



The Solid Protocol and the Decentralized Web: Transforming Data Ownership, Interoperability, and Control

¹Dr. K. Chitra, B.Tech, MBA, Ph.D

¹Research Associate, The Mother's Service Society, Pondicherry

Abstract – The contemporary Web is decentralized at the level of protocols but highly centralized at the level of data custody, identity, application ecosystems, and economic power. Solid (Social Linked Data), initiated by Tim Berners-Lee and developed through the Solid community, proposes a different application architecture: data is stored in independently hosted Personal Online Datastores (Pods), identified and linked through Web standards, while applications obtain permissioned access rather than becoming the permanent custodians of user data. This article provides a detailed analysis of Solid as both a technical protocol family and a data-governance model. It updates earlier descriptions of Solid by distinguishing WebID-based identification, Solid-OIDC authentication, and authorization through Web Access Control (WAC) and Access Control Policy (ACP). It examines RDF and Linked Data, data portability, interoperability, security, privacy, user experience, decentralized discovery, performance, governance, and regulatory alignment. Recent evidence is incorporated from the Flemish data-vault ecosystem, which targets approximately 6.5 million citizens; 2024-2026 studies of Solid usability and decentralized search; and developments in European data regulation, including the European Health Data Space (EHDS), which entered into force in March 2025. The analysis finds that Solid offers a technically credible mechanism for separating data from applications and enabling reusable, permissioned data, but that decentralizing storage does not by itself solve semantic interoperability, usability, trust, cybersecurity, discovery, governance, or platform-network effects. Solid is therefore best understood not as a replacement for the Web or as a blockchain system, but as an evolving standards-based layer for user-centric data spaces. Its long-term significance will depend on interoperable implementations, understandable permission models, institutional governance, viable business incentives, and evidence that decentralized data control can operate reliably at population scale.

Keywords - Solid, decentralized web, Personal Online Datastore, Pod, Linked Data, RDF, WebID, Solid-OIDC, WAC, ACP, data sovereignty, interoperability, privacy, personal data stores, data spaces.

I. INTRODUCTION

The original Web was designed around open protocols that allowed independently operated servers and clients to interoperate. Yet the dominant application economy that developed on top of those protocols increasingly concentrated personal data inside platform-specific databases. The result is a structural asymmetry: users can move between websites, but their identity, social relationships, preferences, histories, credentials, health information, and behavioral traces are frequently difficult to move or reuse across services. The uploaded draft correctly identifies privacy, data silos, and interoperability as central motivations for Solid, but a contemporary account requires a more precise distinction between decentralizing infrastructure and decentralizing control [1]-[4].

Solid addresses this asymmetry by separating application functionality from persistent data custody. Instead of an application being both the interface and the authoritative database for a person's information, a Solid application can interact with resources stored in a Pod under explicit authorization. A Pod is therefore not merely "cloud storage." It is a Web-accessible data environment whose resources can be addressed with URLs, represented using standard Web formats, linked semantically, and protected through authorization mechanisms [1], [5]. This

architectural separation creates the possibility that several applications can operate on the same underlying user-controlled data without each maintaining an isolated copy. The research problem is consequently broader than whether Solid can technically store files. The important question is whether an open, decentralized data layer can simultaneously provide usable identity, granular authorization, semantic interoperability, scalable discovery, security, legal accountability, and viable institutional adoption. This article examines those dimensions using specification analysis, recent peer-reviewed research, regulatory evidence, and contemporary deployment cases.

II. RESEARCH QUESTIONS AND METHOD

The study is organized around five research questions: (RQ1) How does the current Solid architecture redistribute data custody and control compared with conventional platform-centric systems? (RQ2) Which standards and protocol components enable identity, authentication, authorization, Linked Data interoperability, and application-data separation? (RQ3) What do recent empirical studies reveal about scalability, search, usability, transparency, privacy, and security? (RQ4) How is Solid moving from research prototypes toward institutional deployments and regulated data spaces? (RQ5) Which



unresolved technical, economic, governance, and human-factors barriers are most likely to determine adoption?

