PERSONAL SUBMISSION

Comments on NIST SP 1800-42A mDL

SUBMITTED BY	**Zen Dodd**
SUBMITTED	**03 May 2026**
RECIPIENT	**National Institute of Standards and Technology**
DOCUMENT	**NIST SP 1800-42A mDL**

I make these comments in my personal capacity.

I attempted to submit these comments through the NCCoE web comment form but the form repeatedly returned “Antibot verification failed, please try again” across multiple browsers, with or without my proxy disabled. I am therefore providing the same comments by email.

These comments concern NIST SP 1800-42A, Digital Identities - Mobile Driver’s licence (mDL): Accelerating Development and Adoption of Digital Identity for Financial Institutions.

My comments focus on privacy, selective disclosure, data minimisation, relying-party implementation, biometric handling, telemetry risk, credential presentation flows, session binding, recovery, assurance metadata, wallet behaviour, verifier trust and threat-model completeness.

Comment 1

Line: 1093

Page: 64

The guide should make relying-party data minimisation an explicit implementation objective, not only a privacy benefit.

This comment relates to lines 1093-1100.

The document correctly notes that mDLs can support selective disclosure and that only data required for opening the account should be shared. Implementers would benefit from a clearer operational checklist.

Suggested change: add guidance that financial institutions should document each requested attribute, map it to a specific CIP/account-linkage requirement, avoid collection of full credential images or unnecessary attributes where attribute-level proof is sufficient and periodically review whether each requested attribute remains necessary.

Rationale: mDL adoption will be more privacy-preserving only if relying parties request and retain the minimum set of attributes required for the transaction. Otherwise, mDLs risk reproducing the overcollection problems of document-upload identity proofing in a newer format.

Comment 2

Line: 1107

Page: 65

The privacy-risk section should more explicitly distinguish between retention obligations of financial institutions and retention practices of verifiers, vendors, identity-management systems and other ecosystem participants.

This comment relates to lines 1107-1111 and 1112-1141.

Suggested change: add a concise implementation recommendation that third-party processors and verification vendors should delete customer data promptly where they do not independently have a legal retention obligation, should not reuse mDL-derived data for secondary purposes and should be contractually bound by purpose limitation, deletion, breach notification and audit requirements.

Rationale: the document notes that financial institutions may have seven-year retention duties, while not all ecosystem parties need the same retention schedule. That distinction should be converted into concrete implementation guidance because unnecessary third-party retention creates avoidable privacy, breach and linkability risk.

Comment 3

Line: 748

Page: 28

The guide should preserve and strengthen the position that biometric verification should be local to the user's device wherever possible.

This comment relates to lines 748-751 and 1093-1100.

Suggested change: add a clear best-practice statement that financial institutions and vendors should avoid server-side biometric matching or biometric-template collection unless there is a specific, documented, legally required and risk-assessed need. Where local biometric authentication is used only to unlock the wallet or approve release, biometric templates should not be transmitted to the financial institution or verification provider.

Rationale: the guide correctly recognises that server-side biometric matching can create risks that outweigh marginal security benefit and that local device authentication better preserves user privacy. This should be made a prominent implementation principle because biometric identifiers are difficult or impossible to replace if compromised.

Comment 4

Line: 957

Page: 51

The threat model should more explicitly include correlation, linkability and telemetry risks across issuers, wallets, verifiers, identity-management systems and financial institutions.

This comment relates to lines 957-962 and 987.

Suggested change: add examples of privacy and security threats where repeated mDL use enables behavioural profiling or cross-transaction correlation including through verifier logs, wallet telemetry, issuer callbacks, device identifiers, IP addresses, session identifiers, analytics tooling or shared vendor infrastructure.

Rationale: the guide correctly notes that the threat model is a starting point and that organisations should conduct their own threat assessments. Because mDLs may become a high-assurance identity rail, linkability risk should be treated as a first-order threat, not only a general privacy consideration.

Comment 5

Line: 1176

Page: 70

The usability section should more clearly state that mDL adoption must preserve practical fallback pathways for users without compatible devices, mDL availability, sufficient digital literacy, reliable mobile access or comfort using digital credentials.

