STRUCTURED FINANCE & DIGITAL CREDIT INFRASTRUCTURE

WERI*TAS*

A structured finance architecture for Sub-Saharan African credit,
with governance coordination mechanisms described in the WERITAS Whitepaper.

DID · VERIFIABLE CREDENTIALS · CREDIT TRUST GRAPH

Credit as a *Continuously Verified* Identity Graph

A Web5-aligned credit identity infrastructure for emerging markets where identity, not institution, is the root primitive of financial access.

Protocol-first. Institution-agnostic. User-sovereign.
Built for Kenya, engineered for regional scalability.

> *"The fundamental move is the separation of identity and data from applications. When a user owns their identity and their data, the application becomes a lens on that data not a prison for it."*
>
> Web5 Architectural Principle (Adapted)

PRIMITIVE	CORE INNOVATION	PHASE 1	STATUS
DID + VC	**Credit Trust Graph**	**Kenya**	**Architecture Review**
Identity Root	(CTG)	Live Deployment	Pre-Audit

Scope, Purpose & Reviewer Notice

PURPOSE. This paper presents the technical architecture of the WERITAS credit identity infrastructure, with specific emphasis on its alignment with Web5 design principles decentralised identifiers (DIDs), verifiable credentials (VCs), and the separation of identity and data from applications. It is intended as a technical reference document for architecture review, protocol-level collaboration discussions, and peer critique by engineers and researchers working in decentralised identity and credit systems.

WHAT THIS PAPER IS. A protocol-level description of how identity, credentials, and a credit trust computation layer compose into an institutionally integrable credit infrastructure. It addresses the technical design: data models, trust topology, consent flows, credential schemas, and the separation of concerns between identity, credential, computation, and settlement layers.

WHAT THIS PAPER IS NOT. This is not a token whitepaper, a business plan, an investment memorandum, or a legal document. Corporate structure, jurisdictional architecture, special-purpose vehicles, and token economics are addressed in separate documents principally WERITAS Whitepaper and are referenced here only where they are necessary to explain operational context. No financial projections or yield representations appear in this document.

WEB5 ALIGNMENT CLAIM. WERITAS is not a product of the TBD Web5 stack. It is a Web5-aligned implementation meaning it adopts the core architectural primitives (DIDs, VCs, user-owned data) and the underlying principle of separating identity and data from applications, while operating within the regulatory constraints of the jurisdictions in which it originates credit.

REGULATORY CONTEXT. All credit origination referenced in this document is conducted exclusively by independent licensed entities operating under applicable regulatory frameworks. The WERITAS infrastructure does not participate in, control, or execute credit origination activities. In Kenya, the Phase 1 environment, credit is originated by Central Bank of Kenya (CBK) licensed Digital Credit Providers. The identity and credential infrastructure described here does not itself originate credit, issue tokens, or custody funds. It is an infrastructure layer.

IMPLEMENTATION STATUS. Components described at present exist at varying levels of maturity from production-operational (DCP-originated credit rails) to specification-stage (DWN integration, full CTG graph persistence). Each section identifies the implementation status of the component it describes. Readers should not infer uniform production-readiness across the stack.

REVIEWER INVITATION. This paper is published to invite architectural critique, not capital. The sections that would most benefit from external review are those addressing the Credit Trust Graph (Section V), the credential schema architecture (Section IV), and the consent and disclosure model (Section IV.4). Correspondence may be directed to the engineering contact referenced in Appendix B.

NON-DISCLOSURE. Nothing in this document is confidential. The architecture is described publicly because public technical review is the principal means by which identity infrastructure becomes trustworthy.

CONTENTS

Table of Contents

SECTION I

Core System Thesis: Identity as the Root Primitive of Financial Access

THESIS

WERITAS models credit as a **continuously verifiable identity graph** derived from cryptographically attested financial events, rather than as institution-bound records maintained in proprietary systems.
The identity rooted in a user-controlled Decentralised Identifier is the durable primitive. Loans, repayments, incomes, and institutional attestations are appended to that identity as verifiable credentials.
The graph that emerges is the credit record. The user owns it. Institutions may read from and write to it, subject to user authorisation and applicable regulatory frameworks.

1.1 The Inversion

In the legacy model of consumer credit, the institution is the durable primitive. A bank issues a loan, holds the record of that loan on its own systems, and reports selected data to a centralised credit bureau. If the borrower moves to another institution, the new institution must either pay to access the bureau's summary, or begin the credit relationship effectively from zero. The identity of the borrower exists only as a composite of institutional records held by others, which the borrower cannot carry, correct, revoke, or withhold.

This architecture worked reasonably well in markets where credit bureaus were mature, coverage was universal, and institutional records were reliable. It works poorly often not at all in emerging markets where bureau coverage is partial, records are fragmented, and the majority of economic activity occurs outside the reporting perimeter. In Kenya, over **85% of adults** have no formal credit file usable for a meaningful underwriting decision. They have nonetheless built long histories of disciplined financial behaviour through M-Pesa transactions, informal cooperatives, mobile money repayments, and employer relationships none of which compose into a usable credit identity in the legacy architecture.

WERITAS inverts the primitive. The user is issued a Decentralised Identifier at onboarding. From that moment forward, every institution that transacts with the user a licensed lender issuing a loan, an employer paying wages, a mobile money operator confirming transaction volume, a cooperative attesting to group standing issues a cryptographically signed Verifiable Credential to the user's identity. The user accumulates these credentials in a user-controlled store. The set of credentials, and the graph structure that emerges from their relationships, *is* the user's credit identity.

1.2 Architectural Commitments

This thesis produces architectural commitments that shape every design decision in the system:

Identity Sovereignty

The user's DID and the credentials bound to it are controlled by the user. No institution including WERITAS itself can unilaterally revoke, suppress, or exfiltrate the identity record.

Institution Agnosticism

Credentials are readable by any counterparty the user authorises. A borrower's credit identity is portable across lenders without requiring a bilateral data-sharing agreement between them.

Verifiability Without Custody

Credentials carry their own proofs. A lender can verify an attestation's authenticity cryptographically without holding, processing, or trusting the custody of the underlying data.

Composability Over Monolith

The identity, credential, computation, and settlement layers are independent modules communicating via well-defined interfaces. Each can be upgraded, substituted, or deprecated without cascading rewrites.

1.3 Why This Matters for Credit

The implication of these commitments is that the unit economics of underwriting change. In the legacy architecture, the cost of acquiring a credit-worthy customer is dominated by the cost of verifying their identity and reconstructing their financial history cost that falls heaviest on the first lender and is rarely recoverable. In the WERITAS architecture, that cost is paid once, at credential issuance, and amortises across every subsequent lender the user chooses to share with. The marginal cost of underwriting the 10th loan against a mature identity graph approaches zero. This is the condition under which credit for thin-file populations becomes structurally profitable rather than subsidised and therefore scalable.

THE WEB5 CONNECTION

The principle separating identity and data from applications is not a WERITAS invention. It is the central architectural move of Web5 and related self-sovereign identity frameworks. What WERITAS contributes is a concrete, regulated-market implementation of that principle in a vertical (consumer credit) where the economic case is urgent and the institutional counterparties are willing.

SECTION II

Architecture Overview: The Four Technical Layers

The **WERITAS** technical architecture resolves into four layers, each with a defined responsibility and an explicit interface to its neighbours. The layers are composable: a change to the AI Underwriting Layer does not require modification to the Identity Layer, and a new Regulated Financial Layer integration does not alter how credentials are stored in the Credential Layer. The diagram below presents the layers in their functional relationship.

2.1 System Diagram

FIGURE 2.0 WERITAS WEB5-ALIGNED CREDIT IDENTITY INFRASTRUCTURE

LAYER 01

User Identity Layer

User-controlled Decentralised Identifier (DID) with cryptographic keys. Hosts credential wallet. Issues and signs presentations.

▸ DID DOCUMENT ▸ KEY MATERIAL ▸ CREDENTIAL WALLET ▸ DWN ENDPOINT

LAYER 02

Credential Layer

Cryptographically signed Verifiable Credentials issued by institutions, IncomeAttestation, LoanIssuance, RepaymentEvent, GroupStanding, KYCAnchor.