Methodologically, this is a structured narrative and technical review rather than a systematic meta-analysis. Evidence was triangulated from the Solid technical reports and editor's drafts, peer-reviewed literature published primarily from 2023-2026, official regulatory sources, and documented deployments. Vendor case studies are used only for factual deployment descriptions and are explicitly distinguished from independent empirical evidence. Because several Solid documents remain Community Group drafts or editor's drafts, specification maturity is treated as a variable rather than assuming that every component is a finalized W3C Recommendation [1], [6].

Table 1: Analytical Dimensions Used In This Study

Dimension	Central Question	Evidence Used
Architecture	Where does data reside and how do apps interact with it?	Solid Protocol, LDP/RDF, Pod model
Identity & access	How are agents identified, authenticated, and authorized?	WebID, Solid-OIDC, WAC, ACP
Interoperability	Can data be reused across independent applications?	Linked Data, vocabularies, semantic constraints
Performance	Can decentralized resources be discovered and queried efficiently?	SolidBench/LTQP and decentralized-search research
Human factors	Do users understand and trust data-vault models?	2024-2025 usability and perception studies
Governance	Who is accountable for hosting, policy, security, and standards?	GDPR, EU Data Act, EHDS, DPDPA, deployment cases

III. WHY THE EXISTING WEB BECAME DATA-CENTRALIZED

1. Protocol Decentralization Versus Data Centralization

The Internet and Web do not require a single owner, but successful digital services benefit from network effects, centralized analytics, vertically integrated identity systems, and economies of scale. Consequently, a technically

decentralized network can support highly centralized application markets. In a conventional software-as-a-service model, the service provider typically controls the application interface, database schema, authentication relationship, data-processing logic, and export mechanism. This coupling lowers engineering complexity for the provider, but it also makes switching expensive and cross-service reuse difficult.

This distinction matters because Solid does not propose that every user run a physical server at home. A Pod can be self-hosted, institution-hosted, or commercially hosted. The decentralization objective is therefore primarily architectural and governance-oriented: storage and identity should not be inseparably bound to one application vendor. Users or accountable institutions should be able to select data-service providers while applications remain substitutable [3], [7].

2. Data Portability Is Not the Same as Interoperability

A downloadable archive satisfies only a weak form of portability. True application substitutability requires syntactic and semantic compatibility: the receiving application must know what the exported fields mean, how entities relate, and which identifiers persist. GDPR Article 20 gives qualifying data subjects a right to receive data in a structured, commonly used, machine-readable format and, where technically feasible, to transmit it directly between controllers [18]. The EU Data Act further strengthens switching and interoperability obligations for data-processing services [19]. Solid approaches the same policy problem from an architectural direction: rather than repeatedly exporting copies from application silos, data can remain in a user-associated store and be reused by authorized applications.

IV. SOLID ARCHITECTURE IN DETAIL

1. Pods and Resource-Oriented Storage

A Solid Pod is a Web-accessible storage space containing resources and containers. Resources are addressed through HTTP(S) URIs and can include RDF documents as well as other media types. The protocol builds on established Web concepts rather than requiring a proprietary transport network. This is important: Solid is not synonymous with blockchain, peer-to-peer storage, or a cryptocurrency. A Pod may run on ordinary cloud infrastructure; decentralization arises from the ability to distribute custody across independent providers and preserve interoperable access conventions [1], [5].

The architectural consequence is application-data decoupling. In a conventional service, changing applications often means migrating a database. Under the Solid model, a user could theoretically grant a new application access to the same authorized resources. This changes the unit of competition from "who possesses the



dataset” toward “who provides the most useful service over permissioned data.” That possibility is one of Solid’s most consequential economic implications, although realizing it requires shared semantics and stable authorization behavior.

2. WebID: Identification, Not Authentication

WebID is frequently described too loosely. A WebID is an HTTP(S) URI denoting an agent such as a person, organization, group, device, or software entity. Dereferencing it can return an RDF profile document. Its primary role is global identification. Possession of a WebID alone does not prove that a requester controls that identity; authentication requires a separate mechanism [1], [2].

