Age Assurance Technology Trial

# PART A
# Main Report

*August 2025*

Funded by

Australian Government
Department of Infrastructure, Transport,
Regional Development, Communications, Sport and the Arts

Project by

Age Check
Certification Scheme

## Headline Findings

These are our headline findings. They are explained in slightly greater depth in Section A.3 and then explained in more detail for each of the technology types that we have explored.

1. **Age assurance can be done** in Australia privately, efficiently and effectively.

2. **No substantial technological limitations** preventing its implementation to meet policy goals.

3. Provider claims have been **independently validated** against the project's evaluation criteria.

4. A wide range of approaches exist, but there is **no one-size-fits-all solution** for all contexts.

5. We found a **dynamic, innovative and evolving** age assurance service sector.

6. We found robust, appropriate and **secure data handling practices.**

**7** Systems performed broadly **consistently across demographic groups**, including Indigenous populations.

**8** There is scope to **enhance usability, risk management** and system interoperability.

**9** Parental control tools can be effective but may **constrain children's digital participation** and evolving autonomy.

**10** Systems generally align with **cybersecurity best practice**, but vigilance is required.

**11** **Unnecessary data retention** may occur in apparent anticipation of future regulatory needs.

**12** Providers are aligning to **emerging international standards** around age assurance.

**Accessibility Statement:**

This report was produced in line with the accessibility guidelines found in the Australian Government Style Manual.

**Legal Clearance Statement:**

This report includes references to ISO standards through summarisation, referencing and reproduction of definitions only. While the material is not reproduced verbatim, ISO asserts copyright over its publications. For any further use or licensing queries, contact Standards Australia at: copyright@standards.org.au.

**Published By:**

Age Check Certification Scheme
Unit 321B Broadstone Mill, Broadstone Road
Stockport, United Kingdom, SK5 7DL

**www.accscheme.com**

ISBN 978-1-0681646-0-6

# Table of contents

## Analysis of Age Verification

C



## Analysis of Age Estimation

D



## Analysis of Age Inference

E

## Analysis of Successive Validation

F



## Analysis of Parental Control

G



## Analysis of Parental Consent

H

## Analysis of the Tech Stack

J



## Glossary, Bibliography and Literature Review

K

Age Assurance Technology Trial

1

# Introduction & Key Findings

## A.1 Preamble

**A.1.1** This document presents the official report of the Age Assurance Technology Trial, offering a comprehensive overview of its findings, methodologies and key observations. It brings together the conclusions and detailed analyses of the range of age assurance technologies assessed during the Trial. All of this was evaluated within the Australian context; more on this can be found in section A.4.

**A.1.2** The Trial is not an exercise in conformity assessment with Australian law or international standards. Our observations are not considered to be sufficient to meet legal thresholds for compliance checks, which properly remain a function of regulators. In examining safeguards, like privacy and security, whilst we have had regard to the relevant statutory provisions and guidance, we cannot provide the level and depth of analysis that would be required to provide any kind of clearance across all 48 of the Trial participants.

**A.1.3** While the report is neutral on policy matters and does not relate to a specific regulatory regime, this does not mean there are not additional complexities, operational challenges or requirements that will arise in particular policy or regulatory contexts. The report could be used as a basis for regulators to provide more detailed information relating to their remit, including compliance expectations and challenges.

## What to expect from the document

**A.1.4** The report offers a comprehensive evaluation of age assurance technologies, assessing their performance against a wide range of internationally recognised criteria. These include accuracy, interoperability, reliability, ease of use, minimisation of bias, protection of privacy and data security and readiness for deployment.

**A.1.5** It covers technologies from 48 Age Assurance providers, offering age verification, age estimation, age inference, successive validation, parental control and parental consent solutions. The report also examines how these technologies operate within the broader technology stack and how age assurance is integrated across different layers of the digital ecosystem.

## What this report is not

**A.1.6** The report is not a set of policy recommendations or endorsements for certain types of age assurance technology. That is not within scope of the Trial and not what the Trial set out to achieve. The report does not determine whether age assurance technology should be implemented or mandated in specific contexts. Instead, it provides factual and validated observations about the practical capabilities, limitations and potential of age assurance technologies based on structured evaluation processes. It is key to point out that any decision regarding the adoption or regulation of these technologies is ultimately a matter for policy makers and stakeholders beyond the scope of this Trial.

**A.1.7** The Trial is also not intended to test if every individual product works as claimed but rather to consider if the technologies as a whole work. We did not produce league tables, not least because not all vendors were evaluated in an identical way – as an example, to know if age estimation works, we only needed to confirm it did with a sample of age estimation solutions. Vendors will need to subject their solutions to independent audit and certification for an individual record.

## Vendor classification accuracy

**A.1.8** In the detailed analysis of the individual vendors tested, we have provided some data about their classification accuracy (i.e. how likely their system is to give a correct answer). It is not within the scope of the Trial to set the policy for acceptable levels of accuracy (that is a policy question) and we would urge caution about comparing the classification accuracy across different age assurance methodologies as the inputs, processes and measurement methodologies are not the same.

## The intended audience

**A.1.9** This report is designed for a broad range of stakeholders, including government agencies, industry participants, technology developers, civil society organisations and other parties who have an interest in age assurance technology within the Australian context. The report aims to inform these stakeholders about the current state of age assurance technologies, supporting evidence-based discussions and decision-making in this rapidly evolving field.

## Disclaimer on the status of analysis

**A.1.10** We would like to emphasise that while this report aims to reflect the best available insights at the time of publication, it is based on data and analysis conducted within a specific timeframe and set of conditions, within a technological landscape that is continually evolving. As such, the report should be understood as a snapshot of the state of the art rather than a final or exhaustive assessment of the long-term performance or policy implications of age assurance systems and technology solutions.


Calibration
Move hand back.

## A.2 Summary of the Trial

**A.2.1** The **Age Assurance Technology Trial** (the Trial) was commissioned by the Australian Government through the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts (DITRDCSA) to evaluate the effectiveness, reliability and privacy impacts of a range of age assurance technologies. The Trial was initiated in response to growing concerns about protecting children from online harms, including exposure to pornography, age-restricted services and social media-related risks.

**A.2.2** The Trial assessed the feasibility of various Age Assurance Systems – such as age estimation via AI, facial age analysis, parental consent and control mechanisms and identity document verification. These technologies were evaluated for their accuracy, usability and ability to safeguard personal data in real-world applications.

**A.2.3** The core aim was to understand **if age assurance can be done** without compromising Australian citizens privacy and security, as well as to inform consideration of best practice and potential regulatory approaches. This aligns with global efforts to create safer digital environments for young users, while balancing technological innovation with strong data protection and ethical standards.

Importantly, the Trial was conducted independently of DITRDCSA and regulators. It was not designed to make policy recommendations but to determine whether age assurance technologies are technically feasible and operationally deployable. It did not seek to decide whether such technologies should be deployed – this remains a policy and political decision. The Trial instead focused on whether the technology exists to support such decisions with confidence.

**A.2.4** The Trial team would like to thank the many technology providers, industry stakeholders, researchers and experts who participated in the Trial and contributed their time, insights and tools to this complex evaluation. Their collaboration and transparency were critical in enabling a comprehensive and balanced assessment of the current state of age assurance technology.

## A.3 Introduction to the Key Findings

**A.3.1** In this section, we explore the **12 findings of the Trial** in a little more detail before then expanding on each of them in the context of the individual technologies included within the Trial (age verification, age estimation, age inference, successive validation, parental control and parental consent).

1

### Age assurance can be done

Age assurance **can be done** in Australia – our analysis of age assurance systems in the context of Australia demonstrates how they can be private, robust and effective. There is a plethora of choice available for providers of age-restricted goods, content, services, venues or spaces to select the most appropriate systems for their use case with reference to emerging international standards for age assurance.

2

### No technological limitations

Our evaluation **did not reveal any substantial technological limitations** that would prevent age assurance systems being used in response to age-related eligibility requirements established by policy makers. We identified careful, critical thinking by providers on the development and deployment of age assurance systems, considering efficacy, privacy, data and security concerns. Some systems were easier for initial implementation and use than others, but the systems of all technology providers with a technology readiness level (TRL)[1] 7 or above were eventually capable of integration to a user journey.

1. *Discover more about Technology Readiness Assessments on p.48-49.*

3

## Independently validated

We found that the **practice statements** provided by age assurance providers with a TRL of 7 or above fairly reflected the technological capabilities of their products, processes or services (to the extent applicable to the Trial's evaluation criteria). Some of the practice statements provided have needed to be clarified or developed during the course of the Trial, but we observed that they offer a useful option for transparency of the capabilities of the available age assurance systems. Those with a TRL below 7 will need further analysis when their systems mature.

4

## No one size-fits-all

We found **a plethora of approaches** that fit different use cases in different ways, but we did not find a single **ubiquitous solution** that would suit all use cases, nor did we find solutions that were guaranteed to be effective in all deployments. The range of possibilities across the Trial participants demonstrate a rich and rapidly evolving range of services which can be tailored and effective depending on each specified context of use.

## 5 Evolving age assurance services

We found a **vibrant, creative and innovative** age assurance service sector with both technologically advanced and deployed solutions and a pipeline of new technologies transitioning from research to minimum viable product to testing and deployment stages indicating an evolving choice and future opportunities for developers. We found private-sector investment and opportunities for growth within the age assurance services sector.

## 6 Secure data handling practices

We found robust understanding of and internal policy decisions regarding the **handling of personal information** by Trial participants. The privacy policies and practice statements collated for the Trial demonstrate a strong commitment to privacy by design principles, with consideration of what data was to be collected, stored, shared and then disposed of. Separating age assurance services from those of relying parties was useful as Trial participants providing age assurance services more clearly only used data for the necessary and consented purpose of providing an age assurance result.

**7**

### Broad demographic consistency

The systems under test **performed broadly consistently across demographic groups** assessed and despite an acknowledged deficit in training age analysis systems with data about Indigenous populations, we found no substantial difference in the outcomes for First Nations and Torres Strait Islander Peoples and other multi-cultural communities using the age assurance systems. We found some systems performed better than others, but overall variances across race did not deviate by more than recognised tolerances.

**8**

### Scope for technological improvement

We found **opportunities for technological improvement** including improving ease of use for the average person and enhancing the management of risk in age assurance systems. This could include through one-way blind access to verification of government documents, enabling connection to data holder services (like digital wallets) or improving the handling of a child's digital footprint as examples.

**9**

## Limitations to parental control systems

The Trial found that both **parental control and consent systems can be done** and can be effective, but they serve different purposes. Parental control systems are pre-configured and ongoing but may fail to adapt to the evolving capacities of children including potential risks to their digital privacy as they grow and mature, particularly through adolescence. Parental consent mechanisms prompt active engagement between children and their parents at key decision points, potentially supporting informed access.

**10**

## Cybersecurity

We found that the systems were **generally secure** and consistent with information security standards, with developers actively addressing known attack vectors including AI-generated spoofing and forgeries. However, the **rapidly evolving threat environment** means that these systems – while presently fairly robust – cannot be considered infallible. Ongoing monitoring and improvement will help maintain their effectiveness over time. Similarly, continued attention to privacy compliance will support long-term trust and accountability.

**11**

### Unnecessary data retention

We found some concerning evidence that in the absence of specific guidance, service providers were apparently **over-anticipating the eventual needs of regulators** about providing personal information for future investigations. Some providers were found to be building tools to enable regulators, law enforcement or Coroners to retrace the actions taken by individuals to verify their age which could lead to **increased risk of privacy breaches** due to unnecessary and disproportionate collection and retention of data.

**12**

### Accreditation and certification

The standards-based approach adopted by the Trial, including through the **ISO/IEC 27566 Series**[2], the **IEEE 2089.1**[3] and the **ISO/IEC 25000**[4] series (the Product Quality Model) all provide a strong basis for the development of accreditation of conformity assessment and subsequent certification of individual age assurance providers in accordance with Australia's standards and conformance infrastructure.

2. *This Series of International Standards relates to Information security, cybersecurity and privacy protection – Age Assurance Systems. Part 1, referenced throughout this suite of documents. It is the Framework document, at Final Draft International Standard Stage. 27566-2 is the Technical approaches and guidance for implementation document and 27566-3 is the Comparison or Analysis document.*
3. *The IEEE 2089.1-2024 standard establishes a framework for the design, specification, evaluation, and deployment of age-verification systems. It was published internationally in 2024.*
4. *The series of standards ISO/IEC 25000, also known as SQuaRE (System and Software Quality Requirements and Evaluation), has the goal of creating a framework for the evaluation of software product quality.*


UNDER 18
No alcohol can be sold or supplied to anyone under 18. It's against the law.
It is against the law to sell or supply alcohol to, or to obtain alcohol on behalf of, a person under the age of 18 years.

Age Assurance Technology Trial

II

# Trial Context

## A.4 Developing the Trial in the Australian Context

**A.4.1** The Trial was developed specifically to reflect the unique regulatory, social, cultural and technological environment in Australia. While informed by international developments and global standards, the Trial was grounded in the needs, expectations and rights of Australian users – particularly children, young people, parents and guardians – as well as the responsibilities of Australian industry and government stakeholders.

**A.4.2** Australia presents a diverse digital landscape, with high levels of internet usage, a strong mobile-first culture and a growing ecosystem of digital platforms, services and infrastructure. The Trial took this diversity into account by including technologies deployed across a range of settings – urban and remote, commercial and public and across different age assurance methods, such as age verification, estimation, inference methods and parental involvement.

**A.4.3** Crucially, the Trial was designed to operate independently from government and regulators, while still aligning with the policy ambitions of entities such as the eSafety Commissioner, the Office of the Australian Information Commissioner (OAIC) and various legislative frameworks relating to online safety, privacy and child protection. It seeks to provide evidence-based insights to inform future decision-making by governments, platforms and service providers – without prescribing specific regulatory outcomes.

**A.4.4** Stakeholder input from Australian civil society, industry leaders, academic experts, First Nations people voices and child rights advocates helped shape the Trial's methodology, ethics approach and interpretation of findings. As such, the Trial offered a locally grounded, globally informed assessment of how age assurance technologies could be implemented in Australia in a way that is proportionate, privacy-respecting and inclusive of all users.

The Project Team respectfully acknowledge the Traditional Custodians of the lands and water in Australia where the Trial has been conducted and the team pay respects to Ancestors and Elders past, present and emerging.

We have been proud to support their communities through their inclusion and careful consideration throughout the design, implementation, communication and reporting of the Trial. We were particularly grateful for the input of Aboriginal and Torres Strait Islander Peoples in the evaluation activity of the Trial.

*Trial Cultural Advisor, John Fejo and his family.*

## A.5 Stakeholder Advisory Board

**A.5.1** The Trial established a Stakeholder Advisory Board (SAB) to create a forum for representatives of key stakeholder groups to provide input. Due to the independence of the Trial, the SAB was only advisory but provided the opportunity for a wide range of experts and individuals with an interest in age assurance technology and its applications to offer advice and challenge to the Trial team.

**A.5.2** Part B of this report contains a detailed analysis of the work of the SAB, including an analysis of the issues raised and how the Trial addressed them.

**A.5.3** The Stakeholder Advisory Board played a crucial role in ensuring the Trial's transparency, inclusivity and alignment with public interest. While the detailed Terms of Reference are available on the Trial's official website, the SAB's primary functions can be summarised as follows:

- Advisory Role: The SAB provided strategic advice to the Trial's team on effective stakeholder engagement and communication strategies, ensuring that the Trial's processes and outcomes were effectively disseminated and understood across diverse audiences.

- Composition: Chaired by Jon Rouse APM, Professor at AiLECS Labs, Monash University, the SAB comprised a diverse group of experts from various sectors, including government, regulatory bodies, civil society, industry and academia ensuring a comprehensive range of perspectives were considered throughout the Trial.

