

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: October 27, 2026

R.J. Coetzee  
Violet Sky Security SEZC  
April 27, 2026

Sovereign Policy Token Transactions (SPT-Txn)  
draft-coetzee-oauth-spt-txn-tokens-01

## Abstract

This document specifies the Sovereign Policy Token Transaction (SPT-Txn) Framework, a token-based architecture for policy-bound, tamper-evident authorization across trust boundaries in agentic and multi-organizational systems. The framework extends IETF Transaction Tokens (RFC 9700) with Capability Acquisition Tokens (CATs) -- cryptographically scoped authorization tokens that bind capability grants to specific transaction contexts, propagate human-origin identity commitments across delegation chains, and support offline verification without issuer interaction.

SPT-Txn addresses the gap between coarse-grained bearer token authorization and the fine-grained, auditable, scope-limited authorization required in regulated financial infrastructure, AI agent systems, and cross-organizational transaction chains. The framework provides cryptographic guarantees that every hop in a transaction chain was authorized by a verified human principal, operated within a declared capability scope, and can be audited without exposing personally identifiable information.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 27, 2026.

## Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

## Table of Contents

1. Introduction .....	1
1.1. Terminology .....	4
1.2. Changes from -00 .....	5
1.3. Changes from -03 .....	5
2. Problem Statement .....	6
2.1. Limitations of Current Txn-Tokens .....	6
2.2. Requirements .....	6
3. TBAC Capability Token Model .....	7
3.1. The ABAC-to-TBAC Boundary .....	7
3.2. Capability Token Structure .....	7
3.3. Eight-Step Enforcement Engine .....	8
3.4. Scope Invariants .....	9
3.5. Token Serialization Formats .....	9
4. SPT-Txn Token Structure .....	10
4.1. Base Claims (Inherited) .....	10
4.2. SPT-Txn Extension Claims .....	10
5. ZK Circuit Specification .....	11
5.1. Proof System .....	11
5.2. Circuit 1: Attribute Proof .....	11
5.3. Circuit 2: Compliance Claim .....	12
5.4. Circuit 3: Biometric Commitment .....	12
6. Multi-Chain Trust Registry Model .....	13
6.1. Chain-Agnostic Trust Registry .....	13
6.2. Chain-Specific Deployment Notes .....	13
6.2.1. Ethereum / EVM-Compatible Chains .....	13
6.2.2. XDC Network .....	13
6.2.3. Algorand .....	13
6.2.4. Hedera Hashgraph .....	13
7. Transaction Flow .....	14
8. Security Considerations .....	15
8.1. Capability Token Security .....	15
8.2. Delegation Chain Security .....	15
8.3. ZK Proof Security .....	15
8.4. Trust Registry Security .....	15
9. Privacy Considerations .....	16
9.1. Human Anchor Privacy .....	16
9.2. Compliance Reference Privacy .....	16
9.3. Selective Disclosure .....	16
9.4. Auditability vs. Privacy .....	16
9.5. GDPR Compliance .....	16
9.6. Human Anchor Escrow .....	16
9.6.1. Escrow Function and Position in the Architecture .....	16
9.6.2. Escrow Envelope Construction .....	17
9.6.3. Escrow Key Material Requirements .....	17
9.6.4. Threshold Authorization Requirements .....	18
9.6.5. Deanonymization Request Interface .....	19
9.6.6. Escrow Audit Log .....	20
9.6.7. Escrow Key Lifecycle .....	20
9.6.8. Compromise and Recovery .....	21
9.6.9. Implementation Note (Non-Normative) .....	22
10. Security Properties .....	23
10.1. Signature Algorithm Requirements .....	23
10.2. Key Management .....	23
10.3. Transport Security .....	23
10.4. Replay Prevention .....	23
10.5. Token Format Security Properties .....	23
10.6. ZK Proof System Security .....	23
10.7. Multi-Chain Trust Registry Integrity .....	23
11. IANA Considerations .....	24
12. References .....	25
12.1. Normative References .....	25

12.2. Informative References .....	26
Author's Address .....	26

Coetzee Expires October 27, 2026 [Page 2]

Internet-Draft SPT-Txn Tokens April 2026

## 1. Introduction

Authorization in distributed and agentic systems requires answering three questions with cryptographic certainty at every transaction hop:

- (a) Was this action authorized by a verified human principal?
- (b) Does the acting entity operate within an explicitly declared capability scope?
- (c) Can the full chain of authorization be audited without accessing personally identifiable information?

Existing approaches cannot answer all three simultaneously. Bearer tokens (OAuth 2.0 [RFC6749]) answer none of them at the cryptographic layer. Transaction Tokens (RFC 9700) address workload identity propagation but not capability scoping or privacy-preserving human traceability. Demonstrating Proof of Possession (DPoP, RFC 9449) binds tokens to a key pair but not to a transaction context or a human-origin commitment.

The consequence is that when an AI agent or automated service takes an action that causes harm, the audit record is a log. Logs are not proof. They are mutable, selectively produced, and unavailable to external auditors without system access. What regulated environments require is cryptographic proof of authorization scope at every hop -- immutable, verifiable by any party holding the appropriate public key, and auditable without PII exposure.

This document specifies the Sovereign Policy Token Transaction (SPT-Txn) Framework, which provides this capability through two coordinated mechanisms:

**Capability Acquisition Tokens (CATs):** Cryptographically signed authorization tokens issued by an Attribute-Based Access Control (ABAC) Policy Decision Point (PDP), binding a specific capability scope to a holder identity, transaction context, and delegation depth. CTs are verifiable offline and cannot be used outside their declared scope.

**SPT-Txn Transaction Tokens:** An extension of IETF Transaction Tokens that adds CT references, scope hashes, human-origin commitments (humanAnchor), and compliance proof references to every hop in the transaction chain. The chain is immutable and carries cryptographic evidence of its full authorization history.

Together, these mechanisms enable regulated financial systems, AI agent orchestration platforms, and cross-organizational transaction chains to provide cryptographic proof that every action was authorized by a verified human, within declared scope, with a complete and auditable chain of custody -- without exposing identity information to intermediate parties.

The framework is designed for direct alignment with IETF standards. CTs are normatively serialized as JSON Web Tokens [RFC7519] with JSON Web Signatures [RFC7515]. SPT-Txn Transaction Tokens extend the base claims defined in RFC 9700. The ZK proof system uses Groth16 over BN254, with circuit specifications at the

normative level. Trust registries are chain-agnostic and support Ethereum, XDC, Algorand, and Hedera deployments.

Coetzee

Expires October 27, 2026

[Page 3]

Internet-Draft

SPT-Txn Tokens

April 2026

The SPT-Txn Framework directly addresses the requirements of the EU AI Act Article 14 (human oversight of high-risk AI systems). The humanAnchor claim provides cryptographic traceability to the authorizing human; CT scope provides bounded authorization that the AI system cannot exceed; the revocation mechanism provides the override capability Article 14 requires; and the HCS audit chain provides monitoring evidence without PII exposure.

Formal cryptographic security definitions and proofs for the token binding security property are provided in the companion theory paper [SPT-THEORY]. The theory paper defines transaction binding security as a game-based security property and proves that the SPT-Txn construction achieves it under EUF-CMA assumptions in the random oracle model.

## 1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document uses the following terms:

### Capability Acquisition Token (CAT):

A signed token issued by an ABAC Policy Decision Point, binding a capability scope to a holder identity and transaction context. Normative serialization: JWT [RFC7519]. SD-JWT, Biscuit, and CWT are alternative formats (Section 3.5).

### Capability Token (CT):

Synonym for Capability Acquisition Token. Used interchangeably in this document.

### ABAC PDP:

Attribute-Based Access Control Policy Decision Point. The issuing authority for Capability Tokens. Evaluates policy against subject attributes and issues scoped CTs.

### TBAC:

Token-Based Access Control. The enforcement model in which resource servers verify presented tokens rather than querying a central policy engine at each request.

### humanAnchor:

A zero-knowledge commitment to the zkDID of the human principal who authorized the root transaction. Propagated immutably through all delegation hops. Enables human traceability without PII exposure.

### zkDID:

A zero-knowledge Decentralized Identifier commitment that proves identity membership in a verified set without revealing the underlying identifier.

### delegation\_depth:

The remaining number of sub-delegation levels permitted from the current CT holder. Decrement at each delegation.  
A CT with depth 0 MUST NOT issue further delegation CTs.

