

FarmTrustee: A Finest-Grained Filter Service for Agriculture Data Spaces

Sascha Kober
0009-0008-6182-988X
University of Leipzig,
Grimmaische Str. 12, 04109
Leipzig, Germany
Email: kober@wifa.uni-leipzig.de

André Ludwig
0000-0002-0358-3470
University of Leipzig,
Grimmaische Str. 12, 04109
Leipzig, Germany
Email: ludwig@wifa.uni-leipzig.de

Bogdan Franczyk
0000-0002-5740-2946
University of Leipzig,
Grimmaische Str. 12, 04109
Leipzig, Germany
Email: franczyk@wifa.uni-leipzig.de

Abstract— Digital agriculture increasingly depends on cross-organizational data sharing, yet many agricultural actors remain reluctant to disclose data because they lack enforceable control over how individual data elements are exposed. Existing governance and interoperability initiatives provide important foundations, but they do not sufficiently address the runtime enforcement of fine-grained disclosure decisions in practical machine-to-machine exchanges. This paper presents FarmTrustee, a finest-grained filter service for the agricultural data space. Building on a Finest-Grained Attribute-based Policy Machine (FisGAP), FarmTrustee extends attribute-based control to digit-level disclosure and applies policy-driven transformations to outgoing JSON payloads. Rather than merely deciding whether a field is visible, the service can suppress, generalize, reduce, or partially reveal values such as stock levels, timestamps, and geocoordinates. FarmTrustee is implemented as a Policy Decision Point and Policy Enforcement Point pair embedded in a synchronization pipeline. The approach is demonstrated in the MIRO project and integrated into a data platform for synchronization between Farm Management Information Systems and external systems such as webshops and reporting interfaces. Initial prototype results indicate that the approach can preserve data sovereignty and support data minimization while introducing only limited operational overhead. FarmTrustee therefore provides a practical runtime mechanism that connects data space governance with enforceable disclosure control.

Index Terms— agricultural data spaces, access control, attribute-based access control, finest-grained filter service, JSON filtering, policy enforcement.

I. INTRODUCTION

DIGITAL agriculture depends increasingly on the controlled exchange of data across organizational boundaries [1]. Farm Management Information Systems (FMIS), e-commerce platforms, advisory services, public authorities, and processors rely on data to improve efficiency, transparency, and sustainability [2]. Yet agricultural actors remain reluctant to share data because they fear misuse, loss of negotiating power, privacy violations, and opaque downstream processing [3]. This tension is well documented in research on agri-food data governance [3], [4]. However, existing approaches do not suffi-

ciently address how data can be selectively shared at a fine-grained level in real-world systems.

European initiatives such as GAIA-X and the European Data Space for Smart Agriculture provide important governance and interoperability foundations [5]. Despite this, practical enforcement mechanisms are still required at the level of concrete data exchange [6]. In many operational settings, conventional field-level authorization is too coarse [7]. A farmer may wish to share a stock value but not an internal identifier, a field location but not exact coordinates, or an application record but only at reduced temporal precision. In these cases, the decisive question is not merely whether an attribute can be accessed, but which part of the value may be disclosed to which requester and for which purpose [8], [9].

FarmTrustee addresses this gap as a finest-grained filter service for agricultural data exchange platforms. Building on a Finer-Grained Attribute-based Policy Machine [10], the service extends attribute-based control to the digit level and applies runtime transformations to JSON payloads. This enables a policy-based middle ground between complete disclosure and complete suppression. Instead of exposing entire values, the filter service can generalize dates, reduce coordinate precision, reveal selected digits of numeric values, or suppress sensitive fragments while preserving the utility of the remaining data.

The paper makes three contributions. First, it conceptualizes a finest-grained extension of the policy machine for digit-level disclosure control. Digit-level disclosure refers to controlling which parts of a value (e.g., individual digits or precision levels) are exposed. Second, it presents the filter service as an implementation of that extension using a Policy Decision Point and Policy Enforcement Point pair embedded in a synchronization pipeline. Third, it demonstrates the approach in the MIRO project and discusses its application in the data platform for e-commerce synchronization and regulatory reporting. Taken together, these contributions transform the conceptual principles of data sovereignty and data minimization into a concrete runtime

This paper is funded by Federal Ministry of Agriculture, Food and Regional Identity (Research Project ID: 2822ZR0007)

service that can be enforced in practical agricultural data exchange scenarios.

This perspective is important because discussions of agricultural data spaces often remain at the level of governance principles, contractual fairness, or high-level platform architecture [1], [11]. While such discussions are essential, they alone fail to specify the details of generating a tailored API response for a particular recipient. *FarmTrustee* focuses precisely on that operational moment. It asks how a platform can derive a recipient-specific, purpose-limited representation of a source object without modifying the source record itself. In this sense, *FarmTrustee* is both a technical enforcement mechanism and a bridge between data space governance and runtime system behavior.

The agricultural domain is especially suited to this problem definition. Agricultural data is heterogeneous, locally contextual, commercially sensitive, and increasingly interconnected across the value chain. The same record may matter simultaneously for inventory management, web-shop synchronization, traceability, compliance, and advisory services. A binary access decision is, therefore, often too coarse-grained. What is needed is a technically enforceable way to expose useful views of data while withholding the parts that create unnecessary risk. By delegating the interpretation of access rights to an automated transformation layer, this architecture aligns with contemporary research into policy-carrying data and interoperable frameworks designed to harmonize sovereignty with collaborative data utility [12], [13].

