NIST 800-63-4 Comment

Name of submitter: Gabriel Zucker

Contact email: [redacted]

This is a high-level conceptual comment, rather than a proposed line edit. The comment pertains to the entirety of 800-63A, and most specifically to Section 3.1.2.

Broadly, it seems that an identity assurance system could be plausibly approached in two distinct ways: as a discrete categorical system, or as a probabilistic/continuous system. That is, guidance could designate a small and well-defined set of assurance levels; or it could define a system to calculate a probabilistic and continuous risk score, without specific proposed thresholds, letting system owners choose the appropriate threshold for their systems.

Both approaches have advantages and disadvantages. The endless variety in different use cases and risk profiles means that the probabilistic approach makes it more plausible that the assurance system can be precisely well-targeted to the use case, as system owners can select the threshold that best suits their system. Moreover, the probabilistic approach is computationally more flexible. Since the output is a continuous number rather than a binary determination, it can easily incorporate arbitrarily many pieces of data with varied weights, including those that indicate potential risk without definitively demonstrating it.

On the other hand, the continuous approach takes much power out of the hands of the guidance itself and puts it in the hands of system owners, who may prefer to rely on external and objective standards, rather than stand by a specific threshold point set wholly by the program. The continuous approach also hampers cross-program reuse. In a categorical system, an applicant who has proofed at a given level can immediately use all systems requiring that level of proofing; in a continuous system, they may not be able to do so.

Both approaches are, in this sense, defensible — and in the government context, the categorical one may have significant advantages, despite its lower flexibility. (It is also possible different use cases are better suited to different paradigms. A small number of highly-resourced and very widely-used systems might prefer the customization of the continuous approach, while the majority of smaller systems conform to one of the reusable categorical levels.)

In 800-63-4, NIST continues to broadly rely on the categorical approach, defining three specific proofing levels on the basis of precise rules rather than a probabilistic risk score. This comment does not object to this choice, but rather to the implications of Section 3.1.2 for that choice: Section 3.1.2 appears to retrofit a continuous/probabilistic approach onto the categorical approach of the rest of the volume without accounting for or providing clear guidance on how to accommodate this inconsistency.

Section 3.1.2 defines a fraud risk management process that CSPs should implement. The indicators provided in 3.1.2.1, especially (5) and (6), are all sensical indicators, which do indeed indicate heightened risk that an applicant is not who they say they are. But, despite the implications of the draft, these risks do not exist on an axis orthogonal to the primary categorical standards. Rather, they reflect the same risk a system owner is already addressing when they select the IAL level and enforce the associated controls.

The interaction of these signals becomes more obvious with a couple of examples:

- Suppose a system is designated IAL3, and the applicant passes proofing as established in 4.3. Suppose then that the applicant is flagged as risky according to one of several account tenure checks. Surely, in this case, the far stronger evidence provided in the IAL3 process outweighs the weak signal of a single account tenure check.
- Suppose conversely that a system is designated as IAL1, and the applicant passes proofing as established in 4.1, but is likewise flagged as risky according to one of several account tenure checks. In this case, the tenure check may indeed provide an additional signal of risk not ameliorated by the IAL1 proofing process. But, on the other hand, this system has been intentionally set at IAL1, which is intended to "support the real-world existence of the claimed identity" (800-63 3.3.2.1), and not more than that. At the point that the IAL1 applicant must also clear fraud checks, as well as the checks inherent in IAL1, this would appear to inherently raise the identity standard to which the applicant is being held.
- Suppose, meanwhile, that an applicant successfully proofs according to the standards of IAL2, but fraud considerations under 3.1.2 identify the claimant as plausibly fraudulent. Is the applicant considered to be proofed at IAL2 for reuse under a different program using the same CSP? They have, after all, met all requirements for such proofing.

The fraud determinations indeed exist along the same axis as, and thus must cogently interact with, the categorical assessments that make up the bulk of the 800-63A guidance. But the current draft does not provide clear guidance about the interaction of the categorical and the risk-based determinations. It merely requires the imposition of risk-based calculations on top of the categorical system, without clarifying the implications for categorical determinations or suggesting what to do in the case of conflicting signals.

Of course, categorical and continuous approaches are highly distinct, and it is not clear that there *is* a simple and sensical way to combine them. Moreover, it is understandable that the draft wants to highlight a set of fraud checks that are wise and prudent on their own terms. But, ultimately, the draft is insufficiently clear on how to handle these two inconsistent paradigms.

There is no obvious solution. At a high level, the guidance could: (1) explicitly note that the risk-based approach in 3.1.2 represents a different paradigm that CSPs may choose to use *instead* of the overall categorical approach, (2) it could provide concrete guidance on how to interact the two approaches, (3) it could drop the risk-based/continuous indicators entirely, leaving only the categorical approach that represents the bulk of the draft.