LIGHTPAPER

OTCP

Global Trust Infrastructure for People and AI Agents

A standards-compatible layer for the cross-border, cross-organization use of existing identity schemes, and for the verifiable delegation of authority to software agents.

v3.0 - 2026-08-22

Conformant by design.
Not yet certified.

Contents

Abstract

Digital identity is moving from repeated account proofing toward wallet-mediated, cryptographically verifiable credentials, and the move is regulatory as much as technical. National and regional programmes are maturing in parallel: mobile driving licences built on ISO 18013-5, bank-operated identity schemes, open-source foundational identity platforms deployed across multiple countries, machine-readable travel credentials, and the European Digital Identity Wallet framework established by Regulation (EU) 2024/1183, with implementation deadlines in 2026–2027. These programmes converge on the same cryptographic building blocks and diverge on identifiers, trust anchors, certification paths, legal roles, and supervisory expectations.

That divergence is not a defect to be centralized away. Identity is an attribute of sovereignty and will remain plural. What is missing is the layer above the schemes: a neutral, compliance-grade means for a verifier in one jurisdiction to put a narrow policy question to evidence rooted in another, obtain a defensible answer, retain only what the law requires, and demonstrate afterwards what was checked.

A second requirement has emerged more abruptly. AI agents now act — filling carts, booking travel, moving money, signing documents. The interoperability layer for agents matured quickly, but its protocols standardize how agents communicate rather than whose authority they carry. No public identity framework yet standardizes agent delegation; the European framework leaves representation and mandate attestations pending.

OTCP addresses both requirements in one layer, and introduces no credential format to do so. It assembles SD-JWT VC and ISO mdoc credentials, OpenID4VCI and OpenID4VP, the High Assurance Interoperability Profile, and the IETF Token Status List; resolves trust against plural federated anchors rather than a registry of its own; and adds the policy, minimization, and audit functions a regulated verifier requires. Above this sits the element no existing scheme covers: the **VDAC constitution** — seven principles and seven invariants that root agent authority in a KYC/KYB-verified principal, express it as typed risk classes, attenuate it at every hop, bind sensitive presentations to the exact action, and permit identity recovery only through a gated, due-process path.

The architecture is settled and has been verified end to end on a physical device. The remaining work is conformance testing, certification under each applicable regime, and the open problems set out in Section 13.

Standards-maturity legend. External building blocks carry a maturity tag: `standard` (final regulation, law, RFC, ISO standard, or final OpenID specification) · `profile` (an official or ecosystem profile constraining standards for a legal or deployment context) · `draft` (active or non-final specification) · `implementation-specific` (a national system, registry, or product integration) · `product` (an OTCP design or policy choice).

---

1. Introduction

Three developments have matured concurrently, and they cannot be addressed separately.

Verifiable identity has reached a usable standards base. Selective-disclosure formats, issuance and presentation protocols, and revocation mechanisms are now final specifications with independent implementations. The question is no longer whether credentials work, but how they are governed, assured, and made interoperable.

Public identity programmes are being deployed in parallel across jurisdictions — mobile driving licences on ISO 18013-5 `standard`, Ukraine's Diia, NBU BankID, and electronic trust services `implementation-specific`, Georgia's eID and qualified trust services `implementation-specific`, open-source foundational

identity platforms in service across several countries `implementation-specific`, ICAO's digital travel credential work `draft`, and the European Digital Identity Wallet framework `profile`. They share cryptography and differ in law.

Software agents have begun to act. Systems that plan and execute multi-step tasks now transact with counterparties selected at runtime. Tool-calling and agent-to-agent protocols standardized quickly `draft` and address how agents communicate rather than whose authority they carry.

Each development is unsafe without the other two. An agent that can spend but cannot prove whose authority it carries is a liability; an identity that cannot be delegated is of no use to an agent; a credential scheme that stops at a border returns every cross-border verifier to document collection. The layer that connects them must be neutral, because no single jurisdiction or vendor can serve as the trust root for all the others.

---

2. The Trust Gap

2.1. Structural Failures of Current Practice

The prevailing account-and-upload pattern produces three failures. **Users disclose more than the transaction requires**, because minimization is treated as policy rather than as an architectural property. **Verifiers duplicate assurance work** — each rebuilding onboarding, screening, document review, status refresh, consent capture, and audit — which yields inconsistent assurance, high cost, and weak interoperability rather than better trust. And **the rails are fragmented, and will remain so**: official wallets, bank identity schemes, mobile driving licences, and private ecosystems share no identifiers, formats, certification paths, or trust anchors, so a cross-border verifier must put a single policy question to several rails without treating those rails as equivalent.

A wallet addresses part of this, but does not remove the need for a **trust and compliance layer**. A verifier must still establish whether the issuer is authorized for the claim in question, whether the credential is current, whether the presentation is bound to the holder, whether the relying party has a lawful purpose, what retention applies, and what audit evidence to retain.

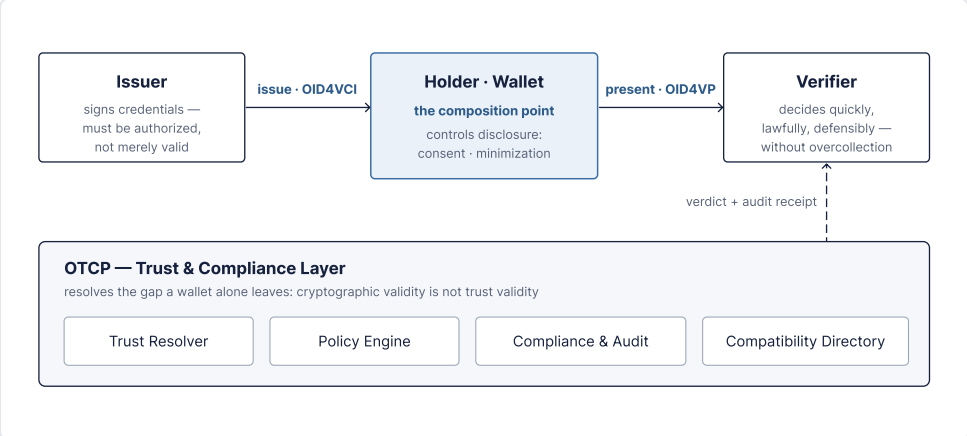

Figure 1. Actors and requirements. The wallet is the holder's composition point; OTCP resolves what a wallet alone cannot — issuer authorization, status, policy, retention, and audit — and returns a verdict with a minimal audit receipt.

2.2. Actors and Requirements

Three parties operate under distinct constraints (Figure 1). **The verifier** must decide quickly, lawfully, and defensibly, without overcollecting and without converting every check into a permanent data store; OTCP presents it with a small surface — a request carrying purpose, claims, assurance level, retention class, and jurisdiction context — and returns a structured verdict and a minimal audit receipt. **The holder** requires control that is visible in the product: the identity of the requester, the purpose, what is required as against optional, whether a predicate can substitute for a raw value, and what will be stored and for how long. **The issuer** requires a means of publishing credentials that verifiers can trust, which is more than a valid signature. A bank can sign a KYC-status credential, but the verifier must still establish that the issuer is authorized for that claim. This is the distinction between **cryptographic validity and trust validity**, and resolving it is the function of trust resolution against federated registries.

2.3. Delegated Authority for Software Agents

The term "agent identity" covers five distinct concerns: naming and discovery, agent attestation, **delegation of authority**, per-call authorization, and payment authority. Development has concentrated at the top and bottom of that stack and left the middle thin. Verifiable delegation of a principal's authority is where the unsolved problems are concentrated, and its root cause is the same as that of the human trust gap: there is no neutral means of proving that an accountable principal stands behind an action while disclosing only what the action requires.

---

3. Scope and Positioning

OTCP is a composition layer rather than a competing scheme. Three commitments define its scope.

It assembles existing standards rather than introducing new ones. OTCP defines no credential format, no issuance or presentation protocol, and no transport. Its components are drawn from specifications that already exist and are converging. Its contribution is the assembly: the profile, the policy layer, the compliance evidence model, and the constitution governing delegated authority.

It facilitates existing schemes rather than replacing them. Sovereign identity systems remain sovereign. OTCP integrates with official rails only through authorized routes, and cannot become an official issuer through technical design. Its value to a national scheme is extended reach: a credential issued under one jurisdiction's rules becomes usable, under jurisdiction-appropriate policy, by a verifier that would otherwise revert to collecting documents.