Trust Registry:

An on-chain smart contract registry that maps ABAC PDP identities to capability types they are authorized to certify. Used by resource servers to verify CT issuer authority without contacting the issuer.

Transaction Token (Txn-Token):

As defined in RFC 9700. A short-lived signed token representing a workload's identity and context within a transaction chain. SPT-Txn extends the base Txn-Token claim set.

SPT-Txn Token:

An SPT-Txn-extended Transaction Token carrying the additional claims defined in Section 4.2.

Coetzee	Expires October 27, 2026	[Page 4]
Internet-Draft	SPT-Txn Tokens	April 2026

## 1.2. Changes from -00

The following changes were made in this revision:

- (a) Section 9 (Privacy Considerations) reorganized into numbered subsections 9.1 through 9.5. Text content of these subsections is unchanged from -00.
- (b) New Sections 9.6.1 through 9.6.9 added, specifying Escrow Function and Position, Escrow Envelope Construction, Escrow Key Material Requirements, Threshold Authorization Requirements, Deanonymization Request Interface, Escrow Audit Log, Escrow Key Lifecycle, Compromise and Recovery, and a non-normative Implementation Note.
- (c) No changes to wire format, token claims, enforcement engine, ZK circuits, or trust registry interfaces.
- (d) No changes to IANA Considerations.

## 1.3. Changes from -03

The following changes were made in this revision:

- (a) Section 1.1 (Terminology): CT definition updated from "signed JWT" to "signed token" with a normative note: "Normative serialization: JWT. SD-JWT, Biscuit, and CWT are alternative formats (Section 3.5)."
- (b) Section 3.2: Added "normatively" qualifier to the CT serialization statement.
- (c) Section 3.3 Step 1: Signature verification now specifies JWS for JWT/SD-JWT, public-key block verification for Biscuit, and COSE verification for CWT.
- (d) Section 3.5 (new): Token Serialization Formats. Full section with RFC 2119 normative language covering SD-JWT selective disclosure and data minimisation, Biscuit cryptographic attenuation with Ed25519 and Datalog, CWT for high-throughput

agent-to-agent, and format selection criteria.

- (e) Section 6.2.4: Hedera Hashgraph section numbering corrected (was duplicate 6.2.3 in -03).
- (f) Section 10.5 (new): Token Format Security Properties. Ed25519 and COSE alongside ECDSA; post-quantum vulnerability note; Biscuit structural attenuation defence. Existing Sections 10.5 and 10.6 renumbered to 10.6 and 10.7.
- (g) References: Added SD-JWT [I-D.ietf-oauth-selective-disclosure-jwt], RFC 8392 (CWT), and Biscuit [BISCUIT]. Updated SPT-WP reference from v3.0 to v6.0.

Coetzee	Expires October 27, 2026	[Page 5]
Internet-Draft	SPT-Txn Tokens	April 2026

## 2. Problem Statement

### 2.1. Limitations of Current Txn-Tokens

RFC 9700 Transaction Tokens solve the workload identity propagation problem: a downstream service can verify that a call originated from a known, authenticated workload rather than an arbitrary external caller. This is a significant security improvement over bearer token architectures.

However, RFC 9700 Txn-Tokens have four limitations that make them insufficient for regulated agentic systems:

No capability scoping: A Txn-Token proves that a workload made a call; it does not prove that the workload had authorization to make that specific call with that specific scope. The workload's authorization is evaluated separately by a policy engine at each hop, creating a distributed policy evaluation problem in cross-organizational deployments.

No human traceability: Txn-Tokens propagate the "sub" claim identifying the originating user. In cross-organizational chains, this propagates PII to every service in the chain, creating GDPR data minimization violations. Alternatively, organizations strip the subject claim, losing human traceability entirely.

No cross-organizational trust: RFC 9700 assumes a shared authorization server within a single trust domain. When a transaction crosses organizational boundaries, services in different organizations cannot verify each other's Txn-Tokens without a pre-established trust relationship at the authorization server level.

No compliance propagation: Regulated transactions require compliance attestations (KYC, AML, jurisdiction verification) at the root. RFC 9700 provides no mechanism for propagating compliance evidence through the chain without re-verification at each hop.

### 2.2. Requirements

The SPT-Txn Framework is designed to satisfy the following requirements:

REQ-1: A resource server MUST be able to verify that the

presenting workload holds a valid capability grant from an authorized issuer, scoped to the requested operation, without contacting the issuing authority at request time.

REQ-2: The capability grant MUST be cryptographically bound to a specific transaction context such that it cannot be reused in a different transaction without detection.

REQ-3: Every hop in a transaction chain MUST carry a cryptographic commitment to the human principal who authorized the root transaction. This commitment MUST NOT reveal the human's identity to intermediate services.

REQ-4: Capability scope MUST be monotonically non-increasing through the delegation chain. No delegatee MAY hold greater capability than its delegator.

REQ-5: The maximum delegation depth MUST be set at root issuance and enforced cryptographically at each delegation step.

REQ-6: Resource servers in different organizations MUST be able to verify capability grants issued by issuers in other organizations using a shared trust registry, without bilateral pre-configuration.

REQ-7: Compliance attestations established at transaction root MUST be propagatable to downstream services by reference, without re-exposing the underlying compliance data.

REQ-8: The complete authorization chain MUST be auditable by an authorized auditor without requiring access to PII.

Coetzee

Expires October 27, 2026

[Page 6]

Internet-Draft

SPT-Txn Tokens

April 2026

### 3. TBAC Capability Token Model

#### 3.1. The ABAC-to-TBAC Boundary

The SPT-Txn Framework operates at the boundary between Attribute-Based Access Control (ABAC) and Token-Based Access Control (TBAC).

ABAC evaluation is expensive: it requires policy retrieval, attribute collection, and real-time evaluation at each access decision. In high-throughput transaction chains, ABAC evaluation at every hop introduces unacceptable latency and creates availability dependencies on the ABAC Policy Decision Point.

The SPT-Txn approach performs ABAC evaluation once, at the transaction root, and materializes the result as a Capability Token -- a signed, scope-limited, cryptographically verifiable authorization artifact. Downstream enforcement is TBAC: verify the token, check the scope, make the decision. No policy engine call required.

This boundary is the fundamental architectural claim of SPT-Txn: ABAC at the root, TBAC at every subsequent hop.

#### 3.2. Capability Token Structure

A Capability Token (CT) is normatively a JSON Web Token [RFC7519]

signed with JSON Web Signature [RFC7515]. Alternative serialization formats are specified in Section 3.5.

CTs MUST contain the following claims:

**iss:** The DID or HTTPS identifier of the ABAC PDP that issued this CT. REQUIRED.

**sub:** The identity of the CT holder. SHOULD be a workload identity (SPIFFE SVID, DID, or equivalent). REQUIRED.

**iat:** Issued-at time. REQUIRED.

**exp:** Expiry time. REQUIRED. CTs MUST have a finite lifetime.

**jti:** JWT ID. A unique identifier for this specific CT. REQUIRED. Used for revocation registry lookups.

**ct\_type:** The capability type granted. MUST be a URI identifying the capability. REQUIRED.

**ct\_scope:** The capability scope granted. A structured object defining the specific operations, resources, and constraints permitted. REQUIRED.

**human\_anchor:** A zero-knowledge commitment to the zkDID of the authorizing human principal. REQUIRED for CTs in agentic chains. Format: bytes32 (Poseidon hash of zkDID commitment).

**delegation\_depth:** The remaining delegation depth permitted from this CT holder. REQUIRED. MUST be a non-negative integer. A value of 0 means the holder MAY NOT issue further delegation CTs.

**max\_depth:** The maximum delegation depth set at root issuance. REQUIRED. MUST equal the delegation\_depth of the root CT.

**parent\_ct:** The hash of the parent CT in the delegation chain. REQUIRED for all non-root CTs. OPTIONAL for root CTs.

**compliance\_ref:** A hash reference to the ZK compliance proofs established at transaction root. REQUIRED for regulated transaction contexts. Format: bytes32.

