Open Identity Exchange

NIST V4 Feedback

OIX Response

14th April 2023

Version 1.0

Author: Nick Mothershaw, Chief Identity Strategist
Email: nick.mothershaw@openidentityexchange.org

1 Feedback on Specific Requests for Comments

1. **Identity Proofing and Enrollment:**
    a. 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:
        i. 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?
        ii. Are these technologies supported by existing or emerging technical standards?
        iii. 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)?

The following table lists proofing validation and verification methods, along with known standards, and a comment as to whether a new technical standard is required for this method. OIX's work on interoperability between trust frameworks, as part of our Global Interoperability analysis is identifying that trust frameworks around the globe support the techniques listed below and would benefit from global standards for these techniques:

Proofing Process Element	**Technique**	**Sub Technique**	**Known Standards**	**Global Standard Required?**
Validation	Document Validation via Face to Face session with user	Visible Security Features by appropriately trained person	PRADO	
		haptic/tactile security features (if present)		
		UV / IR Features by appropriately trained person	PRADO	
		UV / IR Features by Machine	PRADO	
Validation	Document Validation by Video Session with User	Liveness Check	ISO/IEC 30107-3 – presentation Attack. ISO/IEC 19989-3:2020 – presentation attack	
		Visible Security Features checked by person, remotely		
		Visible Security Features by machine		
Validation		Liveness Check	ISO/IEC 30107-3 – presentation Attack.	

	Document Validation via Static Picture		ISO/IEC 19989-3:2020 – presentation attack	
		Visible Security Features checked by person, remotely		
		Visible Security Features by machine		
Validation and Verification	Cryptographic Validation	via logon to account e.g. PSD2/SCA logon		
Verification	Image against Photo from ID Document	Face to Face		
		By Person Remotely		
		Biometric by Machine	ISO/IEC TR 29156:2015 ISO/IEC 19795-1:2021 (reference in EU stds). NIST FRVT	

To enable trust frameworks around the globe to express their policies in a way that can be shared with other trust frameworks, OIX is creating an Open Policy Rules Exchange Framework (**OPREF**) through its Global Interoperability working group. The OPREF will allow collaborating frameworks to assess alignment and agree interoperability. OIX has already assessed the NIST V4 requirements as part of this analysis and can share how NIST aligns with UK and EU trust frameworks from our progress so far. We are now moving on to analyse trust frameworks from Singapore, Bank ID Sweden, Thailand and Canada.

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

In order to accept a credential, be it in mDL standard, or as a verifiable credential, the proofing methods applied to connect the credential to the end user need to be understood, such as the methods listed in the table above. The credential needs to carry the proofing techniques used as policy attributes so that the value of the credential can be determined as part of an Assurance Policy (e.g. NIST SP 800-63A, UK GPG45). The OIX OPREF will enable the exchange of common proofing techniques.

OIDC for ID Assurance can then be used to communicate which verification and validation methods have been used using the 'check_method' tag. OIX is calling for the 'verification' object in OIDC for ID Assurance to become a 'protocol independence evidence standard' that could be carried in a OIDC ID Token, a Verifiable Credential or an MdL standard credential.

c. 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?

i. 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? If so, at what assurance levels could these be applied?
ii. How might emerging methods such as fraud analytics and risk scoring be further researched, standardized, measured, and integrated into the guidance in the future?
iii. What accompanying privacy and equity considerations should be addressed alongside these methods?

Without robust fraud controls to defend against identity theft and synthetic identity Digital IDs will be vulnerable to fraud and will suffer a poor reputation with users and relying parties as a result.

OIX has been running a working group on Fraud Control and Signals for Digital ID from the last 3 years. The working group has created two Guides:

- Guide to Fraud Controls
- Guide to Shared Signals

These contain comprehensive recommendations on which fraud controls to implement and what signals should be generated and distributed amongst the parties, principally ID Providers, in the ID Ecosystem. The basis of these guides is the fraud controls and signals that are already successfully implemented to protect again ID fraud in the UK financial services industry.

