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:	*1Kosmos*
Name of Submitter/POC:	*Christine Owen*
Email Address of Submitter/POC:	[redacted]

	Publication				**Comment**	
Comment #	(Base, 63A, 63B, 63C)	**Section**	**Page #**	**Line #**	(Include rationale for comment)	**Suggested Change**
	63A	title page	title page		typographical error	Ryan is listed as an author twice on the first and second pages
	63A		7	540	typographical error	Verification is misspelled
	63A	1.3.2	6	534	typographical error	remove space after Organizations
	63A	2.1.3	8	611	The requirement to provide a remote unattended path is overly restrictive. Programs that routinely interface with applicants in-person may not have need for such a path.	Change to "CSPs that offer services at IAL1 or IAL2 **SHOULD** provide a Remote Unattended identity proofing process and SHALL offer at least one attended identity proofing process option."
	63A	2.4.1.1	11	674	typographical error	The list contains two number 1 items
	63A	2.5.1	14	798	Since the top algorithms now perform facial comparisons more accurately and equitably then humans, an option should be provided to allow an algorithmic comparison during an in person proofing session.	Add this sentence: "Optionally, the photo on the identity evidence can be scanned and algorithmically compared to a photo of the applicant taken by the proofing agent."
	63A	3.1.2.1	19	939	Organizations should continuously monitor performance not only for disparate performance but also for effectiveness.	Add"...**, and to ensure continued design and operational effectiveness in mitigating fraud risks.**"
	63A	3.1.3.2	22	1039	Suggest highlighting other privacy preserving techniques, such as encryption or hashing, which can be of value.	Include - e.g., transmitting and accepting hashed, encrypted, or partial values rather than values transmitted in clear text or full attribute values
	63A	3	23	959	Legitimate users may face barriers to identity proofing. For example, where proof of address is required as a fraud prevention measure, that will create a barrier to a homeless individual with a pre-paid phone. Bias is a type of barrier that involves prejudice. It is therefore covered under the concept of a barrier, so should only be explicitly called out if there is evidence of prejudice in government identity proofing processes. If such prejudice existed that would be a serious issue, so if there is evidence specific examples should be provided.	Remove the phrase "including biases" which implies that prejudice is part of the identity proofing processes used by agencies.
	63A	3	23	964-965	It is understandable that usability issues can be a barrier to some individuals presenting an authenticator successfully. For example, authenticator apps may change numbers too quickly for some users to successfully enter them as a second factor. However, what 'bias' or prejudice could be at play during a failed authentication attempt? Not mentioned is that the *availability* of an authenticator may be an issue, such as when someone loses their phone or fido token.	Replace 'including biases' with 'including availability'
	63A	2.1.2	7	555-558	If a proofing agent is also making a determination of the linkage of the claimed identity to the applicant (Sec 2.5.1), then the agent should also be trained to detect deception and signs of social engineering; similarly, (Sec. 3.1.2.1) requires proofing agents to detect indicators of fraud and flag suspected events, which means their fraud training should be more compatible with trusted referees and (Sec. 4.1.7) requires them to be trained to detect deception and fraud.	Either increase the level of training of the proofing agent or change Sec. 2.5.1, pg. 14, line 799-808, from proofing agent to trusted referee.
	63A	2.4.1.3	12	720-721	Enrollment should be onsite attended only, not remote attended, for superior evidence. Given the rise of deepfake technology and how easy it is for bad actors to impersonate their victims and socially engineer live service desk workers, superior evidence should only be given to those people who must go onsite to obtain the evidence (e.g., DMV, PIV office, etc.)	"The issuing source had the subject participate in an onsite attended enrollment..."
	63A	2.4.1.3	12	729-730	The evidence should be difficult to reproduce regardless of whether it is physical or not. While there is also a cryptographic requirement for Superior evidence, reiterating the need that reproduction should be difficult regardless of credential type is not too repetitive	Remove "If the evidence is physical"
	63A	2.5.1	14	784-786	The ability to return a confirmation code is not a strong enough piece of evidence to prove that the applicant is more than likely than not the claimed identity at IAL2 or IAL3; it could be acceptable for IAL1, as that level of assurance is considerably weaker. Because the section requires "one or more" method, this evidence alone should not be considered enough to verify the linkage of the applicant to the identity	Remove these lines
	63A	2.5.1	14	799-808	A visual facial image comparison should only be conducted at the time the person presents the evidence (so for attended sessions only). Proofing agents need to be trained enough in social engineering to be able to detect potential fraud, ask the applicant questions, and observe the applicants facial movements to detect possible fraudulent activities. This would not be available in unattended sessions, and without the ability to biometrically compare the photographic evidence, a deep fake could easily pass a proofing agent who reviews the evidence after capture.	"Remote (Attended) visual facial image comparison. The (*trusted referee (see comment above on necessity of social engineering training fo*
	63A	3.2.1.1	18	936-938	The amount of training required in this section is more than required in 2.1.2	Either increase the level of training of the proofing agent in section 2.1.2 or remove proofing agent from this section.
	63A	3.1.4	23	1085-18	While the results of an equity assessment should be made available to any and all RPs that are customers / potential customers of a CSP, having an assessment and associated mitigations publicly available could invite bad actors to find the ability to commit fraud against the CSPs.	Suggest limiting the availability to RPs
	63A	A.2 Table 5	78	2728	When not validated, the passport is on par with a driver's license (Sec. 4.2.4). Keep the passport in SUPERIOR only when NFC chip is read.	Add Passport to "STRONG", when only visual inspection is used.
	63A	3.1.12	30	1325-13	The shall statement begins with "when setting allows for it", suggest removing that and changing the entire statement to should. Reads stronger that way.	change shall to should; remove "when the setting allows for it ...events)"
	63A	3.1.12	31	1336-13	"assessed and certified" means a training and a certificate, of which there is none for proofing agents/trusted referees today. This requirement puts undue burden on the CSPs and RPs, who will have to spend additional money for certification bodies to certify employees.	Remove "and be assessed and certified with at least annual evaluations"
	63A	3.1.13.1	31	1358	This language implies that entire categories of individuals will find identity proofing challenging.	Change "such individuals and demographic groups includes:" to "such individuals may include"

	63A	3.1.13.1	31	1360	Individuals with no access to online services to not need to be identity proofed since they will not be able to then use online services.	Change "individuals with **little or no** access to online services" to "individuals with **limited** access to online services"
	63A	3.1.13.1	32	1368	"train and certify" - means a training and a certificate, of which there is none for trusted referees today. This requirement puts undue burden on the CSPs and RPs, who will have to spend additional money for certification bodies to certify employees.	Remove "and certify"
	63A	3.1.13.1	32	1379	"annual recertification of the trusted referee's capabilities" - means a training and a certificate, of which there is none for trusted referees today. This requirement puts undue burden on the CSPs and RPs, who will have to spend additional money for certification bodies to certify employees.	Remove "annual recertification of the trusted referee's capabilities"
	63A	4.1.6	37	1561-15	A visual facial image comparison should only be conducted at the time the person presents the evidence (so for attended sessions only). Proofing agents need to be trained enough in social engineering to be able to detect potential fraud, ask the applicant questions, and observe the applicants facial movements to detect possible fraudulent activities. This would not be available in unattended sessions, and without the ability to biometrically compare the photographic evidence, a deep fake could easily pass a proofing agent who reviews the evidence after capture.	Remove "or an asynchronous process (i.e., visual comparison made by a proofing agent at a different time)"
	63A	4.1.7	38	1576-15	While we agree that anyone who is overseeing the video sessions should have extensive fraud training, but we are now unsure what the actual difference between a proofing agent and a trusted referee is, considering they need the same amount of fraud training	Review the other comments on proofing agents/trusted referees; think about removing proofing agents from the guidelines altogether
	63A	4.1.7	38	1580	typographical error	"records **the** session"
	63A	4.1.8	39	1600-16	While we agree that anyone who is overseeing the onsite sessions should have extensive fraud training, but we are now unsure what the actual difference between a proofing agent and a trusted referee is, considering they need the same amount of fraud training	Review the other comments on proofing agents/trusted referees; think about removing proofing agents from the guidelines altogether
	63A	4.1.10	40	1637	typographical error	"Return of **a** confirmation"
	63A	4.2	40	1648	Biometrics improve security. Requiring the capture of a facial image during identity proofing is a powerful deterrent for community and family-level bad actors. If biometrics are optional at IAL2 then an agency that requires biometrics for security reasons will not be able to accept an IAL2 credential that was established by an agency with a higher risk tolerance. This breaks the consistency and trust that enables federation.	Either require biometrics for IAL2 or split IAL2 so that the use of biometrics during identity proofing is clearly captured and transmitted to all RPs so they can make a decision on its use or absence, and make sure the pathways are clearly marked such as IAL2-B (biometrics) and IAL2-O (other). These two pathways are *not* equivalent from a security and fraud-deterrence perspective. The non-biometric pathway is highly vulnerable to attack by family members, caregivers, and acquaintances, which can lead to devastating financial and life consequences for disabled beneficiaries and the elderly who rely on their benefits. Capturing the facial image of the individual who is applying for benefits is a strong deterrent to impersonation, particularly for individuals who are personally acquainted with a victim. There is no equivalent deterrent in the non-biometric pathway. Individuals with common names are also highly vulnerable to attacks when address verification is used for proofing without sufficient additional controls. Records may show that a number, email address, or home address is strongly associated with a James Smith, for example. There are 38,313 James Smith's in the United States. This is a common attack that is happening at scale today. https://www.statista.com/statistics/279713/frequent-combinations-of-first-and-last-name-in-the-us/
	63A	4.2	40	1649-16	These options DO have different security and assurance outcomes, which effectively waters down the security of IAL2 to the least security option. Also, digital evidence will either be part of the biometric or non-biometric pathway. Why create a third pathway just for digital evidence? Wouldn't it really need to be four pathways in that case? Digital non-biometric and digital biometric?	Acknowledge that these pathways are not equivalent from a security and assurance perspective and rank them, perhaps: IAL2-High (biometrics), IAL2-Moderate (non-biometrics), IAL-low(?). Otherwise, trust and interoperability and therefore the ability to leverage federated credentials will break - an agency that requires IAL2 with biometrics will not be able to accept an IAL2 credential that originated with another agency.
