Self-Sovereign Identity Blockchain-based Architecture for Personal Data Exchange

Silvio Langer · Fabiano Hessel

School of Technology, Pontifical Catholic University of Rio Grande do Sul (PUCRS) · Porto Alegre, Brazil

BLOCKCHAIN FOR INFORMATION SYSTEMS ENGINEERING — BC4ISE26

Problem

Personal data exchange today already depends on centralized intermediaries, creating single points of failure, loss of user control, and an absence of verifiable consent.

Self-Sovereign Identity (SSI) shifts control from centralized entities back to individuals, using blockchain as a trust anchor rather than a trusted third party.

Centralized Model

Single authority holds identity data; users have no direct control.

SSI Model

Identity derived from cryptographic keypairs; blockchain as trust anchor.

Regulatory demands and Data ownership sovereignty

Three converging forces make decentralized identity infrastructure both timely and necessary.

GDPR

Mandates data confidentiality, transparency, and privacy by design, requiring architectures that enforce user rights at the protocol level.

W3C Standardization

Decentralized Identifiers (DIDs) and Verifiable Credentials are now recognized W3C standards, establishing a stable foundation for interoperable SSI systems.

Data Economy

Stakeholders across various sectors need secure, verifiable data sharing with authenticated provenance, a requirement that centralized systems cannot consistently fulfill.

The Gap in the Literature

Four research streams address adjacent problems: blockchain-based access control, SSI frameworks, decentralized storage, and dataspace governance.

DIDs are widely adopted, but limited to identity resolution. The literature leaves a critical gap: **what data is governed, under what conditions, and how identity evolves with exchange events** remain unaddressed as an integrated architecture.

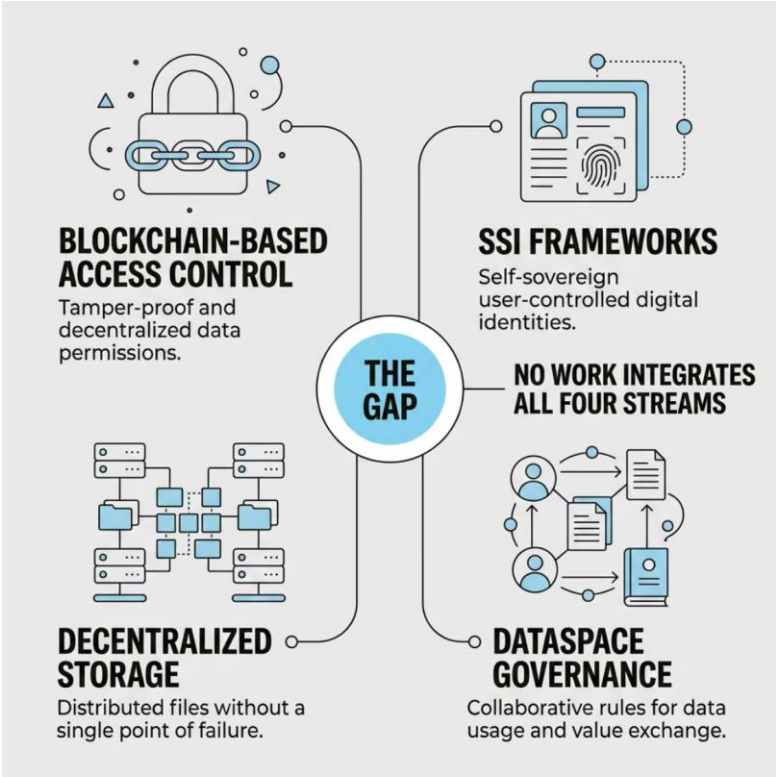

Source: Generated by Gamma (https://gamma.app/)

> **The Gap**
>
> There is little evidence in the literature that unifies DID governance, consent state, metadata provenance, and data exchange into a single coherent architecture.

Research Questions and Contribution

Three research questions drive the architectural design, each mapped to a concrete contribution and mechanism.

Research Question	Contribution	Architectural Mechanism
RQ1: DID Documents as governance layer for personal data exchange.	Extends DID Document role beyond identity resolution to active data governance.	Service endpoints + consent-linked DID Document updates via *VDRRegistry.*
RQ2: Metadata separation from raw personal data with cryptographic provenance and non-repudiation.	Separates governance metadata from raw data while preserving verifiable provenance.	On-chain compact commitments; off-chain IPFS content addressing.
RQ3: Consent as verifiable, tamper-evident state machine whose lifecycle transitions are cryptographically linked to updates of the DID Document..	Models consent as an auditable on-chain lifecycle linked to DID Document updates.	*AgreementRegistry* smart contract enforcing Proposed → Accepted → Revoked transitions.

Architecture: Three Premises

The architecture rests on three foundational design commitments that determine every subsequent technical choice.

