Economic Security Project

MAY 2026

The Future of Tax Filing

Part 2: Building Direct File: Policy and Strategy

Chapter 5: Direct File and identity verification

Gabriel Zucker, Chris Given

Contents

Summary

- Direct File was classified as an IAL2 system, which meant that taxpayers had to undergo rigorous identity verification procedures to use it. Identity verification was a barrier in 2024 and 2025.

- Direct File's classification was, on its face, an accurate implementation of NIST 800-63-3, the government-wide guidance on identity verification at the time. The fact that it was the correct implementation suggests issues with the NIST 800-63-3 guidance in practice (5.1). The Direct File classification also introduced de facto inconsistencies between IRS's treatment of public and private tax filing software (5.2).
- Because of its IAL2 classification, Direct File had to use the private service ID.me to identity-proof taxpayers, as ID.me was the only service then able to implement IAL2 at scale. Details of ID.me's implementation introduced further barriers on top of the inherent barriers of IAL2, which might have been somewhat ameliorated by the use of Login.gov, if that service had been ready at the time. (5.3)
- Next steps for government digital identity in general:
    - NIST 800-63-4, the newest revision to the government-wide guidance, was released in 2025. It makes improvements to NIST 800-63-3 on some fronts, though does not fully resolve the issues. (5.4.1)
    - Culturally, the relationship between NIST and other federal agencies lends itself to buck-passing of risk assessments, which tends to lead information systems to be overly restricted. This culture must change. (5.4.2)
    - Login.gov is a vital service, which needs to provide robust offline fallbacks for users who fail online identity verification, and must continually improve as it makes verification as easy and equitable as possible. (5.4.3)
- Next steps for Direct File in particular:
    - Direct File should consider reclassifying portions of functionality at IAL1, in light of NIST 800-63-4 (5.5.1), and creating alternative non-logged-in service patterns. That said, the growth of pre-population functionality makes this choice less clear cut than it otherwise might be (5.5.2).
    - Direct File should adopt Login.gov as an option for identity verification. (5.5.4)
- This entire section is a point-in-time snapshot of the identity verification space; generative artificial intelligence is very likely to vastly disrupt the entire ecosystem. This may well make many of our specific comments here irrelevant. (5.6)

Direct File required taxpayers to clear a high standard of identity verification, which was a meaningful barrier to usage (the Direct File team wrote bluntly in 2025 that "identity verification deterred taxpayers") and a point of some controversy. This chapter explores why Direct File was classified the way it was and what steps could be taken to ameliorate the impact of the barrier in the future.

But while there are some steps that can and should be taken within the confines of Direct File, **the problems with identity verification and digital government services extend far beyond Direct File or the tax system**. As we will lay out below, **many of these solutions must occur at a government-wide level**, and as such are beyond the scope of tax access work in general, or this report in particular.

5.1 NIST 800-63-3

Online authentication protocols throughout the federal government are governed by NIST Publication 800-63 — as of Direct File's birth, the then-seven-year-old revision 3, or 800-63-3. Publication 800-63 lays out three axes along which an online interaction may be authenticated: (1) identification assurance level (IAL; is this user the person they say they are?), (2) authentication assurance level (AAL; is the person who is currently using the account the same person who created the account?), and (3) federation assurance level (FAL; a dimension that comes into play when verification is mediated through a third party). For each of these, 800-63-3 defines three levels of assurance — level 1 (low/no assurance), level 2 (medium assurance), and level 3 (high assurance). We are interested here in IAL, verifying whether a person is who they say they are.

When it comes to IAL and Direct File, NIST 800-63-3 has two notable issues:

- **IAL1 and IAL2 are in practice quite far apart.** IAL1 essentially constitutes no identity assurance whatsoever. Donald Duck can create an account using Donald Duck's address and the email donald.duck@gmail.com and assert he is Mickey Mouse — and this is perfectly fine under IAL1. IAL2, meanwhile, generally[1] requires showing a government-issued photo identification in such a way that the Identity Provider (IdP, the service that is performing IAL verification) can confirm that it matches the face of the person holding it, which can be challenging in a digital context. Even if Donald Duck is using an address and email address well-associated with himself from prior interactions, even if Donald Duck is providing tax information about himself that only Donald Duck has, and even if the risk of fraud is vanishingly low — none of this sufficiently proves Donald Duck is Donald Duck under IAL2, unless he can also show an ID. IAL3, meanwhile, is for in-person verification routes — ultra-sensitive use cases, in which the identity verification cannot be mediated online at all. In practice, the ultra-low bar of IAL1 and the ultra-high bar of IAL3 tend to mean any online service with any risk whatsoever was classified as IAL2 under 800-63-3.
- The authors of the NIST guidance might respond that the above is a fundamental misreading. Indeed, **NIST outlines a concept called "compensating controls," which essentially give agencies latitude to amend the margins of IALs according to the specifics of the use case. But this is incredibly challenging in practice**. For one thing, individual information security officers at agencies are understandably loath to be seen as playing around with government-wide NIST guidance, formulating an agency posture that is below the government-wide standard. For another, the categorical approach of IAL1, 2, and 3 is fundamentally inconsistent with the probabilistic, risk-based approach implied by the "compensating controls." In practice, IdPs proof users to specific levels of the categorical approach, and these proofings are re-used across applications. The idea that separate applications might have bespoke proofing levels (according to compensating controls) has its own

[1] This is a simplification; IAL2 allows other methods of authentication. But in practice this is usually the method, and the other methods share similar levels of rigor.

advantages, but it upends the entire logic of the categorical system that NIST itself created. If agencies relied heavily on compensating controls, then a user who proofed their identity at IAL2 for one use case might still have to start from square one for another use case, and it would be essentially impossible for IdPs to say, for example, that they "support IAL2." Which IAL2, with which controls?

5.2 Direct File's IAL classification in context

The IRS determined in 2023 that Direct File should be classified as IAL2 under NIST 800-63-3. (This meant it would use ID.me, for reasons discussed in the next section.)

The rationale did not have to do with confirming the return was being filed by the person they said they were; indeed it had nothing to do with all of the functionality up through clicking submit. Rather, it had to do with what happens after submission. After a taxpayer submits a return to MeF, their return is either accepted or rejected according to one of thousands of business rules, which may validate the submitted return against other IRS data about the filer. For example, MeF may reject a return because the taxpayer already filed a return or because records show the taxpayer had received Advanced Premium Tax Credit payments (which were not adequately represented on the return). Through these reject codes, Donald Duck may file a return pretending to be Mickey Mouse, and thereby receive back information from MeF about whether Mickey has filed, or whether Mickey received the Advanced Premium Tax Credit. This would constitute, according to the IRS, an unauthorized disclosure of Mickey's federal tax information to Donald Duck. And so, in order to access these MeF reject codes, the taxpayer must prove they are indeed Mickey Mouse.

Now, given that the sensitivity of data being released via MeF reject codes is relatively low, and given the amount of information a person provides via their tax return is relatively high, a sensible implementation might use other information from the tax return to establish a sufficient level of confidence in the user's identity. The IRS in fact requires prior-year adjusted gross income (AGI) or self-select PIN as an authentication attribute on the return. The IRS could assess that a return with the correct AGI or PIN has a low probability of being fraudulent, and a reject code, a low-risk piece of data, could be issued in response.[2]

But, it is simply not how the logic of IAL works. IAL does not live in a world of risk assessments and probabilities; it lives in a world of categories. **Since any information was being released at all via MeF reject codes, it stood to reason under 800-63-3 there ought to be *some* identity control, and so the only sensible category for Direct File was IAL2.**

Critics noted then and would note now that, though this logic may be a valid interpretation of NIST 800-63-3, it is inconsistent with the IRS's actual policy toward other software providers. Taxpayers using private tax software like TurboTax or TaxSlayer can use the

[2] AGI and self-select PIN have their own challenges as authentication measures — but, if they were to be removed, there are others. For example, perhaps if a taxpayer is filing using the same contact information as in a number of past returns, or if the income data a taxpayer provides matches IRS records, the IRS can be confident enough to return MeF reject codes.

products and receive their MeF reject codes without going through IAL2 verification. Nominally, in classifying Direct File, the IRS drew a distinction: the software companies themselves were the ones accessing MeF, and officers of those companies were proofed at IAL2; what the companies did with each taxpayer's data was the companies' problem. With Direct File, on the other hand, the government was giving individuals access to MeF, and so those individuals, too, needed to be proofed at IAL2. But this is obviously a distinction in search of a difference, at best. In practice, if Donald Duck wants to know whether Mickey Mouse filed, and Donald has enough information to spoof Mickey's return, Donald can file that return through TurboTax and receive the information he wants back from MeF. But the mediation of data access via third parties made it easier to obfuscate the nature of the implicit risk assessment.

