ESI Preprints

Not Peer-reviewed

Self-Sovereign Identity Architecture for National Use with Wallet Proofs Zero-Knowledge and the VWR Framework

Md Abul Mansur
Nuspay International Inc., United States

Doi: 10.19044/esipreprint.12.2025.p28

Approved: 06 December 2025
Posted: 08 December 2025

Cite As:
Mansur, M.A. (2025). *Self-Sovereign Identity Architecture for National Use with Wallet Proofs Zero-Knowledge and the VWR Framework.* ESI Preprints. https://doi.org/10.19044/esipreprint.12.2025.p28

Abstract

National identity must deliver fast, fair decisions without exposing personal data. This article advances a Self-Sovereign Identity (SSI) architecture, reinforced by a Verify-Without-Reveal (VWR) framework, to achieve that goal at national scale. SSI places credentials in a citizen wallet and enables selective disclosure and zero-knowledge proofs, so services can verify attributes without seeing underlying records. VWR adds the policy and accountability spine: yes/no attribute APIs for holder-absent cases, purpose-bound and zero-trust enforcement on every call, and an immutable audit layer on a permissioned ledger. The study synthesises current standards and leading implementations in Europe and worldwide and formulates a deployable blueprint with clear roles, consent and lawful-override flows, per-agency pseudonyms, and regulator and citizen visibility. It details reference APIs, wallet and verifier user experience, and performance targets compatible with national workloads. Privacy-preserving AI strengthens biometric liveness, fraud detection, and anomaly response without centralising sensitive data. The framework aligns with GDPR data minimisation and purpose limitation, supports the European Digital Identity Wallet, and meets high-risk AI governance requirements. Results show how SSI proofs and VWR controls reduce unconsented disclosure and cross-agency browsing, while keeping latency low and interoperability high. The contribution is both conceptual and operational: a phased migration path that turns verify-without-reveal into the default mode for government and regulated services, improving security, inclusion, and public trust.

Keywords: Self-sovereign identity; SSI; verify-without-reveal;zero-knowledge proofs; zero-trust architecture; permissioned blockchain; immutable audit; national digital identity; privacy-preserving AI

Introduction

National identity systems must deliver fast, fair decisions for millions of people every day. They also carry the risk of exposing personal data across agencies and sectors. Many current platforms still default to broad lookups of full records when a simple truth value would do. This behaviour erodes trust and increases legal risk under data-protection rules. It also creates technical debt, because copied records spread and are hard to control later (European Commission, 2024). This article proposes a different path. It places Self-Sovereign Identity (SSI) at the centre and reinforces it with a Verify-Without-Reveal (VWR) policy and audit spine. SSI moves credentials to a citizen wallet and enables selective disclosure and zero-knowledge proofs. VWR ensures that, even when the holder is not present, verifiers receive only a yes/no answer bound to a declared purpose, and every access is logged in an immutable audit trail. Together they turn privacy by design into routine system behaviour (Dunphy & Petitcolas, 2020; Camenisch & Lehmann, 2021; Allen & Hess, 2022). The motivation is both legal and practical. The GDPR requires purpose limitation and data minimisation. eIDAS 2.0 promotes cross-border wallets and selective disclosure. The EU AI Act treats biometric identification as high-risk and demands governance, testing, and human oversight. These instruments encourage verification without unnecessary revelation. They also require strong accountability. Logs must be trustworthy, and people must be able to see who accessed their data and why. Programmes in Estonia, the EU wallet pilots, India's Aadhaar, and municipal SSI trials show that these aims are realistic when engineering and governance align (European Commission, 2024; Dunphy & Petitcolas, 2020; UIDAI, 2025). Technically, the core question is simple. How can a verifier decide "eligible or not" without loading a dossier? SSI answers part of the question with holder-present proofs. The wallet presents a cryptographic statement such as "over 18," "licence valid," or "resident in district X" that reveals no extra attributes. VWR answers the rest for holder-absent flows. A secure API returns only yes/no, enforces purpose via zero-trust policy, and writes a tamper-evident audit for every call. The result is consistent behaviour across channels: minimal disclosure by default, consent for richer data, and complete accountability either way (Camenisch & Lehmann, 2021; Allen & Hess, 2022).

Motivation and scope of the Research

This work is motivated by three pressures. First is scale. National workloads demand low latency and high throughput. Predicate proofs and yes/no answers are small and fast, which makes them fit. Second is trust. People accept digital services when they keep control and can see the audit trail. SSI wallets and citizen portals satisfy that need. Third is interoperability. Services must work across borders and sectors. Standards for DIDs, verifiable credentials, revocation, and OpenID for Verifiable Credentials make this possible (Dunphy & Petitcolas, 2020; European Commission, 2024). The scope is national identity as the primary identity. The framework covers government-to-citizen, government-to-business, and regulated private use. It treats SSI as the preferred mode when the holder is present. It treats VWR APIs as the safe default for background and back-office flows. It includes consent capture, lawful overrides with oversight, and per-agency pseudonyms to reduce linkability. It integrates privacy-preserving AI for liveness, anomaly detection, and fraud scoring without centralising raw sensitive data (Paredes-García et al., 2023; Kaul, 2021; Ren et al., 2025). Sectoral IDs appear only as comparators where they inform national design. We do not attempt a cost–benefit model. We focus on technical and governance design, compliance mapping, performance targets, and migration from legacy systems. We assume a heterogeneous environment with existing registries and mixed network quality. We design for low-bandwidth regions, assisted channels, and recovery for people who lose devices or cannot use biometrics, so inclusion is practical and not just stated (Dunphy & Petitcolas, 2020).

Contributions and significance (EU and global)

This article offers four contributions. First, it integrates SSI and VWR into a single national architecture. It shows how wallet-based selective disclosure and zero-knowledge proofs combine with yes/no APIs, zero-trust policy, and immutable audit. It explains how to keep behaviour consistent across holder-present and holder-absent flows. It provides concrete patterns for per-agency pseudonyms, purpose codes, consent artefacts, and regulator and citizen visibility (Camenisch & Lehmann, 2021; Allen & Hess, 2022). Second, it delivers a deployable blueprint. It defines reference APIs for verify, present, consent, and audit. It sets latency and anchoring targets. It describes policy-as-code enforcement and testing. It details wallet and verifier user experience, including delegated consent, accessibility, and recovery. It gives a phased migration plan: audit first, then yes/no by default, then SSI augmentation, then legacy retirement. It aligns these steps with real constraints of national operations (European Commission, 2024; UIDAI, 2025). Third, it maps compliance to engineering. It turns GDPR data minimisation and purpose limitation into defaults. It embeds AI Act governance for high-risk biometrics. It adopts eIDAS 2.0 wallet profiles and cross-border trust services. It shows how to make legal duties measurable through immutable logs, purpose mappings, and public transparency reports (European Commission, 2024). Fourth, it grounds the design in evidence. It draws lessons from Estonia's transparency model, the EU wallet pilots, Aadhaar's at-scale yes/no authentication, and municipal SSI deployments. It explains which choices improve adoption convenience plus control and which choices fail over-collection, weak logging, and implicit internal trust (Dunphy & Petitcolas, 2020; UIDAI, 2025). The significance is direct. For the EU, the design operationalises wallet-centric identity with enforceable purpose limits and visible accountability. For other regions, it offers a path to modernise without expanding surveillance risk. It raises assurance with liveness and anomaly detection while keeping sensitive data distributed. It improves inclusion by providing assisted flows and simple recovery. It strengthens legitimacy by giving people a clear view of who accessed what and why. In short, it aims to make verify-without-reveal the normal way identity is used in government and regulated markets, not an exception reserved for pilots or niche services (European Commission, 2024; Allen & Hess, 2022; Dunphy & Petitcolas, 2020).

Background and Related Work

Self-Sovereign Identity (SSI) arose from the need to restore user control in digital identity while keeping high assurance and interoperability. In SSI, trusted authorities issue verifiable credentials to the holder; the holder stores them in a wallet and presents proofs to verifiers when needed. Decentralized Identifiers (DIDs) provide resolvable identifiers and public keys without a single, central directory. Together, DIDs and VCs let a verifier check the authenticity and freshness of claims without contacting the issuer each time, which reduces linkability and improves resilience (Dunphy & Petitcolas, 2020; Sporny, Longley, & Chadwick, 2022; Sporny et al., 2022). Wallets add policy and UX: the holder can select which attributes to disclose, set consent preferences, and manage recovery. Modern wallets support mobile secure elements or trusted execution, remote revocation checks, and presentation of cryptographic proofs that are compact enough for web and in-person flows (Preukschat & Reed, 2021; Meylan & Sabadello, 2021). National identity practice shows both progress and gaps. Estonia's model demonstrates how separated registries, strong authentication, and full-stack logging enable safe data exchange across government. Every lookup is policy-checked and time-stamped, and citizens can later see who accessed what and when. This transparency improves trust and reduces silent misuse (Dunphy & Petitcolas, 2020). Across the European Union, eIDAS 2.0

introduces the European Digital Identity Wallet, which standardises selective disclosure and cross-border verification so that residents can prove attributes abroad without sharing full records (European Commission, 2024). Outside Europe, India's Aadhaar separates yes/no authentication from consented e-KYC to limit data spread, proving that minimal answers can work at population scale when audit and consent are enforced (UIDAI, 2025). These programmes reveal the core gap VWR addresses: holder-present SSI proofs are strong, but governments also need safe holder-absent verification for back-office processes. Without a zero-trust policy layer and immutable audit, background checks tend to revert to dossier pulls and broad internal access.

The cryptographic building blocks for minimal disclosure are mature. Selective disclosure allows a holder to reveal exactly one or a few attributes from a credential without exposing the rest. Zero-knowledge proofs (ZKPs) go further by proving predicates over attributes "over 18," "licence valid," "resident of district X" without revealing the values or the identifier. Efficient constructions, such as BBS+ signatures for unlinkable selective disclosure, and accumulator-based revocation, make verification fast enough for web and mobile at scale (Camenisch & Lehmann, 2021; Khovratovich & Law, 2020). Revocation lists and status endpoints prevent use of stale credentials, while caching keeps latency low. In parallel, domain standards such as ISO/IEC 18013-5 for mobile driving licences show how signed attributes can replace photocopies and still pass inspection, which aligns with SSI and VWR aims (ISO/IEC, 2021).

Policy enforcement in large public systems is moving from implicit trust to zero-trust. Traditional role-based access on a "trusted network" allows broad internal browsing and weak oversight. Attribute-based access control (ABAC) evaluates purpose, role, legal basis, consent state, and risk on each API call, so every access is a decision linked to declared intent. Combined with least-privilege defaults and rate limits, ABAC reduces cross-agency "surfing" and turns policy into code that auditors can test (Allen & Hess, 2022; Zhang & Li, 2020). Accountability depends on logs that cannot be silently changed. Simple database logs help, but they are editable by insiders. Permissioned blockchain or hash-chained audit systems provide append-only, time-stamped records with cryptographic integrity and multi-party control, so tampering becomes evident. Governments can record events, not data who asked what, for which purpose, and the policy outcome while keeping personal data off-chain (Juels & Oprea, 2020; Vukolić, 2021; Gencer & Basu, 2021). Estonia's integrity anchoring shows this model is practical at national scale (Dunphy & Petitcolas, 2020). The legal context in Europe favours verification without revelation. GDPR requires data minimisation and purpose limitation; systems should ask for only what is necessary and prove necessity in records. eIDAS 2.0 defines cross-border wallets, trusted issuers, and selective disclosure profiles that make predicate proofs portable across the Union. The EU AI Act treats biometric identification as high-risk and requires risk management, testing, documentation, and human oversight. Together these instruments push identity programmes toward minimal outputs, strong consent capture, and auditable operations. They also require visible accountability through logs that citizens and regulators can scrutinise (European Commission, 2024; NIST, 2020). Background checks that pull dossiers by default are hard to justify under these rules. Holder-present SSI proofs and holder-absent yes/no APIs meet the legal tests more directly, because they reduce exposure by design and connect each access to a declared purpose and lawful basis.

Finally, related strands in privacy-preserving computation support the same direction. Federated learning lets agencies train anomaly detectors and liveness improvements without centralising raw logs. Encrypted inference in trusted execution or with homomorphic methods enables selected risk scoring without exposing inputs. When used with documented bias testing and human appeal routes, these tools increase assurance while preserving privacy (Paredes-García et al., 2023; Kaul, 2021; Ren et al., 2025). The literature converges on a simple principle: most services need a decision, not a dossier. SSI provides cryptographic, holder-centred proofs; VWR adds policy, purpose, and immutable audit to make the same minimal-disclosure behaviour the default even when the holder is not present. This article builds on that consensus and shows how to deploy it at national scale within European legal and operational constraints.

Research Questions and Methodology

This study investigates how a self-sovereign identity model, reinforced by a verify-without-reveal framework, can meet national requirements for speed, legality, and trust. The central research question asks how a state can decide eligibility without exposing underlying records in both holder-present and holder-absent scenarios. From this, three subsidiary questions follow. The first asks how DIDs, verifiable credentials, selective disclosure, and zero-knowledge proofs can be composed so that most interactions return a predicate or a yes/no answer rather than a dossier (Camenisch & Lehmann, 2021; Sporny et al., 2022). The second asks what policy and control plane is required to stop cross-agency browsing while keeping latency low; in particular, whether attribute-based access control and zero-trust enforcement at the API layer can encode purpose limitation and least privilege in routine operation (Allen & Hess, 2022; Zhang & Li, 2020). The third asks how assurance can rise without new pools of sensitive data, focusing on liveness detection, anomaly detection, and fraud scoring delivered through federated learning and encrypted inference where feasible

(Paredes-García et al., 2023; Kaul, 2021; Ren et al., 2025). A fourth question considers compliance and adoption: whether an SSI-VWR design aligns with GDPR data minimisation and purpose limitation, with eIDAS 2.0 wallet profiles, and with EU AI Act obligations for high-risk biometric systems, and whether it is compatible with cross-border services and legacy registries in practice (European Commission, 2024; NIST, 2020).