DIDs as Stable Coordination Anchors

Decentralized Identifiers provide persistent, resolvable references that remain stable across infrastructure changes.

Consent as On-Chain Agreement State

Smart contracts manage the full consent lifecycle, ensuring consent is explicit, verifiable, revocable, and tamper-evident by construction.

Governance On-Chain; Data Off-Chain

Governance commitments and references are anchored on-chain; raw data artifacts remain off-chain under identity-controlled storage.

Functional Requirements (FR)

Five functional requirements derived from SSI principles govern the architecture's behavior.

ID	Name	SSI Principle	What it requires	How the architecture meets it
FR01	Identity Existence	SSI P1	Identity derived from cryptographic keypairs only; no third-party provider.	secp256k1 keypair generates did:ethr address; no registration authority required.
FR02	Owner Control	SSI P2	Owners update, revoke, and manage DIDs and VCs without intermediary.	VDRRegistry and AgreementRegistry callable only by DID owner's keypair.
FR03	Data Access	SSI P3	VC binds consumer identity to channel, scope, and validity window.	VC issued per channel; access layer verifies VC proof before resolving endpoint.
FR04	Consent	SSI P8	Explicit, verifiable, revocable consent enforced via smart contract.	AgreementRegistry manages full consent state machine with immutable event log.
FR05	Minimalization	SSI P9	Raw data off-chain only; compact references on-chain	IPFS CIDs and terms digests stored on-chain; no personal data in contract state.

Non-Functional Requirements (NFR)

Five non-functional requirements ensure that the architecture maintains SSI principles across various operational aspects.

ID	Name	SSI Principle	What it requires	How the architecture meets it
NFR01	Persistence	SSI P5	Identity and dataset references resolvable over time, with key rotation.	Mutable pointers decouple CID changes from service endpoint references.
NFR02	Protection	SSI P10	User rights prevail across all architectural layers.	Consent revocation immediately blocks access; IPFS unpinning removes raw data.
NFR03	Transparency	SSI P4	Governance, consent, and DID updates publicly auditable on-chain.	All state transitions emit on-chain events; full audit trail without trusted indexer.
NFR04	Interoperability	SSI P7	W3C DID and Verifiable Credential conformance.	did:ethr method and VC data model conform to W3C specifications.
NFR05	Portability	SSI P6	Identity transferable across W3C-compliant infrastructure.	Standard did:ethr and W3C VC format enable cross-platform migration

Architecture Overview

The architecture is organized into four layers, each with a distinct responsibility in the identity and data exchange lifecycle.

01

Identity Layer

Keypair generation, DID derivation, DID Document management

02

Blockchain Layer

Smart contracts: VDRRegistry and AgreementRegistry.

03

Data Layer

IPFS for immutable content; mutable stream state for the current dataset pointer for each channel

