National Institute of Standards and Technology (NIST)

ID.me's Comment to NIST Special Publication 800-63-4

14 April 2023

COMPANY DETAILS

Company Name: ID.me, Inc.

Address: 8280 Greensboro Drive, Suite 800
McLean, VA 22102

Respondent: Blake Hall

Respondent's role: CEO & Founder

Email: Blake@ID.me

14 April 2023

VIA E-MAIL to dig-comments@nist.gov

Re: Comments on the draft fourth revision to the four-volume suite of Special Publication 800-63-4, *Digital Identity Guidelines*.

ID.me appreciates everything NIST does to promote secure and equitable access and prevent fraud when consumers interact with their government. Thank you for everything you are doing to accept public feedback as you respond thoughtfully and in a data-driven way.

After winning NSTIC grant funding 10 years ago, we believe that ID.me has evolved to be the largest Credential Service Provider (CSP) in America by volume of Identity Assurance Level 2 credentials issued and maintained with over 40 million credentialed users. These credentials have been issued in support of 12 federal agencies, 36 state agencies, and over 50 healthcare organizations. Our healthcare deployments are particularly meaningful as we are helping providers and states combat the opioid epidemic by enhancing the security of prescribing flows.

Based on these experiences, we would like to share with NIST our perspectives on digital identity verification to help NIST make more informed decisions about the final guidance to be included in the -4. In doing so, we will focus on what matters most to the end-user: secure and equitable access, fraud prevention, and user control over their data and experience.

At a high level, there are three things ID.me wants NIST to consider:

1. Security and access are compatible. We applaud NIST for upgrading normative requirements for IAL2 with mandatory Presentation Attack Detection, a critical tool to prevent identity theft and fraud. Identity theft and fraud reduces access and equity, and it is important to require CSPs to innovate to expand access and equity at the baseline level of security that protects Americans and agencies from rampant identity theft and fraud.
2. Industry and government agencies need a standardized set of metrics by which to assess outcomes and customer experience. **NIST can start by requiring auditors to measure the outcomes the CSP is driving** - specifically the overall pass rate net of fraud and time spent verifying – prior to drilling down into components of a single pathway. A focus on access, equity, and security should apply to outcomes and to all components, not just one.
3. Research suggests the best algorithms often perform more accurately than humans, but the optimal solutions feature the best algorithms and humans working together so there are multiple relief valves and escape hatches to mitigate any real or potential bias. For that reason, we applaud NIST for making Trusted Referees a normative capability.

This comment starts with an overview of these three points and then gets into specific feedback and response to comments invited by NIST.

NIST's current guidance enabled solutions that provided agencies with more security and more equitable access relative to legacy identity verification methods.

In a 2019 report, the Government Accountability Office found that until "agencies take steps to eliminate their use of knowledge-based verification, the individuals they serve will remain at increased risk of identity fraud."[1] Despite this warning, many government agencies continue to use data brokers and Knowledge Based Verification methods. During the pandemic, weaker methods of identity proofing were overwhelmed with fraud.

On August 27, 2020, Pew reported, "States that were generous and quick to help workers were also quick to be targeted by scammers. In response, states have had to slow down the processing of claims, delaying payouts to people supposed to be getting them."[2] In December 2020, USA Today interviewed an attacker who claimed he was successful "about one in six times" when filing a claim even when the system used "additional verification."[3]

The connection between security and access could not be more clear. When security proves inadequate to protect systems, Americans and government agencies suffer from elevated rates of identity theft and fraud. In turn, elevated rates of identity theft and fraud reduce access and equity as noted in Pew's reporting along with many anecdotal examples in other media outlets.[4]

While we applaud NIST for driving to increase flexibility through concepts such as IAL1 and "tailoring," NIST should be as specific as possible when it comes to the conditions that would be appropriate for IAL2 or tailoring. Given that many government agencies still use KBV solutions, weakening the IAL2 guidelines could lead to more instances where solutions prove ineffective to do the very thing agencies need them to do - separate good users from fraudsters. Current NIST requirements that mandate agencies document deviations is wholly appropriate. If there are mitigating controls at the systems level that compensate for specific NIST technical and policy controls, then agencies can indeed tailor requirements accordingly but they must justify and document equivalency to the controls they are bypassing. We strongly recommend specific guidelines and risk criteria to separate situations where IAL1 is appropriate from IAL2.

[1] https://www.gao.gov/assets/gao-19-288.pdf
[2]https://www.pewtrusts.org/en/research-and-analysis/blogs/stateline/2020/08/27/fight-against-fraud-slows-payments-to-unemployed
[3]https://www.usatoday.com/in-depth/news/investigations/2020/12/30/unemployment-fraud-how-international-scammers-took-36-b-us/3960263001/
[4] https://www.propublica.org/article/how-unemployment-insurance-fraud-exploded-during-the-pandemic

Without data based rationale, an agency avoiding the use of one technology or control may inadvertently cause harm because they are focused on a single method or setting rather than on the Key Performance Indicators (KPIs) that measure outcomes of the system as a whole. KPIs like pass rates, fraud rates, and the impact on customer experience of both successful identity proofing as well as the time spent remediating instances of identity theft provide vital context.

ID.me's approach has relied upon NIST establishing baseline requirements that protect all Americans from scalable events of identity theft and fraud. NIST IAL2 proved to be the critical control that reduced an unprecedented wave of identity theft and fraud in unemployment benefits during the pandemic. In July 2022, Deputy District Attorney for Sacramento County Nick Johnson said, "The mass volume, what we were seeing, thousands and thousands of victims on an almost daily basis coming in, that's done because of the (ID.me) verification."[5] Importantly, this comment is a testament to the effectiveness of NIST guidelines.

Once those security requirements are met, ID.me has worked with NIST over the years, starting with our NSTIC grants, to innovate to improve access, equity, privacy, and security for more people to access digital services in order to improve customer experience. Based on agency testimony, we are driving significant access and equity gains even while we meet a higher security standard, IAL2, versus other solutions.

When former IRS Commissioner Rettig testified before the Senate Appropriations committee on May 3, 2022, he noted ID.me **nearly doubled** IRS's pass rates. *The Washington Post* later reported this doubling of pass rates included "low income earners and minorities." At the same time, ID.me also moved multiple agencies and applications from a legacy KBV or LOA3 configuration to the more secure NIST SP 800-63-3 IAL2/AAL2 policy.

Today, more than 40 million people have already verified their identity to IAL2 and enrolled into AAL2 – so the identity proofing event is portable to other agencies. Of these 40 million IAL2 credentialed users, more than six million Americans verified through ID.me's Trusted Referee path. With a portable credential, users that verify via Supervised pathways are able to then engage with other agencies in a secure fashion, in parity with those users that verify via Unsupervised pathways. This approach is consistent with OMB M-19-17 guidance that calls on agencies to reduce the identity proofing burden on the public through the use of shared services for authentication.

ID.me's empirical data shows that preverified users convert at the highest rates of any method, relative to Unsupervised Remote or Trusted Referees, at rates up to 99% while also enjoying a superior customer experience as releasing verified status requires simply logging in and

[5] KCRA, "EDD fraud 2 years and counting: Criminals still stealing identities to cash in on benefits," 5 July 2022, https://www.kcra.com/article/edd-fraud-2-years-and-counting-criminals-still-stealing-identities-to-cash-in-on-benefits/40519380#

providing permission which takes users about 20 seconds. Agency testimony and ID.me's performance data makes it clear that under OMB's current shared service policy and NIST's current framework for IAL2, which balances protections against identity theft and a fraudster claiming your identity, with innovations to enable greater access and equity, that government agencies can have both security *and* greater access at the same time – without compromise.