3. Solid-OIDC Authentication

Current Solid authentication is centered on Solid-OIDC, which builds on OAuth 2.0 and OpenID Connect. The specification addresses a distinctive decentralized condition: a resource server, identity provider, and application may have no pre-established trust relationship. Solid-OIDC therefore combines WebID with established authorization-code flows and contemporary mechanisms such as PKCE and Demonstration of Proof-of-Possession (DPoP). DPoP binds tokens to cryptographic keys, reducing the usefulness of a stolen bearer token [2]. Earlier Solid implementations also used WebID-TLS, but the current protocol documentation identifies Solid-OIDC as the recommended approach [1].

4. Authorization: WAC and ACP

Authentication answers “who is making the request?”; authorization answers “what may that agent do?” The distinction is fundamental. Web Access Control (WAC) uses an ACL model to describe authorization conditions for HTTP resources. The May 2024 WAC report describes decentralized, cross-domain authorization over Linked Data resources [6]. Access Control Policy (ACP) provides a policy language for describing, controlling, and granting access and remains an evolving editor’s draft [8]. The Solid Protocol permits servers to support WAC, ACP, or both. Consequently, the uploaded draft’s presentation of ACLs as the sole authorization mechanism should be treated as historically incomplete rather than as the full contemporary architecture.

Table 2: Core Solid Components and Their Functions

Component	Primary Function	Important Clarification
Pod	Stores Web-addressable resources	Can be self-, institution-, or provider-hosted
WebID	Globally identifies an agent	Identifier; not itself proof of authentication
Solid-OIDC	Authenticates	Builds on OAuth

	users/clients	2.0/OIDC; uses modern security mechanisms
WAC	Resource authorization	ACL-based decentralized access control
ACP	Policy-based authorization	Alternative/evolving authorization policy model
RDF/Linked Data	Semantic representation and linking	Interoperability still depends on shared vocabularies
HTTP/LDP concepts	Resource operations and containers	Reuses ordinary Web architecture

V. LINKED DATA AND SEMANTIC INTEROPERABILITY

Solid’s use of Linked Data is more than a file-format choice. RDF represents information as subject-predicate-object statements, enabling identifiers to be shared across datasets and relationships to be machine-readable. If two applications agree that a URI denotes the same person, organization, concept, or property, data produced in one context can be meaningfully connected in another. This is a different interoperability strategy from bilateral API integration, where every pair of systems may require a custom mapping.

However, RDF does not magically create semantic agreement. Two applications can represent the same concept using different vocabularies, shapes, units, or assumptions. Solid therefore moves part of the interoperability problem upward: transport and resource access can be standardized while communities still need domain vocabularies, shape constraints, validation rules, and application interoperability conventions. The Solid Application Interoperability work explicitly addresses these higher-level concerns and references mechanisms such as SHACL, Shape Trees, and Shape Expressions [9]. This distinction explains why “data portability” and “semantic interoperability” should not be conflated. A user may possess their data and still be unable to use it elsewhere if receiving software cannot interpret it. The strongest form of data sovereignty therefore requires both control and intelligibility.

VI. SECURITY, PRIVACY, AND THE LIMITS OF DECENTRALIZATION

1. Security Benefits

Solid can reduce some risks created by mass centralization. If data is distributed across many independently controlled Pods, a single application database need not contain complete records for millions of users. Fine-grained authorization can also support data



minimization: an application can be granted access only to the resources necessary for a defined function. Separation between identity providers, applications, and storage providers can further reduce architectural lock-in.

2. Security Risks That Remain

Decentralization redistributes risk rather than eliminating it. A poorly secured Pod provider can still be compromised; misconfigured permissions can expose sensitive resources; compromised identity providers can undermine authentication; malicious applications can misuse legitimately granted access; and users can make permission decisions they do not fully understand. Revocation is also not equivalent to “taking data back” if an authorized application already copied information while access was valid. Security therefore requires secure implementations, transport protection, token handling, logging, policy transparency, vulnerability management, and governance beyond the protocol itself [13], [14].

3. Privacy and Legal Control Are Not Identical to Technical Control

The language of “ownership” is intuitively powerful but legally imprecise. Data-protection regimes typically allocate rights and obligations among data subjects, controllers/fiduciaries, processors, recipients, and regulators rather than treating personal data as ordinary property. Solid can provide technical affordances for access control and portability, but compliance still depends on lawful purpose, transparency, retention, accountability, security, and rights handling. Research evaluating Solid against GDPR has similarly emphasized that technical control mechanisms do not automatically satisfy every legal obligation [12], [13].

