. Finally, a multitude of use cases where usage control enforcement is required such that the system as a whole can be tested.

5. Evaluation

In the previous section, the method for answering RQ-I consists of creating multiple test cases with as output a conformance report. As ODRL contains three classes for rules representing the deontic concepts of permission, prohibition and obligation, a sub-task is to make a bottom-up test-suite for an ODRL permission evaluation. It allows us to define the necessary and sufficient conditions for this single deontic concept. Using this knowledge, a revision of the ODRL Evaluator can be created to support prohibitions. Any limitations to the conformance report can be picked up to adapt the model. Finally, a last iteration will be executed for obligations. This iterative approach, where the intermediate version of the conformance report is continuously refined, ensures that all aspects of the proof are included. This allows us to provide an answer to RQ-I and additionally results in an implementation for an ODRL Evaluator and accompanying a reference test suite that can be reused for the implementation of other evaluators.

To evaluate a usage control system in a network, it must not only be possible to measure all the different components in that system, but the communication must also be evaluated. For this, the plan is to have different scenarios. First, a baseline must be created against which each other different scenario will be evaluated. Then for each next scenario, variety will be introduced in one domain. At this time, the following scenarios are envisioned where a range is introduced to vary the number of nodes in the network, the number of constraints per policy, the number of policies per node, and the number of revocations. After measuring everything, it should be possible to answer which of those metrics are important and thus be able to answer RQ-II. However, if the data is not conclusive, new scenarios need to be investigated.

6. Preliminary Results

The results of my current research span both decentralized and federated systems. Initially, my focus was on Solid and its interoperability. It began with research on usage control agnostic approaches

for storing real-time data over Solid pods [32]. This was under the assumption that services could inherently have the appropriate access rights, recognize that they possess these rights, and encode the purpose of data use. The next phase evolved creating an ecosystem over these pods using Web agents [33], operating under the same assumptions regarding access rights and data use. However, this assumption proved incorrect, revealing amongst others the limitations of Solid current access control protocols and the tight coupling to its interface [34].

To introduce usage control enforcement in Solid, I initially attempted to do so externally using personal agents [25]. However, I discovered that not all types of usage control policies, such as the purpose of data use, can be enforced as a service over the Solid protocol. Given that full usage control via external services is unfeasible, I have shifted my focus to exploring its feasibility within a Solid server while still adhering to the Solid specification. By implementing the Request Flow from the Solid Open ID Connect primer[9], we can separate a Solid server into an Authorization Server[10] following the User-Managed Access (UMA) specification[11] and a Resource server [35]. This approach seems promising but requires more research.

My research on usage control for federated systems began with a project centered around sharing sensitive data among multiple actors within a dataspace. This led to the design of a Usage Control Framework [20], which each actor can incorporate to continuously enforce usage control policies. While this approach envisions continuous enforcement, it currently lacks implementation, preventing evaluation of performance and scalability.

Continuation of these research efforts relies on the proof of data use at any point in time, emphasizing the necessity for an ODRL Evaluator and addressing RQ-I. Since the Usage Control Framework provides a blueprint for the creation of a system that can be benchmarked it might provide answers for RQ-II. Once the framework is in place, I plan to integrate it into the UMA server. Since the UMA specification is agnostic of the Resource Server it integrates with, this should enable continuous usage control enforcement for both Solid (through integration with the Community Solid Server [36]) and dataspaces (by integration into an IDS connector).

7. Conclusion

I have elaborated the research plan for my Ph.D. where I aim to contribute to the formalisation of ODRL evaluations and building the modular continuous usage control enforcement system.

Future work entails legal and usability to be combined with the UMA server making it possible to manage policies, and change policies that are compliant to the legislation of the country. This will make it possible to have true data sovereignty and control facilitated in both decentralized and federated systems.

Acknowledgments

Supported by SolidLab Vlaanderen (Flemish Government, EWI and RRF project VV023/10). I thank my supervisers Pieter Colpaert and Ruben Verborgh. Furthermore, I thank Beatriz Estevez, Julian Rojas, Ruben Dedecker and Jos De Roo for the discussions.

References

[1] T. O'reilly, What is web 2.0, " O'Reilly Media, Inc.", 2009.

[2] H. Berghel, Malice domestic: The Cambridge analytica dystopia, Computer 51 (2018) 84–89. URL: http://www.berghel.net/col-edit/out-of-band/may-18/oob_5-18.pdf, publisher: IEEE Computer Society.