For example, federal agency asked ID.me to do a deep-dive on one vulnerable demographic – users from Puerto Rico – and found that a NIST conformant solution that offers both Unsupervised and Supervised Remote pathways resulted in a 6x increase in access rates for users from Puerto Rico, when compared to legacy records-based solutions. Users from Puerto Rico have one-third the median income of the overall US population, are up to 7x as likely to be credit invisible, and are primarily Spanish speakers. We expect that other demographic groups that share these characteristics would experience similar access gains.

Trusted Referees as a safety valve drive substantial equity gains for historically underserved communities. There are 45 million consumers who have thin, incomplete, incorrect, or unscorable credit. These users have traditionally been left behind by solutions powered by credit bureaus and data brokers that do not offer any recourse or relief valves for remote identity proofing. By offering Supervised Remote as an alternative pathway, NIST is making it possible for users without a large digital footprint to attain the same outcomes as users that do. NIST's Supervised Remote pathway has enabled CSP's and agencies that adopt an omnichannel model to "level the playing field" for consumers.

Furthermore, NIST leveled the playing field permanently for disadvantaged populations by enabling portability through standardized identity assurance levels. For example, because IAL2 is a standard that is well-defined, widely understood, and accepted broadly, the use of portable credentials is stripping friction out of a consumer's entire "digital life." NIST guidelines and standardizations have turned verification into a 1-time event to create a credential that can then be used quickly, safely, and securely at other agencies, when needed. This means users who formerly struggled with each agency's unique and different verification process can have frictionless access at every agency with which they need to engage. **This is only possible if agencies have a standardized structure of IAL's that allow them to trust credentials issued at other agencies.**

Industry and government agencies need a standardized set of metrics by which to assess performance in terms of outcomes and customer experience.

ID.me is concerned that focusing on any one component – and removing or including it – is too narrow a view when it comes to identity assurance. Digital identity solutions providers are responsible for including equity and access "safety valves" to mitigate the risk of bias for ALL

potential component failures tied to identity resolution, identity validation, and identity verification. Additionally, digital identity solutions providers should account for populations such as international users, who may not be present in an issuing or authoritative source, Americans without credit history or an accurate presence in records, unhoused populations, etc.

More transparency in the market as to which CSP can serve a particular type of user e.g. international, English as a Second Language (ESL), no credit history, data listed inaccurately in records (e.g. recently moved or name change), will provide agencies with data to compare access and equity capabilities across CSPs. Coupled with performance data that measures outcomes by demographic group and community, agencies will be able to procure the most effective CSPs.

One challenge facing agencies is that neither they nor industry have a standardized set of metrics or common language by which to (1) assess performance of CSPs and (2) articulate impact on their user bases. While pass rates are easy to conceptualize, they can be misleading if not taken in the right context, such as the NIST policy in-use. In other words, without a common set of metrics or "scoreboard," agencies risk trying to compare an IAL2 pass rate to a less-rigorous KYC or account onboarding pass rate. As another example, comparing initial verification pass rates for a point solution or component vendor to that of a portable solution across its entire network is like comparing apples to oranges: one is for a single transaction at a single agency, and the other could include access at multiple sites across a user's entire digital life.

ID.me recommends NIST implement the following normative requirements to Base Volume Section 5, page 23, and supplement them through regular audits, to ensure consistent and accurate measurement of performance.

1. **Digital identity solution providers SHALL be required to report the following metrics with independent auditing to validate CSP reported metrics:**

- **Pass Rate:** Overall percentage of unique users who complete proofing
- **Fail Rate:** Overall percentage of unique users who try but fail a given step e.g. unable to pass document authentication despite two attempts
- **Abandon Rate:** Overall percentage of unique users who abandon at a step in the flow without attempting that step
- **Fraud Rate:** Overall percentage of estimated fraudulent users

Research suggests that using top-performing algorithms with trained human agents to mitigate any real or potential bias leads to optimal outcomes.

The limitations of individual components are why ID.me has always taken the approach of combining best-in-breed algorithms with human reviewers. The way we like to think of it, an algorithm can say, "Yes" or "Maybe", but only a human should be able to say, "No."

Recognizing that biometrics are in the spotlight, we empathize with NIST's need to write practical guidance when it comes to the use of biometrics for comparing the claimant of Personal Identifiable Information (PII) to the strongest piece of evidence.

That being said, we also understand that, based on NIST and DHS data, the top-performing biometric algorithms have demonstrated consistent performance across different demographics in government tests. Rather than avoiding facial recognition technology, NIST should provide guidance on how to use it responsibly, through the lens of an entire system. This can be done in two ways: (1) use only the best performing algorithms and/or (2) use them in collaboration with human reviewers.

- **Using the best-performing algorithms.** As far as we are aware, two government agencies study and test face matching technology rigorously – NIST and Department of Homeland Security. In recent, publicly available test results, the top-performing algorithms tested by each agency demonstrated high degree of accuracy (*e.g.*, 5 algorithms with over 99.9% accuracy) as well as consistent performance across demographics (*i.e.*, true match rates of over 99.3% for all tested demographics, at a 1 in 100,000 False non-match rate). The same tests show weaker technology does not perform as consistently across demographics. NIST enforcing a performance threshold for the use of biometrics is exactly the right thing to do. It will promote broader adoption of the best algorithms and encourage weaker ones to innovate and improve in order to be more competitive. Additionally, because these tests are repeated over time, we can see from the public results that facial recognition technology accuracy continues to improve with each subsequent test.

- **Using them in collaboration with human reviewers**. Performance of technology should always be viewed within the context of how it compares to alternatives. The National Academy of Sciences (NAS), an authoritative source of original and independent scientific research, did just such a comparison when it assessed performance of forensic facial examiners to face identification algorithms. They published the results of their study in the Proceedings of the National Academy of Sciences, a peer reviewed journal. Their findings show that leading algorithms are more accurate than even forensic examiners, who specialize in facial comparison, at comparing two human faces. The study also found that "collaboration among humans and between humans and machines

offers tangible benefits to face identification accuracy in important applications." Even the most well-intentioned humans can exhibit bias, suggesting that solutions that combine humans and algorithms working in collaboration achieve "the most accurate face identification possible."

Use of biometrics has a critical role in Unsupervised Remote verification. When weighing concerns raised over the use of it, **NIST should provide guidance on how CSPs can best deploy it responsibly, through (1) the deployment of algorithms that have demonstrated consistent performance across demographics in government testing and (2) use of biometrics in a system that also includes backing them up with human reviewers.**

Specifically, we recommend NIST implement the following normative requirements in 63A, Section 5.1.8, page 23, line 948:

1. CSP **SHALL** offer human review or alternative verification pathways to attain the same IAL level in the event a user is unable to successfully complete a biometric check

2. **CSPs SHALL submit themselves to Independent audits conducted by a reliable third party, such as the Kantara Initiative, to validate the digital solution's provider's calculations of its performance.** For example, an auditor should look into users marked as fraudulent to determine if those individuals are in fact fraudulent actors or if the vendor had a false positive. This independent check would force reconciliation of vendor reported data to ensure accurate measurement.

3. **Digital identity providers SHALL be required to monitor the following KPIs:**

- **True Pass Rate:** Overall Pass Rate minus the Fraud Rate
    - If attempted fraud blocked is 3% one month and 20% the next month, it doesn't make sense to show a 17% drop in performance when the solutions provider is actually effectively preventing identity theft and fraud. Rather, the key metric for accessibility should focus on the experience of non-fraudulent users.
- **Customer Experience:** Time spent identity proofing. This time should be the unique user average of new identity proofing events at each agency:
    - **Federated Users** who arrived at the agency with an IAL2 credential, logged in at AAL2 and provided consent to verify at a different agency.
    - **Unsupervised Remote** users who proofed via self-serve flows.
    - **Supervised Remote** users who proofed via video chat.
    - **In-Person** users who verified at an in-person location.

