

Age Assurance Technology Trial

Document Sensitivity: Public

Stakeholder Advisory Board Minutes

08/07/2025

Becky Burgess

© Commonwealth of Australia

Date: 29/01/2025

Doc. Version: 1

Page 1 of 7

Funded by



Australian Government

**Department of Infrastructure, Transport,
Regional Development, Communications and the Arts**

Project by





Minutes of the Second Stakeholder Advisory Board Meeting held Via Teams on Tuesday, 8 July 2025, at 4 pm AEST

In Attendance:

Professor Jon Rouse APM	Chair
Iain Corby	AATT Engagement Lead
Jenny Duxbury	INK & Digi
Andrew Hammond	KJR
Tony Allen	ACCS
Ajoy Ghosh	Cyber Alchemist
Kirra Pendergast	Safe on Social
Simon Callaghan	RegTech
Cheryl Seeto	Meta
Colm Gannon	ICMEC
John Pane	Electronic Frontiers Australia Inc.
Ben Au	Snap
Yvette D’Ath	Independent
Susan McLean	Cybersafety Solutions
Ron Curry	IGEA
Melinda Tankard Reist	Collective Shout
Conrad Townson	IFYS
George Billinge	Illuminate Tech
Amber Hawkes	Blue Lantern Consulting
(Observers) Liz Miller Sally O’Donnell Sarah Vandenbroek	e-Safety

Apologies:

Maree Crabbe	It’s time we talked
--------------	---------------------



Introductions

Jon Rouse welcomed attendees to the sixth meeting of the Stakeholder Advisory Board. A quorum was confirmed. A reminder was issued regarding confidentiality until the report's public release.

Feedback from the minutes from the 17th of June meeting was sought. Iain confirmed that amendments had already been made based on feedback received from Ajoy, which included two specific textual changes. Jon Rouse then invited the group to raise any further proposed amendments. Ben offered general feedback, noting that while the minutes captured the presentations well, the discussions were summarised too generally. Jon invited Ben to submit his thoughts within the week, and Iain confirmed that the team would wait a week before finalising the minutes.

Presentation of the Report Drafts

Tony Allen presented an update on the report and its structure. He reiterated that the project was delayed from its intended start, beginning only in mid-to-late November, which has had a knock-on effect on delivery timelines. Nevertheless, the team's goal remained to complete the final report by the end of July, with a publication date in August to be determined by the Minister. The report is structured in ten parts (A through K), with each section spanning approximately 60–70 pages in Word format.

Tony provided an overview of the document structure. Part A serves as the executive summary of the entire report, which had been circulated to the board ahead of the meeting. Part B, covering methodology and ethics, was also provided, although most of its contents had already been shared as part of the earlier Evaluation Proposal. Part K includes the glossary, bibliography, and references, also now circulated. Parts C through J, which contain in-depth analysis of specific technologies and testing results, had not yet been shared with the board due to the inclusion of sensitive vendor-specific information. Each vendor was sent extracts relevant to their service, including their interview notes, privacy policies, testing outcomes, and a copy of their practice statement. Vendors have been given until the July 18 to submit their feedback.

Questions and comments

Colm Gannon asked how ISO/IEC 27566 is referenced throughout the document. Colm noted that while Part K properly described it as an "emerging standard," other sections failed to consistently use that qualifier. Given the importance of distinguishing between formal standards and those still in draft stages, especially in research and regulatory contexts, he recommended this be harmonised throughout the report.

Tony accepted this, stating that while the report does note that the standard is in FDIS (Final Draft International Standard) stage, that alone might not be understood by all readers, he will increase consistency throughout the report. Iain added that no significant content changes were anticipated between the FDIS version and the final one, and that the standard had remained consistent throughout the trial period.

Ben Au asked about the report's statistical findings, particularly the use and communication of Mean Absolute Error (MAE) in age estimation technologies. He expressed concern that the figure of 2.7



years might become a focal point for public scrutiny and media coverage, especially since facial estimation technologies have already drawn significant press attention, namely the [ABC piece](#).