II. LITERATURE REVIEW

A. Related Work and Research Gap

Research on agricultural data sharing consistently identifies trust, data sovereignty, and interoperability as decisive prerequisites for adoption [3], [14]. Farmers often perceive asymmetric contracts, weak transparency, and unclear downstream use of their data as major risks [3], [15], which reinforces existing data silos and slows digital transformation [1], [3]. The literature therefore emphasizes that agricultural data spaces must offer not only technical interoperability but also credible governance structures that make the terms of exchange understandable and enforceable for primary data providers [4], [12]. The AgriTrust framework offers a coherent pathway toward a transparent and equitable agricultural data economy by addressing technical and economic barriers through integrated governance and semantics [14]. Such frameworks endeavor to overcome the "AgData Paradox," where the recognized value of agricultural data is often stifled by a pervasive lack of trust and interoperability [14]. While these initiatives aim to foster trust through transparency and clear data flow identification [16], the practical implementation of fine-grained control over data disclosure, particularly for sensitive environmental and soil conditions, remains a significant challenge [17]. Specifically, current frameworks

often lack the granular mechanisms required for usage control and data provenance that can adequately address the nuances of selective disclosure, which is essential for mitigating concerns about data misuse and ensuring data sovereignty [3], [18]. Existing data-sharing frameworks for digital agriculture often focus on high-level privacy and security, proposing access control mechanisms and privacy-preserving data aggregation schemes; however, their applicability to small and medium-sized farming operations remains limited and do not fully address the granular requirements for selective data disclosure [17]. This limitation highlights a significant research gap at the intersection of data governance, semantic technologies, and distributed systems, indicating a need for a foundational layer to achieve a cohesive, operational framework for agriculture [14].

Several strands of research contribute relevant building blocks. Work on privacy-preserving data sharing investigates protected aggregation, secure linkage, and disclosure reduction. Research on access control explores attribute-based authorization, dynamic policy evaluation, and trusted intermediaries. The AgriTrust framework proposes a governance model for smart farming ecosystems, built on pillars of data sovereignty, transparent data contracts, equitable value sharing, and regulatory compliance [14].

Despite this progress, a practical gap remains. Most existing approaches continue to make access decisions at the level of datasets, objects, or complete attributes. For many agricultural integration scenarios, that granularity is insufficient because the useful and sensitive parts of a value may coexist within the same field. A timestamp can be useful at the day level but sensitive at the minute level; a coordinate may be operationally adequate when rounded but commercially risky when exact; a stock or price field may be partially shareable but not in full precision. Consequently, existing frameworks face challenges in consistently enforcing fine-grained masking or generalization policies dynamically, which can prevent the effective implementation of data sharing agreements that depend on context-sensitive disclosure [19], [20]. A runtime mechanism capable of expressing and enforcing nuanced disclosure decisions in machine-to-machine exchanges is still needed [21].

The finer-grained policy machine already advances fine-grained access control in data trusts by combining inheritance mechanisms with attribute-based policy definitions. *FarmTrustee* builds on this foundation and targets the remaining research gap: the operationalization of finest-grained disclosure decisions in near real-time synchronization workflows. The core novelty lies not merely in defining more precise permissions, but in converting those permissions into policy-driven transformations on outgoing JSON payloads so that data minimization becomes part of the exchange pipeline itself [10].

This gap matters for two reasons. First, without operational enforcement, governance promises remain difficult

to verify. Second, without sufficiently granular enforcement, data owners are forced into inefficient trade-offs between usefulness and confidentiality. *FarmTrustee* responds to both issues by combining policy semantics with a filtering mechanism that can act directly on exchanged objects.

B. Research Method and Design Requirements

This study follows a design-oriented research approach. Requirements were derived from two sources: first, the literature on agricultural data governance, privacy, and trusted data sharing; second, the practical integration context of the MIRO project, where data must be synchronized between Farm Management Information Systems and external systems such as webshops. Rather than treating the problem as a purely conceptual governance question, the study focuses on the design of a technical artifact that can be embedded in an existing data exchange platform (*DEP*).

Five design requirements emerged from the analysis: (1) data minimization, ensuring that only information adequate, relevant, and necessary for a specific purpose is disclosed; (2) machine-readable policy specification to enable automated decision-making; (3) compatibility with heterogeneous JSON-based interfaces, reflecting the prevalence of structured API payloads in agricultural systems; (4) flexible, attribute-based authorization considering requester, resource, and contextual factors; and (5) low runtime overhead to avoid performance bottlenecks in synchronization workflows [4], [8], [20].

The filter service was therefore designed as a policy-driven filter layer inside the *DEP* rather than as a separate manual review process. The prototype was evaluated using multiple predefined policy scenarios and requester profiles based on synthetic product and stock data. Requests were executed under varying requester profiles, while the system applied *FisGAP*-derived policies and returned filtered JSON payloads. The evaluation assessed both the functional correctness of the filtering results and the operational overhead introduced by policy evaluation and payload transformation. To improve reproducibility, the prototype benchmark was conducted under a defined deployment context, including specified request volume, payload structure, and container resource limits. However, broader performance testing with varying payload sizes, policy complexity, and concurrent request scenarios remains part of future work.