The research design blends conceptual synthesis with comparative case analysis. The synthesis assembles cryptographic and architectural elements DIDs and verifiable credentials for issuer–holder–verifier trust, selective disclosure and zero-knowledge proofs for minimal revelation, and permissioned ledger audit for tamper-evident accountability into a coherent architecture that can operate at national scale. The comparative analysis tests that architecture against programmes that already operate or pilot adjacent ideas. Estonia contributes evidence on transparency, separated registries, and integrity anchoring for auditability. EU wallet pilots contribute evidence on selective disclosure and cross-border portability. India's Aadhaar contributes evidence on yes/no authentication at population scale with consented e-KYC for richer cases. Municipal SSI pilots such as Zug contribute evidence on holder-present proofs with user-held credentials. Together these cases provide a realistic yardstick for security, privacy, interoperability, governance, and performance (Dunphy & Petitcolas, 2020; European Commission, 2024; UIDAI, 2025).

Sources include peer-reviewed articles on SSI, DIDs/VCs, selective disclosure, zero-knowledge proofs, zero-trust access control, permissioned audit ledgers, and privacy-preserving machine learning, with preference for work published from 2020 onwards to reflect the maturing standards and deployments. Standards and official materials include eIDAS 2.0 communications, European Digital Identity Wallet profiles, W3C specifications for verifiable credentials and decentralized identifiers, NIST identity guidance, and national documentation for at-scale systems. Inclusion criteria require technical specificity, deployment or pilot evidence, and clear links to national identity use. Exclusion criteria remove marketing papers without protocol detail and opinion essays without sources (Sporny, Longley, & Chadwick, 2022; European Commission, 2024; NIST, 2020). Data extraction focuses on seven elements that recur in every credible design. These are the binding between person and identifier, the credential model and revocation, the verification path for holder-present and holder-absent flows, the policy and logging layer, the security services for keys and biometrics, the interoperability profile, and the operational and legal controls. For each source, the study records the architectural choices, performance claims, failure modes, and governance arrangements. It also notes how consent is captured, how overrides are justified and limited, and how citizens can see access history. Coding groups these observations under minimal disclosure, user control, zero-trust enforcement, immutable audit, privacy-preserving AI, and legal alignment. The same codes are then applied to the proposed SSI-VWR design to test completeness and to expose trade-offs (Dunphy & Petitcolas, 2020; Camenisch & Lehmann, 2021; Allen & Hess, 2022).

Evaluation proceeds against five criteria that reflect national realities. Security measures whether the design resists spoofing, replay, and insider abuse, and whether keys and biometrics are handled with liveness, rotation, and recovery. Privacy measures whether selective disclosure and zero-knowledge proofs replace dossier pulls, whether per-agency pseudonyms reduce linkability, and whether consent and lawful basis are tied to each access. Interoperability measures whether proofs and yes/no checks travel across borders and sectors using recognised profiles and revocation methods. Governance measures whether roles are clear and audit is visible to both regulators and citizens, and whether sanctions and transparency reports deter misuse. Scalability and performance measure whether yes/no endpoints meet tight latency budgets, whether audit writes and anchoring remain reliable under load, and whether wallet proofs verify fast enough for web and counter use (European Commission, 2024; UIDAI, 2025). Validity relies on triangulation between academic work, standards, and programme documentation. Where claims conflict, official statistics and peer-reviewed measurements take priority. When metrics are absent, the study reports the gap and uses conservative assumptions drawn from analogous systems. Reliability is supported by a transparent coding framework and by consistency checks across sources. Bias is considered in two places. First, in the literature that reports biometric accuracy and risk scores; results are weighed by dataset diversity and by the presence of human oversight. Second, in governance materials that may under-report misuse; auditability and citizen visibility are therefore treated as essential rather than optional features (European Commission, 2024; Paredes-García et al., 2023). Ethical considerations are integral rather than ancillary. The study handles no personal data. It maps recommendations to GDPR principles and to AI Act duties for high-risk systems. It assumes that people must have meaningful control over disclosure and a clear view of access history. It also assumes that programme success depends on inclusion. This is why the design favours holder-present proofs that run on commodity devices, assisted channels for people who need help, and recovery that is simple but safe. It also proposes low-tech alternatives for those without smartphones, such as smart cards or printed codes paired with in-person checks, because equal access is part of legitimacy as well as policy (Dunphy & Petitcolas, 2020; European Commission, 2024).

Scope is limited to national identity as the primary identity, with government as the main issuer of authoritative credentials. Sectoral identities are treated as comparators where they inform national rollout. The analysis does not attempt a cost–benefit model, because costs vary by legacy estate and institutional capacity. Instead, it specifies performance targets, governance processes, and migration steps that any national programme can adapt. The proposed path begins with audit everywhere, then moves to yes/no by default, then adds wallet-based selective disclosure, and finally retires unsafe bulk interfaces. Each phase has measurable gates, such as minimal-disclosure ratios, latency targets, audit completeness, and bias thresholds (Allen & Hess, 2022; European Commission, 2024).

This methodology positions SSI as the default for holder-present interactions and VWR as the enforcement and accountability spine for all interactions. It treats policy as code and audit as a first-class deliverable. It evaluates success by how little data crosses boundaries, by how clear the purpose and evidence of each access are, and by how easy it is for a person to see and control what happened. With these foundations, the next section analyses the specific problems that a national identity programme must solve unconsented disclosure, cross-agency surfing, linkability, biometric risks, and model governance and shows how an SSI-VWR design addresses each in practice (European Commission, 2024; Camenisch & Lehmann, 2021; Dunphy & Petitcolas, 2020; UIDAI, 2025).

Problem Analysis

National identity systems face a cluster of problems that reinforce one another. Unnecessary disclosure occurs when front-line systems pull whole records for simple checks. Cross-agency "surfing" occurs when insiders browse identity data beyond a lawful purpose. Linkability grows when a single identifier appears in many domains or when metadata ties events together. Biometric risk rises when liveness is weak or when models embed bias. These issues are technical and institutional at once. They arise from defaults that favour dossiers over decisions, from policy that trusts networks instead of purposes, and from logs that can be edited or that citizens never see. An SSI-centred design, reinforced by a verify-without-reveal policy spine, must therefore change both the cryptography and the control plane so that minimal disclosure, purpose enforcement, and immutable accountability become routine (Dunphy & Petitcolas, 2020; Allen & Hess, 2022; European Commission, 2024).

Unconsented disclosure and cross-agency "surfing"

In many legacy platforms the fastest way to resolve a case is to fetch a person's full record. Teams accept this because performance budgets are tight and because interfaces were built for batch exports, not for precise questions. The result is routine over-collection. Welfare clerks view tax fields, police operators see health hints, and private verifiers receive more than they need to decide eligibility. Copies proliferate across caches, inboxes, and data lakes. Each copy becomes a new attack surface and a legal risk. Worse, people cannot see where their data went, so trust falls even when no breach occurs (European Commission, 2024). Cross-agency "surfing" follows the same pattern. Internal roles are broad and network gates are soft. Staff who can connect to the internal network can often browse records with little friction. Logs exist, but they are mutable or hard to query, so deterrence is weak. When a scandal occurs, programmes respond with training and memos. The technical conditions remain unchanged. The GDPR requirement of purpose limitation is thus honoured in policy but not in code (Zhang & Li, 2020; Campbell & Weitzner, 2022).

SSI on its own does not fix this background behaviour. Wallet proofs help when the person is present. They do not stop a back-office process from pulling a dossier when no one is at the counter. What changes the behaviour is a verification posture that makes minimal answers the norm. A purpose-bound yes/no API turns most checks into narrow decisions. An ABAC engine ties each call to a declared purpose, role, and legal basis. An immutable audit writes who asked, for what, and what the policy decided. These three controls reduce the surface for unconsented disclosure and make surfing costly, because every look leaves a tamper-evident trail that regulators and citizens can see (Allen & Hess, 2022; Juels & Oprea, 2020; Vukolić, 2021). A concrete illustration shows the point. A transport inspector needs to know whether a licence is valid today. The legacy instinct is to fetch the licence file, which includes address, birth date, and photo. The minimal path is a predicate: licenceValid = true/false. The system returns the truth value and a short-lived token. No extra fields leave the registry. The inspector can finish the task. The audit notes the purpose code and method used. If the inspector tries to view the dossier, the policy engine denies the call unless a lawful basis allows it. This is not only more private; it is also faster to build and easier to govern because the call surface is small and testable (European Commission, 2024; UIDAI, 2025).

Linkability, metadata leakage, and dossier creep

Even when attributes are hidden, people can be tracked if the same identifier appears across domains or if metadata is rich and stable. Linkability allows adversaries or curious insiders to reconstruct a person's activity across services. It undermines selective disclosure because the pattern of use becomes a signal. It also turns audit logs into a privacy risk if subject identifiers are reused across agencies (Troncoso et al., 2020;

Camenisch & Lehmann, 2021). Three mechanisms drive linkability. The first is a global identifier reused everywhere. The second is metadata that act like fingerprints: timestamps with high precision, device hints, location, and network details. The third is dossier creep. Once a relying party receives a dossier, it tends to keep it "for convenience," even if policy says otherwise. Dossiers then travel between partners and contractors. Each transfer creates new linkages and weakens control (Wagner & Eckhoff, 2020). An SSI-VWR design reduces these forces. SSI replaces many lookups with holder-present proofs. Proofs carry no global identifier and disclose only a predicate or a minimal attribute. VWR replaces global IDs with per-agency pseudonyms derived from the national identifier and an agency salt. The same person appears as different tokens to different agencies, so linking across domains by API traffic becomes hard. The verification layer strips non-essential metadata and rotates short-lived tokens. The audit layer writes events with pseudonyms and event hashes rather than raw attributes. When richer data must flow, the system sends signed credentials instead of raw rows and sets short retention and clear revocation paths. Over time this suppresses dossier creep because verifiers cannot justify keeping more than what they received for a narrow, logged purpose (Camenisch & Lehmann, 2021; Gencer & Basu, 2021; European Commission, 2024). Designing for low linkability has a developmental cost. Pseudonyms must be deterministic per agency to support reconciliation but resistant to cross-domain joins. Revocation and status checks must work without revealing global keys. Developers must learn to ask for predicates instead of datasets. These costs are one-time. The benefit is permanent: simpler compliance checks, smaller breach impact, and clearer accountability (Sporny et al., 2022; Dunphy & Petitcolas, 2020).

Biometric risks, bias, and model governance

Biometrics improve convenience and help bind people to credentials. They also introduce risks. Replay and deepfake attacks can defeat naive face or voice checks. Poor liveness detection lets printed photos or screen replays pass. Centralised biometric stores invite high-impact breaches. Models trained on narrow datasets misclassify more often for some groups, creating inequity and legal exposure (Paredes-García et al., 2023; European Commission, 2024).A robust approach separates where and how biometrics are used. During enrolment, capture is in controlled conditions and evidence is sealed with strong chain-of-custody. During routine use, liveness checks run on device where feasible, with templates stored in secure elements or trusted execution. Server-side matching, if unavoidable, runs in enclaves, uses minimal retention, and is tied to a specific purpose code. Every capture is logged with a consent or lawful-basis artefact so that oversight is possible. This reduces the blast radius of compromise and clarifies accountability when mistakes occur (NIST, 2020; Allen & Hess, 2022).Bias and drift require continuous governance. Before deployment, teams test accuracy across demographics and lighting or device conditions. They publish error-rate gaps and mitigation plans. After deployment, they monitor false accepts and false rejects, segmented by context, and they provide human appeal routes and fallbacks for people who cannot pass liveness or do not have compatible devices. This turns compliance with the AI Act into daily practice rather than a one-off certification (European Commission, 2024; Paredes-García et al., 2023).Assurance can rise without centralising raw data. Federated learning trains anomaly detection and liveness enhancements across agencies without pooling logs. Encrypted inference or trusted execution evaluates sensitive features without exposing inputs. These techniques reduce privacy risk while improving detection of misuse, such as burst queries from a single terminal, location-device mismatches, or improbable sequences that indicate credential theft. They must be paired with update validation and rollback to prevent model poisoning, and with transparent documentation so regulators can audit how models affect decisions (Kaul, 2021; Ren et al., 2025; Zhou et al., 2021).The final risk is cultural. When biometrics work well, organisations can become comfortable with broad deployment and forget that purpose still governs every capture. A national programme must keep purpose controls in the loop: no capture without a mapped purpose, no reuse beyond scope, and visible logs that citizens can inspect. Biometrics then act as one factor among several, not as a universal key that opens every door (European Commission, 2024; NIST, 2020).

Proposed SSI–VWR Framework

This section specifies a deployable framework that puts Self-Sovereign Identity (SSI) at the centre and uses Verify-Without-Reveal (VWR) as the control spine. SSI provides holder-controlled, cryptographic proofs. VWR guarantees that verifiers receive only what is necessary, that every request is tied to a declared purpose, and that each access is immutably auditable. The result is high assurance with minimal disclosure at national scale (Dunphy & Petitcolas, 2020; Camenisch & Lehmann, 2021; Allen & Hess, 2022; European Commission, 2024).