This comment relates to lines 1176-1198 and 1371-1375.

Suggested change: add implementation guidance that financial institutions should provide non-mDL alternatives that are not punitive, excessively slow or materially inferior for ordinary account access, onboarding or reverification.

Rationale: the guide already recognises that alternatives for users without compatible devices or mDLs are essential for a positive onboarding experience. This point should be elevated because mandatory or de facto mandatory mDL usage could exclude users and weaken trust in digital identity adoption.

Comment 6

Line: 1471

Page: 98

The DCQL examples in Appendix F are valuable because they show how structured requests can support relying-party data minimisation. The guide should make clear that DCQL or equivalent request structures should be human-understandable as well as machine-readable.

This comment relates to lines 1471-1474.

Suggested change: add guidance that wallets should display requested attributes in plain language before release including the requesting party, purpose, whether the attribute is required or optional and whether the relying party intends to retain it.

Rationale: selective disclosure only supports user control if the user can understand what is being requested and why. A technically correct attribute request can still fail from a privacy perspective if the user interface obscures the consequence of approving it.

Comment 7

Line: 444

Page: 17

The discussion of revocation and reissuance should be expanded to include account lockout, recovery and fraud-response consequences.

This comment relates to lines 444-449.

Suggested change: add guidance that financial institutions implementing mDL-based account onboarding or reverification should define fallback recovery flows for lost devices, revoked mDLs, wallet compromise, issuer outages and customer disputes. These recovery flows should avoid excessive friction while preventing social-engineering bypass.

Rationale: mDL revocation and remote reissuance are useful capabilities but financial institutions also need resilient customer-support and recovery processes. Weak recovery flows can become the easiest path for account takeover, while overly strict recovery flows can lock legitimate users out of essential financial services.

Comment 8

Line: 1107

Page: 65

The guide should recommend that financial institutions distinguish between evidence needed to complete a regulated identity process and evidence retained for future audit.

This comment relates to lines 1107-1141.

Suggested change: add guidance that institutions should prefer retaining a verifiable transaction record, issuer assurance metadata, policy decision record and minimal attributes over retaining full mDL payloads, credential images or unnecessary identity attributes.

Rationale: regulated institutions may need to demonstrate that identity verification occurred but this does not always require retention of every attribute disclosed during the mDL transaction. Separating proof-of-process from full identity-data retention would reduce breach impact while preserving auditability.

Comment 9

Line: 957

Page: 51

The threat model should include compromise or misuse of vendor analytics, SDKs and telemetry embedded in wallet, verifier or identity orchestration components.

This comment relates to line 957 and the surrounding threat-model discussion.

Suggested change: add a threat example covering third-party analytics or SDK collection of device identifiers, event logs, session metadata, behavioural telemetry or error data during mDL presentation flows.

Rationale: mDL systems may be privacy-preserving at the credential protocol layer while still leaking sensitive information through ordinary application telemetry or vendor tooling. Implementers should be prompted to review SDKs, analytics and logging paths as part of their mDL threat assessment.

Comment 10

Line: 1176

Page: 70

The guide should address user comprehension during consent and attribute release.

This comment relates to lines 1176-1198.

Suggested change: add guidance that wallets and relying-party flows should avoid dark patterns, confusing consent wording, preselected optional disclosures, bundled release prompts or interface designs that make refusal difficult.

Rationale: selective disclosure depends not only on protocol capability but also on user-interface design. A user cannot make a meaningful privacy choice if optional attributes are presented as mandatory, if the relying party's purpose is unclear or if refusal results in avoidable friction.

Comment 11

Line: 607

Page: 24

The centralised IDMS and verifier pattern is practical for financial institutions but the guide should more explicitly describe the compensating controls needed because centralisation concentrates risk.

This comment relates to lines 607-636.

Suggested change: add guidance that centralised IDMS/verifier deployments should include strong service segmentation, least-privilege service identities, mutual TLS or equivalent service authentication, strict API authorisation, tamper-resistant audit logging, separation of duties for administrative access, key isolation, data minimisation between services and explicit blast-radius analysis.

