Progressive Assurance: Flexible Organizational Identity in the Digital Economy

2026-05-05

Abstract

The digital economy demands reliable, scalable methods for verifying organizational identity. Today, organizations must choose between low-assurance, low-cost tools like X.509 certificates and high-assurance, high-cost systems like the verifiable Legal Entity Identifier (vLEI). Most business interactions fall between these extremes, lacking a fit-for-purpose solution. This paper introduces the Progressive Assurance Model, which defines four incremental levels of assurance—each built on the same cryptographic identifier and credential schema. The model enables organizations to match assurance to risk, scale as needs evolve, and maintain a durable, auditable evidence chain. We analyze the limitations of current approaches, detail the properties of the progressive model, and discuss its implications for open ecosystems, independent verification, and long-term trust in digital business.

1. The Organizational Identity Gap

In an economy run on digital interactions, a basic question often gets a weak answer: *who are we really dealing with?* The inability to verify a counterparty — supplier, customer, partner — drives sophisticated phishing, invoice fraud, and the slow erosion of trust in business communications.

A piece of the problem hides in plain sight. Most digital infrastructure does not distinguish between a brand name like "Coca-Cola" and a specific legal entity like Coca-Cola Europacific Partners Nederland B.V. [a] The brand is what consumers recognize. The legal entity is what can be sued, what owes taxes, what holds contracts, and what carries reputational consequence. When digital systems treat them as the same thing, fraud finds room to operate.

Today, organizations face a binary choice: low-assurance, low-cost identity tools, or high-assurance, high-cost ones. Almost nothing fits between them. This paper argues for a middle path — a Progressive Assurance Model that lets organizations pay for the assurance they actually need, scale up as risk grows, and reuse the same cryptographic identifier across all levels.

1.1 The cost of digital anonymity

When you cannot reliably know who is on the other end, the malicious actors who exploit the uncertainty have an edge [1, 2]. The losses are concrete: wire fraud, payroll diversion, regulatory penalties, brand damage. The current toolkit is a patchwork. Each tool forces a trade between security, cost, and friction, and none of them is satisfying.

1.2 Today's spectrum: a flawed choice

Two poles dominate the landscape, with a wide and underserved middle.

The low-assurance floor: X.509 certificates For four decades, X.509 PKI has secured the web [3]. Its application to organizational identity, however, is an accident of history. X.509 was designed to authenticate web servers for TLS, not to prove the identity of a durable legal entity [1]. The fit is poor in several ways.

Lifespans don't match. A legal entity's identity is stable, measured in decades. X.509 certificates are ephemeral by design and shrinking. The CA/Browser Forum has approved a phased reduction of public TLS certificate lifetimes to 47 days by 2029 [4]. Short lifespans are a sensible response to the risk of key compromise on a server. They are an awkward fit for proving that a corporation incorporated in 1975 is still the same corporation today. A business does not renew its articles of incorporation every 47 days [1, 5].

Governance is opaque. Trust in X.509 bottoms out in administrators: Certificate Authorities, browser root programs, and the CA/Browser Forum [5]. Administrators fail. DigiNotar was breached by state actors in 2011 and used to spy on Iranian dissidents [6]. Symantec's CA business was distrusted by Google in 2017 after years of compliance failures [7]. In 2024, Chrome, Firefox, and Apple all removed Entrust from their trust stores citing a consistent pattern of incidents [5]. The pattern recurs. And the practices that matter most — how the certificate *holder* manages its private key — are invisible to verifiers and unaudited by anyone [1].

Verification often phones home. To check whether a certificate has been revoked, clients query an Online Certificate Status Protocol responder. If the responder is unreachable, browsers typically accept the certificate anyway. Liu et al. measured this empirically and found revocation in the web's PKI to be unreliable in practice [8]. The system is weakest exactly when it matters most: at the moment of compromise.

The audit trail is thin. When a key is rotated, the old certificate is revoked and a new one is issued. The cryptographic link proving the new holder is the same entity as the old holder is weak; verifiers must trust the CA's assertion of continuity [1]. Reconstructing the state of the world at a past moment is difficult or impossible.

The high-assurance ceiling: the verifiable LEI At the other end sits the verifiable Legal Entity Identifier (vLEI), built on ISO 17442-3 by the Global Legal Entity Identifier Foundation (GLEIF) [9, 10]. The vLEI is a digital counterpart to the LEI — a 20-character identifier already required for financial reporting in many jurisdictions [10] — that supports automated verification without human intervention.