Fraud control types that should be considered for baseline normative requirements are as follows. More details on specific controls and OIX's recommendations as to which are implemented as a minimum by a trust framework can be found in the Guide to Fraud Controls.

Control Type	**Description**
Known Fraud	Checks to establish any known links fraud pre/post registration or in new evidence being provided in the form of trusted credentials. Should include a check on if the individual has been a victim of fraud.
Device Risk	Can the device be linked to other registrations, or to known fraud/suspicious device characteristic?
Anomaly	Are the discrepancies or patterns in key information (inc. liveness test of biometric image capture).
Velocity	Is there repeated use of key information
Evidence Check Failures	User fails an evidence check that they should pass. This could be indicative of a fraud attempt. Especially repeated failure. Mitigating action is required.
Behavioural discrepancies and Risks	Behaviour of individual not 'realistic', not 'normal' for that users or linked to other registrations/known fraud
Risk Indicators	Checking of external data sets to highlight any areas of increased fraud risk – e.g. mortality, redirection, email address age etc.

Privacy is a key consideration when implementing fraud controls. In particular when sharing information on risk or suspected frauds between parties, such as ID Providers, in the ecosystem. The Guide to Fraud controls considers several different approached to information sharing:

Collaborative Approach to Fraud Prevention	Description
1. Informal Intelligence Sharing	Sharing non-individual-user-specific fraud intelligence between ecosystem participants through round robin or centrally managed secure email and regular meetings to discuss new attack vectors.
Sharing Actual PII Data for Fraud Detection and Prevention:	
2. Identified Fraud Outcome Sharing	Sharing information about actual frauds that have been identified, for example the PII used, devices. (e.g. CIFAS in the UK, SAFPS in South Africa)
3. Shared Fraud Indicators	Sharing fraud indicators between parties as that may indicate a fraud risk (e.g., inconsistent data provided as part of ID creation attempt, repeated attempts to create an ID with similar information).
4. Centralised Fraud Analysis	Pooling all information used as part of the any registration or account update across all participants in the eco-system. This allows anomalies and patterns to be found across the participants that may not be found with an individual participants data alone (e.g. the fraudster attacking multiple participants with the same device or PII element such as email or address). There are two common approaches to centralised fraud analysis databases: 4a) data is pooled and then automated risk referrals are sent back to the contributor for analysis (e.g. in the UK National Hunter, National SIRA, National Fraud Initiative) and 4b) data is pooled and analysis is undertaken centrally with centrally selected referrals being sent back to the contributors. (e.g. the UK's National Fraud Intelligence Database).

Consideration needs to be given to how options 2, 3 and 4 are implemented to best effect.

Central databases are seen as a security risk and also go against the tenet of distributed identity solutions.

Options 2 - Identified Fraud Outcome Sharing - and 3 - Shared Fraud Indicators - can be combined into a "Shared Signals" implementation that share minimal information thus being seen as more privacy respecting. OIX has produced a supporting Guide to Shared Signals that explores this information sharing approach in more detail. To be effective, shared signal implementations need to persist some types of signal (e.g. an Identified Fraud Outcome) so that new ID applications can be checked against them. Depending on implementation choice this may result in a partial database of centralized information to defend against fraud.

However, it can be argued that option 4, Centralised Fraud Analysis is more effective in finding fraud than simply sharing fraud indicators per option 3. With a view of data from

across the entire ecosystem fraud can be found that would not be visible through shared fraud indicators alone. This advantage needs to be offset against drivers not to create centralised databases of ID information. Where Centralised Fraud Analysis capabilities already exist in a market, consideration should be given to leveraging or collaborating with the existing implementations, otherwise the fraud defence capability of a Digital ID ecosystem might be inferior to fraud defences deployed by relying parties today, which is unlikely to be acceptable to those relying parties.

In practise, trust frameworks might implement combinations of:

- Options 1, 2 and 4a
- Options 1, 2, 3 with options 2 and 3 delivered through shared signals.
- Options 1, 2, 3 and 4b with options 2 and 3 delivered through shared signals.