Methodologically, the work is best understood as artifact design and prototype demonstration. The contribution lies not in statistical generalization, but in the design and demonstration of a practically grounded architecture that can be instantiated, tested, and extended. This positioning is appropriate for a domain where trust-related challenges are well understood, while concrete runtime enforcement mechanisms are still evolving [1].

III. FISGAP AND FILTER SERVICE

A. Introduction

The Finest-Grained Attribute-based Policy Machine (*FisGAP*) provides the conceptual foundation for the filter service. In the introduced previous version of the policy machine, access rights are modeled in a data trust context using attribute-based policies and inheritance mechanisms [10], [22]. *FisGAP* extends this logic from attribute-level visibility to digit-level disclosure control (see Fig. 1). When a request is received, the policy machine evaluates it by identifying the requested object ID and the corresponding user (identity) ID. Based on these identifiers, the system determines the shortest path within the policy graph to the relevant *object access node* associated with the requester. At this node, the applicable access rules are retrieved. If additional digit-level access nodes are linked to the *object access node*, these rules are also evaluated and included in the decision. This mechanism ensures that both attribute-level and digit-level permissions are consistently enforced. Policies can be dynamically extended or modified by submitting change requests for specific controlled *object access nodes*. This allows authorized users to define, update, or refine access rules without altering the underlying system logic, thereby supporting flexible and scalable policy management. At this level, access can be defined at the granularity of individual digits. For example, a policy may allow only the first four digits of a data field to be disclosed. The shift is important because many data-sharing conflicts arise not from a binary need to show or hide a field, but from the need to reveal only a constrained representation of the underlying value [23], [24].




Fig. 1. Finest-Grained Attribute-Based Policy Machine

In summary, the Policy Machine (*FisGAP*) is responsible for the definition, storage, and evaluation of access control policies. Its primary function is not to process or transform data, but to manage and enforce rules that determine how data may be accessed and disclosed. By systematically combining subject, object, and contextual attributes, the Policy Machine enables consistent and automated decision-making [25]. The system operates via an API, through which policies are retrieved and delivered in JSON format [26]. This interface allows seamless integration into existing systems, while enabling standalone operation as a public API [26]. Such a rule-based architecture forms the foundation for finest-grained access control, as it allows precise specification of which parts of a data value may be revealed under which conditions [25].

TABLE I.
ENFORCEMENT ACTIONS IN *FarmTrustee*

Enforcement Action	Description	Example
Full Disclosure	Returns the original value without modification	Exact stock level is shared with an authorized partner
Suppression	Removes the value entirely from the response payload	Internal product identifier is omitted
Generalization	Replaces a value with a less specific representation to reduce sensitivity	Full timestamp is reduced to month or day level
Reduction	Reveals only a selected part of a value while masking the rest	Coordinates are rounded; price shown with limited decimals
Noising	Applies controlled perturbation to a value to protect sensitive information	Slight random variation added to a quantity value

This finest-grained control is broadly applicable across domains wherever sensitive data must be shared at the digit level under controlled delegation, including agricultural data spaces, where commercially sensitive details and operationally useful information are often encoded within the same record [27].

With *FisGAP* integrated, the *FarmTrustee* model supports five principal enforcement actions that enable fine-grained control over data disclosure. These actions define how attribute values are transformed before being shared with a requester and are summarized in Table I.

Together, they establish a pragmatic policy vocabulary that allows data providers to balance data utility and confidentiality within heterogeneous data exchange scenarios [20].

A policy in *FarmTrustee* combines subject attributes, resource attributes, and contextual attributes with transformation directives. In operational terms, a webshop, regulator, logistics partner, or advisory service can each receive

a different representation of the same source object. The source data remains unchanged in the originating system, while the outgoing representation is tailored to the permitted disclosure scope [27].

This architecture preserves data sovereignty because decisions about what leaves the trusted boundary remain under the control of the data owner or governing organization.

The refinement from fine-grained to most fine-grained control also improves explainability. Policies can be described in terms that align with practical business decisions: reveal stock but round quantity, disclose field location but truncate coordinates, or share treatment information but suppress supplier details. This alignment matters because policy adoption depends not only on technical expressiveness but also on whether farmers and intermediaries can understand how the resulting disclosure behavior matches their trust expectations [3], [4].

In this respect, *FarmTrustee* can be understood as a data-space service platform that makes governance actionable. It does not replace contracts, consent procedures, or institutional trust; rather, it complements them with a concrete enforcement layer that implements what those governance instruments specify.

B. System Architecture and Runtime Filtering

FarmTrustee follows an Attribute-Based Access Control (ABAC) architecture with a Policy Administration Point (PAP), a Policy Decision Point (PDP), and a Policy Enforcement Point (PEP). Policies defined in the *FisGAP* are managed within the Policy Machine, exposed via an API in a machine-readable JSON representation, and evaluated by the PDP against subject, resource, and contextual attributes. The PEP intercepts outgoing data flows, recursively traverses JSON payloads, and transforms each relevant field before the response is forwarded to the requester. This separation of policy evaluation from business logic improves maintainability and avoids hard-coding disclosure rules in integrated systems [28].