04

Access Layer

Consent-gated retrieval with sequential cryptographic verification

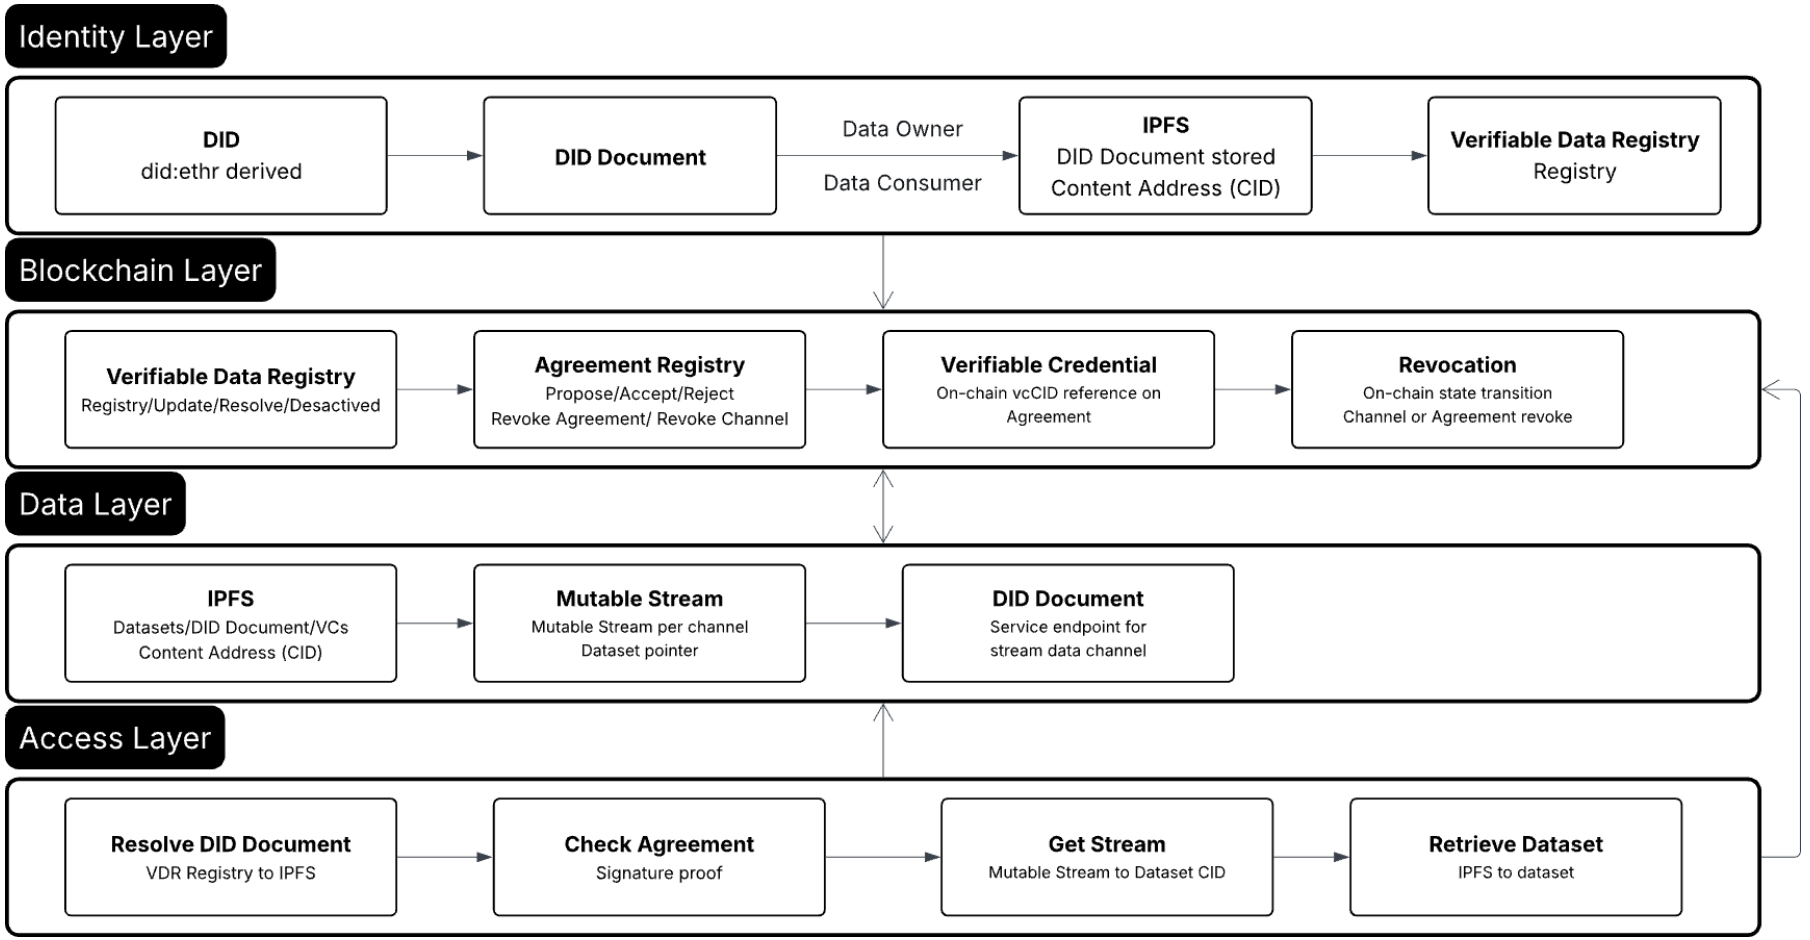

Identity Layer

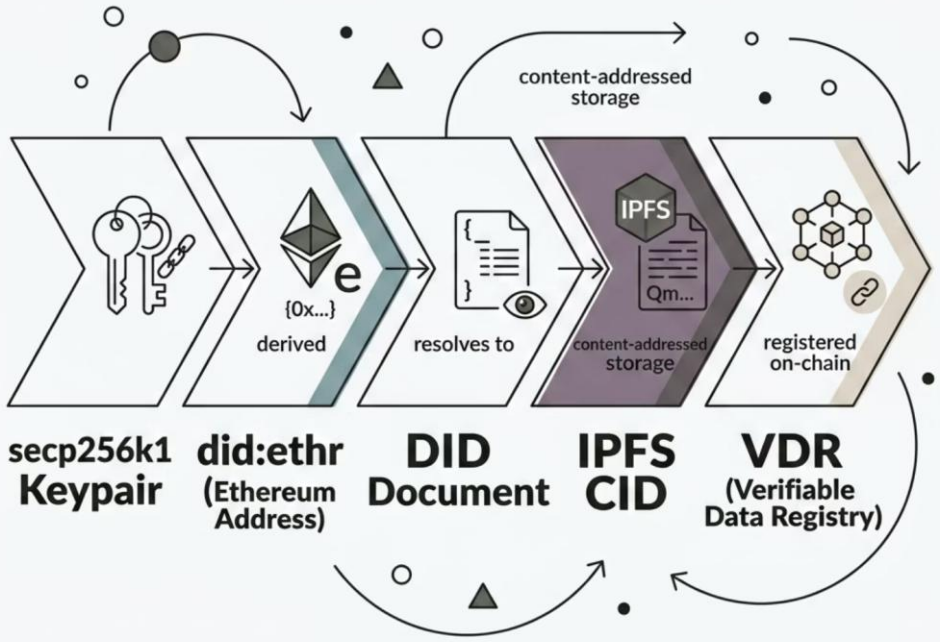

Source: Generated by Gamma (https://gamma.app/)

Key Design Decisions

- A **keypair** derives the did:ethr identifier directly from the Ethereum address;
- The same keypair signs blockchain transactions, asserts DID ownership, and produces VC proofs, satisfying FR01 and NFR04;
- The DID Document is stored on IPFS; its CID is anchored on-chain via VDRRegistry;
- Resolution is dynamic: the current DID Document is recovered by replaying on-chain events. No manual registry maintenance required.

Blockchain Layer:
Smart Contracts and Consent Lifecycle

Two smart contracts implement the entire on-chain governance surface.

1

VDRRegistry

Maps did:ethr to current DID Document CID. Operations: **register · update · deactivate · resolve**. Each operation emits an auditable on-chain event, thereby directly supporting FR01 and NFR03

