Comment Template for: NIST SP 800-63-4 Suite (Second Public Draft)

Please submit responses to dig-comments@nist.gov by October 7, 2024.

Organization:	*Google*
Name of Submitter/POC:	*Omid Ghaffari-Tabrizi*
Email Address of Submitter/POC:	[redacted]

Comment #	**Publication** (Base, 63A, 63B, 63C)	**Section**	**Page #**	**Line #**	**Comment** (Include rationale for comment)	**Suggested Change**
1	63-Base	throughout			While we appreciate that the main audience of this publication is US government agencies and contractors, we expect many other entities will use this as the basis of their own standards. As such, we'd recommend calling out aspects of these standards that are more about compliance with other governmental policies and less about xAL per se. Specifically, some of the subscriber notification requirements are likely not applicable to most enterprise contexts. Perhaps more use of the "operated by or on behalf of federal agencies" qualifier could be helpful in such instances (e.g. 63b 4.1.2.1, 63c 3.4, 63c 3.4.3, 63c 3.6, 63c 3.7.1, 63c 3.9, 63c 3.15, 63c 4.6.7).	
2	63-Base	2. General	iii	198	"What specific implementation guidance, reference architectures, metrics, or other supporting resources could enable more rapid adoption and implementation of this and future iterations of the Digital Identity Guidelines?" The user controllability of the digital identity solutions, especially the pattern with a subscriber controlled wallet, is understated throughout the guideline. With a subscriber controlled wallet, the user has full control of the digital identity and the exact PII data attributes being presented to RP. This advantage makes a significant difference on the traditional identity solutions and digital identity solutions.	
3	63-Base	2. General	iii	203	"What applied research and measurement efforts would provide the greatest impacts on the identity market and advancement of these guidelines?" 1. Overall the privacy advantages on the federated digital identity model with subscriber-controlled wallet example (Fig. 5.) are not clearly stated within the guideline. Given the CSP-wallet-RP model, it is a significant advantage that RP does not need to actively communicate back to the CSP during the proofing process thus CSPs would not be able to track the subscribers on their actions on the RP side. (The classic phone home problem) 2. Zero-knowledge proof would be able to significantly eliminate the risk of applicant being tracked under a collusion between multiple RPs via unique identifiers being exchanged during authenticated sessions. 3. Zero-knowledge proof would be able to significantly eliminate the risk of applicant being tracked under a collusion between CSP and RP via unique identifiers being exchanged during enrollment and authenticated sessions. 4. Selective disclosure, which is already implemented by ISO 18013-5 mDL, gives user the ability to disclose a subset of the digital identity while the RP can still easily verify the authenticity and integrity of the identity.	
4	63-Base	Fig. 5	20	898	"Federated digital identity model with subscriber-controlled wallet example (Fig. 5.)" The term “federated” on Fig.5. requires additional clarification to avoid confusion. Although the protocol between wallet and RP can be considered as a “federation protocol”, there is no federation between a CSP and RP. A reference or name without “federated” would make it much clearer to the audience.	
5	63-Base	2.3.1	12	702	Recommend consistent use of "type of factor" e.g. "Multi-factor authentication (MFA) refers to the use of more than one distinct factor." => "Multi-factor authentication (MFA) refers to the use of more than one type of factor."	
6	63A	2.1.3 Identity proofing Types	8	595	The standard specifies remote unattended identity proofing vs remote attended identity proofing, where the difference is essentially if there is an agent online through a secure video session. Evidence submitted for personhood and ownership should be categorized as remote attended identity proofing regardless of when the agent reviews it, as long as liveness and ownership verification are passed. Currently, the proofing method is categorized as "unattended' as long as there's no live human agents on the video session with the user. All un-attended verification methods are considered as lower assurance level than attended verification, which does not align with currently accepted best practices. As long as the verifier collects confident enough evidence from the user to prove their personhood (i.e. liveness check) and ownership (e.g., selfie matching), it doesn't matter if the human reviewing agent is interacting with the user online or reviewing the evidence in a separate session.	