As an extension of these normative requirements, agencies will be able to filter metrics through the lens of equity to understand the impact of a solutions provider on a given community:

- Zip Code Level Performance
- Demographic Performance, including but not limited to:
    - International Users
    - Rural vs. Urban Users as measured by population density thresholds
    - Unhoused Populations
    - Age Bands
    - Race/Ethnicity

One challenge NIST will need to work with industry to address is controlling for the exogenous impact of user incentives on performance. The incentives of a user to undergo identity proofing can significantly impact metrics, like Abandon Rate, and are largely outside of the digital identity solution providers control. For example, a doctor who needs to meet the NIST 800-63 requirements to electronically prescribe according to state law will have a much lower abandonment rate than an American who might have a mild curiosity to learn a little bit more about their retirement benefits data.

Standardized metrics and independent review of performance will ensure a fair and impartial scoreboard for access, equity, reliability, security and privacy-protection. With transparency around solution performance at a given level of trust (IAL1, IAL2, etc.) and for the same use case, trust will increase and NIST and other policy makers will have data to further improve.

As NIST and DHS do with facial recognition test results, performance data validated or established by independent third parties could be made public. This sort of standardized and transparent performance metrics will allow comparison of CSP's relative to each other.

I. CALL FOR PATENT CLAIMS.

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication . . . This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to the ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ID.me is not aware of any patent claims that would be required for compliance, nor are we aware of any third-party claims or patent applications relating to this ITL draft publication.

II. INVITED COMMENTS

ID.me's comments and recommendations for the topic areas NIST requested are detailed below.

A. Identity Proofing and Enrollment

1. NIST sees a need for inclusion of an unattended, fully remote Identity Assurance Level (IAL) 2 identity proofing workflow that provides security and convenience, but does not require face recognition. Accordingly, NIST seeks input on the following questions:

a. What technologies or methods can be applied to develop a remote, unattended IAL2 identity proofing process that demonstrably mitigates the same risks as the current IAL2 process?

The crucial factor setting IAL2 apart from other legacy NIST identity verification policies is the establishment of a concrete connection between the Personal Identifiable Information (PII) being claimed and the person asserting the claim during the verification process. To our knowledge, there currently is no technology available that can effectively mitigate this risk in a remote, unattended setting as efficiently as biometric or physical comparison to the strongest evidence presented. In the context of persistent and widespread attacks by fraudsters and crime rings during the pandemic, facial recognition and Presentation Attack Detection proved to be critical controls. ID.me actually **saw that pass rates went up** by single to double digits for other verification components when IAL2 and biometrics were introduced across multiple agencies.

Alternative Technologies

In an effort to assess potential alternative solutions, ID.me has evaluated device fingerprinting and behavioral biometric technology to determine their efficacy in mitigating the risk of a claimant not being the rightful owner of the claimed PII. While ID.me acknowledges the usefulness of these technologies in characterizing the claimant, they do not offer a definitive "yes" or "no" answer, thus failing to reduce the risk to the same extent as biometric or physical comparisons.

- **Device fingerprinting** comes with its limitations, as it necessitates the assumption that the device being used is in the possession of its rightful owner. Additionally, there are many situations where devices are shared within families or at government centers.
- **Behavioral biometric** technologies similarly require the establishment of a reference point before verification, such as confirming the legitimate owner of the PII does indeed type or click at the cadence observed on the day of verification.
- **Risk scores** provide a quantitative metric that is computed from multiple data sources and data points. This technology is often used as a "data point" rather than a "decision"

because it does not provide a definitive "yes" or "no" answer on whether or not a claimant is the rightful owner of a set of claimed PII

Moreover, ID.me points to the findings of the NIST Face Recognition Vendor Test (FRVT), which demonstrates top-performing 1:1 face matching algorithms perform consistently across demographic groups. To our knowledge, there has been no comparable testing conducted on potential alternatives described above. Any alternatives considered should be subject to similar assessment to FRVT and should exhibit equally high performance in terms of accuracy and consistency across demographics prior to use. Again, requiring CSPs to report on KPIs that measure outcomes, versus specific components or inputs, will enable a comparison of performance across solutions that meet IAL2 to measure performance by results versus process.

Control of Digital Accounts

In the current draft of Section 4.4.1 Identity Verification Methods, NIST provides guidance that CSPs can use demonstration of control of a digital account via AAL2 OR FAL2 as an alternative to physical or biometric comparison. **If NIST is looking to implement an alternative method for achieving "STRONG verification," demonstrating control of a digital account is not comparable to a biometric or physical comparison of the applicant to a piece of identity evidence. This is especially true when considered in the context of ever evolving and complex schemes used by fraudsters. As such, ID.me strongly recommends that detailed normative guidance should be established in order to maintain the integrity and effectiveness of the existing IAL2 standard. If such normative guidance is not included, IAL2 standards will not in fact be standardized in practice which will undermine security.**

ID.me recommends that NIST develop the following related to use Digital Identity Evidence:

- **General**
    - Clear definitions should be established for issuers, verifiers, strengths of digital identity evidence.
    - A table that explicitly lists the strengths of digital identity evidence and the characteristics of each strength
    - What are the strengths or acceptable methods for *validating* digital identity evidence?
    - What are the strengths or acceptable methods for *verifying* digital identity evidence?
        - What identity proofing measures were established while creating the external account?
        - What constitutes sufficient tenure to trust the external account?
        - What constitutes sufficient activity to trust the external account?

- **Issuers**
    - What are the requirements that issuers need to follow to ensure the proper issuance of digital identity evidence? Do these requirements change at all depending on the strength of identity evidence that is being issued?
    - Are there any baseline security standards or certifications that they need to adhere to (e.g., Can a FedRAMP low system be used to issue STRONG digital identity evidence like an mDL?)
    - What responsibility do issuers have to ensure that the digital identity evidence was properly bound to the subscriber's device?
    - Are there are minimum identity proofing requirements for the issuer to perform before the digital evidence is bound to the subscriber's device (e.g., IAL1 minimum for the issuance of FAIR digital evidence, IAL2 for STRONG, and IAL3 for SUPERIOR?
    - Are issuers permitted to issue a single piece of digital evidence to multiple subscriber devices?
    - What processes should issuers follow around the potential loss, theft, and/or unauthorized duplication of digital evidence?
    - How are CSPs/Verifiers informed if previously validated digital evidence is subsequently revoked or suspended?
- **CSPs/Verifier**
    - What standards should CSPs/verifiers follow to ensure the proper remote presentation and verification of digital evidence?
    - Are there minimum authentication requirements for the subscriber to present digital evidence remotely (i.e., What if the subscriber's device doesn't require a PIN or Face/TouchID to present digital evidence)?
    - How does the CSP/Verifier verify that the applicant is presenting a valid piece digital evidence that they are the true owner of and not just a valid piece of digital evidence that could be owned by someone other than the applicant?
        - For example, how does the CSP confirm that the applicant is the true owner an mDL that was validated?
            - Attribute Comparison: A minimum set of validated attributes would need to be sent back to the CSP/verifier for comparison/matching to ensure that the applicant going through the identity proofing process presented evidence that they own.
            - Biometric or Physical Comparison: Biometric or physical comparison could be another method, but if the purpose of this new verification method is to provide an alternative to biometric/physical comparison, then it seems that attribute comparison is the only viable option.

If this modality is too immature in terms of measurement and data to establish equivalency, then ID.me recommends that NIST should remove this as an acceptable method.

IAL2 as a Fraud Deterrent

It is also important to note that based on ID.me's experience, quantitative evidence shows IAL2 in its current form is a deterrent to fraud. Alternative technologies should be required to demonstrate similar effectiveness prior to consideration in replacing the physical or biometric comparison step.