▸ W3C VC ▸ SD-JWT ▸ STATUSLIST2021 ▸ PRESENTATION EXCHANGE

LAYER 03

Credit Trust Graph (CTG)

Dynamic computation layer materialises user subgraph per query; emits 6-dimensional feature vector; computes counterparty trust weights.

▸ FEATURE VECTOR ▸ TRUST WEIGHTS ▸ CONTINUOUS REASSESSMENT

LAYER 04

Regulated Financial Layer

CBK-licensed DCPs, banks, mobile money operators. Independent licensed institutions originate loans under applicable regulatory frameworks and may issue RepaymentEvent credentials to user-controlled identities.

▸ CBK-LICENSED DCPs ▸ M-PESA ▸ PAYROLL PARTNERS

KEY OUTCOMES: PORTABLE CREDIT IDENTITY · WOMEN-LED LENDING · FINANCIAL INCLUSION

Figure 2.0 The four-layer composable architecture. Identity (Layer 1) anchors the user's sovereign DID. The Credential Layer (Layer 2) accumulates institutionally signed Verifiable Credentials. The Credit Trust Graph (Layer 3) the computation layer builds a dynamic risk profile from the credential set. The Regulated Financial Layer (Layer 4) is the set of licensed institutions that issue credit against the graph. The CTG does not produce credit decisions or approvals; it produces generalised, interpretable signals derived from credential relationships.

FIGURE 2.1 — WERITAS WEB5-ALIGNED CREDIT IDENTITY INFRASTRUCTURE — SYSTEM TOPOLOGY

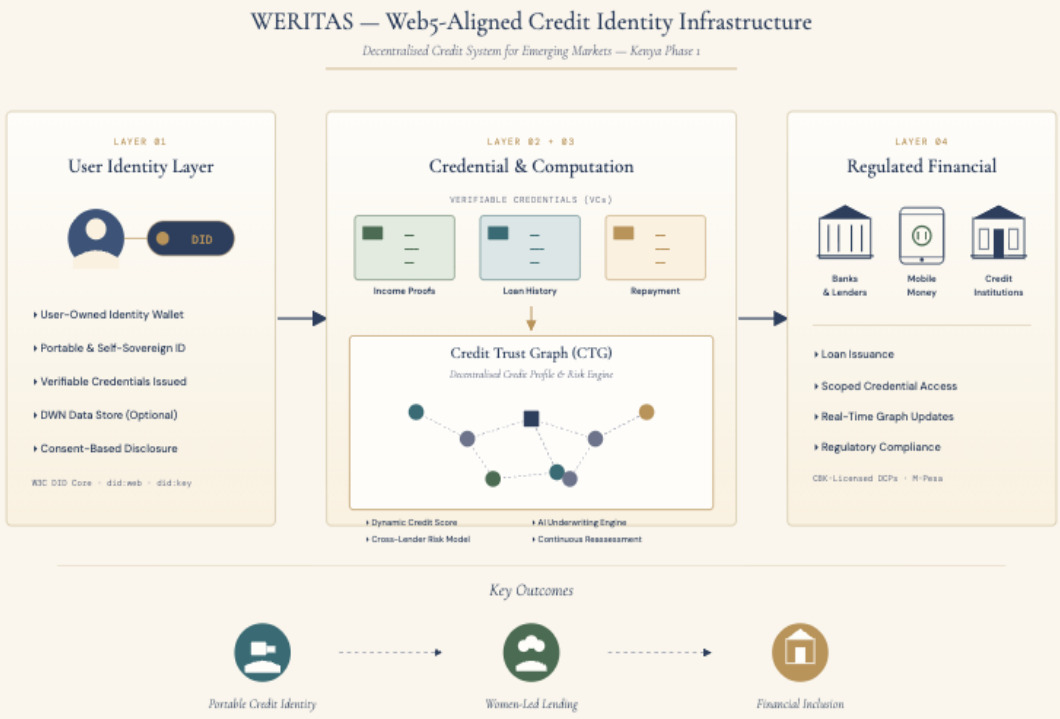

Figure 2.1 — The four-layer composable architecture. Identity (Layer 1) anchors the user's sovereign DID. The Credential Layer (Layer 2) accumulates institutionally signed Verifiable Credentials. The Credit Trust Graph (Layer 3) — the computation layer — builds a dynamic risk profile from the credential set. The Regulated Financial Layer (Layer 4) is the set of licensed institutions that issue credit against the graph. Each layer is independent and communicates through well-defined interfaces.

2.2 Separation of Concerns

Each layer has a narrow, explicit responsibility. A design change confined to one layer does not require modification of any other. This is the central property that enables the system to evolve without fragmentation.

FIGURE 2.2 LAYER RESPONSIBILITIES AND INTERFACES

LAYER	RESPONSIBILITY	WHAT IT DOES NOT DO	INTERFACE TO NEIGHBOURS
Identity Layer	Issues and resolves DIDs; manages key material; hosts the user's credential wallet	Does not score, underwrite, or originate credit; does not hold loan balances	Exposes DID documents via resolver; signs VC presentations
Credential Layer	Standardises VC schemas; stores credentials (user-side or DWN); manages revocation state	Does not compute credit scores; does not make lending decisions	Accepts signed VCs from issuers; presents selectively disclosed VPs to verifiers

LAYER	RESPONSIBILITY	WHAT IT DOES NOT DO	INTERFACE TO NEIGHBOURS
Credit Trust Graph	Builds the user's financial identity graph; computes dynamic trust and risk signals	Does not store PII (reads from credentials on authorised access); does not underwrite	Consumes VPs; emits risk features to the underwriting layer; writes no institutional data
AI Underwriting	Generates decision-support outputs derived from CTG signals for use by licensed financial institutions in their underwriting processes	Does not issue loans; does not hold user credentials; does not write to the graph directly	Consumes CTG features + lender policy; emits a decision-support output object for use within the licensed institution's independent underwriting process
Regulated Financial	Licensed institutional origination, settlement, collection; issues VCs back to the user	Does not own the user's identity record; does not control the credential store	Consumes decision-support outputs and applies its own independent underwriting process prior to disbursement; issues post-event VCs to DID

2.3 What This Separation Buys

SUBSTITUTION COST

Low

Replacing any single layer e.g., swapping the AI underwriting model does not require migration of identity or credential data.

AUDIT BOUNDARY

Clean

A security audit of the credential layer is independent of an audit of the underwriting model. Regulators can review each layer's controls in isolation.

BLAST RADIUS

Contained

A compromise at the underwriting layer cannot exfiltrate identity credentials; a key compromise at one DID does not cascade across the graph.

REVIEWER ATTENTION OPEN QUESTION FOR CRITIQUE

The separation between the Credential Layer and the Credit Trust Graph is the most load-bearing design decision in this architecture. The graph must compute meaningful risk signals without custodying the underlying credentials. Section V addresses how this is achieved via scoped, short-lived verifiable presentations rather than persistent data replication. We invite critique on whether the VP-refresh model scales at the transaction volumes implied by Kenya Phase 1.

This separation enables three enforceable system guarantees:

(i) compromise containment a breach in any single layer does not expose full user identity or financial history;

(ii) independent auditability each layer can be evaluated in isolation under regulatory or security review; and

(iii) substitution resilience components may be replaced or upgraded without requiring system-wide migration.

SECTION III

Identity Layer: DID Architecture

The Identity Layer anchors every user interaction with the WERITAS system in a Decentralised Identifier that the user controls. The layer is narrow in scope: it issues DIDs, manages their key material, hosts the user's credential wallet, and signs presentations. It holds no credit history, no loan balances, and no underwriting logic.

3.1 DID Method Support & Resolution

The W3C DID Core specification defines an abstract DID syntax, but real implementations must commit to one or more concrete DID methods. WERITAS deliberately supports multiple methods rather than a single proprietary one consistent with the Web5 principle that identity must not be locked to a platform.

