ESJ Social Sciences

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/esj.2026.v22n1p1

Submitted: 05 November 2025
Accepted: 20 January 2026
Published: 31 January 2026

Cite As:
Mansur, M.A. (2026). *Self-Sovereign Identity Architecture for National Use with Wallet Proofs Zero-Knowledge and the VWR Framework.* European Scientific Journal, ESJ, 22 (1), 1. https://doi.org/10.19044/esj.2026.v22n1p1

Abstract

National identity systems require efficient, equitable decision-making that safeguards personal data. This article proposes a Self-Sovereign Identity (SSI) architecture, supported by a Verify-Without-Reveal (VWR) framework, designed for national-scale implementation. 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. The study outlines reference APIs, user experiences for wallets and verifiers, and performance metrics suited for 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

1. 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. Current platforms often retrieve complete records when a simple binary verification is sufficient. 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). The fundamental technical challenge is enabling verifiers to determine eligibility without accessing comprehensive personal dossiers. 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. For holder-absent flows, VWR provides secure APIs that return binary responses, enforce zero-trust policies, and generate tamper-evident audit logs. 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).

Contributions and significance (EU and global)

This article presents four primary contributions, beginning with the integration of SSI and VWR into a unified national architecture. It demonstrates how wallet-based selective disclosure and zero-knowledge proofs work alongside binary APIs, zero-trust policies, and immutable auditing. The framework ensures consistent operational behavior for both holder-present and holder-absent interactions. 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 minimization 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).

2. Background

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). 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).

3. Research Methodology

This study investigates how a self-sovereign identity (SSI) model, reinforced by a verify-without-reveal (VWR) framework, can meet national requirements for speed, legality, and trust. The central question examines deciding eligibility without exposing records. Subsidiary questions address composing DIDs and zero-knowledge proofs to return predicate answers rather than dossiers (Camenisch & Lehmann, 2021; Sporny et al., 2022); implementing policy controls and zero-trust enforcement to prevent cross-agency browsing (Allen & Hess, 2022; Zhang & Li, 2020); and enhancing assurance through privacy-preserving fraud scoring and encrypted inference (Paredes-García et al., 2023; Kaul, 2021; Ren et al., 2025). Compliance analysis assesses alignment with GDPR, eIDAS 2.0, and the EU AI Act (European Commission, 2024; NIST, 2020). The research design blends conceptual synthesis with comparative case analysis. The synthesis integrates cryptographic elements—DIDs, verifiable credentials, and permissioned ledgers—into a scalable national architecture. This is tested against active programs: Estonia's integrity anchoring, EU wallet pilots on selective disclosure, India's Aadhaar on population-scale authentication, and municipal pilots like Zug. These cases provide a yardstick for security, interoperability, and performance (Dunphy & Petitcolas, 2020; European Commission, 2024; UIDAI, 2025). Sources include peer-reviewed literature (2020 onwards) and official standards (eIDAS, W3C, NIST), prioritizing technical specificity over marketing materials (Sporny, Longley, & Chadwick, 2022; European Commission, 2024; NIST, 2020). Data extraction focuses on seven recurring elements, including binding, revocation, verification paths, and policy layers. Architectural choices and governance controls are coded under minimal disclosure, zero-trust enforcement, and privacy-preserving AI to expose trade-offs in the proposed design (Dunphy & Petitcolas, 2020; Camenisch & Lehmann, 2021; Allen & Hess, 2022). Evaluation uses five criteria: Security (resistance to abuse), Privacy (selective disclosure and unlinking), Interoperability (cross-border proofing), Governance (audit visibility), and Scalability (latency and load reliability) (European Commission, 2024; UIDAI, 2025). Validity relies on triangulation between standards and program documentation, prioritizing empirical statistics where available. Reliability is ensured by a consistent coding framework. Bias is managed by weighing biometric claims against dataset diversity and treating auditability as essential for detecting misuse (European Commission, 2024; Paredes-García et al., 2023). Ethical considerations are integral, mapping recommendations to GDPR and AI Act duties. The design prioritizes user control and inclusion, favouring holder-present proofs on commodity devices while proposing low-tech alternatives for equal access (Dunphy & Petitcolas, 2020; European Commission, 2024). Ultimately, this methodology positions SSI-VWR as the enforcement spine for all interactions, evaluating success by data minimization and audit transparency (European Commission, 2024; Camenisch & Lehmann, 2021; Dunphy & Petitcolas, 2020; UIDAI, 2025).