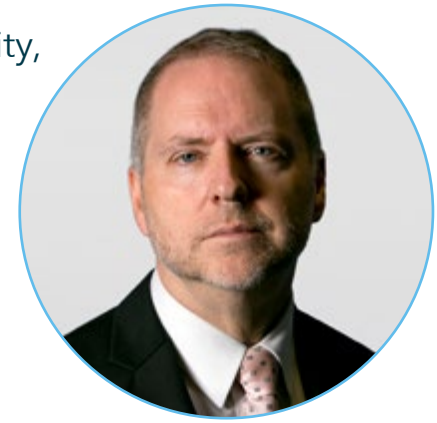

*Jon Rouse, Chair of the SAB*

- Independence: While the SAB offered guidance and recommendations, it operated independently of the evaluation activities. This separation ensured that the Trial's assessments remained impartial and unbiased, maintaining the integrity of the evaluation process.

- Transparency and Communication: To uphold the principles of openness and accountability, the SAB committed to publishing minutes of its meetings and disclosing its membership details on the Trial website. This practice fostered trust and allowed stakeholders to stay informed about the board's deliberations and contributions.

**A.5.4** By fulfilling these roles, the Stakeholder Advisory Board significantly contributed to the Trial's mission of evaluating age assurance technologies in a manner that was transparent, inclusive and reflective of the diverse interests and concerns of all stakeholders involved.

*Project Director, Tony Allen, presenting at one of the Stakeholder Engagement Events held during the Trial.*

Age Assurance Technology Trial

# Structure of the Report

## A.6 Parts of the Report


**A.6.1** The Trial Report consists of ten separate but interlinked reports, each examining a different aspect of age assurance technology. Together, they provide a comprehensive view of the Trial's findings, observations and methodology.

### Part A: Main Report

**A.6.2** Provides the overarching summary of findings across all technology types, including the Australian context, stakeholder engagement and forward-looking insights. It includes a summary of the analysis of age verification, estimation, inference, successive validation, parental consent, parental control and tech stack deployment.

### Part B: Methodology and Ethics

**A.6.3** Outlines the research design, data collection, analysis methods, ethical framework and risk management strategies that underpinned the Trial. Emphasises rigour, transparency and integrity in data handling and evaluation.

### Part C: Age Verification

**A.6.4** Explores technologies used to verify a user's age through access to authoritative records of their date of birth. Includes technical, privacy and security analysis, as well as provider practice statements and readiness assessments.

### Part D: Age Estimation

**A.6.5** Evaluates systems that estimate age based on biometric or behavioural features. Assesses functional and performance characteristics, data protection implications and contextual suitability for Australian users.

## Part E: Age Inference

**A.6.6** Focuses on systems that infer age from user behaviour, digital footprints, facts about them or contextual signals. Reviews effectiveness, ethical considerations and potential risks around profiling and inclusion.

## Part F: Successive Validation

**A.6.7** Examines layered or waterfall models where multiple age assurance methods are applied in sequence. Analyses their adaptability, escalation logic and proportionality in different risk settings.

## Part G: Parental Control

**A.6.8** Investigates pre-configured restrictions across devices, platforms or services. Analyses the effectiveness, overreach risks and alignment with children's evolving rights and capacities.

## Part H: Parental Consent

**A.6.9** Reviews mechanisms for obtaining guardian permission at the point of access. Considers legal integrity, usability and inclusion across diverse family and care arrangements.

## Part J: Tech Stack

**A.6.10** Assesses how age assurance is embedded across layers of the digital ecosystem – device, operating systems, network or application-level. Explores scalability, integration challenges and future potential.

## Part K: Glossary, Literature Review and Bibliography

**A.6.11** A comprehensive review of academic research, standards, laws, media and advocacy literature relevant to age assurance technologies in Australia and internationally. We have also included a comprehensive glossary of technical terms used throughout the report.

## A.7 Navigating the Reports

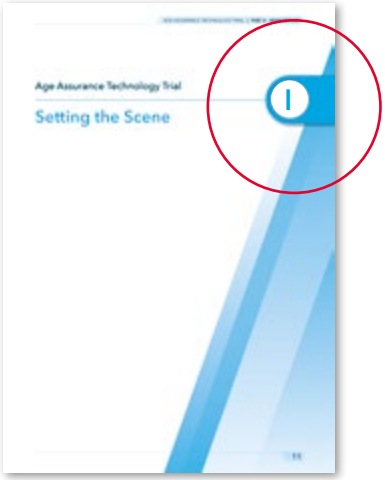


### | Chapters

**A.7.1** Chapters are numbered sequentially (I, II, III, etc.) and are easy to locate, each beginning with a dedicated chapter divider page for quick visual reference.

### | Subheadings

**A.7.2** Subheadings are numbered using a two-part code: The letter of the report e.g. A, followed by the subheading number e.g 1:

**A.÷ Developing the Age Assurance Technology Trial in the Australian Context**

the second subheading in Report A will appear as A.2 and so on.

### | Paragraphs

**A.7.3** Paragraphs are numbered using a combined three-part code: Report number e.g. A, Subheading number e.g. 5, and the paragraph number e.g. 3.

**A.1.3** Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua.

### | Infographics and data visuals

**A.7.4** Infographics and data visuals such as figures, charts and graphs are clearly labelled with a descriptor (e.g. Figure, Chart, Graph) followed by a code that corresponds to the relevant subheading – for example, Figure A.1.3 refers to, Report A, Subheading 1, the 3rd figure in that sub section. B.2.5 would link to Report B, Subheading 2 and be the 5th graphic to appear in that sub section. This ensures each visual element is tied directly to the supporting text.

## Callouts, spotlights & links

**A.7.5** In addition to the main text structure, the report includes a variety of call-out boxes designed to draw attention to specific types of content. Each box type features an identifying icon so it can be quickly recognised as you move through the document. These include:

### *Callout boxes*

Call Out boxes – used to highlight important definitions, clarifications or notes.

### *Spotlight boxes*

Spotlight boxes – used to showcase standout insights, examples or standout findings.

**Link boxes** – used to direct readers to additional sources, tools or related content elsewhere in the report, look out for these icons:

| Icon | Description | Icon | Description |
|---|---|---|---|
|  | Website |  | Interview |
|  | Practice Statement |  | Privacy policy |
|  | Test report & documents |  | Report cross reference to another report |
|  | News articles and segments | | |


27 CHILTERN
39 WODONGA
WANGARAT
POST

Age Assurance Technology Trial

B

# PART B
# Methodology & Ethics

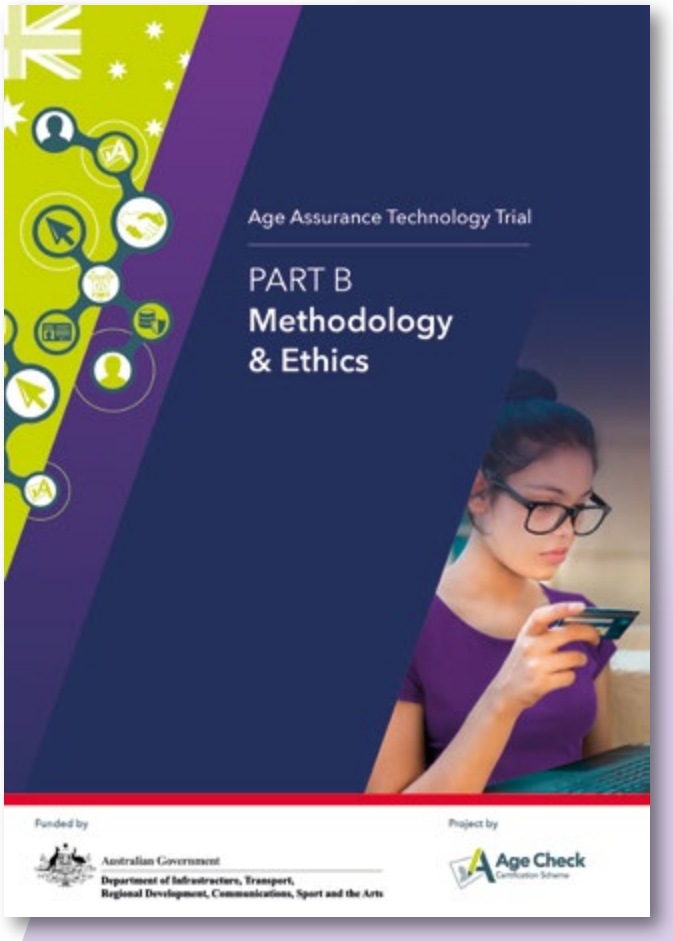


***See full report: Part B** – Methodology & Ethics*

## A.8 Our Core Principles

**A.8.1** These principles guided every stage of the Trial. They reflect the ethical standards we applied in assessing technologies and engaging participants.

### 1 Respect

We honour the inherent worth, autonomy and diverse backgrounds of all participants – particularly children – through culturally sensitive, age-appropriate engagement.

### 2 Transparency

We commit to open communication about the Trial's purpose, scope, methods and outcomes – empowering trust, understanding and public confidence.

### 3 Accountability

We uphold clear governance and independent oversight – enabling concerns to be raised, reviewed and acted on with integrity.

## 4 Fairness

We pursue equity and inclusivity – actively addressing bias to ensure impartial treatment and representation across all demographics.

## 5 Privacy

We safeguard participant privacy through data minimisation, secure handling and respectful collection aligned with human dignity.

## 6 Safeguard Children

We prioritise child safety and wellbeing – ensuring informed participation, adherence to rights and protection through every Trial phase.

## A.9 Introduction to Part B: Methodology and Ethics

**A.9.1** Recognising the increasing global and domestic demand for effective age assurance solutions, the Trial's research methodology was built on a foundation of strong ethical principles – respect, transparency, accountability, fairness, privacy and safeguarding children.

**A.9.2** To ensure robust and replicable results, the research methodology aligned with leading international standards and frameworks, including ISO/IEC 25040[5] (for quality evaluation), ISO/IEC FDIS 27566-1[6] (for age assurance systems) and IEEE 2089.1[7] (for online age checking systems). The methodology also considered unique Australian regulatory, cultural and social considerations, with specific attention to the participation of Aboriginal and Torres Strait Islander Peoples and alignment with Australia's privacy and online safety frameworks.

**A.9.3** The Age Assurance Technology Trial was an initiative led by the DITRDCSA to evaluate the effectiveness, reliability and privacy impacts of various age assurance technologies. The Trial was set up in response to growing concerns about protecting children from harmful content such as pornography and other online age-restricted services, as well as harms on social media. By evaluating a range of age assurance systems – including age analysis, AI-based estimation, parental consent/control and identity document verification – the Trial assessed the feasibility of these technologies in real-world applications, ensuring they were accurate, user-friendly and privacy preserving.

5. *All references to ISO/IEC 25040 throughout this report are referring to ISO/IEC 25040: 2024 Systems and software engineering - Systems and software Quality Requirements and Evaluation (SQuaRE) - Quality evaluation framework.*
6. *All references to ISO/IEC FDIS 27566-1 Standard throughout the suite of reports are referring to ISO/IEC FDIS 27566-1 - Information security, cybersecurity and privacy protection - Age assurance systems - Part 1: Framework.*
7. *All references to IEEE 2089.1 throughout the suite of reports are referring to IEEE 2089.12024 - IEEE Standard for Online Age Verification.*

**A.9.4** The Trial explored how different methods perform in verifying a user's age without compromising their personal data, helping Australia establish best practices and potential regulatory frameworks for age assurance. This effort aligned with global movements towards safer digital environments for young users, as Australia seeks to balance technological advancement with robust data protection and ethical standards.

**A.9.5** Ethical considerations were at the forefront of the Trial and this section of the report seeks to explore the Methodology and Ethics behind the Trial and its evaluation.

## A.10 Research and Evaluation Design

**A.10.1** The Trial was designed to address three key challenges identified in the evidence base: the reliance on theoretical evaluations, the absence of comprehensive technical assessments of age assurance solutions and the underrepresentation of Australian subpopulations in global studies. The Research and Evaluation Design of the Trial directly responded to these challenges and tailored it to Australia's unique regulatory, social and cultural context.

**A.10.2** The Trial's evaluation framework was structured around **four interdependent pillars**:

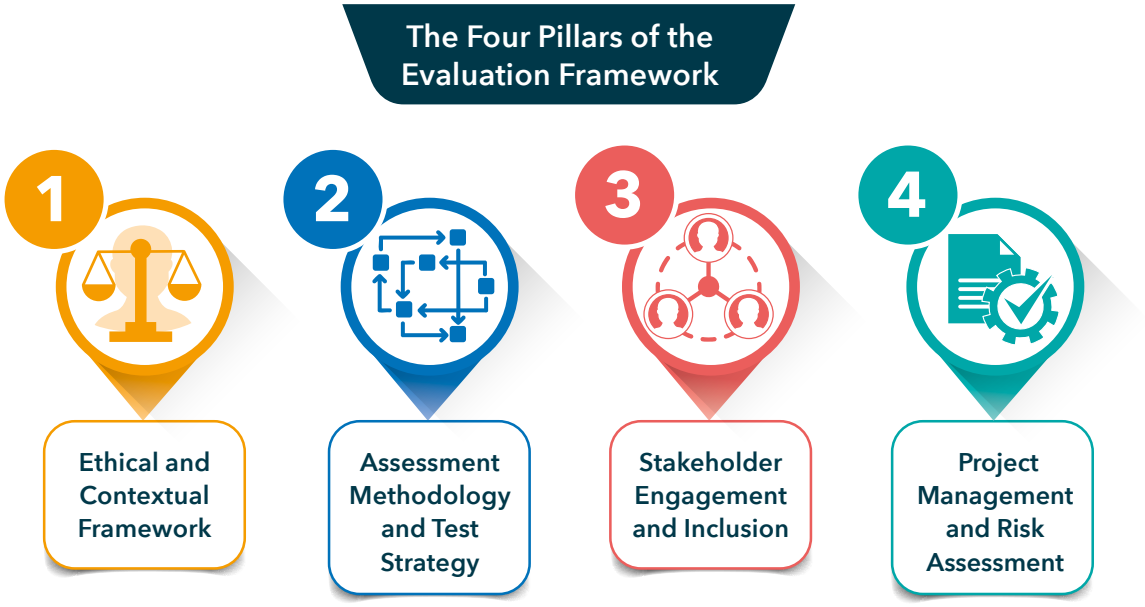


***Figure A.10.1*** *The Four Pillars of the Evaluation Framework*

1. **Ethical and contextual framework**
   The design of the Trial was grounded in explicit ethical principles – respect, transparency, accountability, fairness, privacy and safeguarding children – operationalised through the Data, Ethics and Impartiality Work Package 1[8]. This ensured that all Trial activities prioritised the protection of vulnerable users, aligned with Australian legal and cultural norms and reflected the diversity of the Australian community, with a focus on Aboriginal and Torres Strait Islander Peoples. Oversight by the project's Ethics Committee[9], including monthly meetings, reinforced the impartial and accountable conduct of the research.

*8. More information about Work Package 1 can be found:*

***Part B: Section VII*** *- Management of the project*

*9. More on the Ethics Committee can be found:*

***Part B: Section III*** *- Ethics*

2. **Assessment methodology and test strategy**
   The evaluation methodology was built around globally recognised standards (ISO/IEC 25040, ISO/IEC FDIS 27566-1, IEEE 2089.1) to ensure a rigorous, transparent and replicable approach. Technologies were evaluated against a set of clear and comprehensive criteria.

3. **Stakeholder engagement and inclusion**
   Central to the Trial's design was the inclusion of diverse stakeholders[10] – government, industry, academia, civil society and user groups, including children and parents. This approach sought to ensure the evaluation was not only technically robust but also socially informed and culturally respectful. In particular, the Trial's recruitment of participants and technology providers aimed to reflect the diversity of Australia's population, addressing previous gaps in representation in global studies.