VII. RECENT EMPIRICAL EVIDENCE: WHAT HAS CHANGED SINCE EARLY SOLID RESEARCH?

1. User Experience and Explainability

A 2024 exploratory usability study involving ten participants examined Solid-enabled data vaults in a job-application scenario. It identified mental models, trust and brand concerns, and explainability as important UX challenges [15]. The significance of this finding is that permission dashboards alone are insufficient: people need understandable models of what a Pod is, who hosts it, what an application is requesting, what will happen after access is granted, and whether data may be copied elsewhere.

A larger experimental line of research published in 2025 compared conventional cookie-enabled personalization with Solid-Pod-enabled personalization. The reported results found significant improvements in perceived transparency and control, but not automatic improvements in trust, usability, or behavioral intention [16]. This is a

crucial empirical correction to simplistic claims that giving users technical control necessarily produces trust. Trust is socio-technical: it depends on institutions, interfaces, incentives, reliability, and comprehension as well as architecture.

2. Decentralized Search and Query Performance

Centralized platforms can index their own databases globally. A decentralized Pod ecosystem cannot assume that a single search engine is permitted to crawl every private resource. This creates a difficult information-retrieval problem: discovery must respect per-user authorization while remaining fast enough for interactive applications. Research on link-traversal query processing has shown that structural assumptions about Solid environments can enable accurate results within seconds for non-complex queries, while complex queries remain challenging [10].

The ESPRESSO framework attacks the related problem of decentralized search by integrating access control with distributed Pod indexes. A 2024 evaluation reported an enhanced indexing design tested at a prototype scale extending to 50 servers and roughly 475,000 Pods, illustrating that decentralized search can be engineered beyond toy deployments, while still leaving broader Internet-scale questions open [11]. In 2026, work in ACM Transactions on the Web further reported that decentralized retrieval models could achieve competitive top-item relevance under tested access conditions [17]. These results indicate progress, but they should not be interpreted as proof that arbitrary private-data queries across billions of resources have been solved.

Table 3: Selected Recent Empirical Evidence

Study/Year	Evidence	Main Implication
Theys et al., 2024	Usability test, n=10	Mental models, trust, and explainability affect Pod UX
Ragab et al., 2024	Distributed search prototype	Access-aware decentralized search is feasible but architecturally complex
Theys et al., 2025	Comparative user experiment	Pods improved perceived transparency/control; trust/usability gains were not automatic
Bahrani et al., 2026	Decentralized information retrieval	Competitive relevance is possible under evaluated conditions
Taelman & Verborgh, 2023	Link-traversal query benchmark	Simple decentralized queries can be fast; complex planning remains challenging



VIII. REAL-WORLD DEPLOYMENT: FLANDERS AND THE MOVE TOWARD POPULATION-SCALE DATA VAULTS

One of the most important developments for Solid is the movement from university prototypes toward public digital infrastructure. In Flanders, the data utility athumi selected Solid-based technology for a data-vault ecosystem. Inrupt reports that the architecture is intended to provide government-hosted personal data vaults to approximately 6.5 million citizens, with public and private organizations able to request permission for specific data [20]. Because this information comes partly from a technology supplier's case study, it should be treated as deployment evidence rather than independent proof of outcomes.

The case is nevertheless analytically important. Population-scale deployment changes the research questions. Instead of asking whether one Pod can store RDF, the questions become: Can permissions be understood by millions of people? How are guardianship, deceased users, accessibility, identity recovery, fraud, and organizational liability handled? What service-level guarantees apply? Can competing applications rely on stable semantics? Who finances storage and support? How are security incidents coordinated across providers? These governance questions become as important as protocol conformance.

A 2024 information-systems study examining decentralized and interoperable storage in the Flemish ecosystem used stakeholder interviews to identify business affordances and barriers. It highlights revocable access, linked resources, semantic meaning, and cross-organizational data sharing as important capabilities, while showing that value creation depends on an ecosystem of participants rather than a protocol alone [21].

