



Amplivy Submission to the Scams Prevention Framework

Exposure Draft Sector Codes and Rules

Identity establishes who opened the account. Presence establishes who executed the event.

Consultation: Scams Prevention Framework: sector codes and rules (c2026-765133)

Submitted by: Amplivy, Inc. (United States)

Date: 25 June 2026

Status: Public

Amplivy, Inc. welcomes the opportunity to respond to the exposure-draft sector codes and rules. Amplivy is a United States company founded by an Australian. Amplivy has built and patented trust infrastructure for the digital economy we call Verified Presence Infrastructure: infrastructure that establishes, deterministically, that a real human being is physically present and in voluntary control at the moment a digital event executes. The technology is protected by a granted patent family in the United States, Europe and Canada, with an associated continuation estate, and is being prepared for deployment in Australia.

Executive summary

Australia's Scams Prevention Framework makes verification its central preventive lever, then leaves the method undefined. The consultation guide asks what specific actions are appropriate for verification. This submission answers that question.

The method the codes implicitly assume, document-custody identity verification (the know-your-customer model), is failing. It was being defeated by synthetic identities before generative AI; AI has now industrialised that failure. And because it requires third parties to capture and retain complete identity records, it has become a fraud accelerant: that retained data is the raw material of the breaches that fund the next generation of impersonation fraud. A standard built on retaining more identity data cannot be the answer to a fraud problem driven by the theft of identity data.

There is a different standard, and a class of infrastructure now exists to deliver it: event-time verification of human presence. It does not verify a document; it verifies, deterministically and in real time, that a real human being is present and in voluntary control at the moment a protected event executes. Identity establishes who opened the account; presence establishes who executed the event. The codes ask the first question and leave the second unanswered.

This standard eliminates the category of fraud that scales: the absent-human fraud that AI has weaponised, including account takeover, impersonation, synthetic counterparties, deepfakes, and automated or AI-agent action. It does not prevent a real, present human from being deceived

or coerced; that harm does not scale, and it remains the domain of the institution's own controls and of law enforcement.

We ask one thing: that the codes specify the verification outcome rather than any particular technology, so that discharging the verification obligation requires establishing, at the time a protected event executes, that a real human is present and in control, and not merely that documents were checked at onboarding. Draftable language is set out in the appendix.

This submission is confined to policy; it names no product and proposes no commercial terms.

1. The question the record leaves open

Across the first-round submissions the debate divided almost entirely between liability and reimbursement: who bears the loss, in what proportion, and through which dispute process. That debate matters, but it begins after the money has moved. Not one submission proposed verifying that a real human being is present at the moment a protected event executes. Presence, as distinct from identity, does not appear in the record as an architectural proposition.

This matters because the draft codes make verification the central preventive lever, and then leave the method undefined. The banking code (section 3-3) requires banks to verify the identity of consumers “to prevent any misuse or manipulation by scammers,” and the guide asks whether it is clear how banks would implement this and “what further specific actions are appropriate for verification.” The digital platform code (sections 5-3 and 5-4) requires platforms to verify new users and advertisers, and to re-verify on a detected change “to address potential account takeover”; again asking what specific verification actions are appropriate. The telecommunications code asks what additional checks should be undertaken before higher-risk services, such as the overseas use of numbers or the sending of bulk messages.

Treasury has, in effect, asked the market to supply the verification standard. The first round did not. This submission does.

2. Why the method is the whole question

A verification obligation is only as strong as the method that discharges it. Two things bear on that: the scale of the problem the method must address, and the fact that the method the codes implicitly assume, document-custody identity verification (the “know your customer” or KYC model), is failing on two independent fronts, the second of which is structural.

2.1 The scale is larger than the record shows, and the record is a floor

Australians reported A\$2.18 billion in scam losses in 2025 (National Anti-Scam Centre, Targeting Scams Report 2025). Globally, financial-sector fraud losses, meaning bank fraud and consumer scams combined, reached US\$579.4 billion in the same year, part of an estimated US\$4.4 trillion in total illicit financial activity (Nasdaq Verafin Global Financial Crime Report 2026). That figure excludes fraud against government programs, which is additive: in the United States alone, the Government Accountability Office estimates annual fraud against the federal government at between US\$233 billion and US\$521 billion. Two things should be said about the Australian figure before it is relied upon.