	63A	4.2.6.1	42	1716-17	FAIR evidence doesn't require facial images, and therefore should not require biometric or non-biometric methods to verify the evidence. However, a verification method to match the facial image on a STRONG or SUPERIOR piece of evidence is required.	Remove the requirement for FAIR evidence
	63A	4.2.6.1	42	1720-1724 & 1730-1734	The likelihood of a successful impersonation varies significantly between these options, but they are listed as though they provide comparable security, which can be misleading.	Add a sentence that indicates that these are not equivalent in terms of fraud prevention.
	63A	4.2.6.1	42	1727-17	A confirmation code is not only an onerous process for the applicant, but it is also onerous on the CSP. This does not provide a strong ability to link the applicant to the facial image on the evidence. While this could prevent unknown bad actors from committing fraud, it does not prevent family members or close friends from committing fraud.	Remove the requirement for STRONG/SUPERIOR evidence
	63A	4.2.6.1	42	1733-17	A visual facial image comparison should only be conducted at the time the person presents the evidence (so for attended sessions only). Proofing agents need to be trained enough in social engineering to be able to detect potential fraud, ask the applicant questions, and observe the applicants facial movements to detect possible fraudulent activities. This would not be available in unattended sessions, and without the ability to biometrically compare the photographic evidence, a deep fake could easily pass a proofing agent who reviews the evidence after capture.	Remove "or an asynchronous process (i.e., visual comparison made by a proofing agent at a different time)"
	63A	4.2.6.2	43	1745	What is meant by a 'validated account'? Does it need to be an account owned by someone with the same name? That is not sufficiently granular for someone with a common name. Does it need to be an account associated with someone's name and SSN? Does length of ownership matter? Note: Association with an address is insufficient. Bad actors can provide a target's legitimate home address and then opt to only receive communications via email.	Provide guidance, or a reference to guidance.
	63A	4.2.6.2	43	1757-17	AAL3/FAL2 is too onerous of an authentication protocol. Looking ahead to digital passports which will include cryptographic elements within a TPM on the phone that will be able to be validated by the CSP, an applicant will more likely be utilizing an AAL2 authenticator over an AAL3 authenticator, as the protocol to unlock the digital passport will likely be WebAuthn. As such, AAL2 or higher should be accepted	change AAL3 to AAL2
	63A	4.3.7	47	1866-18	While we agree that anyone who is overseeing the onsite sessions should have extensive fraud training, but we are now unsure what the actual difference between a proofing agent and a trusted referee is, considering they need the same amount of fraud training	Review the other comments on proofing agents/trusted referees; think about removing proofing agents from the guidelines altogether

	63A	4.4	49	Table 4.	Why does IAL1 attended require FAIR with an image? It is the only FAIR evidence that requires an image.	Remove "w/ image"
	63A		54	Table 2	The example given in 'False Claims' may be committed by a legitimate user. Providing evidence to fraudulently claim a privilege that one is not entitled to is out of scope. Change this example to one that applies only to identity proofing.	Change to something like: "An attacker registers an address they control with the attributes of a legitimate user."
	63A	7.4	59	2112	Redress mechanisms can be highly vulnerable to impersonation attacks	Change to "provide effective **and secure** mechanisms for redressing applicant complaints or problems"
	63A	7.4	59	2124	Bad actors are *known* to have used identity proofing processes to verify PII. To prevent this, the 'should not inform' should be changed to SHALL NOT inform.	
	63A	9.3	72	2560-25	This is an overly broad and inaccurate statement of the landscape of identity vetting technology today. While GSA's "A large-scale study of performance and equity of commercial remote identity verification technologies across demographics" did find that all the technologies but one "were equitable across most demographics with [a few] exceptions," and that one technology "was equitable across demographic groups with all groups following within a 95% confidence interval."	change to "Some facial image capture technologies lack the ability to capture certain skin tones or facial features of sufficient quality to pe
	63A	9.3	72	2567	"Sex assigned at birth" is a highly political term that only one party uses; to allow these Guidelines to survive executive party changes, suggest changing to just "sex"	Change "sex assigned at birth" to "sex"
	63A	9.3	72 & 7	2572 & 2	Technological limitations accurately and completely captures any challenges with image capture without anthropomorphizing.	Remove "residual bias"
	63A	9.3	73	2587-25	False non-matches can occur not only in 1:1 matches, but also in 1:N matches. Suggest changing the language slightly, especially given the term "biased facial comparison algorithms" is in and of itself biased against algorithms.	change to "When using poor quality facial image comparison technology, there could be a false non-match result."
	63A	9.3	73	2598	People have limitations when it comes to accurate facial verification, which is very different than implying that people who are unable to accurately perform verifications are prejudiced.	Replace "biases" with "limitations"
	63A	9.3	73	2597-25	In the Guidelines discussion of equity, it assumes that not giving access to a user is the only non-equitable result, whereas the reality of fraud and identity theft creates an extremely unequitable result for those who need public benefits but cannot because their identity was stolen. While false non-matches are an unequitable result of poorly trained algorithms, identity theft due to the rise of deep fakes and social engineering are an even worse result given the amount of hardship a person has to go through to prove that they are, indeed, their own identity.	Stronger discussion around the real-life, unequitable hardships to those identity theft victims who are suffering due to an unproven belief that facial image comparison technology are inherently biased and wholly inadequately untrained, and therefore allowing for a default of a non-biometric approach to IAL2.
	63A	9.3	73	2601	When high quality images are used, best of breed algorithms now perform facial verifications more accurately than trained human agents. So, providing an automated option for individuals who have failed verification by a human would reduce false non-matches.	Replace the first mitigation (which is, in all practical ways, identical to the more concisely worded second mitigation), with "1. Provide the option for applicants to have a photo taken which will be algorithmically compared to the portrait on their strongest piece of evidence."