Tony agreed that the timing of the news report wasn't ideal and noted the added complexity to the Trial. He said that this average obscures a significant range of vendor performances, from as low as 1.03 years to as high as 4.6 years in MAE. He shared that the team was now considering replacing the average with a range in the final report, which would more accurately represent vendor capabilities.

Ben stressed the importance of ensuring that any high-level findings, particularly those that will be quoted by journalists and policymakers, are clearly supported by the underlying data. Iain echoed this concern, reiterating that age estimation technologies were never intended to identify precise ages, only to establish general age eligibility.

Ajoy questioned how the government's commitment to non-ID-based verification methods squared with the findings, as most methods appeared to rely, directly or indirectly, on some form of government-issued documentation.

Tony clarified that verification depends on authoritative sources of birth data, which could include bank or school records, and not necessarily identity documents like passports and noted the buffer age to aid accuracy. Tony noted that he is unable to comment on Government decisions. Ajoy remained sceptical, arguing that those alternative sources often originate from ID-based processes. Tony reiterated that there are many places where date of birth is recorded, ultimately stemming from a government issued birth certificate.

He also asked whether the testing included circumvention techniques discovered through related research in his own work on biometric systems. Tony confirmed that key circumvention methods were tested and are discussed in Part J, but that the report intentionally omits details that could serve as a how-to guide for malicious actors. He offered to follow up with Ajoy privately to discuss technical details further.

Ajoy also expressed concern that certification under ISO 27001 was accepted without a detailed evaluation of the vendors' Statements of Applicability. Tony responded that the report applied the more targeted criteria from ISO 27566 and that the level of analysis was appropriate for a trial, not a certification program. He emphasised that individual vendor certifications should follow as a next step.

Jenny Duxbury submitted her written feedback ahead of the meeting and summarised her points verbally. She cautioned against referencing or praising government policy, which could compromise the perceived independence of the report. She also requested clearer definitions of accuracy and its measurement across different age assurance methods, especially given the public's likely interest in effectiveness with children. Jenny found the term "live testing" vague and sought clarification on its meaning.

Tony noted that the term "government policy" was only used once in the report, noting this was out of scope for the Trial. Tony responded to her accuracy question by saying that there is no single, universally accepted benchmark for what constitutes "sufficient accuracy", largely because different types of age assurance operate using different mechanisms and require different evaluative approaches. Tony reassured her that the term referred to real-world user interactions (e.g., using a device at home), as opposed to lab-based photo testing. Iain added that children were included in the mystery shopper program, primarily through their parents. Jenny was surprised by this, having



previously understood that children would not be involved. She maintained that clarity around accuracy measurement and methodological terminology is crucial to preserving the report's credibility. Iain highlighted it's important to recognise the different methods and noted the technology was reported on but not judged. Jenny reiterated the importance of transparency in the report to maintain credibility.

Ajoy recommended that, given the absence of clearly stated pass/fail thresholds for accuracy in the trial, the team should consider drawing on benchmarks from other standards, such as ICAO (International Civil Aviation Organization) benchmarks for passport and biometric verification, and standards for driver's license authentication. He proposed that referencing such benchmarks could help contextualise what "acceptable accuracy" might look like in the age assurance context. Jon thanked Ajoy for his suggestion.

John Pane raised concerns about how the report presents privacy practices, especially in its headline claims that age assurance can be done "privately, efficiently, and effectively." He cautioned that such phrases, if not well-defined and evidenced, risk being misleading, particularly when cited by policymakers or regulators. John argued that the presence of a privacy policy or mention of "privacy by design" is insufficient; what matters is the actual governance behind those policies. He was especially critical of vendors who appeared to store personal or biometric data "just in case" it might be needed by regulators or law enforcement, suggesting this showed a misunderstanding of privacy principles such as data minimisation and purpose limitation.

Tony acknowledged the difficulty of summarising complex findings for public audiences but defended the report's underlying rigor. He emphasised that the team's analysis of privacy went beyond reading policy documents, it included review of practice statements, real-world implementation, and data flow analysis. Tony noted that while many vendors showed a strong grasp of privacy, the report did identify instances of unnecessary data retention, which are flagged as concerns. He stressed, however, that the report is not a certification or compliance assessment. Final responsibility for regulation and enforcement lies with bodies like the OAIC and the eSafety Commissioner. Tony agreed that clearer regulatory guidance is needed to prevent vendors from over-retaining data out of uncertainty.