FIGURE 3.1 SUPPORTED DID METHODS, PHASE 1

METHOD	USE CASE	PROPERTIES	STATUS
did:key	Lightweight, ephemeral identity for pre-onboarding and low-assurance flows	Derived directly from public key; no registry required; fast, cheap	Operational · First-touch
did:web	Institutional issuers (DCPs, employers, mobile money operators) publishing signing keys	Resolved via HTTPS from issuer-controlled domain; human-auditable; easy rotation	Operational · Primary issuer
did:ion	High-assurance user identities intended to outlive any single custodian	Sidetree-based, anchored to Bitcoin; censorship-resistant; higher complexity	Pilot · Integration underway
did:jwk	Interop bridge for parties preferring JWK-native key exchange	Single-key DID derived from a JSON Web Key; trivially verifiable	Operational

Each user is issued a persistent **did:ion** (or, in Phase 1 interim mode, a locally-anchored equivalent) as their primary identity, with **did:key** sub-identifiers derived for specific contexts for example, a one-time did:key used to interact with a lender during a single loan application, without exposing the primary DID. This follows the principle of minimal correlation.

3.2 DID Document Structure

```
// Representative DID document for a user's primary did:ion
{
  "@context": [
    "https://www.w3.org/ns/did/v1",
    "https://w3id.org/security/suites/jws-2020/v1"
  ],
  "id": "did:ion:EiClkZMDxPKqC9c-umQfTkR8...",
  "verificationMethod": [
    {
      "id": "did:ion:Ei...#sig-1",
      "type": "JsonWebKey2020",
      "controller": "did:ion:Ei...",
      "publicKeyJwk": { "kty": "OKP", "crv": "Ed25519", ... }
    }
  ],
  "authentication": ["did:ion:Ei...#sig-1"],
  "assertionMethod": ["did:ion:Ei...#sig-1"],
  "keyAgreement": ["did:ion:Ei...#kex-1"],
  "service": [
    {
      "id": "did:ion:Ei...#dwn",
      "type": "DecentralizedWebNode",
      "serviceEndpoint": {
        "nodes": ["https://dwn.weritas.io", "https://dwn.user-host.example"]
      }
    }
  ]
}
```

Two properties are worth reviewer attention. First, the **service endpoint of type DecentralizedWebNode** advertises the user's credential store. Users may elect to run their own DWN, use a community-operated node, or host with **WERITAS** the choice is the user's and is revocable. Second, the **separation of *authentication*, assertionMethod, and *keyAgreement* keys** enforces that the key used to sign a credential presentation is distinct from the key used to authenticate to a service and from the key used to establish encrypted channels. Compromise of one does not imply compromise of the others.

3.3 Key Management & Recovery

Key management is the hardest operational problem in any self-sovereign identity system, and it is the problem on which naïve deployments most often fail. WERITAS adopts a tiered approach calibrated to the device, literacy, and connectivity profile of the target population.

FIGURE 3.2 KEY CUSTODY TIERS

TIER	KEY CUSTODY MODEL	RECOVERY MECHANISM	TARGET USER PROFILE
Tier 0 Custodial	WERITAS holds keys in HSM on user's behalf; explicit consent recorded	Re-authentication via KYC-linked channel (SIM + document)	First-time users; low-friction onboarding
Tier 1 Assisted	Client-side key with server-held encrypted backup (user passphrase derives KEK)	Passphrase + secondary factor; HSM-assisted decryption	Users with basic smartphone familiarity
Tier 2 Social Recovery	Client-side key; Shamir-split backup distributed across 3–5 trusted contacts	Threshold reconstruction (e.g., 3 of 5); no central party can recover	Established users with community ties
Tier 3 Self-Custodial	Fully user-held; hardware wallet or secure enclave; no WERITAS-side custody	User-managed seed phrase or hardware backup	Technical users; high-assurance use cases

The design principle is that **self-sovereignty is an outcome on a spectrum, not a prerequisite for onboarding**. Forcing Tier 3 custody at first interaction excludes precisely the populations WERITAS exists to serve. A user may onboard at Tier 0, graduate to Tier 1 as they become comfortable, and migrate to Tier 2 or Tier 3 over time. The migration is user-initiated and does not require re-issuance of credentials the underlying DID is preserved.

> DESIGN DISCIPLINE
>
> At every tier, the invariant is that the *user's credential set is addressable by the user's DID, not by any WERITAS-internal identifier*. A user who migrates from Tier 0 to Tier 3 takes their credentials with them. WERITAS cannot withhold them, and the counterparty institutions that issued them do not need to re-issue. This is the concrete meaning of "identity sovereignty" in this system.

The custodial tiers (Tier 0 and Tier 1) introduce an explicit trust dependency on WERITAS-managed infrastructure. While this is necessary for initial user accessibility, it represents a temporary security trade-off.

The system is therefore designed to progressively migrate users toward higher self-custody tiers, reducing centralized risk exposure over time.

3.4 Device Constraints & Reality Check

A significant share of the Kenya Phase 1 population does not own a smartphone capable of running a full DWN node. The architecture accommodates this without compromising the long-term target:

- **USSD fallback.** Basic feature-phone users interact via USSD. Signing operations occur on a server-side HSM bound to the user's KYC anchor; the DID and credentials are real, the user can migrate to client-side custody when their device permits. This is explicitly Tier 0.
- **Progressive enhancement.** The same DID that was custodially managed at USSD onboarding is the DID the user brings forward when they upgrade to a smartphone and install the WERITAS app. No loss of credit history.
- **Offline-tolerant VC verification.** VC proofs are verifiable without network access to the issuer, provided the verifier has cached the issuer's DID document. This is essential in low-connectivity retail points.
- **SIM-bound recovery.** Because the Kenya mobile ecosystem pairs every mobile money account to a registered SIM, SIM-binding is a pragmatic secondary factor for Tier 0 and Tier 1 recovery. This is a regulatory strength, not a weakness.

3.5 What the Identity Layer Does Not Do

It is worth stating plainly.

The Identity Layer does not score credit, does not approve loans, does not hold fiat or tokens, does not record loan balances, and does not make underwriting decisions.

It is narrow by design. Every component above it that does these things operates against the DID and the credential set, not against any proprietary user record. This is the discipline that lets the layer be replaced, audited, or migrated without touching the rest of the stack.

SECTION IV

Credential Layer: Verifiable Credit History

If the Identity Layer answers "who", the Credential Layer answers "what is verifiably true about them". Every financial event that is relevant to credit an income deposit, a loan issuance, a repayment completion, an institutional attestation of good standing is represented as a **Verifiable Credential**: a cryptographically signed, structured statement issued by one party (the institution) about another party (the user, identified by DID).

4.1 Credential Taxonomy

WERITAS standardises a taxonomy of credit-relevant credentials. The taxonomy is open institutions can issue credentials outside the core set, and the system is designed to consume them gracefully but the core set is what the Credit Trust Graph explicitly recognises.

FIGURE 4.1 CORE CREDENTIAL TAXONOMY (PHASE 1)

CREDENTIAL TYPE	TYPICAL ISSUER	CONTENTS (SUBJECT TO SCHEMA)	TYPICAL VALIDITY
IncomeAttestation	Employer; payroll processor	Amount bracket, frequency, payer DID, attestation period	1–3 months; renewable
LoanIssuance	Licensed DCP; bank	Principal bracket, tenor, rate bracket, product class, currency	Loan lifetime + archival
RepaymentEvent	DCP; mobile money operator	On-time flag, period, loan reference (one-way hash), cumulative count	Permanent; append-only
TransactionBehaviour	Mobile money operator (M-Pesa)	Volume bracket, frequency bracket, regularity score, period	Rolling 1–3 months
GroupStanding	Cooperative; chama; SACCO	Membership tenure, contribution consistency, peer-attested standing	3–12 months
KYCAnchor	Licensed KYC provider	Assurance level, document class (hash), issuing jurisdiction	12–24 months
DefaultEvent	DCP (with regulatory obligation)	Severity class, resolution status, period	Per Kenya DCA provisions
DisputeFlag	User-initiated; institution-countersigned	Flagged credential reference; dispute status; resolution	Until resolved

Two properties are worth noting. First, credentials carry **brackets and flags, not raw values wherever possible**. An IncomeAttestation states that income falls within a bracket (e.g., KES 30,000–50,000/month), not the exact figure. This is a deliberate minimisation: the CTG does not need exact values to score risk, and brackets reduce the disclosure surface when a credential is presented. Second, the credential set explicitly includes **negative** signals DefaultEvent, DisputeFlag because a credit infrastructure that cannot represent adverse outcomes is not a credit infrastructure.

The integrity of the credential layer is contingent on issuer quality. The system assumes that credential issuers operate within regulated or auditable frameworks. To mitigate the risk of low-quality or adversarial issuers, downstream computation layers apply issuer-specific trust weighting and may incorporate additional validation heuristics.