IX. SOLID, HEALTH DATA, AND THE EUROPEAN HEALTH DATA SPACE

Health data illustrates both the appeal and difficulty of personal data stores. Health information is longitudinal, highly sensitive, distributed among institutions, and valuable for both primary care and secondary research. A 2024 review of personal data store ecosystems in health and social care argues that such systems can improve meaningful individual control and visibility while raising implementation and governance challenges [22]. Other

2024-2025 research proposes decentralized personal health data-space architectures and examines how personal stores can support interoperability and user control [23], [24].

The policy environment has also changed substantially. Regulation (EU) 2025/327 establishing the European Health Data Space was published on March 5, 2025 and entered into force on March 26, 2025. The EHDS is designed to increase individuals' access to and control over electronic health data, enable secure secondary use for research and policy, and create a more interoperable market for electronic health-record systems. Major obligations phase in over several years, including priority categories beginning in 2029 [25]. Solid is not mandated by the EHDS, but its architectural goals—user control, interoperability, permissioned data exchange, and separation of applications from data—are directly relevant to the broader policy direction.

X. REGULATORY ALIGNMENT BEYOND EUROPE

The GDPR's data-portability right and the EU Data Act's switching/interoperability provisions create legal pressure against irreversible lock-in [18], [19]. India's Digital Personal Data Protection Act, 2023 similarly emphasizes consent and permits a Data Principal to give, manage, review, or withdraw consent through a registered Consent Manager [26]. These legal regimes differ in scope and terminology, and Solid should not be presented as automatically compliant with any of them. Nevertheless, they reveal a converging policy interest in user agency, transparent permission, accountable processing, and portability.

This convergence creates a potential role for decentralized data architectures: technical infrastructure can make legal rights easier to exercise. But the reverse is equally important—law can provide accountability that protocols cannot. A user interface that exposes a revoke button is not a substitute for legal duties concerning misuse, discrimination, retention, security breaches, or unauthorized onward transfer.

XI. SOLID COMPARED WITH BLOCKCHAIN AND FEDERATED ARCHITECTURES

Table 4: Conceptual Comparison of Data Architectures

Characteristic	Conventional Cloud Platform	Federated Service	Blockchain-Centric System	Solid
Primary data custody	Application/provider	Independent service instances	Replicated/distributed ledger or off-chain stores	Independent Pods/storage providers



Core interoperability	Provider APIs	Federation protocol	Ledger/smart-contract rules	HTTP + Linked Data + Solid specifications
User data control	Provider-mediated	Instance-dependent	Protocol/wallet dependent	Resource-level permission model
Mutability/privacy	Provider controlled	Instance controlled	Public ledgers can complicate deletion/privacy	Ordinary mutable Web resources with access control
Cryptocurrency required	No	No	Often, but not always	No
Main strength	Operational simplicity	Service decentralization	Shared tamper-evident state	Application-data separation and reusable personal data

Solid is often grouped under “Web3,” but this can obscure important differences. Blockchain systems establish shared state through consensus and replicated ledgers; Solid primarily relies on Web resources, identifiers, identity providers, authorization, and linked data. Putting personal data directly on an immutable public ledger can conflict with confidentiality and erasure requirements. Solid instead allows ordinary mutable resources to remain off-chain and access-controlled. Blockchain could be combined with Solid for particular audit or credential functions, but it is not a prerequisite of the protocol.

Federated systems such as email or ActivityPub decentralize service operation: users belong to different servers that communicate through a federation protocol. Solid goes further in a different direction by attempting to separate data storage from the application itself. These models are complementary rather than mutually exclusive.

XII. ECONOMIC AND PLATFORM IMPLICATIONS

If applications no longer possess exclusive custody of user data, competitive dynamics could change. Switching costs may fall because a new application can request access to existing data rather than rebuilding a user profile from zero. Small developers could theoretically innovate over permissioned datasets without first accumulating a proprietary data moat. Users could also maintain continuity across services.

However, incumbent platforms derive value not only from stored data but also from network effects, recommendation systems, proprietary models, integrated workflows, brand recognition, advertising infrastructure, and enormous engineering capacity. Solid therefore does not mechanically dissolve platform power. A decentralized data layer needs compelling applications, low-friction onboarding, reliable hosting, sustainable provider economics, and governance institutions. Without these, users may rationally prefer convenient centralized services.