4. **Project management and risk assessment**
   The Trial included rigorous risk management and quality control processes[11] to ensure the integrity of findings and to manage risks specific to the Australian environment, including cybersecurity, privacy and data protection concerns. These processes ensured that the Trial's outputs – these ten detailed reports – were delivered to the highest standards of quality and independence.

*10. More information about the Stakeholder Advisory Board can be found:*

***Part B: Section IV -*** *Peer Review and Stakeholder Engagement*

*11. More information about project management and control can be found:*

***Part B: Section VII -*** *Management of the Project*

| Recognised Standards | Key Criteria |
|---|---|
| **Accuracy** | How well the technology could detect a user's age. |
| **Interoperability** | How well the technology could be used across multiple online platforms. |
| **Reliability** | How consistently the technology could produce the same result. |
| **Ease of use** | How simple the technology was to operate, including how the system offered functionality appropriate to the capacity and age of a child or adult. |
| **Minimisation of bias** | How well the technology avoided racial or other bias, recognising that the complete elimination of bias was unattainable. |
| **Protection of privacy and data security** | How well the technology protected users' personal information. |
| **Human rights and accessibility protections** | Including people with disabilities, as well as applicable rights under the UN Convention on the Rights of the Child.[12] |
| **Circumvention** | Resistance to certain kinds of attacks including Biometric Presentation Attacks and Spoofing attacks. |
| **Technology Readiness Level (TRL)** | Ensuring the technology was sufficiently mature for meaningful testing. |

12. *The UNCRC is a legally binding agreement which outlines the fundamental rights of every child, regardless of their race, religion or abilities. Australia became a signatory to the UNCRC on 22 August 1990 and ratified it on 17 December 1990.*

## A.11 Ethical Considerations

**A.11.1** Ethical considerations were at the heart of the Trial and underpinned each phase of its research, evaluation and reporting, acknowledging the sensitive and complex nature of age assurance technologies, particularly their impact on children's rights, privacy and safety. Recognising this, the Trial adopted a set of guiding ethical principles:

1. **Respect**
   The principle of respect was central to the Trial, recognising the inherent dignity and autonomy of every individual. It required that participants were provided with clear and accessible information about the Trial and its purpose, that their decisions to participate were voluntary and free from coercion and that they were treated with sensitivity to their cultural, social and historical backgrounds. This also extended to special care for children, ensuring age-appropriate communication and safeguards to prevent undue pressure or harm.

2. **Transparency**
   Transparency underpinned the Trial's credibility and trustworthiness. The Trial team prioritised clear communication about the goals, processes and outcomes, making information accessible to stakeholders, including participants and the wider public. Transparency also involved clarifying the scope of the Trial – emphasising that it was not intended to develop or endorse any specific technology but to independently evaluate existing approaches to age assurance. Open data handling and clear accountability mechanisms further bolstered this principle.

3. **Accountability**
   Accountability was built into the Trial's governance through an independent Ethics Committee whose purpose was to scrutinise Trial activities and decision-making. Detailed analysis of the Ethics Committee mechanism can be found in Part B. All Trial members had a responsibility to uphold ethical standards, address any concerns raised by stakeholders and ensure prompt corrective action if ethical challenges emerged. Mechanisms for participants, including children, to raise concerns or withdraw consent reinforced this principle in practice.

4. **Fairness**
   Fairness guided the Trial's design and delivery to ensure inclusivity and equitable treatment of all participants and technology providers. The Trial actively worked to identify and address any risks of bias – such as demographic, racial or gender bias – in both the technologies under review and the evaluation methods themselves. Fairness also meant ensuring that diverse Australian populations, including Aboriginal and Torres Strait Islander Peoples, could participate meaningfully and share their unique perspectives.

5. **Privacy**
   Privacy was recognised as a fundamental human right and a key pillar of ethical research practice. The Trial followed principles of data minimisation – collecting only what was strictly necessary – and implemented strong data security measures to protect personal information. Privacy considerations were not only about data handling but also about upholding participants' dignity and autonomy, especially in the face of intrusive or sensitive data collection methods such as biometric analysis.

6. **Safeguarding children**
   As children's safety and wellbeing were central to the Trial's mission, child safeguarding was a dedicated ethical principle. The Trial drew on national frameworks, including the Australian Government's National Principles for Child Safe Organisations[13], to ensure children were protected and empowered throughout their involvement. This meant prioritising children's rights, responding swiftly to any concerns and ensuring that those working with children were properly supported and trained.

## | Operationalising the principles

**A.11.2** The Trial ensured that the chosen principles were not just theoretical – they were actively integrated into each stage of the project. What follows is how these principles were operationalised:

| Principle | Operationalisation Summary |
|---|---|
| **Data Protection and Privacy** | Data handling followed the Australian Privacy Principles[14], using data minimisation, strong security, and privacy-by-design approaches. |
| **Inclusion of First Nations Peoples** | The framework ensured respectful inclusion and engagement of First Nations peoples, aligning with cultural commitments. |
| **Child Safeguarding and Rights** | The Trial prioritised children's rights under the UNCRC, ensuring their best interests were central, especially online. |

13. *In February 2019, the National Principles for Child Safe Organisations were endorsed by all state and territory governments and the Australian Government. The principles aim to provide a nationally consistent approach to creating organisational cultures that foster child safety and wellbeing.*
14. *The Australian Privacy Principles are a set of 13 Principles that are the cornerstone of the privacy protection framework in the Privacy Act 1988.*

| | |
|---|---|
| **Impartiality and Accountability** | An independent Ethics Committee and Impartiality Panel oversaw the Trial. Transparency in stakeholder roles reduced bias. |
| **Minimising Bias and Promoting Equity** | Technologies were assessed for fairness across race, gender and age, with a focus on equitable treatment of all users. |
| **Transparency and Open Data** | The Trial shared methods, conflicts of interest and findings where possible to support public trust and policy use. |
| **User-Centric and Rights-Respecting Design** | Technologies were evaluated for usability, accessibility and respect for dignity – especially for children and marginalised groups. |
| **No Endorsement or Policy Mandate** | The Trial provided neutral, evidence-based insights without promoting specific technologies or policy outcomes. |

## A.12 School Field Trials

**A.12.1** Geographic diversity: Schools were located across five Australian states and territories – Western Australia, Australian Capital Territory, Queensland, Victoria and the Northern Territory.

**A.12.2** School types: Included a mix of government and independent schools, covering co-educational settings in both urban and regional areas.

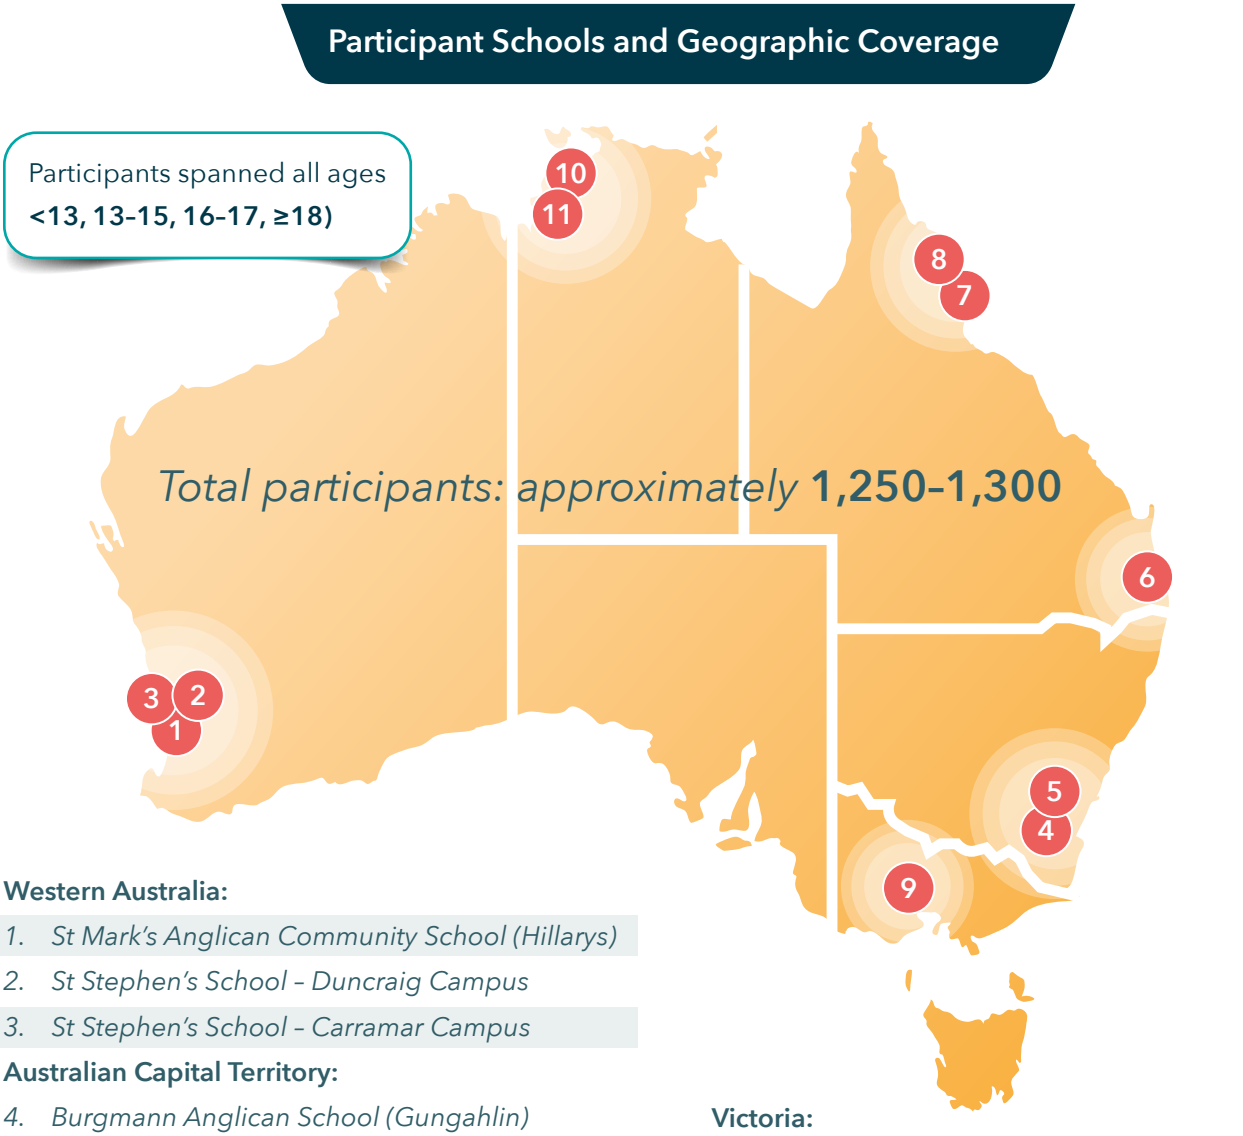


**Western Australia:**

1. *St Mark's Anglican Community School (Hillarys)*
2. *St Stephen's School – Duncraig Campus*
3. *St Stephen's School – Carramar Campus*

**Australian Capital Territory:**

4. *Burgmann Anglican School (Gungahlin)*
5. *John Paul II College (Nicholls)*

**Queensland:**

6. *Parklands Christian College (Park Ridge)*
7. *Radiant Life College (East Innisfail)*
8. *AFL Academy (Cairns)*

**Victoria:**

9. *Kyneton High School (Kyneton)*

**Northern Territory:**

10. *Nightcliff Middle School (Nightcliff)*
11. *Good Shepherd Lutheran College (Howard Springs)*

***Figure A.12.1*** *Participant Schools and Demographic Coverage*

## A.13 Independent Validation

**A.13.1** As a part of ensuring confidence and credibility of the Trial, the approach, methodology and testing was subject to independent validation by Prof, Toby Walsh.

### Professor Toby Walsh Validation Statement

The proposal does a very good job of scoping out a trial to evaluate the effectiveness of age assurance technologies in Australia. The proposal is especially strong with respect to: (1) the comprehensive evaluation criteria; (2) addressing evidence gaps; (3) explicit ethical principles; (4) a standards-based approach; (5) a commitment to open scientific reporting; (6) and recognition of children's rights.

I identified a few minor issues in the initial draft where I recommended some attention such as addressing combinations of age assurance methods, sample sizes for minority groups, and child friendly project outputs (given this group will be directly impacted by age assurance).

All these issues have been adequately addressed in the final evaluation proposal.

In summary, the trial has been scoped out well and looks set to deliver high quality results on the capabilities of age assurance technologies. I commend the work that the team has put in so far.

**Professor Toby Walsh**
FAA FAAAI FAAAS FACM FEurAI FRSN

*Scientia Professor of Artificial Intelligence*
*University of New South Wales*

## A.14 Technology Readiness Assessments

**A.14.1** When submitting an Expression of Interest at the start of the Trial's process, prospective participants were asked to state the Technology Readiness Level (TRL) of their solution. TRLs are based on a scale from 1 to 9 with 9 being the most mature technology.

**A.14.2** The New South Wales (NSW) Government's Invest NSW initiative provides a tool to calculate the TRL level for a technology system. The table of TRL levels is set out on the next page.

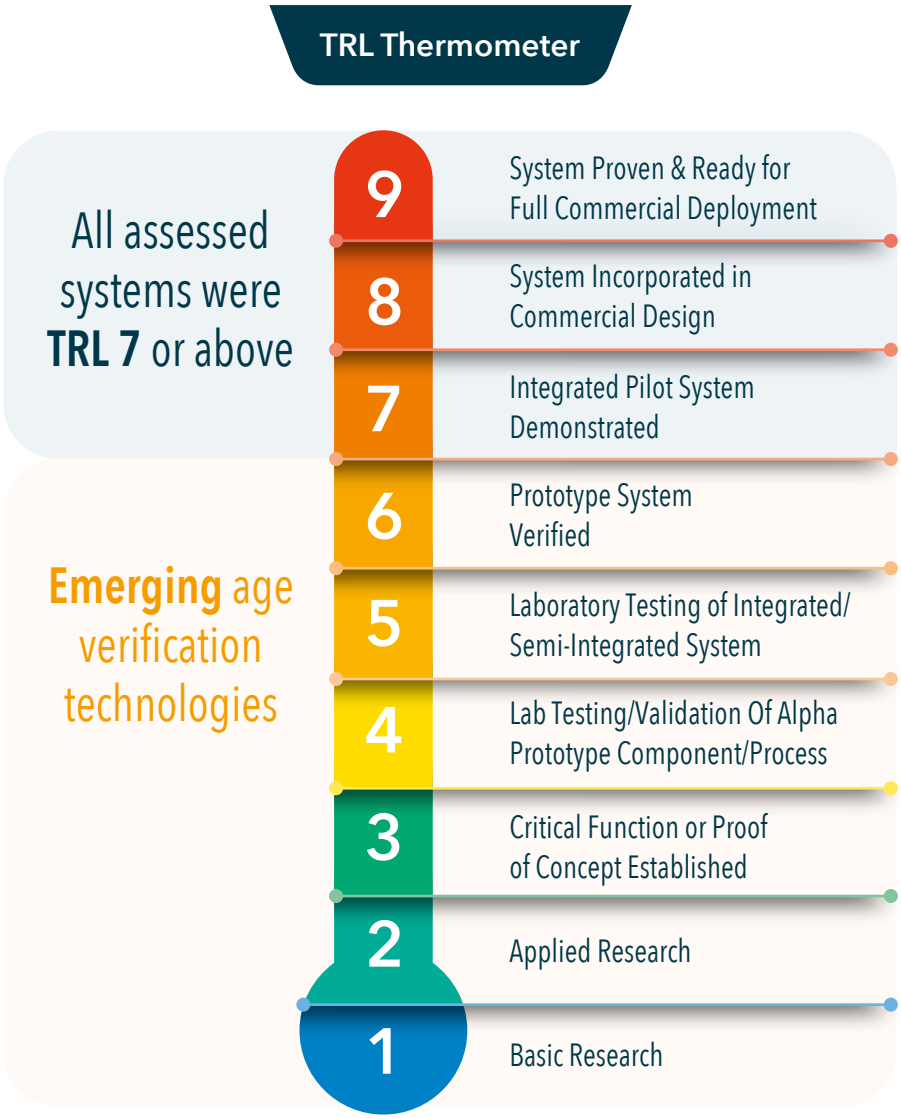


***Figure A.14.1*** *TRL Thermometer*

**TRL 9** **System Proven and Ready for Full Commercial Deployment:** Actual system proven through successful operation in an operating environment, ready for full commercial deployment.