Keep in mind that the IAL2 classification was premised solely on the sensitivity of these MeF reject codes. Once Direct File was implemented with all taxpayers proofed at IAL2, however, it became trivial to import additional data from the online account, like prior-year income, or IP PINs. This information is also protected at IAL2, which is as it should be; the data in the online account is much more sensitive than the MeF reject codes are, and it is by definition provided prior to filing, making it impossible to use the contents of the return to assess fraud risk. But the decision to make Direct File IAL2 predated any notion of using the online account data.

As noted, the IAL2 classification was not premised on preventing refund fraud. An IAL1 product may have needed to introduce some measures to detect and prevent such fraud. Implemented at IAL2, though, the product inherently had very strong fraud protections, and Direct File consistently saw a far lower incidence of fraudulent use than any other DIY filing option (although this is partially a function of bad actors tending to prefer the path of least resistance).

5.3 ID.me, Login.gov, and other IdPs

To implement Direct File at IAL2, the IRS needed an identity provider (IdP) to proof users' identities at the required level.

In late 2023, there was essentially one provider operating at scale and able to provide IAL2 proofing for government services: ID.me. Moreover, this is the provider the IRS was already using for its online account.

IRS's use of ID.me, then and now, has issues. GAO in June 2025 published a report highlighting that IRS's oversight of ID.me was limited, to the point that IRS did not independently establish goals for pass rates and data privacy.

Meanwhile, Login.gov is a public IdP which stood to have a few potential advantages over ID.me: (1) Login's user interface is arguably more accessible, (2) Login provides more in-person fallback proofing mechanisms, including, for example, at post offices, (3) the option of Login means users are not forced to trust a third party with their personal

information, and (4) Login is subject to more government oversight and has greater mission alignment, as a public institution. This last point is probably the most important; the fact that the government owns Login.gov means the government has more power to understand the barriers it creates, iterate upon it, and improve it. That said, the differences between Login and ID.me should not be overstated: using Login.gov IAL2 would still leave verification for Direct File far more onerous than for other tax software, which generally does not impose IAL2 verification at all.

As of 2023, at any rate, Login.gov did not support IAL2. They did by October 2024, and Congressional champions pushed Direct File to switch. But the switch could not happen immediately because of additional technical and security work required by the IRS and GSA teams. However, internal roadmaps expected that Login would be available as an option to authenticate for Direct File by filing season 2026.

5.4 Next steps for government digital identity in general

5.4.1 NIST 800-63-4

As discussed above, a big problem here is the NIST standards.

NIST had been working on NIST 800-63-4, a revision to 800-63, since roughly 2020. After two drafts released for public comment, **the final revision was finally released on August 1, 2025**.

The latest draft makes some improvements from version 3, with regards to the issues discussed above:

- **The draft partially fixes the problem of the gulf between IAL1 and IAL2.** IAL1 is redefined as a new level of assurance higher than the 800-63-3 IAL1 but lower than the 800-63-3 IAL2 — an IAL1.5, as it were. (The 800-63-3 IAL1 is essentially renamed in 800-63-4 as a "no proofing" level, or colloquially, IAL0.) In this world, products are not forced to pick between two highly disjunct and far apart options. They can pick (the new) IAL1. (As we discuss below, this could have been an option for Direct File.)
- **The guidance doesn't, though, make meaningful progress on the conceptual inconsistency of a discrete authentication paradigm (i.e., IAL 1, 2, 3) with a probabilistic risk-based paradigm (e.g., compensating controls) layered on top of it.** In fact, to a degree, it makes it worse. It contains detailed language about fraud mitigation measures (800-63-4 Section 3.2, and especially 3.2.1) that appears to require agencies to layer additional assurances over the categorical approach. Fraud management is described not as an outcome of identity verification, but as if it were a wholly orthogonal issue. As one of us noted in a public comment, this exacerbates the philosophical tension of the whole approach. Categorical and continuous approaches to identity each have their advantages and disadvantages; the NIST

guidance continues to try to have its cake and eat it too, seemingly ignoring the realities of how individual agencies digest this guidance.

