Towards Consistent Policy Enforcement in Dataspaces

Julia Pampus[1][a] and Maritta Heisel[2][b]
[1]*Fraunhofer Institute for Software and Systems Engineering ISST, Dortmund, Germany*
[2]*University of Duisburg-Essen, Duisburg, Germany*

Keywords: Data Spaces, Policies, Requirements Engineering, Data Sovereignty, Design Framework.

Abstract: Data sovereignty refers to the autonomy and self-determination of organisations when it comes to sharing data. The focus, thereby, is on the data usage conditions that are expressed as policies. Current research explores the structure of these policies, the processes related to data offerings and policy negotiations, and their enforcement using access and usage control methods. However, there is still a lack of a consistent and comprehensive understanding of data sovereignty among data-sharing participants across various system landscapes. First, we discuss the reasons for this issue and its significance in the context of dataspaces, then take a position. We present a model-based design framework encompassing different environments for describing sovereign data sharing. To conclude our contribution, we outline an approach for systematically eliciting and analysing data usage requirements, thus strengthening interoperability and trust.

1 MOTIVATION

Dataspaces are the technical foundation for trusted, autonomous, and self-determined data sharing within data ecosystems (Jarke et al., 2019). Their adoption in the industrial sector has gained significant traction with the European strategy for Common European Data Spaces (European Commission, 2025). Today, already various dataspaces, such as Catena-X[1] and Eona-X[2], have been established. Current activities within these dataspaces address the design of shared governance models, the identification of new business models, and the exchange of data. By now, related concepts and technologies are transitioning from a research topic to an industrial standard. This evolution is especially promoted by open-source projects and standardisation activities (Noardo et al., 2024).

Technical policies are the key to ensuring data sovereignty (Hosseinzadeh et al., 2020). By defining, negotiating, and enforcing policies, all data sharing participants agree to ensure compliance with data usage conditions. Today, legal processes are applied to sanction companies for misuse and non-compliance with contractually determined regulations; technical processes should support and replace these within dataspaces (Otto et al., 2019). The dataspace connectors, i.e., data sharing agents, provide the basis for sharing and negotiating policies; the actual enforcement occurs at all points of the data lifecycle and value creation. In recent years, motivated by the idea of dataspaces and open data ecosystems, most of the central components have evolved towards a decentralised architecture. Therefore, the vision of a holistic policy enforcement necessarily involves many independent, technical and non-technical systems, comprising entire hardware and software stacks and human-centric processes. Consequently, *interoperability* is of utmost importance concerning policy exchange and enforcement. However, this is also the biggest challenge.

1.1 Levels of Interoperability

The advanced Levels of Conceptual Interoperability Model (LCIM) by Turnitsa comprises seven levels of interoperability in system engineering (Tolk et al., 2007): First, Level 0 (*No Interoperability*) covers systems in isolation without any interactions. Second, Level 1 (*Technical Interoperability*) defines an established communication using a common protocol. Next, Level 2 (*Syntactic Interoperability*) introduces a common format of information that is exchanged. Then, Level 3 (*Semantic Interoperability*) is reached when the interacting systems share the meaning of

[a] https://orcid.org/0000-0003-2309-6183
[b] https://orcid.org/0000-0002-3275-2819
[1]https://catena-x.net/ (Accessed: 2025-03-06)
[2]https://eona-x.eu/ (Accessed: 2025-03-06)

Pampus, J., Heisel and M.
Towards Consistent Policy Enforcement in Dataspaces.
DOI: 10.5220/0013569100003967
In *Proceedings of the 14th International Conference on Data Science, Technology and Applications (DATA 2025)*, pages 560-566
ISBN: 978-989-758-758-0; ISSN: 2184-285X

this information in a certain context. Level 4 (*Pragmatic Interoperability*) describes that the integrated systems are aware of each other's methods of processing information. Finally, Level 5 (*Dynamic Interoperability*) introduces the handling of state changes, and Level 6 (*Conceptual Interoperability*) is reached when the systems can switch contexts.

1.2 State of the Art

Current research and technical approaches consider different aspects of policies in dataspaces: their structure in data offerings, their negotiation as part of communication protocols, and their enforcement by implementing access and usage control.