First, it is not the cost of fraud; it is a fraction of it. Every A\$1 of fraud costs Australian organisations A\$3.68 once detection, investigation, remediation and false-positive handling are counted (LexisNexis True Cost of Fraud, Asia-Pacific 2024). The headline loss is roughly a

quarter of the economic burden; the remainder is the cost of the detection paradigm itself: the spend that chases fraud after it has occurred.

Second, and more important, the reported figure measures detected and reported fraud. It does not measure successful fraud. Fraud that succeeds without detection is, by definition, absent from the statistics; the Australian Institute of Criminology states that official figures “significantly underestimate” the true scale, and the Commonwealth’s own fraud-prevention body describes reported fraud as the visible tip of an iceberg. The true scale of successful fraud in the Australian economy is unknown.

This is not a quibble about data quality. It is central to the policy question, because the fraud that is hardest to count is precisely the fraud this framework must address. Fraud engineered to pass as legitimate does not announce itself: a synthetic identity that clears onboarding, a deepfake that clears a liveness check, an account takeover that transacts as the real customer. The better the fraud, the less likely it is even recorded. As artificial intelligence improves the quality of that fraud, the gap between what is counted and what occurs widens. Every figure in this submission should be read as a floor.

2.2 Document-custody verification was already being defeated before AI

The assumption underlying the codes’ verification obligation is that checking a person’s identity documents at onboarding establishes that the person is genuine. That assumption was failing before generative AI arrived. Synthetic-identity fraud (assembling a credible identity from a blend of real and fabricated data) was already the fastest-growing financial crime in comparable markets, precisely because document verification was designed to confirm that a real person exists, not to detect a well-constructed fiction (Deloitte Center for Financial Services).

Generative AI did not create this failure. It industrialised it, collapsing the cost and skill required to defeat document verification at scale. The trajectory is not speculative: gen-AI-enabled fraud is projected to grow from US\$12.3 billion to US\$40 billion across comparable markets between 2023 and 2027, a 32% compound annual rate (Deloitte). Material presented at identity verification is now 300% more likely to be entirely AI-generated than a year earlier, and impersonation accounts for 85.6% of verification fraud attempts (Veriff Identity Fraud Report 2026). Human inspection is no longer a defence: average human accuracy at detecting a deepfake is barely better than chance (Veriff Deepfake Report 2026), and 90% of financial institutions report a rise in AI-driven attacks (Nasdaq Verafin Global Financial Crime Report 2026). In a joint statement on 23 June 2026, the heads of the Five Eyes cyber-security agencies, the Australian Signals Directorate among them, warned that AI is reshaping cyber risk in months rather than years and urged leaders to act now; independent testing of a frontier model found it could compromise systems around 73% of the time, described as a step change in capability.

2.3 The document-custody model is not merely failing; it is a fraud accelerant

There is a deeper problem, and it bears directly on the privacy balance the consultation asks about (General Question 3). Document-custody verification requires a third party to capture and retain the consumer’s identity documents; in practice it captures the entire document, far beyond the single attribute the check required. That retained store of complete identity records is the raw material of the data-breach epidemic, and breached identity data is what funds the next generation of impersonation and synthetic-identity fraud.

This is not theoretical. Identity data exposed in Australia’s large-scale breaches has been re-used to defeat document-based identity checks: the exploitation of stolen documents to create fraudulent myGov accounts and divert more than A\$557 million is a direct example

(acknowledged by the Australian Taxation Office; reported by the ABC and through freedom-of-information). The verification industry itself carries the same exposure. A third-party identity provider exposed roughly 70,000 government IDs in October 2025; a category-leading verification provider left its entire government-service dashboard codebase exposed on a public endpoint in February 2026; a separate provider exposed approximately one billion identity records. The Electronic Frontier Foundation's assessment is that such data leaks "through hacks, misconfigurations, or retention mistakes"; the question is not whether, but when.

The implication for the codes is uncomfortable but unavoidable: mandating more document-custody verification as the answer to scams enlarges the very honeypot that produces the breached identities scammers use. A verification standard that increases the volume of retained identity data cannot be the long-term answer to a fraud problem driven by the theft of retained identity data. It sits on the wrong side of the privacy-versus-prevention balance.