At runtime, the requesting identity is first resolved within the surrounding platform context. Subsequently, the PDP evaluates the applicable policy set and determines the permitted operations for each addressed data fragment. These decisions are then enforced by the PEP, which applies the required transformations while traversing the outgoing data object. Because the filtering mechanism operates on nested JSON structures, it can process composite payloads containing identifiers, coordinates, status fields, timestamps, and numerical values within a single synchronized response. A key design aspect is that the recursive traversal mechanism allows finest-grained transformations to be uniformly applied across different object hierarchies. The runtime interaction between the Filter Service and the Policy Machine is illustrated in Fig. 2. When a user submits a data request, it is first directed to the Filter Service,

which acts as the central enforcement component within the architecture.

The Filter Service then fetches the raw data exclusively via standardized APIs, typically in JSON format, from the source systems, rather than accessing databases directly. This API-based integration ensures loose coupling between systems, allowing *FarmTrustee* to operate across heterogeneous data infrastructures without requiring access to internal storage.

Simultaneously, the Filter Service requests the applicable access control policies from the Policy Machine. Based on the combination of subject, resource, and contextual attributes, the Policy Machine returns the relevant policy set for the specific request.




Fig. 2. Filter Service and FisGAP Interaction

The Filter Service then applies these policies by filtering and transforming the retrieved data before returning the response to the user. This process includes finest-grained transformations, such as masking, truncation, aggregation, or digit-level disclosure control. As shown in Fig. 2, the Filter Service integrates data retrieval and policy evaluation into a unified processing pipeline, ensuring that only policy-compliant data leaves the trusted boundary. This approach enables dynamic, context-aware data sharing without modifying the original data stored in the source system.

This design operationalizes the principle of data minimization. Instead of creating multiple pre-modeled data products for each recipient, the system derives the allowed representation dynamically from the source object and the active policy context. As a result, new recipient classes or changed trust relationships can be managed through policy adaptation rather than application redesign. This is particularly valuable in data spaces, where organizational constellations evolve over time and where one source system may need to serve several downstream actors with different disclosure entitlements [29].

The architecture also separates storage from disclosure. Source systems continue to maintain full, authoritative records. *FarmTrustee* only manipulates the outbound representation. This distinction is essential because it preserves the integrity of internal data while allowing recipient-specific views to be generated on demand.

Finally, the recursive approach reduces implementation friction. Instead of writing special-case logic for every nested payload shape, *FarmTrustee* applies a consistent filtering strategy to all traversed objects. This makes the service attractive not only for one demonstrator platform, but also as a reusable architectural pattern for other data-space integrations.

IV. FARMTRUSTEE

In the MIRO project, the data exchange platform was developed for synchronizing agricultural business data between internal management systems and external applications. To enable this integration, APIs from entity resource planning (ERP) systems, FMIS and webshops were connected to the platform, allowing automated data flows across organizational boundaries. However, simply enabling connectivity does not address the underlying trust concerns associated with data sharing. Therefore, an additional requirement emerged: users must be able to explicitly control which data elements are shared and with whom.

Within this project context, FisGAP and the filter service were integrated into the *DEP*, resulting in *FarmTrustee*, a policy-aware data exchange platform with an integrated filter service. This platform provides the operational context in which source data is prepared, synchronized, and delivered to consuming systems. *FarmTrustee* was positioned between the source data handling logic and the outbound synchronization interface, thereby turning the *DEP* into a policy-aware exchange platform rather than a neutral data relay.

This integration has two practical effects. First, it allows producers to automate data sharing with webshops and other recipients without losing control over the granularity of disclosure. Second, it keeps the authoritative source data within the internal environment while only the reduced representation is exported. The approach therefore supports data sovereignty at the point where it matters most operationally: in the last step before data crosses organizational boundaries.




Fig. 3. FarmTrustee Platform

Fig. 3 clarifies the relationship between *FisGAP*, Filter Service, *FarmTrustee* and APIs. *MIRO* is the project context in which the overall solution was conceived and implemented. *FarmTrustee* is the data exchange platform developed within *MIRO*. *FisGAP* is the conceptual policy model that provides the foundation for fine- and finest-grained disclosure control over the filter service. The filter service, which emerged from integrating *FisGAP* into *FarmTrustee*, implements the resulting finest-grained enforcement logic. Making this separation explicit is important for scientific clarity because it distinguishes the project context, the platform artifact, the conceptual model, and the enforcement service.

Further integration shows that finest-grained control can be introduced incrementally. Existing synchronization logic does not need to be discarded; instead, a filtering layer can be inserted at the system boundary. This lowers adoption barriers because it allows organizations to preserve existing workflows while upgrading the level of control they can exercise over outbound data streams.




Fig. 4. Data Process in FarmTrustee

Fig. 4 shows the process of requesting and receiving filtered data through *FarmTrustee*. Before data exchange can occur, both the data owner and the requester must be registered in the system. The data owner can then use a frontend interface to define and manage access permissions for specific data objects. Based on these policies, a registered requester may submit a data request to *FarmTrustee*. The system evaluates the request, retrieves the relevant source data, and applies the integrated filter

service within the data pipeline. Consequently, the requester receives only a policy-compliant, filtered representation of the data rather than the raw source data. This process ensures controlled and automated data sharing while preserving the data owner's control over the original information, thereby enabling *finest-grained attribute-based access control*.