Rationale: centralisation reduces integration complexity and avoids parallel identity silos but it also creates a high-value orchestration point for identity data, verifier outputs, tokens, customer records and policy decisions. Implementers should not read the centralised pattern only as an enterprise convenience pattern; it should be presented as a pattern requiring explicit resilience, containment and abuse-detection controls.

Comment 12

Line: 654

Page: 26

The SaaS architecture discussion should more explicitly translate SaaS drawbacks into implementation requirements for production financial-institution deployments.

This comment relates to lines 654-666.

Suggested change: add guidance that SaaS verifier, IDMS and identity-orchestration components should be assessed for tenant isolation, administrator access controls, support-access logging, subcontractor access, data-location commitments, incident notification timelines, key-management boundaries, logging/telemetry minimisation, deletion guarantees, vulnerability disclosure handling and audit rights.

Rationale: the guide correctly notes that SaaS introduces new threats and boundary crossings. Financial institutions will need concrete supplier-assurance and contractual-control guidance, not only a general statement that SaaS increases exposure.

Comment 13

Line: 850

Page: 38

The temporary verifier-session association and attribute transfer path should include explicit requirements for session binding, expiry and one-time use.

This comment relates to the account application flow, particularly the steps where mDL attributes are transmitted to the verifier, a session identifier is returned to the IDMS and temporary storage associates the verifier session identifier with the active application.

Suggested change: add guidance that verifier session identifiers and temporary mDL attribute associations should be short-lived, one-time-use, bound to the browser session and transaction context, invalidated on cancellation or completion, protected against session fixation and purged from temporary storage after use. The verifier mediation service should also enforce strict authorisation so a session identifier cannot be replayed or queried by another application context.

Rationale: the session identifier becomes a bridge between the verifier transaction and the banking application workflow. If this binding is weak, stale, reusable or insufficiently authorised, it can become a practical target for replay, confused-deputy or account-linkage attacks.

Comment 14

Line: 1183

Page: 70

The QR-code reliability recommendation should more clearly balance usability and security.

This comment relates to lines 1183-1190.

Suggested change: add guidance that QR-code validity periods should be risk-based, that regenerated QR codes should invalidate prior codes, that QR codes should be bound to the requesting origin, transaction, verifier and session and that user-facing error messages should clearly distinguish between expiration, cancellation, restart and failure.

Rationale: extending QR-code validity may improve usability but it can also increase the time window for misuse if binding and invalidation are weak. Implementers need guidance that improves reliability without creating a broader replay or relay window.

Comment 15

Line: 895

Page: 45

The high-risk transaction authorisation section should caution against making mDL re-verification a routine or sole step-up mechanism.

This comment relates to lines 895-902.

Suggested change: add guidance that mDL re-verification should be one risk signal among several, not a default replacement for phishing-resistant authentication, transaction risk analysis, device signals, behavioural risk indicators or out-of-band fraud controls. The guide should also recommend limiting mDL re-verification to cases where identity re-verification is genuinely relevant to the risk decision.

Rationale: mDLs are valuable for identity proofing and certain high-risk actions but repeated use of an identity credential for routine step-up may increase linkability, user fatigue, overcollection and the chance that users become conditioned to release identity attributes too often.

Comment 16

Line: 979

Page: 51

The guide should recommend that downstream relying-party systems receive structured assurance metadata, not just a binary verification result.

This comment relates to lines 979-984 and the broader discussion of conveying mDL verification results to the relying party.

Suggested change: add guidance that verifier-to-IDMS and IDMS-to-banking-system assertions should include structured, machine-readable assurance information such as issuer identifier, credential type, verification time, holder-authentication method, verifier identity, presentation protocol, nonce/challenge binding, audience, purpose, requested attribute set and confidence/caveat information where available.

Rationale: financial institutions need to make risk decisions based on the quality and context of the verification event. A simple “verified” outcome can hide important distinctions between issuer assurance, wallet assurance, holder binding, local biometric verification, server-side biometric verification, protocol path and transaction context.

