Design Features for Data Trustee Selection in Data Spaces

Michael Steinert[a], Daniel Tebernum[b] and Marius Hupperz
Fraunhofer Institute for Software and Systems Engineering ISST, Dortmund, Germany

Keywords: Design Features, Data Spaces, Data Trustee, Design Science Research.

Abstract: As the world becomes increasingly digital, data is becoming a critical resource. When used effectively, it can lead to more accurate forecasts, process optimization, and the creation of innovative business models. The necessary data is often distributed across multiple organizations, and its full value can only be realized through shared collaboration. Data spaces provide organizations with a platform for sovereign and secure data sharing. To enable legally secure data sharing and ensure compliance with regulations, data trustees play a critical role as trusted intermediaries. However, choosing a suitable data trustee that meets the needs of the participants who want to share data with each other is difficult. Our study seeks to elucidate the process by which participants of a data space can choose an appropriate data trustee. To this purpose, we have implemented a whitelist approach. We report on the results of our design science research project, which includes design features to facilitate the integration of our whitelist approach into different data space instantiations. Potential shortcomings were identified and addressed during an expert workshop. By providing verified design knowledge, we help practitioners in the data space community to incorporate the concept of how to select the most appropriate data trustee.

1 INTRODUCTION

This paper describes the design features (DFs) needed for the selection of a data trustee in a data sharing scenario. Therefore, we utilize the Eclipse Dataspace Components[1] (EDC) as connectors for data spaces. The proposed data trustee whitelist extension allows each EDC-based data space connector to maintain a whitelist of trusted data trustees to enable data sharing with other actors trusting the same data trustee.

The development of the data trustee whitelist was driven by the need for a reliable mechanism to manage and control access to data, ensuring compliance with agreed terms during data exchange. This necessity arose from insights gained through discussions with data space experts, highlighting the challenges in existing access and usage control frameworks.

We aim to enable data sharing scenarios utilizing data trustees, like in automotive (Caruso, 2024), mobility (BMDV, 2024), or agriculture (Azkan et al., 2022). In such scenarios, the primary goal of the data trustee is to establish trust (Schinke et al., 2023). Therefore, a data trustee needs to be an independent third party, that coordinates usage rights and can receive, own, and store data, as well as to offer and send data to potential data users (Doan et al., 2022; Ghayyur et al., 2020; Maaß, 2022; Potoczny-Jones et al., 2019). As (Schinke et al., 2023) propose, that a data trustee can be part of a data space or even enable it in the first place. Therefore, we are driven by the following research question:

RQ: *What are design features that assist practitioners in the secure and sovereign selection process of finding a data trustee in a data space?*

To answer this research question, we conducted a Design Science Research (DSR) project in which we created a solution in the form of a *software artifact* and extracted fundamental design knowledge in the form of *DFs* from it.

The remainder of the paper is organized as follows. First, in Section 2, we establish the underlying knowledge required for comprehension. In Section 3, we outline the research design we employed to answer our research question. The

[a] https://orcid.org/0009-0008-3888-2092
[b] https://orcid.org/0000-0002-4772-9099
[1] https://github.com/eclipse-edc

Steinert, M., Tebernum, D. and Hupperz, M.
Design Features for Data Trustee Selection in Data Spaces.
DOI: 10.5220/0012851400003756

In *Proceedings of the 13th International Conference on Data Science, Technology and Applications (DATA 2024)*, pages 559-570
ISBN: 978-989-758-707-8; ISSN: 2184-285X

results are shown in Section 4 in the form of DFs and an EDC extension. Furthermore, in Section 5, we provide our findings in a use case based on a real data space project that focuses on large language model (LLM) training. Details and findings of our evaluation are covered in the following Section 6. In Section 7, we then review our findings and explore the theoretical and practical implications. In Section 8, we finally conclude our work and discuss limitations and future work.

2 BACKGROUND

Data spaces are a novel data management approach to collect large-scale heterogeneous data distributed over various data sources in different formats (Gieß et al., 2024).

Connectors enable peer-to-peer data transfer in a data space and offer storage and processing capabilities, as well as data access for other participants and usage control capabilities (Brost et al., 2018; Hupperz & Gieß, 2024; Munoz-Arcentales et al., 2019).

A data trustee is a natural or legal person or a business partnership that mediates access to data provided or held by data subjects in accordance with contractually agreed or legally prescribed data governance regulations (also) in the interests of third parties (Blankertz & Specht-Riemenschneider, 2021).

In summary, all concepts allow organizations to share their data in compliance with their terms of use, but from different standpoints. A data space is the infrastructure of an ecosystem, a connector is a technical gateway for participation in a data space for organizations, and a data trustee is a middleman in a specific data sharing scenario. In the following, we will describe these three concepts in more detail.

We also discuss various approaches to building trust, such as centralized trust authorities, peer-to-peer trust, and reputation-based trust. These are discussed further to compare their effectiveness and limitations.

2.1 Data Spaces

Data spaces are still a relatively novel infrastructure theory, where data providers and data users are enabled to share their data independently and data hegemony is denied (Braud et al., 2021; Otto & Jarke, 2019). The core competency of a data space is data sovereignty, which means that data providers can allow how their data is used, even if the data leaves their organizational boundaries. In addition to that, data spaces offer a variety of advantages, compared to traditional data sharing infrastructures:

- Increased interoperability and seamless data exchange.
- Enhanced data security and compliance with usage policies.
- Promotion of a collaborative ecosystem among diverse stakeholders.

In conclusion, data spaces are an infrastructure to bring data users and data providers together, without the need for a central keystone (Gelhaar & Otto, 2020). Data spaces are organized by various participants, storing their data decentral, while being able to query other participants' databases (Curry, 2020).