It is horizontal, and designed for cross-jurisdictional use. The same rails serve banking and financial services (reusable KYC, account opening, business onboarding), eGovernment and public services (citizen verification, eligibility, cross-border recognition), regulated commerce, and age-restricted access. Age assurance is the clearest initial application — proving a threshold without disclosing a birth date combines current regulatory pressure, high-volume onboarding, and a direct demonstration of predicate proofs — but it is also the substrate: every sector reuses the same rails, and a jurisdiction is expressed as a profile rather than a fork.

OTCP is **not** an identity provider or state issuer; not a payment processor or settlement layer, holding no funds, taking no share of payment value, and initiating no settlement; not a gatekeeper, operating no proprietary trust list that participants must join; and not a certified wallet, qualified trust service, or accredited body unless and until such a role is formally obtained.

Implementation status. The reference wallet is a native Android application (Kotlin and Jetpack Compose), verified on device through onboarding, issuance, presentation, and an integrated authenticator. The issuer, verifier, and status services are deployed and hosted, and the complete round trip — identity verification, issuance of three credentials, and a data-minimized age presentation verified server-side — has been demonstrated on a physical device. Conformance testing and certification are process steps under each applicable regime; they are not outstanding architectural work.

---

4. Design Principles and Standards Baseline

- **Compatibility takes precedence over novelty.** In a trust market, interoperability is the product. A credential only OTCP can read has no value, and a verifier that accepts only OTCP-issued credentials cannot serve a regulated buyer that must also accept government and bank credentials. Where a standard cannot express a requirement, the rule is to adopt a further standard rather than to define one.
- **Official identity remains sovereign.** Identity regimes, trust services, relying-party registration, and public-sector integrations are governed relationships. OTCP integrates where authorized, and treats alignment as a conformance and certification track rather than as an attained status.
- **Verification is not authorization to provide a service.** A presentation establishes that a signed claim exists, that a wallet controlled a key, and that a status check passed. It does not render a relying party's service lawful: the relying party retains the service decision, the legal basis, and sector policy.
- **Data minimization is a protocol and product property.** Every request carries purpose, required and optional claims, claim mode, retention class, and jurisdiction profile in machine-readable form, so that the wallet can display precisely what is requested and the holder can refuse optional claims or select a narrower mode.
- **Trust anchors are plural.** X.509 chains, qualified certificates and trusted lists, mobile-document issuer authorities, federation entity identifiers, web-origin controls, key thumbprints, and national registries that are not web-native. The resolver accommodates all of them without imposing a single model.
- **Claims are limited to what has been obtained.** Conformance by design, stated without asserting certification not yet held, is both stronger and safer than overstatement. The same standard is applied to the open problems in Section 13.

The standards baseline. SD-JWT VC `standard`, the primary credential format and one of the two formats mandated under the European framework; **ISO/IEC 18013-5 and -7 mdoc** `standard`, the second mandated format, used for proximity and online presentation — OTCP holds and presents both from the outset; **OpenID4VCI** and **OpenID4VP** `standard` for issuance and for presentation with a holder-signed key-binding proof; **HAIP** `profile`, the high-assurance interoperability profile; and the **IETF Token Status List** `standard` for revocation — a compressed, signed bit array (`statuslist+jwt`) published at a stable URL, in which each credential's index encodes valid or revoked. This is the mechanism used by the European reference implementation for SD-JWT VC, and the mechanism OTCP uses, rather than the W3C Bitstring variant.

Interoperability in practice. In implementation, the wallet's independently written SD-JWT presentation code and a third-party verifier library produced byte-identical salted-hash digests at first integration, without adjustment on either side. That outcome depends on both parties implementing the same specification rather than two dialects. Conformance is therefore not a constraint on delivery; it is the property that allows independently built components to interoperate, and it is the condition of entry to a market in which every serious buyer must also accept credentials OTCP did not issue.

---

5. Trust Resolution Across Jurisdictions

5.1. Federated Trust Registries

The central trust question is how a verifier establishes that an issuer is legitimate. OTCP **consumes federated trust registries and operates no trust list of its own.** It bridges to EU trusted lists (ETSI TS 119 612) `standard`, EBSI `implementation-specific`, OpenID Federation `standard`, and mdoc IACA roots `standard`, and resolves any issuer against them. A closed, OTCP-controlled list would constitute a gatekeeping function, and would make a single vendor a point of failure for a system whose purpose is to operate across plural regimes. Identity anchors are method-agnostic: issuers use X.509 chains with `did:web` dual-published against the same keys, while holders and agents use `did:jwk` keys held in the device secure element.

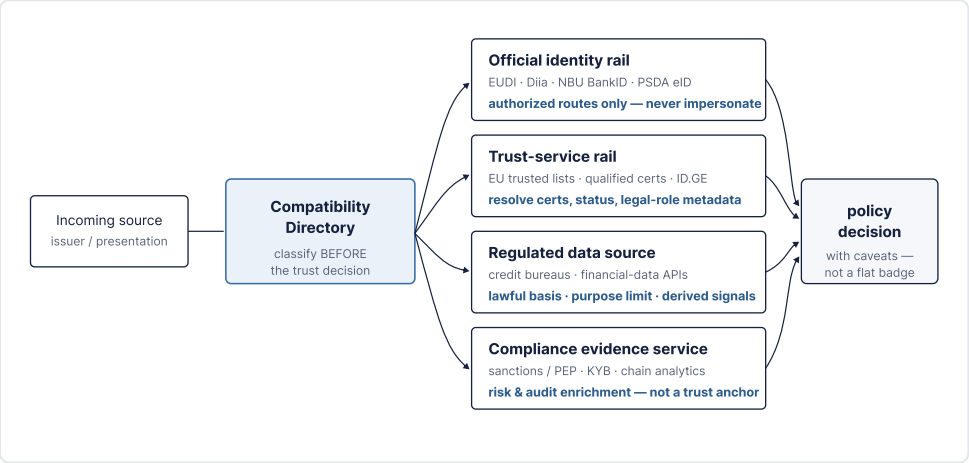

Figure 2. The four-role compatibility directory. A source is classified before any trust decision, so that an official identity rail, a trust service, a regulated data source, and a compliance evidence service are not treated as one generic category of issuer.

5.2. The Compatibility Directory

External sources are not of a single kind. Treating a government identity rail, a qualified trust service, a credit bureau, and a screening vendor as one generic category of issuer is a classification error that produces either unwarranted trust or unwarranted rejection. OTCP therefore classifies a source before any trust decision, into four connector roles (Figure 2):

Role	Examples	Treatment
Official identity rail	National wallets and eID schemes, bank identity schemes, mobile driving licences	Integrate only through authorized routes; no scraping, impersonation, or assertion of official status
Trust-service rail	Trusted lists, qualified certificates, mdoc issuer roots, national trust services	Resolve certificates, lists, status, and legal-role metadata; do not force into a DID-only model
Regulated data source	Credit bureaus, financial-data APIs, company registries	Use with lawful basis, consent, purpose limitation, and retention; prefer derived eligibility signals
Compliance evidence service	Sanctions and PEP screening, KYB, blockchain analytics	Risk and audit enrichment; not identity issuers and not public trust anchors

The output is a policy decision with stated caveats rather than a single trusted indicator. This is also what makes the model portable to a jurisdiction not previously encountered: a new rail is classified into one of four roles rather than accommodated as a special case.

5.3. Jurisdictional Profiles

Localization is expressed in **profiles**: which formats are accepted for a given check, which assurance metadata a verifier must request, which retention class applies, which trust anchors are authoritative, and which relying-party registration is required. Each profile records whether a given rule rests on a public standard, an official profile, a draft, a national implementation, or an OTCP product decision, so that a supervisor can audit the basis of every acceptance rule.

One constraint follows directly: **OTCP publishes no equivalence tables between jurisdictions' assurance levels.** Comparing levels of assurance across regimes is an unresolved governance problem (Section 13), and a table of equivalences would convert a legal judgment into a technical default. Acceptance rules are instead profile-specific and source-cited.

5.4. Key Material and Hardware Binding

The holder's signing key must be **non-exportable and hardware-rooted**: generated within the secure element — Android StrongBox or TEE, or the Apple Secure Enclave — never released from it, with every signature gated by a biometric prompt. This satisfies the sole-control requirement associated with the highest assurance level, expressed in the European framework as the Wallet Secure Cryptographic Device (WSCD).