4.2 Issuance Flow

01	02	03	04	05
Event Occurs	**Issuer Composes VC**	**Delivered to DID**	**User Accepts**	**Revocation Registered**
Loan disbursed, wage paid, repayment made	Against published schema; signs with did:web key	To user's DWN endpoint or inbox	Credential stored in user wallet	Issuer records in revocation registry

The issuance flow is deliberately asymmetric: the issuer signs and pushes; the user accepts or rejects. The user cannot forge a credential (the issuer's signature is required), and the issuer cannot unilaterally remove a credential from the user's store (they can only publish a revocation, which verifiers honour during verification). This asymmetry is the integrity property that makes the graph trustworthy.

4.3 Schema & Proof Format

WERITAS credentials conform to the W3C Verifiable Credentials Data Model. Each credential carries a JSON-LD context pointing to a WERITAS-published schema definition for its type, a credentialSubject identified by DID, and an attached proof. The default proof suite is Ed25519 with a Data Integrity proof (*DataIntegrityProof*), with ECDSA secp256r1 supported for interoperability with institutional HSMs that do not natively handle Ed25519.

```
// Example RepaymentEvent Verifiable Credential
{
  "@context": [
    "https://www.w3.org/2018/credentials/v1",
    "https://schemas.weritas.io/credit/v1"
  ],
  "id": "urn:uuid:4f8a2e1c-...",
  "type": ["VerifiableCredential", "RepaymentEvent"],
  "issuer": "did:web:asap-credit.weritas.io",
  "issuanceDate": "2026-04-15T14:22:17Z",
  "credentialSubject": {
    "id": "did:ion:EiClkZMDxPKqC9c-...",
    "loanReferenceHash": "b3a1f4e2...9c7d",    // one-way hash of loan ID
    "onTime": true,
    "period": "2026-04",
    "cumulativeOnTimeCount": 17,
    "cumulativeDelinquentCount": 0
  },
  "credentialStatus": {
    "id": "https://revoke.weritas.io/credit/list#42",
    "type": "StatusList2021Entry",
    "statusListIndex": "42",
    "statusListCredential": "https://revoke.weritas.io/credit/list"
  },
  "proof": {
    "type": "DataIntegrityProof",
    "cryptosuite": "eddsa-2022",
    "created": "2026-04-15T14:22:18Z",
    "verificationMethod": "did:web:asap-credit.weritas.io#key-1",
    "proofPurpose": "assertionMethod",
    "proofValue": "z5vgY8..."
  }
}
```

Several schema decisions are worth reviewer attention:

- **Loan references are hashes, not identifiers.** A verifier can confirm that two repayment credentials reference the same underlying loan (because the hashes match), but cannot from the hash alone discover the loan's institutional identifier. This preserves linkability for the CTG while minimising leakage.

- **Cumulative counts are carried on each credential.** This allows the CTG to score from any single recent credential without needing to reassemble the full history, a performance property that matters at scale.
- **Periods are coarse.** The RepaymentEvent above records the period as "2026-04", not a timestamp. This frustrates correlation across credentials while preserving the temporal signal the CTG needs.
- **StatusList2021 is used for revocation.** It is efficient (a single bitstring revokes millions of credentials), privacy-preserving (a verifier cannot tell which credential is being checked), and widely implemented.

4.4 Decentralized Web Node Integration

Credentials must live somewhere.

The Web5 reference architecture specifies Decentralized Web Nodes (DWNs) as the user's sovereign data store, a node the user either runs themselves or elects to have hosted.

WERITAS integrates with **DWNs** as the first-class credential storage pattern, with a managed-hosting option for users who do not elect self-custody.

FIGURE 4.2 DWN INTEGRATION PATTERNS

HOSTING PATTERN	USER EXPERIENCE	TECHNICAL PROPERTY
Self-Hosted	User runs DWN locally on device or on a VPS they control	Full user sovereignty; DID service endpoint points to user's node; data never touches WERITAS infra
Community-Hosted	User selects a community-operated DWN (e.g., by a cooperative or NGO)	Node operator is not WERITAS; user can migrate between nodes without re-issuing credentials
WERITAS-Hosted	User opts into WERITAS-operated DWN at onboarding	Default for low-tech users; user-exportable at any time; WERITAS cannot read credentials (end-to-end encrypted at rest with user key)
Dual-Write	User writes to two DWNs for redundancy	CRDT-style reconciliation on conflict; enables "migrate away" without data loss

The critical property across all four patterns is that the DID service endpoint is user-mutable. A user on a WERITAS-hosted DWN who later runs their own node simply updates their DID document.

Counterparty institutions, on their next credential issuance or verification, resolve the DID, discover the new endpoint, and proceed.

The migration is invisible to the rest of the system.

4.5 Consent, Selective Disclosure & Revocation

A credit identity must support selective disclosure to remain usable in real-world financial interactions while minimizing unnecessary data exposure. The Credential Layer implements three disclosure primitives.

PRESENTATION EXCHANGE

When a lender wishes to underwrite a prospective borrower, they do not request "all credentials". They publish a **Presentation Definition** a machine-readable description of exactly which credential types, issued by which parties, with which constraints, they require. The user's wallet evaluates the definition against its local credentials and either constructs a matching **Verifiable Presentation** or informs the user that the request cannot be satisfied without sharing out-of-scope information. The user reviews, and approves or declines.

This mirrors the DIF Presentation Exchange specification. It makes data requests auditable (the definition is a durable artefact) and makes user consent meaningful (the user sees exactly what is being asked for).

SELECTIVE DISCLOSURE VIA SD-JWT

For credentials where even bracket-level disclosure is more than a verifier needs, WERITAS supports the IETF SD-JWT profile. A credential is issued with claims hashed and salted; the user, when presenting, reveals only the hashes for the claims being disclosed. A verifier can confirm that a specific claim was issued without seeing unrelated claims on the same credential.

It should be noted that selective disclosure mechanisms reduce, but do not eliminate, correlation risk across repeated interactions. Further privacy enhancements (e.g., zero-knowledge proofs) are under active research and are not yet part of the production stack.

> CONCRETE EXAMPLE
>
> A user holds an IncomeAttestation credential with claims: *income bracket, payer DID, employment type, region, period*. A ride-hail platform only needs `employment type = "active"` and `period = current`. With SD-JWT, the user presents exactly those two claims; the platform verifies them; the income bracket and payer identity are never disclosed. The verification is cryptographically equivalent to verifying the full credential the other claims simply don't appear in the disclosure.

REVOCATION

Revocation is an issuer capability. An issuer can mark a credential as revoked by flipping the corresponding bit in their StatusList2021 credential. A verifier, during verification, fetches the current status list and checks the bit. Revocation is:

- **Issuer-scoped.** Only the issuer who signed a credential can revoke it. WERITAS cannot revoke a credential issued by a partner DCP, and no partner DCP can revoke a credential issued by another.
- **Non-deletive.** Revocation does not remove the credential from the user's wallet. The user may still see it, and may still present it a verifier who accepts revoked credentials in certain contexts (e.g., a regulator auditing historical conduct) can do so explicitly.

- **Disputable.** The user can issue a DisputeFlag credential (countersigned by an adjudicating party) that verifiers are expected to check alongside the revocation list. This prevents issuer abuse.

4.6 Consent as a First-Class Artefact

Every credential presentation to a verifier generates a **Consent Receipt** itself a *Verifiable Credential*, signed by the user that records what was disclosed, to whom, for what stated purpose, and under what retention terms.

The receipt is stored in the **user's DWN** alongside the original credentials.