There is a further reason the current paradigm cannot prevail. It is built on static trust: stored credentials, long-lived keys, and the assumption that what is computationally hard today stays hard. That assumption is time-bounded. Google Quantum AI's March 2026 resource estimates indicate that, once a sufficiently capable quantum computer exists, a 256-bit elliptic-curve key could be broken in minutes; and "harvest now, decrypt later" means identity data stolen today need not stay encrypted forever. A defence that depends on keeping stored secrets secret, while bound by laws that those weaponising AI simply ignore, is structurally on the losing side of the asymmetry. The only durable answer is to stop depending on stored secrets at all: to verify, in real time, something that cannot be precomputed or replayed, the presence of a real human at the event.

3. The standard we propose: deterministic proof of presence

A class of event-time verification is emerging that does not share these failure modes. It does not verify a document; it verifies presence: that a real human being is physically present and in voluntary control at the moment a protected event executes.

3.1 What it is

Deterministic presence verification operates at the event, not at onboarding. It does not collect or retain identity documents. It does not depend on human inspection of an image. It produces a verifiable attestation that a real human was present at this specific event, which is delivered to the regulated entity for its records. Amplivy has developed one implementation of this class of architecture, built as patented infrastructure; this submission describes what such verification does, not how a particular implementation does it, and does not ask the codes to specify any provider. Because it retains no identity documents and the attestation can be generated within Australian infrastructure, it also avoids the data-sovereignty exposure that arises when verification depends on an offshore custodian of identity records.

The distinction from KYC is not a matter of degree. Identity establishes who opened the account; presence establishes who executed the event. KYC asks: is this person who they claim to be, at account creation? Presence verification asks: is a real human present and in control, at this event? Onboarding identity checks remain necessary for their own purposes (anti-money-laundering among them), but they are insufficient against fraud that occurs at the event, after onboarding, through a genuine account. The two are complementary; the codes can require both, and the interaction with AML/CTF customer due diligence that the guide raises is resolved by treating presence verification as an event-level control sitting alongside, not replacing, that due diligence.

This is not behavioural tracking, device fingerprinting, or a one-time code, and it should not be mistaken for them. Those legacy event controls are themselves defeated by the absent-human attacks at issue: a one-time SMS code by SIM-swap, a device identifier or session cookie by session hijacking, and behavioural analytics infers a probability rather than establishing a fact. Presence verification establishes, in real time, that a human is present, and it can operate in the background without adding a step or a prompt for the user, so prevention need not come at the cost of the experience the regulated entities are protecting.

3.2 What it does, and what it does not do

Deterministic presence verification eliminates the category of fraud that scales: absent-human fraud. Account takeover, impersonation of a known identity, synthetic counterparties, deepfake voice and video, credential abuse, session hijacking, and unauthorised automated or AI-agent action all share one feature: the legitimate human is not present at the event. This is the category artificial intelligence has weaponised into an industrial-scale threat, and it is the category presence verification deterministically defeats.

It does not, and is not designed to, prevent a real, present human from being deceived into authorising a payment themselves, or from acting under coercion. That harm is real, but it does not scale: it is bounded by individual human effort, it is not multiplied by AI, and it is properly the domain of the regulated entity's own controls, of law enforcement, and of the courts. By design, presence verification establishes a fact and delivers it; it does not adjudicate the legitimacy of a present human's decision. That boundary is a feature of the architecture, not a limitation of it: it is why the infrastructure is neutral, makes no automated determination about any individual, and retains no profile or history of any person, a point of direct relevance to concerns about automated decision-making and surveillance.

This does not mean presence verification is absent from the authorised-push-payment problem. Deployed at the carrier layer, as described in section 4.3, it disrupts that problem at its origin: by establishing that a call or message comes from a verified number and a real, present human, it strips trust from the unsolicited scammer contact that initiates many such scams before a victim is ever engaged. Presence verification is therefore additive to the codes' broader prevent, detect and disrupt obligations, not a substitute for them or a safe harbour from them.