V. USE CASES

The following use cases illustrate how the *FarmTrustee* architecture addresses critical data-sharing challenges within the agricultural sector. These scenarios highlight the practical application of finest-grained filtering to reconcile the increasing demand for transparent, high-resolution data exchange with the fundamental need for farmer data sovereignty. By enabling controlled disclosure at the system boundary, the proposed solution facilitates participation in the Agriculture Data Space while mitigating operational risks and reducing the administrative burden associated with manual data preparation. The selected cases demonstrate the flexibility of the filter service in both commercial contexts, such as supply chain synchronization, and regulatory environments, such as mandatory environmental reporting.

A first use case concerns the synchronization of product data and stock levels from Farm Management Information Systems to external webshops. Downstream platforms require product names, descriptions, availability information, and sometimes price signals, but producers may wish to withhold internal identifiers, storage details, supplier metadata, or highly granular location information. *FarmTrustee* enables such differentiated disclosure decisions without interrupting the synchronization workflow. The webshop receives a policy-compliant view of the product record while the original source data remains unchanged in the management system.

A second use case concerns the reporting of plant protection product applications to public authorities. Such records can include product names, dosage information, application dates, and georeferenced field data, and are relevant for environmental monitoring and compliance processes [30], [31]. At the same time, repeated reporting can create administrative burden and intensify farmers' concerns about excessive visibility into operational practices [12], [32], [33]. *FarmTrustee* supports compliant reporting by allowing a farmer to disclose only the legally required parts of the record while suppressing commercial metadata or reducing spatial and temporal precision where appropriate [12], [32], [33], [34].

Typical transformations in this scenario include truncating coordinates, generalizing timestamps to broader intervals [35], [36], and removing unrelated commercial fields such as supplier information or internal purchase prices. The service, therefore, converts raw farm records into policy-compliant views that remain useful for the receiving

authority [12] while protecting sovereignty over the original dataset [37]. These cases illustrate that finest-grained disclosure control [38] is not only a privacy feature but also a mechanism for reducing organizational friction in data-sharing ecosystems [39].

A third, more general use case lies in data sharing with analytics or advisory services [40], [41]. Here, the exact operational record may be less important than a pattern-preserving representation that still supports recommendations or aggregated insights [42]. The same transformation vocabulary that is useful for online shops and regulators [43] can thus also support broader data-space participation while preserving farmer autonomy [44].

Furthermore, the generic filter service is applicable to all data transitions between API requests and responses.

VI. PROTOTYPE EVALUATION AND LIMITATIONS

The prototype confirmed that *FarmTrustee* can process structured JSON objects, apply policy decisions recursively, and generate filtered payloads in real time. Initial tests with synthetic product and stock data indicated only limited computational overhead compared with unfiltered transfer, which supports the feasibility of the approach for near real-time synchronization. Functionally, the service behaved as intended: fields were suppressed, generalized, reduced, or revealed according to the configured policy, and the source objects remained unchanged in the underlying system.

To evaluate the performance overhead, we measured response times for direct API calls versus calls routed through the *FarmTrustee* filter service, which invokes the API and processes the results. The results are depicted in Fig. 5. As expected, filtered responses exhibit longer latencies. The figure compares requests to the original API endpoint with equivalent requests via the filter service. Overall, the filter service incurs only a slight overhead relative to direct calls, with filtering requiring approximately 27 ms for a JSON response comprising approximately 159 fields across four nesting levels.

To make the prototype evaluation reproducible, the benchmark was executed against the same product synchronization endpoint with and without *FarmTrustee*. The *FarmTrustee* filter service was deployed as a containerized workload in a Rancher-managed Kubernetes/RKE2 cluster. The evaluated deployment used one replica of the *FarmTrustee* filter service container and exposed the service through an HTTPS ingress. The container was configured with a CPU limit of 1 vCPU and a memory limit of 1 GiB. The Neo4j database used by the policy machine to persist and retrieve access control rules was deployed in the same Rancher-managed Kubernetes/RKE2 environment and was configured with the same resource limits. At the time of the benchmark, the Neo4j database contained approximately 100 million synthetic test records. Consequently, both the runtime filtering component and the policy storage component were evaluated under constrained

and comparable container resource settings, while policy retrieval was performed against a large-scale test dataset.




Fig. 5. Response Time Comparison (Summarized)

The evaluated JSON response contained approximately 159 fields distributed across four nesting levels. Each measurement series consisted of 50 consecutive GET requests after an initial warm-up phase. For every request, the end-to-end response time was recorded at the client side, including data retrieval, policy retrieval from Neo4j, recursive JSON traversal, transformation, and serialization. The baseline represents direct calls to the original API endpoint, whereas the filtered variant routes the same request through the *FarmTrustee* filter service. The reported overhead therefore reflects the additional runtime cost introduced by policy-based disclosure enforcement.

However, while these initial benchmark results are encouraging, the current evidence base remains preliminary. Although the revised evaluation reports the deployment context, container resource limits, request volume, payload structure, and the size of the Neo4j test dataset, it is still limited to one representative synchronization endpoint and sequential request execution. The benchmark does not yet provide a broader analysis of concurrent access, throughput under load, median and worst-case latency across different payload sizes, or the isolated contribution of individual processing steps such as policy retrieval, recursive traversal, transformation, and serialization. Future performance tests should therefore extend the benchmark design by varying dataset size, policy complexity, payload structure, and request concurrency.