**revocation\_nonce:** A nonce included in the revocation registry commitment. REQUIRED.

Example Capability Token (decoded JWT payload):

```
{
  "iss": "did:web:abac-pdp.example.org",
  "sub": "spiffe://org.example/workload/agent-a1",
  "iat": 1741017600,
  "exp": 1741021200,
  "jti": "ct-7f3a9b2c-4d1e-8f6a-b5c3-2e9d1a7b4f8e",
  "ct_type": "urn:example:capability:financial-transfer",
  "ct_scope": {
    "operations": ["transfer", "query"],
    "max_amount_usd": 50000,
    "currencies": ["USD", "EUR"],
```



```

    "jurisdictions": ["US", "EU"]
  },
  "human_anchor": "0x7f3a9b2c4d1e8f6ab5c32e9d1a7b4f8e...",
  "delegation_depth": 2,
  "max_depth": 3,
  "parent_ct": "0x1a2b3c4d5e6f7a8b9c0d1e2f3a4b5c6d...",
  "compliance_ref": "0x9a8b7c6d5e4f3a2b1c0d9e8f7a6b5c4d...",
  "revocation_nonce": "0x3c4d5e6f7a8b9c0d1e2f3a4b5c6d7e8f..."
}

```

### 3.3. Eight-Step Enforcement Engine

Resource servers MUST implement the following eight-step enforcement procedure when a CT is presented. The enforcement engine is format-agnostic; format-specific parsing is handled by a token adapter layer (Section 3.5).

#### Step 1: Token Signature Verification

For JWT/SD-JWT: Verify the JWS signature using the issuer's public key. The issuer key MUST be retrieved from the Trust Registry (Section 6), not from a URL in the token.  
 For Biscuit: Verify the public-key block signature using the Ed25519 public key from the Trust Registry.  
 For CWT: Verify the COSE\_Sign1 structure using the issuer's key from the Trust Registry.

If signature verification fails, MUST reject with HTTP 401.

#### Step 2: Issuer Trust Verification

Retrieve the Trust Registry entry for `ct.iss`. Verify that the issuer is authorized to certify capability type `ct.ct_type`. If the issuer is not listed for this capability type, MUST reject with HTTP 403.

#### Step 3: Temporal Validity

Verify `ct.exp > now` and `ct.iat <= now`.  
 If either check fails, MUST reject with HTTP 401.

#### Step 4: Revocation Check

Query the revocation registry using `ct.jti` and `ct.revocation_nonce`. If the CT appears in the revocation registry, MUST reject with HTTP 401.

#### Step 5: Scope Verification

Verify that the requested operation and parameters fall within `ct.ct_scope`. Scope checking MUST be performed against the structured scope object, not a free-form string. If the operation is outside scope, MUST reject with HTTP 403.

#### Step 6: Delegation Depth Verification

If this CT was issued by delegation (`parent_ct` is present), verify that `ct.delegation_depth < parent.delegation_depth` and `ct.max_depth == parent.max_depth`.  
 If either check fails, MUST reject with HTTP 403.

#### Step 7: Human Anchor Verification

Verify that `ct.human_anchor` is present and non-zero.  
 Verify that `ct.human_anchor` matches the `human_anchor` claim in the SPT-Txn Transaction Token presented alongside the CT.

If anchors do not match, MUST reject with HTTP 403.

#### Step 8: Scope Hash Binding

Compute hash(ct.ct\_scope) and compare with the spt\_ct\_scope\_hash claim in the accompanying SPT-Txn Token. If hashes do not match, MUST reject with HTTP 403.

Coetzee	Expires October 27, 2026	[Page 8]
Internet-Draft	SPT-Txn Tokens	April 2026

### 3.4. Scope Invariants

The following invariants MUST hold across the delegation chain. Implementations MUST enforce these at issuance time.

Invariant 1 (Scope Monotonicity):

For any CT\_child delegated from CT\_parent:  
ct\_child.ct\_scope is a subset of ct\_parent.ct\_scope

No delegatee MAY hold a scope that is not a subset of its delegator's scope. Scope extension through delegation is explicitly prohibited.

Invariant 2 (Depth Monotonicity):

For any CT\_child delegated from CT\_parent:  
ct\_child.delegation\_depth < ct\_parent.delegation\_depth  
ct\_child.max\_depth == ct\_parent.max\_depth

Invariant 3 (Human Anchor Immutability):

For any CT in a delegation chain rooted at CT\_root:  
ct.human\_anchor == ct\_root.human\_anchor

The human anchor MUST NOT be modified at any delegation step.

Invariant 4 (Compliance Reference Continuity):

For any CT\_child delegated from CT\_parent:  
ct\_child.compliance\_ref == ct\_parent.compliance\_ref

Compliance proofs established at root propagate by reference and MUST NOT be replaced or omitted at delegation steps.

### 3.5. Token Serialization Formats

The TBAC enforcement engine (Section 3.3) operates through a format-agnostic verification interface. Format-specific parsing is isolated behind a token adapter layer. The eight enforcement steps operate identically regardless of serialization format.

#### 3.5.1. Primary Format: JWT with Selective Disclosure (SD-JWT)

JWT [RFC7519] is the normative CT serialization for cross-organizational interoperability. All implementations MUST support JWT.

SD-JWT [I-D.ietf-oauth-selective-disclosure-jwt] extends JWT with selective disclosure of individual claims. Implementations SHOULD support SD-JWT where data minimization requirements apply.

SD-JWT selective disclosure enables a CT holder to present only the subset of claims required by the verifying resource server, without revealing the full CT payload. This is particularly relevant for the ct\_scope claim, where a holder MAY disclose

only the sub-scope relevant to the requested operation, while the issuer's binding covers the full scope.

SD-JWT CTs MUST protect the following claims from selective omission (they MUST always be disclosed):

iss, sub, iat, exp, jti, human\_anchor, delegation\_depth, max\_depth, revocation\_nonce.

The ct\_scope and compliance\_ref claims MAY be selectively disclosed in SD-JWT presentations.

### 3.5.2. Alternative Format: Biscuit

Biscuit [BISCUIT] is an append-only capability token format using Ed25519 signatures and Datalog for policy expression. Biscuit provides cryptographic attenuation: a holder can create a more restricted version of a token by appending a new block signed with a fresh key, without interaction with the original issuer.

Implementations MAY support Biscuit where cryptographic attenuation without issuer interaction is required. Biscuit's append-only structure enforces Invariant 1 (Scope Monotonicity) cryptographically at the format layer.

When using Biscuit:

- The root block MUST contain all REQUIRED CT claims as Datalog facts.
- The human\_anchor and delegation\_depth MUST be in the root block and MUST NOT be modifiable by attenuation blocks.
- The Trust Registry MUST list the Ed25519 public key of the root block issuer for the relevant capability type.
- Step 1 of the enforcement engine MUST verify the root block signature using the Trust Registry key, and verify all attenuation block signatures form a valid chain.

### 3.5.3. Alternative Format: CWT (CBOR Web Token)

CWT [RFC8392] provides the same semantic structure as JWT in CBOR encoding, with COSE [RFC9052] for signing. CWT is RECOMMENDED for high-throughput agent-to-agent transaction contexts where binary encoding reduces token size and parsing overhead.

When using CWT:

- All REQUIRED CT claims MUST be present using their CBOR integer claim keys as defined in the IANA CBOR Web Token Claims Registry, or as defined in Section 11 for SPT-Txn-specific claims.
- The COSE\_Sign1 structure MUST be used for signing.
- Step 1 of the enforcement engine MUST verify using COSE signature verification procedures [RFC9052].

### 3.5.4. Format Selection Criteria

The following criteria SHOULD guide format selection:

JWT/SD-JWT: Cross-organizational deployments; human-readable debugging requirements; broad ecosystem library support; when GDPR data minimization is required (SD-JWT).

Biscuit: Deployments requiring cryptographic attenuation without issuer interaction; edge computing contexts; when Datalog policy expression is operationally preferred.

CWT: High-throughput agent-to-agent pipelines; IoT and

constrained environments; when binary encoding provides measurable operational benefit.

Implementations MUST NOT mix formats within a single delegation chain (e.g., a JWT root CT MUST NOT have a CWT child CT). The format is set at root issuance and MUST be consistent through the chain.

Coetzee	Expires October 27, 2026	[Page 9]
Internet-Draft	SPT-Txn Tokens	April 2026

## 4. SPT-Txn Token Structure

### 4.1. Base Claims (Inherited)

SPT-Txn Tokens inherit and MUST include all REQUIRED claims defined in RFC 9700 for Transaction Tokens:

iss: The authorization server that issued this token.  
iat: Issued-at time.  
exp: Expiry time.  
txn: Transaction identifier from RFC 9700.  
sub: The workload identity of the calling service at this hop.  
azp: The authorized party (original requesting workload).  
rctx: The request context as defined in RFC 9700.