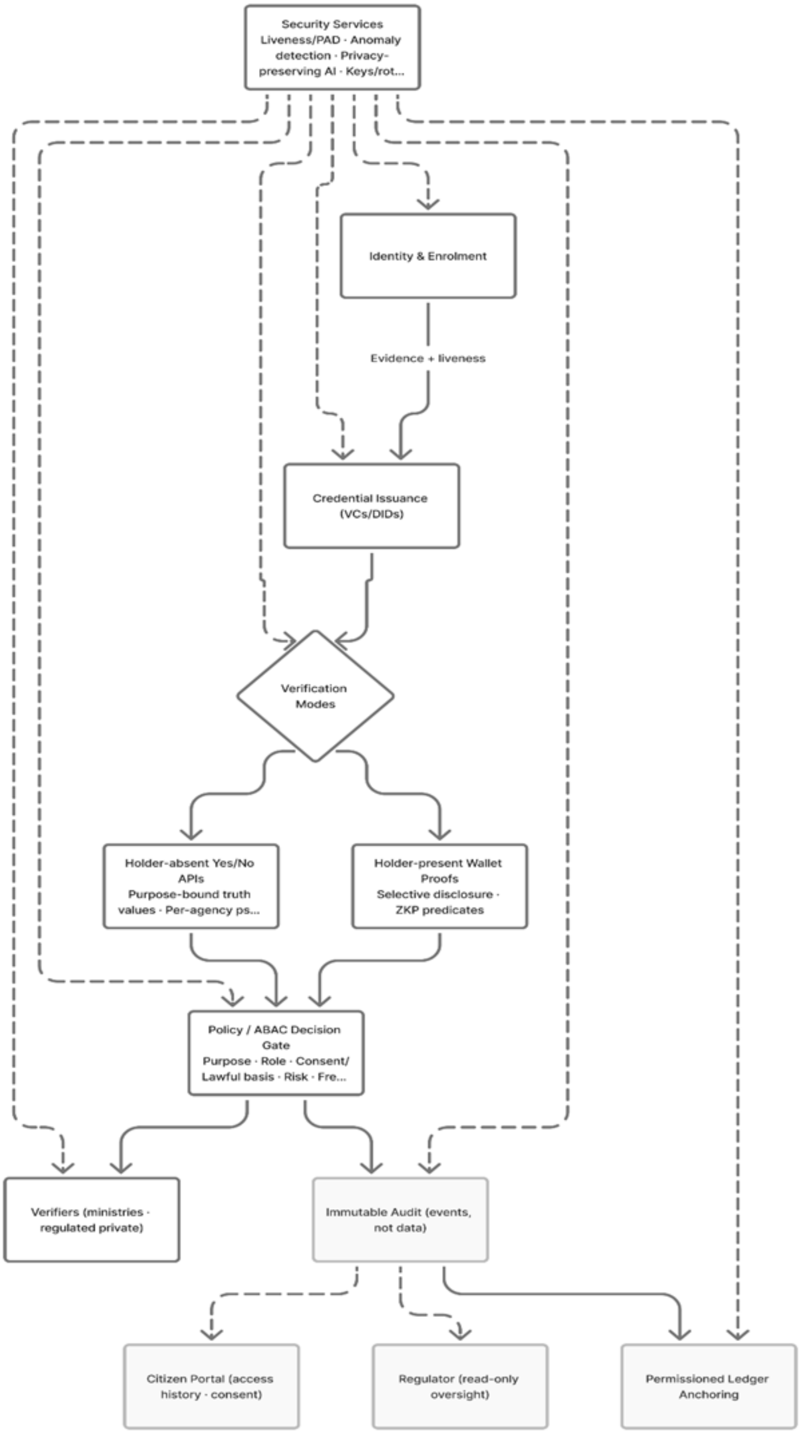
Security Services
Liveness/PAD · Anomaly detection · Privacy-preserving AI · Keys/rot...
Identity & Enrolment
Evidence + liveness
Credential Issuance (VCs/DIDs)
Verification Modes
Holder-absent Yes/No APIs
Purpose-bound truth values · Per-agency ps...
Holder-present Wallet Proofs
Selective disclosure · ZKP predicates
Policy / ABAC Decision Gate
Purpose · Role · Consent/ Lawful basis · Risk · Fre...
Verifiers (ministries · regulated private)
Immutable Audit (events, not data)
Citizen Portal (access history · consent)
Regulator (read-only oversight)
Permissioned Ledger Anchoring

Figure 1: SSI–VWR Architecture Overview

A left-to-right layered view of the SSI–VWR stack showing issuance once, two verification modes (holder-present wallet predicates and holder-absent yes/no APIs), a single Policy/ABAC decision gate, immutable audit events anchored to a permissioned ledger, and security services overlaying all layers; with external read-only views for the regulator and citizen portal

Design principles

i. Control and clarity: The person decides when to disclose; prompts explain who is asking, what will be disclosed, why, and for how long. The person can later see every access in a portal and revoke standing consents.

ii. Inclusion and recovery: Recovery is safe and simple (guardian/assisted recovery, in-person options). Delegated authority supports carers and parents. Alternative channels (smartcards, assisted counters) ensure people without smartphones are not excluded (Dunphy & Petitcolas, 2020).

iii. Consistency across borders: Wallet proofs interoperate across the EU through eIDAS 2.0 profiles, so users carry one experience at home and abroad (European Commission, 2024).

iv. Minimal outputs. The routine result of a check is a yes/no or a predicate proof (e.g., "over 18" without date of birth). Rich data move only when strictly necessary.

v. Low linkability: Per-agency pseudonyms, selective disclosure, short-lived tokens, and metadata hygiene reduce cross-domain correlation (Camenisch & Lehmann, 2021; Troncoso et al., 2020).

vi. Measurable targets: Programmes track a minimal-disclosure ratio (share of transactions resolved with minimal outputs). A realistic target at maturity is ≥ 85% (UIDAI, 2025).

vii. Purpose limitation in code: Every request carries a purpose from a public catalogue. A policy engine maps that purpose to the smallest allowable attributes. Requests outside the map are blocked or require step-up approval with audit evidence (Campbell & Weitzner, 2022).

viii. Data minimisation: Default minimal responses operationalise GDPR's minimisation duty. Rich disclosures require consent or a documented legal basis, both tied to specific events in the audit (European Commission, 2024).

ix. High-risk AI governance: Biometric uses follow EU AI Act obligations: documented risk files, testing, bias monitoring, human oversight, incident reporting (European Commission, 2024).

General overview

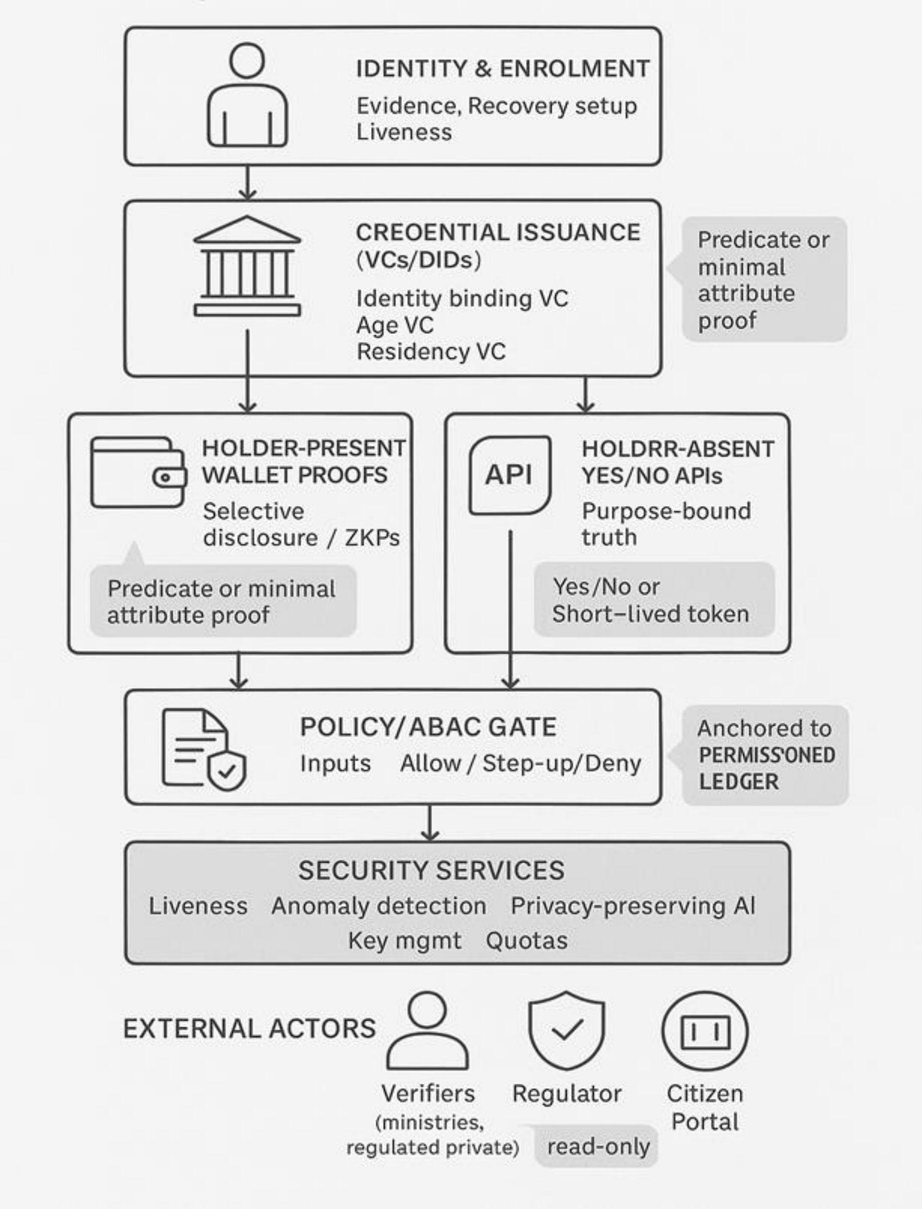
IDENTITY & ENROLMENT
Evidence, Recovery setup
Liveness
CREOENTIAL ISSUANCE (VCs/DIDs)
Identity binding VC
Age VC
Residency VC
Predicate or minimal attribute proof
HOLDER-PRESENT WALLET PROOFS
Selective disclosure / ZKPs
Predicate or minimal attribute proof
API
HOLDRR-ABSENT YES/NO APIs
Purpose-bound truth
Yes/No or Short–lived token
POLICY/ABAC GATE
Inputs Allow / Step-up/Deny
Anchored to PERMISS'ONED LEDGER
SECURITY SERVICES
Liveness Anomaly detection Privacy-preserving AI
Key mgmt Quotas
EXTERNAL ACTORS
Verifiers (ministries, regulated private)
Regulator
read-only
Citizen Portal

a) Identity & enrolment: A person is enrolled once with strong evidence and liveness. They receive a non-meaningful national identifier, recovery options (guardian/assisted), and if they choose a regulated wallet. This prevents meaningful sequence IDs and sets up safe recovery/delegation.

b) Credential issuance (VCs/DIDs): Authoritative bodies issue verifiable credentials to the holder’s wallet: Identity Binding, Age, Residency, Licence Status. Issuers publish keys via DIDs/trust lists and attach status/revocation references so freshness can be checked without central lookups every time.

c) Verification mode selection: Each service chooses the minimal path based on context:

- Holder-present (SSI) when the person is on site or online with their wallet.
- Holder-absent (VWR) for back-office or cross-agency checks.

d) Holder-present flow (wallet predicates / selective disclosure): The verifier sends a purpose-bound request (e.g., “age ≥ 18”, “licence valid today”, “resident of M”). The wallet authenticates the requester and displays a plain-language prompt (who/what/why/how long). On approval, it produces either a predicate proof (ZKP) or the single attribute requested not a dossier. Proofs are audience-bound and time-limited; the verifier checks issuer trust, revocation freshness, audience and time, then acts on the minimal decision.

e) Holder-absent flow (yes/no APIs): A system calls a purpose-declared yes/no endpoint: “is residency current?”, “is licence valid today?”. The call is evaluated by a zero-trust ABAC gate (purpose, role, consent or lawful basis, contextual risk, freshness policy). The registry returns only true/false or a short-lived token. No raw attributes leave the source. Subjects are represented with per-agency pseudonyms to suppress cross-domain linkability; quotas/rate limits deter scraping.

f) Policy/ABAC decision gate (single control point): Both modes pass a common gate that outputs Allow / Step-up / Deny. Step-up can require stronger auth, supervisory co-sign, or an explicit consent artefact if richer data are requested. Requests outside the purpose catalogue are blocked or rewritten to minimal predicates.

g) Consent and lawful overrides: Rich attributes flow only with specific, time-limited consent captured as a signed artefact and bound to the event. Rare lawful overrides (e.g., court order) are narrow in scope, time-boxed, require two-person approval, and trigger regulator alerts; where lawful and safe, the citizen is notified afterwards.

h) Immutable audit and ledger anchoring: Every access produces an event (who, purpose, method, outcome, consent/override reference, freshness) in an append-only log. Batches are anchored to a permissioned ledger for tamper-evidence. Citizens see their own access history in a portal; regulators have full oversight dashboards. The audit stores events, not personal data.

i) Security services overlay: Cross-cutting controls harden the fabric: biometric liveness/PAD for high-risk actions, anomaly detection on access patterns, privacy-preserving AI (federated/enclave) for fraud without centralising raw logs, strong key/crypto management, mutual auth, quotas, rate-limits, and rapid key rotation.

j) Interoperability and performance: Wallet proofs follow W3C VC/DID and EUDI profiles; status uses cacheable lists with risk-based freshness (real-time for safety-critical, day-level for static claims). Minimal payloads keep p95 latencies in the hundreds of milliseconds; clear degradation policies (deny/retry/queue by risk) sustain national-scale operations. Programmes track a minimal-disclosure ratio KPI to prove privacy-by-default is working.