XIII. THREAT MODEL AND FAILURE MODES

A mature assessment of Solid should evaluate failure modes rather than treating decentralization as an intrinsic security property. At least eight classes deserve attention: (1) identity-provider compromise; (2) token theft or client compromise; (3) authorization misconfiguration; (4) insecure Pod hosting; (5) malicious or deceptive consent interfaces; (6) data exfiltration by an application that had legitimate temporary access; (7) semantic attacks in which misleading or poisoned linked data influences downstream applications; and (8) availability failures when a user-selected provider disappears or is unreachable.

Mitigation requires defense in depth: secure OIDC implementations, DPoP/PKCE where specified, least-privilege authorization, auditable permission histories, encryption in transit and where appropriate at rest, secure software supply chains, backups and portability, trustworthy recovery mechanisms, human-readable consent, and clear incident-response responsibilities. Importantly, a decentralized ecosystem may require stronger coordination standards because responsibility is distributed among more actors.

XIV. THE AI ERA: WHY PERSONAL DATA ARCHITECTURE MATTERS MORE

Generative and agentic AI increase the strategic importance of personal data architecture. Useful personal agents need access to calendars, documents, preferences, credentials, health information, purchase histories, and communication context. If all such context must be copied into the database of each AI provider, the result could intensify data concentration. A Pod-like architecture offers an alternative model in which an authorized agent retrieves only the resources necessary for a task and access can later be revoked.



This opportunity also magnifies risk. AI agents can act across multiple resources at machine speed; permission scopes that appear harmless in isolation may become sensitive when combined; and model providers may retain derived information even after raw-data access is revoked. Future Solid research should therefore consider purpose limitation, delegated authorization, provenance, machine-readable usage policies, privacy-preserving computation, and auditable agent actions. The next stage of decentralized data control may depend as much on safe AI delegation as on human-facing applications.

XV. DISCUSSION: WHAT SOLID SOLVES—AND WHAT IT DOES NOT

Solid directly addresses a structural design problem: the Web lacks a universally adopted user-controlled data layer that is independent of applications. By standardizing access to external storage, identity references, and permissioned resources, Solid provides a coherent mechanism for reducing application-data coupling. Recent deployment and research evidence indicates that this concept is technically viable beyond isolated demonstrations.

Nevertheless, five limits are clear. First, technical control is not equivalent to legal ownership or trustworthy processing. Second, Linked Data enables semantic connection but does not guarantee vocabulary agreement. Third, decentralization makes global search and analytics harder because centralized indexes cannot simply assume unrestricted access. Fourth, granular control can overwhelm users unless permission models are explainable and supported by trusted institutions. Fifth, adoption is a coordination problem: applications, storage providers, identity services, governments, businesses, and users must perceive sufficient value simultaneously.

The strongest interpretation of Solid is therefore neither utopian nor dismissive. It is an architectural experiment in changing where persistent data authority resides. Its success should be measured empirically through portability, user comprehension, security outcomes, interoperability across independent implementations, switching costs, service quality, and the ability to sustain ecosystems at scale.

XVI. PROPOSED EVALUATION FRAMEWORK FOR FUTURE SOLID DEPLOYMENTS

Table 5: Proposed Metrics for Evaluating Solid Ecosystems

Category	Example Metrics
Control	Time to grant/revoke access; permission error rate; percentage of resources with

	understandable policies
Portability	Time/cost to switch Pod provider or application; percentage of data preserved semantically
Interoperability	Cross-vendor conformance; vocabulary/shape compatibility; successful reuse rate
Security	Unauthorized-access incidents; token compromise rate; recovery time; patch latency
Privacy	Data minimization; onward-copy controls; auditability; consent comprehension
Performance	Read/write latency; query completion; discovery time; availability
UX & trust	Task success, perceived control, comprehension, trust, abandonment
Economics	Hosting cost/user; developer integration cost; switching cost; ecosystem participation
Governance	Incident accountability, dispute resolution, certification, regulatory compliance

Future studies should report these metrics rather than relying only on architectural claims. Comparative experiments should test Solid-enabled and conventional architectures using equivalent tasks and realistic populations. Longitudinal research is particularly important because the value of personal data stores emerges over time as multiple applications reuse accumulated data.