That constraint determines a cryptographic choice that warrants explanation, because the strongest general-purpose option is not the one selected. The natural modern choice is EdDSA/Ed25519, but the Apple Secure Enclave supports only the NIST P-256 curve. The holder key is therefore **P-256 / ES256** `standard`, which is also the curve mandated by ISO 18013-5 and the most broadly supported across wallet profiles. P-256 is not selected on cryptographic merit; it is the only curve that satisfies hardware binding and interoperability simultaneously. Every credential declares its algorithm, so the system remains agile: EdDSA where hardware supports it, and post-quantum schemes (ML-DSA, SLH-DSA) `draft` subsequently, without a schema change.

---

6. Integration and Participation

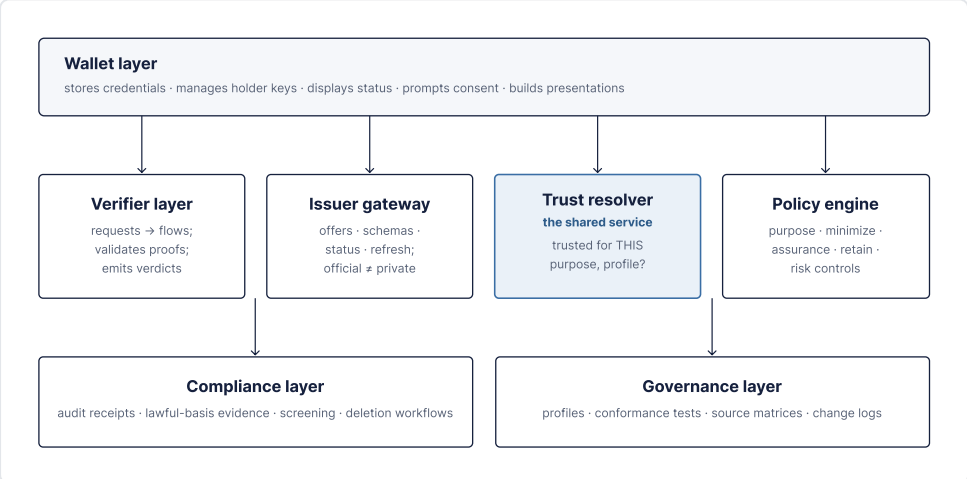

Figure 3. The component model. OTCP comprises interoperating services — wallet, verifier, issuer gateway, trust resolver, policy engine, compliance and governance layers — which a deployment instantiates and a participant integrates against, rather than a monolith adopted whole.

6.1. Participants and Integration Points

OTCP is designed to be integrated into existing systems rather than to displace them. The component model (Figure 3) comprises interoperating services — wallet layer, verifier layer, issuer gateway, trust resolver, policy engine, compliance layer, and governance layer — which any deployment instantiates. Externally, a participant encounters a small surface:

Participant	Integrates	Receives	Applicable vetting
Relying party / verifier — bank, operator, marketplace, public portal	Verifier API over plain HTTP, or a hosted OpenID4VP endpoint	A structured verdict and a minimal audit receipt	Relying-party registration and purpose declaration where the jurisdiction operates one; existing sector licensing unchanged
Issuer — bank, employer, university, licensing body, authorized agency	Issuer gateway (OpenID4VCI), schema and status publication, refresh	Credentials any conforming verifier can resolve and check	The authorization its jurisdiction requires for the claim it signs; OTCP resolves that status and does not confer it
Trust service provider	Certificate, trusted-list, and status bridges into the trust resolver	Machine-readable presence in verifiers' trust decisions	Its existing qualified or accredited status under its own regime
Wallet / platform provider	Holder SDK (native Android first), W3C Digital Credentials API `draft`, mdoc and SD-JWT libraries	Interoperation with OTCP issuers and verifiers, in both formats	Wallet certification where the regime requires it
Agent platform / enterprise	Agent-issuance and VDAC verification APIs; AP2 mandate export `draft`	Presentable, verifiable, revocable delegated authority	KYC/KYB of the principal; no separate agent registry
Regulator / supervisory body	Read surfaces: profiles, conformance tests, audit-receipt schemas, source matrices	Supervisable evidence without a central data store	Defines the applicable vetting rather than undergoing it

6.2. Protocol Surface

No element of the integration path is proprietary. A verifier creates a presentation request and receives a request URI; the wallet retrieves a signed authorization request object carrying a query in the standard claims-query language and returns the presentation over the standard direct-post channel; the verifier validates the issuer signature against published keys, the key-binding proof against the credential's confirmation key with nonce and audience, and the status against the Token Status List; the relying party reads a verdict. Issuance follows the OpenID4VCI pre-authorized-code flow. Revocation is a status-bit flip at a stable URL. Agent presentations add a verification call executing the eleven-step algorithm of Section 10.4.

External providers integrate behind **provider-agnostic connector interfaces** — identity verification, entity and wallet screening, travel-rule exchange, business verification and registry lookup, government digital-ID import, and settlement — each with normalized result shapes, so that substituting a provider is a configuration change rather than a rewrite. Risk oracles and data providers inform a verifier's decision; they do not constitute an identity gate operated by OTCP.

6.3. Authorization and Vetting Model

Participation rights are determined externally to OTCP, in three layers.

1. **Authorization to issue** is determined by each jurisdiction through its accreditation, notification, or qualified-status process. OTCP reads that determination from the federated registries and enforces it in policy. It does not confer issuing authority, and does not override a jurisdiction that withholds it.
2. **Authorization to request** is determined by the relying-party regime where one exists. Where a jurisdiction registers relying parties and binds them to declared purposes, OTCP carries that registration into the request, so that the wallet can display the requester and its authority and the policy engine can enforce purpose limitation rather than accept a self-declared string.
3. **Authorization to operate a new role** — an unmasking escrow, an agent-attestation service, a sectoral accreditor — is determined through extensible federation: OpenID Federation entity statements `standard` and accreditation credentials, under multi-stakeholder governance. Trust in a new role is itself expressed as a credential, issued by the body competent to issue it in that jurisdiction.

The resulting property is that **OTCP functions as a policy-enforcement and evidence layer rather than as an admissions authority.** An ecosystem usable across jurisdictions does not require a global gatekeeper; it requires each local gate to be machine-readable, and each verifier to be able to require the gates its own regulator recognizes.

---

7. Regulatory Conformance and Unlinkability

7.1. Conformance Posture

OTCP's conformance rule is to build to the strictest requirements imposed by any target regime, so that certification in a given jurisdiction is a process step rather than a re-architecture. The most fully specified of the current regimes is the European one, and its requirements are largely a superset of those imposed elsewhere; it therefore sets the engineering bar, while jurisdiction-specific obligations are carried in the profiles of Section 5.3.

Against that bar, the wallet is built as **conformant and certification-track on both the holder and verifier sides.** In concrete terms: both mandated credential formats held and presented; selective disclosure; **batch issuance for unlinkability**; secure-element keys under sole control (WSCD, corresponding to the high assurance level); an explicit level-of-assurance vocabulary; and issuer-unlinkability. Credentials are structured to be conformant with Qualified Electronic Attestations of Attributes. The corresponding restraint holds: formal certification and any qualified status are pursued under the applicable timetable, which is the gating factor rather than the architecture.

One asymmetry should be stated. Conformance applies in full to the human credential layer. **The agent layer cannot be described as conformant, because no public framework yet standardizes agent delegation** — in the European framework, representation and mandate attestations remain out of scope pending member-state input. Delegation credentials are therefore aligned in direction, designed to converge on representation attestations as those are standardized, rather than asserting a conformance that does not yet exist.

7.2. Unlinkability Requirements and Batch Issuance

Where a regime requires unlinkability — that two verifiers cannot determine by collusion that they encountered the same person — the requirement is binding rather than advisory. In the European framework it is set out in **Article 5a(16) of Regulation (EU) 2024/1183** `standard`: paragraph (a) prohibits attestation providers from obtaining data that permits transactions to be tracked, linked, or correlated; paragraph (b) requires the wallet to enable privacy-preserving techniques ensuring unlinkability where the attestation does not require identification of the user.

Three scoping facts determine its effect: it binds certified wallets; it is scoped to non-identifying attestations, since full-identity flows such as account opening are linkable by nature; and the technical bar is met by **batch issuance** rather than by zero-knowledge proofs, which appear only in a non-binding recital.

Batch issuance (Figure 4) addresses a specific correlation channel. An SD-JWT credential's issuer signature is identical on every presentation of that credential, so two verifiers comparing logs can establish that they encountered the same person even though neither learned a name; the signature functions as a persistent identifier. Under batch issuance the issuer mints many single-use copies in one flow, each asserting the same fact under different random salts, producing different hashes, a different issuer signature, and a different holder key, so that neither signature nor public key is persistent. The wallet consumes each copy once and retrieves a further batch as the pool is depleted. The cost is operational rather than cryptographic, which is the appropriate trade for a trust-critical layer.