4. Problem Analysis

National identity systems face a cluster of reinforcing problems: unnecessary disclosure, cross-agency "surfing," linkability, and biometric risk. These issues stem from defaults that favour dossiers over decisions and trust networks rather than purposes. An SSI-centred design, reinforced by a verify-without-reveal (VWR) spine, must therefore change 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" Legacy platforms often fetch full records for simple checks, creating legal risk and enabling insider surfing where staff browse records without friction. GDPR purpose limitation is honoured in policy but not in code (Zhang & Li, 2020; Campbell & Weitzner, 2022). SSI-VWR changes this posture by making minimal answers the norm. A purpose-bound API returns predicates (e.g., "licenceValid = true") rather than dossiers (European Commission, 2024; UIDAI, 2025). An Attribute-Based Access Control (ABAC) engine ties each call to a legal basis, while immutable audits make surfing costly by leaving a tamper-evident trail visible to regulators and citizens (Allen & Hess, 2022; Juels & Oprea, 2020; Vukolić, 2021). Linkability, metadata leakage, and dossier creep Linkability undermines privacy when global identifiers or metadata fingerprints allow tracking across domains (Troncoso et al., 2020; Camenisch & Lehmann, 2021; Wagner & Eckhoff, 2020). The proposed design mitigates this by replacing global IDs with per-agency pseudonyms derived from the national identifier and an agency salt. The verification layer strips non-essential metadata and rotates short-lived tokens. This prevents dossier creep, as verifiers cannot justify keeping more data than received for a narrow, logged purpose (Camenisch & Lehmann, 2021; Gencer & Basu, 2021; European Commission, 2024; Sporny et al., 2022; Dunphy & Petitcolas, 2020). Biometric risks, bias, and model governance Biometrics introduce risks of replay attacks, deepfakes, and algorithmic bias (Paredes-García et al., 2023; European Commission, 2024). A robust approach separates usage: capture occurs in controlled conditions, while routine liveness checks run on-device or in secure enclaves (NIST, 2020; Allen & Hess, 2022). Assurance is enhanced via federated learning and encrypted inference, which detect misuse without pooling sensitive data (Kaul, 2021; Ren et al., 2025; Zhou et al., 2021). Continuous governance monitors error rates across demographics and mandates human appeal routes, turning AI Act compliance into daily practice (European Commission, 2024; Paredes-García et al., 2023).

Comparison of Traditional Identity Systems vs. SSI–VWR Framework

Feature	Traditional Identity Systems	Proposed SSI–VWR Framework
Data Disclosure	Defaults to retrieving complete personal dossiers for simple verification.	Provides "decisions over dossiers" through yes/no APIs and zero-knowledge predicates.
Trust Model	Relies on implicit trust within internal agency networks.	Enforces a zero-trust architecture where every request is independently verified.
Access Control	Uses broad Role-Based Access Control (RBAC), often allowing internal "surfing".	Implements Attribute-Based Access Control (ABAC) tied to a mandatory purpose catalogue.
Accountability	Logs are often mutable, editable by insiders, or inaccessible to citizens.	Features immutable, append-only audit trails anchored to a permissioned ledger.
User Privacy	High traceability due to the reuse of stable, global identifiers across domains.	Suppresses traceability using per-agency pseudonyms and metadata hygiene.
User Visibility	Citizens have little to no visibility into who accessed their records and why.	Offers full transparency via a citizen portal showing every access event in plain language.
Biometrics	Often involves centralized storage of sensitive biometric templates.	Prioritizes on-device template storage and privacy-preserving liveness checks.