This latter point is a government-wide guidance problem that will ultimately require a government-wide solution approach. Will agencies simply typologize their use cases into discrete boxes of IAL0, IAL1, IAL2, IAL3; or will they do this and *also* define some analogous level of assurance along the essentially-equivalent question of fraud controls? If the latter, how can these assessments be made standard across uses? How does an agency determine if IAL2 plus X fraud control is sufficient for a use case that requires IAL2 and Y fraud control, when X and Y are not standardized?

5.4.2 NIST, agencies, and responsibility

Regardless of what the NIST guidance actually says, there is the cultural and anthropological problem of how agencies interpret it.

Currently, for a combination of reasons, security decisions often end up getting stuck in a circle of buck-passing. NIST will say that 800-63 cannot substitute for decisions made with the full context of an agency's threat model, and agencies are free — and indeed encouraged — to adapt it for their use cases. Agencies will say their hands are tied because they are bound by NIST, and really, how would an agency official defend unilateral straying from NIST? Within agencies, information security officers are asked to prioritize security, but executives are not empowered to make judgment calls accepting risk on the agencies' behalf where needed. Hanging over all of this, inspector generals stand ready to assess decisions from a standpoint weighed heavily on the side of procedural compliance rather than outcomes.

It is this constellation of issues that causes systems like Direct File to adopt a posture that is misaligned with real risks, often substituting having checked the box for actions that could more effectively protect data security. **Addressing these issues will require cultural realignment across the federal government, probably driven centrally by OMB and the White House.**

5.4.3 Continuously improving IdP with offline fallbacks

For whatever set of IAL standards are promulgated, it is critical to have an IdP that implements the proofing standards as easily and equitably as possible. The IdP should be constantly assessing whether users are making it through the identification process, and, if not, why and where they are failing. The IdP should be continuously improving its design and its functionality to ensure more legitimate users can pass, with a particular eye on marginalized populations. This likely includes creating a variety of pathways to verification, including offline fallbacks — like in-person options at Post Offices or other government locations — for those who may struggle with the online process. Issuing cryptographically verifiable credentials like mobile drivers licenses is also a promising avenue of exploration.

And the process needs to be transparent, with honest assessment of progress and areas of weakness.

Login.gov, as a public service, is the most sensible place for this work to occur. It will be an ongoing project requiring ongoing investment to ensure it lives up to its mandate.

5.5 Next steps for Direct File in particular

Section 5.5.1 considers the possible reclassification of Direct File at a lower IAL. For reasons discussed in Section 5.5.2, though, such a reclassification would become less and less important over time. Sections 5.5.3 and 5.5.4 broadly assume that such a reclassification does not occur, and explore what else might be done and what the implications would be for Direct File.

5.5.1 Reclassifying portions of Direct File

As discussed above, NIST 800-63-4 redefines IAL1 far above its 800-63-3 assurance level. This means, almost by definition, there are plenty of systems that were previously classified as IAL2, that should be reclassified at IAL1. **There is a strong case that the core Direct File functionality (not including functionality based on non-MeF IRS-held data, like Data Import and One-Step Signature)** is one of them, and **should be reclassified at IAL1**, given the relatively low sensitivity of disclosures via MeF reject codes. (As discussed in Section 5.5.2, though, the fact Data Import and One-Step Signature would not be included in the IAL1 product lower the impact of such a change.) We believe it would in principle be worth reassessing Direct File's IAL level if the service were reanimated in an 800-63-4 world.

If this baseline Direct File functionality were reclassified, **it would introduce an interesting design challenge within the product: taxpayers would have to select between an IAL2-verified experience with pre-population, and an IAL1 (unverified) experience without it**. One possible approach to this problem was explored in Code for America's early 2023 notional prototype of a Direct File system. That prototype encourages taxpayers to select the IAL2 experience, but gives them the option of continuing unlogged-in. Throughout the experience, when taxpayers are asked for information that would in the IAL2 state be pre-populated or automated, they are prompted again to consider logging in, and can escalate to the IAL2 path. Perhaps a taxpayer is being asked to enter their dependents, and does not want to go through the trouble of entering this information, which the IRS already has from last year. Well, on the dependents start page, they can indicate this, and switch over to the IAL2 path.