4. How the standard answers each sector's code question

4.1 Banking (section 3-3)

The banking sector's own first-round submissions identify the gap this standard fills. The Commonwealth Bank told Treasury that "if a consumer has given a scammer their password or remote access, re-verifying their identity will not help." National Australia Bank described the "first-party mule", a person who opens an account with genuine identity documents and then hands control to a criminal, and concluded that "identity verification is unlikely to aid prevention or detection in these cases." Both describe the same thing: identity checked at onboarding is silent at the event. Deterministic presence verification speaks precisely where identity verification is silent; it establishes whether the legitimate human, rather than a remote operator or an automated agent, is present and in control at the moment of the transaction. It is the "further specific action appropriate for verification" the banking question asks for, and it complements the customer due diligence banks already perform under AML/CTF law.

4.2 Digital platforms (sections 5-3 and 5-4)

The platform sector's central objection to the verification obligation is that it forces the collection of more personal data. The Digital Industry Group warned that the framework "encourages the collection of more PII" and "may inadvertently put consumers at greater risk"; major platforms argued they cannot practically verify, at scale, advertisers who reach them through agencies and resellers rather than directly. Deterministic presence verification answers both. Because it verifies presence without capturing or retaining identity documents, it satisfies the verification and re-verification (account-takeover) obligations in sections 5-3 and 5-4 without enlarging the data store the platforms do not want to hold. It is the rare verification approach that can strengthen prevention while reducing privacy exposure at the same time.

This holds even for end-to-end encrypted messaging services such as WhatsApp and Signal. The platforms' objection to content-based obligations is that they cannot, and should not, read encrypted messages. Presence verification does not ask them to. It verifies that a real human is present and in control of the account sending a message; it does not inspect the message. A platform can therefore satisfy a presence-verification obligation without weakening encryption and without reading private content, removing the conflict the sector has identified between scam prevention and message confidentiality.

4.3 Telecommunications

The telecommunications sector told Treasury, candidly, that carriers "do not detect scams or fraud; they only detect suspicious activity at a systemic level rather than for individual communications" (Australian Telecommunications Alliance), and that they cannot inspect message content without breaching the privacy provisions of Part 13 of the Telecommunications Act and the Telecommunications (Interception and Access) Act. Deterministic presence verification gives carriers an individual-level signal they currently lack: confirmation that a real, present human, rather than a hijacked number or an automated origination, is behind a high-risk action, without inspecting content, and therefore without crossing the privacy boundary the sector correctly identifies. It converts the carriers' strongest argument against liability into a workable basis for prevention.

It also completes work the government has already begun. The SMS Sender ID Register verifies that a sender identifier is registered; it does not establish that a real person, rather than an automated system or a compromised number, is behind a given call or message. Presence verification adds that layer. With it, a recipient can be shown that a call or message genuinely originates from the party it claims to, and that a real human is present behind it: the verified signal a consumer can actually trust.

It is also the layer at which industrial origination fraud fails. SIM farms and comparable bulk-messaging operations depend on automated origination at scale, with no human present at any individual event. Because presence verification requires a real, present human, automated bulk origination cannot satisfy it.

5. Why this standard serves every objective in the consultation

The first-round record contains a problem every participant circled and none resolved: the "reasonable steps" test is self-assessed and, therefore, unverifiable. Consumer advocates fear regulated entities will "self-assess as having taken reasonable steps" to avoid liability. The Australian Securities and Investments Commission told Treasury that requirements "to merely have processes or methods in place are not sufficient"; obligations must be effective in achieving outcomes. The heads of the Five Eyes cyber-security agencies made the same point in June

2026, stressing that leaders need confidence their defences would hold in a real attack rather than simply having controls in place. Banks, for their part, want liability set by deterministic rules rather than subjective judgement. These positions look opposed. They are reconciled by the same thing: a deterministic, auditable record of whether a real human was present at the event converts “reasonable steps” from an assertion into evidence.

