Transparent Accountability for Personal Data Sovereignty: Blockchain-Based Verification of Policy Compliance

VIJON BARAKU, Department of Computer Science, University of York, United Kingdom
IRAKLIS PARASKAKIS, SEERC – South East European Research Centre, Greece and University of York Europe Campus, Greece
SIMEON VELOUDIS, SEERC – South East European Research Centre, Greece and University of York Europe Campus, Greece
POONAM YADAV, Department of Computer Science, University of York, United Kingdom

Personal data sovereignty frameworks enable individuals to discover and govern their distributed personal data. Yet these frameworks face a practical limitation: policy expression does not guarantee policy enforcement. Individuals can define governance rules for their data, but they lack mechanisms to verify whether organisations actually respect these preferences. This paper presents a blockchain-based verification component that addresses this accountability gap within a data sovereignty framework. Building on ontology-based data federation and ODRL policy control, we introduce immutable audit logging through Ethereum smart contracts. The blockchain component records policy definitions and access decisions, creating an auditable trail that allows data subjects to review how their data is being accessed. We acknowledge that organisations can bypass such systems by querying their databases directly. The blockchain cannot prevent this. What it provides is accountability infrastructure: transparent, tamper-proof records of all access that does pass through the system. We present the design of this verification mechanism, and show that this combination of semantic federation, machine-readable policies, and blockchain accountability does not exist in current personal data sovereignty frameworks.

CCS Concepts: • **Security and privacy** → **Access control**; • **Information systems** → *Data management systems*;

Additional Key Words and Phrases: Personal Data Sovereignty, Blockchain, Policy Verification, Audit Logging, ODRL

ACM Reference Format:
Vijon Baraku, Iraklis Paraskakis, Simeon Veloudis, and Poonam Yadav. 2018. Transparent Accountability for Personal Data Sovereignty: Blockchain-Based Verification of Policy Compliance. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 9 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

In the digital economy, personal data functions as key input to value creation. Organisations collect and process information about individuals to improve services and inform decision-making. Regulatory frameworks like the General

Authors' Contact Information: Vijon Baraku, Department of Computer Science, University of York, York, United Kingdom, vibaraku@seerc.org; Iraklis Paraskakis, SEERC – South East European Research Centre, Thessaloniki, Greece and University of York Europe Campus, Thessaloniki, Greece, iparaskakis@seerc.org; Simeon Veloudis, SEERC – South East European Research Centre, Thessaloniki, Greece and University of York Europe Campus, Thessaloniki, Greece, sveloudis@seerc.org; Poonam Yadav, Department of Computer Science, University of York, York, United Kingdom, poonam.yadav@york.ac.uk.

Manuscript submitted to ACM

Data Protection Regulation (GDPR), the Data Governance Act, and the AI Act have established principles for data protection and data sharing. Yet, individuals still lack practical means of controlling how their personal information is handled across organisational boundaries. The problem grows worse as personal data becomes fragmented across service providers, each with their own database systems and storage formats.

Research in personal data sovereignty has explored various approaches to enabling individuals to and govern its use. Policy languages such as the Open Digital Rights Language (ODRL) [9], which is an ontology-based information model that enables interoperable specification of data usage policies, provide standardised mechanisms for specifying usage rules. Our prior work [1, 4] introduced ontology-based data federation to enable individuals to discover their personal data across heterogeneous sources without moving the data itself, combined with ODRL policies for usage governance. These capabilities provide visibility and control specification. They do not provide verification.

To illustrate this gap, consider Alice, a patient whose medical records are held by a hospital that has adopted a data sovereignty framework. She has visibility: through the federated view, she sees her health data across the hospital's systems. She has control specification: she applies a policy prohibiting AI training on her records, and the policy is stored. The hospital acknowledges it. But acknowledgment is necessary, not sufficient. How does Alice know the hospital actually respects this prohibition? The policy exists, but she has no mechanism to confirm whether it is being enforced. We need to close the loop and ensure Alice that her expressed wishes are adhered to.