Policy Syntax. In dataspaces, data offerings are represented as Data Catalog Vocabulary (DCAT) (W3C, 2024) datasets. The specified information model (Koen et al., 2025) describes the metadata of data, focusing primarily on how it can be accessed (referred to as 'distribution') and what data usage conditions apply (marked as 'hasPolicy'). These data usage conditions are expressed as Open Data Rights Language (ODRL) (W3C, 2018) policies. Unlike the definition of the data offerings, the information model for dataspaces does not prescribe the vocabulary of the shared data. However, in many dataspace projects, organisations have domain-specific requirements for managing data in particular formats, such as the Asset Administration Shell used in the manufacturing domain.

Policy Negotiation. During the negotiation process, two data sharing participants (data provider and consumer) agree on common usage conditions for shared datasets (Jung and Dörr, 2022). The Dataspace Protocol (Koen et al., 2025) specifies this process as a state machine that determines the message types and the permitted sequences of interactions. Both data sharing participants may initiate the process with a policy offer, negotiate their requirements according to the policy rules, and conclude with a policy agreement.

Policy Enforcement. The implementation of policy enforcement involves both access and usage control. While access control restricts who can access data, usage control takes it a step further by continuously monitoring and managing how that data is used (Schütte and Brost, 2018). In current dataspaces, we observe that access control mechanisms address most policies. For example, access to data or data offerings is often limited to individual organisations or validated members of the dataspace. Overall, the technologies used for policy enforcement are quite diverse. A typical architecture for designing policy engines, systems that intercepts data flows, validate policies, and apply rules, utilises the Extensible Access Control Markup Language (XACML) (OASIS, 2013) (Jung and Dörr, 2022). Additionally, there is a growing trend in attribute-based access control towards self-sovereign identity management through the use of decentralised identifiers (Čučko and Turkanović, 2021). A first step towards usage control is realised with traceability and observability approaches that take effect where access control ends (Akaichi and Kirrane, 2022).

1.3 Problem Statement

Intervenability, security, and interoperability are essential for end-to-end data sovereignty (Pampus and Heisel, 2024). However, at the time of writing, holistic policy enforcement is not implemented in any dataspace (Hellmeier et al., 2023).

In current dataspace initiatives, such as Catena-X, we observe approaches to reduce the scope for semantic interpretation of policies by addressing syntactic equality. Consequently, some syntactical expressions that are semantically correct are excluded. One example of different policy expressions with the same semantics is the following rules: 'Data may be shared with companies based in the EU' versus 'Data may not be shared with companies based in non-EU countries'. From a technical perspective, the first statement is a permission along with an equality operator (see Figure 1), and the second one is a prohibition with an inequality operator (see Figure 2). The meaning of both policy expressions is equal and, thus, must be enforced similarly. A policy validation method must handle this semantic equality despite syntactic differences.

To address formalisation, the ODRL Community Group (2025) is working on harmonising the semantic interpretation of their policy language, thus addressing the aforementioned example. However, dataspace specifications increase the complexity of semantic interpretation by allowing external contexts in addition to ODRL. As described in Section 1.2, the Dataspace Protocol defines the vocabulary as an ODRL profile to express data offerings and attached policies. Its application allows and requires the domain-specific adaptation of the provided vocabulary by extending it with a formalisation of permissible policies.

Other research addresses interoperability in policy enforcement by focusing on compliance mod-

```
{
 "@context": "http://www.w3.org/ns/odrl.jsonld",
 "@type": "Offer",
 "uid": "http://example.org/rules/1",
 "permission": [{
  "target": "http://example.org/data/resource",
  "assigner": "http://example.org",
  "action": "use",
  "constraint": [{
   "leftOperand": "location",
   "operator": "eq",
   "rightOperand": "EU"
  }]
 }]
}
```

Figure 1: Sample Policy of Type *Permission*.

```
{
 "@context": "http://www.w3.org/ns/odrl.jsonld",
 "@type": "Offer",
 "uid": "http://example.org/rules/2",
 "prohibition": [{
  "target": "http://example.org/data/resource",
  "assigner": "http://example.org",
  "action": "use",
  "constraint": [{
   "leftOperand": "location",
   "operator": "neq",
   "rightOperand": "EU"
  }]
 }]
}
```

Figure 2: Sample Policy of Type *Prohibition*.

els and technical approaches to systematically evaluate ODRL policies and test policy implementations (Slabbinck et al., 2025). Thereby, the assumption is that a technical verification could be established across dataspaces and, thus, solve the problem of missing interoperability as a basis for the end-to-end establishment of data sovereignty. For example, the Eclipse Dataspace Components[3] provide a framework with a modular policy engine that supports the development of dataspace technologies that implement any policy. Multiple implementations could be verified against a test suite to prove their conformity.