[9] Solid-OIDC primer: https://solidproject.org/TR/oidc-primer#request-flow

[10] UMA server integrated with a Solid server: https://github.com/SolidLabResearch/user-managed-access/

[11] User-Managed Access specification: https://docs.kantarainitiative.org/uma/wg/rec-oauth-uma-grant-2.0.html

[3] S. Zuboff, The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power, 2018.
[4] J. Theissen-Lipp, M. Kocher, C. Lange, S. Decker, A. Paulus, A. Pomp, E. Curry, Semantics in Dataspaces: Origin and Future Directions, in: Companion Proceedings of the ACM Web Conference 2023, ACM, Austin TX USA, 2023, pp. 1504–1507. URL: https://dl.acm.org/doi/10.1145/3543873.3587689. doi:10.1145/3543873.3587689.
[5] C. Meurisch, B. Bayrak, M. Mühlhäuser, Privacy-preserving AI Services Through Data Decentralization, in: Proceedings of The Web Conference 2020, ACM, Taipei Taiwan, 2020, pp. 190–200. URL: https://dl.acm.org/doi/10.1145/3366423.3380106. doi:10.1145/3366423.3380106.
[6] R. Verborgh, Re-Decentralizing the Web, For Good This Time, in: O. Seneviratne, J. Hendler (Eds.), Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web, ACM, 2023, pp. 215–230. URL: https://ruben.verborgh.org/articles/redecentralizing-the-web/. doi:10.1145/3591366.3591385.
[7] K. U. Fallatah, M. Barhamgi, C. Perera, Personal data stores (PDS): A review, Sensors 23 (2023) 1477. Publisher: MDPI.
[8] J. Park, R. Sandhu, Towards usage control models: beyond traditional access control, in: Proceedings of the seventh ACM symposium on Access control models and technologies, SACMAT '02, Association for Computing Machinery, New York, NY, USA, 2002, pp. 57–64. URL: https://doi.org/10.1145/507711.507722. doi:10.1145/507711.507722.
[9] A. Pretschner, M. Hilty, F. Schütz, C. Schaefer, T. Walter, Usage Control Enforcement: Present and Future, IEEE Security & Privacy Magazine 6 (2008) 44–53. URL: http://ieeexplore.ieee.org/document/4588229/. doi:10.1109/MSP.2008.101.
[10] J. Wright, B. Esteves, R. Zhao, Me want cookie! Towards automated and transparent data governance on the Web, 2024. URL: http://arxiv.org/abs/2408.09071. doi:10.48550/arXiv.2408.09071, arXiv:2408.09071.
[11] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data Privacy Vocabulary (DPV) – Version 2, 2024. URL: http://arxiv.org/abs/2404.13426. doi:10.48550/arXiv.2404.13426, arXiv:2404.13426.
[12] P. A. Bonatti, S. Kirrane, I. M. Petrova, L. Sauro, Machine Understandable Policies and GDPR Compliance Checking, KI - Künstliche Intelligenz 34 (2020) 303–315. URL: https://doi.org/10.1007/s13218-020-00677-4. doi:10.1007/s13218-020-00677-4.
[13] S. Steyskal, A. Polleres, Towards Formal Semantics for ODRL Policies, in: N. Bassiliades, G. Gottlob, F. Sadri, A. Paschke, D. Roman (Eds.), Rule Technologies: Foundations, Tools, and Applications, Springer International Publishing, Cham, 2015, pp. 360–375. doi:10.1007/978-3-319-21542-6_23.
[14] N. Fornara, M. Colombetti, Using semantic web technologies and production rules for reasoning on obligations, permissions, and prohibitions, Ai Communications 32 (2019) 319–334. URL: https://content.iospress.com/articles/ai-communications/aic190617, publisher: IOS Press.
[15] I. Akaichi, S. Kirrane, Usage Control Specification, Enforcement, and Robustness: A Survey, 2022. URL: http://arxiv.org/abs/2203.04800. doi:10.48550/arXiv.2203.04800, arXiv:2203.04800 [cs].
[16] W3C Working Group, The Open Digital Rights Language (ODRL), 2018. URL: https://www.w3.org/TR/odrl-model/.
[17] M. G. Kebede, G. Sileno, T. Van Engers, A critical reflection on ODRL, in: International Workshop on AI Approaches to the Complexity of Legal Systems, Springer, 2018, pp. 48–61.
[18] M. De Vos, S. Kirrane, J. Padget, K. Satoh, ODRL policy modelling and compliance checking, in: Rules and Reasoning: Third International Joint Conference, RuleML+ RR 2019, Bolzano, Italy, September 16–19, 2019, Proceedings 3, Springer, 2019, pp. 36–51.
[19] R. Zhao, J. Zhao, Perennial semantic data terms of use for decentralized web (2024). Publisher: Association for Computing Machinery.
[20] I. Akaichi, W. Slabbinck, J. A. Rojas Meléndez, C. Van Gheluwe, G. Bozzi, P. Colpaert, R. Verborgh, S. Kirrane, Interoperable and Continuous Usage Control Enforcement in Dataspaces, in: Proceedings of the Second International Workshop on Semantics in Dataspaces, 2024.