### 4.2. SPT-Txn Extension Claims

SPT-Txn Tokens MUST include the following additional claims:

spt\_ct\_ref: bytes32. The SHA-256 hash of the Capability Token presented by the caller at this hop. Binds the transaction token to the specific capability grant used. REQUIRED.

spt\_ct\_scope\_hash: bytes32. The SHA-256 hash of the ct\_scope claim of the presented CT. Used by the enforcement engine in Step 8 for scope binding verification. REQUIRED.

spt\_human\_anchor: bytes32. The humanAnchor value from the presented CT. MUST equal ct.human\_anchor. Propagated immutably. REQUIRED.

spt\_delegation\_depth: uint8. The delegation\_depth value from the presented CT at this hop. REQUIRED.

spt\_compliance\_ref: bytes32. The compliance\_ref from the root CT. MUST be identical at every hop. REQUIRED for regulated transaction contexts.

spt\_parent\_txn: The jti of the parent SPT-Txn Token in the chain. OPTIONAL for root tokens. REQUIRED for all non-root tokens.

Example SPT-Txn Token extension claims:

```
{
  "spt_ct_ref":
    "0x7f3a9b2c4d1e8f6ab5c32e9d1a7b4f8e9c0d1e2f...",
  "spt_ct_scope_hash":
    "0x1a2b3c4d5e6f7a8b9c0d1e2f3a4b5c6d7e8f9a0b...",
  "spt_human_anchor":
    "0x7f3a9b2c4d1e8f6ab5c32e9d1a7b4f8e9c0d1e2f..."
}
```

```

"spt_delegation_depth": 2,
"spt_compliance_ref":
  "0x9a8b7c6d5e4f3a2b1c0d9e8f7a6b5c4d3e2f1a0b...",
"spt_parent_txn": "ct-4a5b6c7d-8e9f-0a1b-2c3d-4e5f6a7b8c9d"
}

```

Coetzee	Expires October 27, 2026	[Page 10]
Internet-Draft	SPT-Txn Tokens	April 2026

## 5. ZK Circuit Specification

### 5.1. Proof System

SPT-Txn uses Groth16 [GROTH16] over the BN254 pairing-friendly elliptic curve. Proof generation is performed by the ABAC PDP at root transaction time. Proof verification is performed by resource servers and auditors.

The proof system provides:

- Completeness: Valid witnesses always produce valid proofs.
- Soundness: Invalid witnesses cannot produce valid proofs (computationally, under the BN254 discrete log assumption).
- Zero-knowledge: Proofs reveal nothing about witnesses beyond the verified public inputs.

Circuit parameters MUST be generated through a trusted setup ceremony. The SPT-Txn reference implementation uses the Hermez ceremony outputs for BN254 [HERMEZ].

### 5.2. Circuit 1: Attribute Proof

The Attribute Proof circuit proves that a principal holds specified attributes issued by a trusted issuer, without revealing which attributes or their values.

Public inputs:

```

attr_root: Merkle root of the attribute set.
issuer_pk: Public key of the trusted attribute issuer.
attr_type_commitment: Commitment to the attribute type
                     being proved.

```

Private inputs (witness):

```

attr_value: The actual attribute value.
attr_merkle_path: Merkle path proving inclusion.
issuer_signature: Issuer's signature over the attribute.
nullifier_key: Holder's nullifier key.

```

Outputs:

```

nullifier: Hash(nullifier_key, attr_root). Prevents
          double-use of the same attribute proof.
proof: Groth16 proof of the above relation.

```

The circuit proves:

1. attr\_value is included in the Merkle tree with root attr\_root (via attr\_merkle\_path).
2. issuer\_signature is a valid signature over (attr\_type, attr\_value) under issuer\_pk.
3. nullifier == Hash(nullifier\_key, attr\_root).

Coetzee	Expires October 27, 2026	[Page 11]
Internet-Draft	SPT-Txn Tokens	April 2026

### 5.3. Circuit 2: Compliance Claim

The Compliance Claim circuit proves that a principal satisfies a compliance requirement (KYC, AML, jurisdiction) without revealing the underlying compliance data.

Public inputs:

- compliance\_type: URI identifying the compliance requirement.
- compliance\_issuer\_pk: Public key of the compliance certifier.
- jurisdiction\_set\_root: Merkle root of permitted jurisdictions.
- timestamp\_bound: Maximum age of the compliance attestation.

Private inputs (witness):

- compliance\_credential: The underlying compliance credential.
- issuer\_signature: Certifier's signature.
- jurisdiction\_merkle\_path: Proof of jurisdiction inclusion.
- credential\_timestamp: Timestamp of the credential.

Outputs:

- compliance\_commitment: Hash(compliance\_credential, compliance\_type). Used as compliance\_ref in CTs.
- proof: Groth16 proof of the above relation.

The circuit proves:

1. compliance\_credential is validly signed by compliance\_issuer\_pk.
2. The jurisdiction in compliance\_credential is included in the set with root jurisdiction\_set\_root.
3. credential\_timestamp > (now - timestamp\_bound).
4. compliance\_commitment == Hash(compliance\_credential, compliance\_type).

### 5.4. Circuit 3: Biometric Commitment

The Biometric Commitment circuit proves that a biometric measurement matches a committed template, producing a zkDID commitment without revealing the biometric data.

Public inputs:

- template\_commitment: Pedersen commitment to the biometric template.
- liveness\_proof\_root: Root of the liveness attestation tree.

Private inputs (witness):

- biometric\_measurement: The raw biometric measurement.
- biometric\_template: The enrolled template.
- template\_randomness: Randomness in the commitment.
- liveness\_attestation: Signed liveness attestation.
- liveness\_merkle\_path: Merkle path in attestation tree.

Outputs:

- zkdid\_commitment: Poseidon(biometric\_template, template\_randomness). This is the humanAnchor.
- proof: Groth16 proof of the above relation.

The circuit proves:

1. biometric\_measurement matches biometric\_template within a defined distance threshold.
2. template\_commitment == Pedersen(biometric\_template, template\_randomness).
3. liveness\_attestation is valid and included in the liveness attestation tree.
4. zkdid\_commitment == Poseidon(biometric\_template, template\_randomness).

## 6. Multi-Chain Trust Registry Model

### 6.1. Chain-Agnostic Trust Registry

The Trust Registry is an on-chain data store mapping issuer identities to the capability types they are authorized to certify. It is the root of trust for Step 2 of the enforcement engine.

Trust Registry entries MUST contain:

`issuer_id`: The DID or HTTPS identifier of the ABAC PDP.  
`capability_types`: A list of capability type URIs the issuer is authorized to certify.  
`issuer_pubkey`: The current public key of the issuer, used for CT signature verification.  
`valid_from`: Unix timestamp from which this entry is valid.  
`valid_until`: Unix timestamp at which this entry expires.  
`status`: ACTIVE, SUSPENDED, or REVOKED.

Trust Registry entries MUST be signed by a Trust Anchor -- a higher-level authority that has authorized the issuer to operate for the specified capability types. The Trust Anchor public key MUST be distributed out-of-band at deployment time.

Resource servers MUST cache Trust Registry entries for performance. Cache TTL SHOULD be configured based on the `valid_until` field minus a safety margin. Implementations MUST treat expired Trust Registry entries as equivalent to revoked.

### 6.2. Chain-Specific Deployment Notes

#### 6.2.1. Ethereum / EVM-Compatible Chains

Trust Registry SHOULD be implemented as an upgradeable proxy smart contract using the EIP-1967 transparent proxy pattern. Access control for Trust Registry updates MUST use a multi-signature scheme with a minimum threshold of 3-of-5.

Event logs (TrustRegistryEntry events) MUST be emitted for all Trust Registry updates to support auditability.

Gas optimization: Trust Registry reads are view functions and consume no gas for read-only verification.

#### 6.2.2. XDC Network

XDC provides EVM compatibility with Delegated Proof of Stake consensus and native support for trade finance applications. Trust Registry deployment on XDC follows the same EVM contract specification as Section 6.2.1.

XDC's 2-second block finality provides faster Trust Registry update propagation than Ethereum mainnet. Implementations SHOULD use XDC for deployments with high Trust Registry update frequency.