2

AgreementRegistry

Manages the full consent lifecycle. Stores terms digest, VC CID, validity window, and per-channel revocation status. Enforces FR04, FR02, NFR02, NFR03.

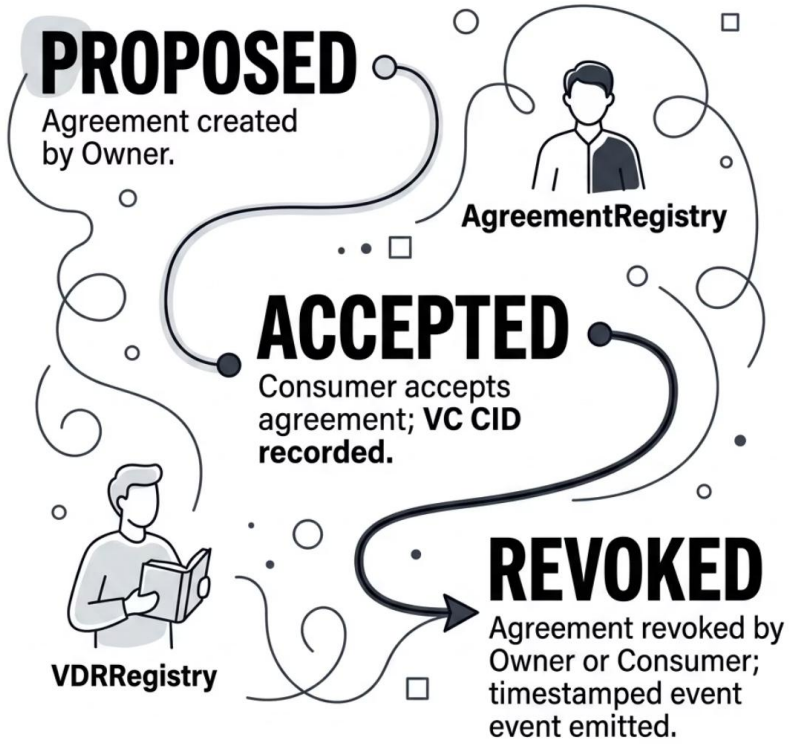

Source: Generated by Gamma (https://gamma.app/)

Revocation is an on-chain state transition: immutable, timestamped, and immediately enforceable at the access layer.

Data Layer:
Solving the Mutability Problem

The Challenge

IPFS provides immutable content-addressed storage, ideal for DID Documents, datasets, and VCs (enforcing FR05). However, a CID changes whenever content changes, breaking service endpoint references stored in DID Documents.

The Solution

Apply a **mutable pointer** per owner-channel pair. Dataset updates change the underlying IPFS CID, but the mutable stream identifier remains stable, decoupling dataset versioning from service endpoint management.