While protecting unemployment benefits during the pandemic, ID.me saw that the *presence* of IAL2 as a security policy had the added benefit of being a *deterrent to attempted fraud*. It follows that government agencies should expect to see a higher rate of attempted fraud against systems that do not comply with IAL2. Fraudsters closely monitor the controls and security policies used by various agencies and openly share such information on the dark web.

During the pandemic, ID.me compared the attempted fraud rate in states that opted to use the legacy NIST SP 800-63-2 LOA3 (no biometric or physical comparison) policy versus the current IAL2 policy. **We observed that *attempted* fraud rates were 11 - 29% higher in states that used the legacy LOA3 policy versus a large state that used the IAL2 policy.** These higher attempted fraud rates were even more concerning because LOA3 is less effective at actively preventing fraud.

ID.me also found that pass rates increased when states shifted to the IAL2 policy versus an LOA3 policy that does not require a physical or biometric comparison of the user to a photo ID. The logical explanation for this increase is that fraudulent actors did not want to go through the IAL2 security controls. As a result, attempted identity theft and fraud decreased significantly.

For these reasons, an **identity verification system that does not adhere to the IAL2 standards will see both higher attempted fraud rates and higher post-verification fraud rates relative to IAL2 security**.

b. Are these technologies supported by existing or emerging technical standards?

ID.me is not aware of existing or emerging technical standards for technologies described above.

c. Do these technologies have established metrics and testing methodologies to allow for assessment of performance and understanding of impacts across user populations (e.g., bias in artificial intelligence)?

ID.me is not aware of established metrics and testing methodologies that would allow for the assessment of performance across user populations for the technologies listed above.

2. What methods exist for integrating digital evidence (e.g., Mobile Driver's Licenses, Verifiable Credentials) into identity proofing at various identity assurance levels?

Integrating digital evidence, such as Mobile Driver's Licenses (mDL) and Verifiable Credentials, into identity proofing can be done using various methods. However, these methods can potentially be subject to or held captive by operating systems like Apple, Android, Windows, etc. configuring the capture and sharing of mDL information. To address these issues, ID.me recommends NIST consider the following when providing guidance on mDLs:

1. **Advocate for Platform-Agnostic mDL APIs:** A platform-agnostic mDL standardized API would enable seamless integration of mobile driver's licenses and other digital evidence across multiple operating systems and platforms. This API would allow CSPs such as ID.me to query and verify the user's mDL without worrying about the specifics of an individual platform's implementation.

2. **Native and Browser-based APIs:** To support both native and web-based applications, platforms should provide APIs that can be accessed natively by apps and through web browsers. This would ensure that users can access services requiring identity verification without having to download a separate app when using a web flow.

3. **Seamless User Experience:** To prevent users from having to leave their current web flow, platforms should design their APIs to enable seamless access to mDL and other digital evidence. This will help create a more streamlined user experience and reduce friction in the identity verification process.

4. **Interoperability:** Platforms should work together to ensure interoperability between their systems, enabling seamless integration of mDLs and other digital evidence across various platforms and operating systems. This will help in creating a unified system for identity proofing, regardless of the user's device or platform (e.g., using Windows on a computer at work, Android on a mobile phone, and Apple on a computer at home).

5. **Privacy and Security:** The integration of digital evidence into identity proofing should be done with privacy and security in mind. This includes encrypting data in transit and at rest, as well as implementing strong access controls to protect sensitive user information. Additionally, the use of other security technologies can help further enhance privacy and security by ensuring that personal data is protected.

6. **Compliance with Standards and Regulations:** Ensuring that the integration of digital evidence into identity proofing complies with relevant industry standards, such as ISO/IEC 18013-5 for mDLs, and data protection regulations, such as CCPA is crucial for building trust and maintaining user privacy.

By incorporating these methods and principles, platforms can effectively integrate digital evidence into identity proofing processes at various assurance levels, providing a seamless and secure experience for users.

3. *What are the impacts, benefits, and risks of specifying a set of requirements for CSPs to establish and maintain fraud detection, response, and notification capabilities?*

It should be NORMATIVE for CSPs to establish and maintain fraud detection, response, and notification capabilities. Such a requirement is critical because in addition to fraudulent use and/or loss of public funds, weak fraud controls lead to decreased access to benefits and services for eligible individuals. When fraudsters exploit stolen identities to claim benefits, they prevent legitimate users from using their own identity to access the benefits and services to which they are entitled.

To address these issues, NIST should go a step further than current guidelines and require fraud programs include the following features:

- Include measures that stop fraud both during and after verification;
- Provide human-powered relief valves and reviews to mitigate the impact of false positives (*i.e.*, an algorithm can say "yes" or "maybe" to verification, but only a human can say "no");
- Include capabilities to stop social engineering, including raising victim awareness that they may be under the manipulation of a fraudster;
    - These controls should also be included via proofing scenarios where Trusted Referees and/or Applicant References are involved.
- Provide notification directly to agencies about users determined to be fraudulent;
- Monitor and report on evolving attack vectors; and
- Assess and deploy new controls in response to evolving attack vectors.

As an extension of the 4th point about notification, and outside of the scope of the 800-63-4, NIST is encouraged to work with industry and other agencies (*e.g.*, DSH, CISA, etc.) to develop a neutral, centralized platform for Credential Service Providers (CSPs) to report and exchange information on fraudulent activities. This sort of "clearinghouse" could mitigate the risks associated with identity theft and fraud and promote tighter collaboration across CSPs and agencies that rely on them. The benefits of doing so include:

- **Enhanced fraud detection:** By prioritizing fraud detection alongside equity and access, ID.me has generated valuable insights on domestic and international threats, identified malware campaigns, and preempted fraud schemes before impacting their partner network.

- **Reduction in improper government benefit payments:** Streamlining fraud detection helps ensure that benefits are distributed to their intended recipients.
- **Mitigation of identity theft-related damages:** With nearly 300 million victims affected by data breaches in 2021, as reported in the 2023 National Cybersecurity Strategy, proactive fraud control is essential.
- **Regulatory compliance:** Fraud control guidelines can assist CSPs in adhering to industry-specific and data regulations, minimizing potential fines and penalties.
- **Effective response and communication:** Clearly-defined processes allow CSPs to promptly address and communicate security incidents, reducing the impact to operations and enhancing timely notification of all those impacted.

While the implementation of robust anti-fraud measures may marginally increase CSP operating costs, the importance of ensuring government benefits are only dispersed to their intended recipients far outweigh these expenses. ID.me's success in providing IAL2 verification, multiple verification pathways, 24/7 member support, and comprehensive fraud capabilities demonstrates the feasibility of this approach.

Moreover, as noted above, when it comes to protecting government benefits, weak fraud controls negatively impact equitable outcomes. Each instance of fraud committed with a stolen identity directly harms a legitimate user, preventing them from accessing much-needed benefits. Because low-income demographics are eligible for and in need of benefit programs, they are more likely to be negatively impacted by such fraud. Moreover, a delay in benefits while the fraud is investigated and resolved is more likely to cause significant distress for the individuals and/or household members. Standardizing fraud detection and prevention measures will also serve to improve trust and guarantee robust security across the board for benefit programs.

a. Are there existing fraud checks (e.g., date of death) or fraud prevention techniques (e.g., device fingerprinting) that should be incorporated as baseline normative requirements?

There are several existing fraud checks and fraud prevention techniques that should be incorporated as baseline **SHOULD** requirements to enhance security and reduce the risk of fraud. These include:

- **Device fingerprinting and device tenure:** This method identifies unique device characteristics and tracks the length of time a user has been using a specific device. This information can be used to determine if a device is associated with fraudulent activities or if it is being used by a new user, which may warrant additional verification steps.