5. 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).

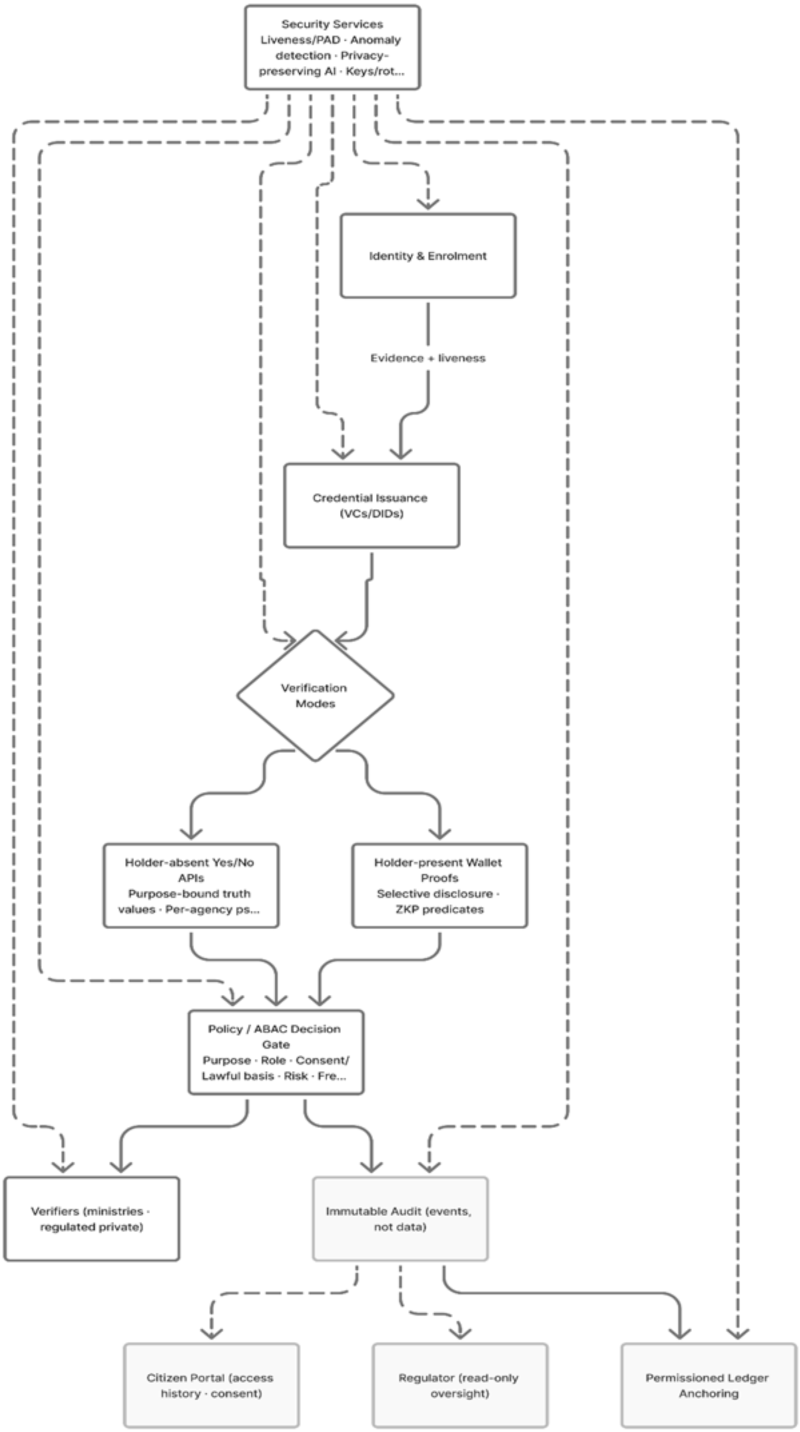
Security Services
Liveness/PAD · Anomaly detection · Privacy-preserving AI · Keys/rot...
Identity & Enrolment
Evidence + liveness
Credential Issuance (VCs/DIDs)
Verification Modes
Holder-absent Yes/No APIs
Purpose-bound truth values · Per-agency ps...
Holder-present Wallet Proofs
Selective disclosure · 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: 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