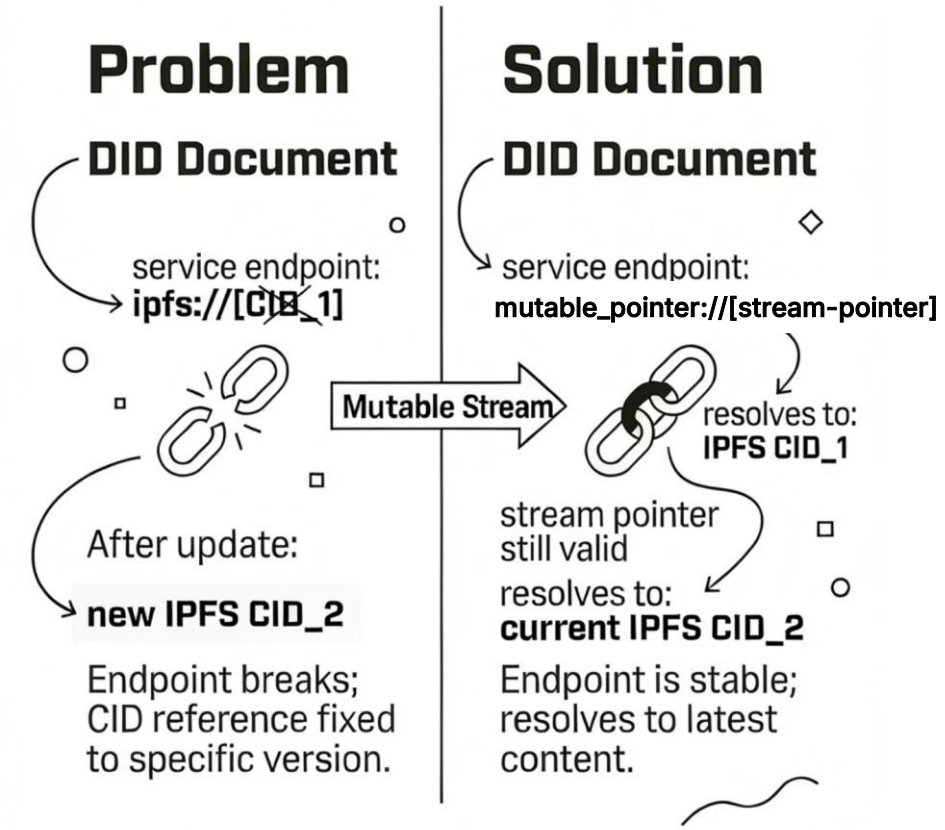

Source: Generated by Gamma (https://gamma.app/)

Access Layer: Consent-Gated Retrieval

Every data retrieval request passes through three sequential verification checks before any content is returned.