- **Synthetic identity detection:** This technique helps identify and prevent the use of fabricated identities, which are often used to commit fraud. By analyzing various data points, organizations can detect patterns that may indicate the use of a synthetic identity.
- **Contextual clues to eliminate Social Engineering:** Providing users with clear instructions and warnings throughout the verification process can help them recognize when they may be falling victim to a scam. For example, on-screen language can alert users not to proceed if they are not applying for unemployment benefits, potentially preventing them from being manipulated by a romance scam, lottery scam, or job application scam.
- **Multi-Factor Authentication (MFA) - phishing resistant:** Requiring phishing-resistant MFA for AAL2 (Authentication Assurance Level 2) can help align with zero-trust strategies and accelerate the adoption of more secure MFA methods. This ensures that multiple forms of verification are used to establish the user's identity, reducing the risk of unauthorized access. NIST could also consider making these NORMATIVE to align with the national Zero Trust strategy.
- **Continuous threat research and dark web monitoring:** By continuously researching and monitoring emerging threats and activities on the dark web, organizations can stay informed about new tactics used by cybercriminals and proactively develop strategies to protect against these threats.

Additionally, ID.me believes that **Presentation Attack Detection should be a NORMATIVE requirement for IAL2 because of its effectiveness in fraud prevention**:

- **Liveness or Presentation Attack Detection**: This is a critical aspect of identity assurance systems that aims to discern genuine biometric data from falsified or spoofed input. Enhancing the robustness of these systems ensures reliable identity verification and mitigates the risk of unauthorized access.

Incorporating these fraud checks and prevention techniques as baseline normative requirements can significantly enhance an organization's security posture and reduce the risk of fraud.

i. If so, at what assurance levels could these be applied?

ID.me's above stated recommendations could be applied to IAL2/AAL2 requirements as a way to further standardize integrity across CSPs and protect the credentials.

b. How might emerging methods such as fraud analytics and risk scoring be further researched, standardized, measured, and integrated into the guidance in the future?

Emerging methods like fraud analytics and risk scoring have the potential to significantly impact cybersecurity and risk assessment practices. To further research, standardize, measure, and integrate these methods into NIST guidance, several steps should be taken:

1. **Benchmarking and testing:** Establish standardized metrics and methodologies to evaluate the performance of fraud analytics and risk scoring models. This will facilitate the comparison of different approaches and help identify best practices. It will also enable NIST to assess performance of different approaches across user populations.

2. **Guidance and frameworks:** Develop guidelines, frameworks, and best practices for implementing fraud analytics and risk scoring in various sectors. This includes the integration of these methods into existing NIST guidance.

3. **Standardization and interoperability:** Develop standards and protocols to ensure that fraud analytics and risk scoring models can be easily integrated and shared across different platforms, systems, and organizations. This will help promote interoperability and allow organizations to leverage the capabilities of these methods more effectively.

4. **Privacy considerations:** Address privacy challenges related to the use of fraud analytics and risk scoring, such as data privacy, ethics, and potential biases in algorithms. Specifically, NIST should endorse a privacy schema that limits CSPs to collecting data reasonably related to the purpose for which it is collected and processed. ***Requiring CSPs to adhere to a justifiable basis for collection and process would inhibit the collection of massive amounts of extraneous data under the umbrella of "security" at the expense of individual privacy.*** This would encourage the appropriate data minimization practices in line with existing and developing privacy legislation. Further, risk scoring should be evaluated, and regulated, in a manner comparable to other scoring methodologies that affect other eligibility determinations - namely considerations of credit as covered by the Fair Credit Reporting Act (FCRA). ***As currently presented, risk scoring related to identity verification is strikingly opaque, with individuals being neither apprised of the criteria by which their "score" is evaluated, nor given any opportunity or avenue by which to address inaccuracies in the underlying data set or the presumptions underpinning the scoring methodology.***

By taking these steps, NIST can promote use of fraud analytics and risk scoring methods in a responsible, privacy-enhancing way.

c. What accompanying privacy and equity considerations should be addressed alongside these methods?

Risk scores that lack a mandated structure and transparency raise serious privacy and equity concerns when used to gate access to public resources (as opposed to voluntary commercial transactions such as credit card application processing). The biggest privacy concern raised by methods relying on risk scoring is understanding what information and scoring rubric is used to develop the risk scores.

The CSPs strongly advocate for the use of risk scores (*i.e.*, data brokers) passively build a profile about an end-user by buying and selling massive data sets and crawling public records, most often without users consent or awareness. These data brokers also have limited opt-out policies because it is difficult to exclude single users when building profiles in this way. The model that underpins most risk scoring approaches treats someone's identity like an independently tradeable commodity rather than a personal asset belonging to an individual.

In order to promote privacy and equity when considering risk scores, NIST should:

- **Include guidelines for transparency around risk scores**. As part of privacy assessments, solutions that use risks scores should be required to post information in their privacy policies related to: the fact that a risk score has been developed, how that risk score was developed, what data sources were used to do so, and what users can do to decrease their risk scores (similar to how credit reporting works today)
- **Create normative guidance requiring CSPs to allow ANY user to opt out at any given time**. This should include direct communication to the end user when their information has been deleted
- **Require an alternative path for user verification** (*e.g.*, video chat or in-person verification) if a "risk score" is decided by an RP as a failed verification.
- **Require notice** similar to FCRA requirements indicating basis of failed verification decision with specificity (*e.g.*, source and nature of information serving as basis) sufficient for the user to challenge/appeal the accuracy of the information or data set used for decisioning.

When risk scoring relies on an undetermined set of broad data collected about users instead of their specific attributes, it becomes akin to credit scoring, which uses aggregated data to determine a "score." While risk scores attempt to evaluate the likelihood of fraud, credit scores assess the probability of debt repayment. Given the similarities, NIST should consider risk scoring as subject to regulations similar to those governing credit scoring, such as the Fair Credit Reporting Act (FCRA). The FCRA and similar state level statutes and regulations aim to eliminate bias and enhance transparency in the scoring process. Similarly, here, individuals should have the right to understand and challenge their risk scores. Notably, these risk scores can potentially marginalize certain demographic groups and make it more difficult to access government benefits for eligible individuals. Furthermore, without transparency, it is unclear how feedback loops for risk scores operate. For example, it is uncertain whether past scores or "failures" are incorporated into future scores. Thus, NIST should address these concerns to promote fairness and accountability in risk scoring systems.

4. Are current testing programs for liveness detection and presentation attack detection sufficient for evaluating the performance of implementations and technologies?

Yes, but testing programs should be refreshed on a periodic basis to account for new and emergent threats (*e.g.*, deep fakes).

5. What impacts would the proposed biometric performance requirements for identity proofing have on real-world implementations of biometric technologies?

The proposed biometric performance requirements would have significant positive impacts on the real-world implementation of biometric technologies. It would limit the number of biometric technologies used to **only those that perform above the bar set by NIST, ultimately lowering fraud and increasing equitable access.** This would also accelerate the pace of innovation and performance improvement as technologies that are "below the bar" would be forced to improve their performance or be shut out of government business opportunities. Furthermore, this would strengthen the need for technologies to participate in and excel during NIST FRVT, DHS biometrics rally, and other official testing programs.

Beyond performance requirements, NIST should require digital identity providers to offer alternative verification pathways, such as Supervised Remote or In-person Verification, if the primary verification pathway includes biometrics. This would promote user choice and fallbacks for users that are unwilling or unable to successfully complete biometric checks for any reason.

B. Risk Management

1. What additional guidance or direction can be provided to integrate digital identity risk with enterprise risk management?

Guidance developed by NIST should include standards that minimize the transmission of unnecessary data and reduce the excess data points that third parties can passively collect. These standards should reduce the risk of "dark data" or data that is collected and stored but not used for regular enterprise activities.