Comment 17

Line: 884

Page: 42

The passkey provisioning discussion should more explicitly bind passkey registration to the verified identity-proofing session.

This comment relates to lines 884-888.

Suggested change: add guidance that passkey registration should occur only after a fresh, successful, transaction-bound identity verification flow, using an origin-bound WebAuthn challenge, with clear account/application binding and audit logging. Adding or replacing passkeys later should require a risk-appropriate recovery or re-verification process rather than relying only on an already-authenticated session.

Rationale: passkeys are phishing-resistant but account takeover risk often moves to enrolment, replacement and recovery flows. The guide would be stronger if it treated passkey lifecycle events as part of the identity assurance chain rather than only as a convenient post-onboarding authenticator.

Comment 18

Line: 1107

Page: 65

The guide should elevate “no issuer notification during presentation” from a privacy benefit to a deployment objective.

This comment relates to lines 1107-1111 and the Table 7 discussion of device retrieval preventing issuer surveillance across financial transactions.

Suggested change: add guidance that implementations should prefer device retrieval or other presentation models that avoid issuer callbacks during ordinary transactions. Any design requiring issuer contact, online status checks or server retrieval during presentation should be assessed as a higher privacy-risk architecture and should include mitigations against issuer-side transaction logging, correlation and behavioural profiling.

Rationale: one of the strongest privacy properties of mDL presentation is the ability to verify attributes without notifying the issuer each time the credential is used. If this property is weakened in production deployments, the ecosystem may introduce a durable transaction-surveillance risk.

Comment 19

Line: 850

Page: 38

The use of a hashed mDL document number and issuing authority as a persisted applicant identifier should include guidance on keyed hashing and correlation risk.

This comment relates to the account application flow where the mDL document number and issuing authority attributes are hashed and persisted as an identifier for the applicant.

Suggested change: add guidance that if document number and issuing authority are used to derive a persistent identifier, implementers should treat the derived value as personal information and use a keyed HMAC or equivalent construction with institution-specific secret material and domain separation rather than an unsalted or globally reproducible hash. The guide should also recommend key rotation planning, access control and review of whether a stable mDL-derived identifier is necessary.

Rationale: deterministic hashes of low-entropy or structured identifiers may be vulnerable to guessing, enumeration or cross-institution correlation if implemented poorly. Hashing does not automatically remove privacy risk; the derived identifier can still function as a durable identity handle.

Comment 20

Line: 821

Page: 33

The SSN/TIN validation flow should include explicit data-handling guidance for sensitive identifiers.

This comment relates to lines 821-824 and the later account application steps where the banking system retrieves and stores the applicant’s SSN in an encrypted token.

Suggested change: add guidance that SSN/TIN values should be collected only where required, masked in user interfaces, excluded from ordinary logs and telemetry, encrypted with strong key management, tightly access-controlled and retained only according to applicable legal and business requirements. Backchannel retrieval of SSN/TIN data should require strict API authorisation, service authentication and audit logging.

Rationale: SSNs and TINs are high-risk identifiers. mDL adoption should not improve one part of identity proofing while leaving sensitive identifier handling under-specified.

Comment 21

Line: 562

Page: 23

The trust-service discussion should address resilience and failure-mode handling for issuer public-key access and trust metadata.

This comment relates to lines 562-566.

Suggested change: add guidance that relying parties should define cache lifetimes, stale-cache policies, outage handling, fail-open/fail-closed behaviour, key rotation handling, revocation/status handling, metadata integrity checks and monitoring for trust-service unavailability or compromise.

Rationale: trust services can simplify relying-party integration but they also become operational dependencies. Financial institutions need predictable behaviour when a trust service is unavailable, stale, inconsistent or suspected of compromise.

Comment 22

Line: 1819

Page: 59

The malicious-verifier threat should include stronger ecosystem mitigations than user education alone.

This comment relates to lines 1819-1839.