2.2 Connectors

Connectors are the gateways into a data space and enable peer-to-peer data transfer (Gieß et al., 2024). They need to be installed by data providers and data users and work as data storage, process engine, and enforcers of usage control (Munoz-Arcentales et al., 2019). Connectors can run various services to produce or consume data, as well as for management functionalities (Hupperz & Gieß, 2024). Currently, there are several organizations offering connectors for data spaces[2], with the EDC being the framework on which most of these connectors are built.

The EDC is the *de-facto* standard for creating data space connectors, providing a framework for sovereign, inter-organizational data sharing, utilizing the IDS Dataspace Protocol (DSP) and relevant protocols from the Gaia-X community (Eclipse Foundation, 2024). The EDC are designed to allow for implementation of different protocols and to be integrated into various ecosystems[3]. To establish the integration into other ecosystems the EDC allow to write extensions within the community to enable customization. Examples of extensions are AWS, Azure, HTTP, SQL, and storage extensions[4].

2.3 Data Trustees

Users can safely delegate the management of their data to a data trustee, rather than exposing it directly to products and services, as shown in Figure 1. Entrusted with this fiduciary duty, the data trustee is

[2] https://internationaldataspaces.org/data-connector-report/
[3] https://github.com/eclipse-edc
[4] https://github.com/eclipse-edc/Connector/tree/main/extensions

obligated to oversee the data with a primary focus on protecting and serving the user's best interests (Fürstenau et al., 2023).

Data trustees are responsible for the ethical handling of personal data, maintaining the integrity of the data rights or beneficial interests assigned by individual data subjects (Lechte et al., 2023). They manage the relationship with companies on behalf of users, ensuring that all data flows are in strict compliance with the permissions and expectations of users, and are compliant with legal and regulatory standards. This responsibility includes not only protective oversight, but also proactive engagement to manage, negotiate, and monitor the use of data to prevent misuse and maintain trust.

2.4 Trust Building Methods

In the following, various mechanisms for establishing trust will be discussed. Whitelists always play a central role in this context. By defining clear criteria for inclusion – whether through centralized verification, peer endorsement, or reputation scores – a whitelist provides a tangible reference for participants of a data space seeking trusted data trustees.

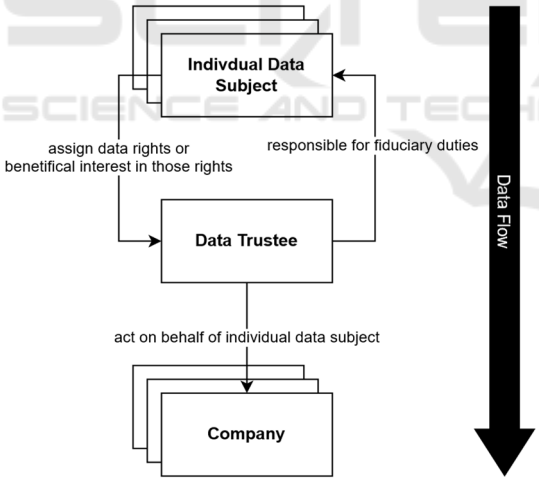

Figure 1: Data Trustee Data Flow.

2.4.1 Centralized Trust Authorities

A centralized authority, similar to a certification authority, would be responsible for verifying and approving participants in a data space before they can become trusted data trustees (Huang & Nicol, 2013). These approved entities would then be added to a central whitelist. This whitelist acts as a catalog of data trustees that have been recognized as trustworthy and capable of handling data responsibly.

While this approach is the easiest to implement, it has the following limitations. Relying on a single authority introduces risks such as a single point of failure – if the central authority's systems are compromised, the entire trust framework could be compromised. In addition, this model can create bottlenecks because the approval process can be slow, and it places a lot of power in the hands of the central authority, raising concerns about impartiality and potential abuse of power.

2.4.2 Peer-to-Peer Trust

In this approach, trust is established through direct interactions and endorsements between participants (i.e., peers) (Haoyang Che, 2006). In this model, participants in a data space recommend or vouch for the trustworthiness of data trustees based on previous exchanges or interactions. A whitelist in this context would be dynamically curated based on participant recommendations and interactions.

The limitations of this approach are that its effectiveness is highly dependent on the size of the data space and the activity level of the participants. There is also a risk of *echo chambers* or collusion, where participants endorse data trustees regardless of their actual trustworthiness. Establishing initial trust can be challenging without prior interactions, and the dynamic nature of the whitelist of trusted data trustees can introduce instability and unpredictability.

2.4.3 Reputation-Based Trust

In this approach, participants in a data space can assign scores or ratings to data trustees based on their behavior, history, and feedback from previous exchanges or interactions (Sänger et al., 2015). These ratings determine their trustworthiness. A whitelist in this approach would include data trustees that meet or exceed a certain reputation threshold, indicating that they are trusted to act as data trustees.

The limitations in this approach are like the peer-to-peer trust, where the trust can be manipulated through fraudulent feedback or coordinated efforts to artificially raise or lower scores. In addition, new or smaller data trustees may find it difficult to compete if the reputation system overly favors long-established reputations. Reputation metrics must be transparent and robust to minimize bias and ensure fairness.

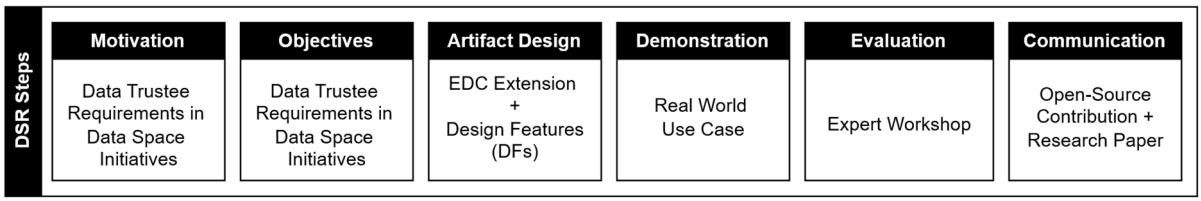

Figure 2: Design science research approach based on (Peffers et al., 2007).