Beyond these performance-related limitations, further aspects require additional investigation. The present prototype focuses on structured JSON payloads and does not yet address other exchange formats in equal depth. Conflict resolution between overlapping transformation rules also deserves formal specification. Finally, user-centered evaluation with farmers, cooperatives, and platform operators

remains necessary to determine whether the policy concepts are understandable and whether the resulting disclosure patterns match stakeholder expectations in practice. These points do not diminish the contribution of the artifact, but they define the next steps required to move from a promising prototype toward a robust data space component.

Nevertheless, the prototype is already useful as a proof of operational feasibility. It demonstrates that a data-space component can be designed to implement nuanced disclosure behavior in a way that is compatible with existing API-centric workflows. This is a meaningful contribution in a field where governance discussions often outpace concrete runtime enforcement mechanisms.

VII. CONCLUSION

FarmTrustee operationalizes the principles of data sovereignty and data minimization by providing an enforceable runtime service for agricultural data spaces. By extending the FiGAP model to digit-level disclosure control and integrating policy evaluation directly into JSON-based synchronization workflows, the approach bridges the gap between abstract governance concepts and practical, enforceable data-sharing mechanisms. The service is particularly valuable in scenarios where full disclosure would be excessive but complete suppression would prevent legitimate data reuse. In these cases, partial disclosure, generalization, and reduction provide a balanced approach that strengthens trust in agricultural data sharing.

The approach addresses the “AgData Paradox” by enabling transparent, controlled data sharing and mitigating concerns related to data misuse, regulatory risks, and power imbalances among agricultural stakeholders [14]. By enabling fine-grained control over what is shared and with whom, *FarmTrustee* empowers farmers to maintain autonomy over their data while participating in data-driven innovation and regulatory reporting [45], [46]. This capability is especially relevant given widespread concerns about privacy, unclear data governance practices, and the lack of robust consent mechanisms in digital agriculture [3], [47].

Initial results from the *DEP* context demonstrate the technical feasibility and practical relevance of finest-grained access control for agricultural data exchange. Beyond technical performance, broader adoption will depend on transparent governance, comprehensible consent frameworks, and demonstrable value for primary data providers. Data intermediaries, cooperatives, and trusted governance structures may play a key role in fostering trust and enabling equitable data use [4]. Consequently, *FarmTrustee* should be viewed not as a standalone security mechanism, but as a foundational component of a broader agricultural data governance ecosystem.

VIII. ACKNOWLEDGMENTS

During the preparation of this work, the authors used GPT-5.4 for language editing and improving transitions between paragraphs. The authors take full responsibility for the content of this publication.

REFERENCES

- [1] M. Šestak und D. Copot, „Towards Trusted Data Sharing and Exchange in Agro-Food Supply Chains: Design Principles for Agricultural Data Spaces“, *Sustainability*, Bd. 15, Nr. 18, S. 13746, Jan. 2023, doi: 10.3390/su151813746.
- [2] K.-O. Wenkel, M. Berg, W. Mirschel, R. Wieland, C. Nendel, und B. Köstner, „LandCaRe DSS – An interactive decision support system for climate change impact assessment and the analysis of potential agricultural land use adaptation strategies“, *Journal of Environmental Management*, Bd. 127, Apr. 2013, doi: 10.1016/j.jenvman.2013.02.051.
- [3] C. S. Sullivan, M. Gemtoui, E. Anastasiou, und S. Fountas, „Building trust: A systematic review of the drivers and barriers of agricultural data sharing“, *Smart Agricultural Technology*, Bd. 8, S. 100477, Jan. 2024, doi: 10.1016/j.atech.2024.100477.
- [4] J. Kaur, S. M. H. Fard, M. Amiri-Zarandi, und R. Dara, „Protecting farmers’ data privacy and confidentiality: Recommendations and considerations“, *Frontiers in Sustainable Food Systems*, Bd. 6, Okt. 2022, doi: 10.3389/fsufs.2022.903230.
- [5] J. Schoenke u. a., „Gaia-AgStream: An Explainable AI Platform for Mining Complex Data Streams in Agriculture“, in *Communications in computer and information science*, Springer Science+Business Media, 2021, S. 71–83. doi: 10.1007/978-3-030-88259-4_6.
- [6] A. Cimmino, J. Cano-Benito, und R. García-Castro, „Open Digital Rights Enforcement Framework (ODRE): from descriptive to enforceable policies“, *arXiv (Cornell University)*, Sep. 2024, doi: 10.48550/arxiv.2409.17602.
- [7] P. Colombo und E. Ferrari, „Fine-Grained Access Control Within NoSQL Document-Oriented Datastores“, *Data Science and Engineering*, Bd. 1, Nr. 3, S. 127–138, Aug. 2016, doi: 10.1007/s41019-016-0015-z.
- [8] H. Jiang und A. Bouabdallah, „JACPoL: A Simple but Expressive JSON-Based Access Control Policy Language“, in *Lecture notes in computer science*, Springer Science+Business Media, 2018, S. 56–72. doi: 10.1007/978-3-319-93524-9_4.
- [9] P. Plebani, D. Rossetto, und F. Tiezzi, „Empowering trusted data sharing for data analytics in a federated environment: A blockchain-based approach“, *Frontiers in Blockchain*, Bd. 6, Apr. 2023, doi: 10.3389/fbloc.2023.1141760.
- [10] S. Kober, M. Koch, A. Ludwig, und B. Franczyk, „Building a Data Trust: Fine(r)-Grained Attribute Based Policy Machine“, *Procedia Computer Science*, Bd. 270, S. 572–581, Jan. 2025, doi: 10.1016/j.procs.2025.09.176.
- [11] C. Atik, „Towards Comprehensive European Agricultural Data Governance: Moving Beyond the “Data Ownership” Debate“, *Deleted Journal*, Bd. 53, Nr. 5, S. 701–742, Mai 2022, doi: 10.1007/s40319-022-01191-w.
- [12] A. Kotal, L. Elluri, D. Gupta, V. Mandalapu, und A. Joshi, „Privacy-Preserving Data Sharing in Agriculture: Enforcing Policy Rules for Secure and Confidential Data Synthesis“, in *2021 IEEE International Conference on Big Data (Big Data)*, Dez. 2023, S. 5519–5528. doi: 10.1109/bigdata59044.2023.10386276.
- [13] J. Padget und W. Vasconcelos, „Fine-Grained Access Control via Policy-Carrying Data“, *ACM Transactions on Internet Technology*, Bd. 18, Nr. 3, S. 1–24, Feb. 2018, doi: 10.1145/3133324.
- [14] I. Bergier, J. G. A. Barbedo, É. L. Bolfe, D. P. Drucker, und F. Soares, „The AgriTrust Framework: Federated Semantic Governance for Trusted and Interoperable Agricultural Data Sharing“, *Automation*, Bd. 7, Nr. 2, S. 57–57, März 2026, doi: 10.3390/automation7020057.
- [15] L. Klerkx, E. Jakku, und P. Labarthe, „A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda“, *NJAS: Wageningen Journal of Life Sciences*, Bd. 90, Nr. 1, S. 1–16, Dez. 2019, doi: 10.1016/j.njas.2019.100315.