#### 6.2.3. Algorand

Algorand's AVM (Algorand Virtual Machine) supports Trust Registry implementation as an Algorand Smart Contract (ASC1). State storage uses global state for Trust Registry entries.

Algorand's 3.9-second block finality and immediate transaction finality (no confirmation delay) make it suitable for Trust Registry deployments requiring low-latency updates.

Algorand Standard Assets (ASAs) MAY be used to represent Trust Registry authority grants in a transferable format.

#### 6.2.4. Hedera Hashgraph

Hedera Consensus Service (HCS) provides a tamper-evident, timestamped message log suitable for Trust Registry audit trails. Trust Registry state SHOULD be maintained in Hedera Smart Contracts (compatible with Solidity/EVM), with all state-changing operations also submitted to an HCS topic for auditability.

HCS provides:

- Cryptographic proof of message ordering and timestamp.
- Gossip-about-gossip consensus with Byzantine fault tolerance.
- GDPR-compliant auditability (HCS messages are public but can reference encrypted off-chain data by hash).

The HCS audit topic MUST record:

- All Trust Registry entry additions and modifications.
- All CT issuance events (by hash reference, not content).
- All revocation events.

Coetzee	Expires October 27, 2026	[Page 13]
Internet-Draft	SPT-Txn Tokens	April 2026

#### 7. Transaction Flow

The following illustrates a complete SPT-Txn transaction flow for a regulated financial agentic system.

Step 1: Human Authentication and Root CT Issuance

Human H presents to ABAC PDP:

- ZK proof of biometric commitment (liveness attested)
- ZK compliance proofs (KYC, AML, jurisdiction)

ABAC PDP evaluates against Policy NFT and issues:

```
CT_H: {
  iss: "did:web:abac-pdp.org-a.example",
  sub: H.zkDID,
  ct_type: "urn:example:capability:financial-transfer",
  ct_scope: { full authorized scope },
  human_anchor: H.zkDID_commitment,
  delegation_depth: 3,
  max_depth: 3,
  compliance_ref: hash(ZK_compliance_proofs)
}
```

```
Txn-Token-root: {
  (RFC 9700 base claims),
  spt_ct_ref: hash(CT_H),
  spt_ct_scope_hash: hash(CT_H.ct_scope),
  spt_human_anchor: H.zkDID_commitment,
```



```

    spt_delegation_depth: 3,
    spt_compliance_ref: hash(ZK_compliance_proofs)
}

```

Step 2: Human Delegates to AI Agent A1

H requests a delegation CT for A1 from ABAC PDP:

```

CT_A1: {
  sub: A1.workload_identity,
  human_anchor: H.zkDID_commitment,  <- unchanged
  ct_scope: { reduced_scope },        <- subset of CT_H
  delegation_depth: 2,                <- decremented
  parent_ct: hash(CT_H)
}

```

H provides A1 with CT\_A1 and an updated Txn-Token identifying A1 as the next hop workload.

Step 3: A1 Calls Service S1 (Organization B)

A1 requests a delegation CT for S1:

```

CT_S1: {
  sub: S1.workload_identity,
  human_anchor: H.zkDID_commitment,  <- unchanged
  ct_scope: { further_reduced_scope },
  delegation_depth: 1,
  parent_ct: hash(CT_A1)
}

```

A1 presents CT\_S1 and SPT-Txn Token to S1.

Step 4: S1 Verifies CT\_S1 (Offline, No Issuer Contact)

S1 runs the eight-step enforcement engine:

- (a) JWS signature valid per Trust Registry key for Org A.
- (b) Org A's ABAC PDP authorized for this capability type.
- (c) CT\_S1.exp > now.
- (d) CT\_S1.jti not in revocation registry.
- (e) Requested operation within CT\_S1.ct\_scope.
- (f) CT\_S1.delegation\_depth (1) < CT\_A1.delegation\_depth (2).
- (g) CT\_S1.human\_anchor == Txn-Token.spt\_human\_anchor.
- (h) hash(CT\_S1.ct\_scope) == Txn-Token.spt\_ct\_scope\_hash.

All checks pass. S1 processes the request.

No call to Org A's authorization server required.

Step 5: S1 Calls S2 (Organization C)

S1 requests a leaf CT for S2:

```

CT_S2: {
  sub: S2.workload_identity,
  human_anchor: H.zkDID_commitment,
  ct_scope: { minimal_leaf_scope },
  delegation_depth: 0,  <- cannot delegate further
  parent_ct: hash(CT_S1)
}

```

Audit Trail (verifiable by any authorized auditor):

Txn-Token chain: H -> A1 -> S1 -> S2

Each hop: CT reference, scope hash, depth, compliance ref.

Any party can verify: human authorized root, scope only reduced at each step, compliance claims valid from root,

chain of custody intact.  
No PII revealed to intermediate parties or auditors.

Coetzee	Expires October 27, 2026	[Page 14]
Internet-Draft	SPT-Txn Tokens	April 2026

## 8. Security Considerations

### 8.1. Capability Token Security

CT Signature Verification: All CT signatures MUST be verified against the issuer's public key retrieved from the Trust Registry. Implementations MUST NOT accept issuer public keys embedded in the token itself or retrieved from a URL in the token. This prevents key substitution attacks.

CT Expiry: CTs MUST have a finite exp claim. Implementations MUST NOT accept CTs with exp in the past. The recommended maximum CT lifetime is 1 hour for transactional CTs. Long-lived CTs increase revocation exposure.

JTI Uniqueness: The jti claim MUST be globally unique across all CTs issued by a given issuer. Implementations MUST reject CTs with a jti that has been seen before (replay prevention).

CT Scope Precision: CT scopes MUST be as narrow as possible while permitting the required operations. Overly broad scopes reduce the security benefit of TBAC enforcement. Issuers MUST validate scope requests against the requesting principal's ABAC policy before issuing.

### 8.2. Delegation Chain Security

Scope Monotonicity Enforcement: Issuers MUST verify Invariant 1 (Section 3.4) before issuing delegation CTs. Implementations MUST reject delegation requests where the requested scope is not a strict subset of the delegator's CT scope.

Depth Enforcement: Issuers MUST verify that the delegator's CT has `delegation_depth > 0` before issuing a delegation CT. Implementations MUST set the delegatee's `delegation_depth` to `delegator.delegation_depth - 1`.

Human Anchor Immutability: Issuers MUST copy the `human_anchor` from the parent CT without modification. Any CT presenting a `human_anchor` different from its parent MUST be rejected.

Parent CT Validity: Before issuing a delegation CT, the issuer MUST verify that the parent CT passes all eight enforcement steps (Section 3.3). An issuer MUST NOT issue delegation CTs on behalf of an expired, revoked, or out-of-scope parent CT.

### 8.3. ZK Proof Security

Trusted Setup: The Groth16 proof system requires a trusted setup. The toxic waste from the setup MUST be destroyed. Multi-party computation ceremonies SHOULD be used to distribute trust. The reference implementation uses the Hermez ceremony.

Proof Replay: ZK proofs are publicly verifiable. The nullifier mechanism in Circuit 1 MUST be used to prevent proof replay. Implementations MUST maintain a nullifier registry and reject proofs with previously seen nullifiers.

Circuit Audit: ZK circuits MUST be independently audited before production deployment. The circuit specification in Section 5 is normative; implementations MUST match the specified input/output structure exactly.

#### 8.4. Trust Registry Security

Registry Immutability: Trust Registry updates MUST require multi-signature authorization. Single-key control of the Trust Registry is a critical security risk.

Registry Monitoring: Implementations MUST monitor Trust Registry events for unexpected entry additions or modifications. An unexpected Trust Registry update MAY indicate a key compromise at the Trust Anchor level.

Entry Expiry: Trust Registry entries MUST have a `valid_until` field. Perpetual Trust Registry entries create long-term exposure if the associated issuer is later compromised.

Coetzee

Expires October 27, 2026

[Page 15]

Internet-Draft

SPT-Txn Tokens

April 2026

### 9. Privacy Considerations

#### 9.1. Human Anchor Privacy

The `humanAnchor` is a zero-knowledge commitment to the authorizing human's zkDID. It does not reveal the human's identity, biometric data, or any other PII to intermediate services. Only the licensed issuer holding the escrow key can deanonymize the `humanAnchor` under lawful process. The operational requirements of the escrow mechanism are specified in Section 9.6.