Architecture overview: SSI holder-present proofs + VWR policy/audit spine

a) Conceptual layers
   - Identity & binding: A non-meaningful national identifier is bound to the person at enrolment with strong authentication and liveness checks. Recovery and re-binding procedures are defined for loss or compromise (NIST, 2020).
   - Credential issuance: Authoritative bodies issue Verifiable Credentials (VCs) for core attributes (age, residency, licence status). Decentralised Identifiers (DIDs) and recognised trust lists make issuers discoverable and revocation verifiable (Sporny, Longley, & Chadwick, 2022).

b) Verification modes.
   - Holder-present (SSI-native): wallet-based selective disclosure and zero-knowledge proofs; no routine registry lookups.
   - Holder-absent (VWR): narrow, purpose-bound yes/no checks and time-boxed tokens; dossiers are not returned (Camenisch & Lehmann, 2021; UIDAI, 2025)

c) Policy & trust: A zero-trust control plane evaluates each request: authentication, authorisation, purpose mapping, consent/legal basis, and contextual risk. Least privilege, quotas, and step-up are normal (Allen & Hess, 2022).

d) Integrity & audit: Every access becomes an event with who/when/why and the policy outcome. Events are chained and anchored to a permissioned ledger operated by multiple state entities and the regulator. Personal data never go on-chain (Juels & Oprea, 2020; Vukolić, 2021).

e) Security services: End-to-end encryption, key management, liveness and behavioural assurance, anomaly detection, and privacy-preserving AI provide defence-in-depth without centralising sensitive data (Paredes-García et al., 2023; Kaul, 2021; Ren et al., 2025).

Lifecycle view

a) Enrolment: strong evidence capture, liveness, issuance of initial credentials, and setup of recovery paths.
b) Routine use: predicates and yes/no checks by default; consent prompts for richer cases; step-up for risk.
c) Oversight: immutable audit, regulator analytics, transparency reports; citizen visibility through a portal.

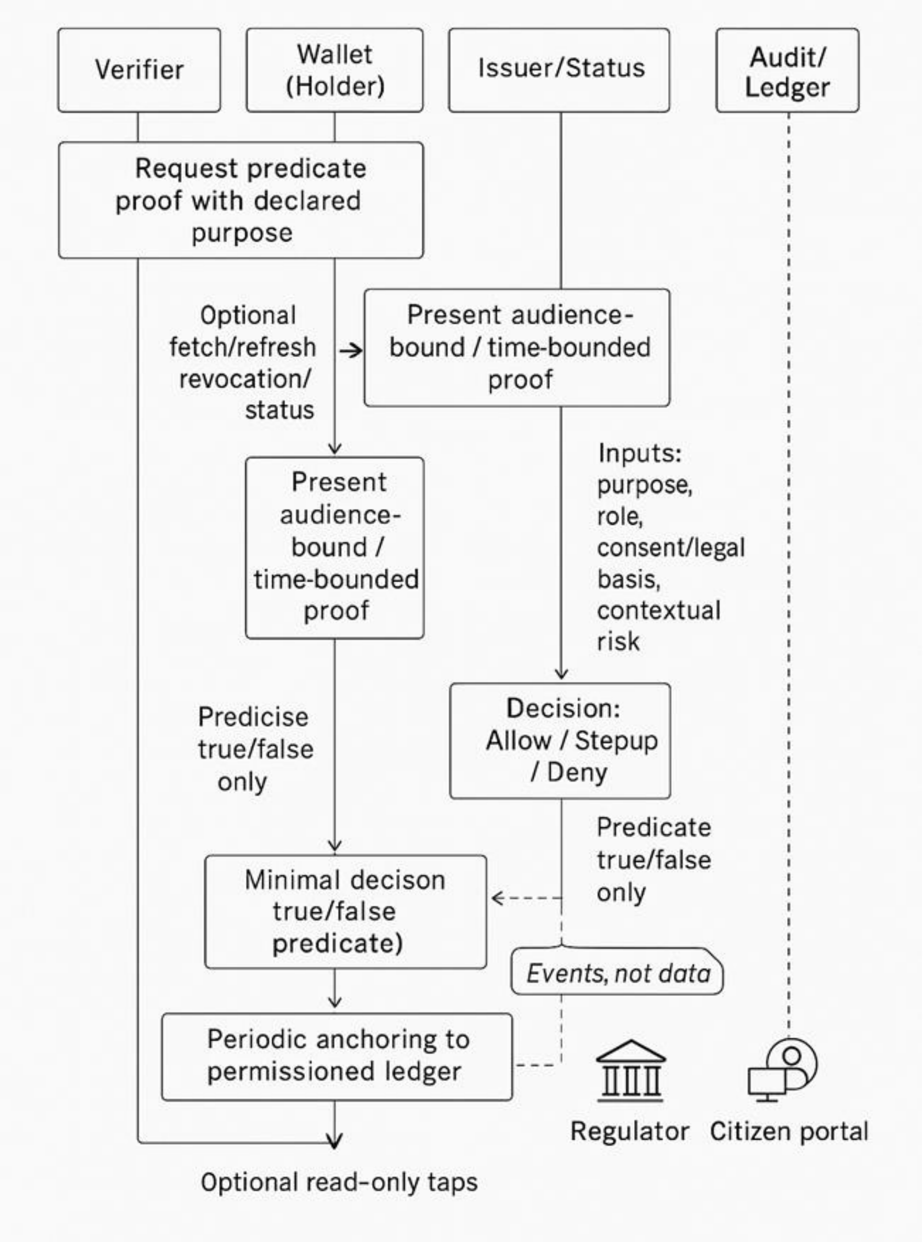
Verifier
Wallet (Holder)
Issuer/Status
Audit/ Ledger
Request predicate proof with declared purpose
Optional fetch/refresh revocation/ status
Present audience-bound / time-bounded proof
Present audience-bound / time-bounded proof
Inputs: purpose, role, consent/legal basis, contextual risk
Predicise true/false only
Decision: Allow / Stepup / Deny
Predicate true/false only
Minimal decison true/false predicate)
Events, not data
Periodic anchoring to permissioned ledger
Regulator
Citizen portal
Optional read-only taps

Figure A1: Holder-Present Selective Disclosure Flow (portrait)

The figure is a vertical swimlane diagram with five lanes Verifier, Wallet (Holder), Issuer/Status Service, Policy/ABAC Gate, and Audit/Ledger Anchor laid out top-to-bottom for A4 portrait. It depicts a complete wallet-based interaction where a person proves a fact without revealing a dossier.

From the Verifier lane, the flow begins with a predicate request carrying the declared purpose. In the Wallet lane, the holder sees a clear prompt (who is asking, what will be disclosed, purpose, and duration) and approves. The wallet refreshes revocation/status from the Issuer/Status lane if needed, then returns a presentation marked "selective disclosure / ZKP predicate." The Verifier performs trust checks (issuer keys, status freshness, audience and time bounds) and forwards the context to the Policy/ABAC Gate, which evaluates purpose, role, consent or lawful basis, risk, and freshness policy. The gate returns Allow / Step-up / Deny. On Allow, the verifier acts on a minimal decision (true/false), not on raw attributes. The Audit/Ledger lane records an event who, purpose code, method (selective disclosure or ZKP), outcome, and any consent/override reference then shows periodic anchoring to a permissioned ledger. Visual callouts emphasise "events, not data," "predicate truth only," and "audience-bound, time-bound proofs.

Holder-absent flow: APIs with consent and purpose controls

Many national processes run in the background: eligibility checks, reconciliations, cross-agency verifications. Holder-present SSI cannot cover every case. The alternative is not bulk data; it is narrow, purpose-bound questions that return minimal answers.

a) Request discipline
  - Purpose declaration is mandatory. The request states the purpose code and method (e.g., "licence validity check for traffic enforcement").
  - Per-agency pseudonym represents the subject so that cross-domain linking by traffic analysis is difficult.
  - Contextual risk (channel, device posture, time, location, request velocity) is part of the decision (Allen & Hess, 2022).

b) Decision logic
  - Allow when the method is permitted for the purpose, role is appropriate, risk is low, and consent or legal basis is satisfied.
  - Step-up when risk is moderate (manager co-sign, second factor, or deferred consent).
  - Deny when purpose mapping does not allow the method, quotas are exceeded, consent is missing, or legal basis is absent.

c) Response discipline
   - Truth value or short-lived token only. No dossiers, no long-lived identifiers.
   - Tight expiry prevents replay; single-use semantics are preferred for high-risk checks.
   - No superfluous metadata (no precise device hints or unnecessary timestamps) to avoid new tracking surfaces (Wagner & Eckhoff, 2020).

d) Consent and lawful basis: If richer data are needed, the process references a consent artefact captured earlier or a lawful basis (statutory duty, court order). Both are narrow and time-boxed, with a visible reference in the audit (European Commission, 2024).

e) Controls against misuse
   - Quotas and rate limits per purpose and per caller prevent scraping.
   - Anomaly detection flags burst patterns, unusual geography, and method–purpose mismatches.
   - Separation of duties prevents a single insider from both elevating access and clearing logs.

f) Interoperability and legacy: Legacy registries integrate behind the minimal interface. They answer a small set of standard questions rather than exposing raw tables. This reduces integration cost and keeps policy review tractable.

Consent, lawful overrides, and citizen visibility

a) Consent (the default for rich disclosures)
   - Specific, informed, time-limited. The person approves a request naming the requester, purpose, field(s), and duration.
   - Bound to events. Each approval produces a verifiable reference tied to subsequent disclosures.
   - Revocable. The person can revoke standing consents; future requests must fail fast or re-prompt.
   - Accessibility. Assisted capture at counters; translated prompts; alternatives to biometrics when needed.

b) Lawful overrides (the exception)
   - Defined in law and narrow by scope. Examples include imminent threats or court-mandated investigations.
   - Two-person rule and time-boxing. An override requires multi-party approval and automatic expiry.
   - Automatic regulator alerting. Overrides are copied to regulator dashboards; post-hoc review is mandatory.

- Citizen notice. Where lawful and safe, affected citizens are notified after the risk window (European Commission, 2024).

c) Citizen visibility

- A portal shows who accessed what, when, why, and how (proof vs. yes/no), and whether consent or an override was used.
- Exportable histories support complaints and appeals.
- Visibility shifts incentives: silent browsing becomes costly because it is seen by the citizen, the regulator, and supervisors (Dunphy & Petitcolas, 2020).

Immutable audit on permissioned ledger (events, not data)

a) What is logged

- Caller identity and role: Which agency, which service, which officer or system.
- Declared purpose: The exact purpose code that justified the request.
- Subject linkage: A per-agency pseudonym, not a global identifier.
- Method and outcome: Predicate proof vs. yes/no; allow/deny/step-up; result of the minimal check.
- Consent or override reference: A verifiable reference to the artefact or legal basis.
- Risk context: Coarse signals used in the decision (e.g., “unusual time,” “quota near limit”).

b) How integrity is protected

- Events are written to an append-only log and hash-chained so local tampering is detectable.
- Batches of events are periodically anchored to a permissioned ledger run by multiple state bodies and the regulator. This provides tamper-evidence and multi-party control without putting personal data on-chain (Juels & Oprea, 2020; Vukolić, 2021).
- Independent verification. Regulator nodes verify anchors and reconcile event counts; periodic public digests strengthen accountability.

c) Privacy of the audit

- Events, not data. Logs contain decision metadata, not raw personal attributes.
- Access views. Citizens see their own histories. Agencies see only their own calls. Regulators have full oversight.
- Retention. Logs are kept for legally defined periods, then archived with integrity proofs.

Security services: liveness, anomaly detection, privacy-preserving AI

a) Biometric liveness and presentation-attack detection
   - Where used: Enrolment, high-risk actions, recovery, and delegated authority.
   - How used: Challenge–response liveness for face or voice; behaviour-based checks after login for continuity.
   - Data handling: Templates held on device where possible; if server-side, they are processed in secure environments and not retained as raw media (Paredes-García et al., 2023; NIST, 2020).

b) Anomaly detection on the access fabric
   - Signals monitored: Request velocity, unusual times or geographies, purpose–method mismatches, repeated denials followed by sudden approval, and clustering of overrides.
   - Actions taken: Step-up for medium risk, temporary blocks for high risk, and alerts to supervisors and regulators for investigation.

c) Privacy-preserving AI for fraud and misuse
   - Federated learning: Agencies train models locally on their logs; an aggregator computes robust updates without centralising raw data (Ren et al., 2025).
   - Encrypted or enclave-based inference: For sensitive features, scoring happens without exposing inputs.
   - Bias governance: Error-rate gaps are tested before deployment and monitored in production; there are human appeal routes and rollback paths for problematic models (European Commission, 2024).

d) Zero-trust hardening around everything
   - Mutual authentication between agencies; strong request signing; replay protection.
   - Least privilege enforced by purpose maps; quotas and rate limits per purpose and per caller.
   - Separation of duties so no single person can both change policy and consume data.
   - Rapid key rotation and revocation to contain compromise (Allen & Hess, 2022).

e) Performance and reliability expectations
   - Latency budgets: ≤ 300 ms p95 for domestic minimal checks; ≤ 600 ms cross-border.
   - Audit completeness: 100% of accesses recorded; anchoring within a few minutes.

- Availability: Degradation modes prefer deny-by-default for high-risk requests and retry for lower-risk ones with operator guidance (UIDAI, 2025).

f) Threats and corresponding mitigations

- Insider "surfing": ABAC on every call, immutable audit, quotas, anomaly alerts, sanctions.
- Deepfakes and replay: Liveness, challenge-response, multi-factor step-up, regular model refresh.
- Linkability creep: Per-agency pseudonyms, rotating tokens, minimal metadata, predicate proofs.
- Ledger tampering: Multi-party consensus, independent regulator nodes, periodic public digests.
- Model poisoning: Update validation and rollback in federated training; provenance tracking for model artefacts (Vukolić, 2021; Ren et al., 2025).

g) What this delivers in practice

- For citizens: faster service with less exposure, a clear access history, and credible remedies.
- For verifiers: stable, small interfaces that answer the real question without legal risk.
- For regulators: truthful, tamper-evident logs and live anomaly insights.
- For the state: lower breach impact, fewer data silos, easier compliance, and a defensible, future-proof identity stack aligned with GDPR, eIDAS 2.0, and the EU AI Act (European Commission, 2024).