7	63A	2.5.1 Identity Verification Methods	14	784	Greater clairification is required around Confirmation Code Verification, especially where an individual is able to demonstrate control of a piece of identity evidence through the return of a confirmation code, consistent with the requirements specified in Sec. 3.1.8. Confirmation code equivalent technology should also be considered here, specifically for multi-factor authentication. E.g., confirm on a different trusted device on file.	

8	63A	2.5.1 Identity Verification Methods	15	817	“Knowledge-based verification (KBV) or knowledge-based authentication SHALL NOT be used for identity verification.” The KBV / KBA shall not be used for identity verification by itself, but KBV/KBA should still be considered an identity verification method along with other methods and the standard should account for this fact. E.g., PIN code on Japan My Number Card, SSN verification question challenge.	
9	63A	3.1.2.1 CSP Fraud Management	18	903	"CSPs SHALL implement the following fraud check for all identity proofing processes: 1. Date of Death Check – Confirm with a credible, authoritative, or issuing source that the applicant is not deceased. Such checks can aid in preventing synthetic identity fraud, the use of stolen identity information, and exploitation by a close associate or relative." CSP shall perform a date of death check (liveness check) only if needed. Performing the liveness check for all identity verification processes is unrealistic, especially for remote unattended identity verification.	
10	63A	3.1.2.1 CSP Fraud Management	18	907	"6. CSPs SHOULD implement - but are not limited to - the following fraud checks for their identity proofing processes based on their available identity proofing types, selected technologies, evidence, and user base: SIM Swap Detection – Confirm that the phone number used in an identity proofing process has not been recently ported to a new user or device. Such checks can provide an indication that a phone or device was compromised by a targeted attack." SIM Swap may not be the best and only way to detect a compromised attack. The method suggestion here should be more generic: Device fraud indicator check, including device wifi, power, movement, rooted, simulator etc.	
11	63A	3.1.2.1 CSP Fraud Management	19	943	"11. CSPs SHALL establish a technical or process-based mechanism to communicate suspected and confirmed fraudulent events to RPs." This should be modified so that a CSP shall suspend or revoke the identity, if possible. If not possible, CSP then shall communicate fraudulent events to RPs. Communication is a post-process mitigation after a fraud event has happened. As the CSP, a better mitigation is to make the identity no longer usable by suspending or revoking it, so the damage would be fully controlled. Given suspending or revoking by a CSP may not always be feasible, communicating with RP could be a back-up mitigation.	
12	63A	2.2. Core attributes	9	626	"Government Identifier: A unique identifier which is associated with the applicant in government records (e.g., SSN, TIN, Driver’s License #)" NIST regulations should account for cases in EUDIW PID credentials where there is no mandated regulation to have a unique identifier.	
13	63A	2.2. Core attributes	9	628	"Physical Address: A physical address to which the applicant can receive communications related to the proofing process;" PO box addresses are not a reliable signal. In those instances where there is no alternative, there should be guidance provided to outline how other signals should be weighted or utilized to make up for this weakness.	
14	63B	2.{1,2,3}.3 (reauth)	5	522	Whether and how frequent reauthentication is necessary should be a function of the strength of the session credential, and not of the Authentication Assurance Level. Thinking of possible threats, the following attacks come to mind: 1) theft of authenticator secrets (through phishing, key extraction, etc.) 2) theft of session credentials 3) local attackers (someone taking over a signed-in user’s device) Reauthentication mitigates (2) and (3), but not (1). (3) can be mitigated with screen locks (which the document should acknowledge). This leaves (2) as the threat that reauthentication can and needs to mitigate. If the session credential is strong (say, a DBSC-bound cookie), there is surely less need for quick reauthentication than if the session credential is weak (a normal cookie). The current treatment of reauthentication ignores this. The discussion of reauthentication should move from Section 2 to Section 5, and should take into account session credential strength	
15	63B	2.3.2 (aal3 verification)	8	606	AAL3 requires FIPS. Additional clarification is required as to whether that is consistent with other areas of the standard or if it should have the federal agencies caveat and/or the "relevant standards" qualifier.	