This paper presents a blockchain-based verification component that addresses this accountability gap. We do not claim to prevent policy violations: organisations control their own databases, and no external mechanism can stop them from accessing data they physically possess. What we provide is a mechanism that paves the way towards greater transparency and non-repudiation over personal data processing, which in turn leads to accountability. Every access request processed through the system is logged immutably. Data subjects can review who accessed their data, when, for what purpose, and whether access was permitted or denied.

This work extends our prior research [1–4] on ontology-based data federation and ODRL-based policy control by introducing a third component: *blockchain-based verification*. Our contribution is threefold. First, we present an architecture that integrates blockchain audit logging with semantic data federation and policy control. Second, we implement smart contracts that record both policy definitions and access decisions, linking governance specifications to enforcement outcomes. Third, we discuss what blockchain verification can achieve in this context and its limitations.

The remainder of this paper is structured as follows. Section 2 provides background on personal data sovereignty and the verification challenge. Section 3 describes the framework architecture. Section 4 details the blockchain verification component. Section 5 discusses limitations and future directions, and Section 6 concludes

2 Background

2.1 Personal Data Sovereignty

Personal data sovereignty refers to the ability of individuals to exercise meaningful control over how their personal data is collected, stored, processed, and shared [8]. Realising this ability depends on technical mechanisms that enable this sovereignty.

The challenge grows with data fragmentation. An individual's personal information is typically spread across many service providers. Each maintains a separate database infrastructure. Without mechanisms for unified discovery, individuals cannot understand the full scope of their digital footprint, let alone govern it.

Two architectural approaches have emerged. The first, represented by the Solid project [11, 16], implements personal data stores called pods where individuals centralise their information under direct control. Solid provides clear ownership boundaries, but requires data migration and ecosystem adoption. The second approach appears in the European data spaces initiative [5, 10], where frameworks like the International Data Spaces Association (IDSA) [14] keep data at its original location while establishing protocols for controlled exchange. These data spaces focus on business-to-business data sharing and industrial ecosystems rather than individual sovereignty.

Our framework [1–4] takes a different path from personal data stores: rather than requiring data migration, personal data remains in organisational databases. In this respect, we share IDSA's principle of keeping data at its source. However, where IDSA focuses on business-to-business data exchange, we centre on individual data subjects, using Schema.org [7] as a common vocabulary to create semantic bridges between diverse database schemas. This preserves organisational data ownership while enabling individual discovery and governance.

2.2 The Verification Gap

Existing approaches to personal data sovereignty share a common limitation: they focus on policy specification and access control, not on enforcement verification or auditing compliance.

Solid's primary goal is giving data subjects control over who can access their data stored in pods [16]. This is implemented through Access Control Lists (ACLs) or, more recently, Access Control Policies (ACP). These mechanisms determine which agents can read, write, or append to resources within a pod. However, Solid focuses exclusively on access control, not usage control. Once an agent is granted access to data, Solid provides no mechanism to specify or enforce constraints on how that data may subsequently be used. There is no support for usage control policies that might, for example, prohibit AI training or require data deletion after a specified period. Consequently, there is no mechanism for verifying or auditing compliance with such policies, because the policies themselves cannot be expressed within the framework.

IDSA takes a different approach. The framework explicitly recommends ODRL for expressing usage constraints in a standardised, machine-readable way, making ODRL essentially standard practice within IDSA-compliant implementations. This enables rich policy specification including purpose limitations, temporal constraints, and usage prohibitions. However, IDSA does not provide explicit support for verifying or enforcing such policies. The Clearing House component offers logging capabilities for data transactions, but this logging remains within the IDSA ecosystem and under data provider control. No independent, immutable audit trail exists that would allow data subjects to verify whether their usage policies are being respected [15].