How the SSI–VWR Framework Works

A. Core mechanics problem solved, step by step

i. Identity & binding without surveillance creep:

A non-meaningful national identifier is issued and bound to the person at enrolment with strong evidence and liveness checks. This avoids meaningful, sequential numbers that leak demographics, and sets up safe recovery so people are not locked out (NIST, 2020).

ii. Authoritative credentials held by the person (SSI):

Government authorities issue verifiable credentials (VCs) for identity binding, age, residency, licence status, etc. The person holds them in a regulated wallet and decides when to disclose turning "check my data" into "prove a fact" (Sporny, Longley, & Chadwick, 2022; Preukschat & Reed, 2021).

iii. Minimal disclosure by default two complementary paths:

- Holder-present: the wallet presents selective disclosure or zero-knowledge proofs so the verifier learns only the predicate (e.g., “over 18,” “licence valid today”) and not a dossier (Camenisch & Lehmann, 2021).
- Holder-absent: back-office systems call yes/no attribute endpoints (e.g., “is residency current in municipality X?”), returning a truth value or short-lived token never full records (UIDAI, 2025).

iv. Purpose limitation implemented in code (zero-trust):

Every request carries a purpose code from a public catalogue (e.g., WELFARE_ENROLMENT, MOTOR_LICENSE_CHECK). An ABAC policy engine evaluates purpose, role, consent or legal basis, and contextual risk (time, location, velocity) before allowing or denying the call (Allen & Hess, 2022; Campbell & Weitzner, 2022).

v. Linkability suppressed at the root:

The same person appears as different per-agency pseudonyms derived from the national identifier and an agency salt. Responses strip non-essential metadata and use short-lived tokens, so cross-domain correlation is hard even for insiders (Troncoso et al., 2020; Camenisch & Lehmann, 2021).

vi. Consent for rich data; narrow lawful overrides for emergencies:

Routine checks proceed with minimal outputs. Rich fields (e.g., address string) need explicit, time-limited consent captured as a verifiable artefact; lawful overrides are rare, time-boxed, co-approved, and auto-notified to the regulator (European Commission, 2024).

vii. Immutable accountability without exposing personal data:

Every access becomes an event who, when, declared purpose, method (proof vs yes/no), policy outcome anchored to a permissioned ledger operated by multiple state entities and the regulator. Logs store events, not raw attributes; citizens can see their own history; regulators can investigate anomalies (Juels & Oprea, 2020; Vukolić, 2021; Dunphy & Petitcolas, 2020).

viii. Security and AI with restraint:

Biometric liveness is used for enrolment, recovery, and high-risk actions; templates stay on-device where possible. Access-pattern anomaly detection runs with privacy-preserving methods (federated learning, encrypted inference), and high-risk AI follows EU AI Act governance with bias monitoring and human appeal routes (Paredes-García et al., 2023; European Commission, 2024; Ren et al., 2025).

ix. Performance & migration that work at national scale:

Predicate proofs and yes/no checks are compact, cacheable, and meet tight latency SLOs (p95 hundreds of ms). Migration is phased: add immutable audit to legacy interfaces, switch high-volume use cases to minimal endpoints, roll out wallet proofs, retire bulk exports keeping services running while privacy improves (UIDAI, 2025; European Commission, 2024).

B. Case verifying a person's identity and residency in the framework

A resident applies for a welfare benefit at a municipal office. The authority must confirm (i) the person is the rightful holder of the national identity, and (ii) the person currently resides in the municipality. Surveillance risks to avoid: revealing full birth date and address to front-line staff, cross-agency "surfing," and silent dossier copies.

i. At the counter (holder-present, SSI):

The officer's screen shows purpose = WELFARE_ENROLMENT and requests two wallet proofs: identity binding and municipal residency. The holder approves. The wallet returns a cryptographic presentation that proves binding to the state-issued identity and a residency predicate ("residentOf = Municipality M") no birth date, no full address. The verifier checks issuer authenticity, revocation, time bounds, and audience binding. Both proofs verify, so the case proceeds without looking up back-end registries (Camenisch & Lehmann, 2021; Sporny, Longley, & Chadwick, 2022).

ii. That night (holder-absent, VWR yes/no):

The case system performs a batch yes/no re-check: "Is residency still within Municipality M for purpose WELFARE_PAYMENT?" ABAC evaluates role, purpose, existing consent artefact from enrolment (or statutory basis), and contextual risk. The civil registry returns true with a short-lived token. No address string is exposed; the token suffices to authorise payment routing (Allen & Hess, 2022).

iii. Rare exception (rich data with consent):

If a delivery vendor needs the address line to drop equipment at home, the agency triggers a clear consent prompt naming the requester, field, purpose, and retention period. On approval, only the address field is released as a signed attribute bundle; retention is short and enforced. If consent is refused or expires, access stops. This keeps rich data exceptional and traceable (European Commission, 2024).

iv. Emergency edge case (lawful override):

If a court later mandates an urgent address disclosure for a criminal investigation, an override is registered with scope and timebox, co-approved by two officers, and auto-alerted to the regulator. The disclosure is narrow and expires quickly. Where lawful and safe, the resident is notified after the risk window. Overrides are visible, not backdoors (European Commission, 2024; Dunphy & Petitcolas, 2020).

v. Accountability and deterrence throughout:

Each action above two wallet proofs, one yes/no re-check, and any consented or overridden disclosure is logged as a separate audit event with purpose, method, and decision. Events are anchored to the permissioned ledger. The resident later opens their portal and sees the exact sequence in plain language; the regulator sees population-level patterns and anomalies. This visibility deters “surfing” and enables sanctions where needed (Juels & Oprea, 2020; Vukolić, 2021).

The authority makes correct, timely decisions while disclosing the least possible information. Insider browsing becomes risky because it is visible and immutable; linkability is constrained by design; and legal duties data minimisation, purpose limitation, and high-risk AI oversight are enforced in system behaviour, not only in policy text (European Commission, 2024; Allen & Hess, 2022).

Interoperability and Standards

Interoperability is the difference between a neat pilot and a national utility. The SSI–VWR framework therefore anchors its cryptography and policy flows in open, widely adopted standards so credentials, proofs, and policy decisions travel across agencies, sectors, and borders without custom integration. This section lays out the standards stack what to adopt, how to combine it, and where profiles are needed to keep implementations aligned.

Alignment with eIDAS 2.0 and the European Digital Identity Wallet

- Core idea. The EUDI Wallet standardises holder-controlled proofs of attributes across the EU. SSI-VWR treats the wallet as the preferred path when the person is present, and uses yes/no verification for holder-absent cases, preserving the same data-minimisation logic in both modes (European Commission, 2024).
- Roles and trust. eIDAS defines recognised issuers, verifiers, and qualified trust services. SSI-VWR maps these directly to the issuer–holder–verifier model and to the policy steward and regulator roles needed for purpose enforcement and audit. Trust lists published under eIDAS become the “roots of trust” for issuer discovery and key roll-over (European Commission, 2024).
- Cross-border use. An age, licence, or residency proof created in one member state must verify in another without new bilateral agreements. The framework insists on profiled credential types and status mechanisms that all verifiers can process, and it documents error handling (e.g., what to do if revocation is temporarily unavailable) so public counters can stay serviceable (European Commission, 2024).

Credentials and Identifiers (W3C VC 2.0, DIDs)

- Verifiable Credentials (VC 2.0). Credentials carry authoritative claims (age, licence status, residency). Using VC 2.0 means a consistent data model, predictable proofs, and a standard way to attach revocation/status references (Sporny, Longley, & Chadwick, 2022).
- Decentralized Identifiers (DIDs). Issuers and verifiers publish DID documents so wallets can resolve keys and service endpoints without a single national directory. EU programmes can prioritise DID methods aligned with eIDAS trust lists (e.g., did:ebsi), while still accepting others where policy allows (Sporny et al., 2022; European Commission, 2024).
- Identifier hygiene. The person's national identifier is non-meaningful and never used as a global correlation key. SSI-VWR derives per-agency pseudonyms for back-office checks, and wallets avoid reusing stable identifiers across verifiers in holder-present proofs to suppress linkability (Camenisch & Lehmann, 2021; Troncoso et al., 2020).

Selective Disclosure and Predicate Proofs

- Mechanisms. Two complementary families are mature: (i) BBS+ / unlinkable selective disclosure in JSON-LD VCs, and (ii) SD-JWT-VC for compact, web-friendly selective disclosure. Both support revealing one attribute from a multi-attribute credential without leaking the rest (Khovratovich & Law, 2020; Sporny et al., 2022).
- Predicates. For "over 18," "licence valid," or "resident in district X," the framework prefers zero-knowledge predicates that reveal only truth values. This lets the verifier decide without receiving birth dates, addresses, or long identifiers (Camenisch & Lehmann, 2021).
- Profile discipline. To avoid fragmentation, the programme publishes national SSI profiles: which signature suites to accept, how to express common claims, what clock skew to tolerate, and how to check status. This prevents "works on my wallet" failures in production.

Status, Revocation, and Freshness

- Status lists. The framework adopts StatusList 2021/2023 style bit-vectors so verifiers can check whether a credential is suspended or revoked without contacting the issuer for every transaction. Lists are cacheable, shardable, and rotate predictably to keep bandwidth low (Sporny, Longley, & Chadwick, 2022).
- Freshness policies. Not every claim needs the same freshness. A licence validity check may require up-to-the-minute status; an age-

over proof can tolerate daily lists. The purpose catalogue encodes these freshness targets so verifiers behave consistently (Allen & Hess, 2022).

- Fallback rules. If status is temporarily unreachable, policy defines safe fallbacks per purpose (deny, retry, or allow with post-event verification) to keep counters moving without eroding assurance (European Commission, 2024).

Presentation Protocols and Federation

- OpenID for Verifiable Presentations (OID4VP) and Issuance (OID4VCI). These standardise how wallets receive requests, how proofs are returned, and how credentials are issued or refreshed, reducing custom glue between government portals, private verifiers, and wallets (European Commission, 2024).
- Cross-ecosystem bridges. Many verifiers already use OpenID Connect for login and FIDO2/WebAuthn for phishing-resistant authentication. SSI-VWR treats them as complementary: OIDC for session and account binding; VCs for attribute assurance with minimal disclosure (NIST, 2020).
- Sector adapters. Health, finance, mobility, and education have their own data dictionaries. The framework defines mapping guides (e.g., mDL → driving entitlement, IBAN/name match for KYC) so each sector can accept minimal proofs without inventing new credential types (ISO/IEC, 2021; Arner, Barberis, & Buckley, 2020).

Interop with Existing Government Credentials (mDL, ePassports, Cards)

- Mobile Driving Licence (ISO/IEC 18013-5). mDLs are signed attribute containers. The framework maps their driving entitlement and document validity fields to VC claims and yes/no checks, so roadside or portal verifiers can accept both mDL and wallet proofs with the same minimal-disclosure semantics (ISO/IEC, 2021).
- ePassports (ICAO 9303) and residence cards. Where travel or residency proofs are needed, data groups from machine-readable travel documents can be wrapped as verifiable claims for domestic services, avoiding photocopies and reducing fraud, while preserving signature verification paths (Dunphy & Petitcolas, 2020).
- Card coexistence. Physical smartcards remain available as accessibility fallbacks. Their chip-based proofs should map to the same predicates used in wallets so staff procedures and audits remain uniform.

Back-Office Interoperability

- Canonical questions. Holder-absent services converge on a small, stable vocabulary of yes/no questions: identity-binding valid, licence valid, residency current, name-matches, eligibility flag. Designing back-office around canonical questions simplifies integration and audit (UIDAI, 2025).
- Policy-first federation. Each request carries a purpose code and is evaluated under ABAC with context (role, consent/legal basis, risk). This creates a federation of decisions rather than a federation of databases, which is crucial for GDPR data minimisation (Allen & Hess, 2022; Campbell & Weitzner, 2022).
- Metadata hygiene. Responses avoid stable identifiers and unnecessary device hints to prevent accidental cross-domain tracking (Wagner & Eckhoff, 2020).

Cryptographic Agility and Post-Quantum Readiness

- Agility. Keys and signature suites rotate on a documented schedule; verifiers accept a short, bounded set of suites to keep interop predictable while allowing upgrades (Sporny et al., 2022).
- Post-quantum (PQ) roadmap. The programme publishes a migration plan for hybrid classical+PQ signatures in credentials and presentations once standards stabilise. Ledger anchoring and audit proofs also receive a PQ review so integrity remains durable over decades (Juels & Oprea, 2020; European Commission, 2024).
- Evidence provenance. Issuers attach concise provenance notes (how identity was established, liveness methods) to support cross-jurisdiction acceptance without revealing raw biometrics.

Assurance Levels and Cross-walks

- LoA mapping. Member states use different Levels of Assurance (LoA) for identity proofing and authentication. SSI-VWR maintains a cross-walk (e.g., eIDAS High ↔ NIST IAL2/AAL2) so relying parties can accept foreign credentials with confidence while maintaining sector policy (NIST, 2020; European Commission, 2024).
- Predicate-specific assurance. A single wallet can hold claims with different assurance. Verifiers evaluate the assurance of the claim used, not of the whole wallet, which keeps minimal disclosure compatible with strong risk management.

Conformance, Testing, and Change Management

- Test suites and plug-fests. Programmes run public interoperability events and automated conformance suites for wallets, verifiers, and issuers. Passing means not only verifying signatures but also enforcing purpose maps, freshness targets, and minimal-disclosure behaviour in edge cases (Sporny et al., 2022; European Commission, 2024).
- Stable profiles with versioning. The state publishes versioned implementation profiles: credential types, accepted proof types, status/refresh rules, wallet UX norms, and error semantics. Deprecations come with long notice and migration aids.
- Transparency in operations. Quarterly reports include: minimal-disclosure ratio, override counts, revocation freshness compliance, interop failures, and corrective actions. Regulators participate in profile governance to prevent drift (European Commission, 2024).