The user of machine learning / AI models for fraud detection needs to be considered carefully. Most implementations of fraud controls are rules based, with the output of the rules being prioritised using machine learning model. However, Fraud risk scores are increasingly created by the use of machine learning algorithms. The ethics of using machine learning to detect fraud need to be considered carefully, in particular if the score is used to make a decision that will affect the user's ability to complete their transaction. In this case false positives may lead to customer detriment. It is therefore good practise to refer any suspected frauds to a fraud operator for the final decision on whether a matching learning generated fraud alert is genuine or is a false positive.

Equity: The fraud controls proposed in the OIX guides have been assessed for any impact they may have on equity as part of the OIX working groups dialogue with the UK government DSIT team creating the UK Digital Identity and Attributes Trust Framework. Our analysis found that, as the controls run in the background and do not insist on any data items being mandatory, they have no impact on equity.

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

OIX cannot comment on these.

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

The FMR and FNMR rates proposed is V4 look aggressive and will challenge biometrics providers models, resulting in increased true non match rates leading to user frustration. However, we agree these objectives should be aggressive to protect against fraud. Given biometrics is never allowed as a sole authentication factor, the risk of false match allowing access to the user's data to the wrong person is mitigated by the second factor.

2. Risk Management:

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

There are several key considerations in the adoption of a Digital Identity solution that are not covered in these guidelines. Some may be considered as part of risk management, whilst other are more commercial considerations:

Consideration	**Items to Consider**	**Impact on the Relying Party of not addressed**	**Risk / Commercial Consideration**
Availability	When is the Digital ID ecosystem available? Is it 24/7? Is some down time planned? What evidence does the ID Provider have to show they have a robust scalable architecture?	Users cannot access services.	Risk
Liability	What happens if the system is unavailable? What happens if an ID fraudster has control of an ID? What happens if the ID Provider makes an error? Will any losses be recoverable?	Users cannot access services or suffer losses Relying party suffers losses.	Commercial.
Technical Interoperability	When accessing a federation of ID providers, is information delivered from each ID Provider consistent? Is the federation provider using standards or are they veering to proprietary overlays>	Must code to cater for ID Provider Inconsistencies Vendor lock in	Commercial
Fraud	Whilst fraud is addressed elsewhere in the guidance, the risk of fraudsters having control of a Digital ID needs to be considered as part of adoption. Digital ID implementation are unlikely to be 100% free of fraudsters, so it is important to understand how fraud is controlled and managed.	Loss due to fraudster access of services User loss.	Risk
Support	Users and the Relying Party must have somewhere to go when there is a problem.	Users unable to access services	Risk

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

These are baseline requirements for all ID Providers. They should be consistently applied by all ID Providers regardless of the assurance level being afforded to the identity and should be part of a set of General Policy requirements for ID Providers. This is the approach taken by

most trust frameworks, and an approach for which the exchange of policy between frameworks will be enabled by the OIX Open Policy Rules Exchange Framework (OPREF).

c. How might risk analytics and fraud mitigation techniques be integrated into the selection of different identity assurance levels? How can we qualify or quantify their ability to mitigate overall identity risk?

One approach is to implement the highest level of fraud controls required regardless of the use case, ensuring that to whole ecosystem is as robust as possible from a fraud controls point of view. In this instance a consistent set of fraud controls become part of the bassline policy requirements that all ID Providers must implement.

However, this "highest bar" approach may mean that if an ID is created for a lower risk use case in the first instance the cost of applying fraud controls may be prohibitive based on the price that will be paid for the ID for that use case.

Thus, in the same way different levels of ID proofing are applied to different use cases, difference Levels of Fraud Control (LoFC) might be applied based of the use case:

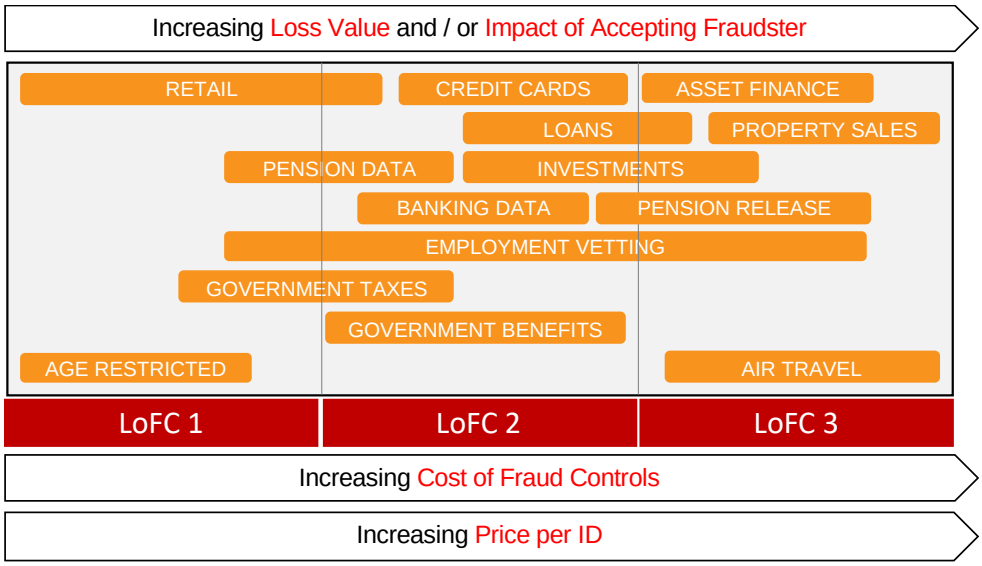

If an ID is initially used for a lower LoFC use case and this then used for a higher LoFC user case, additional "step up" fraud controls would be applied at that point of use. A problem with is approach is that an ID that has already been used for many lower risk user cases may be found to be fraudulent on applying fraud control step up, requiring mitigation action to be taken for all previous transactions made using that ID.

3. **Authentication & Lifecycle Management:**

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

      OIX cannot comment as it is not expert in authenticators.

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

      Yes, these seem clear. Although OIX is not expert in authenticators.

   c. How are session management thresholds and reauthentication requirements implemented by agencies and organizations? Should NIST provide thresholds or leave session lengths to agencies based on applications, users, and mission needs?

      OIX cannot comment on this area.

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

      The FMR and FNMR rates proposed is V4 look aggressive and will challenge biometrics providers models, resulting in increased true non match rates leading to user frustration. However, we agree these objectives should be aggressive to protect against fraud. Given biometrics is never allowed as a sole authentication factor, the risk of false match allowing access to the user's data to the wrong person is mitigated by the second factor.

4. **Federation & Assertions:**

a. 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?

Additional considerations would include:

- Consent History. It is implied in the NIST requirements that this is recorded. However, it should be more explicit.
- Right to be Forgotten. It is clear data should be removed from relying parties if an account is closed, however there is no provision for users to choose to revoke data shared from a specific replying party.
- Relying Party use of data. There are no provisions to require the relying party to register it's intended use(s) of data, or an approval mechanism for these. The EU proposes to operate an approval mechanism for relying parties who use sensitive personal data. Also, provisions to limit the monetisation or profiling of data without the user's explicit consent to the replying party should be considered.

b. Is the updated text and introduction of "bound authenticators" sufficiently clear to allow for practical implementations of federation assurance level (FAL) 3 transactions? What complications or challenges are anticipated based on the updated guidance?

IdP Change: Whilst is it clear that a RP may allow a subscriber to connect multiple IdP accounts to their RP subscriber account, more detail on how this works should be included. For example, can a second account only be added when the user is in an authenticated state with the RP. If a subscriber account at the IdP is deleted, then should users RP subscriber account should only be deleted if it is orphaned? How is this communicated to the user? Should the user be able to move the control of the RP Subscriber Account from one IdP to another? I might choose to move my ID from IdP A to IdP B, when I do so I should be able to move control of any RP Subscriber Accounts I choose form IdP A to IdP B, provided IdP B is acceptable to the RP. This is an important consideration to avoid subscriber lock in to IdPs.