3 DESIGN SCIENCE RESEARCH METHODOLOGY

To address our research question, we aim to develop the design knowledge necessary to adequately build a federated data trustee selection in data spaces. Consequently, we have investigated our research question in the context of a DSR project. Although DSR is originally rooted in business informatics (Österle, 2010), it is now also often used in software engineering (Barcellos et al., 2021; Wohlin & Runeson, 2021) and thus suitable to reach our objectives. The primary goal of DSR is the iterative development of an artifact, which "[...] is something created by people for some practical purpose" (Wieringa, 2014, p. 29), always taking into account relevance and scientific rigor.

(Gregor & Hevner, 2013) describe that artifacts can reach different stages of maturity. In our project, we decided to develop level 2 design knowledge in the form of DFs as well as to demonstrate them in a concrete level 1 instantiation in the form of a software artifact. Level 1 artifacts contain design knowledge in implicit form and are perfectly suited for demonstration purposes, although this knowledge is very context-specific and bound. Level 2 artifacts, on the other hand, represent more abstract operational principles. By providing DFs, we want to make it easier for practitioners to apply our findings to other technology stacks in the data space field.

The formulation of our DFs is based on existing work on design principles (Chandra et al., 2015), which offers abstract, prescriptive design knowledge in the form "if you want to achieve Y in situation Z, then something like action X will help" (van Aken, 2004, p. 227). However, our DFs are more concrete to enable practitioners to implement them. We will therefore also provide a more detailed rationale and description of what the technical implementation might look like.

The basic concept of (Hevner et al., 2004) has led to the development of more specific process models for conducting a DSR project (e.g. (William L Kuechler et al., 2021), (Peffers et al., 2007), or (Sein et al., 2011)). As we follow an objectivist/positivist approach and aim for an adaptation for daily use, we follow the recommendations of (John R. Venable et al., 2017) and structure our DSR project according to the model of (Peffers et al., 2007). We are conducting a single iteration in our DSR project. From a more global perspective, we follow the design knowledge development process of (Möller et al., 2020) by conducting responsive research. In short, this means that we first develop an artifact and then extract our design knowledge from this artifact.

We now explain in detail the steps of the process model proposed by (Peffers et al., 2007) and its application within our research.

The initial **motivation** for our DSR project was gained through a careful examination of existing data space initiatives, from which we derived the need for a federated data trustee integration in data spaces. Our findings clearly showed that there is currently no agreed-upon procedure for participants in data spaces to select a data trustee in a secure and sovereign manner (see Introduction).

Our **objectives** are drawn from a data space initiative, involving secure and sovereign data interchange for training LLMs. There, we are collaborating in partnership with other stakeholders, who are experienced in the conception and development of data spaces. The objectives include a working approach for selecting a data trustee in data spaces, as well as a practical implementation. Furthermore, participants valued the knowledge's reusability and ability to be applied to other data space technology stacks. As a result, an EDC extension should be developed, and the design knowledge gained should be documented in the form of DFs.

Our **artifact design** was carried out in two steps. To generate design knowledge for selecting a data trustee in data spaces, we first built a prototype in the form of an EDC extension, which was successfully tested and used as the foundation for extracting DFs. The process of extracting the DFs was characterized by continuous reflection by the authors. The resulting and final set of DFs as well as information about the EDC extension can be found in Section 4.

To **demonstrate** the DFs' practicability, we instantiated a minimal viable data space that is grounded in the LLM training use case of our data space initiative described earlier. In Section 5, we describe how the EDC extension and thus the DFs function in this context.

The DFs were **evaluated** at a workshop with experts. We based our evaluation strategy on the ontological expressiveness developed by (Recker et al., 2011). We are thus ensuring that no DFs are missing that cause significant phenomena in the environment (*deficit*), as well as that no additional DFs are triggering irrelevant phenomena (*excess*). We also ensure that a DF does not handle multiple phenomena (*redundancy*) and that no more than one DF addresses each phenomenon (*overload*) (visualization in Figure 4). This type of evaluation has already been applied to design principles by (Janiesch et al., 2020). We have adapted this method for our evaluation of DFs because, in our understanding, DFs are on the same level of design knowledge as design principles according to (Gregor & Hevner, 2013), but are just more concrete. See Section 6 for further information about the evaluation, participants, and its outcome.

Finally, we **communicated** our results by making the EDC extension publicly available and presented our DFs and analytical approach in this publication.

4 RESULTS

The proposed concept introduces a novel approach to establishing trust within dataspaces through a *Trusted Participants Whitelist Extension* for EDC data space connectors. Central to this concept is the creation and management of a whitelist of trusted data trustees, individually maintained by each connector. This system differs from traditional models by placing the establishment and verification of trust directly in the hands of data space participants, rather than relying on a centralized authority or the collective assessment of a community as described in Section 2.4. We will now describe in more detail the software artifact we developed.

4.1 Trusted Data Trustees Extension

The *Trusted Data Trustees Extension* for the EDC operationalizes a paradigm in which trust is explicitly determined by the direct actions and decisions of data space participants. This approach avoids reliance on a single authoritative body or collective community judgment, thereby mitigating the risks of centralized power, manipulation, and bias inherent in alternative approaches.

The whitelist is created dynamically based on criteria defined by the data space participants themselves. These include compliance with security standards, data protection compliance, existing ratings, and reliable handling of data in the past. Entities wishing to become trusted data trustees must meet these predefined criteria to ensure that trust is both earned and verifiable. This participant-centered approach offers several distinct advantages:

- By allowing participants to define the criteria for trust, the approach ensures that the whitelist reflects the specific needs and expectations of the data space community.
- Eliminating a central authority reduces vulnerability to single points of failure and avoids bottlenecks in the verification process.
- The whitelist can be updated and adapted as participant needs and external conditions evolve, ensuring its continued relevance and effectiveness.