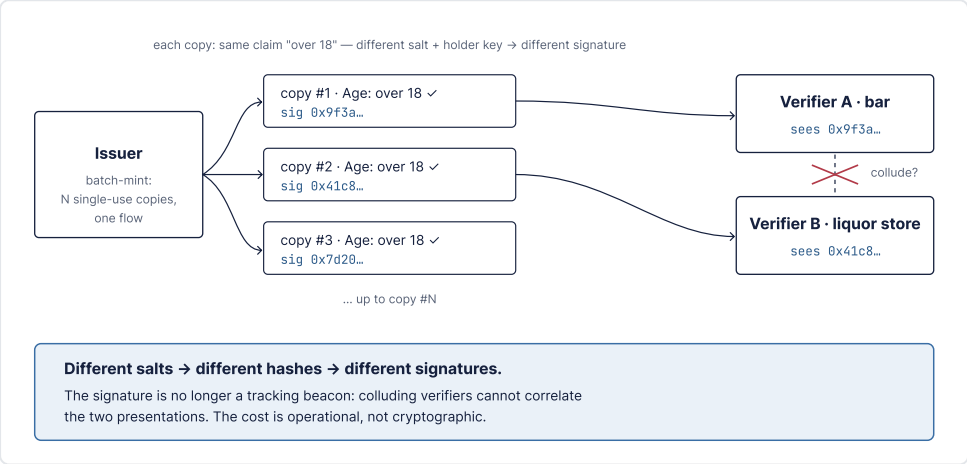

Figure 4. Batch issuance and cross-verifier unlinkability. Each single-use copy carries different salts and a fresh key, producing a different signature, so that neither signature nor public key persists across presentations and colluding verifiers cannot correlate them.

7.3. Assessment of BBS+ and Zero-Knowledge Proofs

Batch issuance defeats correlation between verifiers but not collusion between an issuer and a verifier, which raises the question of stronger constructions. Three considerations govern the assessment. **BBS+ keys cannot reside in a secure element**, since the pairing-friendly curve is unsupported, which breaks hardware binding and the sole-control requirement. **BBS+ is not an accepted credential format** under the European framework, so such a credential forfeits interoperability. And the stronger privacy property is not mandated. The more promising direction is **zero-knowledge proofs over existing ECDSA-signed SD-JWT credentials** `draft`, which would preserve both the hardware-bound key and the credential format while adding collusion resistance; this remains a research line rather than a deployable component. The current position is therefore SD-JWT with batch issuance, with zero-knowledge-over-SD-JWT under review.

---

8. Credentials for Natural Persons

8.1. Credential Portfolio

The reference wallet presents six surfaces: **OTCP ID** (a composite identity card, masked by default and revealed behind a biometric prompt), **Wallet** (the credential portfolio), **Agents** (the delegation surface), **Auth** (an authenticator, RFC 6238 `standard`), **Verify** (which allows the holder to act as verifier and check another party's presentation), and **Settings**. From a single verification session the issuer batch-mints three credentials:

OTCP ID `product`, the basis identity credential, carrying name, birth date, nationality, document type and number, issuing authority, and document images and a liveness selfie as separately disclosable claims; **Age** `product`, an `age_equal_or_over` object with independently disclosable thresholds, allowing a verifier to check a threshold as a single boolean **without disclosure of the birth date**; and **KYC-Cleared** `product`, a reusable verified-person anchor carrying assurance level, date cleared, and provider, whose value lies in reuse, since a second verifier can confirm a prior clearance without repeating the identity check.

The Age credential illustrates the underlying property. A venue confirming that a customer meets an age threshold has no requirement for the birth date, address, or document number, yet a scanned identity document discloses all three. The Age credential reduces the exchange to a single signed boolean, with the source birth date verifiably absent from the presentation.

8.2. Credential Lifecycle

Onboarding is digital-ID-first with an identity-verification fallback: the primary path imports a government digital ID from the device platform wallet through the W3C Digital Credentials API `draft`, and the coverage path performs document capture with biometric liveness. A privacy property is enforced at this point: **personal data does not persist server-side beyond issuance.** The issuer retains a hash commitment and an audit record, sufficient to evidence that an event occurred and to respond to a lawful demand, but not the claims themselves.

Issuance follows the OpenID4VCI pre-authorized-code flow, with a trust panel identifying the issuer, a biometric confirmation, and the credential minted **bound to the wallet's holder key**. Presentation follows OpenID4VP: the verifier names the claims required, the wallet displays the request with per-claim toggles, and returns only the disclosed claims together with a **key-binding proof**, a fresh holder signature over the verifier's nonce and audience, which is the mechanism that prevents replay. The verifier then performs an eight-step validation: format and signature; issuer-trust resolution; holder binding; validity window; revocation status; selective-disclosure validation; compliance hooks; audit-log write. Revocation is a single status-bit flip, and cascades by construction because each credential carries its own status entry.

8.3. Data Minimization Vocabularies

Privacy is expressed as structural properties, each supported by a design rule: selective disclosure and predicate proofs; no personal data at rest; issuer-unlinkability, since the wallet presents directly to verifiers and does not contact the issuer at presentation time; cross-verifier unlinkability through batch issuance; and no public activity log.

Two machine-readable vocabularies carry minimization into the protocol. **Claim modes** state what is transmitted: *predicate* (a boolean over a threshold, disclosing no underlying value); *selective attribute*; *bundle* (a defined set for regulated onboarding); *redacted document* (constrained, for legacy compatibility); and *no-share result* (pass or fail only). **Retention classes** state what a verifier may retain: nothing beyond the session; minimal audit receipt only; receipt with disclosed attributes; regulated retention period; manual-review record; and deletion pending or complete. Both are displayed before the holder consents and enforced by default. A minimal audit receipt records that a check occurred — request and verifier identifiers, credential type, issuer reference, time, purpose, policy version, result, and presentation hash — without raw attributes unless the profile requires them, and must not itself become a correlation channel.

9. Delegated Authority for Agents

A **VDAC** — Verifiable Delegated Agentic Credential — is the artifact by which a verified principal grants an agent scoped, attenuating, revocable authority.

Figure 5. The five layers of agent identity. Implementation has concentrated at the top (naming, attestation) and the bottom (authorization, payment); delegation of authority in the middle is the layer OTCP addresses.

9.1. The Agent-Identity Stack

The agent-identity stack comprises five layers (Figure 5): naming and discovery, agent attestation, **delegation of authority**, per-call authorization, and payment authority. Implementations exist at the top and bottom; the middle layer is thin.

Two technical approaches have emerged. The **OAuth and bearer-token approach** accounts for most deployed code, but its credentials are bearer tokens: replayable if stolen, privacy-blind, and binding the human server-side in a vault rather than within the artifact the verifier examines, with the authority chain opaque to intermediaries and lost at each token exchange. The **verifiable-credential approach** holds the only ratified credential standards, with native selective disclosure and holder binding, but its delegation components are immature. OTCP adopts the second, on the substrate toward which standards work is converging and which national wallet programmes are deploying, because that substrate is simultaneously standards-mandated, privacy-capable, and hardware-bindable. The bearer approach does not provide the three together.

The first commitment follows: **a delegation of authority is a verifiable, portable, signed credential rather than a session, a bearer token, or a server-side access-control entry.** A credential the agent holds and presents preserves the authority chain across intermediaries and permits a verifier to check it offline by signature.