Governance of Standards in Practice

- Public purpose catalogue. The most powerful "standard" is the list of allowed purposes and their minimal attributes. Publishing this catalogue and the policy-as-code that enforces it lets developers, auditors, and civil society test that the system does what the law says (Campbell & Weitzner, 2022).
- Sanctions and incentives. Procurement and accreditation reward verifiers that adopt minimal-disclosure interfaces and wallet proofs; sanctions apply to agencies that keep pulling dossiers where a predicate would do. This aligns market behaviour with the standards posture (European Commission, 2024).
- International coordination. For non-EU partners, the framework relies on mutual recognition anchored in public trust lists, LoA cross-walks, and common VC/VP profiles, avoiding bespoke bilateral gateways that are costly and fragile (Dunphy & Petitcolas, 2020).

By grounding credentials in W3C VC/DID, proofs in selective-disclosure and predicate standards, wallet flows in the EUDI profile, and holder-absent checks in a small set of purpose-bound yes/no interfaces, SSI–VWR turns "minimal disclosure" into portable behaviour. Add explicit LoA cross-walks, revocation freshness policies, cryptographic agility, and regular conformance testing, and the result is an identity fabric that works within a country, across the EU, and with trading partners without reverting to surveillance-prone data sharing (European Commission, 2024; Sporny, Longley, & Chadwick, 2022; NIST, 2020; ISO/IEC, 2021).

Implementation Blueprint

This blueprint turns the SSI–VWR theory into an operational plan that ministries and regulated verifiers can deploy without halting existing services. It focuses on small, stable interfaces, policy baked into runtime, immutable accountability, and measurable gates that prove privacy-by-default is working at national scale (Allen & Hess, 2022; European Commission, 2024; Dunphy & Petitcolas, 2020).

Reference interfaces for minimal verification

The platform exposes a compact set of canonical questions that cover most government and regulated use. These questions are intentionally narrow and map one-to-one to legal purposes.

- Predicate checks: over-18, over-65, licence-valid-today, residency-in-municipality, identity-binding-valid, name-matches-for-bank-account. Each call returns a truth value or a short-lived transaction token. The token exists only to let a downstream step confirm that the check was done; it expires rapidly and cannot be used for tracking (UIDAI, 2025; Allen & Hess, 2022).
- Selective disclosure presentations: for holder-present cases, verifiers ask the wallet to prove a predicate or reveal a single attribute, never an entire dossier. The verifier validates issuer trust, revocation status, time bounds, and audience binding before acting (Camenisch & Lehmann, 2021; Sporny, Longley, & Chadwick, 2022).
- Richer attributes by exception: when a specific field (for example, an address line) is unavoidable, the platform requires a verifiable consent artefact or a documented lawful basis tied to the event and visible in audit (European Commission, 2024).

These interfaces are documented with plain-language semantics (what question is being asked, which purpose allows it, what the possible outcomes are), error-handling rules aligned to risk (deny, retry, or queue), and freshness targets per attribute (for example, licence status must be real-time; age-over can be day-old) (Allen & Hess, 2022).

Data models and identity hygiene

The blueprint assumes a non-meaningful national identifier bound to the person at enrolment through strong evidence and liveness. No birth date, region code, or sequence leaks from the identifier. For holder-absent checks, the platform derives a per-agency pseudonym so that the same person appears differently to different agencies; this is deterministic within an agency and rotates on policy schedule to suppress linkability (Camenisch & Lehmann, 2021; Troncoso et al., 2020). Credentials are issued as Verifiable

Credentials (VC 2.0) with thin, well-scoped schemas: AgeCredential, Residency Credential, License Status Credential, and Identity Binding Credential. Each schema includes a status or revocation reference, an issuance timestamp, and assurance notes that explain how identity was established and what liveness was used at enrolment, without exposing raw biometrics. Schemas are versioned and deprecation is coordinated through a public change calendar so wallets and verifiers stay in lockstep (Sporny, Longley, & Chadwick, 2022; European Commission, 2024). The purpose catalogue is a first-class artefact. Each purpose has a human-readable description, the minimal attributes or predicates it permits, the lawful bases that can justify a request, retention limits, and escalation paths for overrides. The catalogue is published and versioned; every change is reviewed by policy, security, and the regulator to avoid purpose creep (Campbell & Weitzner, 2022).

Policy-as-code and zero-trust enforcement

Every request wallet proof or yes/no check passes through a decision point that evaluates: declared purpose; caller identity and role; consent or lawful basis state; contextual risk (time, location, device posture, request velocity); and freshness of status information. Outcomes are allow, deny, or step-up (extra authentication or supervisory co-approval). This is zero-trust in practice: no implicit internal trust, and least privilege by default (Allen & Hess, 2022).

- Purpose–method matrix that blocks any method not mapped to the purpose; requests for dossiers are refused or automatically rewritten into predicates where feasible.
- Quotas and rate limits per purpose, per caller, and per organisational unit to deter scraping and catch automation gone wrong.
- Contextual risk scoring that raises friction in suspicious contexts (e.g., sudden bursts at night from an unusual location), with clear operator guidance and appeal paths.
- Separation of duties so no one actor can both change policy and consume data; administrative actions are logged and independently reviewed.

Policy is stored in a repository with peer review, automated tests for regressions, and a canary rollout path, so changes are traceable and reversible (Campbell & Weitzner, 2022).

Wallet and verifier experience

Wallet UX is the instrument that makes minimisation normal. Prompts must tell the user who is asking, what will be disclosed (predicate or attribute), for what stated purpose, and for how long approval lasts. Approvals are time-limited and revocable; a single tap shows past approvals and lets the user withdraw standing consent. Delegation is built in for carers and parents, with time-boxed authority and easy revocation (Dunphy & Petitcolas, 2020). Verifier UX is designed to make the minimal path the easiest path. For holder-present flows, the default is to ask the wallet for a predicate or a single attribute; for back-office flows, the default is a yes/no check tied to a purpose. UI and SDKs warn when a request exceeds what the purpose allows and explain lawful alternatives (for example, "request licence-valid predicate instead of full licence record"). Staff training uses realistic scenarios and emphasises that requesting more creates legal risk and slows the case (Allen & Hess, 2022; European Commission, 2024). Citizen portal turns accountability into something people can see. Each access appears in plain language with who/when/why/how, including whether consent or an override was used. Exports support complaints and discovery. Accessibility requirements are applied to both wallet and portal so that disability, language, and bandwidth do not exclude users (Dunphy & Petitcolas, 2020).

Immutable audit and regulator visibility

All accesses are written as events that include the caller identity and role, the declared purpose, the subject's per-agency pseudonym, the method used (wallet predicate, wallet attribute, yes/no check), the policy outcome (allow/deny/step-up), and references to any consent artefact or lawful override. Events are hash-chained locally for order and integrity, then anchored periodically to a permissioned ledger operated by multiple state entities and the regulator. The ledger contains only integrity metadata, not personal data, which preserves privacy while making tampering evident (Juels & Oprea, 2020; Vukolić, 2021). The regulator runs independent nodes, receives automatic alerts on overrides, quota breaches, and anomaly clusters, and publishes quarterly transparency reports: minimal-disclosure ratio; override counts and justifications; fairness and error metrics for any high-risk AI; and corrective actions taken. Agencies receive their own dashboards for internal supervision. Citizens can verify inclusion of their own events via the portal without needing to understand the underlying cryptography (European Commission, 2024; Dunphy & Petitcolas, 2020).

Security services, assurance, and performance targets

Key and channel security rely on mutual authentication between agencies, strong request signing, replay protection, rapid key rotation, and strict logging of administrative actions. Biometric liveness is used at enrolment, recovery, and high-risk actions; templates are kept on device where possible, with server-side processing isolated and non-retentive when necessary (NIST, 2020; Paredes-García et al., 2023). Anomaly detection watches the access fabric for burst patterns, time–location anomalies, method–purpose mismatches, and suspicious sequences (e.g., repeated denials followed by a sudden approval). To protect privacy, models are trained with federated learning or use encrypted/enclave inference where sensitivity warrants it; models are governed under the EU AI Act with risk files, bias testing, human oversight, and rollback procedures (European Commission, 2024; Ren et al., 2025; Kaul, 2021). SLOs keep the system usable at scale: p95 latency of ≤ 300 ms for domestic yes/no checks and ≤ 600 ms cross-border; 100% audit coverage with anchoring within minutes; transparent error handling that prefers deny-by-default for high-risk requests and safe retries for low-risk ones. Caching of revocation and issuer trust data, idempotent request semantics, and circuit breakers on registries maintain service during spikes. Minimal-disclosure ratio is tracked as a primary KPI with a maturity target of ≥ 85% of transactions resolved via predicates or yes/no responses (UIDAI, 2025; Allen & Hess, 2022).

Rollout plan and measurable gates

A phased rollout avoids disruption while improving privacy:

1. Audit-everywhere: wrap existing interfaces with immutable audit; launch the citizen portal showing access history. Success metric: 100% access coverage in audit; regulator node operational (Dunphy & Petitcolas, 2020).
2. Yes/no by default: convert the top high-volume use cases to minimal checks; publish the purpose catalogue; enforce ABAC at the gateway. Success metric: minimal-disclosure ratio ≥ 60% within six months; p95 latency ≤ targets (Allen & Hess, 2022).
3. Wallet-first proofs: issue core credentials (age, licence status, residency); enable selective disclosure for holder-present flows across counters and major portals. Success metric: wallet adoption in target cohorts; reduction in dossier pulls at counters; usability scores above threshold (European Commission, 2024).
4. AI hardening and bias governance: deploy federated anomaly detection; formalise AI risk files and fairness monitoring; publish the first transparency report. Success metric: anomaly detection precision/recall at targets; bias gaps within policy bounds (Paredes-García et al., 2023).
5. Legacy retirement: decommission bulk exports and undocumented lookups; mandate minimal responses and policy checks in procurement. Success metric: zero active endpoints that return dossiers for predicate use cases; external audit attestation.

Case Studies and Mapping to SSI–VWR

Estonia: transparency, separated registries, and integrity anchoring

Estonia shows that digital government can be fast without becoming a surveillance system. The state connects many sector registries through a secure exchange layer, but it does not expose a single super-database to every clerk. Each query is authenticated, checked against policy, and written to an audit trail that citizens and supervisors can read. People can log in and see which agency looked up which record and when. Log integrity is protected with cryptographic anchoring so that neither an insider nor an attacker can silently edit the past (Dunphy & Petitcolas, 2020). This design maps closely to the SSI–VWR framework. The holder-present path is visible in services that accept signed attributes and avoid fresh database pulls when a credential suffices. The holder-absent path appears in narrow back-office checks that return only what a legal purpose allows. The purpose catalogue is implicit in Estonia's service-by-service access rules, and immutable audit aligns with the framework's ledger-based tamper evidence. The remaining gap is selective disclosure at scale across all sectors. SSI adds that capability with wallet-based predicate proofs, reducing the need for staff to view full records even when the citizen is present. The lesson for other states is simple. Separate the registries, make every access visible, and move routine interactions to minimal proofs; trust grows because accountability is a daily experience, not a promise (Dunphy & Petitcolas, 2020).

European Union wallet pilots: selective disclosure and cross-border scale

The European Digital Identity Wallet turns the privacy rule of data minimisation into a user routine. A person proves a fact age over a threshold, licence valid, university status without handing over a dossier. Verifiers check cryptographic proofs against recognised issuers and revocation sources that are shared across borders. Early pilots show that the same wallet flow can work in another member state with no bespoke integration, which is essential for labour mobility, study, and travel (European Commission, 2024). This is the holder-present pillar of SSI–VWR. The mapping is direct. Wallet prompts explain purpose in plain language. Predicate proofs reveal only the truth needed. Revocation freshness is tuned to risk so verifiers do not stall. The framework extends the pilots by adding two elements. First, a back-office yes/no path for holder-absent checks that mirrors the same minimal logic under zero-trust policy. Second, a public purpose catalogue and policy-as-code that make minimisation enforceable for developers. Together they prevent drift back to dossier pulls as services expand, and they give regulators a way to test compliance across many agencies. The grand outcome is interoperability with restraint: proofs travel, but dossiers do not, and every exceptional access is tied to consent or a specific lawful basis that appears in the audit (European Commission, 2024).

India's Aadhaar: yes/no at population scale with consented e-KYC

Aadhaar demonstrates that binary answers can support very large programmes. Many verifiers need to know whether a claim is true, not to see the full record. Aadhaar authentication returns yes/no for identity binding and basic attributes, while richer e-KYC flows occur with consent or specific legal grounds. Separation between these paths limits unnecessary spread of personal data. Each authentication leaves an audit entry that can be reviewed, and operational hardening has added tokenisation features to reduce linkability across relying parties (UIDAI, 2025). This aligns with the holder-absent side of SSI–VWR. The mapping is clear. Narrow questions, fast answers, and strong logging match the framework's minimal APIs and immutable audit. Where India faced pressure privacy concerns, over-collection by some verifiers, and biometric spoofing later reforms moved practice closer to the framework: tighter consent capture, stricter retention limits, more granular tokens, and stronger liveness checks. The key lesson is that small interfaces scale better than broad data pulls. They keep latency low, focus engineering on a few canonical questions, and reduce legal risk. SSI complements this with holder-present selective disclosure, preventing verifiers from copying rich data when the citizen is on-site, and per-agency pseudonyms reduce linkability in background traffic. The combined pattern SSI at the edge, VWR in the back office gives speed without dossier creep (UIDAI, 2025; Dunphy & Petitcolas, 2020).