	63A	2.1.2	7	555-571	More clearly state the difference between a trusted referee and a proofing agent: the trusted referee is empowered by the CSP to make deviations from the required evidenced following pre-established guidelines. Yes, there may be some difference in training, but both need fraud and social engineering detection skills.	Add language about training to end of proofing agent definition. Add language about deviations to definition of trusted referee. Change no
	63A	2.1.3	8	595	Indicate that the session does not have to be wholly attended. For instance, if the applicant struggles with only the selfie, they can do the rest of the process unattended. That fits with the language at the end of this subsection but is not clear in the definitions. Also state that if any part is attended, the entire process should be considered attended.	add second sentence: "Some parts of the process may still be automated. If any portion of the process involves a proofing agent, it is consi
	63A	2.1.3	8	593 and	add "by the applicant"	The location and devices used by the applicant in the proofing process are not controlled by the CSP.
	63A	2.4.1.1	11	674	numbering issue. Two #1s.	
	63A	2.4.2	12	736	It's confusing how we say the document uses valid to not refer to expired but then parenthetically defines valid as unexpired. Can this language be clearer?	
	63A	3.1.1	17	875	Add 15. Fraud management practices to align with 3.1.2.1(1)	The CSP's approach to fraud management, consisted with sec 3.1.2.1
	63A	3.1.2.1	18	888	Be more precise about what capabilities should be included in the practice statements. While the market certainly needs additional transparency, some of these are trade secrets or would advantage fraudsters if they were known	The specific capabilities and details of this program SHALL be documented within their CSP practice statement with sufficient detail to app
	63A	3.1.2.1	18	907	consider including other fraud check techniques: device fingerprinting and correlation, VPNs and proxy servers, predictive analytics	
	63A	3.1.2.1	19	939	continuous monitoring has a specific meaning that perhaps is not intended here (though it is the ideal)	Change continuous to continual
	63A	3.1.2.2	19	954	Is this intended to imply CSPs cannot collect additional attributes to support fraud mitigation? I think that would be limiting. May need to clarify that the CSP can collect additional attributes, even if the RP should only request the core, except for these fraud exceptions in this clause	Add similar language to 3.1.2.1
	63A	3.1.3.1	22	1023	More clearly define what we mean by organization. Is that just RPs or something else? If an end user (subscriber) is an org or is acting on behalf of an org, could they see it? What if the end user is an individual?	Change to RP
	63A	3.1.4	23	1067	typo "are be"	are to be?
	63A	3.1.4	23	1085	Be more specific about results. Raw data? Summary reports? Making all mitigations public may advantage attackers--and may do so in a way that undermines the very groups we hope to benefit from the mitigations	Limit what is publicly available and more comprehensive results available only to RPs, regulators, etc.
	63A	3.1.8	26	1162	24 hours is too long or email. In reality if someone doesn't confirm an email within a few minutes, they've probably forgotten. It only adds risk	Lower 24 hours to 1 hours
	63A	3.1.9	26	1164	make clear that authentication at an appropriate AAL is sufficient as an option as well	Authentication at a corresponding AAL or higher is sufficient to obviate the need for a continuation code
	63A	3.1.9	26	1174	By which methods can the CSP deliver the code? Validated address? Just in session?	
	63A	3.1.11	27	1226	Define "explicit biometric consent." Does this mean a separate consent from that for the service, or that the terms contain an explicit statement.	explicitly state consent for biometric collection before collecting…
	63A	3.1.12	29	1295	What does live capture mean here? I understand it in the context of DL or passport, but it sounds like I can't upload an image of fair evidence; I have to have a physical copy of it to scan in-session	live capture applies to STRONG and above evidence. Alternatively, the entire section only applies to STRONG and above

	63A	3.1.13	31	1349	Trusted referees are only necessary when an RP wants to accept based on its risk determinations. It's not up to a CSP, particularly a 3rd party CSP, to determine if a trusted referee process is sufficient. This could lead to problematic incentives for CSPs. It could also disadvantage small CSPs and make market entry difficult	change SHALL to SHOULD, when requested by the RP
	63A	3.2	35	1490	Subscribers should have to authenticate at the AAL that corresponds to the upgraded identity, not the highest. If you're going from IAL1 to IAL2, you don't need to authenticate at AAL3.	Authenticate at an AAL at or higher than the upgraded IAL
	63A	4.1.10	39	1625	the authenticators can also be bound prior to the proofing event	before, at the time of, or after the proofing event
	63A	4.1.6(5)	37	1561	If you do retain the possibility of an asynchronous process, it needs more explanation and a clearer description of the requirements around it. A still photo is not sufficient. At minimum, a still would need to be captured with liveness. Ditto this comment for other mentions of a synch comparison	
	63A	5.1	50	1927	Are there unintended consequences with requiring "each subscriber as a unique identity"? For instance sole proprietorships may have an individual account and a business account, both with the same identity. This is also not specific to an IAL, so it could apply to pseudonymous applications.	change identity to entity or persona
	63A	5.1	50	1941	The type of proofing should be for each piece of evidence and step in the process as proofing may be multi-modal (got stuck on selfie)	Type of proofing for each type of identity evidence
	63A	5.2	51	1958	not just PII	change PII to personal information
	63A	9.3	74	2627	Reference WCAG. Accessibility challenges should be mitigated consistent with federal requirements. Consider making this normative	
	63A	A.1	78	2728	Consider adding a column for allowable proofing type. For instance, is SSN acceptable if remote so long as it's presented as an image of a card and not just the number?	
	63B	1	1 & 2	403-413	This is an overarching comment for 63B: In its introduction, NIST recommends "applications assessed at AAL1 offer" MFA, which suggests that AAL1 is barely a "basic confidence" assurance level. More strongly, NIST requires "[a]applications assessed at AAL2 must offer a phishing-resistant authentication option." While this is a great mandate, if an application has both phishable and phishing resistant authenticator options, the application continues to be phishable and the phishing-resistant option mandate loses its purpose - to effectively protect against phishing attacks. Additionally, federation assertions don't always assert whether an authenticator was phishing resistant or not, just that it was MFA/2FA, so adding an additional AAL would create a need for vendors to insert greater veracity in their assertions. As such, phishable MFA should become AAL1, AAL2 should be phishing resistant and an AAL0 should be created to allow for single factor authentication scenarios.	Create an AAL0 that is a low assurance and allows for single factor; AAL1 should become phishable MFA, whereas AAL2 should be phishing
	63B	3.1.7.4	28	1192-11	This is an overarching comment for 63B: All cryptographic authenticators are technically syncable, but the type of authenticator NIST is referencing is actually synched authenticators - those that affirmatively decided to synch and remove its device-bound status	Change all "syncable" references to "synced"
	63B	3.2.2	28	1216	100 attempts is very high; suggest reducing the attempts.	Suggest reducing the amount of attempts to below 100.
	63B	8.4	73	2463-24	Generally biometrics will unlock a private/public key pairing, therefore in an "ideal" situation allowing users pick their second factor doesn't make sense.	suggest removing the sentence or noting that most uses of biometrics will be to unlock a cryptographic certificate.
	63B	9	76	2554-2555		remove reference to sexual exploitation, as it is a type of trauma.
	63-Base	1	1	359	change agencies to organizations, as throughout this section you discuss "organizations"	change agencies to organizations
	63-Base	1	2	379	It's unclear what a "culturally appropriate" option would be for individuals, considering every identity needs to be vetted to the appropriate assurance level.	Remove "culturally appropriate" reference.