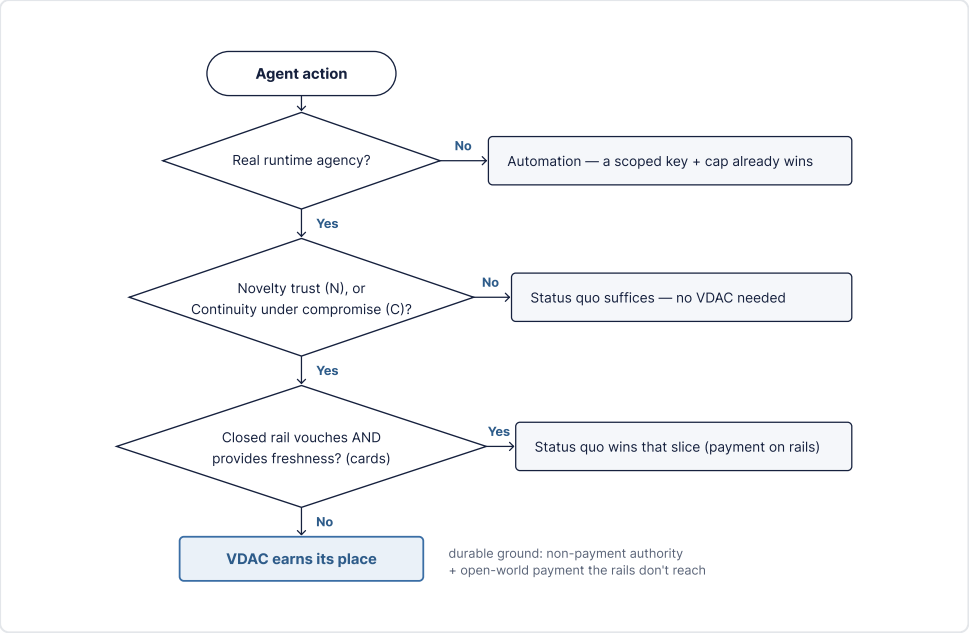

Figure 6. Applicability criteria for a delegation credential: genuine runtime agency as precondition, novelty of trust or continuity under compromise as qualifying condition, and the absence of a closed rail providing both vouching and freshness.

9.2. Applicability Criteria

A delegation credential is warranted only where it improves on existing controls. Many current agent-identity deployments are scripted automation with a spending limit, for which existing controls are adequate. Each scenario is therefore assessed against three criteria (Figure 6).

Precondition: the agent must exercise genuine runtime agency, selecting actions and counterparties at runtime rather than following a fixed script. In the absence of agency the system is automation, and a scoped API key with a limit is sufficient.

Qualifying conditions: a VDAC is warranted where either **N — novelty of trust** applies, in that the verifier must extend trust to a principal with which it has no prior relationship and whose internal controls it cannot inspect; or **C — continuity under compromise** applies, in that the action is sufficiently sensitive that the verifier requires per-request evidence that the delegation is current, unrevoked, and bound to an uncompromised agent, and the incumbent credential is a static bearer token. Agent takeover is the predominant empirical threat, and this second condition, rather than novelty, is the principal justification in practice.

Exclusion: where a closed rail already provides both vouching and freshness — as card networks do for payment, through per-transaction authorization, issuer decline, and step-up authentication — the existing mechanism is sufficient.

These criteria yield a conclusion that determines scope: **payment, where card rails provide both vouching and freshness, is outside VDAC's domain.** The durable domain is **non-payment authority** — proving attributes, operating tools, entering obligations, and signing — where neither a vouching rail nor a freshness mechanism exists, together with **open-world payment** to counterparties the card rails do not reach.

---

10. The VDAC Constitution

The design is governed by a constitution of **principles**, which state the design intent, and **invariants**, which state properties that must hold in all cases (Figure 7).

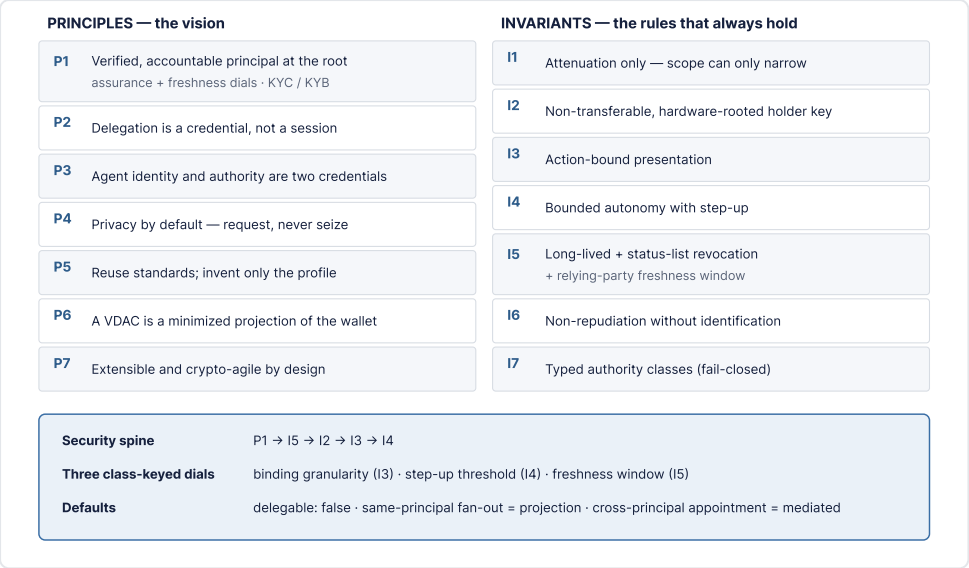

Figure 7. The VDAC constitution: seven principles stating design intent and seven invariants stating properties that must hold in all cases, with the security spine, the three class-keyed policy dials, and the structural defaults.

10.1. Principles P1–P7

	Principle	Statement
P1	A verified, accountable principal at the root	Authority roots in a verified, revocable real-world principal — an individual (KYC), a legal entity (KYB), or an individual acting under an entity (chained). Two parameters are set by the relying party against the stakes of the action: assurance and freshness. No anonymous root is permitted for an action with real-world effect, and no indefinitely standing authority for sensitive actions.
P2	Delegation is a credential, not a session	The grant is a portable signed artifact. Where an ecosystem cannot consume credentials, a VDAC may be exchanged for a short-lived bearer token at a gateway; that token is a derived, disposable artifact and never the root authority.
P3	Agent identity and delegated authority are separate credentials	The identity of the agent and the authority it carries are issued separately and evaluated jointly, and are never combined. OTCP defines the delegation credential and references a pluggable agent identity — device key, name-service name, DID, or workload identity.
P4	Privacy by default	A presentation discloses only what the action requires; cross-verifier unlinkability derives from batch issuance; there is no public activity log. A verifier may request further disclosure but cannot compel it.
P5	Reuse standards; define only the delegation profile	No new format, protocol, or transport. Where a standard cannot express a requirement, a further standard is adopted rather than defined.
P6	A VDAC is a minimized projection of the principal's wallet	At delegation the wallet mints agent-key-bound, scoped, minimized derivative credentials. The agent operates a narrowed projection and never holds copies of the principal's credentials.
P7	Extensible and crypto-agile by design	Every surface extends through a registered namespace, a profile version, and a per-extension critical flag, where critical denotes rejection by any party that does not recognize the extension. Algorithms are never hardcoded, so post-quantum schemes can be introduced without a schema break.

10.2. Invariants I1–I7

	Invariant	Rule
I1	Attenuation only	A delegation cannot convey more authority than the delegator holds; scope narrows only, enforced cryptographically (`child_scope ⊆ parent_scope`). In a dual-root case the reference is the entity's authority. Quantitative over-use and capability composition fall outside attenuation and are governed by caps, minimization (P6), and step-up (I4); attenuation is necessary but not sufficient.
I2	Non-transferable, hardware-rooted holder key	The agent's key is non-exportable and hardware-rooted; authority is never a bearer or transferable asset. This excludes designs in which authority can be sold. Authority relocates only by re-issuance from the principal to a new non-exportable key, which covers rotation, migration, and horizontal scaling.
I3	Action-bound presentation	Every presentation binds to the action: at baseline to verifier or origin and nonce; for state-changing classes the holder additionally signs a digest of the exact action — amount, payee, document hash — so that a modified parameter invalidates the proof.
I4	Bounded autonomy with step-up	Autonomy within scope is the default; step-up, being a fresh human confirmation escalating to a biometric co-signature, is the exception path. The threshold is set independently by principal and verifier and fires on the union, so the agent is autonomous only where neither party requires a human.
I5	Long-lived credentials, status-list revocation, freshness window	VDACs are long-lived and checked against a Token Status List with three cascading levels: revoking the principal root invalidates every derived credential, then the delegation, then the agent key. The freshness window — the maximum acceptable status age for a given action — is set per workflow by the relying party. The registry should be decentralized and principal-controlled, so that a principal can revoke independently of OTCP availability.
I6	Non-repudiation and accountability without identification	Every action is attributable to the named principal, who cannot credibly deny authorizing within scope, while the verifier need not learn the principal's identity. Two grades apply: scope-level for autonomous actions, and action-level for step-up actions.
I7	Typed authority classes	Authority is expressed as a small, closed set of typed risk classes carrying the semantics that drive step-up, attenuation, and accountability, with a fail-closed extension namespace in which an unrecognized type is denied.