The EDC extension is available as open-source software[5] and can be used in data spaces built using the EDC.

4.2 Design Features

In the following parts, we will present the DFs that we created. As described in Section 3, we applied a reflection process to extract the design knowledge from the concrete level 1 software artifact. The DFs thus identified were discussed among the authors. They were also evaluated through a workshop with experts (see Section 6). The final DFs can be found here. The template of (Chandra et al., 2015) will be utilized to describe the individual DFs. We will also give input on additional important information regarding the *how* and *why*.

DF1: Decentralized Trust Architecture. *To avoid a single point of failure and reduce reliance on central authorities (goal), while selecting a data trustee in a data space (context), adopt a decentralized or federated whitelist approach (mechanism).*

Each connector in the data space must maintain its whitelist according to externally defined rules. This adoption of a decentralized or federated architecture avoids single points of failure and reduces reliance on central authorities, mitigating both technical and antitrust risks (Deng et al., 2023).

[5] https://github.com/MichaelSteinert/edc-participants-whitelist-extension/tree/main

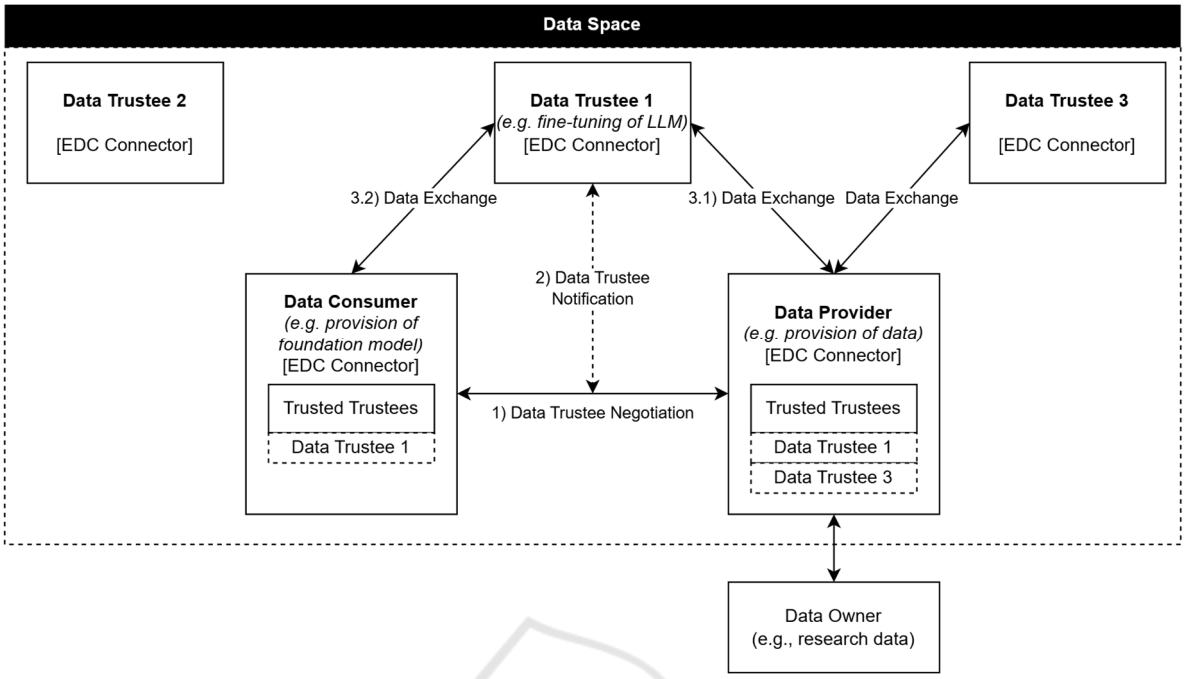

Figure 3: Operational Overview of Trusted Data Trustees Extension.

DF2: Hash-Based Integrity Checks. *To maintain the security and reliability of the whitelist (goal), while selecting a data trustee in a data space (context), its hash value should be computed and transmitted along with the whitelist itself (mechanism).*

This allows the integrity of the whitelist to be verified after transfer, ensuring its authenticity, and protecting against unauthorized changes. Regular integrity checks are essential to maintain confidence in the validity of the whitelist (Ferguson & Schneier, 2003).

DF3: Transparent Whitelist Management. *To promote trust and clarity in the data trustee selection process (goal) within a data space (context), the criteria and processes for whitelisting should be openly documented and accessible (mechanism).*

The process for adding or removing a data trustee to or from the whitelist must be articulated, detailing the necessary conditions, such as cost considerations, performance ratings, compliance with security standards, privacy policies, and required certifications. Transparency in whitelist management fosters an environment of trust, ensuring that all participants understand and can rely on the trustee selection process. This openness is essential to maintaining fair and defensible selection (Lucy et al., 2023).

DF4: Verifiable Credentials for Data Trustees. *To establish the authenticity and reliability of data trustees (goal) in a data space (context), trustees must present credentials that are verifiable and subject to independent audit (mechanism).*

These credentials[6] , which may include certify-cates of compliance, accreditation, or other forms of validation, must be verifiable by data space participants or third-party auditors. This level of verification ensures that the trustees meet the required standards of the whitelist and maintain trust within the data space. The ability to audit credentials is a fundamental aspect of a trusted and resilient data space (European Union Agency for Cybersecurity, 2022, p. 40).

DF5: Interoperability and Open Standards. *To ensure efficient integration and transfer of the whitelist across different data spaces (goal), within the multi-data space ecosystem (context), adherence to open standards and protocols is essential (mechanism).*

Where interoperability is required (Aliprandi, 2011), the whitelist must be carefully managed to ensure that only data trustees approved by the originating data space can be transferred and trusted in the receiving data space. This restriction preserves the integrity of the individual trust paradigms of the data spaces and prevents potential misinterpretation

[6] https://www.w3.org/TR/vc-data-model/

of trust relationships. Such governance is critical for maintaining privacy and consent in the exchange of data trustees between data spaces, thereby enhancing the security of multi-data space interactions.