	63-Base	1	1	362	typographical error	Change "considerations **and** organization "to considerations **an** organization"
	63-Base	2.1	10	646	typographical error	make identity provider lower case
	63-Base	2.1	10	640	typographical error	Remove the hyphen after the em-dash
	63-Base	2.1	10	646	typographical error	should be "relying party" not "relying party"
	63-Base	3			While CSPs absolutely must help their RP clients with risk assessments and the DIRM process, requiring them to also create a DIRM is excessive work. While there are many paths to get to IAL2, alongside many risks and exceptions process, ultimately the onerous is on the RP to understand the CSP's process, ask for enhancements, and accept the risk obtained from the original process alongside the enhancements.	Remove requirements for CSPs to create DIRMs; keep the requirement to aid RPs in creating their DIRMs
	63-Base	3.1	27	1086	it's unclear what a "culturally responsive" communication alternative would be and whether that description actually enhances this section of the guidelines	than values transmitted in clear text or full attribute values
	63-Base	3.1			Scale of impact is a critical missing element from the user group discussion, which focuses exclusively on role. For example, the successful impersonation of a tax preparer who has a handful of small business as clients will not have the same impact as the successful impersonation of a tax preparer who handles the accounts for Walmart and Wells Fargo. Similarly, the impacts on the organization/application itself also needs to be assessed, as well as any other impacted entities. For example, beneficiaries may not be direct users of an application for appointed representatives for those beneficiaries, but should nevertheless be considered since their information would be exposed if a DI error were to occur.	Update impact review
	63-Base	3.1	28	1120	typographical error	"At a minimum, agencies shall document all impacted **entities** when conducting their impact"
	63-Base	3.2.1			Again, it would be difficult for a CSP to determine the harms and the risk acceptance for an RP. Not only does the CSP not have a decent understanding of what the RP is proacting (HIPPA, PIA, basic info, just a file upload system, etc.), but the CSP also doesn't have an understanding of the other products within the RP's technology stack that could mitigate any harms done, and therefore reduce the impact of a breach.	Remove requirements for CSPs to create DIRMs; keep the requirement to aid RPs in creating their DIRMs
	63-Base	3.2.2	30	1205	Section 3.1 defines potential harms to both user groups and impacted entities	Update the sentence to include user groups: "...impacts **on user groups and** entities identified in Sec. 3.1"

	63-Base	3.2.3	33	1290-12	It would be impossible for a CSP to perform risk assessments on user groups, given a particular user can be a part of different types of user groups in different RPs. For example, a police officer might have administrative access within their own organization, read only rights to law enforcement sensitive documents in other, federated organizations, the need for benefits for IRS/SSA, a Global Entry participant for State and TSA, as well as a FOIA requester at GSA. This user's changing groups would create different levels of risk, all of which each individual RP needs to accept or decline.	Remove requirements for CSPs to create DIRMs; keep the requirement to aid RPs in creating their DIRMs
	63-Base	3.2.4	34	1321	How does an organization come to an "overall impact score" for each user group? What goes into the score? How does the score translate to assurance levels?	Clarify what an impact score is and how to calculate it, or remove it from Base.
	63-Base	3.4.4	44		"initially assessed xALs" are required in a DIAS, which is required to be created by a CSP. This requirement doesn't make sense for CSPs, as CSPs should not be assessing what the xAL should be for a given RP. While the CSP can help the RP create the xAL assessment, CSPs must create workflows to meet their customer's needs, which means their assessed xALs, tailored xALs, and compensating controls will be determined by their customers - the RPs - not by the CSPs themselves.	Remove requirements for CSPs to create aDIAS; keep the requirement to aid RPs in creating their DIAS
	63-Base	3.8	50	1822	The term "biased outputs" when describing AI/ML results is biased against the technology.	change to "disparate outcomes or outputs," as the adjective "disparate" encompasses the issue - which is that algorithms that are not well
	63-Base	3.7	50	1809-18	The RP is the customer, and the CSP is a shared/managed service that is providing the customer with services, including data about potential fraudulent activities, extensive activity logs, etc., to provide the RP with a broader, more holistic view of its cybersecurity landscape and threats.	change shall to should
	63-Base	2.1	10	646	you'll want an apostrophe here	relying' party
	63-Base	3	22	932	the system might already be implemented but it's risks are still part of DIRM. This will reflect DIRM as a continuous process. Consider flowthrough changes.	Replace "might be addressed" with "addressed." Replace "to be implemented" with "itself"
	63-Base	3.2	28	1127	The first dimension is defined earlier as "risks to the online service that might be addressed by an identity system." The parenthetical suggests something else, though the rest of the sentence is consistent	remove the parenthetical, replace "and seeks to identify" with "by identifying"
	63-Base	general			in all volumes, there is too much passive voice. Leaving out the subject of a sentence often does not impact the interpretation, but can lead to confusion, particularly when assigning responsibility. This is especially important near normative statements. For instance, some sections and subsections do not clearly state the actor. 63-4 Section 3 is a good example where the document would be well served to state "RPs shall" or the like at the start of each subsection to eliminate potential confusion	
	63-Base	3.2.1	29	1156	prior sections have mentioned RPs as the actor here. This one says "organizations." I don't think there is a different actor here; the RP is best suited for this activity. Ditto line 1218, 1312, 1320, 1323, and others. If there is a value in saying organizations and not RPs, make it clear at the beginning of section 3 that RPs are responsible for this, at least in most cases. There are other areas in which agency, organization, and RP seem interchangeable, which adds confusion	Change "organizations" to "RPs"
	63-Base	3.3.2.1	35	1358	Given the significance of the change, call out that no identity proofing is no longer IAL1. I don't think you have to state that it's a "change" but mentioned it explicitly to remove confusion. Ideally there would be an IAL0 so it can be recorded properly when agencies do assessments	Add a bullet: No identity proofing: Knowledge of the user's real-life identity is not needed and no identity proofing activities are conducted
	63-Base	3.3.3.1	38	1434	holds suggests possession and not existence.	change "holds" to "is"
	63-Base	3.4	40	1509	if they are underserved now, there should be a focus on them	remove historically
	63-Base	3.4	40	1514	move footnote marker to end of sentence	
	63-Base	3.2.2	30-31	1206-12	Some types of impacts are not easily captured by the listed impact categories, such as a loss of Medicare coverage. Such a loss of coverage can lead to financial loss, endanger someone's health, and lead to extreme stress at a time when someone may be suffering from serious health conditions. Delays in receiving disability benefits can also lead to a cascade of negative consequences that can be incredibly detrimental and hard to recover from.	Consider adding an additional Impact Category such as "Quality of Life Degradation"
	63-Base	3.2.2	32	1259 &	While it is possible to anticipate the types of medical treatments that would be required for minor physical injuries, it is not possible to anticipate whether a minor injury may lead to the need for mental health treatment.	Remove "including mental health treatment".
	63-Base	3.3	34	1332	Organizations that use CSPs cannot select individual controls.	Update the guidance to realistically reflect the limited options RPs will have when they are using an external CSP such as Login.gov.
	63-Base		34	1347	AAL is defined later in the document as "the level of assurance that the claimant is the same individual to whom the credential or authenticator was issued." This is a simpler and clearer definition.	Recommend replacing: "The robustness of the authentication process itself, and the binding between an authenticator and a specific individual's identifier." with "the level of assurance that the claimant is the same individual to whom the credential or authenticator was issued."
	63-Base	3.2.2	36	Table 2:	"Support" for MFA is not the same as a requirement for MFA.	Change "Support multifactor authentication" to "Requires multifactor authentication" or "Enforces…"
	63-Base	3.2.2	36	Table 2:	"Providing" (the option of) phishing resistance is not the same as a requirement for phishing resistance.	Change "Provide phishing resistance" to "Requires phishing resistance"

	63-Base	3.4	40	1507-15	Marginalized and historically underserved populations are those which are often most severely impacted by DI errors that can result in consequences such as stolen benefits or identity theft. Yet the tailoring instructions direct agencies to focus exclusively on frustrations with the DI controls themselves.	Add: "These considerations should be weighed against the disproportionate impacts that may be experienced by these same populations in the event of a DI error."
	63-Base	3.4	41	1544	Keeping individuals safe is also critical.	Change to "...while supporting **security**, equity, privacy, and usability for individuals."
	63-Base	3.4.1	42	1574	Marginalized and historically underserved populations are those which are often most severely impacted by DI errors that can result in consequences such as stolen benefits or identity theft.	Extend the final sentence: "The intent of this assessment is to mitigate potential impacts on marginalized and historically underserved groups and limit disproportionate impacts from the requirements of the identity management functions **while providing adequate protections against impacts of the fraud and impersonation that can occur when those functions fail.**"
	63-Base	3.5.3	48	1726	There is a real risk that focusing exclusively on equity & accessibility will result in greater harms being done to vulnerable individuals whose money and data are then stolen due to a reduction in effective controls.	Change to: "A primary purpose of continuous improvement is to improve Equity and Accessibility outcomes for different user populations **in a way that does not result in a substantial increase in fraud or theft of PII or personal or sensitive information."**
	63C	1	1	403-407	While a digital wallet could be used as an IdP, it's not advised to be used as one. A digital wallet holds a user's attributes, uses strong authentication (generally phishing resistant) to open the wallet, and requires user consent to submit attributes to a requesting RP. As a wallet is focused on holding attributes and not authenticators, most wallet providers also provide an IdP that allow verified users (who hold a digital wallet/verified credential) to create strong MFA authenticators to use for authentication with a variety of RPs. Additionally, not all subscriber-controlled wallets are on-device. While the private key to unlock the identity attribute is stored with the subscriber, many wallet vendors have taken a privacy-focused, cloud-storage approach for wallets, which reduces the loss of verified attributes when a device is lost/stolen/broken.	Remove subscriber-controlled device as an IdP.