Future Research Directions

Priority research areas include: scalable privacy-preserving search; semantic alignment and schema evolution; usable authorization for non-experts; delegated access for AI agents; privacy-preserving analytics over distributed Pods; verifiable credentials and selective disclosure; secure identity recovery; cross-provider conformance testing; offline and mobile-first Pods; accessibility; children and guardianship; posthumous data governance; and economic models for sustainable storage. Research should also examine whether decentralized storage actually reduces breach impact in practice, rather than assuming this from topology.

A second priority is independent evaluation of large deployments. Population-scale initiatives provide an opportunity to measure whether data-vault ecosystems improve service delivery, reduce duplicate data collection, lower integration costs, increase citizen control, or create new burdens. Such studies will be more informative than protocol-level demonstrations alone.



XVII. CONCLUSION

Solid represents one of the most sustained attempts to reintroduce user-centric data architecture into the Web without discarding the Web's existing standards. Its central innovation is conceptually simple but structurally profound: applications should not need to own the persistent data they use. Pods provide externally stored resources; WebIDs provide global identifiers; Solid-OIDC supports decentralized authentication; WAC and ACP address authorization; and Linked Data provides a basis for machine-readable relationships and cross-application reuse.

Evidence through 2026 suggests meaningful progress. Solid-based data vaults are being developed for population-scale public infrastructure in Flanders; empirical studies show gains in perceived transparency and control; decentralized search and query research is advancing; and regulatory developments increasingly emphasize portability, interoperability, and individual access to data. At the same time, the evidence rejects simplistic technological determinism. Trust, usability, semantic agreement, security, discovery, institutional accountability, and network effects remain unresolved or only partially resolved.

The future of the decentralized Web is therefore unlikely to be a binary replacement of centralized services. A more plausible trajectory is a hybrid Web in which applications, institutions, and AI agents operate over portable, permissioned data spaces controlled through interoperable standards. If Solid can mature technically while supporting understandable governance and compelling services, it may become an important layer in that transition. Its ultimate contribution will be judged not by whether data is physically decentralized, but by whether people and institutions can exercise meaningful, secure, interoperable, and durable control over digital information.