Ocean Protocol [12] provides blockchain-based transparency through its data marketplace infrastructure. Data providers can tokenise datasets and define access terms via smart contracts, with all transactions recorded on-chain. However, Ocean Protocol operates as a data marketplace where providers actively choose to list their data for exchange or computation. It does not address the scenario where an individual's data already resides in organisational databases without their direct involvement. There is no mechanism for discovering personal data across heterogeneous sources, and the semantic layer required for unified querying across different database schemas is absent. Ocean Protocol solves data monetisation transparency, not personal data sovereignty.

Even the GDPR [6], while establishing comprehensive legal rights for data subjects, relies on reactive enforcement. Violations are discovered after the fact through individual complaints, regulatory audits, or data breaches. The regulation provides no technical infrastructure for proactive verification of compliance.

The underlying challenge is straightforward. Organisations remain custodians of personal data and control access to their systems. No technical framework can completely prevent an organisation from accessing data it physically possesses. What frameworks can do is establish accountability mechanisms that create consequences for non-compliance and enable detection of policy violations.

2.3 Blockchain for Accountability

Blockchain technology offers properties suited to accountability requirements: immutability, transparency, and decentralisation [19]. Platforms like Ethereum [18] extend these properties with smart contracts, enabling programmable logic that executes automatically when conditions are met. Once recorded, entries cannot be modified or deleted. All participants can verify recorded information. No single party controls the audit trail.

These properties have been applied in various data governance contexts. Third and Domingue explored combining Solid pods with blockchain for data verification [17]. Self-Sovereign Identity (SSI) systems use blockchain for credential verification [13]. However, these applications focus on identity management or data trading, not on auditing policy compliance in federated personal data environments.

This shifts the trust model: rather than relying on assurances that policies are enforced, data subjects can verify whether observed access patterns are consistent with their expressed preferences.

3 Framework Architecture

Our data sovereignty framework comprises three components that address complementary requirements:

(1) **Ontology-Based Data Federation (OBDF)**: Enables unified discovery of personal data across heterogeneous organisational databases using Schema.org vocabulary and OBDA mappings.
(2) **Personal Data Policy Control (PDPC)**: Enables governance specification through ODRL policies with extensions for AI training restrictions, purpose limitations, and temporal constraints.
(3) **Blockchain Verification Service (BVS)**: Provides immutable audit logging of policy definitions and access decisions.

The first two components have been presented in prior work [3, 4]. This paper focuses on the third component and its integration with the existing architecture.

3.1 System Overview

Figure 1 illustrates the architecture. Data subjects (individuals) interact with the system to view their personal data and define governance policies. Data controllers (organisations) register their databases, provide schema mappings, and request data access through the framework. The component interactions through which subjects view their personal data and controllers register their databases are omitted here but can be found in [1–4].

3.2 Component Interaction

The components interact through defined workflows:

Policy Creation Flow: When a subject creates a governance policy, the PDPC service generates ODRL representations stored in the triple store. The BVS hashes the policy content and records it on the blockchain via the PolicyRegistry smart contract. This creates an immutable record linking policy specifications to their creation timestamp and subject identity.