The approach of specifying the semantic interpretation of ODRL can certainly be transferred to other languages and contexts. Referring to the LCIM (cf. Section 1.1), we assume that current approaches of implementing policy enforcement in dataspaces target interoperability level 2: Technical interoperability is supported with the Dataspace Protocol and syntactic interoperability with DCAT and ODRL (cf. Section 1.2). In accordance with the findings of (Tolk et al., 2007), **we state that interoperable policy enforcement within dataspaces cannot be achieved by focusing solely on the syntax and semantics of a policy language. To achieve interoperability level 3 and beyond, we must incorporate contextual knowledge into the system design**. A system cannot evaluate and interpret policies without context, especially not if policy enforcement should be language- and platform-independent. Therefore, our work is guided by the research question (RQ): *How can an aligned semantic and syntax-independent interpretation of policies be achieved in data sharing?* We address this RQ by defining a requirements engineering (RE) process adapted to the key aspects of sovereign data sharing within dataspaces.

2 APPROACH

To define every aspect of sovereign data sharing and being able to design interoperable software systems accordingly, we suggest a model-based RE approach.

2.1 Designing Sovereign Data Sharing

To illustrate the holistic approach of policy management and enforcement and relevant key aspects with direct and indirect influences, we present a model-based design framework for sovereign data sharing. The structure of our framework relies on work by Nilsson (1999). He focuses on business modelling and its integration with system modelling and presents a generic architecture for modelling the mutual influences of environment, enterprise, and information systems. This architecture consists of four elements: intentions, actions, resources, and rules. As shown in Figure 3, the *intentions* represent the reasons for performed actions; the *action* describes who executes which operation; the *resources* describe the target of the action; and the *rules* describe the underlying conditions.

We introduce our design framework for sovereign data sharing in Figure 4. It represents all four types previously described, both in its basic structure and within the individual elements. Since this work focuses on interoperability, Figure 4 shows only the detailed structure of the *system environment*. It represents the resources, which are entirely controlled by the business environment, influenced by the regulatory environment, and created and consumed by the contextual environment. The systems (actions) are data sharing agents and/or data processing systems as

[3]https://projects.eclipse.org/projects/technology.edc (Accessed: 2025-03-06)

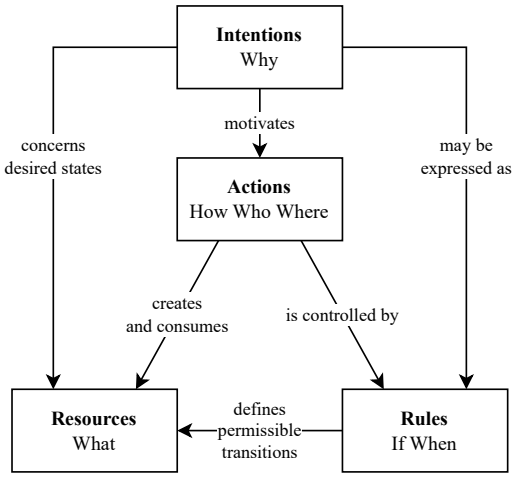

Figure 3: Model Architecture according to Nilsson (1999).

the executing environment in the data sharing context. These systems create and consume data (resources) and are restricted and controlled by policies (rules). These policies represent the defined permissible transitions.

Next, the *business environment* expresses the organisation-internal processes as motivation for data sharing use cases (contextual environment). In organisations, business decisions, represented by processes and models, are motivated by defined strategies. Implementing these strategies creates and consumes business resources, e.g., capital or staff, and is restricted and controlled by business guidelines, including compliance frameworks and standards. Overall, the business environment guides regulations restricting and controlling the data sharing process and concerns modelling the system environment.

Then, the *regulatory environment* presents a data governance motivated by business guidelines and creates and consumes legally binding contractual agreements on executing data sharing. Those are restricted by regulations such as laws defined by legal entities. The regulatory environment defines permissible transitions of the system environment and explicitly restricts the policies.