REFERENCES

1. S. Capadisli et al., "Solid Protocol," Solid Project, Editor's Draft, ver. 0.12.0, Sep. 9, 2026. [Online]. Available: <https://solidproject.org/ED/protocol>
2. A. Coburn, E. Pavlik, and D. Zagidulin, "Solid-OIDC," Solid Project, work in progress. [Online]. Available: <https://solidproject.org/TR/oidc>
3. T. Berners-Lee, "Solid: Re-decentralizing the Web," Solid Project. [Online]. Available: <https://solidproject.org/>
4. N. Shadbolt, W. Hall, and T. Berners-Lee, "The Semantic Web revisited," *IEEE Intell. Syst.*, vol. 21, no. 3, pp. 96–101, 2006.
5. W3C, "Linked Data Platform 1.0," W3C Recommendation, Feb. 26, 2015.
6. Solid Community Group, "Web Access Control," Draft Community Group Report, May 12, 2024.
7. A. Al-Ruithe, E. Benkhelifa, and K. Hameed, "Personal Data Stores (PDS): A Review," *Sensors*, vol. 23, no. 3, art. 1477, 2023, doi: 10.3390/s23031477.
8. Solid Community Group, "Access Control Policy (ACP)," Editor's Draft, ver. 0.9.0, May 18, 2022.
9. Solid Community Group, "Solid Application Interoperability," Draft specification, 2025.
10. R. Taelman and R. Verborgh, "Evaluation of link traversal query execution over decentralized environments with structural assumptions," 2023, arXiv:2302.06933.
11. M. Ragab et al., "ESPRESSO: A framework to empower search on the decentralized web," *Data Sci. Eng.*, vol. 9, no. 4, pp. 431–448, 2024, doi: 10.1007/s41019-024-00263-w.
12. H. Pandit et al., "Making sense of Solid for data governance and GDPR," *Information*, vol. 14, no. 2, art. 114, 2023, doi: 10.3390/info14020114.
13. "Assessing the Solid Protocol in relation to security and privacy obligations," *Information*, vol. 14, no. 7, art. 411, 2023, doi: 10.3390/info14070411.
14. "Who should I trust with my data? Ethical and legal challenges for innovation in new decentralized data management technologies," *Information*, vol. 14, no. 7, art. 351, 2023, doi: 10.3390/info14070351.
15. T. Theys, T. Haegemans, P. Mechant, A. Bourgeois, and J. Saldien, "Solid-enabled personal online data stores: Uncovering UX implications for online data management," in *Human Systems Engineering and Design*, vol. 158, pp. 100–108, 2024, doi: 10.54941/ahfe1005533.
16. T. Theys et al., "Solid Pods: A promising approach to enhance users' perception of data transparency and control," *Interacting with Computers*, vol. 37, no. 6, pp. 504–517, 2025, doi: 10.1093/iwc/iwaf017.
17. M. Bahrani et al., "Rethinking information retrieval in a re-decentralised web: Exploring the feasibility and quality of search across personal online datastores," *ACM Trans. Web*, vol. 20, no. 1, 2026, doi: 10.1145/3777445.
18. European Parliament and Council, "Regulation (EU) 2016/679 (General Data Protection Regulation)," Art. 20, 2016.
19. European Parliament and Council, "Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act)," 2023.
20. Inrupt, "Flanders Government strengthens a trusted data economy with Inrupt's Enterprise Solid Server," case study, accessed Sep. 2026.
21. "Assessing the value of decentralized and interoperable data storage for service providers," *Information Systems and e-Business Management*, 2024, doi: 10.1007/s10257-024-00692-0.
22. L. Carmichael, W. Hall, and M. Boniface, "Personal data store ecosystems in health and social care,"



- Front. Public Health, vol. 12, art. 1348044, 2024, doi: 10.3389/fpubh.2024.1348044.
23. T. Klementi, G. Piho, and P. Ross, “A reference architecture for personal health data spaces using decentralized content-addressable storage networks,” Front. Med., vol. 11, art. 1411013, 2024, doi: 10.3389/fmed.2024.1411013.
 24. T. Klementi, I. Bossenko, and G. Piho, “Personal Health Data Spaces in decentralized content-addressable storage networks,” Stud. Health Technol. Inform., vol. 324, pp. 105–110, 2025, doi: 10.3233/SHTI250169.
 25. European Parliament and Council, “Regulation (EU) 2025/327 on the European Health Data Space,” Feb. 11, 2025.
 26. Government of India, “Digital Personal Data Protection Act, 2023,” Gazette of India, Aug. 11, 2023.
 27. J. Steinbach, G. De Mulder, B. De Meester, B. G. C. Esteves, and R. Verborgh, “Inter-pod credential exchange protocol via linked data notifications,” in Proc. Semantics in Dataspaces (SDS 2024), CEUR-WS, vol. 3705, 2024.
 28. M. Kask, T. Klementi, I. Bossenko, G. Piho, and P. Ross, “Towards Personal Health Data Spaces,” Stud. Health Technol. Inform., vol. 329, pp. 144–148, 2025, doi: 10.3233/SHTI250818.
 29. D. Shin, M. Park, and H.-J. Yoon, “A Web3-based patient-centric health data management system,” Stud. Health Technol. Inform., vol. 336, pp. 1585–1586, 2026, doi: 10.3233/SHTI260483.
 30. H. Oliver et al., “Awakening the web of self-sovereign data with ESPRESSO: A scoping review of Solid’s and Dataswyft’s readiness for decentralized private search,” IET Conf. Proc., pp. 98–109, 2024, doi: 10.1049/icp.2024.2536.
 31. B. Esteves et al., “The Web and Linked Data as a Solid foundation for dataspaces,” in Companion Proc. ACM Web Conf. 2023, pp. 1440–1446, 2023, doi: 10.1145/3543873.3587616.
 32. European Commission, “European Health Data Space Regulation (EHDS),” 2025–2026 implementation information. [Online]. Available: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en