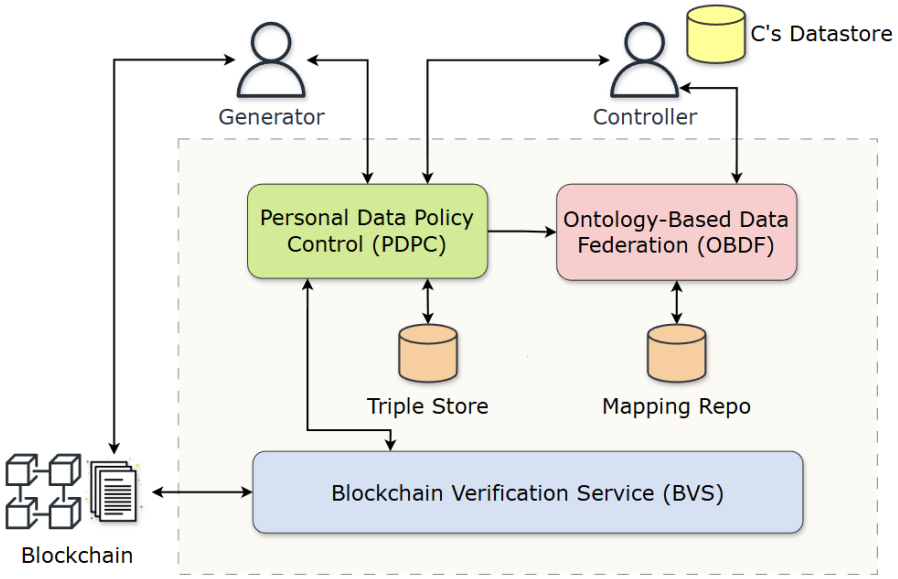

Fig. 1. Framework architecture showing integration of federation, policy control, and blockchain verification.

Access Request Flow: When a controller requests access to personal data, the Policy Decision Point (PDP) within PDPC evaluates the request against applicable policies. The decision, along with contextual information such as controller identity, requested action, purpose, and timestamp, is recorded via the AccessLogger smart contract. Both permitted and denied access attempts are logged.

Verification Flow: Subjects can query their access history through the blockchain, obtaining an immutable record of all access requests to their personal data. Each entry includes the transaction hash, enabling independent verification.

4 Blockchain Verification Component

4.1 Design Rationale

The blockchain verification component is designed around three principles:

Completeness: All access requests processed through the framework are logged, regardless of outcome. Denied access attempts are recorded alongside permitted ones. This provides complete visibility into data access patterns and prevents selective logging where only favourable decisions are recorded.

Independence: The audit trail exists on a blockchain independent of any single organisational system. Neither data subjects nor data controllers can unilaterally modify records. This independence is essential because audit logs controlled by the party being audited have limited evidentiary value.

Linkage: Audit entries are linked to specific policies and policy versions. When a policy is updated, the new version receives a new identifier. Access decisions reference the policy version that was in effect at evaluation time. This enables verification that access decisions were made according to the policies in effect at the time, not policies created or modified after the fact.

4.2 Smart Contract Architecture

The verification component deploys two smart contracts on an Ethereum-compatible blockchain.

PolicyRegistry - This contract maintains an on-chain record of all policies created by data subjects. When a subject defines a new policy through the PDPC interface, the system computes a SHA-256 hash of the complete policy content and submits this hash to the PolicyRegistry. The contract stores the hash alongside the subject's blockchain address, a unique policy identifier, a version number, and the block timestamp. Storing only the hash rather than the full policy content minimises on-chain storage costs while still enabling integrity verification. The full policy remains in the triple store, and any tampering would cause its hash to diverge from the recorded value. Each policy modification increments the version number, creating a complete version history on-chain. The contract emits an event upon each registration, enabling external systems to monitor policy creation activity.

AccessLogger - This contract records every access decision made by the Policy Decision Point. Each log entry captures the controller's blockchain address, the subject's address, a hash of the stated access purpose, the requested action, the decision outcome, the applicable policy group identifier, and the policy version used for evaluation. The purpose is hashed rather than stored in plaintext to avoid placing potentially sensitive information on a public ledger while still enabling verification that the recorded purpose matches what was claimed. The contract maintains mappings that allow efficient retrieval of all access logs for a particular subject or by a particular controller. Events emitted on each logging operation enable real-time monitoring and integration with external alerting systems.

4.3 Integration with Policy Enforcement

The blockchain component integrates with the Policy Decision Point through a defined sequence that ensures every access attempt is recorded.

When a controller submits an access request, the PDP first retrieves all applicable policies from the triple store for the requested data element and subject. The PDP then evaluates whether the requested action is permitted given the policy permissions, constraints such as purpose restrictions and expiration dates, and any AI-specific prohibitions. This evaluation produces a decision: permit or deny.

Regardless of the outcome, the PDP invokes the BlockchainService to record the access attempt. The service constructs the log entry with all relevant contextual information and submits it to the AccessLogger contract. The blockchain processes this transaction and returns a transaction hash. This hash serves as a unique, verifiable reference to the recorded decision.