- [16] B. Lauga u. a., „MULTIPASS: Managing the consents of access to farm data in a chain of trust to make new services emerge for farmers“, *HAL (Le Centre pour la Communication Scientifique Directe)*, Juni 2019, [Online]. Verfügbar unter: <https://hal.archives-ouvertes.fr/hal-02470007>
- [17] O. Zafar, R. S. González, M. Namazi, A. J. Rodríguez-Morales, und E. Ayday, „Empowering Digital Agriculture: A Privacy-Preserving Framework for Data Sharing and Collaborative Research“, *arXiv (Cornell University)*, Juni 2025, doi: 10.48550/arxiv.2506.20872.
- [18] S. Steinbuß u. a., „Usage Control in the International Data Spaces“, März 2021. [Online]. Verfügbar unter: <https://zenodo.org/record/5675884>
- [19] O. Zafar, R. S. Gonzalez, G. A. Wilkins, A. J. Rodríguez-Morales, und E. Ayday, „Privacy-Preserving Data Linkage Across Private and Public Datasets for Collaborative Agriculture Research“, *arXiv (Cornell University)*, Sep. 2024, doi: 10.48550/arxiv.2409.06069.
- [20] M. Zubair, M. Sabzevari, V. Khatri, S. Tarkoma, und K. Hätönen, „Access control for trusted data sharing“, *EURASIP Journal on Information Security*, Bd. 2024, Nr. 1, S. 30, Sep. 2024, doi: 10.1186/s13635-024-00178-z.
- [21] J. Happa, N. Moffat, M. Goldsmith, und S. Creese, „Run-Time Monitoring of Data-Handling Violations“, in *Lecture notes in computer science*, Springer Science+Business Media, 2019, S. 213–232. doi: 10.1007/978-3-030-12786-2_13.
- [22] Kober, Sascha, „Finest-Grained Attribute-Based Policy Machine“. 2025. <https://github.com/KoberSascha/Finest-Grained-Attribute-Based-Policy-Machine>.
- [23] Y. Amar, H. Haddadi, und R. Mortier, „An Information-Theoretic Approach to Time-Series Data Privacy“, Bd. 3, S. 1–6, Apr. 2018, doi: 10.1145/3195258.3195261.
- [24] A. Michel, „Personalising privacy constraints in Generalization-based Anonymization Models“, *HAL (Le Centre pour la Communication Scientifique Directe)*, Apr. 2019, [Online]. Verfügbar unter: <https://tel.archives-ouvertes.fr/tel-02414240>
- [25] D. F. Ferraiolo, S. I. Gavrila, W. Jansen, und P. E. Stutzman, „Policy Machine: Features, Architecture, and Specification“, Okt. 2015. doi: 10.6028/nist.ir.7987r1.
- [26] A. Mimidis-Kentis, „SDN-Enabled Management of Access Infrastructures and Cloudbased Platforms“, *Research Portal Denmark*, S. 141, Jan. 2020.
- [27] H. Pesanhan and C. B. Westphall, „Untangling the Use of Blockchain in Distributed AgriFood Supply Chains: IoT data privacy and ABAC policy smart contract“, *Research Square (Research Square)*, Dez. 2023, doi: 10.21203/rs.3.rs-3750662/v1.
- [28] A. Mohamed, D. Auer, D. Hofer, und J. Küng, „A systematic literature review of authorization and access control requirements and current state of the art for different database models“, *International Journal of Web Information Systems*, Bd. 20, Nr. 1, S. 1–23, Okt. 2023, doi: 10.1108/ijwis-04-2023-0072.
- [29] D. Preuveneers, W. Joosen, und E. Ilie-Zudor, „Policy reconciliation for access control in dynamic cross-enterprise collaborations“, *Enterprise Information Systems*, Bd. 12, Nr. 3, S. 279–299, Juli 2017, doi: 10.1080/17517575.2017.1355985.
- [30] F. M. Fishel, „Applying Pesticides in Compliance with the National Pollutant Discharge Elimination System (NPDES)“, *EDIS*, Bd. 2012, Nr. 1, Jan. 2012, doi: 10.32473/edis-pi239-2011.
- [31] S. Trapp, W. J. Doucette, P. Fantke, und T. Eggen, „A suggested minimum data list for documenting experimental plant uptake studies“, *Research Portal Denmark*, S. 162–163, Jan. 2016.