This is why presence verification is not a concession to one camp at the expense of another. It is the rare measure that advances every stated objective at once:

Stakeholder	Concern raised in the first-round consultation	How deterministic presence verification answers it
Banks (ABA, CBA, NAB)	“Reasonable steps” is too subjective; onboarding identity does not stop control-handover, remote-access or first-party-mule fraud	Provides a deterministic, auditable control at the event, addressing the failure modes the banks say identity verification cannot reach
Telcos (ATA, Telstra, Optus)	Carriers can only detect at a systemic level, and cannot inspect content without breaching Part 13 / TIA privacy provisions	Supplies an individual-level presence signal without inspecting content, and so without crossing the privacy boundary
Digital platforms (DIGI, Meta, Microsoft)	Verification forces collection of more personal data, increasing risk; advertiser verification does not scale	Verifies presence without capturing or retaining identity documents, and works for end-to-end encrypted messaging without decrypting content; prevention up, retained data down
Regulators (ASIC, ACCC)	Obligations must be outcome-based and effective, not process-based; “enhanced identity verification ... will be crucial” (ACCC)	Delivers the outcome-grade, auditable verification the regulators ask for
Dispute resolution (AFCA)	Codes should require entities to “verify whether it is them”; redress needs evidence, not assertion	Produces the evidentiary record needed to resolve complaints at IDR and EDR
Consumers and victims (Joint Consumer, Legal Aid)	Prevention must not rely on warnings; opt-outs are exploited by scammer coaching	Prevents at the event, before money moves; being frictionless, leaves no friction for scammers to coach victims into bypassing
Privacy (General Question 3)	The codes must balance effective regulation with personal privacy	Increases prevention and reduces retained identity data at the same time

The Australian Financial Complaints Authority came closest, in the first round, to describing what is needed, calling for codes that require entities to “verify whether it is them” and to detect remote-access and absent-human indicators. That is a description of the outcome. Deterministic presence verification is the architecture that delivers the outcome; and because it retains no identity documents, it delivers it without the data-retention exposure that other participants, the Digital Industry Group among them, warned against. One participant wants the outcome; another fears the usual method; presence verification provides the first without the second.

6. The framework is already moving in this direction

The logic of attempt-time prevention is not novel to the Australian market; it is where the market is already heading. As National Australia Bank’s Executive Lead for Investigations, Chris

Sheehan, has put it: “If you detect something after it’s happened, the criminality is already occurring. If you had prevented the crime from occurring in the first place, you could have avoided all of that.” He has also observed that banks “can only see a very tiny portion of this ... the payment process is right at the end of the chain” (Nasdaq Verafin Global Financial Crime Report 2026). The Australian Financial Crimes Exchange, the Fraud Reporting Exchange and the Anti-Scam Intelligence Loop are all, in their different ways, attempts to push prevention further up the chain, toward the point of origination.

Deterministic presence verification is the furthest-upstream point on that chain: the event itself, before value moves. It is the same standard the codes are reaching for, expressed as infrastructure. One standard answers the verification question for all three designated sectors, because the question is the same in each; and it will answer that question equally for the sectors the government has flagged for possible future designation, including superannuation, insurance, online marketplaces and digital currency exchanges. Fraud migrates to the weakest verification surface, which is why a standard consistent across sectors matters: it leaves no weakest link for fraud to find. Presence verification should be seen as complementary to those existing investments, not as a replacement for them.

7. What we ask

We do not ask Treasury to mandate any product, including ours. We ask that the codes specify the verification outcome, not a particular technology: that the standard for discharging the verification obligations in the banking, digital-platform and telecommunications codes be that verification deterministically establishes that a real human being is present and in control at the moment a protected event executes, and not merely that identity documents were checked at onboarding. A standard expressed as an outcome is technology-neutral, durable against the next generation of fraud, and measurable in a way that “reasonable steps” alone is not.

Specifying an outcome does not require Treasury to endorse or inspect any particular mechanism. Providers compete to meet it, and the code can require the privacy properties that meeting it must respect, including that no biometric or identity record be retained. This keeps the standard open to any provider and favours no single architecture.

We have been candid about the boundary of what this standard does: it eliminates the absent-human fraud that scales, and leaves the deceived-or-coerced present human to the regulated entity’s own controls and to law enforcement.

This submission is confined to policy. It contains no commercial terms and no pricing. Amplivy would be glad to provide Treasury and the sector regulators with further technical detail on request.

8. Closing

The first-round consultation asked the market what specific actions are appropriate for verification, and the market answered with a debate about who pays when verification fails. The architecture that prevents the failure, proof that a real human is present before the event executes, was missing from the record. Identity establishes who opened the account; presence establishes who executed the event; this submission is directed at the second. Amplivy would welcome the opportunity to assist Treasury further.

Rodric David

Founder & Chief Executive Officer,
Amplivy, Inc.