The controller receives the access decision along with the transaction hash. For permitted requests, the controller can proceed to access the data through the federation layer. For denied requests, the controller receives an explanation of why access was refused. In both cases, the transaction hash provides proof that the interaction was recorded immutably.

This design ensures that the blockchain record reflects actual enforcement decisions, not just policy specifications. A subject reviewing their access history sees what controllers actually requested and what the system actually decided, not merely what policies exist.

4.4 Subject Verification Interface

Data subjects access their verification dashboard through the framework's web interface. The dashboard presents a chronological access history showing all requests to the subject's personal data. Each entry displays the requesting controller's identity, the action requested, the stated purpose, whether access was permitted or denied, and the timestamp

of the decision. Subjects can filter this history by controller, action type, or time period to identify patterns or anomalies in how their data is being accessed.

The dashboard also displays the subject's registered policies alongside their blockchain transaction hashes. When a subject creates or modifies a policy, the system records a hash of the policy content on-chain. This hash serves as a tamper-proof reference point: if the policy stored in the triple store were modified without proper recording, its hash would no longer match the blockchain record.

Each access log entry and policy record includes the associated blockchain transaction hash. Subjects with blockchain familiarity can verify these hashes independently using any Ethereum block explorer, confirming that the records exist on-chain and have not been altered. However, we recognise that many users lack the technical background to navigate block explorers or interpret transaction data. To address this, the framework includes a built-in verification tool. This tool accepts a transaction hash, queries the blockchain directly, and presents the results in a readable format. Users can compare the on-chain data against what the dashboard displays, confirming consistency without requiring blockchain expertise.

This approach balances transparency with usability. Technical users retain the option of independent verification through external tools, while non-technical users can still confirm the integrity of their records through the framework's verification interface.

Table 1 summarises how the verification capabilities of our framework compare against existing personal data sovereignty approaches.

Table 1. Comparison with existing data sovereignty frameworks

Feature	Our Framework	Solid	IDSA	Ocean Protocol
Data Location	Original DBs	Personal Pods	With Provider	With Provider
Data Access Method	OBDA Federation	Pod API	Connectors	Compute-to-Data
Policy Language	ODRL (extended)	WAC	ODRL-based	Smart Contracts
Audit Storage	Blockchain	Local logs	Clearing House	Blockchain
Audit Immutability	Yes	No	Partial	Yes
Subject Verification	Yes	Limited	No	Yes
Individual Focus	Yes	Yes	No (B2B)	Partial
Semantic Federation	Yes	No	No	No

5 Discussion

5.1 Limitations

We acknowledge several limitations in what blockchain verification can achieve:

Cannot Prevent Bypass: Organisations control their databases. A controller determined to violate data sovereignty policies can query their database directly, bypassing the framework. The blockchain cannot record access that does not pass through the system.

Trust in Framework Deployment: The verification mechanism assumes the framework is correctly deployed and that access requests are routed through it. Malicious or negligent deployment could undermine verification guarantees.

Purpose Misrepresentation: Access control decisions may be based on stated access purpose. A controller may declare they are accessing personal data for purpose A while their actual intention is purpose B. The blockchain records

the declared purpose, but cannot verify that this declaration reflects the controller's true intent. This limitation is inherent to any purpose-based access control system.

Gas Costs and Scalability: Recording every access decision on-chain incurs transaction costs. While acceptable for demonstration and limited deployment, production systems handling high volumes of access requests would face prohibitive gas fees under this model. Layer-2 architectures such as rollups (e.g. Optimism or Arbitrum) offer a viable path forward: access log entries could be batched and submitted to the layer-2 network, with only periodic commitments anchored to the Ethereum mainchain. Within the BVS component, this would require the BlockchainService to route transactions to a layer-2 RPC endpoint rather than mainnet, while the AccessLogger and PolicyRegistry contracts would be deployed on the layer-2 network. The immutability and transparency guarantees of the audit trail would be preserved, as layer-2 state is ultimately secured by the Ethereum base layer. The primary trade-off is a slight delay in finality, which is acceptable given that audit logging is inherently asynchronous with respect to the access decision itself.