16	63B	2.5 (summary table)	10	Fig. 1	The "reauthentication" row doesn't match the details in 2.2.3 and 2.3.3. There are no "shall" inactivity requirements at any AAL; only "should". This should be revised to (e.g.) "24 hours overall (required), 1 hour inactivity (recommended)".	
17	63B	3.1.3 (Out-of-Band Devices)	17	857	The document considers only two types of out-of-band (secondary) devices: those that generate (or receive from the CSP) a secret that the user then needs to transfer to the primary device, and those that accept a secret from the user that the user got from the primary device. Another, third (and rather popular), form of out-of-band device use is for the out-of-band device to simply let the user confirm the sign-in (e.g. Duo Push (https://duo.com/product/multi-factor-authentication-mfa/authentication-methods/duo-push) or Google Prompt (https://support.google.com/accounts/answer/7026266?hl=en), etc.). This form has equivalent security properties to the two methods already mentioned in the document and should be documented as well.	

18	63B	3.2.2 (throttling)	28	1211	"verifier SHALL limit consecutive failed authentication attempts using one or more specific authenticators on a single subscriber account to no more than 100" Greater clarification is required around how long this is intended to last (eg, is this "forever") and what resets it (eg, any successful auth or just one using that specific authenticator)?	
19	63B	3.2.9 (restricted authenticators)	35	1431	Greater clarification is required to expand on what this means and whether it should just be informative.	
20	63B	3.2.10 (activation secrets)	35	1461	"For authenticators that are usable at AAL3, verification of activation secrets SHALL be performed in a hardware-protected environment (e.g., a secure element, TPM, or TEE). At AAL2, if a hardware-protected environment is not used, the authenticator SHALL use the activation secret to derive a key used to decrypt the authentication key." Greater clarification is required as this could be deemed excessive given that a biometric activation can satisfy AAL3.	
21	63B	3.2.11.2 (wireless connections)	37	1506	This seems to generally assume "wireless" means RF communication (other than the mention of optical transmission of a secret). Is it sufficiently general or should it explicitly mention the possibility of optical (visible or IR), audio, magnetic, or other potential wireless links? Similarly, would a pairing code on a magnetic stripe be acceptable under bullet (2)?	
22	63B	4.1.2.1 (binding an additional authenticato	40	1600	"When any new authenticator is bound to a subscriber account, the CSP SHALL ensure that the process requires authentication at either the maximum AAL currently available in the subscriber account or the maximum AAL at which the new authenticator will be used, whichever is lower. For example, binding an authenticator that is suitable for use at AAL2 requires authentication at AAL2 unless the subscriber account currently has only AAL1 authentication capability." Greater clarififcation is required related to the process for upgrading a subscriber from AAL1 to AAL2, as this example seems to contradict the "shall" in the preceding sentence. For example, 4.2.2.2 (below) could be read to mean that a password plus a recovery code can establish a session in which the subscriber can bind an AAL2 authenticator, or read another way, they could "just bind one" if they currently only have AAL1 capability.	
23	63B	4.1.2.2 (External Authenticator Binding)	40	1616	This section should explicitly discuss phishing-resistant ways of binding an external authenticator. Right now, the methods described in Section 4.1.2.2 are phishable. But phishing-resistant methods exist: for example, today a user can choose to create a passkey on a primary device (say, a laptop), and choose to have the passkey created on their external authenticator (mobile phone). The hybrid CTAP protocol allows for this to happen in a phishing-resistant manner. This approach is already discussed in section 3.2.11.2 in the context of using an external authenticator, and should similarly be discussed here in the context of binding an external authenticator.	
24	63B	4.2.2.2 (recovery at aal2)	40	1616	Clarify whether IAL1 proofing is required to reach AAL2. It looks like password + a code allows establishment of an AAL2, but see question on Sec 4.1.2.1 above.	
