PCI 2026

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

Vijon Baraku, Iraklis Paraskakis, Simeon Veloudis, Poonam Yadav

SYSTRONLab

Engineering and Physical Sciences Research Council

Department for Science, Innovation, & Technology

REMOTE
Resilient Secure TinyEdge
REDC

CHEDDAR

Overview

1. Introduction
2. Background & Gaps
3. Framework Overview
4. Blockchain Verification Service
5. Comparison
6. Limitations & Conclusion

Introduction - The problem we're looking at

- 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.
- Individuals lack awareness of how their data is used, the kinds of data held about them, or may be unable to exercise their rights.
- Limited ability to influence how their data is used.

The problem statement

- Individuals lack granular control over their personal data. They lack **personal data sovereignty.**
- Imagine Alice, a patient at X Hospital
  - She wants to see all personal data the hospital holds about her
  - She wants to allow research use of her diagnosis, but she wants to prohibit AI training on her mental health records
  - She wants confidence these preferences are actually respected
- **Policy expression ≠ Policy enforcement. We need to close the loop**

Brief Background - Current approaches

- Personal Data Stores (Solid, Mydex, HAT, Databox):
  - The core idea is that instead of organisations holding your data, you have a personal "pod" where your data lives. Applications request access to your storage.
  - Gap: Access Control Lists determine who can access resources, but once access is granted, there's no audit trail of what happened next.

Brief Background - Current approaches

- Organisational Frameworks (IDSA)
  - Data stays where it is (with organisations), but standardised connectors enable controlled exchange. Focuses on B2B scenarios – companies sharing data with other companies.
  - Gap: IDSA's Clearing House offers logging capabilities, but this logging remains under data provider control. The party being audited controls the audit logs.
- Common gap: No framework provides independent, immutable audit trails for data subjects to verify policy compliance

Simplified System Overview

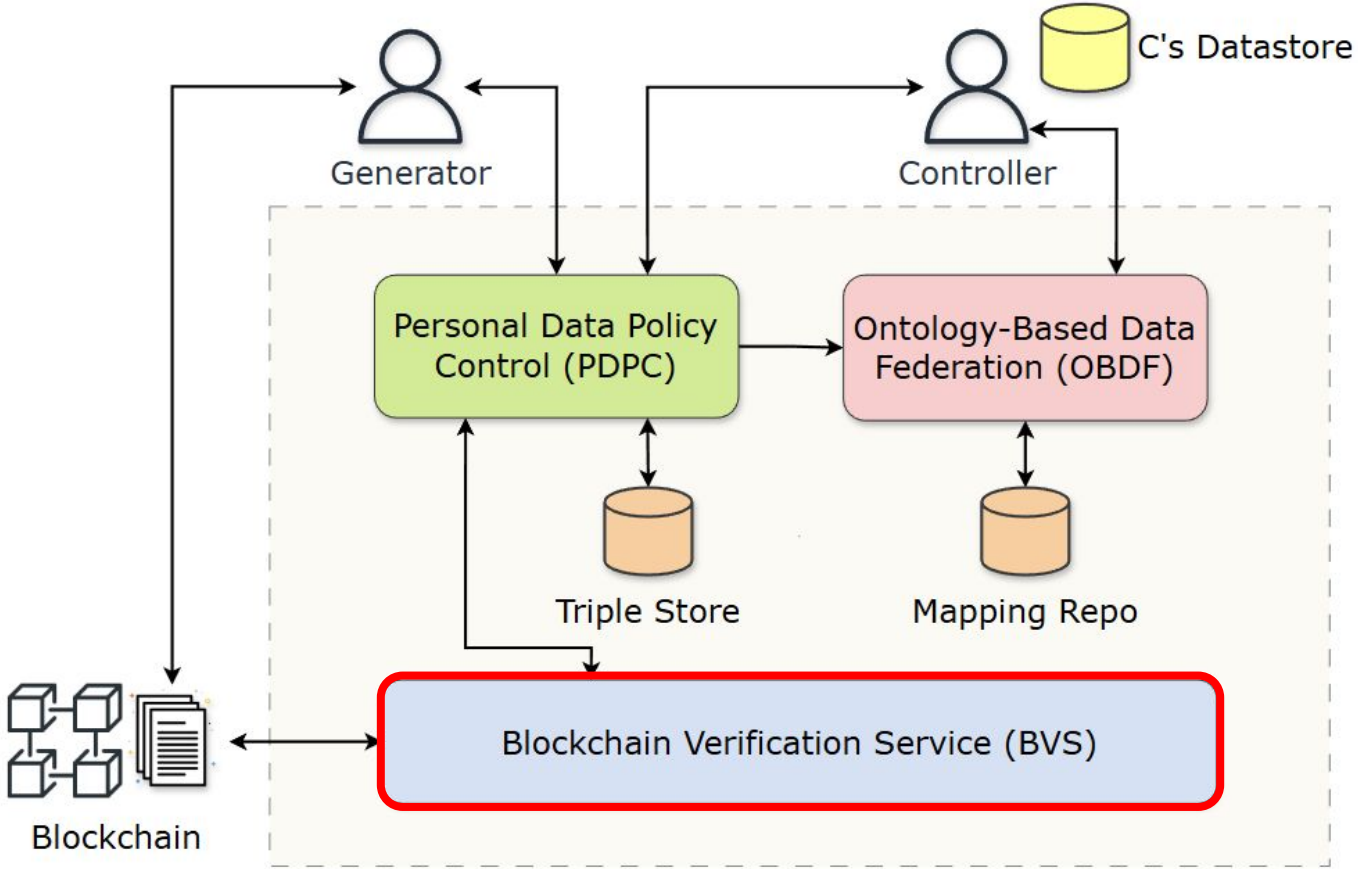

OBDF - What It Does

- Enables unified discovery of personal data across heterogeneous databases
- Virtual federation: data remains at source
- Schema.org vocabulary creates semantic bridges between different database schemas
- OBDA mappings translate SPARQL queries to database-specific SQL
- Result: Subjects see all their data without organisations moving anything