**TRL 8** **System Incorporated in Commercial Design**: Actual system/process completed and qualified through test and demonstration (pre-commercial demonstration).

**TRL 7** **Integrated Pilot System Demonstrated:** System/process prototype demonstration in an operational environment (integrated pilot system level).

**TRL 6** **Prototype System Verified:** System/process prototype demonstration in an operational environment (beta prototype system level).

**TRL 5** **Laboratory Testing of Integrated/Semi-Integrated System**: System Component and/or process validation is achieved in a relevant environment.

**TRL 4** **Lab Testing/Validation Of Alpha Prototype Component/ Process:** Design, development and lab testing of components/processes. Results provide evidence that performance targets may be attainable based on projected or modelled systems.

**TRL 3** **Critical Function or Proof of Concept Established:** Applied research advances and early stage development begins. Studies and laboratory measurements validate the analytical predictions of separate elements of the technology.

**TRL 2** **Applied Research:** Initial practical applications are identified. Potential of material or process to solve a problem, satisfy a need or find application is confirmed.

**TRL 1** **Basic Research:** Initial scientific research has been conducted. Principles are qualitatively postulated and observed. Focus is on new discovery rather than applications.

## A.15 Participants in the Trial























netsweeper


ICG
ONE CLICK
GROUP


Opale.io


persona


PrivateID®


PRIVATELY


PRIVO®
Privacy • Permission • TRUST


Qoria


R2LABS


RIGHTCROWD
EVERYONE ON SITE IS RIGHT


Rigr


sedicii


ShareRing


SHAYYPE


TikTok


TruAnon
Digital Identity


TRUST
ELEVATE


TrustStamp


unissey


VerifyChain


verifymy


YOTI


UNDER 18
Minors are al
in this area a
as they are
responsible

Age Assurance Technology Trial

C

# PART C
# Age Verification


*See full report: **Part C** – Age Verification*

## A.16 Findings on Age Verification

**A.16.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the topic of age verification.

1. Age verification **can be done** in Australia privately, efficiently and effectively.

2. **No substantial technological limitations** preventing its implementation in the Australian context.

3. Providers' claims were independently assessed; are **accurate and reflective of real-world system performance**.

4. **There is no single solution to age verification**; a range of valid models exist, shaped by different contexts, needs and expectations.

5. **The age verification sector in Australia is dynamic and innovative** with active development and communication of verified age information.

**6** We found **robust, privacy-focused and secure** data handling practices.

**7** Age verification systems performed broadly **consistently across demographic groups**, including Indigenous populations.

**8** **Opportunities exist to enhance risk management and system capability**, especially regarding real-time detection of lost or stolen documents.

**9** **Cybersecurity practices were strong across the sector** with various threats addressed; continuous monitoring remains essential.

## A.17 What is Age Verification

**A.17.1** Age verification is an age assurance method based on calculating the difference between a verified year or date of birth of an individual and a subsequent date.

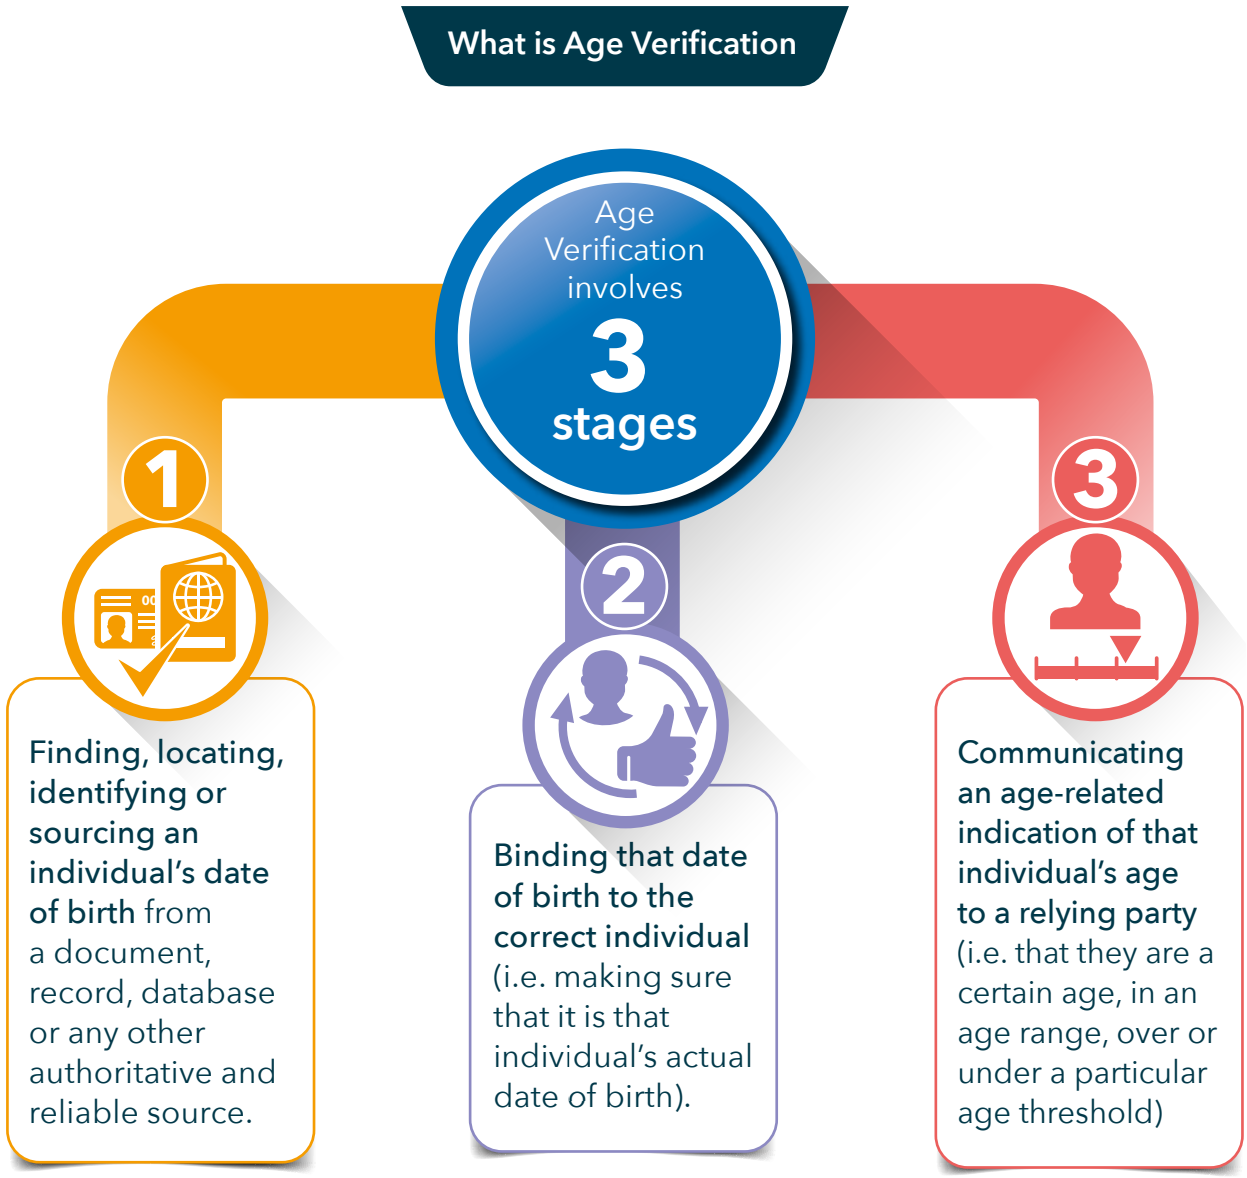


***Figure A.17.1*** *What is Age Verification*

## A.18 Introduction to Part C: Age Verification

**A.18.1** Part C of the Age Assurance Technology Trial focuses specifically on age verification – the process of determining an individual's age by referencing a verified date of birth and calculating their age from that known data point. Age verification represents the most direct and high-assurance form of age assurance and is already in widespread use across many regulated industries.

**A.18.2** This section evaluates how age verification systems perform in the Australian context in terms of technical feasibility, reliability, inclusivity, privacy preservation and security and how they align with emerging international standards such as ISO/IEC FDIS 27566-1 and IEEE 2089.1. The technologies assessed include solutions using official identity documents, secure databases, customer account information and verified credentials, often supported by cryptographic or biometric binding techniques.


Get boarding pass
NEXT
Welcome
QUEST

## A.19 Summary of Age Verification

**A.19.1** Age verification is a high-assurance method of age assurance that determines whether an individual is above or below a specific age threshold by comparing a verified date of birth (DOB) with a point in time – typically the current date.

**A.19.2** Age verification can be done in Australia and is widely used in existing deployments. Australia has a robust foundation for verifying dates of birth, with authoritative government sources, consistent data management practices and secure access to identity records and documents. This framework supports reliable issuance and validation of birth date evidence across services, enhancing trust and integrity in age assurance and identity verification processes. Notwithstanding the robust foundation, there may be cultural and education barriers for age verification.

**A.19.3** Our evaluation did not reveal any substantial technological limitations to the implementation of age verification technologies in Australia. Providers demonstrated compliance with recognised standards and deployed responsible, privacy-conscious approaches, incorporating strong data protection and security. The alignment of these technologies with emerging policy frameworks supports the deployment of effective and trustworthy systems for verifying age based on date of birth.

**A.19.4** Privacy by-design and data minimisation were consistently observed across the participating providers. In most cases, systems were designed to avoid long-term storage of full identity or biometric information. Instead, they returned binary age outcomes or anonymised session tokens that could be reused across services. Several providers supported integration with privacy-focused digital wallets, allowing verified age credentials to be stored locally and reused with explicit consent. These approaches reflect close alignment with the draft ISO/IEC FDIS 27566-1 standard and demonstrate strong readiness for future conformity assessment and certification.

**A.19.5** Demographic consistency was a key area of focus. The Trial found that systems generally performed well across diverse user groups, including First Nations and Torres Strait Islander Peoples. Some providers also made proactive efforts to include users who lack conventional identity documents, by supporting community-issued records or in-person onboarding processes. Nevertheless, gaps persist in remote and very remote communities where digital exclusion and lack of foundational credentials continue to limit access. While technically feasible, exact age verification for children is constrained by limited access to hard data; government-backed blind-access APIs to records (e.g., schools, healthcare) may be needed to improve precision.

**A.19.6** Sector-specific tailoring was a notable strength across the Trial. Different sectors – such as gambling, adult content, education, retail and access to physical venues – require different levels of assurance, privacy and friction. Providers demonstrated flexibility and configurability in their systems, allowing them to be adapted to both high-risk and privacy-sensitive contexts. In high-assurance sectors such as gambling, providers incorporated document checks, facial biometrics and record-matching. In privacy-sensitive sectors like adult entertainment, the focus was on anonymous, one-time checks that avoided any persistent identity linkage.

**A.19.7** Security and fraud resilience were also strong. Most providers operated ISO/IEC 27001-compliant systems, with encryption, multi-factor authentication and tamper detection. Biometric liveness checks were commonly implemented and aligned with ISO/IEC 30107 (presentation attack detection) standards, helping to guard against spoofing and deepfake risks. Systems were also generally effective at identifying document forgeries, including AI-generated fakes. However, several providers lacked the ability to check documents against live government databases to determine whether a document had been reported lost or stolen. The evaluation found that security against injection attacks – where malicious code or media bypasses the biometric capture process – is improving but still emerging.