25	63B	5.1 (session bindings)	48	1879	Greater clarification is required around derived sessions and session bindings as it relates to 63B, 5.1, and more specifically as it relates to the relationship between (e.g.) apps or stateful tools acting on behalf of a user and "sessions". For example, Sec 5.1.2 (access tokens) says "The OAuth access token and any associated refresh tokens could be valid long after the authentication session has ended and the subscriber has left the application" and Sec 5.1 (session bindings) says "A session SHOULD inherit the AAL properties of the authentication event that triggered its creation." For example, if a user signs into an email client, can it check for new mail and notify the user about it after the idle timeout? Also, it is not clear why session secrets such as cookies should be deleted when a device reboots or an app restarts, even if the time limit or inactivity limit has not been reached yet. Furthermore, this recommendation could be improved by taking into account the security of the session credential and providing additional clarification. For example, should a DBSC-bound cookie be subject to the same deletion requirements as a normal bearer-token cookie? This section should be brought in line with accepted practices (and soften the language perhaps to a MAY, or completely remove the non-retention part), and make sessions credential deletion dependent on the security of the session credential.	
26	63B	5.2 (reauth footnote 1)	51	1970	Thank you for addressing managed devices. However, there is no other mention of them. Greater clarification is required as to how they relate to reauth. For example, it would be nice to allow unlocking of a managed device to satisfy reauth after an idle-timeout.	
27	63B	Appendix B4	90	2994	This appendix mentions a Table 5, which doesn’t seem to exist in the document	
28	63C	2.4 (FAL3)	6	548	Request guidance in relation to "the key material used to authenticate the RP and IdP to each other is associated with the identifiers for the RP and IdP in a static fashion using a trusted mechanism". The language above creates uncertainty about the usage of published SAML XML metadata and OIDC key discovery.	Change recommended: strike out 'in a static fashion' given that contradicting guidance regarding OIDC key discovery is provided elsewhere in this document, .e.g, 3.12.2 "Public keys for verifying digital signatures ... MAY be fetched by the RP in a secure fashion at runtime, such as through an HTTPS URL hosted by the IdP."
29	63C	2.5	6	573	Request normative guidance for protocol specific xAL identifiers. For example an RP may signal xAL requirements to the IdP. For example OIDC acr_values parameters such as acr_values=urlencode(ial1 aal2 fal3).	

30	63C	2.5	7	576	Request normative guidance for protocol specific xAL identifiers transmitted to an RP by the IdP including if specific authentication mechanisms or metadata is required for RP validation. For example xAL identifiers met by the subscriber at the IdP may be encoded in the OIDC acr claim such as {"acr": "ial1 aal2 fal3"}. Alternatively specific attributes which may be used by a RP to calculate xAL compliance may be encoded in an assertion such as the "phrh" acr value (reference) or amr values relevant to AALs such as "pwd" or "mfa" (reference). Alternatively alternative OIDC standards [1] for IAL data may be used if field level information must be transmitted to a relying party. If these standards are desired or required please provide normative guidance on if OIDC eKYC-IDA [2] identifiers [3] are sufficient for transmitting IAL verified claims. [1] https://openid.net/specs/openid-connect-4-identity-assurance-1_0-16.html [2] https://openid.net/specs/openid-ida-verified-claims-1_0.html [3] https://bitbucket.org/openid/ekyc-ida/wiki/identifiers	
31	63C	3.3 (federated identifiers)	14	807	"Federated identifiers SHALL contain no plaintext personally-identifiable information (PII), such as usernames, email addresses, or employee numbers, etc." -- Email SAML NameIDs are common in Enterprise SSO. Suggest relaxing this at least in FAL1 and possibly in higher levels when the assertion is transported securely from the IdP to the RP.	
32	63C	3.4	17	894	The trust agreement should be extended to specify if the subscriber should use the same bound authenticator with both the IdP and RP, requiring that the IdP assertion includes bound authentication identifiers such as an attestation certificate.	
33	63C	3.7	24	1107	"Prior to being provisioned, the RP subscriber account does not exist and has no associated data record at the RP." This section contemplates scenarios where a password or other direct credential is associated with a subscriber account at the RP that was initially a federated account. It seems to preclude the case where an existing subscriber account at the RP (using direct authentication) can be later associated with an IDP and a corresponding federated identifier. We suggest that the language be changed to indicate that this is also acceptable. We have similar questions where Sec 4.6.3 says "Pre-provisioned accounts SHALL be bound to a federated identifier at the time of provisioning."	