	63C	2	3	476-479	Trust Agreements can be bilateral or through a trust framework.	Recommend introducing the idea of a trust framework in this section.
	63C	2.3	6	542	typographical error	"At FAL2, the assertion SHALL **be** audience restricted to a single RP."
	63C	2.5	7	585-586	While informing the RP of the IAL of the subscriber and FAL of the transaction is good practice, there might be scenarios where the RP doesn't want the IAL of the subscriber. This should be a SHOULD statement, as the RP should be able to decide what information it needs from the IdP to make a holistic decision of whether the subscriber should gain access. Similarly, both the IdP and RP SHOULD have a record of what the FAL of the transaction is, so requiring an IdP to submit that information seems excessive.	Change IAL and FAL requirements to SHOULD statements
	63C	2.5	7	600	The lowest IAL in the guidelines is IAL0 (no proofing), not IAL1.	Change IAL1 to IAL0
	63C	2.5	7	609-612	This statement is only true for traditional general purpose IdPs, as a privacy-preserving digital wallet/verifiable credential does not have "direct access" to all the details of the subscriber account, only the details the subscriber consented to.	Remove these lines
	63C	3.1.2	11	688-689	Because there are two distinct places where identity evidence is stored on behalf of the subscriber of digital identity wallets (device or the cloud), a wallet should not be protected just by an activation factor (PIN/password) if a subscriber-controlled wallet doesn't require a private key (which is the case for mDLs). Similarly, because wallets are used to store attributes, not authenticators, this section should exclusively be discussing IdPs.	Remove either the presentation of an activation factor and/or of subscriber-controlled wallets in Section 3
	63C	3.2.2	12	720	typographical error	"identity attributes" should be plural, not singular
	63C	3.4	17	896-897	How does a trust agreement "establish usability and equity requirements" in a federation transaction? To the user, they are only operating with the RP and, in some cases, will be able to understand that they are also using another vendor's program. Because a federated transaction is basically based on protocols, how can we ensure that those universally standardized protocols hit "usability and equity" requirements?	Remove or change to a SHOULD statement
	63C	3.4	17	897-899	While the trust agreement SHOULD include details of the proofing process, adding compensating controls and exception processes to the trust agreement could be a security issue, as the trust agreement is required to be public later in this section. To drill down more, while exception processes and compensating controls must be in place to ensure equity and fairness for applications, making those processes and controls public will create a massive attack vector to the CSP with the most relaxed processes, inviting fraudsters to capitalize on those CSPs.	Suggest changing SHALL to SHOULD
	63C	3.4	17	919-920	FALs also need to be included	include FAL in the list.
	63C	3.4	18	930-938	While transparency is very important to give to subscribers, it should not be given when it can compromise security. If, for example, compensating controls are required to be in the trust agreement, that is not something that an ordinary subscriber should be able to see without a clear reason, as it will compromise security of the CSP, IdP, and RP.	Change to "As such, a high-level summary of the terms of the trust agreement..."
	63C	3.4.2	20	965-966	A subscriber-controlled wallet should not also be an IdP, as it holds attributes, not authenticators.	Remove subscriber-controlled wallet reference
	63C	3.6	23	1094	typographical error	add "the" before IdP and RP
	63C	3.6	23	1099	typographical error	add an apostrophe in "subscriber's attributes"
	63C	3.7.1	25	1147	A subscriber should also be notified when the RP terminates their account.	Include a SHALL/SHOULD line for subscriber notification of termination of accounts
	63C	3.9.1	28	1271	Complying with the law or legal process is a requirement of an IdP, whether or not it is in the trust agreement	
	63C	3.9.1	28	1283	Complying with the law or legal process is a requirement of an IdP, whether or not it is in the trust agreement	
	63C	3.11	31	1381-1383	This level of detail was not required under 63A. It is an extremely onerous requirement, and would ultimately lie on the CSP to detail. Those information points would then be used by certain RPs to determine whether the identity vetting met the RP's standards or not, and then might require a subscriber to re-proof certain attributes.	Either make this requirement a part of 63A or remove it from 63C
	63C	3.11.3.1	34	1490-14	How would the RP know definitively that the attributes are allowed to be provided? The shall statement is onerous on the RP	Change from SHALL to SHOULD

	63C	3.11.1	32	1406	The private key is used to sign the assertion, not the public key. The public key is then used to verify the signature. (note - this is an issue throughout C - it allows public keys and signing keys to sign assertions, which is not a normal course of business; see, e.g.. ISO 18013-5 appendix on VICAL key)	Change 'the public key used to sign the assertion" to "the public key required to verify the signed assertion"
	63C	3.14	37	1568-15	As AAL2 allows for both phishing resistant and phishable authenticators, how would an RP know that the assertion has a "phishing resistant" authenticator when technology vendors do not transmit that level of veracity on authentication assertions?	Create a new AAL - AAL0, which would be for single factor authenticators; AAL1 would be for MFA phishable, AAL2 would be MFA phishing
	63C	3.15	38	1594	As AAL2 allows for both phishing resistant and phishable authenticators, how would an RP know that the assertion has a "phishing resistant" authenticator when technology vendors do not transmit that level of veracity on authentication assertions? Additionally, this is at least the second - if not more - time in the guidelines that specifically mandates phishing resistant authenticators. **If there is a higher level of assurance that comes with them (as we know there is), then phishing resistant MFA should become its own assurance level!**	Create a new AAL - AAL0, which would be for single factor authenticators; AAL1 would be for MFA phishable, AAL2 would be MFA phishing
	63C	4		1709	The majority of the digital identity wallets in the US and that will be deployed in the EU will be considered "general-purpose IdPs" based on these draft guidelines. These guidelines do not take into account the concept of verifiable credentials, which are not always a "function of the IdP." W3C verifiable credentials, and similarly most digital wallets deployed to networked systems, bake privacy into their wallets by encrypting a subscriber's attributes and giving only the subscriber the private key. As such, these wallets - even though they are deployed to network systems - are subscriber controlled because only the subscriber can decrypt and share their attributes, which seems to be what subscriber-controlled wallets are focused on.	Remove lines 1709-1710, and include the concept that having a private key in control of the subscriber satisfies subscriber-controlled walle
	63C	4.3.1	47	1807-18	While transparency is very important to give to subscribers, it should not be given when it can compromise security. If, for example, compensating controls are required to be in the trust agreement, that is not something that an ordinary subscriber should be able to see without a clear reason, as it will compromise security of the CSP, IdP, and RP.	Change from SHALL to MAY statement to ensure that Federal security is not compromised by this SHALL statement
	63C	4.3.2	47-48		Transparency is very important to give to subscribers, and many of the bullets in this section should be required even if not requested, **especially if** NIST doesn't change the terminology of subscriber-controlled wallets. However, because this section should be mandatorily given prior to obtaining consent to transmit a subscriber's attributes, the identity API disclosure requirement is unnecessary, as the average subscriber (even me!) wouldn't understand what the ramifications of that is. Additionally, subscribers should not know what other subscriber populations are being asserted, as that can compromise security. Similarly, but not as importantly, only a handful of people understands or comprehends what an xAL is and why it should matter.	Remove "upon request"; remove the bullet that begins with "what if any identity APIs are made available"; remove the bullet that begins
	63C	4.3.2	48	1863	Maybe 300 people (and I feel like I'm being generous here) fully understand the ramifications of xALs. Mandatorily providing this to the whole population will unnecessarily confuse them.	Remove line
	63C	4.6.1.1	51		While this is a MAY statement, identity wallets that are deployed to network systems but are subscriber-controlled as only the subscriber can unlock their accounts cannot pre-populate attributes without a runtime decision from the subscriber. This behavior, however, makes sense for a centralized IdP that is passing identity attributes to an RP that it is enrolling for SSO within an enterprise.	Digital identity wallets that are deployed to network systems should not be considered a "general IdP".