This provides two properties:

(1) the user has an auditable record of every disclosure they have ever made;

(2) a verifier who processes data outside the consented purpose is provably in breach, which is the foundation of any effective regulatory or contractual enforcement.

```
// Example  Consent Receipt
{
  "@context": ["https://www.w3.org/2018/credentials/v1",
               "https://schemas.weritas.io/consent/v1"],
  "type": ["VerifiableCredential", "ConsentReceipt"],
  "issuer": "did:ion:EiClkZMDxPKqC9c-...",      // user signs
  "issuanceDate": "2026-04-19T09:14:02Z",
  "credentialSubject": {
    "verifier": "did:web:partner-lender.example",
    "purpose": "underwriting-decision-single-product",
    "presentedCredentials": [
      { "type": "IncomeAttestation",
        "disclosedClaims": ["bracket", "period"] },
      { "type": "RepaymentEvent",
        "disclosedClaims": ["onTime", "cumulativeOnTimeCount"] }
    ],
    "retentionTermDays": 90,
    "purposeBindingClause": "Decision caching permitted; no secondary use."
  },
  "proof": { ... }
}
```

SECTION V

Credit Trust Graph: The Computation Layer

DEFINITION

The **Credit Trust Graph (CTG)** is the computation layer that materialises a user's credential set dynamically, at query time, with the user's explicit consent into a bounded set of interpretable risk features. It is the novel component in the WERITAS architecture, and the one that most deserves external critique.

5.1 Graph Topology

The CTG does not produce credit decisions or approvals; it produces generalised, interpretable signals derived from credential relationships. The CTG is, at heart, a typed directed graph whose nodes and edges are directly derived from credentials. It is not a separate database; it is a computational view constructed from the credentials the user has authorised to be visible in a given context.

FIGURE 5.1 GRAPH NODE AND EDGE TYPES

TYPE	DESCRIPTION	EXAMPLE
Node: DID	A participant user, issuer, adjudicator	User DID, employer DID, DCP DID
Node: Credential	A signed attestation	A specific RepaymentEvent VC hash
Node: Loan	An abstract loan object (by reference hash)	A loan disbursed to the user identified by hash, not loan ID
Edge: issued-to	Issuer DID → Credential → Subject DID	Employer signs IncomeAttestation to user
Edge: references	Credential → Loan (or other Credential)	RepaymentEvent references a Loan hash
Edge: countersigns	DID → Credential	Chama leader countersigns a GroupStanding credential
Edge: supersedes	Credential → Credential	A newer IncomeAttestation supersedes the previous month's
Edge: disputes	Credential → Credential	A DisputeFlag challenges a DefaultEvent

Two properties follow from this topology. First, **the graph is additive: no edge is ever deleted**. Revocation and dispute are themselves edges, not modifications. This preserves a complete, auditable history. Second, **the graph is not globally stored**. The CTG engine materialises the relevant subgraph for a given query at query time, by consuming the Verifiable Presentation the user has authorised. There is no global CTG database to breach.

5.2 Trust Computation

From the materialised subgraph, the CTG derives a bounded feature vector that summarises the user's credit posture. The features are designed to be interpretable each corresponds to a concept a lender would recognise and to be robust to credential-level noise.

FORMULA 5.1 CTG FEATURE VECTOR (ILLUSTRATIVE)

$$F(user) = \langle \tau, \rho, \sigma, \delta, \gamma, \iota \rangle$$

τ (Tenure): Time since earliest credential; measures identity maturity

ρ (Repayment consistency): On-time ratio over rolling 12-month window; weighted by counterparty diversity

σ (Signal density): Credentials-per-month rate; proxies economic activity

δ (Default exposure): Count and severity of DefaultEvents, adjusted for successful DisputeFlags and resolutions

γ (Group anchoring): Presence and tenure of GroupStanding credentials; proxies community accountability

ι (Income stability): Variance of IncomeAttestation brackets over rolling window

Reviewer Note: The feature vector is deliberately small (6 dimensions). A large, opaque feature space is a disguise for overfitting and a privacy liability. Every element of F can be explained to a borrower in a sentence. This is a deliberate choice against "black-box credit scoring" patterns that have drawn regulatory criticism elsewhere.

The feature vector is intentionally constrained in dimensionality to reduce model overfitting, improve explainability, and limit the privacy surface area exposed during computation. Additional features may be introduced only where they demonstrably improve predictive performance without materially increasing disclosure risk

5.3 Counterparty Weighting

Not every attestation carries equal weight. A repayment credential signed by a CBK-licensed DCP is more informative than one signed by an unregulated counterparty; an IncomeAttestation from an employer with a long track record of consistent issuances is more trustworthy than one from a new issuer. The CTG assigns each issuer a trust weight, which modulates how their credentials contribute to the feature vector.

FORMULA 5.2 CREDENTIAL TRUST WEIGHT

$$w(i) = \lambda(i) \cdot \tau(i) \cdot (1 - d(i))$$

$w(i)$ = Trust weight of credential i (bounded [0, 1])

$\lambda(i)$ = Issuer licensing factor (higher for regulated, audited issuers)

$\tau(i)$ = Issuer tenure factor (saturating function of issuer track record)

$d(i)$ = Dispute rate - fraction of issuer's credentials that have been successfully disputed

Trust weights are published and auditable. A user can inspect why a given credential contributes what it does to their score. Issuers who accumulate disputes see their weight decline which is the mechanism by which bad-actor issuers are structurally de-weighted without any central party needing to adjudicate.

5.4 Continuous Reassessment

Unlike a legacy credit score that is computed at origination and re-computed periodically, the CTG is **evaluated at each query**. When a new RepaymentEvent is issued, the next query materialises the subgraph that includes it and produces an updated feature vector. When a credential is disputed or revoked, the next query reflects the change. There is no batch cycle, no "score refresh window", and no stale data.

This is the property that gives the CTG its name. It is a *trust graph*, not a trust *snapshot* it is re-computed in real time, from the current set of authorised credentials, each time a lender or other verifier needs it.

5.5 Privacy-Preserving Computation

The CTG is the point at which the privacy design must be strongest. The layer must compute meaningful signals over a user's credential set without centralising PII, and must remain verifiable enough that a regulator can audit its behaviour without breaking that property.

FIGURE 5.2 PRIVACY MECHANISMS IN THE CTG

MECHANISM	EFFECT	STATUS
Per-query materialisation	No persistent global graph; subgraph exists only during the query that required it	Operational
Scoped consent	Each CTG invocation is tied to a specific consent receipt; unauthorised invocations are provably invalid	Operational
Differential privacy on aggregates	Portfolio-level statistics emitted to research or policy parties have calibrated noise added	Pilot
Zero-knowledge range proofs	A user can prove their ρ (repayment consistency) exceeds a threshold without disclosing its value	Research
Auditable recomputation	A regulator, given the same credential set under warrant, can reproduce the exact feature vector	Operational

REVIEWER QUESTION WE WANT YOUR CRITIQUE

The privacy-preserving claim of the CTG rests on per-query materialisation combined with scoped consent. The open design question is how this scales when a lender is processing thousands of underwriting decisions per hour. Our current answer is caching at the lender side under a documented retention-and-deletion policy encoded in the Consent Receipt with randomised audit of lender compliance. We welcome critique of whether this is sufficient, and what a stronger mechanism (e.g., verifiable computation, ZK rollups over credential sets) would look like in practice. From a systems perspective, scalability is addressed through a combination of (i) short-term caching of previously computed feature vectors under consent-bound retention policies, (ii) incremental graph updates based on newly issued credentials rather than full recomputation, and (iii) parallelized verification pipelines for credential validation. Target system performance is defined in terms of bounded query latency and throughput under production load, and remains an active area of benchmarking.

SECTION VI

AI Underwriting Layer

The AI Underwriting Layer is the component that converts the CTG's feature vector into a lending decision for a specific loan product and a specific lender's policy. It is deliberately positioned *above* the CTG and *below* the Regulated Financial Layer. It consumes features; it emits decisions; it does not see raw credentials, and it does not hold funds.

6.1 Input / Output Contract

```
// Underwriting Request (schematic)
{
  "subjectFeatures": {           // from CTG
    "tenure_months": 14,
    "repayment_consistency": 0.94,
    "signal_density": 8.2,
    "default_exposure": 0.0,
    "group_anchoring": 1.0,
    "income_stability": 0.87
  },
  "lenderPolicy": {
    "product": "payroll-backed-credit-12m",
    "riskAppetite": "moderate",
    "currency": "KES",
    "principalCeiling": 500000,
    "portfolioConstraints": { ... }
  },
  "consentReceipt": "urn:uuid:..."
}

// Response
{
  "decisionRecommendation": "approve",
  "principalRecommended": 180000,
  "rateBracket": "A",
  "explanation": [
    "Strong repayment consistency (94% on-time, 17 events)",
    "Tenure of 14 months exceeds policy minimum",
    "Group anchoring present (chama, 8 months)"
  ],
  "decisionProof": "did:ion:...#model-v3:signed-proof"
}
```

The underwriting model is versioned and subject to continuous monitoring. Model updates are deployed under controlled release cycles, with each version assigned a unique identifier and cryptographic signature. Historical decisions remain attributable to the model version that produced them, enabling full auditability and post-hoc analysis.