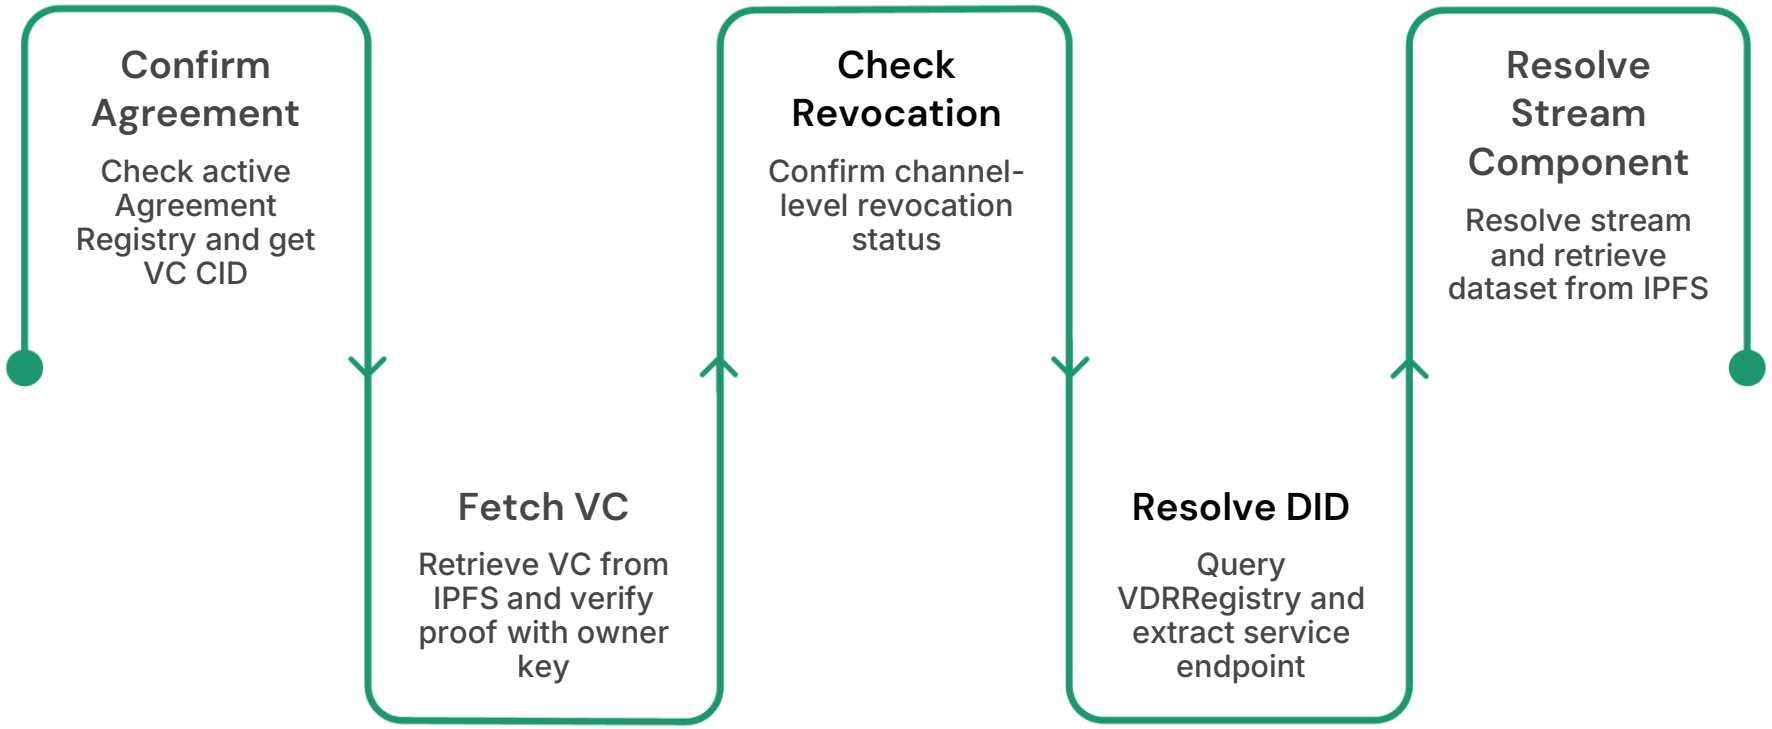

The five-lane sequence spans: **Consumer · AgreementRegistry · IPFS · VDRRegistry · Service Endpoint (through stream component)**. No data is returned unless all three verification checks pass: ensuring consent, credential validity, and revocation status are confirmed at every access event.

Discussion (1/2) — The DID Document

Finding 1 : Governance Artifact

The DID Document functions as a **living governance artifact**, not a static identity record. Service endpoints and consent-linked updates extend its role to active data exchange coordination.
Mechanism: VDRRegistry anchors each DID Document version on-chain; service endpoints reference active data channels per consent agreement.

Finding 2: DID Document stability

A prerequisite for reliable governance, is achieved via mutable pointers. Dataset updates propagate without breaking the service endpoint references embedded in the DID Document.

Mechanism: The mutable stream identifier remains constant across IPFS CID changes; the DID Document requires no update when the underlying data evolves.

Discussion (2/2) — VCs and Architectural Coherence

Verifiable Credentials are the binding mechanism that makes the architecture coherent, connecting identity, consent, and data access into a single verifiable chain.

Issuer DID owner who grants access	**Subject** Consumer DID receiving access
Channel Specific data stream scoped	**Scope** What data may be accessed
Validity Time-bounded access window	**Proof** Cryptographic signature by issuer

A DID alone cannot express what or how much access is granted. VCs carry that authorization semantics. Findings map back to requirements: **governance layer** (FR01·FR02·NFR02); **stability** (FR02·FR05·NFR01); **VC authorization** (FR03·FR04·NFR03·NFR04).

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR01: Identity Existence

Most discriminating requirement. DID-adopting works achieve only partial satisfaction; none unify blockchain transactions, DID ownership, and VC signing under a single keypair.

FR02: Owner Control

Partially addressed by nearly all surveyed works. Only [4] implements a unified control plane covering both DID management and consent operations.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR01: Identity Existence

Most discriminating requirement. DID-adopting works achieve only partial satisfaction; none unify blockchain transactions, DID ownership, and VC signing under a single keypair.

FR02: Owner Control

Partially addressed by nearly all surveyed works. Only [4] implements a unified control plane covering both DID management and consent operations.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR03: Data Access

Only [10] fully satisfied. Most works lack channel-scoped binding, they grant access to a system rather than to a specific, scoped data stream.

FR04: Consent

The least satisfied requirement across the corpus. No surveyed work anchors consent as a verifiable, tamper-evident state machine linked to DID Document updates.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR03: Data Access

Only [10] fully satisfied. Most works lack channel-scoped binding, they grant access to a system rather than to a specific, scoped data stream.

FR04: Consent

The least satisfied requirement across the corpus. No surveyed work anchors consent as a verifiable, tamper-evident state machine linked to DID Document updates.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR05: Minimalization

[5, 11, 4, 1] fully satisfied. Remaining works treat minimalization as a cost optimization rather than a privacy-by-design architectural constraint.