2. How might equity, privacy, and usability impacts be integrated into the assurance level selection process and digital identity risk management model?

As NIST begins to address risk scores directly, we wanted to share some perspectives on their use. Risk scores were originally developed by data brokers and credit bureaus to limit their own liability of potential false positives in the proofing process, reduce the transparency into how data brokers evaluate the risk of an individual, and put the burden on the Relying Parties to establish thresholds for making access decisions based on subjective "risk score" values.

The use of "risk scores" also limits the portability (and therefore, usability) of any credential. There is no way to compare a risk score from one CSP to another. In addition, one agency may consider a "risk score" of 650 sufficient to meet a certain standard (*e.g.*, IAL2) for a use while another agency considers the same score, from the same CSP, insufficient to meet the same standard.

To address the current lack of transparency around how a "risk score" is calculated, NIST could consider:

1. Requiring any scoring system be evaluated against metrics to prove there is no bias in the scoring algorithm based on protected characteristics such as race and gender.
2. Recommending transmission of explicit risk flags to RPs, along with or instead of veiled "risk scores." This will ensure RPs can make informed decisions on whether or not to provision access to the user based on objective facts rather than subjective opinions (styled as a "risk score").

3. How might risk analytics and fraud mitigation techniques be integrated into the selection of different identity assurance levels?

In all cases, the use of risk analytics such as "risk scores" should not be permitted to replace evidence and verification requirements within the IAL framework. Nor should risk scores ever be the only means of preventing fraud. Instead, risk analytics should be integrated only as an additional means of identifying and preventing fraud at various IALs.

a. How can we qualify or quantify their ability to mitigate overall identity risk?

Without a standardization of a uniform risk analytics scoring rubric and output, it is impossible to qualify or quantity risk analytics ability to mitigate risk because CSP1's score of 75, 575, or 975 will never be exactly equal to CSP2's score of 75, 575, or 975 (or CSP3's A+, B-, and F "scoring").

C. Authentication and Lifecycle Management

1. Are emerging authentication models and techniques – such as FIDO passkey, Verifiable Credentials, and mobile driver's licenses – sufficiently addressed and accommodated, as appropriate, by the guidelines?

At this point, yes. The main blocker for passkeys that we see (i.e. cloning of authenticators) is addressed sufficiently.

ID.me encourages NIST to consider the timing of the release of subsequent drafts of these guidelines to incorporate learnings from the mDL NCCOE launched in March.

a. What are the potential associated security, privacy, and usability benefits and risks?

There are some risks associated with the use of mDL's, which are flagged in our response to "Identity Proofing and Verification" Question 2.

2. Are the controls for phishing resistance as defined in the guidelines for AAL2 and AAL3 authentication clear and sufficient?

Yes. NIST could incorporate by reference the national Zero Trust and Cyber Security strategies for consistency as an improvement and method for ensuring guidance evolves in a uniform manner.

3. How are session management thresholds and reauthentication requirements implemented by agencies and organizations?

Generally, this is accomplished using a risk-based approach. Previous versions of the 800-63 had specific session lengths. NIST could consider adding guidelines that allow for less stringent authentication if it is from a trusted device. If it is the same user on the same device on a regular cadence, then could the lesson length be extended to improve UX, so long as access is on the same device.

a. Should NIST provide thresholds or leave session lengths to agencies based on applications, users, and mission needs?

It should do both, *i.e.*, include recommended or industry standard lengths, and also give agencies guidance on what to consider if they want to deviate from standard thresholds.

4. What impacts would the proposed biometric performance requirements for this volume have on real-world implementations of biometric technologies?

Repeat question. See answer to Question 5 in the Identity Proofing and Enrollment section.

D. <u>Federation and Assertions [Part C]</u>

1. What additional privacy considerations (e.g., revocation of consent, limitations of use) may be required to account for the use of identity and provisioning APIs that had not previously been discussed in the guidelines?

While existing guidelines have addressed some aspects of privacy, NIST should incorporate additional guidance to ensure comprehensive protection for users' privacy. Users should have the ability to grant or revoke consent for individual uses of their data, allowing them to maintain control over commercial use (i.e., revenue generating use) of their information. This ensures that

users are aware [and in control] of the precise ways their data is being collected, processed, and shared.

Furthermore, it is essential to scrutinize the practices of data brokers that aggregate and sell user information. Data brokers often operate with limited transparency, making it difficult for users to understand how their data is being used or shared. Regulatory oversight and stricter guidelines are needed to hold data brokers accountable for their practices. For example, the Consumer Financial Protection Bureau (CFPB) recently launched an investigation into data brokers.[6] The CFPB's stated focus is "to understand the full scope and breadth of data brokers and their business practices, their impact on the daily lives of consumers, and whether they are all playing by the same rules."

2. Is the updated text and introduction of "bound authenticators" sufficiently clear to allow for practical implementations of federation assurance level (FAL) 3 transactions?

Yes. The intent is clearly communicated. There could be issues that arise if different entities start inventing their own methods to facilitate this federation, and so any specificity or examples provided by NIST will help.

a. What complications or challenges are anticipated based on the updated guidance?

Complications or challenges may arise if various CSP develop different systems and/or methodologies. For example, CSPs creating their own unique systems and methodologies could lead to a lack of standardization, which could also result in varying levels of security and privacy measures, impacting equity and accessibility for users. In addition, future regulation and compliance requirements may become more complex if different CSPs implement their own systems and methodologies prior to the adoption of standardized protocols.

E. <u>General</u>

1. Is there an element of this guidance that you think is missing or could be expanded?

NIST should reintroduce a decision tree in the guidelines to aid agencies in selecting the appropriate level of assurance for identity proofing requirements. In 800-63-3, Figure 6-1, there was an IAL Decision Tree that combined the results from the risk assessment with additional considerations related to identity proofing services to allow agencies to select the most appropriate identity proofing requirements for their digital service offering. This was incredibly

[6]https://www.consumerfinance.gov/about-us/newsroom/cfpb-launches-inquiry-into-the-business-practices-of-data-brokers/

valuable for agencies to determine the appropriate level of assurance for various transactions and/or applications. A decision tree was removed for the draft of -4. ID.me highly encourages NIST to include the prior or an updated version of Figure 6-1 in the final publication of the -4 guidelines. Without a decision tree, there will be a lack of uniformity and confusion from agencies on what assurance level to select for various transactions thereby undercutting the uniform application of the -4 guidelines.

Additional Guardrails and Standardization

NIST should also provide more specificity, standardization, and examples to new concepts like "tailoring" and "credible sources." As written, we worry that it will weaken portability and federation of issued credentials. Specifically, if one agency disagrees with another agency's approach to tailoring or use of compensating controls, there will be a lack of trust created between those two agencies, effectively negating the point of having guidelines and standards. A lack of standardization will add complexity and erode trust, ultimately hurting the end-user. Additionally, NIST should provide guidance to CSPs on how tailored credentials should be labeled or categorized (i.e. IAL2- or IAL1+) such that relying parties have full transparency into which credentials strictly adhere to guidelines and which have been tailored. Also, if a user were to try to use the same credential at both agencies and be rejected at one, NIST should provide guidelines on who is responsible for supporting that end-user and help them understand why they can access one agency but not the other.

Similarly, the current definition of credible sources encourages broader buying and selling of user data, which runs counter to NIST's desire to promote privacy enhancing technology. For example, allowing use of data that can be traced back to an authoritative source enables CSPs to "move downstream" and complete verification using data that has been provided by an authoritative source to a credible source, presumably through a commercial transaction. This incentivizes CSPs to use data purchased from authoritative sources, rather than encouraging CSPs to move as far "upstream" as possible to the issuing source. As another example, in the first sub-bullet of the definition of credible source in 4.3.4.4, a credible source is one that has access to "information that was validated through an identity proofing process." Does any identity proofing process suffice, or is it one that needs to adhere to NIST guidelines?