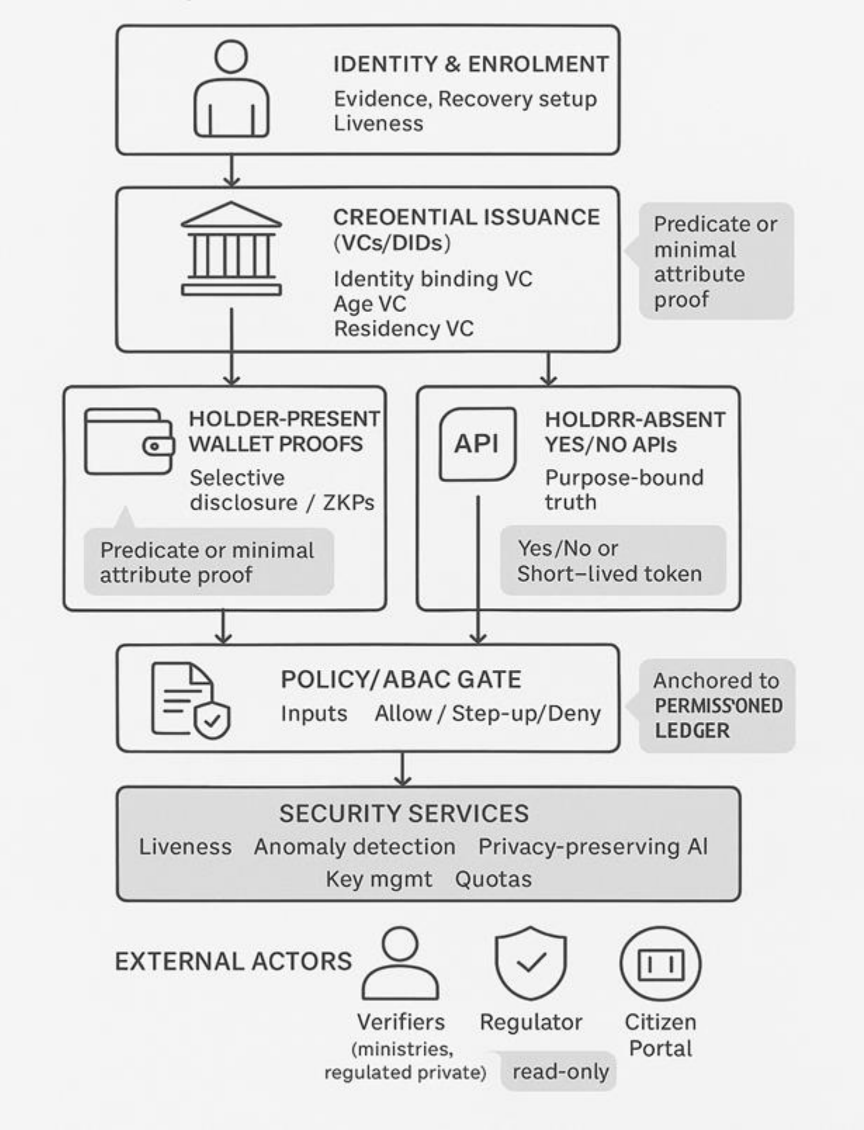
IDENTITY & ENROLMENT
Evidence, Recovery setup
Liveness
CREOENTIAL ISSUANCE (VCs/DIDs)
Identity binding VC
Age VC
Residency VC
Predicate or minimal attribute proof
HOLDER-PRESENT WALLET PROOFS
Selective disclosure / ZKPs
Predicate or minimal attribute proof
API
HOLDRR-ABSENT YES/NO APIs
Purpose-bound 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:
   i. Holder-present (SSI) when the person is on site or online with their wallet.
   ii. 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

Enrolment: strong evidence capture, liveness, issuance of initial credentials, and setup of recovery paths.