The vLEI's strength is its governance. GLEIF acts as the root of trust and qualifies a network of vLEI issuers under a published Ecosystem Governance Framework [11]. Every vLEI credential traces to GLEIF through a verifiable chain. Technically, the vLEI uses Authentic Chained Data Containers (ACDCs) and the Key Event Receipt Infrastructure (KERI), which solve key-management and revocation problems that X.509 does not [12].

The vLEI is rigorous. It is also expensive and slow to obtain. A small business securing its website does not need what a bank executing a cross-border settlement needs.

1.3 The need for a bridge

The market has been forced to choose between flawed-but-cheap and excellent-but-costly. Most everyday business interactions sit in between, and have nowhere to land.

The Progressive Assurance Model fills the gap. It defines four levels of assurance, all built on the same underlying cryptographic identifier and credential schema. An organization can start where it can afford to start, then climb the ladder by adding verifiable claims as needs evolve. The same identifier persists; the same evidence chain persists; only the level of assurance changes.

---

2. A Progressive Assurance Model

A uniform approach to identity wastes money on low-risk interactions and under-protects high-risk ones. The model below breaks assurance into incremental levels, so that organizations can match the strength of the credential to the value of the interaction.

2.1 Two axes of assurance

Before the levels, a clarification. Identity assurance is not one number. It has at least two axes that should be evaluated independently.

Reference assurance asks: how confidently is the identifier tied to a specific legal entity in authoritative records? This is the axis GLEIF already grades. Every LEI record carries a corroboration level that classifies how thoroughly the entity's reference data — name, registered address, legal form, country, registration ID — has been validated against authoritative public sources [13]:

- **Fully corroborated.** Every reference data element has been validated against a public authoritative source on GLEIF's Registration Authority List. As of February 2026, 87.64% of LEI records sit at this level [13].
- **Partially corroborated.** At least one element could not be validated against a public source — typically because the local registry does not disclose it. About 3.86% of records [13].
- **Entity-supplied only.** No public authoritative source is available; the LEI issuer reviewed the entity's submission but could not verify it independently. About 8.50% of records, mostly entities exempt from public registration or with privacy constraints [13].

Control assurance asks: how strongly is control of the identifier bound to the entity, an accountable human, or a governed program? This is the axis the rest of this paper develops, in four levels (LoA 0 through LoA 3).

The two axes are independent in principle but related in practice. An LoA 0 affidavit issued against a fully corroborated LEI is meaningfully stronger than the same affidavit issued against an entity-supplied-only LEI, even though both occupy LoA 0 on the control axis. At the top end, GLEIF's Ecosystem Governance Framework presupposes a fully corroborated LEI as input [11]; LoA 3 therefore implies maximum reference assurance by construction. LoA 0 through LoA 2 can ride on top of any corroboration level.

The remainder of §2 holds reference assurance fixed and develops the four levels of control assurance.

2.2 LoA 0: Unambiguous reference affidavit

Business need. At the most basic level, an organization needs a globally unique, unambiguous way to refer to a specific legal entity. This is the digital equivalent of pointing at a row in a registry. Use cases range from regulatory filings (where the LEI is already mandatory in many jurisdictions [10]) to social media platforms attributing corporate accounts to database systems tagging records.

Vetting process. LoA 0 should be cheap and fast. A requester proves they are human (CAPTCHA or API-key registration), then uses a lookup service connected to a registry like the GLEIF LEI index, OpenCorporates, or Dun & Bradstreet to find the entity. The service confirms the entity is currently active.

Limitations. LoA 0 provides zero proof of *control*. It is an affidavit — an assertion of fact for the public record — not a credential conferring entitlement. Anyone can use it to refer to the entity, including competitors, critics, or attackers [14]. Its trustworthiness is bounded by the freshness of the underlying registry data and by the security practices of the affidavit's issuer. Its reference assurance is exactly that of the underlying record.