This type of activity is why the Gramm-Leach-Bliley Act was established by the Federal Trade Commission (FTC) to safeguard consumers' non-public personal information. Our worry is that data brokers and other records-based solutions will be able to assert that attributes were validated through an identity proofing process, but NIST, Kantara, or any other accreditation body will not be able to verify if that is the case or understand the rigor and security of the identity proofing process for the credible source. This will erode trust in the long run.

- Recommendation: With respect to credible sources, ID.me recommends:
- Removing the ability to validate evidence against credible sources
    - If NIST believes it is critical, the -4 requirements should
        - Limit such use to IAL1 verifications because of the lack of oversight and transparency into traceability for these sources
        - Align use to the privacy guidelines and practices set forth in the Gramm-Leach-Bliley Act
- Adding back incentives for CSPs to validate information against the issuing source, wherever possible.
    - For example, only one piece of STRONG identity evidence was required to achieve IAL2 if the evidence could be validated against the issuing source. There may be a lack of commercially available services that can provide issuing source validation today, but that could change in the future and as such, NIST should continue to incentivize CSPs to validate against issuing sources if available for a specific type of identity evidence.

Mobile Driver's Licenses (mDL)

NIST should provide more specific guidelines for the use of mDL. Specifically, ID.me recommends NIST address the following items:

- **mDL's strength as a piece of evidence.** Based on expected mDL implementation and 63A requirements, it appears mDLs correlate to at least STRONG if not SUPERIOR. Also, the consideration of REAL ID based mDLs which the AAMVA guidance includes. ID.me recommends use as at least STRONG and in the case of RealID derived mDL's, SUPERIOR.
- **The volume and nature of attributes that must be retrieved for evidence validation.** mDL's contain additional information beyond what is contained on physical driver's license. For example, there are a variety of PII retrievable along with metadata denoting things like name truncation. There is also ancillary data such as Commercial Driver's License (CDL) information. ID.me recommends the minimum amount of PII, portrait image, plus license number since the mDL is already created based on issuance and confirmation of valid driver's license.
- **NIST should provide guidance on whether or not mDL use is limited to a piece of evidence or if they could ever be considered a standalone credential at the IAL2 level (and under what conditions).** ID.me's recommendation is to restrict use to a piece of evidence rather than a credential in and of itself. The reason for this is that doing so increases consumer choice, promote innovation, de-risks fraud, and enable accountability:
    - **Avoiding vendor lock-in.** Operating systems, browsers, and other applications that currently house mDL's natively are not incentivized to give users control over how and where they store their mDLs and the information contained within.

There is no common API infrastructure for users to move mDLs into subscriber accounts and cross-platform digital wallets developed by existing, accredited Credential Service Providers (CSPs). This means that the ability to an mDL is "locked-in" to a handful of large vendors that ship device operating systems. ID.me believes NIST should consider guidelines that promote a vendor-agnostic implementation, thereby putting end-users more in control of their mDLs. This more open market will force providers to continue to innovate to provide exceptional products and services to end-users and allow more space for emerging technologies.

  - **Understanding fraud prevention performance relative to IAL2.** Furthermore, evidence is already surfacing that mDL's in and of themselves may not have sufficient fraud controls to protect high-risk services: usage of mDLs to defraud banks. NIST should ensure it understands how mDL and IAL2 perform relative to each other as fraud controls before considering mDL as a standalone credential
  - **Holding issuing sources accountable.** Additionally, this would effectively make mDL issuing sources and the applications / operating systems that house them Credential Service Providers (CSPs) and should be subject to independent review and accreditation.

- **NIST should consider providing guidance on what technical features an mDL needs to have in order to be considered STRONG or SUPERIOR evidence**. Most issuers of mDL's have coalesced around ISO/IEC 18013-5 (Attended transaction) and the soon-to-be-published -7 (Unattended transactions). Additionally, AAMVA is moving forward as a guidance provider and certificate aggregator/disseminator for North America. ID.me recommends providing guidance on the features it would like to see factored into ISO/IEC and Aamva guidelines.

Trusted Referees

In the 63A, Section 5.1.9.1, NIST states that "The CSP **SHALL** train its trusted referees to make risk-based decisions that allow applicants to be successfully identity proofed based on their unique circumstances." **ID.me believes that retaining this language is the single most important NORMATIVE requirement NIST can include to advance equitable access for consumers.** NIST should include it and also provide further governance on what types of decisioning Trusted Referees can make. For example, would a Trusted Referee be able to verify a user to IAL2 based on a single STRONG piece of evidence and a robust personal interaction? Or is the risk-based decisions NIST reference more along the lines of whether or not a submitted piece of evidence should be accepted? **Additional guidance on what types of decisions NIST envisions Trusted Referees making would be helpful in managing programs that increase equitable access while balancing fraud risk.**

Applicant Reference

NIST should offer comprehensive guidance on the definition and use of an "applicant reference," including acceptable relationships with the applicant and safeguards to mitigate fraud concerns. Such guidance is crucial to ensure the appropriate use of "applicant references" while promoting access, equity, and security in various situations.

ID.me recommends that NIST consider adding the following requirements to help mitigate potential fraud concerns associated with the use of applicant references:

- Limit the number of times a subscriber can serve as an applicant reference
- Require CSPs to disclose whether a subscriber was verified via an applicant reference enabling RPs to make an informed decision on whether to permit the subscriber access to their resources. Some RPs may not be willing to accept the risk of provisioning access to these types of users.
- Require RPs to perform risk assessments on whether to accept applicant references and determine permissible applications. This should be incorporated into agencies existing Digital Identity Risk Assessment (DIRA) processes.
- Establish strict requirements for suitable and unsuitable applicant references. For example, NIST should consider adopting a set of requirements similar to the requirements the UK has established for passport application references (*available at* https://www.gov.uk/confirm-identity-online-for-passport-application).
- Provide specific guidance on circumstances in which the use of an applicant reference is appropriate (e.g., minors).
- Depending on the type of user, NIST should also consider additional requirements for the use of an applicant reference. For example, if the user is a minor:
    - Can anyone serve as an applicant reference or does it need to be a parent or guardian?
    - In addition to verifying the identity of the applicant reference, what are the acceptable ways for CSPs to verify the applicant reference is the parent or guardian of the minor?
        - There are many potential ways for CSPs to prove guardianship but NIST should provide examples of acceptable methods for doing so.

In addition, NIST should include additional guidance on applicant references by defining:

- **Who the "applicant reference" can be** (e.g. in the UK, an applicant reference must be from one of a list of "recognised professions")?
- **Requirements on how the applicant reference is associated with the applicant** (e.g., nature and duration of relationship, potential conflicts of interest, etc.).
- **Unique requirements for minors.** NIST should provide for the various unique aspects of an applicant reference for a minor applicant. For example,

    - If an applicant is a minor and the applicant reference does not need to be a parent or legal guardian, what level of relationship is sufficient (e.g. a sibling over the age of 18, a school guidance counselor, etc.)?
    - What should a CSP do when a verified user turns 18 (e.g., sever relationship with applicant reference, require a handover, etc.)?
- **What validation, if any, is needed for the information** for which the reference vouches?

2. Is any language in the guidance confusing or hard to understand? Should we add definitions or additional context to any language?