Last, the *contextual environment* includes actors and use case-specific conditions. The contextual environment is motivated by the business environment, restricted and controlled by the regulatory environment, and creates and consumes the system environment.

In summary, our design framework depicts which conceptual elements impact the design of sovereign data sharing and how they influence each other. Referring back to our motivation (cf. Section 1.2), we can conclude that policies make up an essential but only one aspect of sovereign data sharing and are subject to both direct (system, data) and indirect influences (regulatory, business, and contextual environments).

2.2 Requirements Engineering Process

Considering our design framework for sovereign data sharing, a comprehensive RE process must consider different types of requirements: business requirements, regulatory requirements, and system requirements in a given context (data sharing use case). We outline their interrelations in Figure 5. The business requirements are based on the regulations and lead to the system requirements. We can derive a system design for the data provider's and consumer's data sharing agents from these system requirements.

A structured RE process must consider their dependencies and progress through the types of requirements layer by layer, aiming for interoperability and aggregating the requirements of multiple stakeholders (not restricted to two).

2.2.1 Stakeholders

In Section 1, we have motivated the technical processes around policies to complement the business processes and that, currently, many existing processes are still human-centred. Undergoing an RE process also includes humans, stakeholders respectively. We illustrate the different roles and elements (system, data, policies) they address with their requirements in Figure 6. Referring to the system environment in Figure 4, we consider the data providers and consumers deeply involved in describing the system, data, and policies. As one unique role, the data sharing participant acting as data rights holder concentrates on the data and policies. In addition, service providers, e.g., dataspace operators, play a significant role in the design of the system enforcing policies. Finally, governing bodies, e.g., a Dataspace Governance Authority[4], may predefine rule sets and, thereby, policies that need to be adhered to within a governed dataspace.

2.2.2 Scoping

Referring to Figure 4, we consider the system environment as the *solution domain*, whereas the business, regulatory, and contextual environment form the *problem domain*. The goal of a structured RE process is to systematically analyse the problem domain to define the solution domain (data, policies, systems). In this process, first, all stakeholders must agree on common information and processes, achieving technical and syntactic interoperability (cf. Section 1.1). This

[4]https://dssc.eu/space/bv15e/ (Accessed: 2025-02-14)

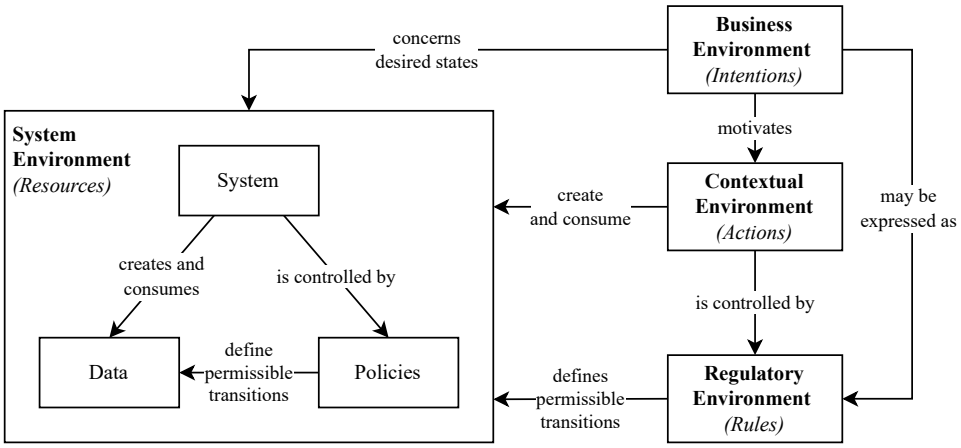

Figure 4: Model-based Design Framework for Sovereign Data Sharing.