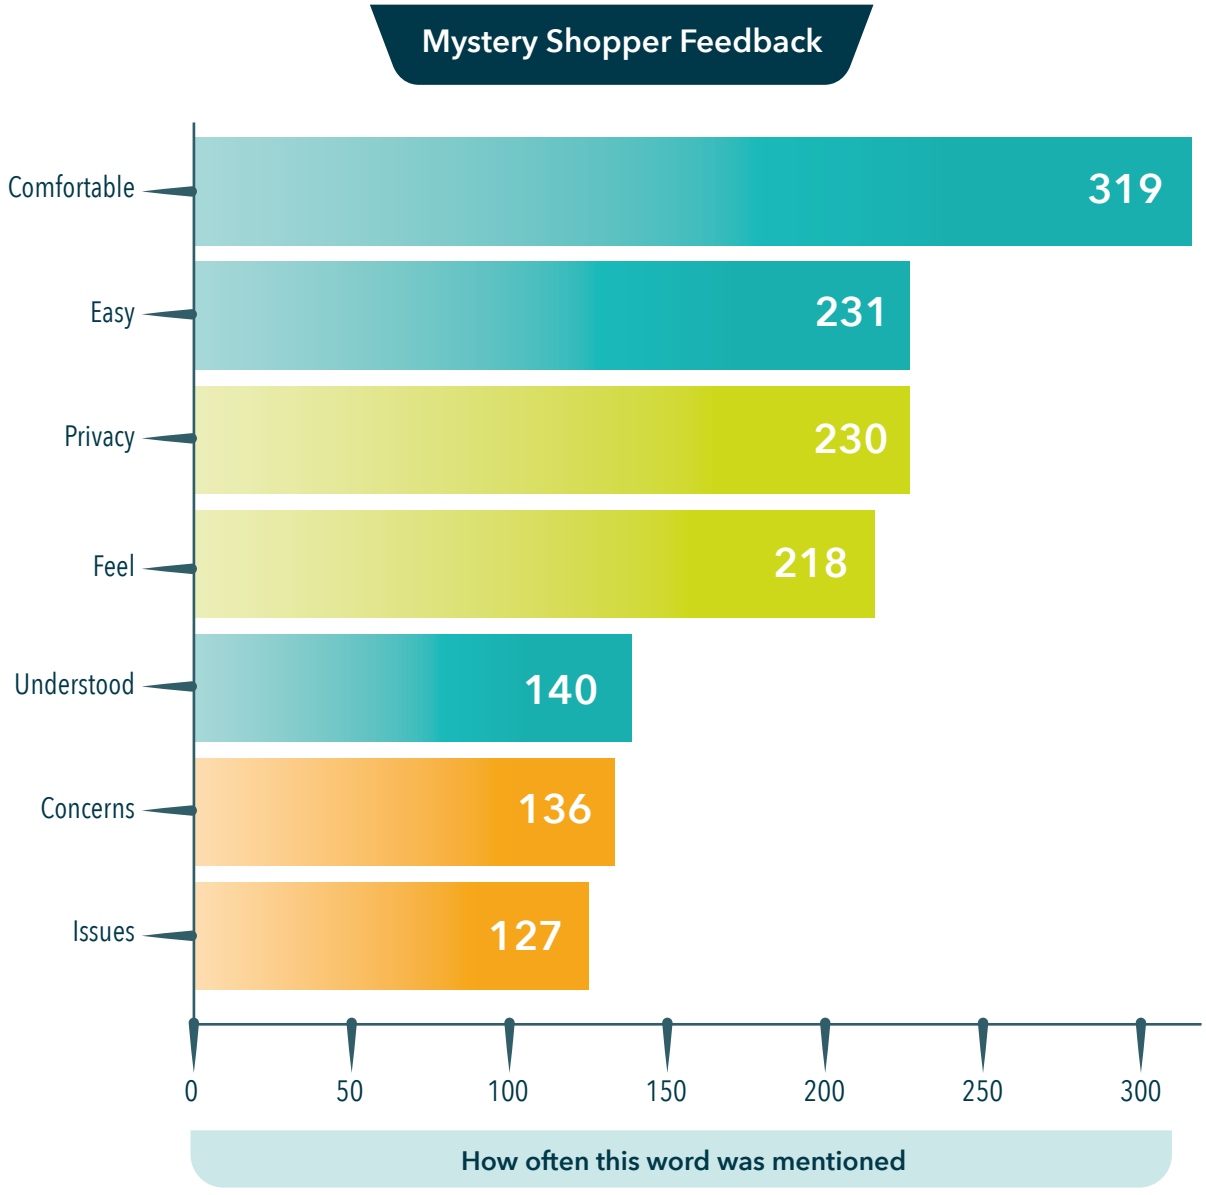


***Figure A.19.1*** *Mystery Shopper Feedback*

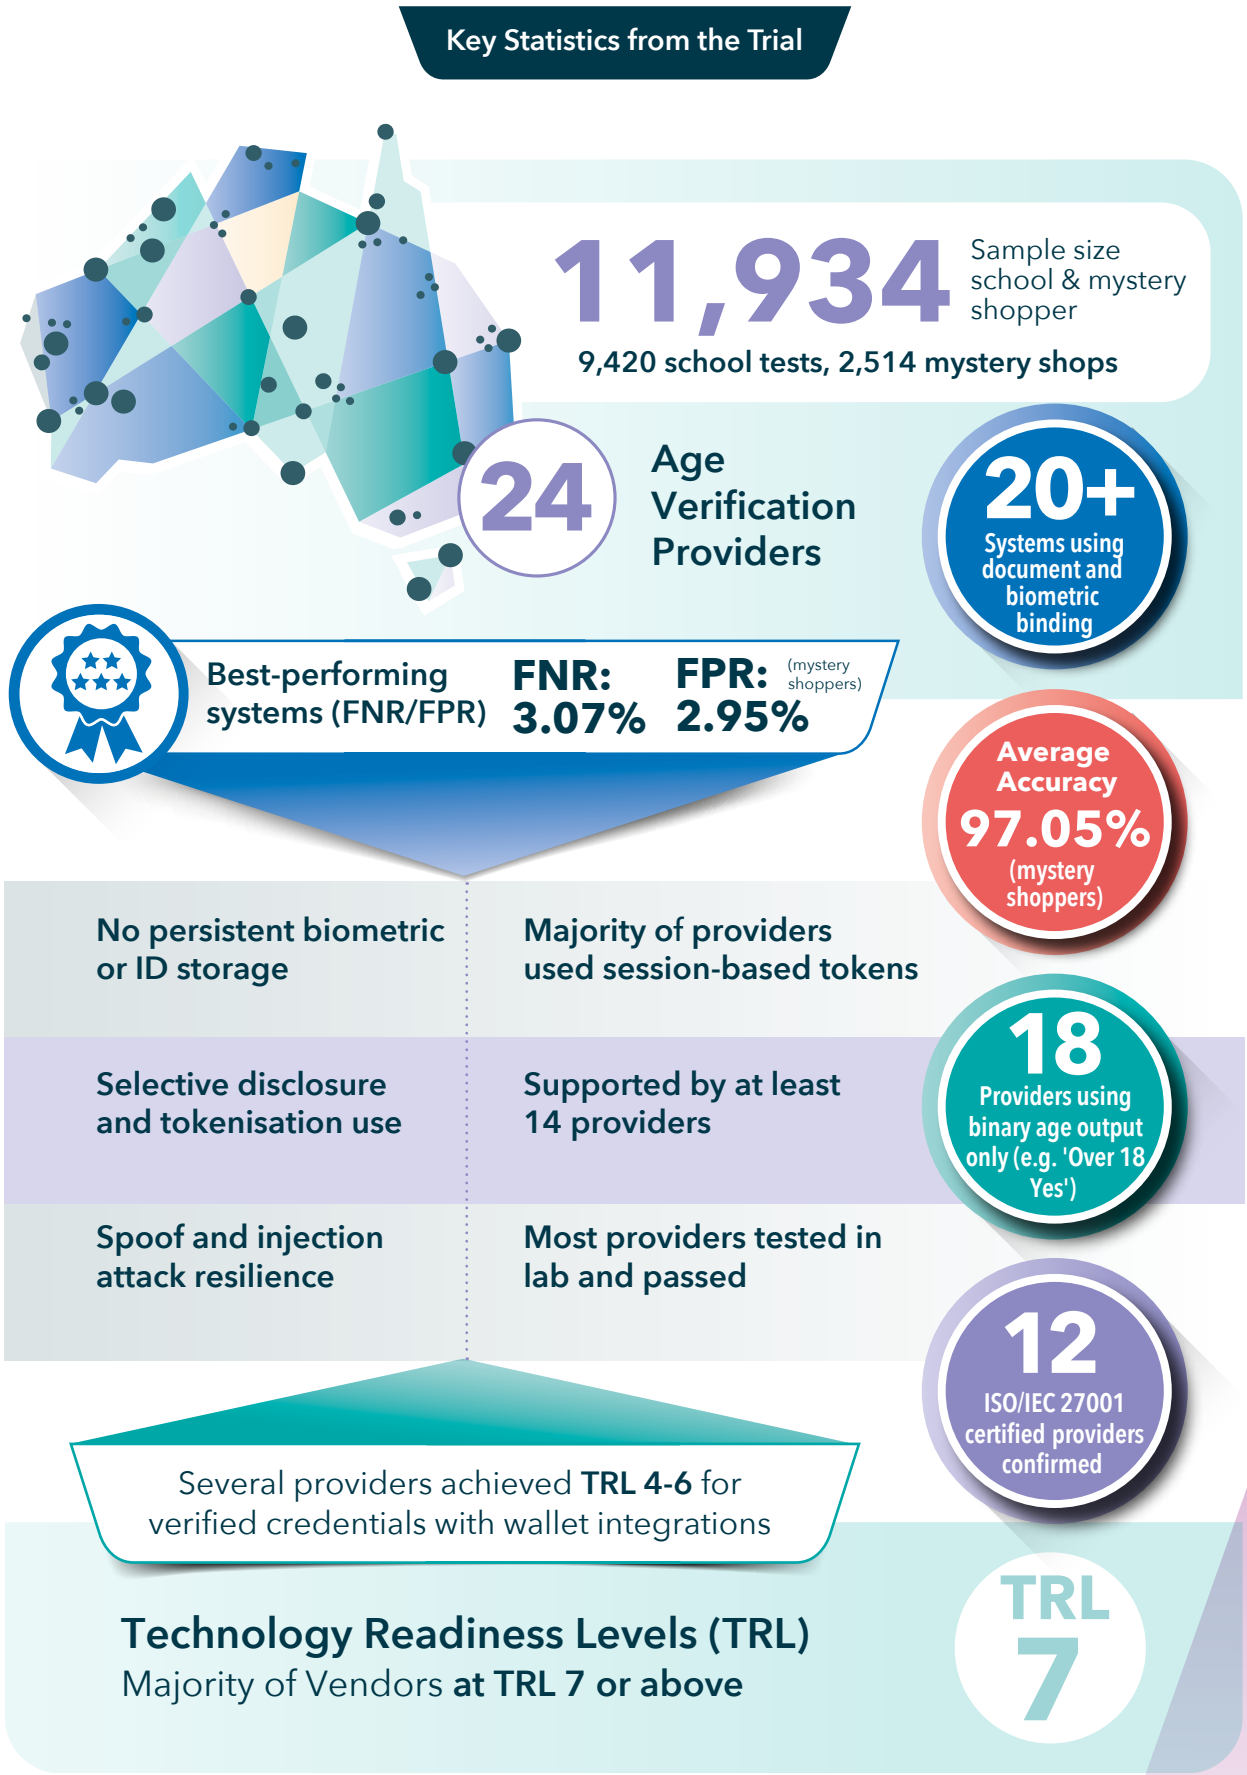


***Figure A.19.2*** *Key Statistics from the Trial on Age Verification*

**A.19.8** While most providers followed clear data minimisation practices, the Trial identified a concerning trend among a minority of providers toward over-preparing for investigatory or forensic requests. This included the retention of full biometric or document data for all users, even when such retention was not required or requested. While these practices may be motivated by a desire to assist regulators or coroners in rare and serious circumstances, they carry significant privacy risks and require clearer regulatory guidance to ensure proportionality.

**A.19.9** The age verification sector in Australia is highly dynamic and marked by innovation. Providers are actively developing new ways to verify age while reducing user friction and improving inclusivity. These include privacy-preserving cryptographic methods, reusable verified credentials, integration with mobile digital wallets and emerging support for blind-verification APIs that enable checks against government-held data without exposing user identity. Although some of these models remain at lower technology readiness levels, they signal a shift toward greater interoperability, reusability and user control.

**A.19.10** In summary, age verification is a technically mature, privacy-conscious and inclusive method of age assurance. When implemented with strong safeguards, ethical oversight and adherence to international standards, it offers a viable and trustworthy solution for protecting children and enforcing age-based access controls in Australia's digital environment. Continued investment in inclusion, standardisation and user-centric innovation will help ensure that age verification systems remain fair, effective and widely accepted.

## A.20 Who Participated in the Trial of Age Verification Technology





























## A.21 Observations About Age Verification

A.21.1 **Age verification based on calculating age from a verified date of birth is technically and operationally feasible in Australia** and can be implemented privately, securely and effectively in line with emerging international standards.

A.21.2 **No substantial technological limitations were identified** that would prevent age verification systems from meeting policy or regulatory requirements in the Australian context.

A.21.3 Providers' claims about age verification capabilities were independently assessed and **found to be accurate and reflective of real-world system performance**, including in lab and field testing.

A.21.4 **There is no single solution to age verification;** a range of valid implementation models exist, shaped by sector-specific needs, risk profiles, data availability and privacy expectations.

A.21.5 **The age verification sector in Australia is dynamic and innovative,** with providers actively developing more efficient and user-centric ways to retrieve, bind and communicate verified age information.

A.21.6 **Most providers implemented robust, privacy-focused data handling practices,** securely binding DOB to individuals, minimising retention and returning binary age signals (e.g., "Over 18") – though some configurations retained more data than strictly necessary.

A.21.7 **Age verification systems performed broadly consistently across demographic groups,** including First Nations and Torres Strait Islander Peoples. Some providers took proactive steps to include users without conventional identity documents.

**A.21.8 Opportunities exist to enhance risk management and system capability,** especially regarding real-time detection of lost or stolen documents and improved access to authoritative government data through privacy-preserving APIs.

**A.21.9 Cybersecurity practices were strong across the sector,** with many providers addressing threats such as biometric spoofing, AI-generated forgeries and injection attacks. However, continuous monitoring and resilience updates remain essential.


IL SSING
ROS WAY
2
TRACKS
STOP
ON RED
SIGNAL

Age Assurance Technology Trial

D

# PART D
# Age Estimation

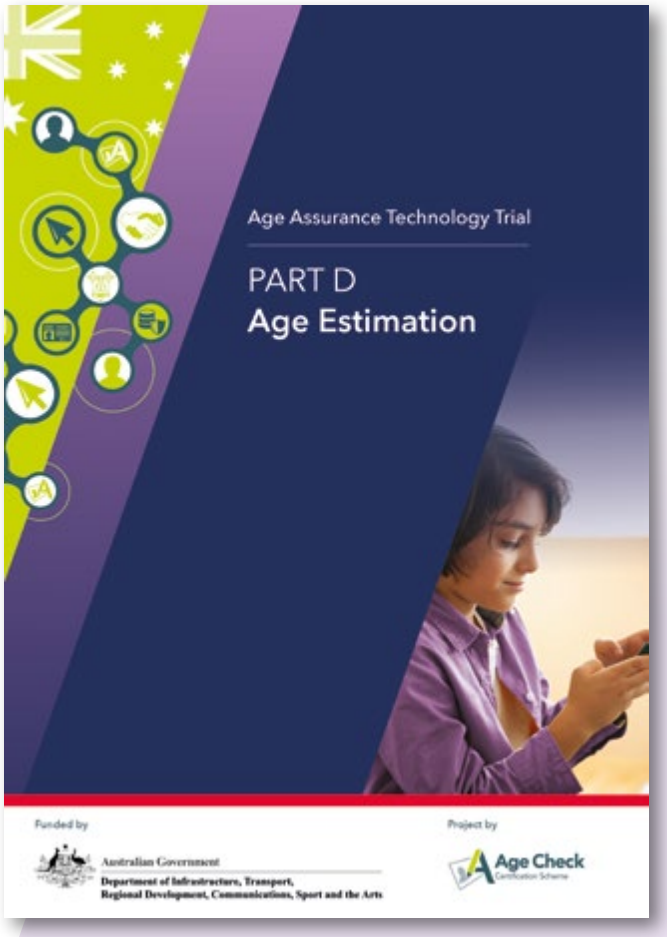


***See full report: Part D** – Age Estimation*

## A.22 Findings on Age Estimation

**A.22.1** These are our headline findings. In line with the overall findings of the Trial, they relate specifically to the topic of age estimation.

1. **Age estimation can be done** in Australia; is being deployed effectively and across multiple sectors.

2. Most **systems are technically deployable** across standard devices and environments, though edge-case limitations remain.

3. Provider **claims regarding performance were generally substantiated** through independent evaluation; some early-stage systems lacked complete transparency.

4. There is no single approach to **age estimation; must be configured to context**.

5. The age estimation sector in Australia is **dynamic, innovative and responsive to privacy and fairness challenges**; providers are iterating rapidly.

6 **Demographic performance consistency is improving**; underrepresentation of Indigenous populations remains a challenge that vendors are beginning to address.

7 **Accuracy varies in suboptimal conditions**, highlighting the need for robustness improvements.

8 **Vendors are actively mitigating adversarial threats**, including spoofing and injection attacks.

9 **Age estimation decisions remained based on real-time**, independently derived evidence – not on self-declared, inferred or parental assertions.

10 Providers are aligning with emerging international standards and **demonstrating readiness for certification**, meeting expectations set out in ISO/IEC FDIS 27566-1.