Zug, Switzerland: municipal SSI and citizen control

Zug's municipal pilot proved that government-issued credentials can live in a citizen wallet and still deliver assurance. The city attested to residency; people used a wallet to prove that fact for local services without a fresh registry lookup each time. The pilot was modest in scale, but it showed two important things. First, selective disclosure can handle routine administration without exposing extra fields. Second, citizens accept digital flows when prompts are clear and recovery is practical (Dunphy & Petitcolas, 2020). In SSI–VWR terms this is the holder-present path working as intended. The missing pieces for national scale policy enforcement, ledger-anchored audit, and yes/no back-office checks are what VWR supplies. The mapping therefore suggests a path for municipalities and agencies: start with wallet-based proofs for the person-facing parts of a service, wrap legacy access with immutable audit, and convert background checks to narrow yes/no calls bound to a published purpose. The result is consistent behaviour across channels and levels of government, with one access history that a citizen can read.

Comparative lessons and the framework's fit

Across these contexts a pattern emerges. Adoption rises when convenience and control move together. Estonia's access history portal is as important as its cryptography. EU wallet prompts make selective disclosure normal. Aadhaar's yes/no checks keep lines short while consented e-KYC handles exceptions. Zug's wallet gives people visible control over disclosure. The SSI–VWR framework collects these elements into one deployable plan. It adds per-agency pseudonyms to suppress linkability, policy-as-code to turn purpose limitation into runtime behaviour, and permissioned-ledger anchoring so audit trails are tamper-evident and reviewable by regulators and citizens (Dunphy & Petitcolas, 2020; European Commission, 2024; UIDAI, 2025). It also addresses the pitfalls seen in practice. Over-centralised designs invite silent browsing; the framework answers with zero-trust decisions on every call and quotas that deter scraping. Weak logging allows disputes to devolve into opinion; the framework answers with immutable events and public transparency reports. Centralised biometrics create high stakes breaches; the framework keeps templates on device where possible and uses liveness and privacy-preserving AI with bias governance to maintain assurance without building new data pools (Paredes-García et al., 2023). Finally, the framework offers a phased migration that matches institutional capacity: audit first, minimal yes/no checks next, wallet proofs at the edge, and retirement of bulk interfaces last. In short, these case studies do not just inspire the design; they validate that its parts work in the real world and show how to combine them to reduce unconsented disclosure, stop cross-agency surfing, and keep national services fast, lawful, and trusted.

Evaluation and Risk

Security, privacy, interoperability, and equity outcomes

The SSI–VWR design measurably reduces attack surface by removing routine dossier pulls and replacing them with predicate proofs or yes/no checks. Security gains appear in three places. First, data minimisation lowers breach impact because fewer attributes traverse networks or sit in verifier systems. Second, zero-trust evaluation on every call collapses the old perimeter model; abuse now requires defeating a decision point that binds purpose, role, consent/legal basis, and contextual risk (Allen & Hess, 2022). Third, immutable audit raises the cost of insider misuse because every access is verifiably recorded and regulator-visible (Juels & Oprea, 2020; Vukolić, 2021). Programmes should track (i) the rate of blocked requests due to purpose mismatch, (ii) time to detect and contain insider anomalies, and (iii) blast-radius metrics average number of attributes exposed per incident expecting steady decline as minimal responses dominate (European Commission, 2024). Privacy. Privacy outcomes hinge on two measurable effects. The first is the Minimal-Disclosure Ratio (MDR) the share of all identity transactions resolved with a predicate proof or yes/no answer. The target in national steady state is ≥ 85%, with critical sectors (licensing, border, welfare eligibility) exceeding 90% (UIDAI, 2025). The second is cross-domain linkability. SSI–VWR suppresses it through holder-present proofs that avoid stable identifiers and holder-absent per-agency pseudonyms derived from a non-meaningful national identifier; responses strip non-essential metadata and use short-lived tokens, which reduces long-term correlation (Camenisch & Lehmann, 2021; Troncoso et al., 2020). Programmes audit MDR monthly and run periodic linkage tests on access logs to confirm that correlation risk remains below policy thresholds (Wagner & Eckhoff, 2020). Interoperability is assessed along technical and policy axes. Technically, the framework aligns with W3C Verifiable Credentials and DIDs and the European Digital Identity Wallet profiles for presentations and revocation/status, so proofs verify across borders without bespoke wiring (Sporny, Longley, & Chadwick, 2022; European Commission, 2024). Policy-wise, a purpose catalogue and policy-as-code travel with services: the same purpose name maps to the same minimal attributes in every agency and member state, improving predictability for developers and auditors. Interop KPIs include cross-border verification success rate, revocation-freshness conformance, and share of verifiers certified against the national profile. Equity improves when strong assurance does not require high disclosure and when recovery is safe and simple. The wallet path supports assisted channels and delegated consent; the back-office path avoids copying attributes that later disadvantage people if leaked. High-risk AI (liveness, anomaly detection) is governed by pre-deployment bias testing, drift monitoring, and human appeal routes as required by the EU AI Act (European Commission, 2024). Equity KPIs include false-reject rates for liveness by device class or demographic proxy (where lawful), successful recovery rates, and time-to-resolution for appeal workflows (Paredes-García et al., 2023).

Threat model, limitations, and mitigations

Insider “surfing.” The classic risk is broad staff browsing under vague role permissions. Mitigation is structural: every access is purpose-bound and ABAC-evaluated; quotas and rate limits apply per purpose and per caller; events are ledger-anchored and regulator-visible. Sanctions and public transparency reports sustain deterrence (Allen & Hess, 2022; Juels & Oprea, 2020). Linkability through metadata. Even minimal proofs can leak patterns if stable identifiers or rich metadata persist. Mitigation is per-agency pseudonyms, rotating tokens, coarse timing in non-critical responses, and periodic privacy reviews that test real-world linkability using audit telemetry (Camenisch & Lehmann, 2021; Troncoso et al., 2020). Biometric spoofing and capture risk. Presentation attacks (photos, replays, deepfakes) and poor capture conditions can defeat naive systems or exclude users. Mitigations are challenge–response liveness, device-side templates where feasible, secure execution for server-side matching, documented fallback paths (e.g., possession + knowledge or assisted in-person recovery), and ongoing model refresh with red-team drills (NIST, 2020; Paredes-García et al., 2023). Model bias and drift. Liveness and anomaly models can underperform for some groups or decay over time. Mitigations include curation of diverse training data, pre-deployment fairness tests, runtime drift alarms, and human oversight for contested decisions, as codified by the AI Act (European Commission, 2024). Over-centralisation pressure. Large programmes tend toward convenience interfaces that reintroduce dossiers. Mitigation is governance: procurement and accreditation reward minimal interfaces; the purpose catalogue is public; a change-control board rejects new data scopes unless no predicate solution exists; MDR becomes a hard KPI for agency leads (Campbell & Weitzner, 2022). Ledger and audit integrity. If audit writers are compromised, they may attempt to suppress or reorder events. Mitigations are append-only hash-chains, frequent multi-party anchoring to a permissioned BFT ledger, independent regulator nodes, and external attestations; logs contain events, not data, to minimise privacy exposure (Vukolić, 2021; Juels & Oprea, 2020). Supply-chain and key compromise. Wallets, SDKs, and issuer keys are attractive targets. Mitigations include key rotation, issuer trust-list monitoring, reproducible builds for wallet components, secure enclaves for sensitive material, and revocation drills. DID/VC agility prevents hard lock-in to one crypto suite (Sporny et al., 2022). Operational limits. National workloads stress revocation freshness, cache consistency, and audit throughput. Mitigations: graceful degradation policies (deny, retry, or allow with post-event check according to purpose risk), CDN-distributed status artifacts, and audit pipelines sized for peak load with back-pressure and lossless queues (European Commission, 2024; UIDAI, 2025). Legal/organizational capture. Purpose creep can arrive via policy, not code. Mitigation is plural governance: regulator seats on change boards, civil-society input, public minutes for purpose changes, and mandatory post-hoc reviews of overrides with sanction powers (European Commission, 2024; Dunphy & Petitcolas, 2020).

Monitoring, bias audits, and red-team practice

Operational monitoring. The platform tracks a small set of leading indicators: MDR by agency and purpose; policy-denial rate and top denial reasons; override counts with time-boxing compliance; revocation-freshness conformance; p95/p99 latency; and audit anchoring delay. Spikes in denials, sudden drops in MDR, or anchor delays trigger incident workflows (Allen & Hess, 2022). Privacy monitoring. Programmes conduct periodic linkage studies on audit telemetry (with privacy safeguards) to measure real-world linkability across agencies. A rising trend prompts review of pseudonym derivation schedules, token lifetimes, and metadata hygiene (Troncoso et al., 2020; Wagner & Eckhoff, 2020). Fairness and bias audits. For liveness and anomaly models, teams run pre-deployment tests (error-rate parity, ROC shifts under lighting/device variations) and post-deployment monitoring (false accept/reject by context). Results and remediation are published in quarterly transparency reports, with human appeal routes and rollback plans for underperforming models (European Commission, 2024; Paredes-García et al., 2023). Red-team exercises. Regular adversarial drills cover (i) deepfake and replay attempts against liveness; (ii) insider scraping masked as normal traffic; (iii) attempt to alter or suppress audit events; (iv) model-poisoning of federated updates; and (v) mis-scoped policy changes that expand attributes silently. Findings feed into policy hardening, SDK defaults, and staff training (NIST, 2020; Ren et al., 2025). Transparency and public trust. Monitoring culminates in public transparency reports that disclose MDR, overrides, incident summaries, fairness metrics, and corrective actions. Reports show not only outcomes but also governance behaviour: which purpose expansions were rejected, how many policy PRs were rolled back, and how often regulators exercised stop powers. This level of candour is central to legitimacy and is consistent with GDPR's accountability and the AI Act's documentation duties (European Commission, 2024; Dunphy & Petitcolas, 2020). Evaluation in the field. Beyond dashboards, agencies should run before–after studies when converting a use case from dossier pulls to predicates (e.g., licence checks, residency confirmation). Expected effects are lower average attributes transferred, shorter handling times, lower breach impact estimates, and unchanged or improved decision accuracy. Where feasible, A/B pilots can compare staff behaviour under legacy vs. SSI–VWR defaults, measuring browsing attempts, step-ups, and appeals. Publication of these results, including negative findings, builds scientific credibility and guides iteration (Allen & Hess, 2022; European Commission, 2024).

Discussion

Policy compatibility and governance models

The SSI–VWR design translates European legal obligations into ordinary system behaviour rather than exceptional paperwork. GDPR's data minimisation and purpose limitation become defaults because the platform's normal outputs are yes/no answers or predicate proofs; rich attributes flow only with explicit consent or a documented lawful basis, each bound to a declared purpose and a time window that the audit can show later (European Commission, 2024). Accountability is operationalised through the append-only audit that records who asked what, under which purpose, and with what outcome; the regulator's independent anchoring and dashboards make those records verifiable and actionable instead of theoretical (Juels & Oprea, 2020; Vukolić, 2021). Data protection by design and by default is reflected in per-agency pseudonyms, metadata hygiene, and short-lived tokens, which suppress linkability even inside government networks (Camenisch & Lehmann, 2021; Troncoso et al., 2020). The framework also complements eIDAS 2.0 and the European Digital Identity Wallet. Holder-present interactions rely on wallet-based selective disclosure that conforms to recognised VC/DID and EUDI profiles, so a proof accepted in one member state will verify in another without ad-hoc agreements (Sporny, Longley, & Chadwick, 2022; European Commission, 2024). Where the holder is not present, the same data-minimisation logic persists through purpose-bound yes/no checks. This continuity removes the familiar gap in which privacy-preserving proofs at the counter are later undone by dossier pulls in the back office. For high-risk AI (notably biometric liveness and some anomaly detection), the SSI–VWR governance track aligns with the EU AI Act: documented risk files, pre-deployment testing for accuracy and bias, human oversight, incident logging, and periodic public reporting (European Commission, 2024; Paredes-García et al., 2023). Importantly, the framework's policy-as-code allows regulators and internal auditors to test whether rules enforce lawful bases and purpose mappings at runtime instead of relying only on policy manuals (Allen & Hess, 2022; Campbell & Weitzner, 2022). Institutionally, the design supports plural governance. The policy steward publishes and versions the purpose catalogue, with change control that includes security, legal, sector regulators, and ideally civil society observers. The regulator operates independent integrity nodes, receives automatic alerts for overrides and anomaly clusters, and can impose sanctions that range from fines to temporary suspension of an agency's access. Parliamentary committees and data protection authorities gain a clearer line of sight because the artefacts they need purpose maps, audit extracts, transparency metrics are generated by the platform itself. Internationally, mutual recognition rests on trust lists for issuers, level-of-assurance cross-walks, and public conformance profiles rather than bespoke bilateral plumbing, which reduces diplomatic and technical friction for cross-border services (Sporny, Longley, & Chadwick, 2022; European Commission, 2024).

There are limits. National-security and criminal-procedure exemptions can dilute privacy by default if not time-boxed and co-approved. The framework counters by constraining overrides: narrow scope, short validity, two-person rule, automatic regulator alerting, and post-hoc notice to the person where lawful and safe. Purpose creep can also arrive via procurement or policy, not engineering. Making the purpose catalogue public, tying funding to minimal-disclosure ratios, and publishing reject/rollback statistics for proposed expansions help hold the line (European Commission, 2024; Allen & Hess, 2022).