34	63C	3.14 (HoK)	37	1570	"the RP SHALL ensure that the authenticator can be uniquely resolved to the RP subscriber account" The use of "uniquely" here (and the requirement of associating it with a specific subscriber at all) is likely to be problematic for machine certs. Greater clarification is required related to when uniqueness improves the FAL. For example, shouldn't it be sufficient at least in some cases for the RP to be confident that the subscriber used the same computer for the IdP auth as for the RP auth? E.g., a factory floor employee uses a shared terminal in a secured facility.	
35	63C	3.15 (bound authenticators)	38	1592	"A bound authenticator SHALL be unique per subscriber at the RP". As above, greater clarification is required as to why this is a "SHALL" since it seems overly strict; using the same hardware key for two accounts is likely to be satisfactory, at least at some FALs.	
36	63C	3.15.1	39	1619	features in FIDO2/WebAuthn that have been implemented in browsers [1] that enable individual attestation of enterprise enrolled authenticators (with enterprise controlled subscriber consent behavior [2]). IdPs such as Microsoft Entra ID support authenticator attestation [3] and vendors such as Yubico support enterprise attestation [4] in addition to batch attestation. This proposal provides a middle ground between subscriber managed bound authenticator issuance and kiosk based issuance. The trust agreement may require an IdP can ensure a subscriber uses a specific bound authenticator known to the IdP at the RP when the RP is present in both the IdP and authenticator allowlist. Where enterprise attestation is not possible, batch attestation still serves to reduce the likelihood that an unknown bound authenticator is issued for the subscriber account at the RP as bound authenticators are restricted to a specific batch of vendor authenticators. If this proposal is accepted normative or informative guidance for implementers is appreciated as the IdP and RP must agree on a method of conveying attestation certificate information and existing standardized claims for holder of key assertions such as the OIDC cnf [5] and x5c [6] claims may be misinterpreted by a RP as a holder of key assertion instead of conveyance of bound authenticator attestation certificate. Note: this proposal is intended to be considered alongside feedback about trust agreements (3.4), RP requirements (13.6.1), and bound authenticator informative examples (10.7). [1] https://www.chromium.org/security-keys/ [2] https://chromeenterprise.google/policies/#SecurityKeyPermitAttestation [3] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-fido2-hardware-vendor [4] https://developers.yubico.com/WebAuthn/Concepts/Enterprise_Attestation/ [5] https://datatracker.ietf.org/doc/html/rfc7800	The trust agreement may require an IdP MAY identify the bound authenticator or a batch of bound authenticators used by a subscriber to authenticate with the IdP as a part of RP bound authenticator issuance using an enterprise or batch attestation certificate. The IdP MUST NOT identify a specific bound authenticator to the RP if the RP is not present in the IdP's enterprise attestation allowlist or if disallowed by the trust agreement and instead MUST present the bound authenticator batch attestation certificate. When indicated by the IdP the RP SHALL attest generated credentials from the bound authenticator and verify the attestation object using the attestation certificate provided by the IdP. The RP must store the attestation certificate or attestation certificate thumbprint in addition to the bound authenticator credential public key such that the RP bound authenticator may be identified using IdP provided data in future FAL3 authentication flows.
37	63C	3.15.2	40	1646	Request protocol specific normative guidance for identifiers used in IdP indication of bound authenticator use.	