Yes. Additional guidance should be provided related to tailoring. Specifically, tailoring should be permitted to increase requirements of an Identity Assurance Level only (never to decrease requirements). It should not be a mechanism or loophole by which agencies relax their controls. Failure to define "tailoring" in this way, would result in entity specific credentials not aligned to any specific NIST standard and therefore, not portable. Thereby defeating the purpose of the standards. For example, if the IAL2 standard is "tailored" down by one agency using a control that was assessed and rejected by a second agency, the agency may not accept the credential as IAL2, limiting the credential's for a single agency only. Also, if a user were to try to use the same credential at both agencies and be rejected at one, who is responsible for supporting that end-user and help them understand why they can access one agency but not the other. In contrast, if an IAL1 credential is "tailored" up by one agency, the IAL1+ credential would be portable to other agencies as an IAL1.

3. Does the guidance sufficiently address privacy?

No, it is not strict enough. Users should have control over their own data, should be able to review at any time which orgs have access to their data, and should be able to delete subscriber accounts and biometrics at their discretion. Additionally, organizations and agencies should gain consent from the end user before collecting or transmitting data.

While the draft guidance provides a comprehensive framework for digital identity management, it does not sufficiently address privacy concerns from the perspective of individual users. A key aspect of privacy is empowering users to have control over their own data, which the current guidance does not fully address.

In order to enhance privacy protection, the NIST 800-63-4 guidance **should explicitly require organizations and agencies performing verification activities to obtain consent from users before collecting their information**. This would allow users to make informed decisions about sharing their personal data and foster trust in the digital identity ecosystem.

Additionally, users should be granted the ability to **review, at any time, the agencies to which users have consented to sharing their data and the data elements that have been shared**. This enables users to better understand what data has been shared, with whom, and when. This level of transparency promotes trust between users and the organizations handling their personal information.

Lastly, the guidance should promote the ability for users to delete specific pieces of data, such as selfies, from their digital identity records. This would give users greater control over their personal information and enhance their privacy protection.

While the NIST guidance offers a robust framework for digital identity management, it does not fully address privacy concerns. By requiring organizations to obtain user consent, granting users the ability to review data access, and allowing users to delete specific pieces of data, the guidance could better protect user privacy and promote trust in the digital identity ecosystem.

4. Does the guidance sufficiently address equity?

The NIST 800-63-4 guidance does not sufficiently address equity. Solutions in industry intend well and want to build platforms that help as many legitimate users as possible participate in the digital economy. However, there is a need to increase transparency and standardization in how we talk about equity.

At ID.me, we have been working with NIST and government agencies to do just that by offering users multiple pathways by which to verify their identity, should they not be able to successfully verify using algorithms alone.

ID.me is constantly looking for ways to improve our platform to (1) reach as much of the population as possible and (2) generate consistent outcomes regardless of the demographics of the end-user. Our work supporting users in Puerto Rico has demonstrated that our omnichannel solution increased pass rates for Puerto Ricans by more than 3x, when compared to solutions based on online records. As we continue to level the playing field for vulnerable demographics, we believe there are two areas where NIST can provide industry with more guidance that will advance further understanding and innovation:

- **Collection of demographic data at the user level.** In order to adhere to NIST guidance on data minimization, CSPs currently verify the identity of individuals without the collection of certain user demographic information such as skin tone, gender, age, etc. These attributes are not required to verify that someone is who they say they are. However, in the absence of collecting these attributes, it is difficult to assess performance of a CSP's platform at the individual user level. Techniques are available to do so either via sampling or developing proxy metrics; however, each of these techniques have

limitations (e.g. they are manual and can't be scaled or they characterize clusters of users rather than individual users).

- Recommendation: To address this, we recommend NIST provide guidance on how CSPs could collect limited sets of demographic data in order to better analyze and study their platform. For example, recommending CSPs offer users an **optional** survey at the end of verification to collect demographic data would (1) give CSPs flexibility they need to collect data necessary to study their platforms, (2) promote user control of their data and user choice by making it optional, and (3) provide industry with a standard / common way of collecting this information

- **Standardized metrics, relevant to the level of security imposed**. At present, there are no common metrics by which to compare performance across platforms, across security policies (e.g. IAL1, IAL2), and across user populations. "Pass rates" make sense conceptually, but there are a few challenges:
  - **They need to be looked at holistically across a user's digital life**. For example, portable and federated models of digital identity turn proofing into a 1-time event. Looking only at pass rates on initial verification overlooks the fact that verified users have a 100% pass rate at each subsequent agency where they assert their IAL2 (because they don't have to repeat verification).
  - **They need to be linked to the security policy**. At present, some players in industry conflate pass rates for lower-assurance activities like account opening, know-your-customer (KYC) and try to compare them to high-assurance activities like government benefits application at IAL2.
    - Recommendation: NIST should provide guidance that when discussing performance metrics, they need to be put in the context of the security policy in use.

- **Responsible use of biometrics.** ID.me believes that NIST and DHS have studied biometrics at length via NIST FRVT and DHS Biometrics Rally. In both of these rigorous government testing programs, the top performing algorithms perform consistently across demographics. We also see in our own data that there are commonly-used *non-biometric* checks that have more variation in pass rates across demographics than biometric controls do (see below). Therefore, we would like to see NIST advance responsible use of specific solutions that have been studied by the government, rather than move away from them. An example of how this could work would be consistent with the White House Office of Science and Technology Policy "Blueprint for an AI Bill of Rights", which includes "Human Alternatives, Consideration, and Fallback" as one of its 5 key principles.
  - Recommendation: NIST should continue to include the use of biometrics as an option for verification, and provide NORMATIVE guidance that biometrics used

are (1) tested by either NIST or DHS, (2) clear a NIST-determined performance threshold, and (3) are back-stopped by human reviewers or alternative verification methods involving humans

Pass rates for this step	Analyzed vulnerable population	Overall U.S.
Financial Records	56%	91%
1:1 Selfie Match	99%	99%

- **Mitigations in the Equity Section.** NIST should also provide examples of alternative processes to compensate for residual bias and technological limitations. For example, when you write "Supporting alternative processes to compensate for residual bias and technological limitations," please provide examples of what alternative processes you had in mind. Examples can be framed as "including but not limited to."

a. What equity assessment methods, impact evaluation models, or metrics could we reference to better support organizations in preventing or detecting disparate impacts that could arise as a result of identity verification technologies or processes?

Please see section in our Executive Summary related to possible methods for assessing performance across demographics without collecting any unnecessary user-level demographic information (e.g. American Communities Survey data, Harvard Geocoding Project, etc.)

5. What specific implementation guidance, reference architectures, metrics, or other supporting resources may enable more rapid adoption and implementation of this and future iterations of the Digital Identity Guidelines?

The most useful assets that NIST could provide to accelerate adoption (and make sure its is done consistent with NIST's intentions) include:

- A decision matrix for DIRA assessments
- A scoreboard / dashboard of standardized metrics that CSPs know they will be evaluated against
- References to the NCCOE findings to make sure CSPs have access to significant learnings from those efforts

6. What applied research and measurement efforts would provide the greatest impact on the identity market and advancement of these guidelines?

Please see section in our Executive Summary related to deployment of standard metrics by which to measure performance of CSPs

III. RECOMMENDED REDLINES

Below is a summary of the proposed redlines submitted separately using the NIST comment template.

A. 800-63-4 Base

5.2 Select Initial Assurance Levels be revised so that tailoring can be used to increase but not decrease IAL requirements.

5.3.2 Document Results be revised to require agencies SHALL include this information in the system authorization package described in [SP800-37].

B. 800-63-4A

Section 4.3.3.2, Page 11, Line 578, Bullet 6.

- *Add the following language*: "A grace period of up to 12 months from expiration is allowable, including for STRONG evidence, subject to risk determination by the CSP"
- *Rationale*: During the pandemic, some Americans who had expired identity documents were unable to get an updated identity document because issuance services were not available. ID.me proposes a grace period of 12 months beyond the expiration date of the identity document, including at a Strong level, as we assess the security risk to be low while the corresponding gains to access and equity to be high, particularly during a crisis.