- [32] N. E. Benni, C. Ritzel, K. Heitkämper, C. Umstätter, A. Zorn, und G. Mack, „The cost of farmers’ administrative burdens due to cross-compliance obligations“, *Journal of Environmental Planning and Management*, Bd. 65, Nr. 5, S. 930–952, Juli 2021, doi: 10.1080/09640568.2021.1920376.
- [33] L. Reissig, A. Stoinescu, und G. Mack, „Why farmers perceive the use of e-government services as an administrative burden: A conceptual framework on influencing factors“, *Journal of Rural Studies*, Bd. 89, S. 387–396, Jan. 2022, doi: 10.1016/j.jrurstud.2022.01.002.
- [34] S. Spycher u. a., „Linking chemical surface water monitoring and pesticide regulation in selected European countries“, *Environmental Science and Pollution Research*, Bd. 31, Nr. 30, S. 43432–43450, Juni 2024, doi: 10.1007/s11356-024-33865-y.
- [35] S. Lestyán, G. Ács, und G. Biczók, „In Search of Lost Utility: Private Location Data“, *Proceedings on Privacy Enhancing Technologies*, Bd. 2022, Nr. 3, S. 354–372, Juli 2022, doi: 10.56553/popets-2022-0076.
- [36] D. Tomaras, M. Tsenos, und V. Kalogeraki, „Practical Privacy Preservation in a Mobile Cloud Environment“, in *2022 23rd IEEE International Conference on Mobile Data Management (MDM)*, Juni 2022, S. 188–197. doi: 10.1109/mdm55031.2022.00047.
- [37] F. Berisha, P. Mooney, Z. Pourzolfaghar, und M. Helfert, „Shared Yet Owned: The Dual Path of Data Ownership in Agriculture: A Systematic Review“, *Data & Policy*, Bd. 8, Dez. 2025, doi: 10.1017/dap.2025.10050.
- [38] G. Bender, L. Kot, J. Gehrke, und C. Koch, „Fine-grained disclosure control for app ecosystems“, S. 869–880, Juni 2013, doi: 10.1145/2463676.2467798.
- [39] J. Holstein, M. Schemmer, J. Jakubik, M. Vössing, und G. Satzger, „Sanitizing data for analysis: Designing systems for data understanding“, *Electronic Markets*, Bd. 33, Nr. 1, Okt. 2023, doi: 10.1007/s12525-023-00677-w.
- [40] A. Durrant, M. Marković, D. R. Matthews, D. May, G. Leontidis, und J. Enright, „How might technology rise to the challenge of data sharing in agri-food?“, *Global Food Security*, Bd. 28, S. 100493–100493, Jan. 2021, doi: 10.1016/j.gfs.2021.100493.
- [41] M. Ryan, „Agricultural Big Data Analytics and the Ethics of Power“, *Journal of Agricultural and Environmental Ethics*, Bd. 33, Nr. 1, S. 49–69, Nov. 2019, doi: 10.1007/s10806-019-09812-0.
- [42] A. Galozy, S. Alawadi, V. R. Kebande, und S. Nowaczyk, „Beyond Random Noise: Insights on Anonymization Strategies from a Latent Bandit Study“, *arXiv (Cornell University)*, Okt. 2023, doi: 10.48550/arxiv.2310.00221.
- [43] Z. Li, K. Gao, und G. Qiao, „From Online Markets to Green Fields: Unpacking the Impact of Farmers’ E-Commerce Participation on Green Production Technology Adoption“, *Agriculture*, Bd. 15, Nr. 14, S. 1483–1483, Juli 2025, doi: 10.3390/agriculture15141483.
- [44] C. Xie, „The role of modern agricultural technologies in improving agricultural productivity and land use efficiency“, *Frontiers in Plant Science*, Bd. 16, S. 1675657–1675657, Sep. 2025, doi: 10.3389/fpls.2025.1675657.
- [45] S. van der Burg, L. Wiseman, und J. Krkeljas, „Trust in farm data sharing: reflections on the EU code of conduct for agricultural data sharing“, *Ethics and Information Technology*, Bd. 23, Nr. 3, S. 185–198, Juli 2020, doi: 10.1007/s10676-020-09543-1.
- [46] L. Wiseman, M. A. Jouanjan, F. Casalini, und E. Gray, „Issues around data governance in the digital transformation of agriculture“, *OECD food, agriculture and fisheries working papers*, Okt. 2020, doi: 10.1787/53ecf2ab-en.
- [47] J. Kaur und R. Dara, „Ensuring Privacy in Smart Farming: Review of Regulations, Codes of Conduct and Best Practices“, 2023, S. 1–16. doi: 10.1007/978-3-030-89123-7_272-1.