## A.23 What is Age Estimation

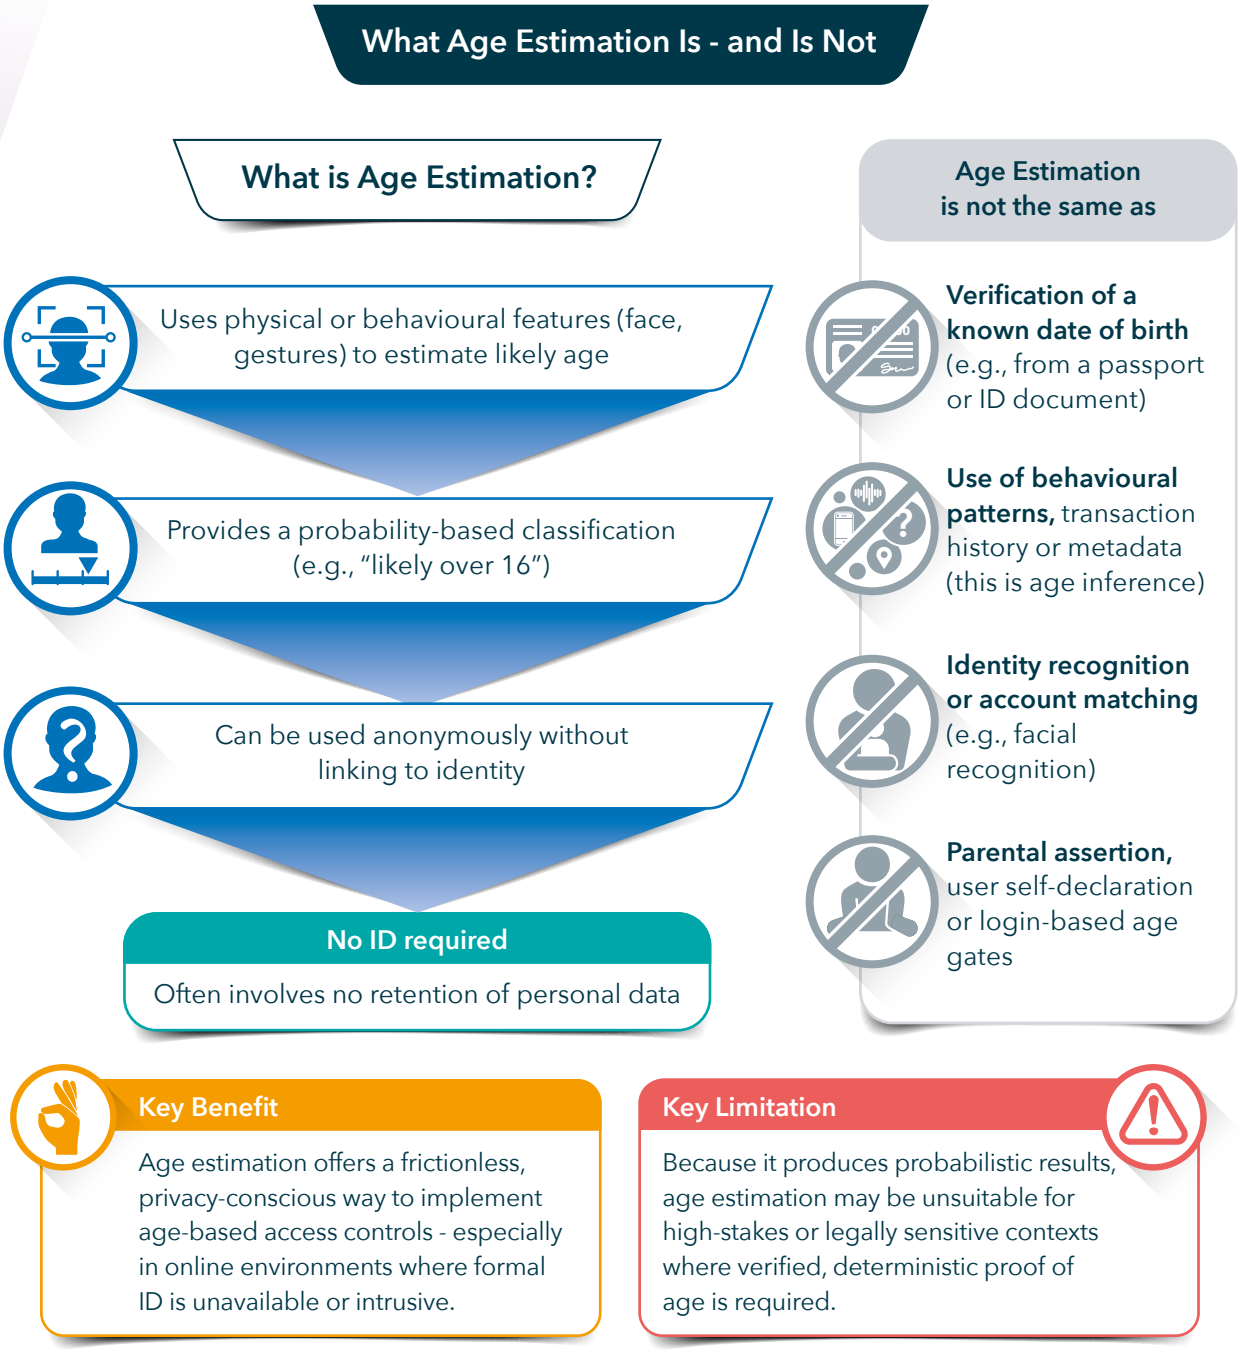


***Figure A.23.1*** *What Age Estimation Is – and Is Not*

## A.24 Introduction to Part D: Age Estimation

**A.24.1** Part D of the Age Assurance Technology Trial focuses specifically on age estimation – a method of determining an individual's likely age or age range by analysing physical or behavioural characteristics using artificial intelligence or machine learning models. Unlike age verification, which relies on known and validated dates of birth, age estimation applies biometric or statistical techniques (such as facial analysis, voice modelling or motion pattern recognition) to predict age without the need for formal identity documents.

**A.24.2** This section evaluates how age estimation systems perform in the Australian context, including their technical feasibility, statistical accuracy, demographic fairness, privacy protection and resistance to manipulation. The Trial assesses alignment with relevant international standards – particularly ISO/IEC FDIS 27566-1, which provides a functional and privacy framework for age assurance systems and IEEE 2089.1, which outlines performance expectations for demographic consistency.

**A.24.3** In Part D of the report, we present our findings on age estimation systems, including their accuracy across age thresholds, performance across demographic groups, ability to operate in low-friction environments and effectiveness when used alongside other age assurance methods. This analysis supports the development of evidence-based standards, best practices and potential pathways for certification in Australia's evolving digital safety landscape.

## A.25 Summary of Age Estimation

A.25.1 Age estimation is a method of estimating a user's likely age based on observable characteristics such as facial features, voice or behavioural patterns. Unlike age verification, which relies on official identity documents, age estimation uses statistical models to estimate age without identifying the user. It is increasingly used to enforce age-based access controls in digital and in-person environments where document-based identity is unavailable, inappropriate or unnecessary.

A.25.2 As part of the Trial, age estimation technologies were evaluated for their accuracy, security, inclusivity, usability and alignment with emerging international standards. The evaluation focused on high-readiness systems (Technology Readiness Level 7 or above) and included, as appropriate, structured technical testing, school-based trials, mystery shopper deployments, practice statement reviews and vendor interviews. Systems were assessed against key benchmarks such as ISO/IEC FDIS 27566-1 (age assurance requirements), IEEE 2089.1 (interoperability and assurance rules), ISO/IEC 27001 (information security) and ISO/IEC 25010 (software quality attributes).

A.25.3 The Trial confirmed that age estimation can be deployed effectively in the Australian context. Many systems are already live in sectors such as social media, retail and content platforms. Most solutions demonstrated low-friction user experiences, fast estimation times (typically under 20 seconds) and high accuracy outside threshold "buffer zones" (e.g. 13+, 16+, 18+). Some systems achieved mean absolute errors (MAE) of approximately one year in controlled conditions and provided reliable threshold classification when estimated ages exceeded configured buffers. However, it is a fundamental misunderstanding of the capabilities of age estimation to test whether it can implement exactly a specific age-restriction without either accepting there will be a margin of error or applying a buffer age to reduce that margin to an acceptable level, acknowledging that false negatives will then be inevitable and alternative methods will be required to correct them.

**A.25.4** Vendors demonstrated strong alignment with privacy and security expectations, including:

- Temporary biometric processing with no image retention
- On-device or edge estimation architectures
- Secure capture pipelines and encrypted data transmission
- ISO/IEC 27001-certified information security practices
- Presentation attack detection (PAD) and emerging defences against injection and deepfake manipulation

**A.25.5** Inclusivity and demographic fairness were active areas of development. While systems generally performed well across diverse user groups, some showed reduced accuracy for older adults, non-Caucasian users and female-presenting individuals near policy thresholds. Underrepresentation of Indigenous populations in training data remains a challenge, particularly for First Nations Peoples, though vendors acknowledged these gaps and committed to remediation through fairness audits and dataset diversification.

**A.25.6** The Trial also explored innovative and emerging modalities, including gesture-based age classification, voice analysis and motion pattern detection. While these approaches are promising – especially for privacy-sensitive, ambient or child-first environments – they are at earlier readiness levels and require further validation before widespread deployment.

**A.25.7** Critically, while age estimation is highly effective for real-time, contextual age checks, it is not currently suitable for generating verifiable digital credentials (e.g. for use in digital wallets or holder services). Probabilistic age estimates lack the fixed, attestable properties required for credential-based identity systems. However, estimation can support layered or progressive assurance models and serve as a valuable pre-check or fallback when ID-based verification is unavailable or declined.

**A.25.8** Age estimation has emerged as a mature, secure and adaptable tool for enforcing age-based access in a wide range of digital and physical contexts. When configured responsibly and used in proportionate, risk-based scenarios, it supports inclusion, reduces reliance on identity documents and enhances user privacy. Its alignment with evolving international standards – combined with continuous innovation in model accuracy, fairness and spoof resistance – positions age estimation as a key component of modern, privacy-respecting age assurance infrastructure.

## Key Statistics from the Trial on Age Estimation

**28,514** Total sample size
facial image tests across controlled settings

**13** Participant Age Estimation Vendors

Best-performing systems (MAE) **1.3 - 1.5 years** mean absolute error

Typically **40** secs. or under to complete estimation

| | |
|---|---|
| No biometric retention | Majority of providers (on-device or temporary inference only) |
| Emerging modalities observed | Gesture, voice, behavioural motion (early TRL stages) |
| Vendors using on-device processing | 2 Providers |

**7/11** ISO/IEC 27001 certified providers confirmed

Most systems achieve at least **92% True Positive Rate (TPR)** at **18+ threshold** when estimated **age ≥ 19**

**Technology Readiness Levels (TRL)**
Majority of Vendors **at TRL 7 or above**

TRL 7+

***Figure A.25.1*** *Key Statistics from the Trial on Age Estimation*

**A.25.9** Age estimation is a flexible, low-friction method of age assurance that offers a practical, privacy-preserving way to assess whether an individual is likely to meet a given age threshold – without the need for formal identity documents or declared dates of birth. It provides service providers, regulators and users with a rapid, non-intrusive tool for assessing age eligibility, particularly in lower-risk or high-volume environments such as social media, app stores and content access gateways.

**A.25.10** While it does not offer the binary certainty of verified date-of-birth checks, age estimation can achieve high levels of accuracy, especially when applied to clear thresholds (e.g. under/over 13, 16 or 18). When configured appropriately (e.g. with a buffer age) and supported by transparent confidence scoring, it allows systems to make probability-based age decisions that are contextually appropriate and scalable.

**A.25.11** When deployed using privacy-preserving, bias-aware and standards-aligned practices, age estimation strikes a meaningful balance between:

- Risk-appropriate compliance
- User autonomy and privacy
- Operational scalability and efficiency

**A.25.12** Its adaptability makes it particularly well suited to real-time use cases and it is increasingly being integrated into interoperable ecosystems – such as platforms exploring in-device estimation, in-app gating or signal-based assurance within digital identity frameworks. As confidence in its accuracy and fairness continues to grow, age estimation plays an important role in the broader ecosystem of age assurance methods.

## A.26 Who Participated in the Trial of Age Estimation Technology













## A.27 Observations About Age Estimation

A.27.1 **Age estimation can be effectively implemented in the Australian context** and is already in active use. It provides a fast, low-friction, document-free method of age assurance well-suited to binary threshold decisions (e.g. 13+, 16+, 18+). Several systems are already deployed across sectors such as social media, e-commerce and youth platforms. These implementations align with emerging international standards, including ISO/IEC FDIS 27566-1 and IEEE 2089.1.

A.27.2 Most systems are technically deployable across standard devices and environments, though edge-case limitations remain. The Trial found **no substantial technological limitations** to adoption. However, performance may degrade under poor lighting, occlusion, extreme angles or low-resolution input – conditions requiring ongoing optimisation for reliable real-world use.

A.27.3 Provider **claims regarding performance were generally substantiated** through independent evaluation. System outputs under test conditions aligned with stated model accuracy, confidence thresholds and demographic performance metrics. A minority of early-stage systems lacked complete transparency, but most vendors provided sufficient documentation and verification.

A.27.4 **Age estimation must be configured to context** – there is no one-size-fits-all approach. Vendors offered different deployment models (e.g. real-time or asynchronous), with configurable thresholds, fallback methods and policy tuning to match the risk profile and regulatory context of the relying party.

A.27.5 The age estimation sector is **innovative, fast-moving and responsive to privacy and fairness challenges**. Providers are iterating rapidly – introducing on-device AI, synthetic data augmentation and lower-latency models – while integrating privacy-preserving architectures such as edge inference and federated learning.

A.27.6 **Demographic performance consistency is improving** but requires continued focus. While many systems showed broadly fair results, some exhibited reduced accuracy for non-Caucasian users, older adults or female-presenting users near age thresholds. Underrepresentation of Indigenous populations, particularly First Nations and Torres Strait Islander Peoples, remains a challenge that vendors are beginning to address through dataset expansion and fairness auditing.

A.27.7 **Accuracy varies in suboptimal conditions**, highlighting the need for robustness improvements. Test scenarios involving occluded faces, substandard lighting, unusual angles or low-quality cameras showed increased false rejections or misclassifications. Technical refinements are needed to maintain reliability in these edge cases and avoid excluding eligible users.

A.27.8 **Vendors are actively mitigating adversarial threats**, including spoofing and injection attacks. Most systems implemented ISO/IEC 30107-aligned **presentation attack detection** and began preparing for injection risk countermeasures (as defined in ISO/IEC AWI 25456). Project DefAI and other certification initiatives will further support resilience to manipulation and deepfakes.

A.27.9 Contextual signals (e.g. parental controls or device settings) were sometimes integrated but not treated as authoritative. These elements were typically used as supplementary inputs to support or refine user journeys. **Age estimation decisions remained based on real-time**, independently derived evidence – not on self-declared, inferred or parental assertions.

A.27.10 Providers are aligning with emerging international standards and **demonstrating readiness for certification**. Many systems reflected the privacy, transparency and proportionality expectations in ISO/IEC FDIS 27566-1 – especially Clauses 5.3 (Privacy), 5.7 (Fairness), 6.2 (Friction Minimisation) and 6.4 (Confidence Expression) – as well as information security standards like ISO/IEC 27001 and biometric assurance practices under IEEE 2089.1.

Age Assurance Technology Trial

E

# PART E
# Age Inference

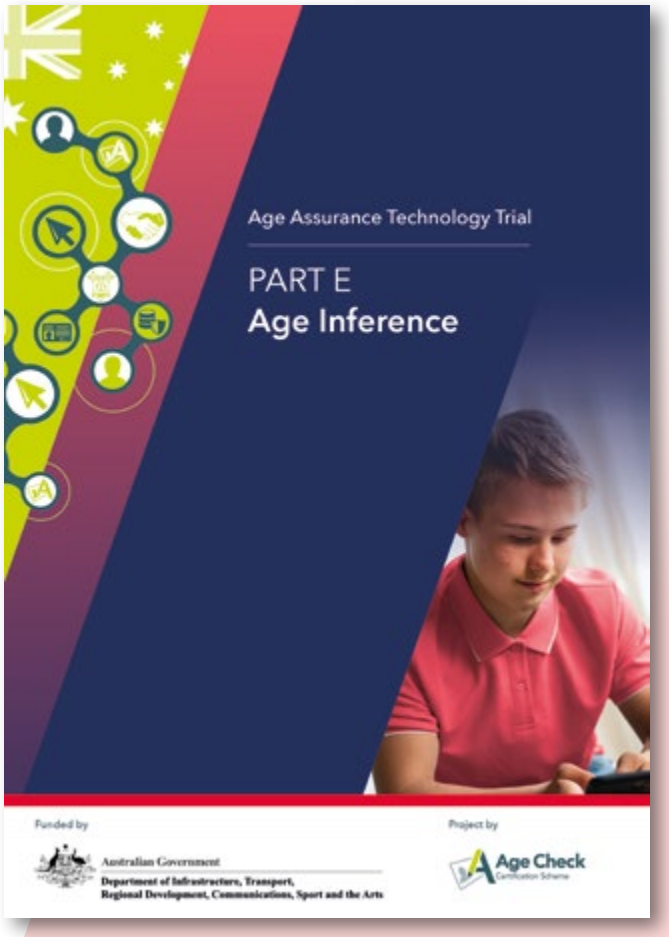


***See full report: Part E** – Age Inference*

## A.28 Findings on Age Inference

**A.28.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the topic of age inference.

1. **Age inference can be done in Australia, is viable and effective** in a variety of use cases.