6.2 Explainability is a system requirement

Every underwriting decision is accompanied by an explanation a short list of the factors that drove it, expressed in the same interpretable feature names that the borrower can inspect in their own wallet.

This is not a courtesy; it is a design requirement for three reasons:

- **Regulatory.** Kenya's Digital Credit Regulations and most comparable frameworks in adjacent jurisdictions require that adverse credit decisions be explainable. An unexplainable decision is, operationally, not a decision a licensed DCP can issue.
- **Contestability.** A borrower who sees that a decision turned on, for example, a DefaultEvent they dispute can file a DisputeFlag. Without the explanation, they would have no route to redress.
- **Model-governance.** Explanations are the primary artefact through which the model can be audited for bias. If a model's explanations systematically cite different features for similar feature vectors, that is a bias signal. Governance frameworks may review model performance, system parameters, and aggregate outcomes, without participating in or influencing individual lending decisions.

6.3 Decision Signatures

Each underwriting decision is signed by the model that produced it, using a DID assigned to that model version. A lender who relies on a decision can, at any point, verify which model version signed it, what the signed model weights were at that moment (by hash), and whether the model has since been retired or superseded.

This provides a full audit trail without requiring the model weights themselves to be public.

> RELATIONSHIP TO EXISTING WERITAS COMPONENTS
>
> The underwriting model is currently implemented as **Pulse Score** in the WERITAS operational stack, with origination routed through licensed DCPs including the Kenya-deployed ASAP Credit partner. The contract described here is the *interface* to that model. The interface is stable; the model behind it is versioned and replaceable.

SYSTEM RISKS & FAILURE MODES

The WERITAS architecture is designed with explicit failure isolation across layers. However, several operational risks remain:

- Issuer compromise: A compromised issuer key may result in fraudulent credential issuance. Mitigation includes revocation mechanisms and trust weight adjustments.
- Credential availability: Temporary unavailability of credential stores (e.g., DWNs) may impact real-time computation but does not result in data loss.
- CTG service interruption: Computation layer downtime may delay underwriting decisions but does not affect underlying identity or credential integrity.
- Custodial risk (Tier 0/1): Centralized key custody introduces exposure to infrastructure compromise, mitigated through progressive migration to self-custody tiers.

These risks are actively monitored and inform system evolution.

SECTION VII

Web5 Principles Mapping

This section addresses directly the question a Web5-literate reviewer will ask first: in what sense is WERITAS Web5-aligned, and what is the honest boundary of that alignment? The table below maps each core Web5 primitive to its concrete WERITAS implementation, and where there is a compromise, states it.

FIGURE 7.1 WEB5 PRINCIPLES MAPPING

WEB5 PRIMITIVE	WERITAS IMPLEMENTATION	HONEST BOUNDARY
Decentralised Identifiers (DIDs)	Every user is issued a persistent DID at onboarding. Multiple DID methods supported (did:ion, did:web, did:key, did:jwk).	In Tier 0 custodial mode, key material is held by WERITAS on user's behalf until they elect to migrate. The DID is the user's; the keys, at Tier 0, are not yet.
Verifiable Credentials (VCs)	All credit-relevant events are represented as W3C-compliant VCs, signed by institutional issuers. Full taxonomy published.	Current issuers are a curated set of licensed institutions (DCPs, employers, mobile money operators). The open-issuer model that lets any party issue a VC to any user is aspirational for Phase 2.
Decentralized Web Nodes (DWNs)	DWN service endpoints are advertised in the DID document. Self-hosted, community-hosted, and WERITAS-hosted options supported.	In Phase 1, most users default to WERITAS-hosted DWN due to device constraints. Migration tooling exists; critical mass of self-hosted users is Phase 2+.
App-agnostic identity	A user's DID and credential set are not bound to the WERITAS application. A competing lender can resolve the same DID and consume the same credentials.	The CTG computation is currently WERITAS-operated. A competing verifier today can consume credentials, but would need to implement their own trust computation. This is the principal area where decentralisation is still to deepen.
Data sovereignty	Credentials are addressed to the user's DID. The user controls presentation, revocation consent, and retention terms via Consent Receipts.	Event-level data (e.g., raw M-Pesa transaction logs) remains with the originating operator. WERITAS never holds that raw data; the VCs it consumes are bracketed abstractions.
Presentation Exchange	Verifiers publish Presentation Definitions; user wallets evaluate and present selectively.	Operational for lender and employer flows. Cross-ecosystem presentation (e.g., a Rwandan lender consuming a Kenyan user's credentials) awaits regulatory framework maturation.
Separation of identity from data	Credential store is addressable via DID; not via any WERITAS-internal identifier. No WERITAS user-ID shadows the DID in the critical paths.	Regulatory reporting to CBK requires mapping to KYC-attested legal identity. This is an intentional compromise credit is a regulated activity.

STANCE HONEST FRAMING

WERITAS is not a Web5 product.

It is a Web5-aligned credit identity implementation operating in a regulated vertical.

It adopts the architectural primitives and the design philosophy that gives them meaning. It does not overstate the degree of decentralisation at Phase 1.

It is explicit about where centralised convenience is currently traded for accessibility, and about the roadmap that reduces that trade-off over time.

The goal is not to be maximally decentralised at launch;

it is to be structurally capable of decentralising as the population it serves moves up the self-custody ladder.

SECTION VIII

End-to-End System Flow

This section traces a single user's journey from onboarding through repeated credit access, showing how the four layers interact. It is presented as an operational narrative because the architecture is easier to evaluate against a concrete path than against abstract diagrams.

8.1 Onboarding First Touch

1	**KYC Anchor** *Identity Establishment*	User presents KYC documents to a licensed partner. A KYCAnchor credential is issued. A DID is created (default did:ion, Tier 0 custody).
2	**DID Published** *To DWN Endpoint*	DID document is published; DWN service endpoint is registered (WERITAS-hosted by default in Tier 0). User's wallet is operational.
3	**Initial Credentials** *Seed the Graph*	With user consent, baseline credentials are issued: an IncomeAttestation (if employed), an initial TransactionBehaviour VC (from mobile money, if history exists).

8.2 First Credit Application

4	**Lender Request** *Presentation Definition*	A licensed DCP publishes a Presentation Definition describing the credentials required for the target product. User's wallet evaluates.
5	**User Consents** *Receipt Issued*	User reviews exactly which claims are being disclosed, to whom, for what purpose. Approves. A Consent Receipt is issued and stored.
6	**CTG Materialises** *Feature Vector Computed*	The CTG engine materialises the subgraph from the presented credentials; computes the 6-dimensional feature vector; returns it to the underwriting layer.
7	**Decision + Support** *Returned*	Decision-support output generated and returned to the originating licensed institution
8	**Disbursement** *Via M-Pesa*	If approved by the licensed DCP under its own independent decision-making process, disburses via M-Pesa. A LoanIssuance VC is signed by the DCP and delivered to the user's DID.

8.3 Repayment & Graph Growth

9	**Repayment Events** *Monthly Cadence*	Each scheduled repayment typically via automated payroll deduction or M-Pesa generates a RepaymentEvent VC signed by the DCP, delivered to the user's DID.
10	**Graph Evolves** *Continuously*	The user's credential set grows; the CTG feature vector improves with each on-time event; the user becomes underwritable for progressively larger and longer-tenored products.

8.4 Second Lender Portability in Action

11	**New Lender** *Same DID*	The user applies to a second lender. The new lender publishes a Presentation Definition. The user reviews and presents using the same DID, from the same wallet, without any re-onboarding.
12	**No Data Migration** *No Bureau Call*	The new lender needs no bilateral data-sharing agreement with the first lender. No bureau pull. The CTG emits the same feature vector to any authorised verifier.
13	**Competitive Pricing** *Because Cost to Serve Fell*	The second lender's underwriting cost on this user is near zero they trusted institutionally signed credentials instead of reconstructing history. Competition, not monopoly, becomes the equilibrium.