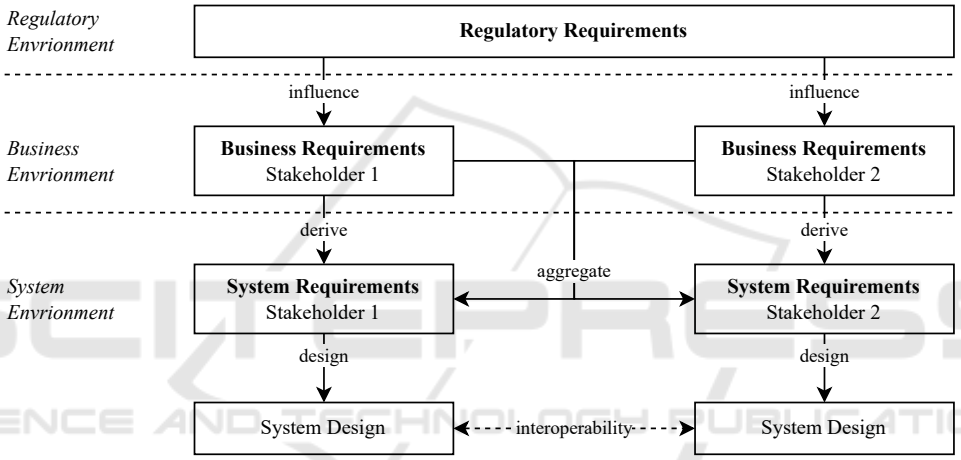

Figure 5: Composition of Requirements in Sovereign Data Sharing.

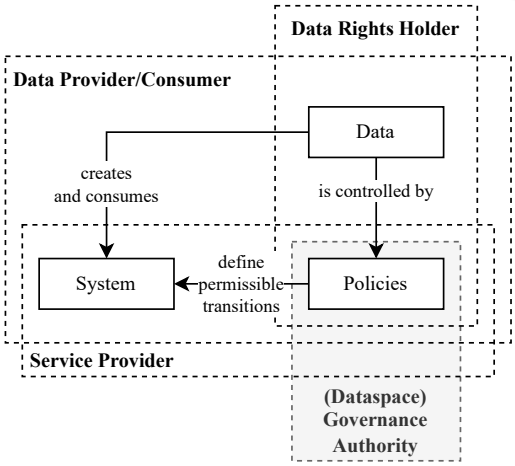

Figure 6: Types of Stakeholders.

agreement includes (1) the identification of the actors involved, (2) the definition of the data to be shared, including its format, protocol, and the access systems, and (3) the specification of the attached metadata, including the data usage conditions (Pampus and Heisel, 2025a).

However, the definition of data usage conditions needs to be achieved without referring to a policy language such as ODRL or a specific architecture such as XACML to allow for similar interpretations across different technology landscapes (addressing the problem described in Section 1.3). Therefore, we propose the development of a consistent policy interpretation according to two core aspects, the (1) point of validating a policy and (2) trust anchors.

- *Point of Validation:* Policy enforcement heavily depends on the location (local and temporal) of policy validation. For example, if data must be anonymised, this can be done on the data provider side, during data transfer, or on the data consumer side, during processing or storage. In the course of the RE process, all stakeholders must agree on a common derivation of validation scopes from

their requirements.

- *Trust Anchors:* Every attribute that is specified by a policy must be verifiable. For this purpose, often, third-party systems are integrated. For example, a policy may specify that data can only be accessed by participants of a specific dataspace. As part of the identity management, the membership proof would be obtained on the sender side and checked on the receiver side. A decentralised approach avoids relying on a central system. Instead, each data sharing participant can decide individually which third-party system or vendor they consider trustworthy for the validation of incoming information (not restricted to identities).

As illustrated in Figure 5 and indicated in Figure 6, some elicitation and analysis steps take place in isolation, while others are collaborative. For example, an essential part of the RE process for sovereign data sharing is aggregating requirements from different stakeholders, thus resolving potential conflicts. Finally, the resulting requirements may lead to sovereignty-specific software features of a data sharing system (Pampus and Heisel, 2025b).

3 DISCUSSION & FUTURE WORK

Addressing our RQ, a systematic RE process facilitates interoperable policy enforcement by enabling (1) a syntactic alignment of policies and (2) semantic equivalence by means of describing the environment (points of validation, trust anchors), regardless of the underlying policy language and architecture. This forms an essential basis for *conceptual interoperability* (cf. Section 1.1). A model-based approach helps understanding the key mechanisms of establishing sovereignty in data sharing, gathering required contextual information, and making them processable and reusable.

In general, such an RE process can be used for designing data sharing systems, but also supporting the onboarding of new participating organisations in a dataspace. Overall, a pattern-based approach allows for reusing requirements of all types in various data sharing use cases. In decentralised ecosystems, implementation is nevertheless challenging: In the simplest case, a system would need to map syntactic rules and their semantics. For this purpose, a centralised or decentralised policy registry could provide information about policy validation points and trust anchors.