#### 9.2. Compliance Reference Privacy

The `compliance_ref` is a hash reference to ZK compliance proofs established at root. The proofs themselves are not propagated in the transaction chain. Resource servers verify that `compliance_ref` matches a known proof commitment; they do not receive or store the underlying compliance data.

#### 9.3. Selective Disclosure

When SD-JWT format (Section 3.5.1) is used, CT holders MAY limit claim disclosure to the specific claims required by the verifying resource server. This reduces the PII surface per hop. The set of claims that MUST always be disclosed is enumerated in Section 3.5.1.

#### 9.4. Auditability vs. Privacy

The SPT-Txn audit chain (Section 7) is designed to be auditable by authorized auditors without revealing PII. The Txn-Token chain carries `human_anchor` (ZK commitment), `spt_ct_ref` (token hash), and `spt_compliance_ref` (proof hash) -- none of which reveal PII directly. Identity deanonymization requires invocation of the escrow mechanism specified in Section 9.6 and is subject to the lawful process requirements defined therein.

#### 9.5. GDPR Compliance

The `humanAnchor` mechanism is designed to satisfy GDPR Article 25

(Data Protection by Design) and Article 5(1)(c) (Data Minimization). Intermediate services receive only cryptographic commitments, not personal data. Implementations operating in EU jurisdictions SHOULD obtain legal review of their specific humanAnchor escrow arrangement against the requirements of Section 9.6 and applicable regional law.

## 9.6. Human Anchor Escrow

### 9.6.1. Escrow Function and Position in the Architecture

The Human Anchor Escrow is the cryptographic and operational mechanism by which the holder of an authorized escrow capability, acting under lawful process, recovers the zkDID of a human principal from a humanAnchor commitment present in a Capability Token or SPT-Txn Token.

Coetzee	Expires October 27, 2026	[Page 16]
Internet-Draft	SPT-Txn Tokens	April 2026

The escrow is OUT OF BAND with respect to all token verification operations. No step of the eight-step enforcement engine (Section 3.3), no Capability Token issuance step, and no SPT-Txn Token issuance step interacts with the escrow. The escrow is invoked solely in response to a Deanonymization Request (Section 9.6.5) and does not appear in the request-time critical path of any resource server.

At root Capability Token issuance, the ABAC PDP MUST construct an Escrow Envelope binding the humanAnchor to material sufficient to recover the zkDID. The Escrow Envelope MUST be retained by the ABAC PDP indexed by human\_anchor. The Escrow Envelope MUST NOT be transmitted in any token, included in any audit chain entry visible to non-escrow parties, or replicated to any party other than the ABAC PDP and (at the PDP's discretion) escrow custody backups under equivalent protection.

### 9.6.2. Escrow Envelope Construction

The Escrow Envelope is an authenticated ciphertext encrypting the zkDID under the Escrow Public Key (Section 9.6.3). The ABAC PDP MUST construct the envelope as follows:

```
envelope = AuthEnc(K_escrow_pub, plaintext, AAD)
```

where:

- plaintext is the binary encoding of the zkDID concatenated with a unique nonce.
- AAD is the binary concatenation of the humanAnchor, the issuer identifier (iss), and the issued-at timestamp (iat) of the root CT.
- AuthEnc is an authenticated public-key encryption scheme satisfying IND-CCA2 against active adversaries, such as ECIES with HKDF-SHA-256 key derivation and AES-256-GCM, or equivalent. Implementations MUST NOT use unauthenticated or non-CCA-secure schemes.

Implementations MUST verify that decryption of the envelope yields a plaintext whose first component, when committed via the Section 5.4 circuit's commitment function, equals the humanAnchor in the AAD. This binding prevents an adversary in possession of an Escrow Private Key share threshold from substituting an envelope and obtaining a

zkDID for a humanAnchor it did not authorize.

### 9.6.3. Escrow Key Material Requirements

The Escrow Public Key is the public component of an asymmetric keypair used to construct Escrow Envelopes. Implementations MUST satisfy the following requirements:

Coetzee Expires October 27, 2026 [Page 17]

Internet-Draft SPT-Txn Tokens April 2026

ESC-1. The Escrow Private Key MUST NOT exist in unsplit form on any single system at any time after generation, except transiently during distributed key generation if a non-DKG construction is used.

ESC-2. Implementations SHOULD use a Distributed Key Generation (DKG) protocol such that no single party ever observes the full Escrow Private Key. Implementations using Shamir Secret Sharing of a centrally generated key MUST destroy the centrally generated key immediately after share distribution and MUST document the generation ceremony in the deployment policy (Section 9.6.7).

ESC-3. The Escrow Public Key MUST be published in the Trust Registry (Section 6) as a separate registry entry distinct from the ABAC PDP's CT signing key. The entry MUST identify its purpose as escrow encryption and MUST NOT be accepted by resource servers for CT signature verification.

ESC-4. The Escrow keypair MUST use an algorithm whose security is not weaker than the CT signing algorithm requirements of Section 10.1. X25519 with HKDF-SHA-256 and AES-256-GCM is RECOMMENDED. P-256 with the equivalent ECIES construction is permitted. RSA key transport MUST NOT be used.

ESC-5. Escrow keypair material MUST be stored in hardware security modules or equivalent tamper-resistant hardware on every escrow custodian's system. Software-only escrow share storage MUST NOT be used in production deployments.

### 9.6.4. Threshold Authorization Requirements

Decryption of an Escrow Envelope MUST require the cooperation of a quorum of Escrow Custodians. Implementations MUST satisfy:

THR-1. The deployment MUST define a threshold  $(t, n)$  where  $n$  is the number of Escrow Custodians and  $t$  is the minimum number whose participation is required to decrypt an envelope. Implementations MUST configure  $t$  such that  $t$  is greater than or equal to  $\text{ceil}((n+1)/2)$ . Deployments handling regulated financial transactions SHOULD configure  $t$  such that  $t$  is greater than or equal to  $\text{ceil}(2n/3)$ .

THR-2. The threshold scheme MUST be implemented such that no proper subset of fewer than  $t$  custodians can produce a valid decryption, recover any portion of the Escrow Private Key, or recover any portion of the plaintext. Schemes failing to provide threshold security against  $t-1$  colluding custodians MUST NOT be used.

THR-3. Each individual Escrow Custodian MUST be a distinct legal entity with an independent operational, technical, and legal control surface. Multiple custodians under common control of a single legal entity MUST NOT be counted as separate participants

for threshold purposes.

Coetzee	Expires October 27, 2026	[Page 18]
Internet-Draft	SPT-Txn Tokens	April 2026

THR-4. The deployment policy MUST document the identities of all Escrow Custodians, the threshold (t, n) in effect, and the legal basis under which each custodian operates. The policy MUST be updated within thirty (30) days of any change to custodian composition.

THR-5. Threshold ElGamal, threshold ECIES, or equivalent threshold decryption protocols providing the property in THR-2 are RECOMMENDED. Schemes that reconstruct the Escrow Private Key in a single location during decryption (including Shamir reconstruction) provide weaker security and SHOULD be avoided in deployments handling regulated transactions.

#### 9.6.5. Deanonymization Request Interface

A Deanonymization Request is the protocol-level invocation by which an authorized requestor initiates escrow decryption of a specific humanAnchor. The request is processed by the ABAC PDP, which mediates between the requestor and the Escrow Custodians. The request format MUST include:

request\_id: A globally unique identifier for the request.

target\_human\_anchor: The humanAnchor whose deanonymization is requested.

requestor\_id: The identifier of the requesting authority.

lawful\_basis\_ref: A reference (URI or hash) to the legal instrument authorizing the request (court order, regulatory production order, or equivalent).

requestor\_signature: A signature over the preceding fields using a key listed in the Trust Registry as authorized for escrow request issuance, distinct from any key used for CT issuance or verification.

The ABAC PDP MUST verify the requestor's signature against the Trust Registry, MUST verify that the signing key is currently authorized for escrow request issuance, MUST locate the Escrow Envelope corresponding to target\_human\_anchor, and MUST forward the envelope and the request metadata to the Escrow Custodians for threshold decryption.

Each Escrow Custodian MUST independently evaluate the deanonymization request against its custodian policy before participating in decryption. Custodian policy SHOULD include verification of the lawful basis reference, applicable jurisdictional review, and any internal authorization workflow. Custodians MUST NOT be required by protocol to participate; threshold non-participation by custodians who determine the request is not authorized is the protocol's intended behavior, not a failure mode.