THE ECONOMIC PAYOFF

Step 13 is the one that justifies the architecture. In the legacy model, the second lender's cost to serve approaches the first lender's each institution bears the full acquisition cost anew. In WERITAS, the acquisition cost is paid once, at Step 1–3, and amortises across every subsequent lender. This is the unit economics that lets credit reach populations the legacy model prices out.

SECTION IX

Deployment Environment: Kenya Phase 1

Kenya is not chosen as Phase 1 for convenience. It is chosen because it satisfies a very specific convergence of technical and regulatory preconditions the infrastructure a Web5-aligned credit system requires to function, and the regulatory clarity that makes operating in it legible to institutional counterparties.

9.1 Why Kenya Satisfies the Preconditions

FIGURE 9.1 KENYA PRECONDITIONS FOR WEB5-ALIGNED CREDIT

PRECONDITION	KENYA STATE	RELEVANCE TO ARCHITECTURE
Universal settlement rails	M-Pesa penetration exceeds 90% of adults	Loan disbursement and repayment are frictionless; TransactionBehaviour credentials have meaningful signal
Regulated digital lending	CBK Digital Credit Provider licensing regime operational since 2022	A clear licensed counterparty exists to originate credit and issue authoritative credentials
Data protection framework	Data Protection Act 2019; Office of the Data Protection Commissioner active	Consent-Receipt model has enforceable legal backing; a breach is not merely a contractual matter
SIM registration	Mandatory KYC-bound SIM registration	SIM-bound recovery (Tier 0/1) is a pragmatic secondary factor with regulatory support
Partner availability	Active CBK-licensed DCP ecosystem; established KYC providers	Credential issuers are not hypothetical; they exist, operate, and are integrable
Identity foundation	National ID system; Huduma integration ongoing	KYCAnchor credentials have a robust underlying identity substrate
English-language interop	English is an official language; technical documentation is accessible	Reduces translation friction for schema definitions, consent receipts, explanations

9.2 Phase 1 Scope

Phase 1 is deliberately narrow. The scope is: **credential issuance and consumption for payroll-linked consumer credit, originated by a single CBK-licensed DCP partner, priced in KES, with repayment via M-Pesa, and accompanied by a running evaluation of the architecture's behaviour at production volumes.**

The aim is not to demonstrate every feature described in this paper in Phase 1, it is to establish that the **fundamental identity, credential, and CTG primitives** operate correctly against real borrowers, real lenders, and real money.

Components explicitly *outside* Phase 1 scope, documented here for reviewer clarity:

- **Cross-border credential portability.** Kenya-only in Phase 1. Regional interop is Phase 2+.
- **Non-DCP originators.** Only licensed DCP credentials are consumed in the CTG during Phase 1. Cooperative- or peer-issued credentials are collected but de-weighted until pilot calibration completes.
- **ZK proof presentation.** SD-JWT selective disclosure is Phase 1; full ZK range proofs are Phase 2.
- **Full self-custody default.** Tier 3 custody is available in Phase 1 for users who opt in; the default for new users is Tier 0/1.
- **Secondary markets / structured finance.** Handled separately in the WERITAS Whitepaper v4.0. Not part of this technical paper's scope.

Operational metrics for Phase 1 including credential issuance volume, underwriting request throughput, and system latency are tracked continuously and will inform scaling decisions in subsequent phases

> DEPLOYMENT POSTURE
>
> Phase 1 is a correctness exercise, not a scale exercise. We are explicitly optimising for "does the architecture behave as specified under real load" rather than for maximum volume. Protocol correctness, cleanly measured, is the foundation on which Phase 2's regional expansion either rests or fails.

SECTION X

Structured Finance: A Separate Abstraction

A WERITAS reader coming from the finance side will recognise that this paper has, until this point, said remarkably little about capital formation, tokenisation, or structured credit.

This is deliberate. **The identity and credit-graph architecture is the primary protocol**. Structured finance is a consuming application a powerful and economically consequential one, but architecturally downstream.

This section establishes the boundary between the two cleanly, so that a reviewer focused on the identity architecture can evaluate it without the noise of the capital stack, and a reviewer focused on the capital stack knows which document to read.

10.1 The Clean Interface

The CTG emits feature vectors per user. Licensed DCPs originate individual loans against those feature vectors.

A *separate* compliant and regulated capital architecture documented in **WERITAS Whitepaper**, pools those originated loans, transforms them into tranched structured credit instruments, and the distribution of such instruments to qualified institutional investors is conducted through separate regulated structures and entities.

The identity, credential, and CTG layers are explicitly decoupled from capital formation and asset ownership structures. Their function is limited to producing verifiable underwriting signals, independent of how underlying loans are funded, held, or securitized.

FIGURE 10.1 INTERFACE BETWEEN IDENTITY PROTOCOL AND STRUCTURED FINANCE LAYER

IDENTITY PROTOCOL (THIS PAPER)	INTERFACE	STRUCTURED FINANCE LAYER (WHITEPAPER V4.0)
Emits: per-loan underwriting signals, portfolio-level CTG statistics (differentially-private)	**Signed feature vectors + aggregate graph statistics**	Consumes: portfolio composition, loss-given-default inputs, pool-level risk metrics
Enables: better unit economics per loan, lower acquisition cost, dynamic reassessment	**Better pool diversification, dynamic risk re-tranching capability**	Produces: Senior / Mezzanine / Junior tranches, institutional instruments
Does not emit: user PII, raw credentials, individual identity data	**Anonymised pool statistics only institutional layer never sees individual users**	Does not consume: user identities, credentials, wallets

10.2 Why This Boundary Matters

A failure in the structured finance layer a tranche mispricing, an SPV re-structuring, a regulatory intervention in institutional markets does not corrupt the identity protocol. A user's credentials remain valid and portable regardless of whether the loan tranches above them are performing. Conversely, the collapse of a single DCP does not destroy user identities; they are mobile, held in the user's wallet, and readable by the next licensed originator.

This clean separation is the principal reason the structured finance ambition which at full scale is consequential does not compromise the protocol correctness argument for the identity layer. They are separate systems, separately auditable, separately upgradeable, and separately failure-isolated.

> POINTER FOR INSTITUTIONAL READERS
>
> The detailed architecture of the structured finance layer $WASAP instruments, Mauritius SPV mechanics, Senior/Mezzanine/Junior tranche design, Big 4 audit architecture, credit pool mathematics (Basel-grade EL/UL/RAROC framework) is presented in full in ***WERITAS Whitepaper***, **Sections V–VI.** That document is the authoritative reference for the capital formation model. This document is the authoritative reference for the identity protocol that makes the capital formation model operate at unit economics that would otherwise be unavailable.

SECTION XI

Strategic Significance & Open Questions

11.1 What This Architecture, If Correct, Enables

If the architecture described here operates at specification, three consequences follow that are not currently available in any other credit infrastructure for emerging markets.

Portable Credit Identity

A borrower's credit history travels with them across institutions, across products, and eventually across borders. The bureau-and-bilateral-agreement pattern is displaced by credential-and-wallet.

Structural Inclusion

Populations priced out of the legacy model women borrowers, informal workers, thin-file adults are underwritable at sustainable unit economics because acquisition cost amortises across the ecosystem rather than falling on the first lender.

Institutional Interop

Licensed institutions write to a common identity and credential substrate. Competition between them does not require fragmenting borrower data, and cross-institutional risk visibility becomes possible.

Regulatory Legibility

The consent-receipt and signed-decision model makes every step of the credit lifecycle auditable to regulators without requiring them to custody the underlying data. This is a stronger regulatory posture than is currently achievable in the legacy model.

11.2 Open Questions We Are Explicit About

The architecture is not finished.

The following are the open design questions the team considers most consequential, stated plainly for the benefit of external reviewers.