URL: https://raw.githubusercontent.com/woutslabbinck/papers/main/2024/Interoperable_and_Continuous_Usage_Control_Enforcement_in_Dataspaces.pdf.

[21] S. Verbrugge, F. Vannieuwenborg, M. Van der Wee, D. Colle, R. Taelman, R. Verborgh, Towards a personal data vault society: an interplay between technological and business perspectives, in: 2021 60th FITCE Communication Days Congress for ICT Professionals: Industrial Data–Cloud, Low Latency and Privacy (FITCE), IEEE, 2021, pp. 1–6.

[22] 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 (2016) 16.

[23] B. Esteves, V. Rodríguez-Doncel, H. J. Pandit, N. Mondada, P. McBennett, Using the ODRL profile for access control for solid pod resource governance, in: European Semantic Web Conference, Springer, 2022, pp. 16–20.

[24] L. Debackere, P. Colpaert, R. Taelman, R. Verborgh, A Policy-Oriented Architecture for Enforcing Consent in Solid, in: Proceedings of the 2nd International Workshop on Consent Management in Online Services, Networks and Things, Association for Computing Machinery, 2022, pp. 516–524. URL: https://dl.acm.org/doi/pdf/10.1145/3487553.3524630. doi:10.1145/3487553.3524630.

[25] W. Slabbinck, J. A. Rojas Meléndez, R. Verborgh, Enforcing Usage Control Policies in Solid Using a Rule-Based Software Agent, in: Proceedings of the Second Solid Symposium, 2024. URL: https://pod.woutslabbinck.com/WIP/24-03-26_SoSy2024___Solid_Agent_for_UCP.pdf.

[26] G. Havur, M. Vander Sande, S. Kirrane, Greater control and transparency in personal data processing, in: Proceedings of the 6th International Conference on Information Systems Security and Privacy, SciTePress, 2020, pp. 655–662.

[27] L. Nagel, D. Lycklama, Design Principles for Data Spaces - Position Paper, Technical Report, Zenodo, 2021. URL: https://zenodo.org/records/5105744. doi:10.5281/zenodo.5105744.

[28] V. Siska, V. Karagiannis, M. Drobics, Building a Dataspace: Technical Overview, 2023. URL: https://www.gaia-x.at/wp-content/uploads/2023/04/WhitepaperGaiaX.pdf.

[29] H. Drees, D. O. Kubitza, J. Lipp, S. Pretzsch, C. S. Langdon, Mobility data space–first implementation and business opportunities, in: ITS World Congress, 2021. URL: https://www.researchgate.net/profile/Johannes-Theissen-Lipp/publication/351519610_Mobility_Data_Space_-_First_Implementation_and_Business_Opportunities/links/610101882bf3553b29174ee6/Mobility-Data-Space-First-Implementation-and-Business-Opportunities.pdf.

[30] 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, ACM, Austin TX USA, 2023, pp. 1440–1446. URL: https://dl.acm.org/doi/10.1145/3543873.3587616. doi:10.1145/3543873.3587616.

[31] S. Schmid, D. Schraudner, A. Harth, The Rights Delegation Proxy: An Approach for Delegations in the Solid Dataspace, in: Proceedings of the Second International Workshop on Semantics in Dataspaces, 2024.

[32] W. Slabbinck, R. Dedecker, S. Vasireddy, R. Verborgh, P. Colpaert, Linked Data Event Streams in Solid containers, in: Proceedings of the 8th Workshop on Managing the Evolution and Preservation of the Data Web, 2022. URL: https://raw.githubusercontent.com/woutslabbinck/papers/main/2022/Linked_Data_Event_Streams_in_Solid_containers.pdf.

[33] W. Slabbinck, R. Dedecker, J. A. Rojas Meléndez, R. Verborgh, A Rule-Based Software Agent on Top of Personal Data Stores, in: Proceedings of the 22nd International Semantic Web Conference: Posters, Demos, and Industry Tracks, 2023.

[34] 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: M. Saleem, A.-C. Ngonga Ngomo (Eds.), Proceedings of the 6th Workshop on Storing, Querying and Benchmarking Knowledge Graphs, volume 3279 of *CEUR Workshop Proceedings*, 2022, pp. 81–96. URL: https://solidlabresearch.github.io/WhatsInAPod/, iSSN: 1613-0073.

[35] W. Termont, R. Dedecker, W. Slabbinck, B. Esteves, B. D. Meester, R. Verborgh, From Resource Control to Digital Trust with User-Managed Access, 2024. URL: http://arxiv.org/abs/2411.05622.

doi:10.48550/arXiv.2411.05622, arXiv:2411.05622.

[36] J. Van Herwegen, R. Verborgh, The Community Solid Server: Supporting research & development in an evolving ecosystem, Semantic Web (2024) 1–15. URL: https://www.medra.org/servlet/aliasResolver?alias=iospress&doi=10.3233/SW-243726. doi:10.3233/SW-243726.