Ben recommended that accuracy be highlighted in the executive findings and that more attention be paid to user and community perspectives, particularly regarding identity requirements. He also underscored the importance of explicitly discussing circumvention in the findings, even if only to acknowledge its limitations. Ben emphasised that findings should reflect both the strengths and weaknesses of the technologies tested to avoid undermining public confidence.

Tony confirmed that circumvention, ID use, and user feedback, including usability, were addressed in the full report and that findings would be revised with these comments in mind. He also noted the report the Australian government have published on user perceptions on age assurance, which is now in the public domain.

Colm offered pointed observations on the report's readability and consistency. He identified specific examples where ISO standards were clearly explained (Section 11.1), and others where terminology and references were confusing (Section 8.1). He recommended that abbreviations such as MAE be fully spelled out and contextualised, particularly given the contentious and technical nature of the report. Tony welcomed the feedback, acknowledged the varied writing styles across the team of 50 contributors, and committed to improving clarity and standardising terminology. He noted that a glossary had been included in response to previous board feedback, and that editing and design were still ongoing. Colm singled out page 27 of the report for particular praise, commending its use



of explanatory text after each referenced ISO standard. He said this section effectively balanced technical depth with accessibility and should serve as a model for the rest of the report. Tony thanked Colm for his feedback and will feed it back to his team.

Tony responded directly to a comment raised by Susan in the chat, which had suggested writing the report to suit the “lowest common denominator.” Tony respectfully disagreed, stating that while accessibility is important, over-simplifying the report would lead to a loss of granularity and scientific precision. He emphasised that the report must strike a delicate balance, being readable to informed laypeople and policymakers, while still offering the necessary depth and technical accuracy expected by researchers, regulators, and practitioners.

Conrad Townson also addressed readability and data presentation. He recommended the report aim for a reading level consistent with Year 9 education, as this is the standard journalistic threshold for general public comprehension. He requested more transparency around the demographic breakdown of test subjects, particularly age ranges.

Tony responded that these breakdowns are included in the detailed reports not yet circulated, and that all non-personal statistical data would be published through the Open Science Foundation for independent review and further research. He also noted that hyperlinks in the final report would enable easy navigation to relevant sections and datasets.

Colm recommended that the team proactively surface this kind of breakdown early in the report or in a clearly marked section, showing how technologies performed for different age brackets (e.g., 10–12, 13–15, 16–17) as well as for adults. Doing this would not only support clearer interpretation but would also help pre-empt criticism or confusion once the report is publicly released.

Tony said explaining concepts like confidence intervals, statistical variance, or error distribution to a general audience, especially without losing their meaning, remains a difficult task. That said, feedback from the board, including Conrad’s comments, would be carefully considered as the final editorial phase proceeded.

Closing Remarks and Timeline

Tony provided further clarity on the report’s completion schedule and the procedural steps that would follow the session.

He confirmed that the team was currently in the middle of the procedural fairness process, a stage in which all vendors involved in the trial had been given the opportunity to review and respond to the portions of the report pertaining to their own technologies. Tony noted that they anticipated finalising this process between the 20th and 23rd of July. This would allow for a short period to incorporate vendor feedback, whether factual corrections, technical clarifications, or challenges to characterisation, before proceeding to the final edit.

Tony confirmed that the final report would be submitted to the government on July 31, in line with the established schedule. However, he clarified that public release would be a separate process, dependent on the Department’s internal handling. He mentioned that certain legal and filing checks would need to occur before the report could be made public. While the intention is still for the report to be released sometime in August, the exact publication date will be determined by the Minister and her office. Tony emphasised that the project team is in regular communication with departmental officials and promised to keep the board informed as soon as any decisions on release timing are confirmed.



Iain addressed a concern raised earlier, that without visibility of Parts C through J (the detailed vendor analyses and technical breakdowns), board members might worry that their feedback could be limited to the summary alone. Tony and his team would ensure that every comment submitted on Part A, whether during the meeting or via email, would be systematically reviewed and carried through to the relevant sections of the full report.