2. **No substantial technological limitations** preventing its implementation in the Australian context.

3. Inference methods most **accurate when grounded in clearly modelled reasoning** and when drawing from well-labelled behavioural signals.

4. **Age inference is inherently context specific** and must be tailored to the sector, risk profile and digital behaviours of the user group.

**5** The age inference sector in Australia is **dynamic and innovative,** with a range of techniques being explored by providers.

**6** **Security and governance of inference systems were generally strong,** particularly among independent providers and those using in-session logic.

**7** Inference quality depends on the **transparency and reasonableness of the underlying logic;** increases effectiveness of system performance.

**8** **Fairness and demographic sensitivity** remains active areas for improvement; some systems risked bias.

## A.29 What is Age Inference

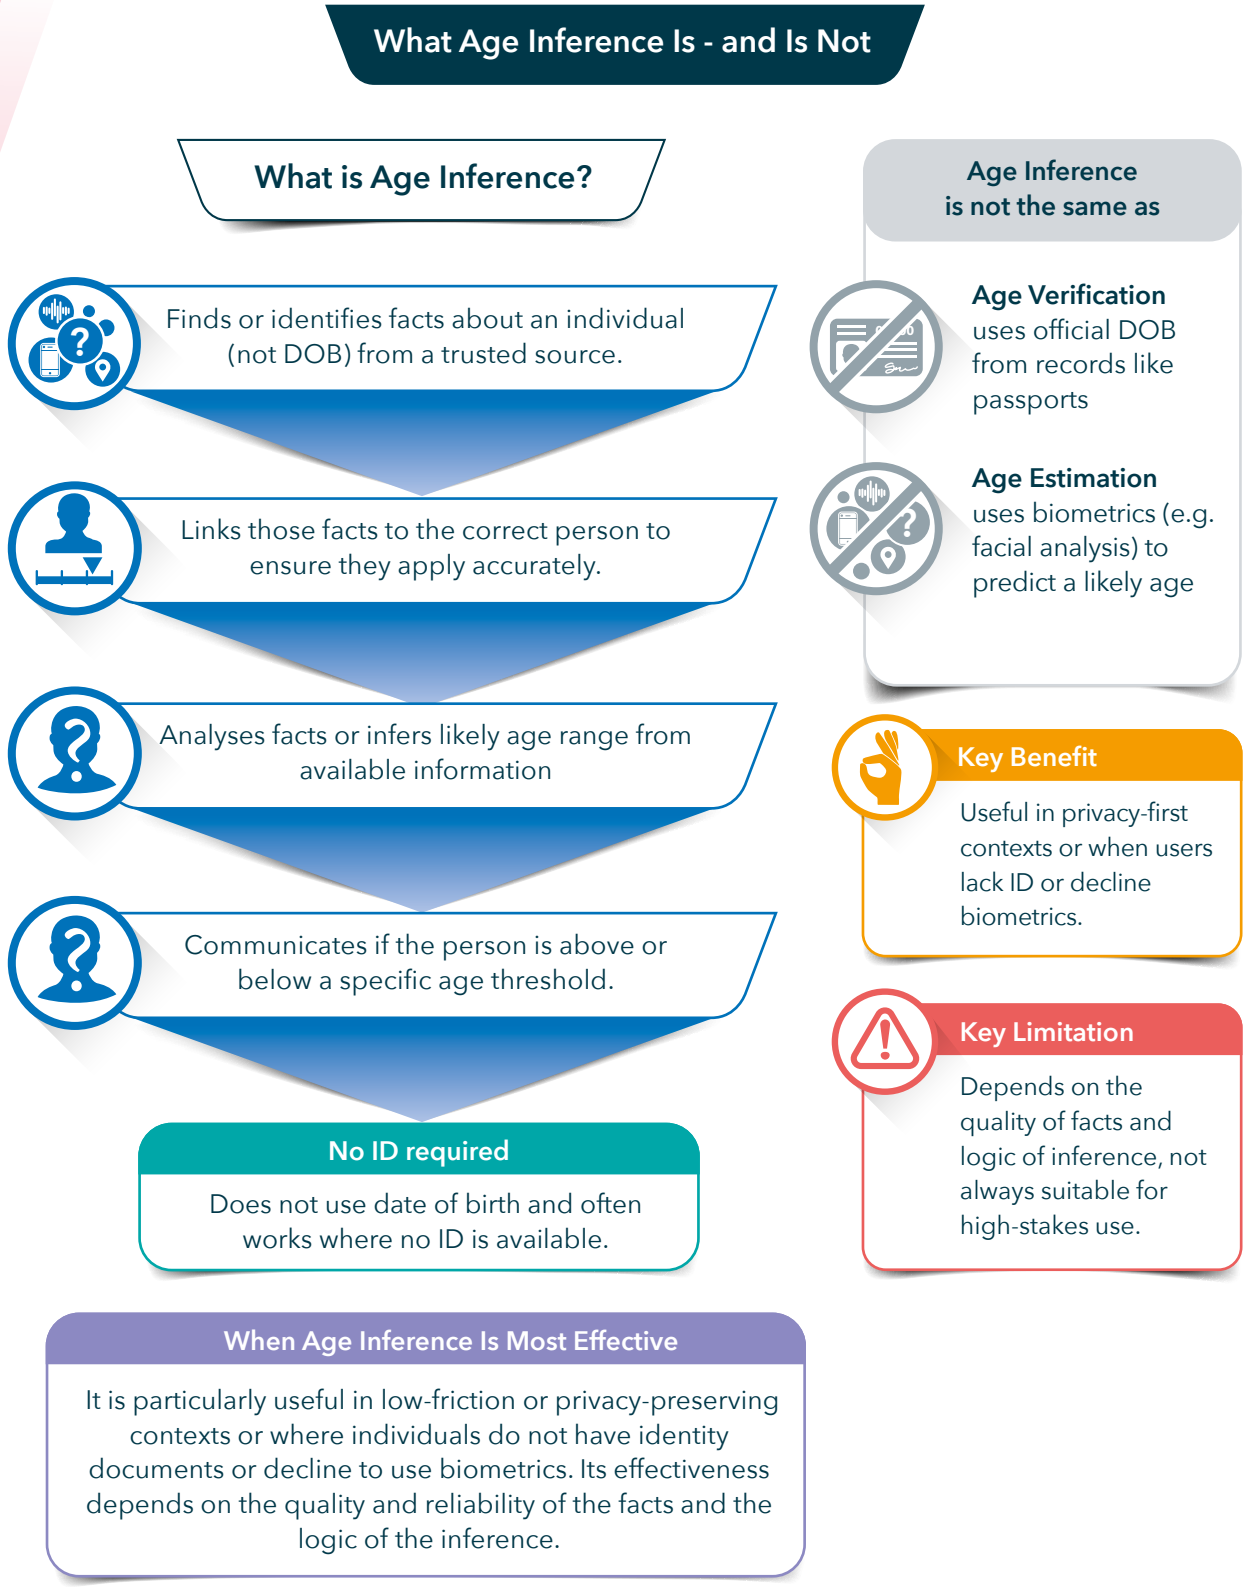


***Figure A.29.1*** *What Age Inference Is and Is Not*

## A.30 Introduction to Part E: Age Inference

**A.30.1** Part E of the Age Assurance Technology Trial focuses specifically on age inference – a method of determining an individual's likely age or age range based on verifiable contextual, behavioural, transactional or environmental signals, rather than biometric data or identity documents. Unlike age verification, which relies on a known and validated date of birth or age estimation, which uses biometric characteristics to predict age, age inference draws reasonable conclusions about age by analysing facts such as school enrolment, financial transactions, content barring settings, service usage or participation in age-specific activities.

**A.30.2** This section evaluates how age inference systems perform in the Australian context, including their technical feasibility, contextual appropriateness, demographic inclusivity, privacy alignment and overall resilience to manipulation or circumvention. The Trial assessed alignment with relevant international standards, particularly ISO/IEC FDIS 27566-1, which defines the privacy, purpose limitation and effectiveness expectations for age assurance systems and IEEE 2089.1, which outlines performance and interoperability requirements for age-related signals.


opal

## A.31 Summary of Age Inference

**A.31.1** Age inference is an approach to age assurance that implies a user's likely age or age range based on behavioural patterns, contextual data, digital interactions or metadata – without requiring direct identity verification or biometric estimation. It is especially valuable where formal documents are unavailable, disproportionate or culturally inappropriate.

**A.31.2** The Trial found that age inference can be effectively and ethically implemented in Australia. A wide variety of verifiable life-stage indicators – such as electoral enrolment, school year, transaction history, email metadata or device usage patterns – can support accurate, session-specific age classification. When deployed close to the point of risk (e.g. accessing age-restricted content or making a regulated purchase), inference systems support proportionate, low-friction user journeys while upholding privacy.

**A.31.3** Independent providers demonstrated mature, standards-aligned implementations, with most systems discarding raw input signals after inference and avoiding persistent tracking or profiling. These approaches reflected strong alignment with ISO/IEC FDIS 27566-1, particularly its principles on data minimisation, footprint control and contextual relevance. Most providers operated under certified security frameworks such as ISO/IEC 27001 and showed robust safeguards against spoofing, signal injection or false behavioural profiles.

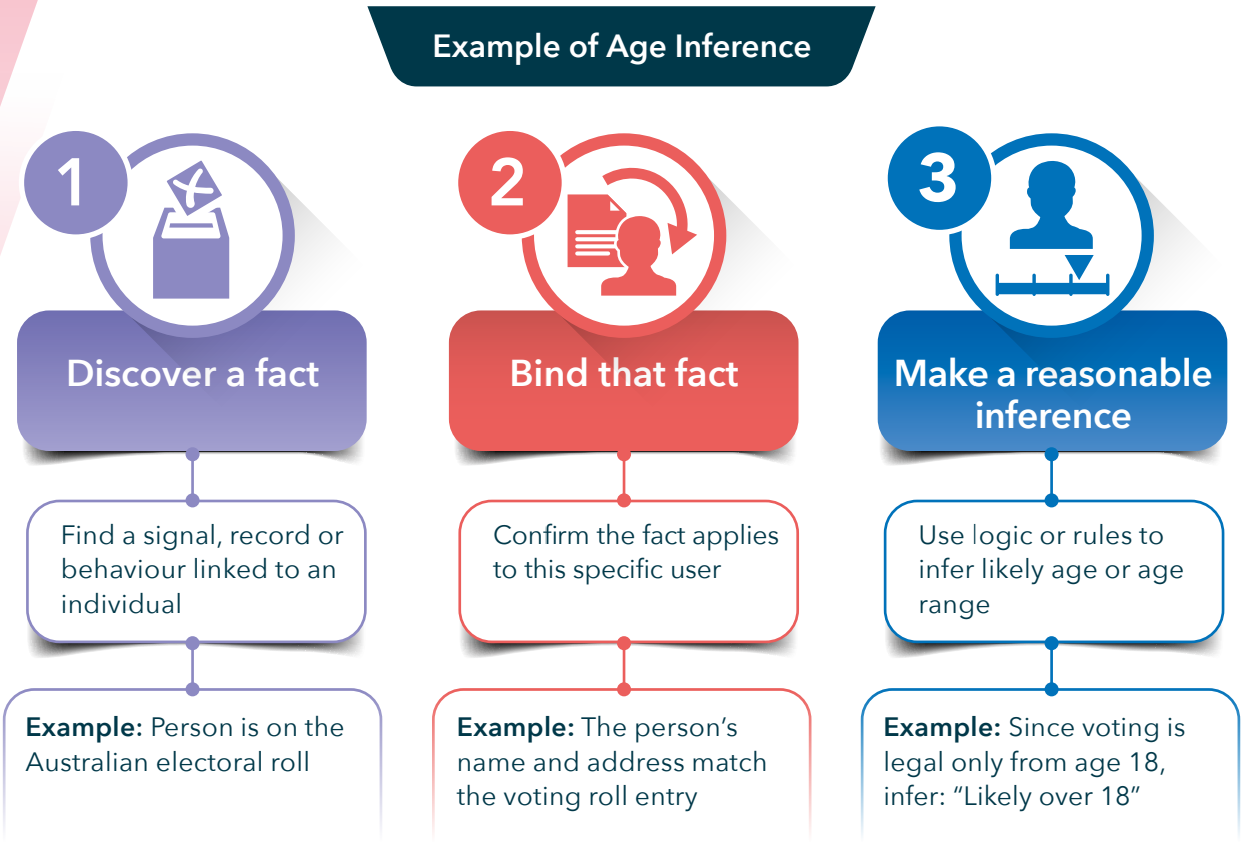


***Figure A.31.1*** *Example of Age Inference*

A.31.4 Providers used diverse inference methods – including email domain recognition, session metadata analysis, interaction patterns, credit eligibility and content engagement. Several participants demonstrated high accuracy in real-world use cases (e.g. detecting likely under-13 or over-18 users) and applied conservative thresholds or fallback logic to minimise misclassification. Some also explored early-stage, culturally grounded inference approaches, including use of knowledge markers or community roles relevant to First Nations contexts.

A.31.5 While session-based inference models offer strong privacy protections, the Trial also identified concerns where inference becomes persistent or platform-wide, particularly in account-based environments. Continuous behavioural monitoring may lead to digital profiling or cross-context inference reuse, which can undermine transparency and user autonomy. In such cases, regulatory clarity may help shape how inference is applied – ensuring it remains proportionate, aligned to risk, and respectful of user expectations.

## | Future opportunities for age inference

**A.31.6** Inference systems were also shown to have potential in future verifiable credential frameworks, issuing temporary, cryptographically signed assertions (e.g. "Likely Over 18") for use in digital wallets. While promising, these innovations should be carefully governed to prevent credential misuse, persistent tracking or cross-service linkage.

**A.31.7** In summary, age inference is a flexible, scalable and privacy-conscious tool for digital age assurance – especially in low-risk, child-facing or successive validation scenarios. When implemented with clear logic, contextual boundaries and transparent governance, it provides a valuable complement to document-based and biometric approaches. Ongoing developments in innovation, inclusion, and standards-aligned oversight may play a key role in maintaining public trust and helping age inference remain safe, fair and fit for purpose.

## Key Statistics from the Trial on Age Inference

**9** Age Inference Providers

Signals included email domains, device metadata, session behaviour, transaction history in the Trial

**6** providers used zero raw data retention

Typical inference latency **Under 5 seconds** for most real-time systems

**7** providers used spoofing mitigation systems

| | |
|---|---|
| Contextural deployment models | Triggered by access attempt, transaction or content engagement |
| Use of fallback or escalation pathways | All providers supported layered or conservative inference configurations |
| Potential for credential issuance (e.g., "Likely Over 18") | 3 providers piloting inference-backed credentials for digital wallets |

**2** providers exploring early-stage approaches for First Nations people

All providers demonstrated formal standard alignment **(ISO/IEC FDIS 27566-1)**, with 4 using Practice Statements

**5** ISO/IEC 27001 certified providers confirmed

***Figure A.31.2*** *Key Statistics from the Trial on Age Inference*

## A.32 Who Participated in the Trial of Age Inference Technology

**A.32.1** Participation in the age inference aspect of the Trial was limited, partly due to the relative novelty of the term "age inference" as defined in ISO/IEC FDIS 27566-1. As a newly formalised concept, age inference is not yet widely understood or differentiated from related age assurance methods such as age verification (based on declared date of birth) or age estimation (using biometric traits).