FIGURE 12.1 OPEN TECHNICAL QUESTIONS

QUESTION	CURRENT STANCE
CTG scalability under production VP volumes	Per-query materialisation works in pilot. Whether it holds at steady-state production throughput, or whether an intermediate cache with strong freshness guarantees is required, is an open engineering question (see Section V.5).
Appropriate trust weight calibration	The formula in Section V.3 is a reference. How w(i) is calibrated in practice what the weights for a first-year DCP should be relative to a ten-year bank is a governance question whose answer has not been determined.
Cross-jurisdictional credential recognition	Kenya-issued credentials are usable in Kenya by design. How they should be treated by a Rwandan or Ugandan verifier in Phase 2 is a policy question that interacts with local regulatory frameworks.
Bureau integration vs. displacement	Whether the system integrates with existing credit reference bureaus (consuming their data as VCs) or targets displacement (rendering them redundant) has near-term operational implications we have not yet settled.
Default event representation	A DefaultEvent VC that cannot be reliably rehabilitated after resolution traps users in a negative record. The current design permits DisputeFlag and Resolution credentials. Whether this is sufficient or whether the graph needs stronger redemption semantics is an open design question.
Adversarial credential farming	An economically rational attacker could attempt to farm positive credentials (small paid loans, fast repayments) to inflate their feature vector before a large application. Trust weights and graph density signals detect this probabilistically, not deterministically. Stronger mechanisms are under study.

SECTION XII

What We Are Seeking

This paper exists to invite specific kinds of engagement. It is worth stating, plainly, what those are and what they are not.

12.1 What We Are Seeking

Architecture Validation *Technical Critique*	A rigorous external review of the DID / VC / CTG design by engineers and researchers who have built decentralised identity infrastructure elsewhere. We want to be told where the design is wrong.
Web5 Ecosystem Feedback *Interop Review*	Feedback from builders in the TBD / Web5 community on how WERITAS should interop with the emerging Web5 tooling ecosystem DWN specifications, credential exchange protocols, key management patterns.
CTG Design Critique *Most Load-Bearing Question*	Specifically: critique of the Credit Trust Graph design, the per-query materialisation model, the feature vector choice, and the privacy-preserving computation claim. This is the most novel component and deserves the hardest review.
Protocol-Level Collaboration *Where Interests Align*	Pathways for collaboration with other protocols building DID-aligned infrastructure in adjacent verticals (payments, health records, educational credentials) where shared primitives could reduce duplication.

12.2 What We Are Not Seeking

OUT OF SCOPE FOR THIS PAPER

We are **not seeking token investment, fundraising, or capital allocation decisions** through this document. $WRTH token details are in the WERITAS Whitepaper and are governed by separate disclosures; this paper is about the architecture that sits beneath the ecosystem, not the token that coordinates it. Readers who would like to discuss capital participation in the structured credit layer should refer to that document and the institutional engagement channels referenced there.

12.3 How to Engage

Correspondence on the technical architecture including critique, questions, and collaboration proposals may be directed to the technical contact in Appendix B. The most useful form of engagement is a concrete, bounded critique of a specific section of this document, accompanied by a reference to prior work or an alternative design if the reviewer has one. We will respond substantively. We will not respond by updating this paper quietly and claiming to have always thought that revisions will be versioned, dated, and changelogged.

Identity is the root primitive. Credit is what grows from it.

weritascouncil.org | weritastoken.io

APPENDIX A

Glossary of Technical Terms

TERM	DEFINITION
DID	Decentralised Identifier a URI-syntax identifier that resolves to a DID Document, controlled by the subject rather than a registrar. Defined in W3C DID Core.
DID Document	A JSON-LD document containing verification methods, service endpoints, and controllers for a given DID. Resolved via a DID Method.
DID Method	A specification describing how DIDs of a given scheme (did:ion, did:web, did:key, etc.) are created, resolved, updated, and deactivated.
VC	Verifiable Credential a tamper-evident, cryptographically signed statement issued by one party (issuer) about another (subject), conforming to W3C VC Data Model.
VP	Verifiable Presentation a credential or derived disclosure, presented by the holder (typically the subject) to a verifier, potentially with selective disclosure applied.
DWN	Decentralized Web Node a user-controlled data store, specified within the Web5 architecture, that hosts the user's credentials and other personal data.
SD-JWT	Selective Disclosure JWT an IETF profile enabling a credential to be issued with claims hashed and salted, such that the holder can selectively disclose individual claims.
StatusList2021	A W3C credential status mechanism using a compressed bitstring to represent revocation state for large populations of credentials efficiently and privacy-preserving.
CTG	Credit Trust Graph the WERITAS-specific computation layer that materialises a user's credential set into a dynamic risk feature vector per query.
Feature Vector	The 6-dimensional output of the CTG ($\langle \tau, \rho, \sigma, \delta, \gamma, \iota \rangle$) that summarises a user's credit posture for consumption by the underwriting layer.
Consent Receipt	A user-signed VC recording a specific disclosure event which credentials, to which verifier, for what purpose, with what retention terms.
Presentation Exchange	A DIF specification for verifiers to describe credential requirements (Presentation Definition) and for wallets to respond with matching presentations.
Proof Suite	A specific cryptographic proof algorithm used with Data Integrity Proofs; WERITAS defaults to Ed25519 (eddsa-2022) with ECDSA secp256r1 as an interop fallback.
KYC Anchor	A VC attesting to the user's completion of Know-Your-Customer checks at a defined assurance level; the foundational credential from which other credentials derive trust.
DCP	Digital Credit Provider a Central Bank of Kenya licensing category for digital lenders; the principal credential issuer for credit-related VCs in Phase 1.

TERM	DEFINITION
CBK	Central Bank of Kenya primary financial regulator for Kenya-domiciled financial institutions and the licensor of DCPs.
Pulse Score	The AI underwriting model currently implementing the underwriting contract described in Section VI; versioned and signature-bound per decision.
did:ion	A Sidetree-based DID method anchored to Bitcoin, providing censorship-resistance and long-term resolvability without a central registry.
did:web	A DID method that resolves to a document served from an HTTPS-reachable domain; commonly used by institutional issuers.
did:key	A DID method where the DID is deterministically derived from a single public key; used for lightweight, ephemeral identifiers.
Phase 1	Kenya-only deployment of the identity, credential, CTG, and underwriting layers with a single licensed DCP partner; 2026.
Phase 2	Regional expansion, cross-border credential portability, and additional DID methods including fuller self-custody defaults; 2027+.

APPENDIX B

References & Contact

B.1 Specifications & Normative References

SPECIFICATION	PUBLISHER	RELEVANCE
DID Core	W3C Recommendation	Core syntax and DID Document structure (Section III)
Verifiable Credentials Data Model	W3C Recommendation	VC and VP structure; proof attachment (Section IV)
Data Integrity Proofs	W3C Working Draft	Default proof format for WERITAS credentials (Section IV.3)
StatusList2021	W3C CCG Draft	Revocation mechanism for issued credentials (Section IV.5)
Presentation Exchange	DIF	Verifier-wallet credential negotiation protocol (Section IV.5)
SD-JWT	IETF Draft	Selective disclosure profile for JWT-format credentials (Section IV.5)
Web5 Architecture	TBD / Open Community	DWN patterns and architectural principles (Section IV.4, VII)
Sidetree Protocol	DIF	Substrate for did:ion (Section III.1)

B.2 Companion WERITAS Documents

DOCUMENT	CONTENTS
WERITAS Whitepaper	Full protocol whitepaper covering governance, token economics, structured finance layer, and business context
WERITAS One-Pager	Institutional executive summary
Technical Paper v3.0 (this document)	Identity & credit-graph architecture, Web5 alignment, deployment context

B.3 Official Links & Contact

RESOURCE	URL / CONTACT
Weritas Council Primary Hub	weritascouncil.org
Technical & Architecture Review	legal@weritascouncil.org
Web5 / DID Interop Discussions	legal@weritascouncil.org
Legal Contact	legal@weritascouncil.org
Media & Press	media@weritas.io
Partnerships	partners@weritas.io

B.4 Document Metadata

ATTRIBUTE	VALUE
Document Title	WERITAS Web5-Aligned Credit Identity Infrastructure Technical Paper
Version	5.0 Architecture Focus
Publication Date	April 19, 2026
Prior Version	v1-v4 Internal drafts, March 2026 (superseded by this version)
Status	Published for external technical review
Companion to	WERITAS Whitepaper v24 (published April 21, 2026)
Feedback Window	Open. Versioned revisions will be published as substantive critique arrives.

WERITAS TECHNICAL ARCHITECTURE | WEB5-ALIGNED CREDIT IDENTITY INFRASTRUCTURE | TECHNICAL PAPER VERSION 5.0 | PUBLISHED APRIL 21, 2026 | FOR EXTERNAL ARCHITECTURE REVIEW