NFR01: Persistence

[6] is the strongest performer, with explicit key rotation support. Most works do not address the long-term resolvability of identity and dataset references across infrastructure changes.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

FR05: Minimalization

[5, 11, 4, 1] fully satisfied. Remaining works treat minimalization as a cost optimization rather than a privacy-by-design architectural constraint.

NFR01: Persistence

[6] is the strongest performer, with explicit key rotation support. Most works do not address the long-term resolvability of identity and dataset references across infrastructure changes.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

NFR02: Protection

No surveyed work fully satisfies this requirement. [5] comes closest, but does not provide a complete mechanism for enforcing user rights across all architectural layers.

NFR03: Transparency

Only [4] fully satisfied, through explicit on-chain/off-chain separation with auditable event logs. Most works rely on off-chain components that reduce auditability.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

NFR02: Protection

No surveyed work fully satisfies this requirement. [5] comes closest, but does not provide a complete mechanism for enforcing user rights across all architectural layers.

NFR03: Transparency

Only [4] fully satisfied, through explicit on-chain/off-chain separation with auditable event logs. Most works rely on off-chain components that reduce auditability.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[2]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○
[4]	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ●	● ● ●	● ● ○
[5]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ●	● ● ○	● ● ●	● ● ○	● ● ●	● ● ●
[6]	● ● ○	● ● ○	● ○ ○	● ○ ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○
[7]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[9]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[10]	● ● ○	● ● ○	● ● ●	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[11]	● ● ○	● ● ○	● ○ ○	● ● ○	● ● ●	● ● ○	● ● ○	● ● ○	● ● ●	● ● ○
[12]	● ○ ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[13]	● ● ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
[14]	● ○ ○	● ● ○	● ● ○	● ○ ○	● ● ○	● ○ ○	● ● ○	● ● ○	● ○ ○	● ○ ○
This work	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ●	● ● ○	● ● ●	● ● ●	● ● ○

NFR04: Interoperability

[5, 10, 11, 4] achieve full W3C DID and Verifiable Credential conformance. Others use proprietary or partially compliant identity schemes.

NFR05: Portability

Only [5] fully satisfied. Most works bind identity to a specific platform or infrastructure, preventing migration across W3C-compliant systems.

Related Work: Compliance Matrix

Twelve surveyed works were evaluated against all ten requirements (FR01–FR05, NFR01–NFR05). No surveyed work simultaneously satisfies all five functional and all five non-functional requirements. FR04 (consent) and NFR02 (protection) are the least satisfied across the corpus.

● ● ● = Satisfied ● ● ○ = Partially Satisfied ● ○ ○ = Not Satisfied

Paper	FR1	FR2	FR3	FR4	FR5	NFR1	NFR2	NFR3	NFR4	NFR5
[1]	●○○	●●○	●●○	●●○	●●●	●●○	●●○	●●○	●○○	●○○
[2]	●○○	●●○	●●○	●●○	●●○	●○○	●●○	●●○	●●○	●●○
[4]	●●○	●●●	●●○	●●○	●●●	●●○	●●○	●●●	●●●	●●○
[5]	●●○	●●○	●●○	●○○	●●●	●●○	●●●	●●○	●●●	●●●
[6]	●●○	●●○	●○○	●○○	●●○	●●○	●●○	●○○	●●○	●○○
[7]	●○○	●●○	●●○	●●○	●●○	●●○	●●○	●●○	●○○	●○○
[9]	●○○	●●○	●●○	●○○	●●○	●○○	●●○	●●○	●○○	●○○
[10]	●●○	●●○	●●●	●○○	●●○	●●○	●●○	●●○	●●●	●●○
[11]	●●○	●●○	●○○	●●○	●●●	●●○	●●○	●●○	●●●	●●○
[12]	●○○	●●○	●●○	●●○	●●○	●●○	●●○	●●○	●○○	●○○
[13]	●●○	●●○	●●○	●○○	●●○	●○○	●●○	●●○	●○○	●○○
[14]	●○○	●●○	●●○	●○○	●●○	●○○	●●○	●●○	●○○	●○○
This work	●●●	●●●	●●●	●●●	●●●	●●●	●●○	●●●	●●●	●●○

NFR04: Interoperability

[5, 10, 11, 4] achieve full W3C DID and Verifiable Credential conformance. Others use proprietary or partially compliant identity schemes.

NFR05: Portability

Only [5] fully satisfied. Most works bind identity to a specific platform or infrastructure, preventing migration across W3C-compliant systems.

Conclusion and Future Work

Contributions

- Positioning the DID Document as a governance layer for personal data exchange;
- All five functional requirements (FR01–FR05) satisfied by architectural design;
- Consent modeled as a verifiable, tamper-evident on-chain state machine;
- Metadata separation with cryptographic provenance and non-repudiation.

Future Work

- NFR01 persistence and NFR05 portability require empirical validation across infrastructure;
- Cross-platform interoperability testing across W3C-compliant DID methods;
- Key-rotation strategy: formal specification and implementation testing;
- Multi-party dataspace scenarios with multiple owners and consumers.