The Code for America prototype also proposed the creation of one additional service pathway. Suppose a taxpayer is using the IAL1 product. They successfully enter their family information, but they get stuck on the income section — like many new and intermittent filers, they can't find their W-2s. But they also fail to get through the identity verification process. This taxpayer plainly can't file a full return; their income information is missing. But they also have entered valuable information — information that could be used, in concert

with IRS income data, to issue at least some of a taxpayer's refund (in a sort of "second-best" return; see Chapter 14: Non-filers). Perhaps Direct File should create a pathway for this taxpayer to file their partial return, with the intention of additional data being automated behind the scenes. The motivating logic is that getting *some* information from a taxpayer is better than getting nothing.

That said, all these redesigns rely on the reclassification of the Direct File baseline functionality — and the probability of such a reclassification happening is realistically quite low. The "security ratchet" is real, and it is very hard for agencies to be seen as lowering the standard of a system, even if the updated guidance validates the choice. And effecting the culture change so that agencies *can* do so is a problem much bigger than the IRS or Direct File, as discussed in Section 5.4.2.

5.5.2 Pre-population makes the ID verification problem smaller

While the current IAL classification of Direct File is a liability, the growth of pre-population in the product limits the extent of the problem. As noted in Section 5.2, pre-population does in fact reflect a sensitive disclosure of taxpayer data, which is (and should be) classified as IAL2. So, as pre-population grows ever larger, the benefits to using Direct File in the IAL2 logged-in state grow larger, too; the barriers due to ID verification would eventually grow smaller than the barriers *avoided* by using pre-population in the product. As such, the IAL2 path would become the default, standard Direct File experience. Setting up baseline IAL1 functionality is really about setting up a secondary pathway for a special set of taxpayers with particular barriers to identity verification.

It is still worth creating the IAL1 path described above. But the growth of pre-population makes this less of an urgent issue than it might otherwise be, and less urgent than it was, for example, in the first year of Direct File.

5.5.3 Non-viable service patterns

When the IRS first determined that access to MeF reject codes would be classified as IAL2 under Direct File, the Direct File team considered an alternate design, in which taxpayers would use Direct File without IAL2 verification, and would be prompted to verify if and only if they received a reject code. This team declined to pursue this path, and we believe this was the right choice. It is true that most taxpayers do not receive a reject, and this would have created a smoother path for that majority. But the user experience would be *far* worse for the sizable minority whose returns are rejected. Going through a rigorous identity verification process post-submission, having finished a return, is plainly unintuitive to taxpayers. And Direct File, due to the IAL2 restrictions, would barely be able to explain what it means for the return to be rejected, much less what a taxpayer could do about it. It would raise the prospect of lots of taxpayers simply not resolving their rejects, and perhaps even thinking they had filed their returns when they had not done so. However, Direct File's Digital Identity Risk Assessment (DIRA) paperwork is configured to permit this option should it ever make product sense.

5.5.4 Adopting Login.gov

Unlike when Direct File was initially being developed in 2023, Login.gov now offers IAL2 functionality. **It would be sensible for the IRS to adopt Login.gov** for Direct File (and Online Account) — probably as another option alongside ID.me for minimal disruption to existing users. (This is likely to happen for Online Account prior to any foreseeable reanimation of Direct File, although the disruption of operations at the IRS and GSA makes the timeline unclear.) **Login.gov, as a public service subject to public accountability, should be positioned as the default ID verification service for new users.**

As GAO recommended, the IRS should also exercise independent oversight of ID.me.

5.6 Impacts of generative artificial intelligence

All of the above assumes that the overall threat and technology landscape remains broadly as it exists today. As we write this in 2026, however, generative artificial intelligence is poised to disrupt the entire problem space of identity verification. Generative AI is probably very near, or even already at, the point that it can spoof the high-bar standards of remote IAL2 proofing. We do not pretend to know what comes next for remote identity proofing in this changed world, and how it will change the equity, accessibility, and reusability considerations discussed here. We merely offer the comment that all of the above may be moot quite soon, and certainly by the time Direct File may be reanimated in a future administration.

The principles will still be the same, though. Direct File will need to find a way to set up the right level of assurance and consistently iterate to ensure and promote access. Across the government, reforms and additional guidance will be needed to reconcile the categorical and continuous approaches that are both implicit in NIST guidance.