Coetzee	Expires October 27, 2026	[Page 19]
Internet-Draft	SPT-Txn Tokens	April 2026

If the threshold of participating custodians is reached, the resulting plaintext zkDID MUST be returned only to the requestor identified in requestor\_id and MUST NOT be retained by the ABAC PDP after delivery.

The protocol does not define what constitutes a valid lawful\_basis\_ref. Determining the legal sufficiency of a request is a matter of jurisdictional law and custodian policy, not protocol specification.

#### 9.6.6. Escrow Audit Log

The ABAC PDP MUST maintain an Escrow Audit Log recording every Deanonymization Request received, regardless of whether the request resulted in successful decryption. Each log entry MUST contain request\_id, target\_human\_anchor, requestor\_id, lawful\_basis\_ref, the request timestamp, the participating custodian set if decryption proceeded, and the outcome (decryption\_completed, threshold\_not\_reached, request\_rejected, or custodian\_policy\_failure).

The Escrow Audit Log MUST NOT contain the decrypted zkDID or any portion thereof.

The Escrow Audit Log MUST be append-only. Each entry MUST be cryptographically chained to the previous entry such that any modification or deletion of historical entries is detectable. The chained log root SHOULD be published periodically to the Trust Registry chain (Section 6) or to an equivalent public, append-only medium, with a minimum publication frequency of once per twenty-four (24) hours when entries have been added.

The Escrow Audit Log SHOULD be accessible to the human principal corresponding to a humanAnchor, upon presentation of a zero-knowledge proof of correspondence to that humanAnchor, such that the human can verify whether their humanAnchor has been the subject of a deanonymization request. The protocol does not require this access channel; deployments operating under jurisdictions that mandate notice to the data subject SHOULD provide it.

#### 9.6.7. Escrow Key Lifecycle

Escrow keypairs MUST be subject to a documented lifecycle covering generation, rotation, retirement, and destruction. Implementations MUST satisfy the following requirements:

LIFE-1. Escrow Public Keys MUST have a documented validity period reflected in the Trust Registry entry's valid\_from and valid\_until fields. CTs issued during a key's validity period MUST have their Escrow Envelopes constructed under that key. Envelope material MUST be retained by the ABAC PDP for at least the operational lifetime of the longest-lived CT issued under that key, plus the maximum statute-of-limitations period applicable to any transaction the CT may have authorized.

LIFE-2. Escrow keypair generation MUST occur in a documented ceremony attended by all Escrow Custodians or their authorized delegates. The ceremony MUST produce a public artifact attesting to the generation: the Escrow Public Key, the threshold parameters (t, n), the custodian identities, and the ceremony timestamp,

signed by a quorum of custodians.

LIFE-3. Escrow keys SHOULD be rotated on a schedule documented in the deployment policy. Rotation MUST follow a key-overlap procedure where the new keypair is generated, published to the Trust Registry, and used for all new envelopes from a defined cutover date, while the previous keypair remains available for decryption of envelopes constructed under it.

LIFE-4. Retirement of an Escrow keypair MUST NOT occur while any Escrow Envelope constructed under that key is within its required retention period (LIFE-1). Premature retirement renders the corresponding humanAnchors permanently un-deanonymizable, which MAY be desirable in specific privacy-by-design deployments but MUST be an explicit deployment policy decision, not an operational accident.

#### 9.6.8. Compromise and Recovery

This subsection specifies the protocol-level response to compromise events affecting the escrow.

Compromise of fewer than  $t$  Escrow Custodians is below the threshold and does not directly enable unauthorized deanonymization. The ABAC PDP MUST nonetheless treat any confirmed custodian compromise as an event requiring escrow keypair rotation per LIFE-3, with a cutover date no later than thirty (30) days after compromise confirmation. The compromised custodian MUST be excluded from the new keypair's custody set.

Compromise of  $t$  or more Escrow Custodians constitutes an Escrow Compromise. On confirmation of an Escrow Compromise, the ABAC PDP MUST:

- Mark the Trust Registry entry for the compromised Escrow Public Key as REVOKED.
- Cease accepting new Capability Token issuance requests under the compromised escrow until a replacement keypair is generated and published.
- Publish a notice to the Trust Registry chain or equivalent public medium identifying the compromised key, the time window of compromise, and the expected replacement schedule.

Capability Tokens whose Escrow Envelopes were constructed under a compromised key are not automatically invalidated for transaction processing purposes -- the token signature, scope, depth, and human anchor binding remain cryptographically intact. The compromise affects only the escrow recoverability of the human identity, not the integrity of the authorization chain. Resource servers SHOULD continue to honor non-revoked CTs issued before compromise; deployments MAY define a stricter policy in their deployment documentation.

Loss of  $t$  or more Escrow Custodians without compromise -- for example through simultaneous key destruction events -- produces a permanent inability to deanonymize humanAnchors constructed under the affected keypair. This is a deployment failure mode that MUST be addressed through custodian redundancy and documented disaster recovery



procedures. The protocol does not provide a mechanism to recover from total custody loss; this is a deliberate design property reflecting the principle that the human's identity is protected even against the operator of the system.

#### 9.6.9. Implementation Note (Non-Normative)

Two threshold cryptography constructions are particularly suitable for the requirements above. Threshold ElGamal over Curve25519 with a published reference such as the FROST family of protocols provides distributed key generation and threshold decryption without ever reconstructing the private key, satisfying THR-2 and THR-5 in their strongest form. Threshold ECIES with the same primitives is operationally equivalent. Implementations choosing Shamir Secret Sharing of a centrally generated X25519 private key satisfy THR-1 through THR-4 but provide weaker security at the moment of decryption, when shares are reconstructed in one location; this construction is permitted but discouraged for regulated financial deployments.

The Escrow Custodian role is operationally distinct from the ABAC PDP role. A deployment in which the ABAC PDP is also an Escrow Custodian satisfies the protocol but reduces the trust separation that the escrow is designed to provide. Deployments are encouraged to recruit Escrow Custodians from distinct legal jurisdictions and distinct regulatory regimes such that no single jurisdiction's lawful process can compel decryption unilaterally. This composition decision is a matter of deployment policy and is not specified by the protocol.

Coetzee	Expires October 27, 2026	[Page 22]
Internet-Draft	SPT-Txn Tokens	April 2026

## 10. Security Properties

### 10.1. Signature Algorithm Requirements

All CT signatures MUST use one of the following algorithms:

ES256: ECDSA with P-256 and SHA-256. REQUIRED to support.

ES384: ECDSA with P-384 and SHA-384. RECOMMENDED.

EdDSA (Ed25519): Edwards-curve DSA. RECOMMENDED.

Ed25519 is RECOMMENDED for new deployments due to its constant-time implementation properties and smaller signature size.

RS256 and weaker algorithms MUST NOT be used for CT signing.

Post-quantum note: ECDSA and EdDSA are vulnerable to quantum adversaries with access to Shor's algorithm. Deployments with long-term security requirements (>10 years) SHOULD monitor NIST post-quantum standardization and plan migration to PQ-safe signature schemes. The SPT-Txn claim structure is algorithm-agnostic and does not require changes to accommodate algorithm migration.

### 10.2. Key Management

ABAC PDP signing keys MUST be stored in hardware security modules (HSMs) or equivalent tamper-resistant hardware.

Key rotation MUST be supported. Trust Registry entries MUST be updated when issuer keys are rotated. CTs signed with rotated-out keys MUST be considered expired at rotation time

unless a key overlap period is explicitly configured.

Key overlap period SHOULD NOT exceed the maximum CT lifetime to bound exposure of the outgoing key.

### 10.3. Transport Security

All SPT-Txn Token and CT transmission MUST occur over TLS 1.3 [RFC8446] or later. TLS 1.2 MAY be permitted for legacy compatibility with explicit operator acknowledgment.

CT transmission in HTTP headers MUST use the Authorization header with a custom scheme:

Authorization: SPT-CT <base64url-encoded-CT>

### 10.4. Replay Prevention

CTs include a jti claim for unique identification. Resource servers MUST maintain a short-term jti cache to detect CT replay within the CT's validity window.

SPT-Txn Tokens inherit the replay prevention mechanisms defined in RFC 9700, including the txn claim binding.