Suggested change: add guidance that wallets should display authenticated verifier identity, requesting domain or application identity, credential request purpose, requested attributes and retention intent before release. The guide should also encourage development of verifier registration, certification or trust-list mechanisms for high-assurance relying parties, with care to avoid creating unnecessary centralised surveillance.

Rationale: users are poorly positioned to distinguish legitimate and malicious verifiers based only on education. A malicious verifier can harvest high-value identity attributes while appearing legitimate. Wallet-mediated verifier identity and purpose display would materially improve user decision-making and reduce phishing-style identity harvesting.

Comment 23

Line: 797

Page: 33

The guide should include a forward-looking caution that same-device flows require a separate threat model and should not inherit all assumptions from the demonstrated cross-device flow.

This comment relates to lines 797-798 and the earlier statement that same-device scenarios will be common but were out of scope for this phase.

Suggested change: add a short section or caution noting that same-device flows may introduce different risks including mobile-app overlay attacks, malicious WebViews, app-to-app request confusion, deep-link handling risk, mobile malware, user-agent ambiguity and weaker user separation between requesting party and wallet.

Rationale: same-device presentation is likely to be common in production financial-services deployments. Because this draft demonstrates cross-device presentation, implementers should be warned not to assume that cross-device security and usability findings transfer directly to same-device flows.

Comment 24

Line: 1771

Page: 58

The social-engineering and identity-theft threat during issuance should acknowledge that liveness and deepfake detection are not complete mitigations.

This comment relates to lines 1771-1782.

Suggested change: add guidance that issuers should treat liveness and deepfake detection as one layer in a broader anti-fraud process, supported by risk-based enrolment controls, rate limiting, device and session risk checks, human review for high-risk cases, fraud reporting loops and post-issuance anomaly monitoring.

Rationale: liveness and deepfake detection technologies can reduce risk but attackers adapt quickly and may bypass single controls. If an attacker obtains a legitimately issued mDL through enrolment fraud, downstream relying parties may treat the credential as high assurance. Issuance fraud therefore deserves layered controls.

Comment 25

Line: 1702

Page: 57

The wallet-leakage mitigation should more clearly distinguish local transaction history from synchronised or cloud-backed transaction history.

This comment relates to lines 1702-1711.

Suggested change: add guidance that wallet transaction logs should be local by default where possible, encrypted at rest, protected from other applications and understandable to the holder. If transaction logs are synchronised to a back end, they should use strong encryption, minimal metadata, limited retention, clear user notice and controls for deletion/export where appropriate.

Rationale: wallet transaction history can be useful for transparency but it may also become a sensitive behavioural record of where and when a user presented identity credentials. Synchronised logs increase breach, subpoena, analytics and correlation risk unless tightly controlled.

Comment 26

Line: 1244

Page: 72

The wallet certification discussion should include certification of privacy properties, not only security properties.

This comment relates to lines 1244-1252.

Suggested change: add guidance that wallet certification criteria should include privacy requirements such as data minimisation, local processing where feasible, transaction-log protection, no unnecessary telemetry, clear attribute-release display, accessible consent flows, deletion controls and restrictions on secondary use of wallet-derived data.

Rationale: financial institutions need confidence in wallet security but mDL adoption will also depend on wallet privacy behaviour. A wallet can be cryptographically secure while still creating privacy risk through telemetry, logging, poor user display or unnecessary cloud synchronisation.

Comment 27

Line: 1267

Page: 73

The presentation-protocol fragmentation discussion should recommend test suites and interoperability profiles focused on negative and abuse cases, not only successful presentation.

This comment relates to lines 1267-1278.

Suggested change: add guidance that standards bodies and implementers should test rejection paths, malformed requests, overbroad attribute requests, expired or replayed requests, wrong-audience assertions, malicious verifier metadata, user cancellation, partial disclosure and inaccessible fallback cases across OpenID4VP, ISO/IEC 18013-7 Annex C and DC API implementations.

Rationale: interoperability problems often appear in edge cases and failure paths. For financial institutions, safe rejection and predictable failure behaviour are as important as successful happy-path credential presentation.