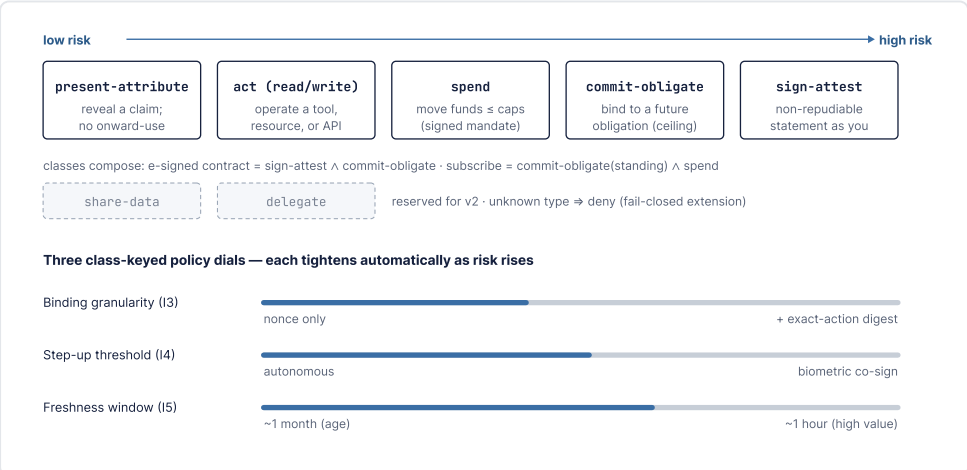

Figure 8. The five composable authority classes and the three policy dials keyed to them. A verifier expresses a workflow's risk appetite by setting binding granularity, step-up threshold, and freshness window, rather than by encoding logic per action.

10.3. Authority Classes and Policy Dials

Two taxonomies are combined — one organized by risk, suited to human consent and accountability, and one organized by action verb, suited to operating tools — yielding five composable classes (Figure 8): **`present-attribute`**, to disclose a verified claim for a decision with no onward use granted; **`act`**, to operate a resource, tool, or API, with a verb (`read` or `write`, extensible) and a resource; **`spend`**, to move funds up to defined caps as a constrained signed mandate, where the rail rather than OTCP moves the funds; **`commit-obligate`**, to bind the principal to a future obligation, carrying a total-liability ceiling distinct from any per-transaction cap, of which a subscription is the standing-recurrence case; and **`sign-attest`**, to produce a non-repudiable statement in the principal's name. Two further classes, `share-data` and `delegate`, are reserved for a subsequent version, each requiring deliberate design in data governance and multi-hop accountability respectively.

Two authorities belong to different classes where they differ in the risk semantics that drive a downstream decision and are separately legible to a consenting human and to a regulator. Legibility is decisive in marginal cases, which is why `spend` remains a first-class named class although it is mechanically a signed mandate: a stated spending limit is the most legible term to which a human consents. A general capability grammar would be more extensible, but the circumstance that justifies a delegation credential — a novel, untrusting verifier — is precisely the circumstance in which an open grammar is least safe, since such a verifier cannot safely interpret an arbitrary capability string it has not encountered. The model is therefore typed classes for legibility and safety, structured constraints for limits, and a fail-closed namespace for extension.

The security spine is **P1 → I5 → I2 → I3 → I4**, and three policy dials are keyed to the authority class: **binding granularity** (I3), **step-up threshold** (I4), and **freshness window** (I5). Each tightens as the class rises in risk, so that a verifier expresses a workflow's risk appetite by setting three parameters rather than encoding logic per action.

10.4. Credential Structure and Verification Algorithm

A VDAC is an SD-JWT VC signed by the principal's wallet key rather than by an OTCP issuer, which keeps OTCP outside the runtime path. It carries the principal's wallet key as issuer; the agent's pluggable identity and its non-exportable holder key; the minimized identity credential, embedded so that the root is verifiable self-contained; platform attestation that the agent key is hardware-rooted; an array of typed authority grants with constraints (per-transaction and cumulative caps, total-liability ceiling, velocity, counterparty allow and deny lists, jurisdictions, recurrence, and step-up threshold), declared as a critical extension; `delegable: false` by default; a freshness default; and three cascading status pointers.

Verification is an **eleven-step algorithm**, in which each step encodes a constitutional rule: parse; check algorithm policy and critical extensions (P7); verify the identity root and its assurance floor (P1); verify the agent's proof of possession with nonce and audience (I2, I3); verify action-binding for state-changing classes (I3); check attenuation (I1); enforce caps (I1, I7); check all three status levels within the freshness window, failing closed for value-bearing classes (I5); require a satisfied step-up assertion where either party's threshold fires (I4); then accept with a private audit record, or return a structured error. Three constructions make the model implementable: **root binding by embedding**, in which the principal's signature on the VDAC serves as proof of possession of the identity credential's holder key, so that the agent never requires the principal's key; **canonical action hashing** over RFC 8785 JSON canonicalization `standard`; and **status-list partitioning** using large shared lists rather than per-principal lists, which would leak information through size and URL.

11. Threat Model, Delegation Topology, and Accountability

11.1. Agent Compromise and Layered Defense

The principal threat is not credential theft in transit but agent takeover: a compromised agent, through prompt injection or a compromised supply chain, acting maliciously with its own legitimate key. The defense decomposes into four layers:

Threat	Defense
Theft or replay — a stolen credential used elsewhere	**I2** — the non-exportable hardware key cannot be exfiltrated
Tamper or repurpose — a channel or intermediary alters or reuses a legitimate proof	**I3** — action-binding; the proof is valid only for the exact action
Manipulation of the agent — injected reasoning causes the agent to author a high-risk malicious action	**I4** — step-up surfaces the actual parameters to the human, who declines
Low-risk residual — a compromised agent performs minor in-scope actions	**I1** — scope and caps bound the exposure

The boundary should be stated explicitly: action-binding prevents a proof from being tampered with or repurposed, but does not prevent a compromised agent from signing a malicious action it has been induced to author. That case is addressed by step-up for high-risk actions and by scope and caps for the residual, which is the limit applicable to any system of this kind. The step-up layer carries the principal weight of that defense: the assertion is signed by the principal's key over the exact action hash together with a liveness attestation, which an agent holding only the agent key cannot produce, and the human, presented with the decoded action, declines a malicious one. The verifier receives the wallet's signed assertion that liveness occurred, and never biometric data.

Failure states are handled explicitly. A missing or stale status fails closed for value-bearing classes and may fail open for `present-attribute`; an unrecognized critical extension is rejected (P7); an out-of-scope action fails the attenuation check and carries no authority, with liability falling on the party that forged it; and an unsatisfiable policy requirement declines the transaction rather than degrading the guarantee. **Recovery is by re-issuance rather than by key movement**: a compromised or rotated agent key is revoked at the agent-key status level and the principal mints a fresh credential to a new non-exportable key, leaving the principal root and the human credentials unaffected.

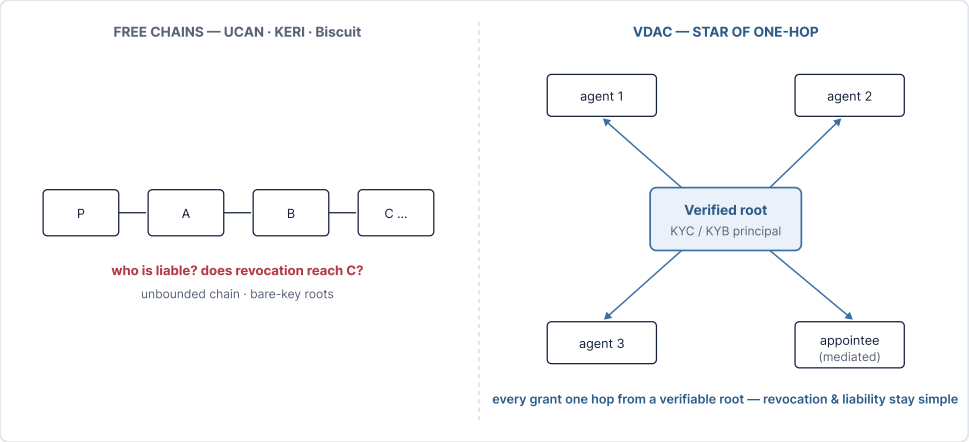

Figure 9. Unbounded chains compared with the star of one-hop. Chain-based systems permit unbounded delegation and subsequently contend with revocation and liability; in OTCP every grant is one hop from a verifiable root, and cross-principal appointment is principal-mediated re-issuance.