38	63C	3.16.1	42	1693	Propose additional RP required behavior for bound authenticator use in FAL3 authentication based on IdP assertion contents which restrict what bound authenticators may be used by a subscriber to authenticate to the RP. This feedback is paired with IdP constrained bound authenticator issuance feedback for section 13.5.1. When supported by the bound authenticator (enterprise attestation enabled) and RP (present in the IdP and authenticator enterprise attestation allowlist) these additional requirements enable an IdP to require authentication using the same bound authenticator at the IdP and RP. If these proposed additions are accepted, we suggest the means of communicating attestation certificates between the IdP and RP follow normative or informative guidance requested as a part of bound authenticator issuance feedback in section 13.5.1. Additional normative or informative guidance for implementers is requested for the means by which IdPs may disallow subscriber registered bound authenticator use.	5. The trust agreement may require an IdP to include a bound authenticator attestation certificate or certificate thumbprint in the assertion presented to the RP. When required by the trust agreement the RP MUST limit the acceptable bound authenticators for the subscriber account using the provided attestation certificate or certificate thumbprint by comparing the IdP supplied value with attestation certificates or certificate thumbprints stored by the RP during subscriber bound authenticator issuance. The RP SHALL reject authentication attempts using subscriber bound authenticators which do not match the IdP supplied attestation certificate or certificate thumbprint value. The RP MAY initiate bound authenticator issuance using the IdP provided enterprise attestation data after rejecting authentication with an unknown bound authenticator.
39	63C	4.3.1	46	1779	"a priori" should always be two words	
40	63C	4.9 (assertion contents)	62	2348	"All assertion SHALL contain sufficient... [IAL, AAL, and intended FAL]" and "At FAL3, the assertion SHALL include ... or ... An indicator that verification of a bound authenticator is required to process this assertion." Greater clarification is required as to the intent as this is challenging unless there's an explicit proposal for both OIDC and SAML in terms of how to do this. Consider allowing "the terms of the trust agreement" to declare this as per Sec. 2.5, especially for lower FAL levels, since this requirement would render any existing implementation that fails to explicitly indicate non-assertion of IAL, AAL, or intended FAL noncompliant for even FAL1.	All assertions contain sufficient ... unless the trust agreement specifies an obligation for the RP and IDP to always meet the required xAL for specified subscribers or for otherwise clearly identifiable assertion workflows.
41	63C	4.10 (authn requests)	65	2400	"A cryptographic nonce". Clarification is required as to whether the SAML request "ID" can satisfy this requirement if implemented to carry sufficient randomness (e.g. as in Sec. 3.12.1).	
42	63C	4.11 (front-channel)	68	2491	"an IdP that uses HTTP redirects for front channel presentation of assertions that contain PII SHALL encrypt the assertion as discussed in Sec 3.12.3." Greater clarification is required around the intent as this seems excessive for FAL1 since front-channel SAML is very common in industry. See also comment about federated identifiers in Sec 3.3	For SAML at FAL1, RPs and IDPs are deemed to meet this requirement if they only expose HTTPS-protected endpoints.
43	63C	5.3 Trust agreements	71	2557	Under "The following terms SHALL be disclosed to the subscriber during the runtime decision:" is does not say that additional information implicitly shared (like the credential issuer or any reused salted-hash values) should be disclosed to the user. Several of us are worried that users may be deceived by the consent screens for salted-hash selective disclosure schemes into believing that the interaction is more private than it really is (especially in the context of proof-of-adulthood). Perhaps this is covered elsewhere in the document, but it seems like a potential omission here to mention disclosure of the subscriber attributes but not the implicit attributes. Also note that there are multiple sections named "trust agreements" (eg, 3.4 and 4.3) which could cause problems for a #trust-agreements anchor.	
44	63C	5.5	73	2604	typo: "attribute bundle >>>singing<<< public key"	
45	63C	10	96	3258	Request informative examples of trust agreements for FAL1, FAL2, and FAL3 including xAL requirements.	
46	63C	10	96	3258	Request informative examples of RP specified xAL requirements for OIDC and SAML.	
47	63C	10.7	101	3423	Expand informative examples of bound authenticator use to include IdP specified bound authenticator use via enterprise attestation.	
48	63C	10.8 (fal3 with referred token binding)	102	3448	"the RP can use [the token binding headers] to associate the contents of the assertion with the subscriber's bound authenticator." Greater clarification is required as to whether is this considered a "bound authenticator". It could be read that this term was only used in this revision for RP-managed things and that's why "or HoK" is used pervasively (e.g. Sec. 3.16). Perhaps this should say "... with the subscriber's authenticator" (omitting "bound")?	