5. General:

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

From our initial analysis of the NIST V4 requirements, many policy areas addressed by a typical trust framework are covered. Although several areas are quite light touch. Areas to consider for more detailed policy requirements are:

- Fraud Control and Management (as discussed above)
- Trustmark
- User Agreement for sharing data
- Relying party obligations / restrictions
- Presentation Protocol Support
- Data portability, tell us once and right to be forgotten
- Quality Management
- Incident Management and responding to data breach
- Routes to join and routes to build up your identity
- An extension of risk / commercial considerations to include: Availability, Technical Interoperability, Fraud, Liability and Support
- Application of risk-based decision – who takes this decision? The relying party? Federation Schemes? Presumably not IdPs.
- What happens if something goes wrong?

OIX is happy to help the NIST team understand how other frameworks are addressing some of these areas.

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

The term credential service provider is confusing as the process this role undertakes is identity proofing. This role could be renamed Identity Proofing Provider per the OIX Guide to Trust Frameworks definition.

c. Does the guidance sufficiently address privacy?

The guidance adopts many elements that we would typically see addressed in a trust framework or addressed by data protection legislation generally. There are several areas we might have expected to see that ate not included:

- Consent History. It is implied that this is recorded. However, it might be more explicit.
- Right to be Forgotten. It is clear data should be removed from relying parties if an account is closed, however there is no provision for users to revoke data shared from a specific replying party.
- Relying Party use of data. There are no provisions to require the relying party to register it's intended use(s) of data, or an approval mechanism for these. The EU proposes to operate an approval mechanism for the use of sensitive data. Also, provisions to limit the monetisation or profiling of data without the users explicit consent to the replying party might be considered.

d. Does the guidance sufficiently address equity?

   i. 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?

There does not seem to be a clear statement or objective that all CSPs should be equally equitable. The draft guidance requires CSPs: “In assessing equity risks, a CSP starts by considering overall user population served by its identity proofing and enrolment service“.

If I am a CSP who specialises in passport scans and selfie cross match, my target user population is those who hold passports. This is not an equitable service for those who do not hold passports. Is that acceptable under these guidelines?

Whilst the guidance requires the CSP to undertake a risk assessment, it does not require them to act to mitigate the risks they identify. A consideration is whether to require a CSP to document a mitigation plan and then show progress in following this.

Another approach is to recognise that not all CSPs (and therefore IdPs) will offer a completely equitable service and therefore to place an obligation on federations to ensure they address equity as far as possible by requiring the operator of the federation to:

- Select a range of IdPs with diverse services to give equity coverage. Including specialists IdPs who support Trusted Referee services.
- Publish users guidance on which IdPs can best support them, based on what documents and data the users might possess.
- Ensure the federation has a plan for improving inclusion.

The use of Trusted Referee services is also cited in the guidelines. OIX has explored how this might work when the Trusted Referee has a Digital ID themselves, resulting in a more robust and traceable process. See the OIX paper on Digital Vouch with Photo

e. 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?

More detailed implementation guidance should be considered on:

- Technical Interoperability, to ensure IdPs deliver information to RPs in a consistent format.
- Fraud Controls, to ensure CSPs are all consistent in their fraud defences. Fraudster will quickly identity and exploit any inconsistencies they find in the ecosystem, targeting weak CSPs as an entry point.

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

OIX has produced a market segmentation, in conjunction with its member Experian, that identifies the demographic profiles of the “ID Challenged” – those users with no passport, driving license and a thin credit bureau file. In the UK 12% of the adult population is ID Challenged. This can be used to model the user population of applicants for a CSPs services to validate whether a CSP is being equitable and help produce a plan to improve equity. A similar ID Challenged set of profiles could be created for the US. Detail of this segmentation can be found in: OIX ID Inclusion & Data Sets Project Report.

2 OIX Reference Documentation

Throughout this feedback OIX has referenced several of its own publications that the NIST team might find useful. These guides and reports are produced by OIX as free and open resources to help those seeking to create and understand trust frameworks:

OIX Guide to Trust Frameworks
Guide to Fraud Controls
Guide to Shared Signals
OIX ID Inclusion & Data Sets Project Report.
Digital Vouch with Photo