11.2. Projection and Cross-Principal Appointment

The principal unsolved problem in agent identity is recursive delegation accountability: where agent A sub-delegates to agent B, the questions of liability and of whether the authority chain survives are unresolved in the general case. OTCP addresses this by constraining the topology rather than by solving the general case (Figure 9).

Where a planner agent spawns workers for parts of a single task, those workers all trace to one principal. This is not delegation across a trust boundary but internal multiplication of the wallet's projection within one revocation envelope, and attenuation guarantees that a compromised sub-agent can perform no action its parent could not perform directly.

The substantive case is cross-principal appointment, in which agent A, acting for principal P, grants authority to an agent B operated by a different party. The rule is that **cross-boundary appointment is always principal-mediated re-issuance**: rather than A minting a sub-credential, the appointment is routed so that a fresh, one-hop, KYC/KYB-rooted credential is issued to B. The resulting property is that **delegation forms a star of one-hop grants from verifiable roots rather than an unbounded chain**, which resolves recursive accountability structurally, in contrast to chain-based systems that permit unbounded delegation and subsequently contend with revocation and liability. The cost is that appointment requires the issuer to be reachable, which is an availability requirement rather than a security compromise. By default every VDAC is non-delegable.

Appointment is also infrequent. Appointment, in which B must exercise P's authority toward a third party, is commonly conflated with a service call, in which B performs work under its own authority; most agent-to-agent interaction is the latter. Appointment is necessary only where B must act as the principal toward a further verifier, was selected at runtime, and holds a capability the principal's own agent cannot retain — the licensed-intermediary case. Because that band is narrow and high-stakes, the cost of mediation is acceptable, and the mechanism is deferred to a subsequent version with the required hooks reserved, so that its introduction is non-breaking.

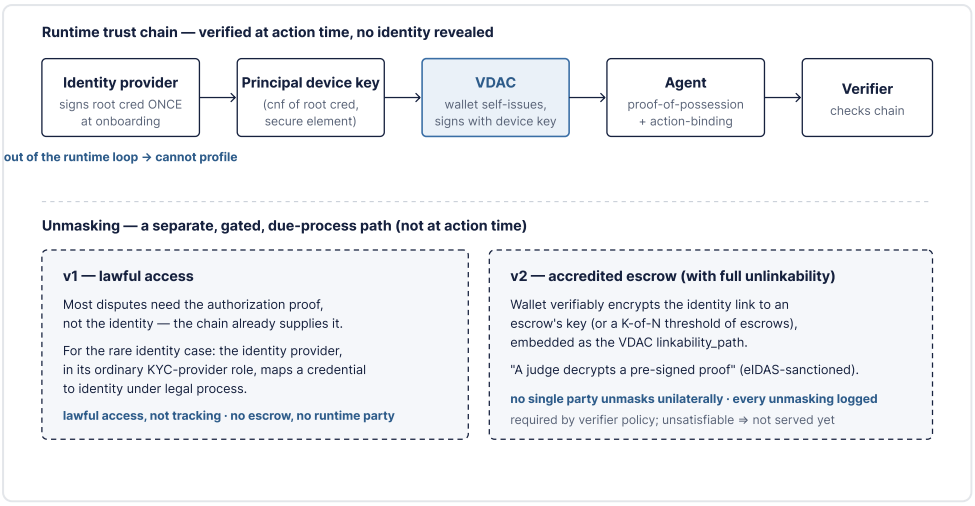

Figure 10. The wallet-centric trust chain and the gated identity-recovery path. The identity provider signs the root once and is thereafter outside the runtime path; the wallet self-issues the agent's credentials. Identity recovery is a separate, due-process event: lawful access in v1, and an accredited threshold escrow as the ecosystem matures.

11.3. Accountability Without Identification

A verifier requires assurance that an accountable principal stands behind an action, and disputes may subsequently require identification of that principal, while privacy requires that the verifier not learn the principal's identity. The resolution separates the two by time and by party: **accountability in principle is present at all times and is non-repudiable, while identification is conditional and gated.** At the time of action the verifier obtains cryptographic evidence that an accountable and identifiable-on-due-process principal authorized the action, and not an identity.

The mechanism is wallet-centric (Figure 10). The identity-verification provider issues the root credential once, at onboarding, and is thereafter outside the runtime path; the principal's wallet self-issues delegation credentials to agents. Because the provider observes no delegations and no presentations, it cannot construct a usage profile, so issuer-unlinkability holds structurally rather than by policy. The verifier's chain runs from the provider's signature on the root, to the principal's device key, to the principal's signature on the VDAC, to the agent's proof of possession and action-binding.

Identity recovery is constrained by the same regulation that governs unlinkability: issuer-held unmasking is precluded, since Article 5a(16) prohibits the issuer from being able to link presentations to identity. The sanctioned model for conditional traceability is verifiable encryption of the identity to a judicial or external authority rather than custody by the issuer. Two phases follow. In **v1**, most disputes require the authorization proof rather than the identity, which the non-repudiation chain already supplies; where identity is genuinely required, recovery proceeds through the identity provider's ordinary role under legal process, which is lawful access rather than tracking. In **v2**, once batch issuance removes the issuer's ability to link, conditional traceability is restored through an opt-in accredited escrow: the wallet verifiably encrypts the identity link to an escrow key, or to a threshold of several, embedded as a linkability path. No single party can unmask unilaterally, every unmasking is logged, and a verifier may require a linkability path as a matter of policy; where no accredited escrow yet exists, that requirement is unsatisfiable and the corresponding use cases are not served.

The trade-off should be stated directly: enforceable non-repudiation requires that the principal cannot prevent unmasking, so the capability must reside with a party other than the principal. This is inherent to accountability. The design objective is that it be gated, auditable, and distributed.

12. Applied Scenarios

Reuse of a verified identity. An individual completes verification once and receives a KYC-Cleared credential. On opening an account at a different institution some months later, they present that credential rather than re-uploading documents: the second verifier resolves the issuer against the federated lists, confirms that the assurance level meets its policy floor, checks status freshness, and accepts the prior clearance. Verification performed once is reused rather than repeated, which is the economic basis of the human credential layer.

Delegation to an agent. A principal instructs an assistant agent to organize a catered event involving food and alcohol. This is not a single transaction but a set of sub-transactions across several verifiers, and the wallet mints a minimized projection accordingly: a `spend` credential with caps and a step-up threshold, a `present-attribute` credential for age, an address credential, and a `commit-obligate` credential for the catering contract, which the principal may narrow but not widen on an approval screen, confirmed by biometric. The set then decomposes, and only some elements require a credential. The caterer's deposit runs on card rails, which already provide vouching and freshness, so the credential adds nothing. Committing the principal to the catering contract, with its cancellation penalty, creates open-ended future liability for which no card network vouches, and falls within the domain of delegated authority. The licensed alcohol retailer, having no prior relationship with the principal, requires a regulated attribute and a deliverable address, and learns that the principal is verified above the age threshold at that address and nothing further. The physical handover, where law requires an adult present with identification, is not delegable: the credential authorizes the order and not physical presence.

Combined use by a person and their agents. A principal operating a small import business is KYC-verified as the authorized representative of a KYB-verified company, chained. That single onboarding serves both roles. As an individual, she reuses her clearance to open a second business account without re-proofing. As a principal, she delegates from the same wallet: a procurement agent receives a `commit-obligate` credential rooted in the company's authority, and a logistics agent receives authority to appoint a licensed customs broker through mediated re-issuance. When the procurement agent presents to a supplier with no prior knowledge of the company, the supplier verifies a chain terminating in the same accredited identity credential used at the bank; when an agent is offboarded, one status flip cascades through every credential it held, leaving the human credentials unaffected. The result is one verified identity, one revocation surface, and one privacy model. Where an agent's authority must root in a verified person or company, a separate agent-identity stack collapses into the human one: an agent's authority is not a distinct form of identity but a projection of a verified human one.

Settlement remains on existing rails. OTCP's relationship to payment is defined by one constraint: **it does not settle, hold funds, or take a share of payment value.** It authorizes; settlement occurs on the verifier's own rail. This keeps the platform outside payment-services, money-transmission, and crypto-asset-service classifications, which is a deliberate constraint, since a trust layer that touches settlement value acquires financial regulation incompatible with its neutrality. Card networks already provide vouching and freshness for closed-loop payment, so OTCP interoperates rather than competes: a `spend` grant is a constrained signed mandate, and a VDAC exports a deterministic AP2 intent mandate `draft`, which is a superset of the emerging payment-mandate format. The domain in which delegated authority does govern payment is the open-world corridor the card rails do not reach `draft`. There too OTCP authorizes while the agent's own wallet moves value, and a monetary cap is denominated rather than routed, preserving asset and chain neutrality.