5 DEMONSTRATION

The *Trusted Participants Whitelist Extension*, designed to strengthen the security and privacy of collaborative data sharing, has been instantiated into a real-world use case. This use case involves the process of fine-tuning foundation models using research data, a critical step in the advancement of machine learning and artificial intelligence. Specifically, the scenario described here involves the fine-tuning of an LLM under the guidance of a selected, mutually trusted data trustee. This selection process is a safeguard to ensure that the resulting fine-tuned foundation model, as well as the sensitive research data, both of which require significant computational resources, are accessible only to a mutually agreed upon and trusted data trustee. The goal is clear: to prevent valuable intellectual property from being compromised or misappropriated. The following prototype describes the workflow of the EDC extension we implemented and demonstrates its practicality in the real world.

As described in Figure 3, the core operational processes of the extension are illustrated, showing how trust in data trustees is decentralized, managed, and maintained within a data space. Each step is numbered to provide a clear reference to the detailed explanation that follows, capturing the intricacies of trust establishment in data trustee selection.

The implementation of the DFs described in Section 4.2 within the extension, which allows the whitelist to be dynamically adapted so that the inclusion of trusted data trustees can evolve with the changing landscape of the data space, is described below.

DF1: Decentralized Trust Architecture. The whitelist is based on a decentralized trust architecture, as evidenced by the multiple connectors – Data Trustee 1, Data Trustee 2, and Data Trustee 3 – that operate within the data space. DF1 mitigates single points of failure and antitrust risk, which is critical to maintaining a robust and equitable data sharing environment. If a data trustee fails or is removed from the data space, the next matching data trustee is selected from the whitelist. Redundancy is built into the whitelist by having multiple data trustees, allowing other data trustees to take over the responsibilities of the failed one without disrupting the requested data exchange.

DF2: Hash-based integrity checks. The hash (checksum) of the data trustee whitelist is calculated and sent along with the whitelist. This is done after the data trustee negotiation (Step 1: Data Trustee Negotiation) and before any data exchange (Steps 3.1 and 3.2: Data Exchange). The use of a checksum allows receivers to verify the integrity of the whitelist, maintaining trust and security throughout the data space. By using cryptographic hash functions, any tampering of the whitelist during transmission can be detected, as changes would alter the checksum, indicating a breach of integrity. This ensures that only authenticated and unaltered whitelists are accepted and processed by our extension.

DF3: Transparent Whitelist Management. Whitelist management is done with full transparency. The process for negotiating and adding data trustees is openly communicated (Step 1: Data Trustee Negotiation), as is the notification system that informs relevant participants, such as Data Trustee 1, of their status (Step 2: Data Trustee Notification). This ensures that everyone can participate in the trust environment with confidence and clear understanding of the selection. In addition, counterparties can see which data trustees are on each other's whitelists, increasing transparency. The matchmaking process to combine both whitelists is also transparent, ensuring that all parties are aware of how data trustees are matched and maintained within the whitelist.

DF4: Verifiable Credentials for Trustees. To participate in the data space as a trusted data trustee, they must provide verifiable credentials. These credentials are issued by a certification authority that is trusted within the data space. They are verified during the negotiation process (Step 1: Data Trustee Negotiation) and can be audited at any time by other participants or third parties, adding a layer of trust through transparency and accountability. The verification process uses authentication methods from the data space in use to ensure the authenticity and validity of the credentials presented. Regular updates and checks are required to maintain the integrity of the credentials over time, providing confidence in the security and compliance of the whitelist.

DF5: Interoperability and Open Standards. The whitelist supports interoperability and the use of open standards to allow different data spaces and connectors to integrate seamlessly. DF5 is key to ensuring that data exchange between Data Trustee 1 and other participants (Steps 3.1 and 3.2: Data Exchange) is efficient and compatible across different

data spaces. This interoperability is facilitated using the JSON data format for transfer within and between data spaces. In addition, DF5 uses and provides data models that meet the minimum requirements for data trustee selection, ensuring consistent and reliable data trustee integration into the whitelist. These models include attributes such as the name of the participant, its unique identification within the data space, the list of trusted participants, and the selection criteria used to qualify participants for the whitelist.

6 EVALUATION

As described in detail in Section 3, we will first provide an overview of the participants in our workshop. We will then report on our findings.

Four experts, whose roles and areas of expertise are outlined in Table 1, participated in the workshop. The experts were identified through personal contact. The low number of participants is owing to the general difficulty of acquiring developers and experts in the data space field. The session lasted 60 minutes and promoted a focused and in-depth analysis of the DFs in question.

Table 1: Overview of experts in the workshop.

Research Institution	Research Focus	Years of Data Space Experience
University	Data Intermediaries	4
Research Organization	Data Trustees	2
University	Business Models	4
Research Organization	Healthcare	1

To conduct a summative evaluation of the proposed DFs emerging from our implementation (see Section 4.2 Design Features), an online workshop was organized and facilitated via Microsoft Teams. Advance notice of the workshop was distributed to potential participants, along with an overview of the methodology from Figure 4 and the DFs to be discussed.

This workshop forms the basis of our evaluation, in which the data trustee experts critically reviewed and evaluated the success and relevance of the five proposed DFs. Their feedback is valuable in assessing the potential for the DFs to be effectively implemented in real-world data spaces with multiple data trustees.