5.2 What Blockchain Verification Achieves

Despite these limitations, blockchain verification provides meaningful benefits:

Accountability Infrastructure: Immutable records of access patterns establish accountability even when enforcement is imperfect. Organisations that route access through the framework cannot later deny or hide their data access patterns.

Detection Capability: Regular review of access logs by data subjects can reveal unexpected access patterns, potentially indicating policy violations or framework bypass.

Regulatory Foundation: The technical infrastructure could support future legislation requiring organisations to route personal data access through transparent logging mechanisms. Our implementation shows that such requirements are technically feasible.

Behavioural Incentives: Transparency may influence organisational behaviour. Controllers aware that their access patterns are recorded immutably may be more likely to respect subject preferences.

6 Conclusion

This paper presented a blockchain-based verification component for personal data sovereignty frameworks. By integrating immutable audit logging with ontology-based data federation and ODRL policy control, the framework enables data subjects to verify how their personal data is being accessed across organisational boundaries.

We have been explicit about what blockchain verification can and cannot achieve. It cannot prevent organisations from bypassing the system. They control their databases. What it provides is accountability infrastructure that makes access patterns transparent and auditable. This changes the trust model from relying on policy enforcement to verifying evidence of data access.

The combination of semantic federation, machine-readable policies, and blockchain accountability represents a contribution not present in existing personal data sovereignty frameworks. As regulatory requirements evolve and individuals seek greater transparency over their personal data, such technical infrastructure becomes increasingly relevant.

Our implementation shows that meaningful accountability is achievable within current technological constraints. Future work should explore how verification can complement technical enforcement mechanisms and how the approach can scale to production environments.

Acknowledgments

This work is supported by the University of York and SEERC - South East European Research Centre. This work is also supported, in part, by EPSRC and DSIT TMF-uplift: CHEDDAR: Communications Hub For Empowering Distributed ClouD Computing Applications And Research (EP/X040518/1), (EP/Y037421/1), EPSRC IAA (EP/X525856/1), and EPSRC REMOTE (EP/Y019229/1).

References

[1] Vijon Baraku, Iraklis Paraskakis, Simeon Veloudis, and Poonam Yadav. 2024. Personal Data Sovereignty in Virtual Enterprises: Implementing Data Capsules for Enhanced Privacy and Compliance. In *Navigating Unpredictability: Collaborative Networks in Non-linear Worlds (PRO-VE 2024) (IFIP Advances in Information and Communication Technology, Vol. 726)*. Springer, 447–461. doi:10.1007/978-3-031-71739-0_29

[2] Vijon Baraku, Iraklis Paraskakis, Simeon Veloudis, and Poonam Yadav. 2024. Responsible Information Sharing in the Era of Big Data Analytics Facilitating Digital Economy Through the Use of Blockchain Technology and Observing GDPR. In *Proceedings of the 14th International Conference on Cloud Computing and Services Science (CLOSER 2024)*. SCITEPRESS, 257–264.

[3] Vijon Baraku, Iraklis Paraskakis, Simeon Veloudis, and Poonam Yadav. 2026. Extending Personal Data Sovereignty by Enabling Governance of AI Training on Personal Data. In *Hybrid Human-AI Collaborative Networks*, Luis M. Camarinha-Matos, Angel Ortiz, Xavier Boucher, and Antonio Lucas Soares (Eds.). Springer Nature Switzerland, Cham, 19–35.

[4] Vijon Baraku, Edon Ramadani, Iraklis Paraskakis, Simeon Veloudis, and Poonam Yadav. 2025. Defining Personal Data Sovereignty: An Ontologically-Based Framework Facilitating Subject Privacy Control. *Data and Information Management* 9, 3 (2025), 100066. doi:10.1016/j.dim.2025.100066

[5] Edward Curry. 2020. *Real-Time Linked Dataspaces: Enabling Data Ecosystems for Intelligent Systems.* Springer. doi:10.1007/978-3-030-29665-0

[6] European Parliament and Council of the European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). https://eur-lex.europa.eu/eli/reg/2016/679/oj Official Journal of the European Union L 119.