### 10.5. Token Format Security Properties

This section specifies security properties specific to each token serialization format (Section 3.5).

#### JWT/SD-JWT Security:

JWS algorithm confusion attacks MUST be prevented by explicitly specifying the expected algorithm at verification time. Implementations MUST NOT accept the "none" algorithm.

SD-JWT holder binding MUST use Ed25519 or P-256 key binding. Unbound SD-JWT presentations MUST NOT be accepted for CT presentation.

#### Biscuit Security:

Biscuit's append-only structure provides structural defence against scope escalation: attenuation blocks can only restrict, never expand, the token's authority. This cryptographically enforces Invariant 1 (Scope Monotonicity) at the format layer, independent of the enforcement engine.

The root block's Ed25519 key MUST be registered in the Trust Registry. Attenuation block keys are ephemeral and MUST NOT be registered.

Biscuit tokens MUST be treated as opaque by intermediate services that are not the intended verifier.

#### CWT/COSE Security:

COSE\_Sign1 MUST be used (not COSE\_Sign for multi-signer). Algorithm parameters MUST be in the protected header. Unprotected header algorithm parameters MUST be ignored.

The COSE kid (key ID) parameter MUST NOT be used to select verification keys. Verification keys MUST be retrieved from the Trust Registry based on the iss claim.

### 10.6. ZK Proof System Security

See Section 8.3 for ZK proof security considerations.

Additionally:

- Groth16 proofs are computationally sound under the BN254 discrete logarithm assumption.
- Proof size is constant (192 bytes for Groth16 on BN254) regardless of circuit complexity.
- Verification time is constant and fast (< 5ms on commodity hardware), making ZK proof verification practical in the enforcement engine hot path.

#### 10.7. Multi-Chain Trust Registry Integrity

See Section 8.4 for Trust Registry security considerations.

Cross-chain deployments: When Trust Registries are deployed on multiple chains, implementations MUST ensure consistency between chain deployments. A Trust Registry entry on one chain MUST NOT be considered authoritative for resource servers configured to use a different chain.

Chain reorganizations: For EVM chains, Trust Registry reads SHOULD wait for a minimum of 12 block confirmations to protect against reorganization-based attacks on Trust Registry state. Hedera and Algorand provide immediate finality and do not require confirmation delays.

Coetzee	Expires October 27, 2026	[Page 23]
Internet-Draft	SPT-Txn Tokens	April 2026

#### 11. IANA Considerations

This document requests registration of the following JWT claims in the IANA JSON Web Token Claims Registry [IANA.JWT.Claims].

Claim Name: spt\_ct\_ref  
Claim Description: SHA-256 hash of the Capability Token presented at this transaction hop.  
Change Controller: IETF  
Specification Document(s): Section 4.2 of this document.

Claim Name: spt\_ct\_scope\_hash  
Claim Description: SHA-256 hash of the ct\_scope claim of the presented Capability Token.  
Change Controller: IETF  
Specification Document(s): Section 4.2 of this document.

Claim Name: spt\_human\_anchor  
Claim Description: Zero-knowledge commitment to the zkDID of the authorizing human principal (humanAnchor).  
Change Controller: IETF  
Specification Document(s): Sections 3.2 and 4.2 of this document.

Claim Name: spt\_delegation\_depth  
Claim Description: Remaining delegation depth from the Capability Token at this transaction hop.  
Change Controller: IETF  
Specification Document(s): Section 4.2 of this document.

Claim Name: spt\_compliance\_ref  
Claim Description: Hash reference to the ZK compliance proofs established at transaction root.

Change Controller: IETF  
Specification Document(s): Sections 3.2 and 4.2 of this document.

Claim Name: spt\_parent\_txn  
Claim Description: JWT ID of the parent SPT-Txn Token in the transaction chain.  
Change Controller: IETF  
Specification Document(s): Section 4.2 of this document.

This document also requests registration of the following Capability Token claims (in a new SPT-Txn CT Claims Registry to be established by IANA, or in the JWT Claims Registry):

Claim Name: ct\_type  
Claim Description: URI identifying the capability type granted by this Capability Token.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Claim Name: ct\_scope  
Claim Description: Structured object defining the capability scope, operations, and constraints granted.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Claim Name: human\_anchor  
Claim Description: Zero-knowledge commitment to the zkDID of the authorizing human (humanAnchor), carried in the CT.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Claim Name: delegation\_depth  
Claim Description: Remaining delegation levels permitted from this CT holder.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Claim Name: compliance\_ref  
Claim Description: Hash reference to ZK compliance proofs at transaction root.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Claim Name: revocation\_nonce  
Claim Description: Nonce included in the revocation registry commitment for this CT.  
Change Controller: IETF  
Specification Document(s): Section 3.2 of this document.

Coetzee	Expires October 27, 2026	[Page 24]
Internet-Draft	SPT-Txn Tokens	April 2026

## 12. References

### 12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web

- Signature (JWS)", RFC 7515,  
DOI 10.17487/RFC7515, May 2015,  
<<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519,  
DOI 10.17487/RFC7519, May 2015,  
<<https://www.rfc-editor.org/info/rfc7519>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174,  
DOI 10.17487/RFC8174, May 2017,  
<<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8392] Jones, M., Wahlstroem, E., Erdtman, S., and H. Tschofenig, "CBOR Web Token (CWT)", RFC 8392,  
DOI 10.17487/RFC8392, March 2018,  
<<https://www.rfc-editor.org/info/rfc8392>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446,  
DOI 10.17487/RFC8446, August 2018,  
<<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC9052] Schaad, J., "CBOR Object Signing and Encryption (COSE): Structures and Process", STD 96, RFC 9052,  
DOI 10.17487/RFC9052, August 2022,  
<<https://www.rfc-editor.org/info/rfc9052>>.
- [RFC9449] Fett, D., Campbell, B., Bradley, J., Lodderstedt, T., Jones, M., and D. Waite, "OAuth 2.0 Demonstrating Proof of Possession (DPoP)", RFC 9449,  
DOI 10.17487/RFC9449, September 2023,  
<<https://www.rfc-editor.org/info/rfc9449>>.
- [RFC9700] Tulshibagwale, A., Fletcher, G., Kasselmann, P., Burgin, K., and D. Richer, "Transaction Tokens", RFC 9700, 2025,  
<<https://www.rfc-editor.org/info/rfc9700>>.
- [I-D.ietf-oauth-selective-disclosure-jwt]  
Fett, D., Yasuda, K., and B. Campbell,  
"Selective Disclosure for JWTs (SD-JWT)", Work in Progress, Internet-Draft,  
draft-ietf-oauth-selective-disclosure-jwt-13,  
2024,  
<<https://datatracker.ietf.org/doc/html/draft-ietf-oauth-selective-disclosure-jwt>>.

## 12.2. Informative References

- [BISCUIT] Felte, C., et al., "Biscuit: A Capability-Based Authorization Token", 2022,  
<<https://github.com/biscuit-auth/biscuit>>.
- [GROTH16] Groth, J., "On the Size of Pairing-Based Non-interactive Arguments", Advances in Cryptology EUROCRYPT 2016, LNCS 9666, pp. 305-326, 2016,  
DOI 10.1007/978-3-662-49896-5\_11.

- [HERMEZ] Hermez Network, "Trusted Setup Ceremony",  
<<https://ceremony.hermez.io/>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749,  
DOI 10.17487/RFC6749, October 2012,  
<<https://www.rfc-editor.org/info/rfc6749>>.
- [SPT-THEORY] Coetzee, R.J., "Transaction Binding Security for Policy-Bound Authorization Tokens: Formal Definitions and Proofs for the SPT-Txn Framework", Zenodo,  
DOI 10.5281/zenodo.18917439, 2025,  
<<https://doi.org/10.5281/zenodo.18917439>>.
- [SPT-WP] Coetzee, R.J., "The SPT-Txn Framework v6.0: Sovereign Policy Token Transactions", Zenodo, 2025,  
<<https://doi.org/10.5281/zenodo.18917439>>.
- [IANA.JWT.Claims] IANA, "JSON Web Token Claims",  
<<https://www.iana.org/assignments/jwt>>.

#### Author's Address

Rudolf Jacobus Coetzee  
Violet Sky Security SEZC  
Cayman Islands

Email: [rudi@violetskeysecurity.com](mailto:rudi@violetskeysecurity.com)  
URI: <https://www.violetskeysecurity.com>