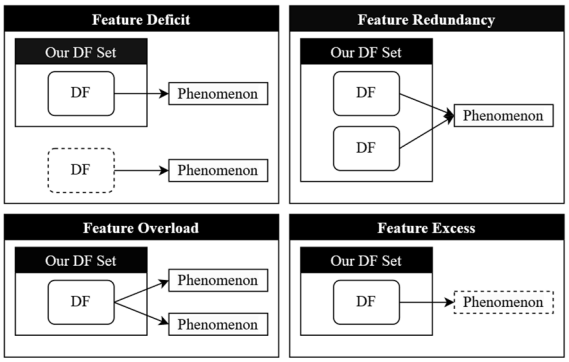

Figure 4: Ontological Expressiveness Assessment.

DF1: Decentralized Trust Architecture. The experts unanimously agreed that DF1 is a significant strength of the proposed concept. It mitigates central points of failure and reduces reliance on a central authority, which is consistent with the current push toward distributed systems in data management. There was no evidence of feature deficiencies, redundancies, overloads, or excess. The experts emphasized that this approach is well suited to modern data ecosystems, such as data spaces, that value resiliency and autonomy.

DF2: Hash-Based Integrity Checks (Modified). DF2 was originally identified as a *feature excess*. The experts were of the opinion that the original approach to integrity checking – based on the checksum of the data trustees on the whitelist and transferring of the individual checksums and, in the case of the counter-whitelist, hashing the data trustees and then comparing the checksums so as not to draw any conclusions about the selected data trustees as long as they are not on their own whitelist – did not meet user requirements and was in conflict with DF3. The modified DF2 now calculates and sends the checksum of the whitelist along with the whitelist itself to verify its integrity after transfer. This change was well received as it simplifies the process while maintaining the security and trustworthiness of the whitelist.

DF3: Transparent Whitelist Management. DF3 was well received, indicating clear processes and criteria for adding and removing data trustees from the whitelist. Experts appreciated the emphasis on transparency, which is paramount to building trust among participants in the data space. The simplicity and clarity of DF3 ensure that the selection process is accessible and understandable to all participants, a key aspect of effective governance and trust.

An additional insight provided by the experts on DF3 highlighted the automated functionality for selecting data trustees. This mechanism allows participants to automatically trust all data trustees that

meet the criteria they specify in DF3, eliminating the need to select specific data trustees. Such automation could prove beneficial to participants without a preference for specific trustees, ensuring a seamless and efficient trust establishment process across the data space.

DF4: Verifiable Credentials for Trustees. DF4 also received positive feedback. The experts recognized the need for verifiable credentials that allow data trustees to demonstrate their trustworthiness and adherence to standards in a verifiable manner. The strength of DF4 lies in its ability to provide an audit trail and assurance, fostering an environment where trust is not only stated, but also demonstrated. The experts also emphasized that verifiable credentials are an essential part of data spaces, realized by one or more trusted authorities within the data space that issue these credentials. This structure ensures that credentials are issued and recognized within an established data space, thereby increasing the reliability and integrity of whitelist verification.

DF5: Interoperability and Open Standards (Enhanced). The original design of DF5 was found to have a *feature deficit* with respect to ensuring the secure and appropriate transfer of whitelisted data trustees to other data spaces. To address this, a constraint was added that only those data trustees that have been approved by the previous data space should be transferred and used in an interoperable manner. This addition strengthens DF5 by ensuring that interoperability does not compromise the customized trust configurations of individual data spaces. In the case of an interoperable whitelist, it must be ensured that only the data trustees in the whitelist are transferred and used in another data space that the previous data space also offered. This is because if all previously trusted data trustees were added, the participants in the other data space could infer whom the whitelist previously trusted if their data trustee is not in the new data space. This inference could lead to unintended privacy breaches or strategic vulnerabilities, as participants could identify not only the relationships, but also the levels of trust or specific security measures associated with each data trustee, and potentially exploit this information to gain undue advantage or undermine existing trust.

7 DISCUSSION AND IMPLICATIONS

This research contributes to the field of data management within data spaces. By focusing on the DFs necessary for selecting a data trustee, it not only addresses a practical concern within data spaces, but also contributes to the theoretical knowledge base by formalizing these practices into design knowledge. The development of the *Trusted Participants Whitelist Extension*, based on the Eclipse Dataspace Components (EDC), provides an extension that establishes trust within these ecosystems, which is critical for enabling secure and efficient data sharing across different domains.

The use of connectors supports the secure and sovereign transfer and management of data across different stakeholders. The connector implemented by the EDC supports the technical implementation of our proposed DFs, thereby enabling data spaces to effectively manage data trustees. This practical application of design knowledge helps bridge the gap between theoretical constructs and real-world applicability, which is critical for the adoption and scalability of data space technologies.

The **theoretical implications** of our research extend the understanding of data spaces by showing that decentralized trust architectures can be effectively implemented using whitelists and connectors. The DFs derived from both the theoretical background and practical insights contribute to the academic discourse on data governance and management in data spaces. Our work underscores the importance of design science research in developing actionable and theoretically informed solutions that address complex problems in data sharing. Using a uniform formulation increases the DFs' reusability, allowing them to be employed in other study areas as well.

The **practical implications** of our research provide a structured, transparent method for selecting data trustees within dataspaces. By integrating the proposed DFs into the EDC, we introduce an extension that enhances trust in data sharing and data reuse by data trustees. This approach is directly aligned with the goals of the EU's Data Governance Act, which aims to increase trust in data sharing and promote data reuse. In addition, the creation of this extension allows each connector within a data space to independently manage a whitelist of trusted data trustees. This architecture significantly reduces reliance on central authorities, mitigating risks such as single points of failure and bias in the selection process. As a result, practitioners are empowered with a decentralized, transparent mechanism that provides greater flexibility and responsiveness to evolving regulatory and operational environments.

8 CONTRIBUTIONS, LIMITATIONS AND FUTURE WORK

In this paper, we have investigated how the selection of a suitable data trustee must look like in a data space. For this purpose, we conducted a DSR project in which we developed an EDC extension and extracted design knowledge in the form of DFs that a data space connector must implement. The DFs were validated within the context of an expert workshop and have already been successfully demonstrated in a data space initiative grounded in a real-world use case.