13. Governance, Design Decisions, and Open Problems

13.1. Ecosystem Governance

It is sometimes assumed that agent delegation requires a registry of trusted agents, which would place OTCP in a gatekeeping role. Three mechanisms make such a registry unnecessary. First, **trust in a delegation credential is transitive**: a verifier accepts it not because the principal appears on a list but because it chains to a root identity credential whose issuer is resolvable on existing lists, so delegation credentials require no list of their own. Second, **new trust-service roles**, such as an escrow or an agent-attestation service, are admitted through extensible federation — entity statements, accreditation credentials, and multi-stakeholder governance — rather than through an OTCP-operated gate. Third, what is excluded is a *gatekeeping* list, being single-source, mandatory, and OTCP-controlled, rather than publication as such: an open, signed, transparent, forkable, and verifier-replaceable registry, offered as one option among several, remains permissible. The only registry OTCP itself publishes is a vocabulary registry covering authority classes, policy dials, and critical-extension members.

13.2. Design Decisions

Decision	Rationale
Revocation via **IETF Token Status List**, not W3C Bitstring	The mechanism used by the European reference implementation for SD-JWT VC, and by the OTCP build
Holder key **ES256 / P-256**, not EdDSA	The only curve backable in hardware in the most constrained common secure element, which sole control requires; also mandated by mdoc. Algorithm agility retained
Batch issuance rather than BBS+ or ZKP at present	BBS+ cannot reside in a secure element and is not an accepted format; the unlinkability requirement is satisfied by batch issuance. Zero-knowledge-over-SD-JWT under review
SD-JWT VC and mdoc, both from the outset	Both are mandated formats; mdoc is supported holder-side from the first release
Conformance and certification track from the outset	Retrofitting conformance costs more than building to it, and the requirements are largely a superset of those imposed by other current regimes
No proprietary trust list	Neutrality is a structural requirement: consume federated lists, and federate for new roles
No open delegation chains; cross-principal appointment is mediated	Resolves recursive accountability structurally
Settlement remains on existing rails; delegated authority covers non-payment and open-world payment	Card rails already provide vouching and freshness; OTCP does not settle, which keeps it outside payment regulation
Native Android reference wallet	mdoc support, secure-element keys, and proximity presentation are native-first, and the available reference libraries are Kotlin
`delegable: false` by default, with hooks reserved	Re-delegation is the least settled area of the field; reserving the mechanism keeps later introduction non-breaking

13.3. Open Problems

- **Cross-border assurance equivalence.** Comparison of assurance levels across regimes is unresolved; OTCP declines equivalence tables in favor of profile-specific, source-cited acceptance rules.
- **Wallet certification.** Conformance from the outset with a certification track is the approach, but certification itself proceeds under each regulator's timetable.
- **Unlinkability against collusion.** Batch issuance does not defeat issuer–verifier collusion; only zero-knowledge over SD-JWT would eventually do so.
- **Status privacy.** Revocation checks can leak holder activity if poorly designed, which motivates shared-list partitioning and, for delegation credentials, decentralized principal-controlled revocation.
- **Business identity.** Legal-entity identity requires registries, representative authority, and role credentials, and constitutes a distinct product line rather than an extension of personal identity.
- **Offline use.** Proximity presentation raises status-freshness questions that must be resolved per profile.
- **Recursive multi-hop delegation across principals.** OTCP constrains the topology to mediated, one-hop, re-rooted appointment; the general cryptographic-chain problem remains unresolved in the field, and OTCP does not claim to solve it.
- **Cross-verifier cumulative caps, and the escrow ecosystem.** Both depend on a multi-stakeholder trust framework that does not yet exist.

---

14. Conclusion

OTCP addresses two requirements that must be met together. The first — repeated proofing, overcollection, and fragmented rails — is met by a horizontal, standards-compatible trust and compliance layer. The second — agents that can act but cannot demonstrate whose authority they carry — is met by rooting every delegation in a verified, accountable principal and expressing it as a scoped, attenuating, revocable, action-bound credential governed by seven principles and seven invariants. Two consequential positions follow: delegation forms a star of one-hop grants from verifiable roots rather than an unbounded chain, and settlement remains on existing rails while delegated credentials govern non-payment authority and open-world payment authority.

Neither requirement can be met by a product owned by one company in one jurisdiction. Both can be met by infrastructure: schemes remain sovereign, accreditation remains local, verifiers remain accountable to their own regulators, and the layer between them is neutral, open, and standards-conformant.

Participation is invited on that basis. For **standards bodies and researchers**, the delegation profile, the authority classes, and the verification algorithm are offered as input, and the open problems in Section 13 are stated as open. For **regulators and supervisory bodies**, the profile mechanism exists so that acceptance rules, retention classes, and relying-party regimes are expressed as machine-readable policy rather than as vendor assurance. For **trust service providers, issuers, and wallet providers**, the integration surface is deliberately small and is not proprietary. For **operators and enterprises**, the delegation layer is implementable on the same verified root already used to onboard customers.

An accountability layer for the agent economy will be established in one form or another. The question this paper addresses is whether it is rooted in self-asserted keys and replayable tokens, or in verified, accountable, privacy-preserving human and organizational authority, on the rails the world's identity schemes already share.

Appendix A — Glossary and Standards References

Glossary. VDAC — Verifiable Delegated Agentic Credential, the scoped, revocable authority a principal grants an agent. **SD-JWT VC** — a credential format with per-claim selective disclosure. **mdoc** — the ISO/IEC 18013-5 and -7 mobile-document format. **OID4VCI / OID4VP** — OpenID protocols for issuance and presentation. **HAIP** — High Assurance Interoperability Profile. **Token Status List** — the IETF revocation mechanism, a signed compressed bit array. **WSCD** — Wallet Secure Cryptographic Device, the secure element holding keys under sole user control. **KYC / KYB** — verification of an individual and of a legal entity respectively. **LoA** — Level of Assurance. **Attenuation** — the rule that delegated authority narrows and never widens. **Step-up** — a fresh human, commonly biometric, confirmation required for high-risk actions. **Projection** — the model in which the wallet mints minimized, agent-bound derivative credentials rather than copying its own. **Appointment** — a cross-principal grant of authority, realized as principal-mediated one-hop re-issuance. **Batch issuance** — issuance of many single-use copies with fresh salts and keys, so that presentations cannot be correlated. **Predicate proof** — proof of a fact without disclosure of the underlying data.

Credential and protocol standards. SD-JWT VC (IETF) `standard`; ISO/IEC 18013-5 and -7 mdoc `standard`; OpenID4VCI and OpenID4VP `standard`; HAIP `profile`; IETF Token Status List `standard`; RFC 8785 JSON Canonicalization `standard`; RFC 9421 HTTP Message Signatures `standard`; RFC 6238 TOTP `standard`; W3C Digital Credentials API `draft`; W3C Verifiable Credentials 2.0 `standard`; DID Core `standard`.

Regulatory. Regulation (EU) 2024/1183 (eIDAS 2.0, notably Article 5a(16)) `standard`; the EUDI Wallet Architecture and Reference Framework `profile`; the EU AI Act `standard`; the revised Product Liability Directive `standard`; ETSI TS 119 612 trusted lists `standard`.

Agentic and adjacent. Google AP2 `draft`; the x402 open-world payment pattern `draft`; Anthropic MCP `draft`; Google and Linux Foundation A2A `draft`; the DIF Trusted AI Agents work and its delegated-authority and know-your-agent task forces `draft`; the academic authenticated-delegation literature `draft`; UCAN, KERI/ACDC, and Biscuit `draft`; OpenID Federation `standard`; EBSI `implementation-specific`; OpenWallet Foundation Multipaz and the European reference Android wallet libraries `implementation-specific`.

National and sectoral rails (compatibility targets). European wallets, PID and attestation providers, and qualified trust service providers `implementation-specific`; Ukraine's Diia, NBU BankID, ID.GOV.UA, and electronic trust services `implementation-specific`; Georgia's PSDA eID, ID.GE, MY.GOV.GE, and its trust-services law `implementation-specific`; mobile driving licences on ISO 18013-5 `implementation-specific`; open-source foundational identity platforms `implementation-specific`; ICAO digital travel credentials `draft`; GLEIF legal-entity identifiers `standard`.

OTCP Lightpaper v3.0 — 2026-08-22. Conformant by design; certification pursued under each applicable regime. Correspondence and technical review are welcome.