2.3 LoA 1: Foundational control

Business need. LoA 1 moves from affidavit to credential. It binds a legal entity, a cryptographic identifier, and a controllable digital asset — usually an internet domain or a phone number — into a single verifiable claim. It is meant as an alternative to a basic Domain Validated X.509 certificate, with several advantages [1, 5]:

- **Long lifespan across key rotations.** Because credentials are issued to the identifier, not to a specific keypair, key rotation does not invalidate the credential. The constant churn of cert reissuance disappears [1].
- **Native delegation.** Authority can be cryptographically passed to subsidiaries, employees, or services, with the chain auditable end to end [1].
- **Post-quantum readiness.** Pre-rotation hides the next signing key behind a hash, allowing controllers to roll forward to post-quantum algorithms without changing the identifier or invalidating evidence [5]. RFC 9881 defines a path for X.509 to use post-quantum signatures [1], but adoption across the global cert ecosystem will be slow, expensive, and uneven.
- **A solid audit trail.** Issuance is anchored to a key event log, so the question "was this credential valid on a specific date three years ago?" has a cryptographic answer [12, 15].

Vetting process. LoA 1 builds on LoA 0 and adds proof that the requester controls a digital asset associated with the entity. The strongest common method is to require a DNS record change on a domain linkable to the entity through public records. For small businesses without a web presence, control of a known business phone number can substitute. Both are stronger than email-based verification, which proves only that someone reads a mailbox.

Limitations. The binding is only as strong as the link between the legal entity and the digital asset, and that link can be tenuous. DNS itself is a frequent target of attack. LoA 1 proves *some* agent associated with the organization wants the identifier; it does not prove which agent, or whether they are authorized to act for the organization.

2.4 LoA 2: Human accountability

Business need. For higher-value interactions — verifiable trade, large web3 transactions, legally binding digital contracts — proving organizational control is not enough. The verifier needs to know that a specific, accountable human is acting on the organization's behalf with explicitly delegated authority. LoA 2 maps roughly to the rigor of an Extended Validation X.509 certificate, applied to a more flexible identifier [5].

Vetting process. LoA 2 keeps the domain control proof of LoA 1 and adds two layers:

1. **Human identity assurance.** The individual requesting the credential proves their own legal identity non-repudiably — for example, by presenting a government-issued mobile driver's license or digital passport, or by showing a physical passport in a recorded video call [14].
2. **Verifiable delegation.** The organization issues the individual a credential proving they hold authority to request and manage the organizational credential. Two natural mechanisms exist: a GLEIF Official Organizational Role (OOR) vLEI for organizations already in the vLEI ecosystem [11], or a Generalized Cooperative Delegation (GCD) credential, which lets the issuer attach explicit constraints — scope, jurisdiction, expiry, signing thresholds — to the delegation [16]. Either replaces flimsy evidence (a LinkedIn profile, a corporate email signature) with a cryptographic claim.

Limitations. LoA 2 establishes accountability for the human requester. It does not guarantee that the organization itself has good internal governance over its identity strategy. The vetting flow remains vulnerable to social engineering and man-in-the-middle attacks if not run carefully. Delegation is one-way at this level: a properly authorized requester can still go rogue without tripping automated alarms. LoA 2 also makes no claim about transparency, duplicity detection, or recovery from compromise.

2.5 LoA 3: Robust governance (the LE vLEI)

Business need. LoA 3 is for the cases where trust is doing the most work: cross-border finance, M&A, regulated activity, evidence that must remain verifiable for years. It requires cryptographic proof that the organization's entire identity-management program operates under an externally audited governance framework.

Vetting process. LoA 3 includes everything below it and adds adherence to a published governance framework. The reference implementation is GLEIF's vLEI Ecosystem Governance Framework [11], under which issuance happens via a versioned, publicly transparent policy set. The organization must demonstrate practical competence: AIDs configured with weighted multi-signature thresholds, independent witnesses for compromise detection, and pre-rotation of keys [12, 17]. Multi-signature authority for high-stakes acts is not a novel idea — it appears in legal codes from the Code of Hammurabi forward [1] — but X.509 cannot express it.

Limitations. No system is perfectly secure. Even at LoA 3, a zero-day exploit or a sophisticated social engineering attack can succeed. What LoA 3 guarantees is that the system is as transparent, auditable, and resilient as current technology and governance allow.

---

3. Properties of the Progressive Model

The four levels are the visible part of the model. What gives them their value is a set of architectural properties — open standards, independent verification, and durable evidence — that distinguish progressive assurance from the closed and ephemeral systems that dominate today.