Next, we discuss the limitations of our work. First, we discuss the **internal validity** of the results. It should be mentioned that the authors underwent a reflection process to construct the DFs. The possibility of author bias exists within this approach. We cannot completely rule out the possibility that our expectations had a disproportionate impact on the DFs. In addition, the specialists included in the assessment session were chosen based on our existing network. The DFs might be impacted by a population bias because of this. For instance, it's possible that the participants' attitudes toward us were more favorable than we would have liked, which would indicate that some DFs' deficiencies went undetected. Regarding the results' **external validity**, while we can see the findings apply to any data space, it is important to note that we haven't carried out any additional implementations, so we are unable to independently verify the results' generalizability.

When it comes to future work, it makes sense to consider not only how the connectors need to be implemented to select a suitable data trustee between two participants, but also to look at the data trustee itself. It could be worthwhile to examine which design principles or design features a data trustee must implement to further support the selection of a suitable data trustee in data spaces. It is also possible to generally look at which design principles a data trustee must implement. Furthermore, it should be examined whether the design features can also be transferred to data spaces of any domain. Although we assume that this is the case, it requires verification. By developing a software artifact, our solution can be used in practice. In the next step, it would be interesting to observe whether the implemented EDC extension also achieves broad acceptance in practice. Particularly interesting is whether data exchanges via a data trustee between different data spaces can be negotiated.

REFERENCES

Aliprandi, S. (2011). Interoperability and open standards: the key to a real openness. International Free and Open Source Software Law Review, 3(1), 5–24. https://doi.org/10.5033/ifosslr.v3i1.53

Azkan, C., Möller, F [F.], Ebel, M., Iqbal, T., Otto, B [B.], & Poeppelbuss, J. (2022). Hunting the Treas-ure: Modeling Data Ecosystem Value Co-Creation. ICIS 2022 Proceedings(14).

Barcellos, M., Santos, G., Conte, T., Trinkenreich, B., & Matsubara, P. (2021). Organizing Empirical Studies as Learning Iterations in Design Science Research Projects. In E. D. Canedo, D. Viana, V. Garcia, C. Bezerra, I. d. S. Santos, B. Gadelha, I. Machado, S. Soares, U. Kulesza, B. de França, T. Conte, J. C. Maldonado, S. Reinehr, A. Malucelli, A. B. Albuquerque, G. Santos, M. P. Barcellos, R. Santos, C. Lima, L. Rocha (Eds.), 2021 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES) (pp. 1–10). ACM. https://doi.org/10.1145/3571473.3571474

Blankertz, A., & Specht-Riemenschneider, L. (2021). Neue Modelle ermöglichen: Regulierung für Datentreuhänder.

BMDV. (2024). Mobilithek. https://mobilithek.info/

Braud, A., Fromentoux, G., Radier, B., & Le Grand, O. (2021). The Road to European Digital Sovereignty with Gaia-X and IDSA. IEEE Network, 35(2), 4–5. https://doi.org/10.1109/MNET.2021.9387709

Brost, G. S., Huber, M., Weiß, M., Protsenko, M., Schütte, J., & Wessel, S. (2018). An Ecosystem and IoT Device Architecture for Building Trust in the Industrial Data Space. In D. Gollmann (Ed.), ACM Conferences, Proceedings of the 4th ACM Workshop on Cyber-Physical System Security (pp. 39–50). ACM. https://doi.org/10.1145/3198458.3198459

Caruso. (2024). Caruso Dataplace. https://www.caruso-dataplace.com/

Chandra, L., Seidel, S., & Gregor, S. (2015). Prescriptive Knowledge in IS Research: Conceptualizing Design Principles in Terms of Materiality, Action, and Boundary Conditions. In 2015 48th Hawaii International Conference on System Sciences (pp. 4039–4048). IEEE. https://doi.org/10.1109/HICSS.2015.485

Curry, E. (2020). Real-time Linked Dataspaces: Enbaling Data Ecosystems for Intelligent Systems. Springer.

Deng, W., Huang, T., & Wang, H. (2023). A Review of the Key Technology in a Blockchain Building Decentralized Trust Platform. Mathematics, 11(1), 101. https://doi.org/10.3390/math11010101

Doan, X., Selzer, A., Rossi, A., Botes, W. M., & Lenzini, G. (2022). Context, Prioritization, and Unexpectedness: Factors Influencing User Attitudes About Infographic and Comic Consent. In I. Herman, L. Médini, F. Laforest, R. Troncy, E. Simperl, D. Agarwal, & A. Gionis (Eds.), WWW '22 Companion: Companion proceedings of the Web Conference 2022: April 25, 2022, Lyon, France (pp. 534–545).

Association for Computing Machinery. https://doi.org/10.1145/3487553.3524632

Eclipse Foundation. (2024). Eclipse Dataspace Components. https://projects.eclipse.org/projects/technology.edc

European Union Agency for Cybersecurity. (2022). Digital identity: Leveraging the SSI concept to build trust. Publications Office. https://doi.org/10.2824/8646

Ferguson, N., & Schneier, B. (2003). Practical cryptography. Schneier's cryptography classics library / Bruce Schneier. Wiley.

Fürstenau, D., Gersch, M., & Schreiter, S. (2023). Digital Therapeutics (DTx). Business & Information Systems Engineering, 65(3), 349–360. https://doi.org/10.1007/s12599-023-00804-z

Gelhaar, J., & Otto, B [B.] (2020). Challenges in the Emergence of Data Ecosystems. Pacific Asia Conference on Information Systems (PACIS) 2020.