	63C	4.6.4	56	2102-21	"The IdP has a direct view of the subscriber account's attributes" - again, while this is true for a centralized IdP scenario, it is not true for an identity wallet that is deployed to a network system where the private key is kept with the subscriber. Only the subscriber has access to and can consent to the transmission of their attributes.	Digital identity wallets that are deployed to network systems should not be considered a "general IdP".
	63C	4.9	63	2348	While informing the RP of the IAL of the subscriber and FAL of the transaction is good practice, there might be scenarios where the RP doesn't want the IAL of the subscriber. This should be a SHOULD statement, as the RP should be able to decide what information it needs from the IdP to make a holistic decision of whether the subscriber should gain access. Similarly, both the IdP and RP SHOULD have a record of what the FAL of the transaction is, so requiring an IdP to submit that information seems excessive.	Move IAL and FAL requirements to MAY statements starting in line 2359
	63C	4.9	64	2364-23	There are instances in these guidelines where phishing resistant authenticators are required (this is the third time in C), yet they do not have their own assurance level. As such, having their authentication event be a "may" and not a "shall" doesn't make sense.	Either move the "authentication event" into the SHALL statement of the assertion on line 2348 or do the more sensible change to the document and create a new AAL - AAL0, which would be for single factor authenticators; AAL1 would be for MFA phishable, AAL2 would be MFA phishing resistant, and AAL3 would stay the same. :)
	63C	4.8	62	2297 & 2306	It is critical that RPs and IdPs be informed when either suspects that an account has been compromised, especially when RPs are involved that hold highly sensitive data or allow access to funds. If other means of notification, such as email, are allowed, then changing this to a SHALL should not be a problem.	2297 - Remove this item from the 'SHOULD' list and change to: "The IdP **SHALL** send a signal **or other notification** regarding any subscriber account suspected of being compromised." 2306 - "The RP **SHALL** send a signal **or other notification** regarding any subscriber account suspected of being compromised."

	63C				While the EU's digital wallet requirements have a heavy focus on device-bound wallet, they discuss the wallet being accessible on "smartphones or computers" as well as the wallet being available as mobile applications or cloud services (E.g., “European Digital Identity Wallets should ensure the highest level of data protection and security for the purposes of electronic identification and authentication to facilitate access to public and private services, irrespective of whether such data is stored locally or on cloud-based solutions”). The EU also aims to provide a "user-centric approach" to its wallet, which necessitates accessibility across different devices to meet various user needs. This user-centric approach aligns to the Guideline's discussion of usability and equity - as not everyone has access to their own mobile device, but they do have access to a publicly accessible computer in a public library, vendors that provide wallets must create more flexible solutions, which include both mobile-based and networked solutions. Continuing, a networked wallet reduces the need for subscribers to re-enroll and re-verify their attributes when they have their device lost or stolen. Of course, the EU's wallet must meet the privacy requirements found in GDPR, which can be obtained through either device-bound or cloud-based wallets. Because "general IdPs" are traditionally an organization's centralized IdP from which users SSO into a variety of RPs, subscribers expect general IdPs to share their attributes without just in time consent, as subscribers desire as much birthright access to RPs as possible based on security risks. In the situation of digital identity wallets that are stored in the cloud (which are most digital identity wallets, and are being developed under eIDAS standards), they should be considered subscriber-controlled if the wallet's private key is kept with the subscriber.	Align 63C to eIDAS 2.
	63C	5	69	2496	A subscriber-controlled wallet can be in the cloud if the subscriber is in control over the private key.	change to "When the subscriber has control over the private key that protects their attributes or runs on a device controlled by the subscriber, whether as a digital wallet ..."
	63C	5.1	69	2507	An activation factor should always be required before any operation that is using a wallet. While this does create more friction to the subscriber, wallets contain such sensitive information that it's important that the subscriber understands their actions before they do them (e.g., reissue an attribute bundle)	change SHOULD to SHALL
	63C	5.2	69	2519	While we agree that wallets *could* be used as an IdP for limited purposes, we don't believe that it is a best practice to set a wallet up to always act as an IdP, especially when ISO-based wallets don't contain private keys on the mobile device. Continuing, the wallet should not be providing attribute bundles to every RP at every transaction. Wallets should only be transmitting attribute bundles at time of enrollment and when an RP requests an additional attribute to the subscriber account. Which means, the wallet is acting more like a CSP than an IdP. At all other times, a subscriber should be using an authenticator (preferably a phishing resistant one in its own AAL2 category) to help minimize PII leaving the wallet and to create a more privacy-preserving authentication experience. It's very important for the Guidelines to make this distinction, as the vast majority of the time wallets should not be performing authentication into RPs, rather traditional IdPs should be providing this service.	
	63C	5.3	72	2561	Maybe 300 people (and I feel like I'm being generous here) fully understand the ramifications of xALs. Mandatorily providing this to the whole population will unnecessarily confuse them.	Remove line
	63C	5.4	72	2578	This section discusses creation of wallets for mDLs under ISO standards, but doesn't discuss verifiable credentials under W3C standards. Why?	include W3c verifiable credentials, especially since eIDAS2 has aligned to those standards.
	63C	5.6	74	2631	Key to privacy is the ability to selectively disclose a subset of attributes	change from SHOULD to shall
	63C		73	2609	It is more accurate and understandable to say that the RP 'obtains' the identifier and key rather than it 'learns' them.	Change 'learns' to 'obtains'
	63C		73	2612	Public keys can't 'present' attributes, but they can verify them.	to present' should be 'to verify'
	63C	5.8	76	2692	Key to privacy is the ability to selectively disclose a subset of attributes	change from SHOULD to shall
	63C		75	2674-2675	The assertion can NOT include the same key that was used to sign the assertion. Note: This document has repeatedly confused which key is involved in signing vs validating an assertion. And, unfortunately, the subscriber-controlled wallet doesn't seem to align with any current wallet standards (ISO, W3C) or regulations (eIDAS2). When it comes to digital identity wallets, the US is unfortunately not the leader (while we do have technology vendors who are cutting edge at wallets); instead, APAC and the EU are becoming fast adopters of the technology, which is why our standards should align to them.	Change: "This MAY be the same key that the subscriber-controlled wallet uses to sign the assertion." to **"This MAY be the public ds key that corresponds to the private key used by the subscriber-controlled wallet to sign the assertion."** **Align C to APAC, ISO, and W3C standards, as well as eIDAS2**
	63C		73	2626	Perhaps the writer is confusing how asymmetric cryptography works for signatures vs encryption? Digital Signatures: The *private key* is used to sign a message. The corresponding public key is then used to verify that the message was signed by the expected private key. Encryption: The message is encrypted using the recipient's *public key.* The recipient then uses their private key to decrypt it.	Change 'signed by the CSP's public key' to 'signed by the CSP's private key'.
	63C		73	2615	How does the 'RP introduce its properties'?	Provide an explanation.
	63C		75	2679	Language that again implies that the public key was used for signing…	Change "for the key" to "that corresponds to the key"
	63C		76	2703	PII does not include the entire universe of private and potentially sensitive data.	Change to (addition in bold): "contains PII **or other private or potentially sensitive data**"
	63C		76	2704	Message level encryption should be required whenever PII or other sensitive data is passed through a third party.	Change SHOULD to SHALL

	63C		76	2727-2728	Line 2664 in Section 5.8 states that the assertion from a subscriber-controlled wallet SHALL contain a cryptographic nonce only if it is provided by the RP. Line 2701 in Section 5.9 also indicates that it is optional for the RP to provide a nonce. Line 2727 then implies that the RP is required to provide a nonce.	Either require that the RP provide a nonce (recommended) and update lines 2664 & 2701, or change lines 2727-2728 to indicate that the requirement only applies if the RP had provided a nonce in its request.
	63C		76	2731	It is inevitable that some bad actors will be able to obtain signed attribute bundles from CSPs. It is also inevitable that bad actors will succeed in stealing signed attribute bundles from insufficiently protected wallets. (It is only the eventual scale of this fraud that is currently unknown). It is therefore *critical* that RPs are able to determine whether a particular attribute bundle has been reported as having been fraudulently obtained so as to prevent its use.	Change MAY to SHALL.