Self-Sovereign Identity Blockchain-based Architecture for Personal Data Exchange

Silvio Langer · Fabiano Hessel

silvio.langer@pucrs.br · fabiano.hessel@pucrs.br

School of Technology, Pontifical Catholic University of Rio Grande do Sul (PUCRS) · Porto Alegre, Brazil

THANK YOU!

References (1/2)

1. Alexandrescu, A., Bărbuță, D.E., Buțincu, C.N.: Secure and decentralized data sharing using blockchain and cryptographic access control. In: 2025 7th International Conference on Blockchain Computing and Applications (BCCA), pp. 437–444. IEEE (2025). https://doi.org/10.1109/BCCA66705.2025.11229511

2. Chandran, L., Lundin, L., Padayatti, G.: Transforming personal data transactions with auditable, privacy-preserving data exchange agreements: Fostering transparency and trust in digital wallet ecosystems. In: 2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA), pp. 1–9. IEEE (2023). https://doi.org/10.1109/ETFA54631.2023.10275546

3. Christopher Allen: The path to self-sovereign identity. Life With Alacrity (blog) (Apr 2016). https://www.lifewithalacrity.com/article/the-path-to-self-soverereign-identity/

4. Costagliola, A.R., Mazzocca, C., Bujari, A., Montanari, R., Bellavista, P.: Vespace: A verifiable blockchain-based data space solution to empower the data economy. Computer Communications 239, 108180 (2025). https://doi.org/10.1016/j.comcom.2025.108180

5. Dixit, A., Zarpelao, B.B., Smith-Creasey, M., Rajarajan, M.: A privacy-aware authentication and usage-controlled access protocol for IIoT decentralized data marketplace. Computers & Security 146, 104050 (2024). https://doi.org/10.1016/j.cose.2024.104050

6. Fotiou, N., Siris, V.A., Polyzos, G.C.: Enabling self-verifiable mutable content items in IPFS using decentralized identifiers. In: 2021 IFIP Networking Conference (IFIP Networking), pp. 1–6 (2021). https://doi.org/10.23919/IFIPNetworking52078.2021.9472820

7. Gazsi, J.S., Zafreen, S., Dagher, G.G., Long, M.: Vault: A scalable blockchain-based protocol for secure data access and collaboration. In: 2021 IEEE International Conference on Blockchain (Blockchain), pp. 376–381. IEEE (2021). https://doi.org/10.1109/Blockchain53845.2021.00059

8. Kernstock, P., Harms, C., Hein, A., Krcmar, H.: Establishing and governing data ecosystems at the crossroads of centralization and decentralization. Electronic Markets 35(1), 71 (2025). https://doi.org/10.1007/s12525-025-00810-x

References (2/2)

9. Li, B., Yang, J., Wang, Y., Huang, X., Ren, J., Wang, L.: A blockchain-based privacy-preserving data sharing scheme with security-enhanced access control. In: 2023 26th International Conference on Computer Supported Cooperative Work in Design (CSCWD), pp. 825–830. IEEE (2023). https://doi.org/10.1109/CSCWD57460.2023.10152751

10. Lin, I.C., Yeh, I.L., Chang, C.C., Liu, J.C., Chang, C.C.: Designing a secure and scalable data sharing mechanism using decentralized identifiers (DID). Computer Modeling in Engineering & Sciences 141(1), 809–822 (2024). https://doi.org/10.32604/cmes.2024.051612

11. Lohar, S.N., Babar, S.D., Mahalle, P.N.: A self-sovereign identity framework for context-aware decentralized identifier creation and credential verification. Engineered Science 36, 1629 (2025)

12. Pincheira, M., Donini, E., Vecchio, M., Kanhere, S.: A decentralized architecture for trusted dataset sharing using smart contracts and distributed storage. Sensors 22(23), 9118 (2022). https://doi.org/10.3390/s22239118

13. Thorve, A., Shirole, M., Jain, P., Santhumayor, C., Sarode, S.: Decentralized identity management using blockchain. In: 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), pp. 1985–1991. IEEE (2022). https://doi.org/10.1109/ICAC3N56670.2022.10074477

14. Wang, S., Zhang, Y., Zhang, Y.: A blockchain-based framework for data sharing with fine-grained access control in decentralized storage systems. IEEE Access 6, 38437–38450 (2018). https://doi.org/10.1109/ACCESS.2018.2851611

15. Zheng, Z., Xie, S., Dai, H., Chen, X., Wang, H.: An overview of blockchain technology: Architecture, consensus, and future trends. In: 2017 IEEE International Congress on Big Data (BigData Congress), pp. 557–564. IEEE (2017). https://doi.org/10.1109/BigDataCongress.2017.85