PDPC - What It Does

- Enables governance specification through ODRL policies
- W3C Open Digital Rights Language (machine-readable)
- Permissions: read, use, share, aggregate, modify
- Constraints: purpose restrictions, temporal limits, expiration dates
- AI-specific controls: allow/prohibit AI training, algorithm restrictions
- Policies stored in triple store and enforced by Policy Decision Point

BVS - Design Principles

- Completeness
  - All access requests logged, regardless of outcome
  - Denied attempts recorded alongside permitted ones
  - Prevents selective logging
- Independence
  - Audit trail exists on blockchain, outside organisational control
  - Neither subjects nor controllers can unilaterally modify records
- Linkage
  - Audit entries linked to specific policy versions
  - Access decisions reference the policy in effect at evaluation time

BVS - Smart Contract Architecture

- PolicyRegistry Contract
  - Records SHA-256 hash of policy content
  - Stores: subject address, policy ID, version number, timestamp
  - Hash storage minimises gas costs while enabling integrity verification
- AccessLogger Contract
  - Records every access decision from the PDP
  - Captures: controller address, subject address, purpose hash, action, decision outcome, policy version used
  - Emits events for real-time monitoring

BVS – Access Request Flow

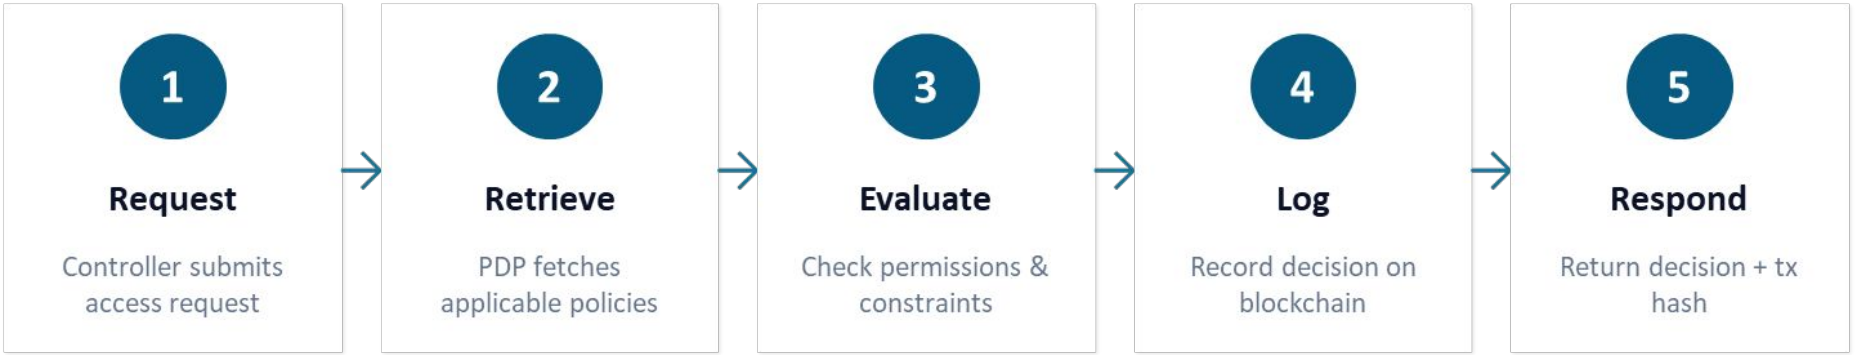

- Both PERMIT and DENY are logged
- The transaction hash serves as a unique, verifiable reference. Controllers receive proof that their interaction was recorded immutably.

BVS - Subject Verification Interface

- Access History Dashboard
  - Chronological log of all access requests
  - Shows: controller identity, action, purpose, decision, timestamp
  - Filter by controller, action type, or time period
- Independent Verification
  - Technical users: verify via any Ethereum block explorer
  - Non-technical users: built-in verification tool detects unauthorised modifications (hash mismatch)

Framework Comparison

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

What Blockchain Verification Achieves

- Accountability Infrastructure
  - Immutable records establish accountability even when enforcement is imperfect
  - Organisations cannot deny or hide their access patterns
- Detection Capability
  - Review of access logs can reveal unexpected patterns
  - Potential indicators of policy violations or framework bypass

What Blockchain Verification Achieves

- Regulatory Foundation
  - Technical infrastructure for future transparency legislation
  - Demonstrates feasibility of mandatory logging requirements
- Behavioural Incentives
  - Transparency may influence organisational behaviour
  - Awareness of immutable recording encourages compliance

What Blockchain Verification Cannot Do

- Cannot Prevent Bypass
  - Organisations control their databases
  - Determined controllers can query directly, bypassing the framework
- Trust in Deployment
  - Assumes framework is correctly deployed
  - Assumes access requests are routed through it
- Purpose Misrepresentation
  - Controller may declare purpose A while intending purpose B
  - Blockchain records declared purpose only

Conclusion

- Presented blockchain-based verification for personal data sovereignty
- Integrates immutable audit logging with semantic federation and ODRL policies
- Shifts trust model: from "policies will be respected" to "access patterns are auditable"
- Combination not present in existing frameworks

Thank you for your attention!

Questions?