	63C		76	2754	Additional common attacks include: interception of the password and 2nd factor with a keylogger or redirecting users to a realistic but fake IDP where the password & 2nd factor are captured and relayed to the IDP. Also see https://githubcom/pushsecurity/saas-attacks	Add information on additional attacks and mitigations. To mitigate against credential theft by fake IDPs and keyloggers, users can be prominently shown logs of their previous visits, or at least the most recent visit, along with instructions for when they see a login that they don't recognize.
	63C		87	2971	Broken link - Account Chooser redirects to a list of WGs	Fix link. Perhaps: https://openid.net/wordpress-content/uploads/2011/12/ac-integration-spec.html
	63C		87	2993	Typo - remove 'as' in "commercial as IdPs" & adjective recommendation -'some'	Change to "**some** users may be less comfortable with commercial IdPs"
	63C		87	2996	There are much better, and perhaps more common, reasons to use commercial IdPs.	Recommended addition in bold: "based on their historical interactions with government services**, or on their knowledge that commercial IdPs provide greater protection against fraud.**
	63C	8.2.1	88	3025-30	In talking with the general public , users don't generally think of identity or how authentication works - they just want it to work and want it to be easy.	
	63C	8.2.1	89	3056	Age is determined based on birthdate, which is NOT dynamic. A better example is an address.	change age to address
	63C	8.2.1	89	3065	Most privacy laws require the ability for users to completely delete their identity account. In all cases where privacy laws are in effect, deactivation would not be appropriate, only deletion. To help prevent fraud, CSPs should have audit trails and other mechanisms in place.	remove suggestion
	63C	8.2.1	89	3056	Non-preference attributes need to be verified before they can be updated by a user. It is common for bad actors to change attributes in a user account to further their purposes, such as replacing the legitimate user's address with one that they control.	"...update **preference** attributes**. Attributes that may be relied upon by RPs, such as postal address and phone number, require validation and verification, and should be subject to fraud prevention analysis, before they are updated in the system."**
	63C	8.2.1	89	3062-3064	What is the use case for this? It would appear to provide bad actors with a way to cover their tracks.	Reconsider including this. If it is retained, provide a concrete example and update the wording so it is not an avenue for exploitation.
	63C		90	3109	Redress methods are exploitable by bad actors seeking to change a legitimate users information.	Addition in bold: "Provide **secure** and effective redress"
	63C	9	93-94	3209-3212	Not only is this paragraph oddly conspiratorial against IdPs with respect to a disadvantaged population, it also assumes IdPs are evil vendors who are trying to somehow gain from the knowledge that a person is disadvantaged. This paragraph is wholly inappropriate, assumes vendors don't care about the entire population, and wrongly targets IdP vendors. For a second time, the document raises the spector of RPs colluding, which is loaded language. The reality is, if a vendor actually did what is suggested here and documented in a government document, the vendor would be canceled by its customers and perhaps face civil or criminal penalties. This paragraph shows a bias from NIST against commercial providers and undermines the role of 800-63 in fostering quality solutions that federal agencies can leverage.	Remove paragraph
	63C		114	3846	This definition contains a list of communities that may not consider themselves 'underserved', excludes other communities that do consider themselves 'underserved', and uses terminology that significant numbers of some of the listed groups themselves find highly offensive.	In all four volumes, shorten the definition to its non-controversial and non-political core meaning "The consistent and systematic fair, just, and impartial treatment of all individuals, including individuals who belong to underserved communities that have been denied such treatment"
	63C	8.2.2	91	3121-31	In talking with the general public, users don't generally think about federation or perceived risks to them - they just want to be able to gain access and want it to be easy.	remove suggestion
	63C	9	95	3253	typographical error	change to "having the RP have its own"
	63C	10			General comment: How does it work for SAML - these guidelines only discuss OpenID Connect.	include SAML examples
	63C	10.4	99-100		This example doesn't reflect today's reality with the current Federal wallet vendors and the soon-to-be operational HHS XMS.	Include a section to reflect today's current examples, which includes digital identity wallets that are cloud-based, not device bound, and an attribute exchange that will standardize the attributes in the wallets and allow for them to be utilize for multiple RPs.
	63C	10.7	101	3423	What is described IS FAL with a PKI Authenticator, as a hardware cryptographic authenticator that speaks the WebAuthn (and also FIDO) protocol IS based on a public/private key pairing, so PKI.	Suggest changing title, perhaps FAL3 with a non-smart card hardware cryptographic authenticator
	63C	1	1	396	any	that does not verify any of the authenticators…
	63C	1	1	404-408	These definitions need improvement. The subscriber account at the CSP can also be "subscriber-controlled." A wallet is not always independent of the CSPs infrastructure. "Onboarded" by the CSP does not have a clear meaning	1) as a verifier for authenticators in which the asserted data originate from the CSPs infrastructure for each assertion 2) as a data structure persisting on a subscriber-controlled device from which assertions are made independent of the CSPs infrastructure
	63C	1	1	410	The last two sentences in this paragraph are very difficult to parse. I believe this conveys the same information more clearly	With rare exceptions, federation is preferred any time the RP and the subscriber are part of the same organizational domain. Even when under the same domain, federation may still be appropriate for centralized account management and technical integration.
	63C	2.2	6	545	By explicitly stating federal agencies, we are implying this doesn't need to be done outside of federal agencies. The scope of applicability (legally) is federal agencies, but we are establishing best practices wider than that and thus shouldn't limit scope in that way. Others can decide on their own whether to model those best practices	remove "by or on behalf of federal agencies"
	63C	2.2	6	545	the inheritance works such at "or higher" is not needed. This section has requirements for FAL2 and should avoid defining requirements for other FALs	remove "or higher"
	63C	2.3	6	566	By explicitly stating federal agencies, we are implying this doesn't need to be done outside of federal agencies. The scope of applicability (legally) is federal agencies, but we are establishing best practices wider than that and thus shouldn't limit scope in that way. Others can decide on their own whether to model those best practices	remove "by or on behalf of federal agencies"
	63C	2.3	6	566	This is not a new requirement as it is inherited from FAL2	remove lines 566-568
	63C	2.4	7	609	The IdP may not have "direct acces" to the details	Restate as "only the IdP has direct knowledge of the xAL achieved by the subscriber. Consequently…"
	63C	3	9	623	The language is ambiguous as to whether the RP can always do these three things	Change lead in sentence to "Depending on the use case, the authenticated session might then be used by the RP to:"

	63C	3	10	651	make conditional. Alternatively, make clear the ones that will be in the subset (federation identifier, etc) and what might not (additional attributes)	change "is" to "may be" made available
	63C	3.1.2	11	683	Should we be referring to IdP as a role given its status with wallets? It might be better described as a function. There may be a third party controlling it, or it may just be a function built into a wallet app, or a standalone app, or could reside with a CSP or some other formulation. Calling it a role (or, worse, a party) boxes it into needing an entity, which it does not	
	63C	3.2.2	11	710	is the authorized party not a role?	
	63C	3.3.1.2	15	860	Should the entropy be future proofed	entropy meeting the latest version of 800-131A, 112 bits as of this publication
	63C	3.6	23	1098	RP can get additional consent	...without specific additional consent from the subscriber
	63C	3.7	24	1118	A federated identifier could be associated with multiple RP subscriber accounts, for instance a small biz owner that has an account as an indivdual and for the business. Sure, an RP should manage this with relationships in their system, but that's not for 63C to decide	remove restriction
	63C	3.7.1	25	1146	missing word?	maybe deleted?
	63C	3.7.2	25	1166	shouldn't the account resolution refer to 63A rather than a generic risk assessment? It's a SHALL, but doesn't accomplish anything without guidelines attached to it.	
	63C	3.7.3	25	1172	Why would there be such requirements on the RP? If I want to create native subscriber access and abandon the federation for that subscriber, that's on me. If the federation wants to block that practice, fine. But I don't think NIST should be in the business of determining that.	
	63C	general			There are too many instances in this volume of ignoring the good work of the other volumes and NIST SPs. It's nearly written as a standalone document and misses opportunities to set real requirments by pointing back to 63A, 63B. A good example of authenticated protected channel, which is left undefined. Another is requiring "a risk assessment" without specifying any conditions around it. There are also requirements that conflict with or duplicate those already met by other volumes.	