Ghayyur, S., Pappachan, P., Wang, G., Mehrotra, S., & Venkatasubramanian, N. (2020). Designing privacy preserving data sharing middleware for internet of things. In ACM Digital Library, Proceedings of the Third Workshop on Data: Acquisition To Analysis (pp. 1–6). Association for Computing Machinery. https://doi.org/10.1145/3419016.3431484

Gieß, A., Hupperz, M., Schoormann, T., & Möller, F [Frederik] (2024). What Does it Take to Connect? Unveiling Characteristics of Data Space Connectors. Proceedings of the 57th Hawaii International Conference on System Sciences (HICSS).

Gregor, S., & Hevner, A. R. (2013). Positioning and Presenting Design Science Research for Maximum Impact. MIS Quarterly, 37(2), 337–355. https://doi.org/10.25300/MISQ/2013/37.2.01

Haoyang Che. (2006). Managing trust in peer-to-peer systems: The last few years have seen an explosive growth of diversified peer-to-peer (P2P) systems over the internet, such as Gnutella, Chord, Freenet, and KaZaA. https://www.bcs.org/articles-opinion-and-research/managing-trust-in-peer-to-peer-systems/

Hevner, March, Park, & Ram (2004). Design Science in Information Systems Research. MIS Quarterly, 28(1), 75. https://doi.org/10.2307/25148625

Huang, J., & Nicol, D. M. (2013). Trust mechanisms for cloud computing. Journal of Cloud Computing: Advances, Systems and Applications, 2(1), 9. https://doi.org/10.1186/2192-113X-2-9

Hupperz, M., & Gieß, A. (2024). The Interplay of Data-Driven Organizations and Data Spaces: Unlocking Capabilities for Transforming Organizations in the Era of Data Spaces. Proceedings of the 57th Hawaii International Conference on System Sciences (HICSS).

Janiesch, C., Rosenkranz, C., & Scholten, U. (2020). An Information Systems Design Theory for Service Network Effects. Journal of the Association for Information Systems, 21, 1402–1460. https://doi.org/10.17705/1jais.00642

John R. Venable, Jan Pries-Heje, & Richard L. Baskerville (2017). Choosing a Design Science Research Methodology.

Lechte, H., Menck, J. H. D., Stocker, A., Lembcke, T.-B., & Kolbe, L. M. (2023). Exploring threat-specific privacy assurances in the context of connected vehicle applications. ECIS 2023.

Lucy, D., Mason, B., Sinclair, A., Bosley, A., & Gifford, J. (2023). Fair selection: An evidence review: Scientific summary. www.cipd.co.uk

Maaß, W. (2022). Contract-based Data-sharing for AI-based Decision Making on the Web. 55th Hawaii International Conference on.

Möller, F [Frederik], Guggenberger, T. M., & Otto, B [Boris]. (2020). Towards a Method for Design Principle Development in Information Systems. In S. Hofmann (Ed.), Lecture Notes in Computer Science: Vol. 12388. Designing for digital transformation: Co-creating services with citizens and industry : 15th international conference on design science research in information systems and technology, DESRIST 2020, Kristiansand, Norway, December 2-4, 2020 : proceedings (Vol. 12388, pp. 208–220). Springer. https://doi.org/10.1007/978-3-030-64823-7_20

Munoz-Arcentales, A., López-Pernas, S., Pozo, A., Alonso, Á., Salvachúa, J., & Huecas, G. (2019). An Architecture for Providing Data Usage and Access Control in Data Sharing Ecosystems. Procedia Computer Science, 160, 590–597. https://doi.org/10.1016/j.procs.2019.11.042

Österle, H. (Ed.). (2010). Gestaltungsorientierte Wirtschaftsinformatik: Ein Plädoyer für Rigor und Relevanz. Infowerk.

Otto, B [Boris], & Jarke, M. (2019). Designing a multi-sided data platform: findings from the International Data Spaces case. Electronic Markets, 29(4), 561–580. https://doi.org/10.1007/s12525-019-00362-x

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. Journal of Management Information Systems, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Potoczny-Jones, I., Kenneally, E., & Ruffing, J. (2019). Encrypted Dataset Collaboration. In ACM Digital Library, Proceedings of the 2nd ACM/EIGSCC Symposium on Smart Cities and Communities (pp. 1–8). Association for Computing Machinery. https://doi.org/10.1145/3357492.3358630

Recker, Rosemann, Green, & Indulska (2011). Do Ontological Deficiencies in Modeling Grammars Matter? MIS Quarterly, 35(1), 57. https://doi.org/10.2307/23043489

Sänger, J., Richthammer, C., & Pernul, G. (2015). Reusable components for online reputation systems. Journal of Trust Management, 2(1). https://doi.org/10.1186/s40493-015-0015-3

Schinke, L., Hoppen, M., Atanasyan, A., Gong, X., Heinze, F., Stollenwerk, K., & Roßmann, H.-J. (2023). Trustful Data Sharing in the Forest-based Sector - Opportunities and Challenges for a Data Trustee. https://doi.org/10.18154/RWTH-2023-08214

Sein, Henfridsson, Purao, Rossi, & Lindgren (2011). Action Design Research. MIS Quarterly, 35(1), 37. https://doi.org/10.2307/23043488

van Aken, J. E. (2004). Management Research Based on the Paradigm of the Design Sciences: The Quest for Field‐Tested and Grounded Technological Rules. Journal of Management Studies, 41(2), 219–246. https://doi.org/10.1111/j.1467-6486.2004.00430.x

Wieringa, R. J. (2014). Design Science Methodology for Information Systems and Software Engineering (Aufl. 2014). Springer Berlin Heidelberg. http://nbn-resolving.org/urn:nbn:de:bsz:31-epflicht-1614141

William L Kuechler, Vijay Vaishnavi, & Bill Kuechler (2021). Design Science Research InInformation Systems. http://www.desrist.org/design-research-in-information-systems/

Wohlin, C., & Runeson, P. (2021). Guiding the selection of research methodology in industry–academia collaboration in software engineering. Information and Software Technology, 140, 106678. https://doi.org/10.1016/j.infsof.2021.106678