Routine use: predicates and yes/no checks by default; consent prompts for richer cases; step-up for risk.

Oversight: immutable audit, regulator analytics, transparency reports; citizen visibility through a portal.

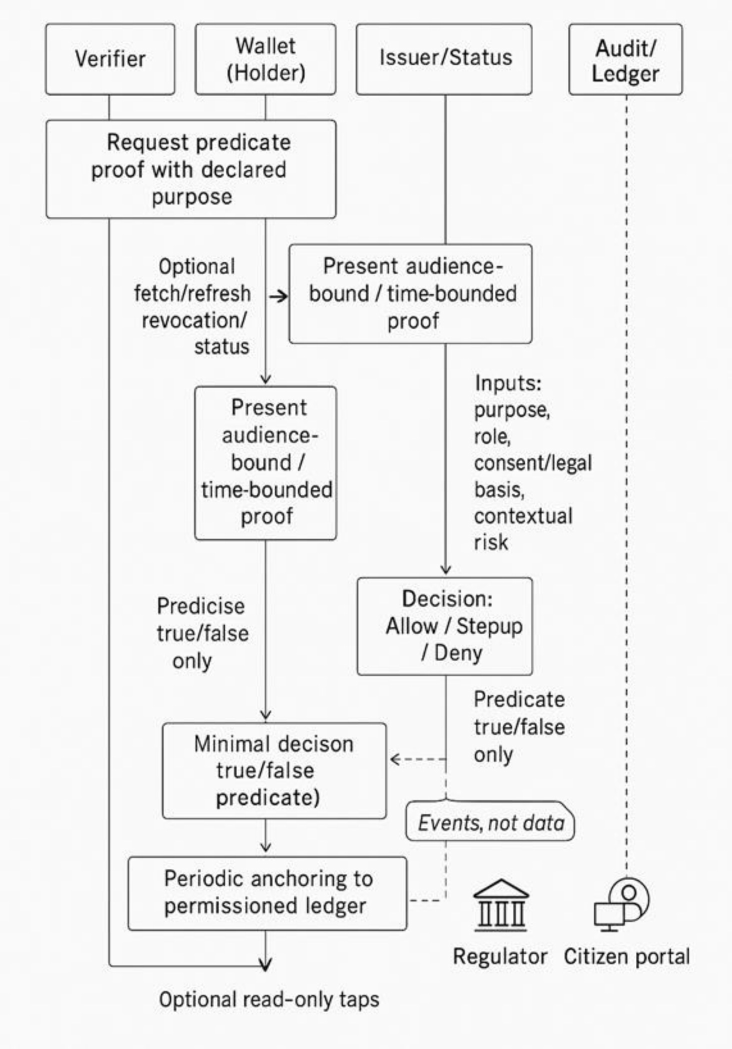
Verifier
Wallet (Holder)
Issuer/Status
Audit/ Ledger
Request predicate proof with declared purpose
Optional fetch/refresh revocation/ status
Present audience-bound / time-bounded proof
Present audience-bound / time-bounded proof
Inputs: purpose, role, consent/legal basis, contextual risk
Predicise true/false 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: 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).

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. 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).

6. 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).

7. 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).

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.

8. Evaluation and Risk

Implementation Limitations While the SSI-VWR framework provides a robust technical blueprint, its successful deployment is subject to several practical constraints. Institutional readiness poses a primary challenge, as transitioning from legacy "dossier-based" systems to a "policy-as-code" environment requires significant upgrades to administrative capacity and the recruitment of specialized personnel, such as policy engineers, to manage the purpose catalogue. User adoption also remains contingent on digital literacy; the framework must ensure that recovery paths and delegated consent models are intuitive enough for those without high-end smartphones or technical expertise. 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). 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).

Conclusion

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). 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.