[7] Ramanathan V. Guha, Dan Brickley, and Steve Macbeth. 2016. Schema.org: Evolution of Structured Data on the Web. *Commun. ACM* 59, 2 (2016), 44–51. doi:10.1145/2844544

[8] Patrik Hummel, Matthias Braun, Max Tretter, and Peter Dabrock. 2021. Data Sovereignty: A Review. *Big Data & Society* 8, 1 (2021), 1–17. doi:10.1177/2053951720982012

[9] Renato Iannella and Serena Villata. 2018. *ODRL Information Model 2.2.* W3C Recommendation. World Wide Web Consortium (W3C). https://www.w3.org/TR/odrl-model/

[10] Matthias Jarke. 2020. Data Sovereignty and the Internet of Production. In *Advanced Information Systems Engineering (CAiSE 2020) (Lecture Notes in Computer Science, Vol. 12127)*. Springer, 549–558. doi:10.1007/978-3-030-49435-3_34

[11] Essam Mansour, Andrei Vlad Sambra, Sandro Hawke, Maged Zereba, Sarven Capadisli, Abdurrahman Ghanem, Ashraf Aboulnaga, and Tim Berners-Lee. 2016. A Demonstration of the Solid Platform for Social Web Applications. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*. ACM, 223–226. doi:10.1145/2872518.2890529

[12] Trent McConaghy et al. 2019. *Ocean Protocol: A Decentralized Substrate for AI Data and Services.* Technical Report. Ocean Protocol Foundation. https://oceanprotocol.com/tech-whitepaper.pdf

[13] Alexander Muhle, Andreas Gruner, Tatiana Gayvoronskaya, and Christoph Meinel. 2018. A Survey on Essential Components of a Self-Sovereign Identity. *Computer Science Review* 30 (2018), 80–86. doi:10.1016/j.cosrev.2018.10.002

[14] Boris Otto and Matthias Jarke. 2019. Designing a Multi-Sided Data Platform: Findings from the International Data Spaces Case. *Electronic Markets* 29, 4 (2019), 561–580. doi:10.1007/s12525-019-00362-x

[15] Boris Otto, Sebastian Steinbuss, Andreas Teuscher, and Steffen Lohmann. 2019. *IDS Reference Architecture Model Version 3.0.* Technical Report. International Data Spaces Association. https://internationaldataspaces.org/publications/

[16] Andrei Vlad Sambra, Essam Mansour, Sandro Hawke, Maged Zereba, Nicola Greco, Abdurrahman Ghanem, Dmitri Zagidulin, Ashraf Aboulnaga, and Tim Berners-Lee. 2016. *Solid: A Platform for Decentralized Social Applications Based on Linked Data.* Technical Report. MIT CSAIL and Qatar Computing Research Institute.

[17] Allan Third and John Domingue. 2020. Towards Complete Decentralised Verification of Data with Confidentiality: Different Ways to Connect Solid Pods and Blockchain. In *Companion Proceedings of the Web Conference 2020 (WWW '20 Companion)*. ACM, 645–649. doi:10.1145/3366424.3385759

[18] Gavin Wood. 2014. Ethereum: A Secure Decentralised Generalised Transaction Ledger. *Ethereum Project Yellow Paper* 151 (2014), 1–32.

[19] Zibin Zheng, Shaoan Xie, Hongning Dai, Xiangping Chen, and Huaimin Wang. 2017. An Overview of Blockchain Technology: Architecture, Consensus, and Future Trends. In *2017 IEEE International Congress on Big Data (BigData Congress)*. IEEE, 557–564. doi:10.1109/BigDataCongress.2017.85

Received 18 January 2026