Adoption, trust, and digital literacy

Trust rises when convenience and control increase together. The wallet gives people obvious control at the moment of disclosure, while the citizen portal gives persistent visibility who accessed what, when, and why which Estonia and municipal SSI pilots suggest is as important to public confidence as cryptography itself (Dunphy & Petitcolas, 2020). On the verifier side, adoption accelerates when the minimal path is the easiest path: SDKs and UI defaults ask for predicates or yes/no checks first; attempts to request dossiers trigger clear warnings about policy and legal risk; the help text shows how to reach the same decision with less data. Digital identity programmes fail when they assume that everyone has the same device, bandwidth, or ability. The framework therefore treats inclusion as a system property. Assisted channels at government counters can initiate wallet proofs with staff support; delegated consent allows parents or carers to act for dependents with time-boxed authority; fallback tokens or smartcards provide a non-smartphone option; and recovery is safe and simple so loss of a device does not translate into loss of identity (Dunphy & Petitcolas, 2020). These measures prevent the classic trade-off in which security and privacy are purchased at the cost of participation by those least served by digital services. Adoption inside government requires habits as much as APIs. Front-line and back-office teams should receive scenario-based training that reframes requests: “What decision do you actually need?” and “What is the minimal proof?” Managers should review minimal-disclosure ratios and denial reasons in regular ops meetings, the way they already review throughput and error rates. Change agents in large agencies can run before–after pilots on high-volume use cases licence checks, residency confirmation so staff see that minimal checks are faster, safer, and easier than dossier pulls. Public communication should focus on two simple ideas: *you choose what to disclose when present*, and *you can see every access later*.

Finally, adoption by the private sector banks, telcos, mobility, education depends on predictable profiles. If a verifier can count on one predicate for "over 18," one for "name-account match," one for "licence valid," integration and compliance shrink from months to weeks. Regulators can help by letting industry codes of practice reference the same predicates and purpose names, reducing audit complexity across sectors (European Commission, 2024; Arner, Barberis, & Buckley, 2020).

Economic and operational implications

Economically, SSI–VWR shifts spend from broad data integration to small, reusable decision interfaces. Narrow predicates and yes/no checks reduce bespoke ETL, data mapping, and privacy impact assessments because the interface surface is tiny and stable. Programmes see savings in three lines: (i) lower integration costs per verifier, since the same handful of predicates serves many services; (ii) reduced breach impact and legal exposure, because fewer attributes sit in verifier systems and audit facts are readily available for defence; and (iii) shorter handling times for routine cases, because a yes/no or a wallet predicate is faster than retrieving and inspecting a dossier (UIDAI, 2025; Allen & Hess, 2022). Operationally, success depends on SRE-style discipline for identity: clear SLOs (p95 latency, audit anchoring delay, revocation freshness), error budgets, and rapid rollback for policy changes that increase denials or lower MDR. Capacity planning focuses on revocation distribution, audit throughput, and wallet-verification peaks; these are easier to scale than general-purpose data exchanges because the payloads are small and cacheable. The regulator needs staffing and tools for continuous supervision: anomaly triage, override review, and transparency reporting. Agencies need a new role: policy engineers who maintain the purpose catalogue and the policy-as-code with legal and security input (Allen & Hess, 2022; European Commission, 2024). Procurement and vendor strategy determine whether the economic benefits persist. Lock-in comes from proprietary proof types and opaque decision engines. The framework's reliance on open standards VC/DID, EUDI wallet profiles and public conformance suites keeps switching costs low and encourages a competitive market for wallets and verifiers (Sporny, Longley, & Chadwick, 2022). Contracts should reward minimal-disclosure behaviour and require MDR and audit metrics as deliverables, not just uptime. Where agencies outsource operations, service credits should be tied to privacy KPIs (for example, revocation freshness and MDR) as well as availability.

There are real costs. Building immutable audit with multi-party anchoring, accessible citizen portals, and federated anomaly detection requires upfront investment and regulator capacity. The mitigation is the phased rollout: wrap existing interfaces with audit first (immediate value for oversight), convert the top 20 use cases to yes/no next (quick privacy and performance wins), then roll out wallet proofs to citizen-facing services (largest trust dividend) before finally retiring bulk interfaces. Each phase has measurable gates MDR thresholds, latency and anchoring SLOs, override governance maturity so funding and confidence grow with evidence (Dunphy & Petitcolas, 2020; European Commission, 2024). In sum, the economics favour a design that verifies without revealing. Agencies spend less integrating and defending sprawling data flows and more on small, secure decisions that are easy to test and govern. Citizens gain speed and control; regulators gain a truthful picture of system behaviour; and governments reduce risk while meeting European legal duties with room to adapt as standards and post-quantum cryptography evolve (European Commission, 2024; Sporny, Longley, & Chadwick, 2022; Juels & Oprea, 2020).

Conclusion and Future Work

This article has argued that national identity can be both fast and private when Self-Sovereign Identity (SSI) is reinforced by a Verify-Without-Reveal (VWR) enforcement spine. The core shift is from dossiers to decisions. Holder-present interactions rely on wallet-based selective disclosure and zero-knowledge proofs, so verifiers learn only what is necessary to act typically a predicate such as “over 18” or “licence valid” without exposing birth dates, addresses, or static identifiers (Camenisch & Lehmann, 2021; Sporny, Longley, & Chadwick, 2022). Holder-absent interactions replace broad queries with purpose-bound yes/no checks that return truth values or short-lived tokens and nothing more. Every access, in both modes, is evaluated by zero-trust policy encoded in a public purpose catalogue and recorded as an immutable event whose integrity is anchored by a permissioned ledger, giving regulators and citizens verifiable visibility without placing personal data on chain (Allen & Hess, 2022; Juels & Oprea, 2020; Vukolić, 2021). Together these elements suppress unconsented disclosure, deter cross-agency “surfing,” reduce linkability through per-agency pseudonyms and metadata hygiene, and keep latency low enough for national workloads. They also transform legal duties into default behaviour: data minimisation and purpose limitation become the ordinary outputs of interfaces; consent and lawful overrides become auditable artefacts; and high-risk AI in biometrics and anomaly detection is governed through testing, documentation, drift monitoring, and human oversight in line with the EU AI Act (European Commission, 2024; Paredes-García et al., 2023).

Evidence from leading programmes shows this approach is not speculative. Estonia demonstrates that separated registries, full-stack logging, and citizen visibility build trust at scale; EU wallet pilots show cross-border selective disclosure works in practice; Aadhaar shows yes/no checks can run at population scale when consented e-KYC handles exceptions; and municipal SSI pilots prove that people accept wallet-based proofs when prompts are clear and recovery is practical (Dunphy & Petitcolas, 2020; European Commission, 2024; UIDAI, 2025). The framework systematises these lessons: policy-as-code turns purpose rules into runtime checks; per-agency pseudonyms cut correlation risk; and ledger-anchored audit gives supervisors, regulators, and citizens a truthful record of access. Migration is deliberately phased audit everywhere, yes/no by default, wallet-first proofs, then legacy retirement so services keep running while privacy improves. Success is measurable through minimal-disclosure ratio, revocation freshness, latency SLOs, override governance maturity, fairness metrics for high-risk AI, and public transparency reports (Allen & Hess, 2022; European Commission, 2024).

Future work falls into five intertwined tracks. First, cryptography and performance need continued maturation: predicate proofs must verify faster on low-end devices; revocation and status mechanisms should deliver “fresh enough” checks with less bandwidth; and programmes should prepare for post-quantum migration with hybrid signatures and integrity proofs that remain verifiable for decades (Sporny et al., 2022; Juels & Oprea, 2020). Second, interoperability requires disciplined profiles and rigorous conformance testing. National and EU-level teams should publish versioned implementation profiles for common predicates (age-over, licence-valid, residency), maintain LoA cross-walks (eIDAS ↔ NIST), and run public plug-fests so wallets, verifiers, and issuers remain compatible across borders and sectors (NIST, 2020; European Commission, 2024). Third, governance must become as operational as the code: formal verification or model checking of ABAC policies for purpose mapping; regulator-operated integrity nodes with live anomaly triage; sanction frameworks that make misuse costly; and transparency reports that include not only outcomes but also governance behaviour (policy changes rejected, rollbacks, override reviews) (Allen & Hess, 2022; Vukolić, 2021). Fourth, privacy-preserving AI needs field validation beyond lab studies: safe federated learning with poisoning resistance, enclave or encrypted inference where latency permits, bias testing using diverse device and context conditions, and human appeal routes that people can actually use (Paredes-García et al., 2023; Ren et al., 2025; Kaul, 2021). Fifth, inclusion and economics merit dedicated study: comparative trials on assisted channels and delegated consent; usability research for recovery without smartphones; cost–benefit analyses that link minimal-disclosure ratio to lower breach impact, faster handling times, and fewer legal disputes; and procurement templates that tie vendor payments to privacy KPIs as well as uptime (Dunphy & Petitcolas, 2020; Allen & Hess, 2022).

Two open questions deserve special attention. The first is linkability measurement. While per-agency pseudonyms and minimal outputs reduce correlation in theory, programmes should run regular, privacy-safe linkage studies on audit telemetry to quantify real-world risk and tune token lifetimes, pseudonym rotation schedules, and metadata minimisation accordingly (Troncoso et al., 2020; Wagner & Eckhoff, 2020). The second is proportionality in lawful overrides. Time-boxing, two-person approval, regulator alerts, and citizen notice are strong safeguards, but their calibration how long, how narrow, how often is a policy choice that should be informed by data from audits and public consultation to maintain legitimacy (European Commission, 2024; Dunphy & Petitcolas, 2020). Addressing these questions will help ensure the framework remains both effective and trusted as it expands.

In closing, SSI provides the means to prove facts without revealing dossiers; VWR supplies the discipline that makes minimal disclosure, purpose checks, and immutable accountability routine also when the person is not present. When combined with transparent governance and inclusion by design, the result is a national identity fabric that people can use with confidence, regulators can supervise with evidence, and engineers can operate at scale. The practical task ahead is straightforward: implement the phased plan, measure the right things, publish what you learn, and iterate. Done well, verify-without-reveal becomes the normal way identity is used in government and regulated markets stronger assurance with less exposure, backed by proofs the public can see (European Commission, 2024; Allen & Hess, 2022; Dunphy & Petitcolas, 2020).

Conflict of Interest: The author reported no conflict of interest.

Data Availability: All data are included in the content of the paper.

Funding Statement: The author did not obtain any funding for this research.

References:

1. Allen, J., & Hess, K. (2022). Purpose limitation as policy-in-code: Attribute-based access control for public-sector APIs. Journal of Identity & Access Management, 6(2), 45–63.

2. Arner, D. W., Barberis, J. N., & Buckley, R. P. (2020). The identity challenge in finance: From KYC to digital identity. Asia Pacific Law Review, 28(2), 257–275.
3. Camenisch, J., & Lehmann, A. (2021). Privacy-preserving attribute-based credentials and zero-knowledge proofs: A survey and outlook. Foundations and Trends® in Privacy and Security, 4(1), 1–104.
4. Campbell, B., & Weitzner, D. J. (2022). Policy as code for data protection: Operationalising legal rules in information systems. IEEE Security & Privacy, 20(6), 39–50.
5. Dunphy, P., & Petitcolas, F. A. P. (2020). Decentralized digital identity and verifiable credentials: The basics and beyond. IEEE Security & Privacy, 18(4), 16–27.
6. European Commission. (2024). European Digital Identity Framework (eIDAS 2.0) and European Digital Identity Wallet Regulation and implementing measures. Publications Office of the European Union.
7. European Commission. (2024). Artificial Intelligence Act Rules on AI systems and obligations for high-risk use cases. Publications Office of the European Union.
8. Gencer, A. E., & Basu, S. (2021). Tamper-evident audit trails with permissioned ledgers: Design patterns and pitfalls. ACM Queue, 19(5), 1–21.
9. ISO/IEC. (2021). ISO/IEC 18013-5:2021 Personal identification ISO-compliant driving licence Part 5: Mobile driving licence (mDL) application. International Organization for Standardization.
10. Juels, A., & Oprea, A. (2020). New directions in tamper-evident logging. Communications of the ACM, 63(4), 38–47.
11. Kaul, S. (2021). Encrypted inference in practice: Trusted execution and homomorphic approaches for privacy-preserving analytics. IEEE Computer, 54(12), 30–41.
12. Khovratovich, D., & Law, J. (2020). BBS+ signatures Unlinkable selective disclosure in practice. IACR Cryptology ePrint Archive, 2020/1416.
13. Meylan, C., & Sabadello, M. (2021). Decentralized identifiers and verifiable credentials for e-government services. In A. De Santis (Ed.), Digital Identity Management (pp. 121–146). Springer.
14. NIST. (2020). Digital Identity Guidelines (SP 800-63 Rev. 3, including 2020 updates). National Institute of Standards and Technology.
15. Paredes-García, W., Marcel, S., Galbally, J., & Fierrez, J. (2023). Face anti-spoofing and liveness detection: A comprehensive survey. IEEE Transactions on Information Forensics and Security, 18, 1234–1267.

16. Preukschat, A., & Reed, D. (2021). Self-Sovereign Identity: Decentralized digital identity and verifiable credentials. Manning.
17. Ren, X., Li, T., Kairouz, P., & McMahan, H. B. (2025). Privacy-preserving federated learning at scale: Systems, robustness, and governance. Foundations and Trends® in Machine Learning, 18(1), 1–189.
18. Sporny, M., Longley, D., & Chadwick, D. W. (2022). Verifiable Credentials Data Model 2.0. W3C Recommendation.
19. Sporny, M., Burnett, D., & Sabadello, M. (2022). Decentralized Identifiers (DID) Core. W3C Recommendation.
20. Troncoso, C., Isaakidis, M., Danezis, G., & Halpin, H. (2020). Systematizing decentralization and privacy: Lessons from cryptographic designs. Proceedings on Privacy Enhancing Technologies (PoPETs), 2020(4), 307–329.
21. UIDAI. (2025). Aadhaar Authentication and e-KYC Specifications and best practices (Ver. 3.x). Unique Identification Authority of India.
22. Vukolić, M. (2021). Permissioned blockchains: Design goals, consensus choices, and performance trade-offs. Communications of the ACM, 64(12), 38–45.
23. Wagner, I., & Eckhoff, D. (2020). Technical privacy metrics: A systematic survey. ACM Computing Surveys, 54(1), 1–38.
24. Zhou, Y., Xu, Y., & Lyu, M. R. (2021). Byzantine-robust federated learning: A comprehensive survey. IEEE Transactions on Big Data, 7(1), 1–20.