3.1 Open ecosystem versus walled gardens

Trust on the internet should not require permission from a gatekeeper. Three contemporary "verified identity" systems illustrate the alternative model:

- **Telecom (STIR/SHAKEN).** The framework for combating caller ID spoofing in voice networks is built on IETF standards, but governance sits with a telecom-specific authority (the STI-GA in the US), and the certificates are issued for the sole purpose of signing phone calls [18]. STIR/SHAKEN does not produce a credential the organization can carry into a contract negotiation, a supply chain audit, or a press release.
- **Social media verification badges.** Meta's verified blue check, X Premium, and LinkedIn's verified-identity program each issue a platform-locked attestation tied to a paid subscription. The verification is visible only inside the platform, cannot be exported, and disappears if payment lapses. It identifies an account, not an entity.
- **Web/e-commerce (EV certificates).** Extended Validation certificates exist to communicate higher organizational assurance in TLS, but their value is gated by browser root programs whose decisions are not subject to external review [5, 7]. Browser vendors have, at various times, removed the special EV UI treatment entirely without warning, leaving the organizations that paid for EV with assurance that no longer renders.

In each case, the verified status lives inside a single platform, channel, or industry consortium. It cannot move. It cannot be reused. It cannot be combined with evidence from elsewhere.

The Progressive Assurance Model uses open identifier and credential standards, so the organization holds and presents its own credentials wherever it transacts. The platform displaying the credential becomes one verifier among many, not the source of truth.

3.2 Independent verification versus phoning home

Credentials in the progressive model carry the cryptographic material a verifier needs to validate them, without contacting the issuer. This is a structural advantage over X.509, where verification typically requires a real-time OCSP query [5]. The OCSP responder is a centralized point of failure. It can be blocked by a network attacker to force soft-fail acceptance of a revoked certificate [8]. It also leaks browsing behavior to the CA: every OCSP query reveals which sites a user is visiting.

Walled-garden verification is even less independent. In RCS or A2C SMS, the user trusts a checkmark rendered by a messaging app, controlled by the carrier or platform [b]. There is no underlying cryptographic proof a third party could verify if they wanted to. The verification is a UI element, not a credential.

Independent verification puts the proof in the verifier's hands. Anyone, anywhere, online or offline, can validate a credential with the same mathematical certainty.

3.3 Durable audit trail versus ephemeral assertions

Because credentials are bound to stable identifiers that can rotate keys without invalidating issued credentials, evidence remains verifiable for as long as the underlying records are preserved [12, 15]. This solves an awkward X.509 problem: when a key rotates, the cryptographic continuity between the old certificate and the new one is weak, and CRLs and OCSP are designed for real-time status, not historical reconstruction [8].

Other systems are more ephemeral still. The PASSporT token used in STIR/SHAKEN is created for a single phone call and is not designed to be stored in a publicly auditable log [18]. Historical analysis of attestations is, in practice, impossible.

Durable evidence matters for compliance, legal discovery, and supply chain provenance. A credential issued years ago can be proven with the same certainty as one issued today — and, equally important, can be proven *to have been valid at issuance time* even if subsequent events have changed the issuer's key state [15].

3.4 Comparison

Framework	Primary use case	Ecosystem	Verification	Key management	Audit trail	Trust anchor
Progressive Assurance	Universal organizational identity	Open, permissionless	Self-contained, cryptographic	Identifier-based, supports rotation	Durable	Self-certifying / governance body
X.509 / PKI	Web server TLS	Federated, governed	Issuer-dependent (phone home)	Key-based, requires reissuance	Ephemeral, disjointed	Certificate authorities
vLEI	High-stakes finance, regulated activity	Governed, qualified	Self-contained, cryptographic	Identifier-based, supports rotation	Durable	GLEIF

---

4. Interoperability and Evidence Formats

The properties above are not free. They depend on credential formats that can support stable identifiers, rotatable keys, and verifiable chaining. Two formats matter for this paper: ACDCs as the foundational evidence, and SD-JWTs (and similar) as derivative evidence for specific contexts.

4.1 Foundational evidence: ACDCs

A progressive assurance model needs a credential that can serve as a stable, long-lived anchor. The credential must remain valid across decades, even as the keys used to manage it are rotated or upgraded. Authentic Chained Data Containers (ACDCs) are designed for this [12].