The definition of shared trust anchors forms an essential part of trust between data sharing participants. In this work, the focus of our design framework and the presented RE process has been on establishing interoperability. However, data sovereignty is also primarily about *trust* (Lohmöller et al., 2022; Hellmeier et al., 2023). For this reason, it is essential to identify which other aspects are part of a trust model and which technical interfaces need to be designed accordingly. The results of this elaboration can then be used to derive additional requirements.

ACKNOWLEDGEMENTS

This work was partially supported by the German Federal Ministry for Economic Affairs and Climate Action (funding number: 13IK040F).

REFERENCES

Akaichi, I. and Kirrane, S. (2022). Usage Control Specification, Enforcement, and Robustness: A Survey.

Čučko, Š. and Turkanović, M. (2021). Decentralized and self-sovereign identity: Systematic mapping study. *IEEe Access*, 9:139009–139027.

European Commission (2025). Common European Data Spaces. https://digital-strategy.ec.europa.eu/en/policies/data-spaces.

Hellmeier, M., Pampus, J., Qarawlus, H., and Howar, F. (2023). Implementing Data Sovereignty: Requirements & Challenges from Practice. In *Proceedings of the 18th International Conference on Availability, Reliability and Security*, pages 1–9.

Hosseinzadeh, A., Eitel, A., and Jung, C. (2020). A Systematic Approach toward Extracting Technically Enforceable Policies from Data Usage Control Requirements. In *Proceedings of the 6th International Conference on Information Systems Security and Privacy*, pages 397–405.

Jarke, M., Otto, B., and Ram, S. (2019). Data Sovereignty and Data Space Ecosystems. *Business & Information Systems Engineering*, 61(5):549–550.

Jung, C. and Dörr, J. (2022). Data Usage Control. In Otto, B., ten Hompel, M., and Wrobel, S., editors, *Designing Data Spaces*, pages 129–146. Springer International Publishing, Cham.

Koen, P., Kollenstart, M., Marino, J., Pampus, J., Turkmayali, A., Steinbuss, S., and Weiß, A. (2025). Dataspace Protocol 2025-1-RC1.

Lohmöller, J., Pennekamp, J., Matzutt, R., and Wehrle, K. (2022). On the need for strong sovereignty in data ecosystems. In *DEco@ VLDB*, pages 51–63.

Nilsson, B. E. (1999). On Why to Model What and How: Concepts and Architecture for Change. In *Perspectives on Business Modelling*, pages 269–303. Springer, Berlin, Heidelberg.

Noardo, F., Atkinson, R., Bastin, L., Maso, J., Simonis, I., Villar, A., Voidrot, M.-F., and Zaborowski, P. (2024). Standards for Data Space Building Blocks. *Remote Sensing*, 16(20):3824.

OASIS (2013). eXtensible Access Control Markup Language (XACML) Version 3.0.

ODRL Community Group (2025). ODRL Formal Semantics.

Otto, B., Steinbuss, S., Teuscher, A., and Lohmann, S. (2019). IDS Reference Architecture Model.

Pampus, J. and Heisel, M. (2024). An Empirical Examination of the Technical Aspects of Data Sovereignty. In *Proceedings of the 19th International Conference on Software Technologies*, pages 112–122.

Pampus, J. and Heisel, M. (2025a). Feature-Oriented Requirements Analysis for Sovereign Data Sharing. Accepted for publication.

Pampus, J. and Heisel, M. (2025b). Pattern-based Requirements Elicitation for Sovereign Data Sharing. *Procedia Computer Science*, 254:147–156.

Schütte, J. and Brost, G. S. (2018). LUCON: Data Flow Control for Message-Based IoT Systems. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, pages 289–299.

Slabbinck, W., Rojas, J. A., Esteves, B., Colpaert, P., and Ruben, V. (2025). Interoperable Interpretation and Evaluation of ODRL Policies.

Tolk, A., Diallo, S. Y., and Turnitsa, C. D. (2007). Applying the levels of conceptual interoperability model in support of integratability, interoperability, and composability for system-of-systems engineering. *Journal of Systems, Cybernetics, and Informatics*, 5(5).

W3C (2018). ODRL Vocabulary & Expression 2.2.

W3C (2024). Data Catalog Vocabulary (DCAT) - Version 3.