Tony invited all members to submit any additional written comments by Monday, July 15, ideally with clear references to page numbers or sections where changes were being suggested. He noted that there was still time to make meaningful edits before final design and formatting began, but emphasised that this window would close quickly.

Iain committed to personally review all other sections of the report (Parts B through J) to ensure that any issues identified at the summary level were consistently addressed throughout the technical and vendor-specific sections. Iain emphasised that the structure of the report had been deliberately built to cascade downward, meaning that there should be no conclusions or assertions in the deeper parts of the report that had not already been flagged or signalled in Part A.

Tony noted that Part B (Methodology and Ethics) and Part K (References and Glossary), which had also been circulated to the board. He reminded participants that Part B was largely identical to the evaluation proposal the board had already reviewed earlier in the process. The main difference, he noted, was that it had now been placed in the context of a final report, rather than a prospective plan.

However, Tony pointed out that there was one new and important addition in Part B: a section documenting the Stakeholder Advisory Board's feedback across the life of the project, and how the project team had responded to each major issue raised. This section had been built from the official minutes of past board meetings. It was laid out as a table, with stakeholder comments in one column and the corresponding project team response in another. Tony encouraged board members to review this new section carefully, to ensure it accurately reflected the board's concerns and fairly described how the team had addressed them. He invited members to send back corrections or additions if anything seemed missing or mischaracterised.

Turning to Part K, Tony clarified that it too contained mostly material the board had seen before. The literature review included there, had been completed earlier in the trial and formed part of the original evaluation proposal. The bibliography was essentially a reference list, not intended to be comprehensive but limited to sources actually cited in the report. If board members noticed any omissions, works referenced in the text but not included in the bibliography, they were encouraged to flag them. However, Tony noted the bibliography was not intended to be a general reading list, so items not directly cited in the report shouldn't be added.

The most important section for board review in Part K, Tony said, was the glossary. He explained that the glossary focused exclusively on technical terms or domain-specific language that couldn't easily be found in a standard dictionary. The team had intentionally excluded definitions of commonplace or general words (e.g., "project," "policy") which anyone could look up online. Instead, the glossary was intended to support readers, especially non-experts, by defining key terms such as those related to biometric verification, privacy engineering, or age estimation metrics. Tony invited board members to review the glossary and suggest additions, particularly if they encountered a term in the report they believed would confuse lay readers and was not currently explained.

Iain confirmed that there would be another meeting timed around the publication of the final report.



Whether that next meeting would take place before or after public release remained uncertain, as there were a number of logistical and legal factors still to be resolved.

Iain addressed the question of whether board members would have the opportunity to make public comments or submit reflections as part of the report's release. He saw no reason why this couldn't happen and proposed that the team could host a dedicated page on the trial's website, where individual statements from board members could be published. Iain added that while it would be ideal if the board could agree on a shared, collective statement, he acknowledged that this might be difficult given the diversity of views and backgrounds represented.

Jon Rouse agreed that a unified board statement would be challenging. He noted that the board had already discussed this informally, and that two members were bound by NDAs (non-disclosure agreements) with their respective organisations. However, Jon believed there was still scope for the board to make some kind of high-level collective comment, such as an endorsement of the process or a general statement about the importance of evidence-based evaluation.

Colm asked when Part C of the report would be available. Tony explained that these parts of the report had not yet been shared with the board due to ongoing legal review, and it was still unclear whether they would be able to be shared before publication.

AOB

Iain said the next meeting was expected to be conducted in a hybrid format, with Tony and himself due back in Australia. Jon thanked Iain and Tony for their work, remarking that he had enjoyed reading the draft. He extended appreciation to the board for its robust and valuable input, which had strengthened the quality of the final product. He noted that the project team had formally documented board contributions in the report.

Jon also reminded members of an informal session with the Department scheduled for Thursday, where questions outside the project team's remit would be addressed. He encouraged members to send any questions for that meeting directly to him and to ensure they were appropriate for government representatives.

Tony Allen reiterated that this was not the final board meeting, and that a follow-up session would be held around publication. He thanked the board for its time, engagement, and constructive feedback, which had shaped the report in meaningful ways.

The meeting concluded at 17:29 AEST