ACDCs differ from W3C Verifiable Credentials and SD-JWTs in three ways that matter here [15]:

- **Identifier-based signing.** ACDCs are issued by and to AIDs — autonomic identifiers from KERI, derived cryptographically from their initial public keys [17]. The credential is bound to the identifier, not to a specific keypair. Issuers can rotate keys, change algorithms, or upgrade governance without invalidating credentials issued under a previous key state. This is what makes a durable audit trail possible.
- **Cryptographic chaining.** Every ACDC is sealed by a Self-Addressing Identifier (SAID) — a cryptographic hash of its content. ACDCs from different issuers can be chained into verifiable graphs. An LoA 2 credential can link to the LoA 1 credential it builds on, which links to the entity's reference data, which links to GLEIF. A verifier traverses the chain without consulting an external trust registry [15].
- **A stronger cryptographic foundation.** AIDs support pre-rotation, weighted multi-signature thresholds, and independent witnesses that detect duplicity. None of these are native to X.509 [1, 5, 17]. A holder who suspects compromise can recover by rotating to the pre-committed next key, which an attacker holding only the current key cannot do.

4.2 Derivative evidence: SD-JWTs and others

Different contexts want different formats. SD-JWTs extend the familiar JSON Web Token format with selective disclosure [19] and are gaining ground in web-native authentication [20]. Other formats — ISO mDL, W3C VCs [c] — have their own niches.

The model treats ACDCs as the source of truth and other formats as derivable presentations. An organization holds its identity as an ACDC. When transacting with a system that consumes SD-JWTs, the holder issues a short-lived SD-JWT derived from the ACDC. The SD-JWT inherits its authority from the foundational credential but does not carry its permanence: SD-JWTs are signed by keys, not by identifiers, and lack native chaining for delegation [12, 21].

The analogy is photographic. A RAW file is the archival source, lossless and reusable; a JPEG is exported for a specific use [15]. You don't archive in JPEG. You don't anchor durable identity in a credential format that cannot reproduce its own history.

---

5. Conclusion

The lack of a flexible, reliable method for proving organizational identity has forced businesses into a bad choice: low-assurance X.509 certificates designed for a different problem, or high-assurance specialty systems too costly for everyday use. The middle ground — most actual business interactions — has been left to operate without a fit-for-purpose tool.

Progressive Assurance addresses the gap with four levels of control assurance, each layered on the same cryptographic substrate, riding on top of GLEIF's three levels of reference assurance.

An organization adopts the level its risk profile justifies and climbs as needs change. The same identifier persists. The same evidence chain extends.

The model rests on three properties that together distinguish it from the systems it replaces: an open ecosystem, independent verification, and durable evidence. None of these is decorative. Each addresses a structural failure of the existing landscape. Each is enabled by credential formats — ACDCs first, derivatives second — that were designed for the post-PKI world rather than retrofitted into it.

Strong organizational identity is a precondition for the trust layer the digital economy lacks. Progressive Assurance is one path to building it.

Notes

[a] Coca-Cola Europacific Partners Nederland B.V. is a real licensed bottler in the Coca-Cola system, distinct from The Coca-Cola Company (the Atlanta-based brand owner). The two entities have different LEIs, different boards, different legal exposure, and different counterparties. Conflating them in a digital interaction is exactly the kind of error organizational identity is supposed to prevent.

[b] Rich Communication Services (RCS) and the US 10-Digit Long Code (10DLC) regime for Application-to-Consumer SMS are two further telecom-flavored examples of the same pattern. RCS sender verification is administered by Google and mobile operators; 10DLC verification is administered by The Campaign Registry through designated vetting providers like Aegis and Numeracle. In both cases, the verification is a permission to use a channel, not a portable credential the business can present elsewhere.

[c] The W3C Verifiable Credentials Data Model 2.0 [22] is the most familiar credential format in this space and would work as a starting point for some implementations, but its structural commitments — RDF semantics, context-by-URL, signing keys rather than identifiers — make it lossy in ways that matter for high-stakes, long-lived evidence. The trade is examined in more detail in [15].