Appendix: proposed drafting

The following is offered as draftable language the codes could adopt. It is expressed as an outcome and is technology-neutral.

Principle: event-time verification of human presence

Where a sector code requires verification, a regulated entity should be able to establish, at the time a protected event executes, that a real human is present and in voluntary control, and should not rely solely on identity checks performed at onboarding. Controls implementing this principle should:

- be technology-neutral, specifying the outcome rather than any particular method or provider;
- establish presence deterministically where practicable, rather than inferring it;
- complement, not replace, existing customer due diligence under AML/CTF law;
- minimise the collection and retention of identity data;
- operate consistently with privacy law and with encryption, without inspecting message content;
- produce an auditable record capable of supporting internal and external dispute resolution.

Sources

1. National Anti-Scam Centre (ACCC), Targeting Scams Report 2025: A\$2.18 billion reported scam losses, 2025. <https://www.nasc.gov.au>
2. LexisNexis Risk Solutions, True Cost of Fraud Study, Asia-Pacific 2024: A\$3.68 total cost per A\$1 of fraud. <https://risk.lexisnexis.com/insights-resources/research/apac-true-cost-of-fraud-study>
3. Australian Institute of Criminology, Cybercrime in Australia 2024 (SR53), and the Commonwealth Fraud Prevention Centre: under-reporting and the iceberg of undetected fraud. <https://www.aic.gov.au/publications/sr/sr53>
4. Deloitte Center for Financial Services, Generative AI and the risk of deepfakes in banking: US\$12.3 billion (2023) to US\$40 billion (2027), 32% CAGR; synthetic identity as the fastest-growing financial crime. <https://www.deloitte.com/us/en/insights/industry/financial-services/deepfake-banking-fraud-risk-on-the-rise.html>
5. Veriff, Identity Fraud Report 2026 (AI-generated material up 300%; impersonation 85.6% of attempts) and Deepfake Detection Report 2026 (human detection near chance). <https://www.veriff.com>
6. Nasdaq Verafin, Global Financial Crime Report 2026: US\$579.4 billion in global fraud losses within US\$4.4 trillion of illicit activity (2025); 90% of institutions report rising AI-driven attacks; Australia spotlight featuring Chris Sheehan, National Australia Bank. <https://verafin.com/nasdaq-verafin-global-financial-crime-report/>
7. United States Government Accountability Office, GAO-24-105833 (2024): estimated annual fraud against the federal government of between US\$233 billion and US\$521 billion. <https://www.gao.gov/products/gao-24-105833>
8. Google Quantum AI (Babbush, Zalcman, Gidney, Neven et al., with Stanford and the Ethereum Foundation), Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities, 30 March 2026: resource estimates indicating that a 256-bit elliptic-curve key could be broken in minutes once a sufficiently capable quantum computer exists. <https://quantumai.google/static/site-assets/downloads/cryptocurrency-whitepaper.pdf>
9. Verizon, 2026 Data Breach Investigations Report: human element present in 62% of breaches. <https://www.verizon.com/business/resources/reports/dbir/>
10. Five Eyes cyber-security agencies, including the Australian Signals Directorate, joint statement on AI and cyber risk, 23 June 2026 (“act now”; AI reshaping risk in months not years), with UK AI Security Institute frontier-model testing. <https://www.smh.com.au/technology/act-now-cyber-spy-chiefs-issue-warning-on-ai-threat-20260622-p608z6.html>
11. Australian Taxation Office (acknowledged; reported by the ABC and through freedom-of-information): more than A\$557 million diverted via fraudulent myGov accounts using stolen identity documents.
12. First-round consultation submissions to Treasury (designation and position-paper package, closed January 2026): Commonwealth Bank, National Australia Bank, Australian Banking Association, Australian Telecommunications Alliance, Telstra, Optus, Digital Industry Group, Meta, Microsoft, Australian Securities and Investments Commission, Australian Competition and Consumer Commission, Australian Financial Complaints Authority, the joint consumer submission, and National Legal Aid. <https://treasury.gov.au/consultation>
13. Treasury, Guide to Scams Prevention Framework rules and sector codes consultation (c2026-765133): the verification consultation questions (banking section 3-3; digital platforms sections 5-3 and 5-4; telecommunications). <https://consult.treasury.gov.au/c2026-765133>