	63C	3.13	36	1552	typo "own its own"	on
	63C	3.15	38	1591	assertion in figure 4 should say bound authenticator ID	add ID
	63C	3.15	38	1606	As a general comment, the 63 suite has common language that is often abandoned in 63C. This should not be the case and 63C needs a complete review with that in mind. Specifically, is "independent of the transaction binding" different than out of band?	change to out of band
	63C	3.15	38	1603	a normative section shouldn't refer to requirements vaguely as "slightly different." Provide the requirements or reference to them	change sentence to refer to 3.15.1
	63C	3.15.1	39	1614	This section buries requirements in a long paragraph of text, yet doesn't include the needed normative statements. This is a weakness of this document overall: it prioritizes the discussion over the requirements, often leaving requirements incomplete. Here, for example, the ways the RP can deliver the authenticator never actually appear in a normative statement.	rewrite section to prioritize normative statements over discussion
	63C	4	43	1704	Change name of general purpose. This is, obviously, too general a descriptor. IdP controlled might be more appropriate.	change general purpose IdP to IdP controlled
	63C	4.2	45	1761	The subscriber is not part of the trust agreement	change all parties to list the specific parties
	63C	4.3.1	46	1779	a priori is two words. This error appears elsewhere in the document as well. Broader point: why use latin when english will do? It feels like an effort to make the document even less approachable	remove a priori throughout document and replace with common english words
	63C	4.3.1	46	1781	having a may nested in a shall like that is very confusing. So those terms are required but the might vary, thus some are not required?	SHALLs can't be MAYed. Rewrite for clarity
	63C	4.3.1	46	1797	is the point how the RP will use them? If so, that's a use specification. Purpose is ambiguous.	rewrite for clarity
	63C	4.6.5	57	2168	The requirement for notification of RPs is an issue. If I delete my google account, google should not be required to notify www.cutekittymugs.com when I delete my account. Moreover, the RP shouldn't be required to delete shared attributes. Should the IRS have to delete my personal information if my idp account is terminated (for non-fraud reasons)?	remove this requirement, particularly for non-fraud reasons
	63C	4.8	62	2295	For privacy reasons these signals may not be desirable. For instance, if my IAL2 expires at the IdP/CSP, I don't want every RP notified. I may just need to come back and re-proof and go about my day. If an RP is relying on the IdP's assertion at logon anyway, they don't need this update ahead of time; they just won't get an IAL2 assertion from the IdP. Likewise, not every RP needs to know if I delete my google account. The RP may not even know my email address, so why would the IdP signal the RP if I change it? There's just too much unneeded attribute sharing going on here.	rethink these and remove all but the essential--which may just be compromise
	63C	4.8	62	2303	Similarly, the IdP does not need to know if I delete my account at www.cutekittymugs.com. Nor does the IdP need to know what authenticators I've bound directly with the RP	rethink these and remove all but the essential--which may just be compromise
	63C	4.8	62	2313	subject to is not normative. Add a SHALL statement about privacy and security review	change subject to to SHALL
	63C	4.9	63	2336	normative statement nested in a list about a different process. Move this requirement to a standalone statement	move requirement after the list
	63C	4.9	64	2379	this is repeated from earlier in the document. Repeating normative statements can cause confusion, particularly when different words are used	delete paragraph
	63C	4.9	64	2383	this is repeated from earlier in the document. Repeating normative statements can cause confusion, particularly when different words are used	delete paragraph
	63C	4.9	64	2394	this is repeated from earlier in the document. Repeating normative statements can cause confusion, particularly when different words are used	delete last sentence
	63C	general			The document repeats requirements in different sections, stating them slightly differently, and putting them both in normative statements. This can create a signficant issue with compliance.	
	63C	4.11.1	66	2435	Stating a required time limit without setting an actual limit has the effect of no time limit. Establish one	establish a required time limit

	63C	5	69	2499	Are we assuming the wallet is associated with a multifactor device and thus has an activation factor? We should be clearer about this assumption. If we're not making it, why are we invoking an activation factor?	Clear up language around authenticators associated with subscriber controlled wallets
	63C	5.4.1	73	2597	Is there a difference between line 2590 and 2597?	
	63C	5.5	73	2604	typo, though I like the idea of an RP singing the public key	signing
	63C	5.8	75	2682	normative statement nested in a list about a different process. Move this requirement to a standalone statement	move requirement after the list
	63C	5.8	76	2696	out of band does not have the same meaning here as elsewhere and could cause confusion. It's not necessary	delete out of band
	63C	5.9	76	2712	not all requirements in the guidelines have to be met, just the ones for holder of key assertions	delete "other requirements in these guidelines are met. For additional requirements for"
	63C	6.1	79	2770	including repudiation as threat makes it sound like the IdP and Subscriber are malicious actors.	Either leave repudiation threats out or find a less hostile way of describing them
	63C	7.1	81	2785	NIST seems to have reverted to using PII instead of personal information. This is a mistake. We should be taking a more expansive view risk management associated with personal informmation as a whole, rather that just PII. This applies in many place in the document.	revert PII to personal information everywhere except where it truly is limited to PII
	63C	7.1	81	2802	this comes awfully close to making a normative statement	change "cannot" to something like IdPs must take care in gaining consent for additional uses
	63C	7.1	81	2808	colluse is a pretty charged term. It may be in some RPs interest to share data to improve their business efforts, but we don't need to put those activities into terms that suggest illegality	rewrite sentence to talk about privacy protetions and not nefarious business practices
	63C	7.1	81	2810	some bound authenticators are recognizable, not all	add some
	63C	7.1	82	2813	this again is close to a normative statement	…an identity API, implying these additional attributes may fall under the privacy risk assessment
	63C	7.1	82	2815	normative statement	change to "The SAOP can typically answer questions about…"
	63C	7.1	82	2828	normative statement	perhaps add "typically"
	63C	7.3	83	2882	normative statement that conflicts with the normative statement in the normative section (3.10.2)	change required to recommended (or the like)
	63C	8.2.1	88	3033	normative statement	replace "ought to" with "may consider"
	63C	8.2.1	88	3036	normative statement	replace "ought to" with "may consider"
	63C	8.2.1	88	3033	the topic of this paragraph is covered in the privacy section and is not really about usability	delete paragraph
	63C	8.2.2	90	3097	this is the usability section	delete sentence about security practices
	63C	8.2.2	90	3095	it seems several items in this list are already covered by normative statements elsewhere in the document	remove repeated recommendations
	63C	8.2.2	91	3124	what's the consideration here? What action is one to take as a result of this paragraph?	delete paragraph
	63C	9	93	3179	normative statement	remove sentence
	63C	9	93	3181	is there a word missing from this sentence? It feels like it is supposed to have a SHOULD in there, but that would be normative and should not be in this section.	review for completeness; do not make normative
	63C	9	93	3183	normative statement. Also this is the equity section, not the usability section, so this last sentence does not need to be here.	remove
	63C	9	93	3186	the first two sentences are covered in 63A and 63B	remove first two sentences of paragraph
	63C	9	94	3224	Let's not call people thoughtless	replace "thoughtless clickthrough" with "users clicking through without fully understanding the implications of their consent."
	63C	9	94	3225	entire paragraph is normative. Moreover it has little to do with equity. These are privacy and usability requirements (again, in an informative section) and while such things can impact equity, in this case their rightful place is in a normative section on privacy	remove paragraph
	63C	9	94	3234	what differences in requirements are we talking about here? I assume by "interests" we mean incentives, but what are the requirements? Also, "has to be addressed" is awful close to normative.	rewrite for clarity; do not make normative
	63C	9	94	3240	motivations makes it sound quite nefarious for both the government and "private" entities. If you mean commercial, say commercial, and be neutral with the language	rewrite for neutrality and clarity
	63C	9	94	3243	awkward sentence. Avoiding passive voice may help	rewrite for clarity
	63C	9	94	3246	normative statement	remove sentence
	63C	9	95	3249	This paragraph is difficult to parse. What inequity? If an RP only accepts one IdP, that impacts everyone equally that wishes it accepted more than one. The solution of an account recovery process is fine, but what follows ("allows for the secure linking...") doesn't make sense given the start of the paragraph established it was difficult to find a second IdP the RP would accept.	remove everything about multiple IdPs and keep the part about recovery