References

[1] Hardman, D. 2024. Why X509 Certs Should Be Secondary Evidence of Org Identity. *Codecraft Papers*. https://dhh1128.github.io/papers/x509-prob.html

[2] Hardman, D. 2025. Verifiable Voice Protocol. IETF Internet-Draft draft-hardman-verifiable-voice-protocol. https://dhh1128.github.io/vvp/draft-hardman-verifiable-voice-protocol.html

[3] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and Polk, W. 2008. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280. IETF. https://doi.org/10.17487/RFC5280

[4] CA/Browser Forum. 2025. Ballot SC081v3: Introduce Schedule of Reducing Validity and Data Reuse Periods. https://cabforum.org/2025/04/11/ballot-sc081v3-introduce-schedule-of-reducing-validity-and-data-reuse-periods/

[5] Hardman, D. 2026. Where Trust Bottoms Out: X.509, Certificate Transparency, and KERI's DPKI Architecture. *Codecraft Papers*. https://dhh1128.github.io/papers/wtbo.html

[6] Fox-IT. 2012. DigiNotar Certificate Authority Breach "Operation Black Tulip." https://docslib.org/doc/2849857/diginotar-certificate-authority-breach-operation-black-tulip

[7] Sleevi, R. 2017. Distrust of Symantec TLS Certificates. *Google Security Blog*. https://security.googleblog.com/2017/09/chromes-plan-to-distrust-symantec.html

[8] Liu, Y., Tome, W., Zhang, L., Choffnes, D., Levin, D., Maggs, B., Mislove, A., Schulman, A., and Wilson, C. 2015. An End-to-End Measurement of Certificate Revocation in the Web's PKI. In *Proceedings of IMC 2015*. https://doi.org/10.1145/2815675.2815685

[9] International Organization for Standardization. 2024. ISO 17442-3:2024 Financial Services — Legal Entity Identifier (LEI) — Part 3: Verifiable LEIs (vLEIs).

[10] GLEIF. 2025. Introducing the verifiable LEI (vLEI). Global Legal Entity Identifier Foundation. https://www.gleif.org/en/organizational-identity/introducing-the-verifiable-lei-vlei

[11] GLEIF. 2025. verifiable LEI (vLEI) Ecosystem Governance Framework v3.0. Global Legal Entity Identifier Foundation. https://www.gleif.org/en/organizational-identity/become-a-vlei-issuer-qvi/vlei-ecosystem-governance-framework

[12] Smith, S. M. 2024. Authentic Chained Data Containers (ACDC) Specification. Trust Over IP Foundation. https://trustoverip.github.io/kswg-acdc-specification/

[13] Manolova, Z. 2026. Transforming Data into Opportunities: Metric in Motion — Understanding Corroboration. *GLEIF Blog*. https://www.gleif.org/en/newsroom/blog/transforming-data-into-opportunities-metric-in-motion-understanding-corroboration

[14] Provenant. 2025. Org-Vet Credential Schema. https://github.com/provenant-dev/public-schema/blob/main/org-vet/index.md

[15] Hardman, D. 2025. ACDCs and W3C VCs: Lossless vs. Lossy. *Codecraft Papers*. https://dhh1128.github.io/papers/acdc-vc-diff.html

[16] Provenant. 2025. Generalized Cooperative Delegation (GCD) Credential Schema. https://github.com/provenant-dev/public-schema/blob/main/gcd/index.md

[17] Smith, S. M. 2024. Key Event Receipt Infrastructure (KERI) Specification. Trust Over IP Foundation. https://trustoverip.github.io/kswg-keri-specification/

[18] ATIS. 2020. Signature-based Handling of Asserted Information using toKENs (SHAKEN): Governance Model. ATIS-1000080.v002. https://cstga.ca/wp-content/uploads/2020/07/ATIS-1000080.v002_SHAKEN-Governance-Model.pdf

[19] IETF. 2024. Selective Disclosure for JWTs (SD-JWT). draft-ietf-oauth-selective-disclosure-jwt. https://datatracker.ietf.org/doc/html/draft-ietf-oauth-selective-disclosure-jwt

[20] IETF. 2024. SD-JWT-based Verifiable Credentials (SD-JWT VC). draft-ietf-oauth-sd-jwt-vc. https://www.ietf.org/archive/id/draft-ietf-oauth-sd-jwt-vc-03.html

[21] Hardman, D. 2024. How SD-JWT and ACDC Are Similar and Different. *Codecraft Papers*. https://dhh1128.github.io/papers/sdjwt-acdc.html

[22] W3C. 2023. Verifiable Credentials Data Model v2.0. https://www.w3.org/TR/vc-data-model-2.0/