**A.32.2** As defined in ISO/IEC FDIS 27566-1, age inference refers to implying a person's likely age or age range based on contextual, behavioural, transactional or environmental signals – such as school enrolment, account tenure or device settings – rather than through biometric or document-based checks. While some providers already undertake similar logic-driven processes within their platforms, many have not yet labelled or structured these as distinct "age inference" offerings, which may have limited the number of formal participants in this category.

## Trial Participants

agechecked

EQUIFAX

frankieone

IDverse™
A LexisNexis® Risk Solutions Company

LUCIDITI®

MyMahi

PRIVO®
Privacy • Permission • TRUST

verifymy

YOTI

## A.33 Observations About Age Inference

A.33.1 **Age inference is viable and effective for age assurance in Australia**, especially when used to flag likely underage access, support early safety interventions or trigger fallback mechanisms. When designed with transparent logic, strong input signals and bounded confidence thresholds, inference systems offer low-friction, proportionate assurance.

A.33.2 There are **no substantial technological limitations to deploying age inference** systems. Many providers demonstrated the ability to operate on existing platform data (e.g. interaction logs, metadata) using mature inference engines or rule-based models. System effectiveness depends not on technology maturity but on the relevance, diversity and quality of input signals.

A.33.3 Inference methods were most **accurate when grounded in clearly modelled reasoning** and when drawing from well-labelled behavioural signals. For example, systems using language complexity, session timing or feature access patterns were able to classify age thresholds (e.g. under 13, under 18) with high reliability in controlled and real-world scenarios.

A.33.4 **Age inference is inherently context specific**. It must be tailored to the sector, risk profile and digital behaviours of the user group. Inference approaches were most successful when adapted to the norms of child-facing platforms, regulated services or role-specific environments (e.g. education, gaming, retail).

A.33.5 The age inference sector is **innovative and rapidly evolving**, with providers exploring techniques like gesture modelling, narrative complexity analysis and contextual metadata synthesis. Some systems demonstrated promising accuracy using only temporary, non-identifying inputs, supporting the shift toward zero-knowledge assurance.

A.33.6 **Security and governance of inference systems were generally strong**, particularly among independent providers and those using in-session logic with no persistent data. However, platforms deploying inference as part of ongoing behavioural monitoring must address the risk of profiling, surveillance or footprint expansion, as flagged in ISO/IEC FDIS 27566-1.

A.33.7 Inference quality depends on the **transparency and reasonableness of the underlying logic**. Systems that disclosed their signal-to-inference mapping, applied confidence thresholds conservatively and included fallbacks or escalation paths performed more ethically and effectively, with fewer classification errors.

A.33.8 **Fairness and demographic sensitivity** remain active areas for improvement. Some systems risked bias where signals correlated with cultural, linguistic or socioeconomic factors. Providers are encouraged to conduct demographic impact assessments, improve calibration and align with inclusion clauses of ISO/IEC FDIS 27566-1 to mitigate false positives.


STOP
END
10
GIVE WAY
NO PARKING BETWEEN SIGNS
DRIVE SAFELY
NEW WORK
SCHOOL BUS ROUTE
PREPARE TO STOP
CREST
60

Age Assurance Technology Trial

F

# PART F
# Successive Validation

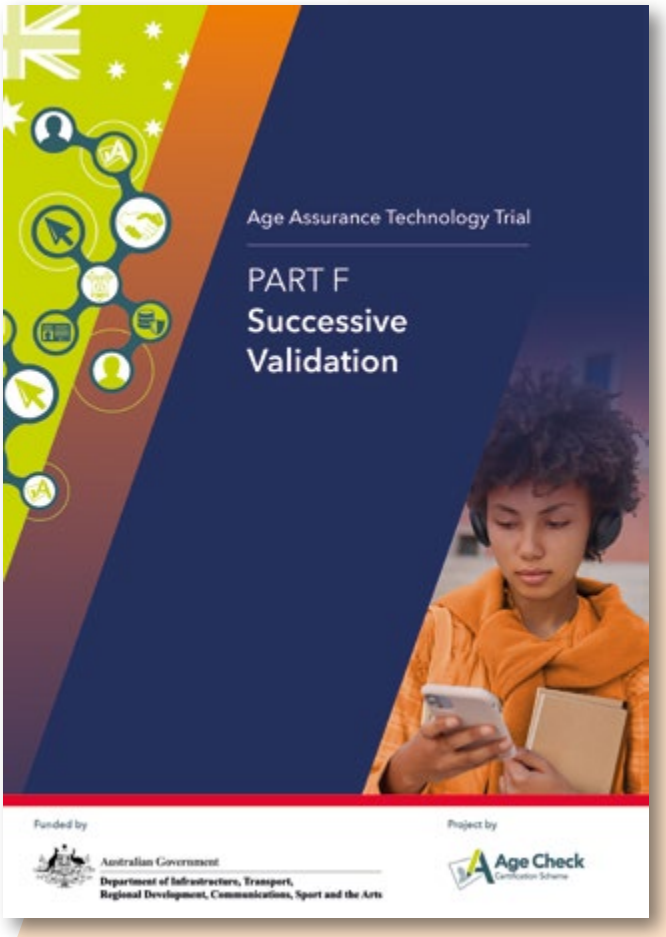


***See full report: Part F** – Successive Validation*

## A.34 Findings on Successive Validation

**A.34.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the topic of successive validation.

1. Successive validation **can be done** in Australia and aligns with emerging international standards.

2. **No substantial technological limitations** preventing its implementation in the Australian context.

3. Successive validation systems demonstrated **internal consistency and standards alignment,** including alignment with ISO/IEC FDIS 27566-1.

4. **There is no single configuration to successive validation;** flexible models exist and approaches varied by risk context and use case.

5. An evolving and innovative sector is **actively exploring layered age assurance models;** an industry focused on inclusion is maturing.

6 **Strong privacy-by-design principles** were observed across successive validation stages.

7 Successive validation can **enhance demographic inclusion and reduce bias,** supporting users without formal ID.

8 **Configuration and escalation logic** would benefit from clearer standardisation and guidance.

9 **Cybersecurity practices aligned with best practice** and addressed emerging attack surfaces; various defences employed to protect against manipulation.

## A.35 What is Successive Validation

**A.35.1** Successive validation is a type of age assurance process where multiple independent methods – such as age inference, estimation and verification – are used sequentially to reach a confident age assurance result.

**A.35.2** Sometimes referred to as a '**waterfall technique**', this process begins with a low-friction method (e.g. age inference or estimation). If the result is inconclusive – particularly near a threshold age (e.g. 18) – the system escalates to another method. This might involve collecting contextual data for inference or requesting biometric estimation. If uncertainty remains, it may culminate in a full documentary age verification.

**A.35.3** This approach is governed by risk and proportionality. The closer a user appears to the threshold, the more likely additional steps are required. When well-designed, successive validation:

- Applies the lightest effective method first;
- Escalates only when necessary; and
- Supports data minimisation by avoiding unnecessary collection and retention.

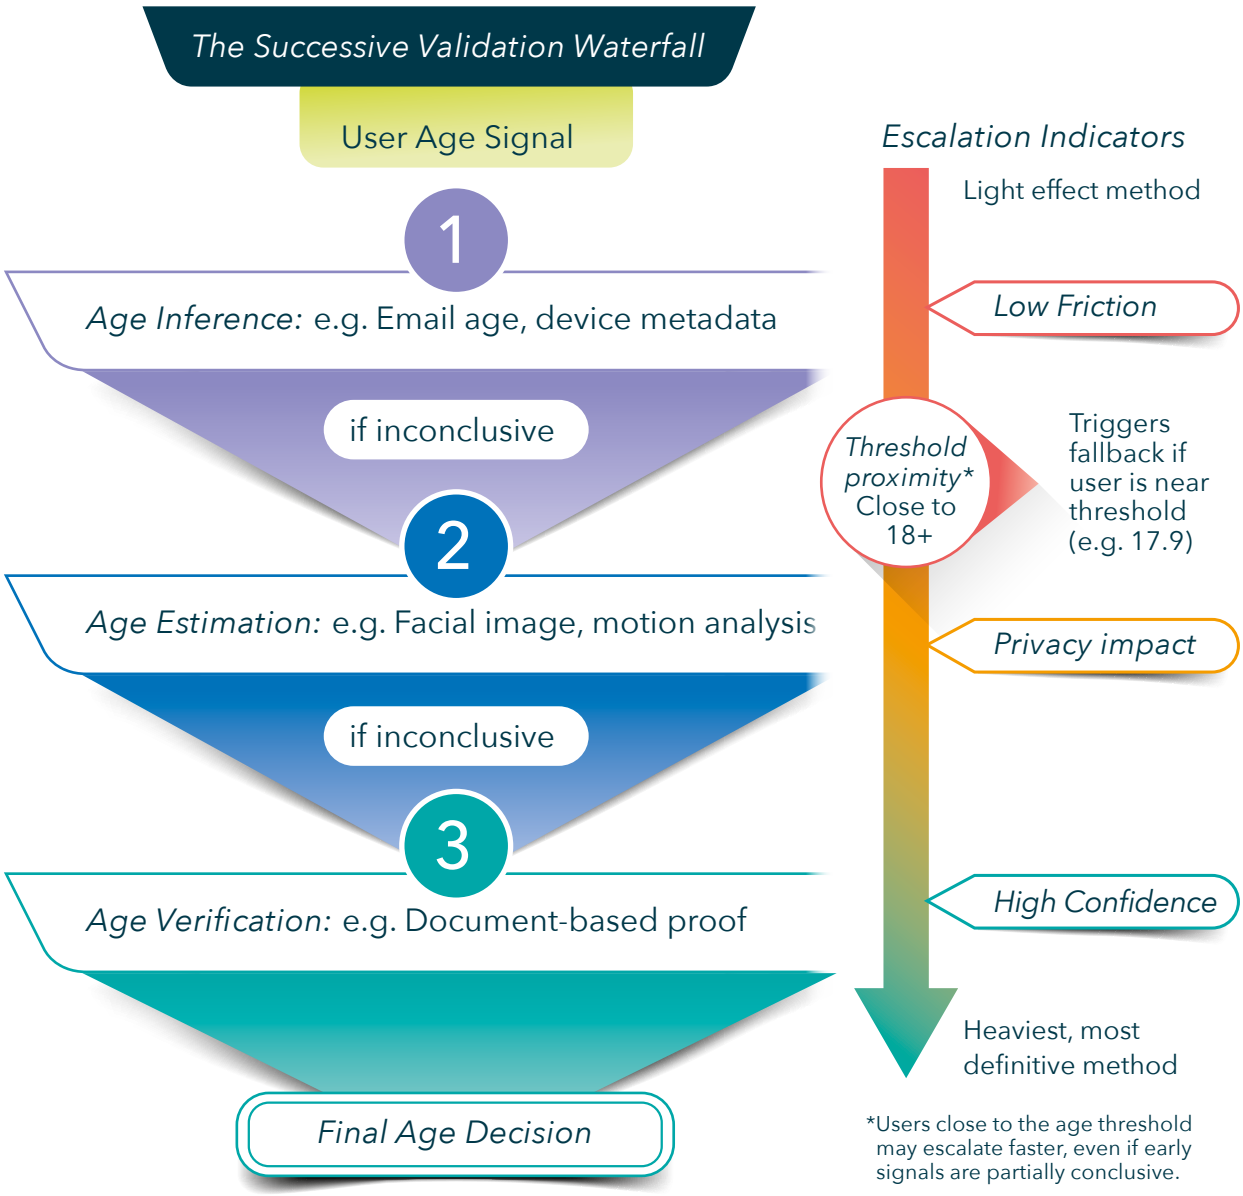


***Figure A.35.1*** *The Successive Validation Waterfall*

## A.36 Introduction to Part F: Successive Validation

**A.36.1** Part F of the Age Assurance Technology Trial focuses on successive validation, the process of combining two or more age assurance methods (such as age inference, age estimation and age verification) to reach a more accurate, risk-appropriate or confidence-boosted age-related decision. Defined in ISO/IEC FDIS 27566-1, successive validation supports the principle that age assurance should be proportionate to risk, enabling layered approaches where no single method alone is sufficient or contextually appropriate.

**A.36.2** Successive validation plays a critical role in real-world deployments by balancing friction, privacy and assurance levels. For example, a platform may initially infer age based on behavioural signals, escalate to biometric estimation if the result is uncertain and offer verification as a fallback only in edge cases. This model allows services to manage trade-offs dynamically using the lightest effective method wherever possible and only requesting higher-assurance inputs when necessary.

**A.36.3** This section of the report evaluates how successive validation has been implemented by Trial participants in the Australian context, examining technical feasibility, data flow, fallback logic, interoperability, privacy handling, demographic consistency and conformance with international standards particularly ISO/IEC FDIS 27566-1, which provides specific guidance on successive validation workflows and IEEE 2089.1, which supports consistency of age-related outputs across methods.


Age Assurance
Technology Trial

## A.37 Summary of Successive Validation

**A.37.1** This section of the Trial report examines the feasibility, implementation and implications of successive validation – a layered approach to age assurance in which multiple methods, such as age inference, age estimation and age verification, are applied in sequence to increase confidence in a user's age. Successive validation enables services to begin with low-friction, privacy-preserving techniques and escalate only when uncertainty remains or the user appears close to a critical threshold. It reflects principles of proportionality and user sensitivity, offering an adaptable model for contexts where no single method alone is sufficient.

**A.37.2** Drawing on practice statements, interviews, technical reviews and international standards – particularly ISO/IEC FDIS 27566-1 and IEEE 2089.1 – the Trial evaluated how successive validation has been deployed by technology providers within the Australian context. The assessment found that successive validation is both technically viable and operationally effective. Providers demonstrated considered and standards-aligned designs that escalated users through assurance steps based on risk, confidence levels and policy-defined thresholds. Configurations varied across services and sectors, but all shared a common emphasis on minimising unnecessary friction while ensuring appropriate assurance.

**A.37.3** The report identifies emerging use cases, including dynamic validation flows embedded in platform-level monitoring systems. In particular, social media services are beginning to apply continuous assurance logic, using behavioural signals – or contra indicators – to detect discrepancies in declared age and trigger additional validation. This approach mirrors real-world escalation (e.g. a shopkeeper requesting ID when unsure of a customer's age), but raises new questions around transparency, data minimisation and user control.

**A.37.4** Privacy-by-design was a consistent theme across provider systems. Early stages of validation typically relied on anonymised, temporary signals that avoided persistent data collection. As validation progressed, providers demonstrated careful separation between operational, training and evaluation datasets and employed clear logic to limit data exposure to what was strictly necessary for each step. Where more intrusive methods – such as document verification or biometric analysis – were required, these were invoked only when prior stages yielded insufficient confidence.

**A.37.5** Security protections were also robust. Systems included defences against spoofing, input manipulation and so-called "hill-climb" attacks. Rate-limiting, session binding and unpredictability in escalation logic helped prevent adversarial circumvention. Providers also addressed cross-method attack surfaces by securing the interfaces between validation steps and ensuring that age assurance outputs could not be tampered with during escalation.

**A.37.6** The Trial found no evidence of systemic demographic bias in the configurations examined. Some providers had begun experimenting with culturally grounded assurance signals to address inclusivity, such as contextual cues that may be more accessible to First Nations users or individuals without conventional identity documents. These early efforts suggest that successive validation may offer a more equitable model of age assurance by enabling alternative pathways to demonstrate eligibility.

### *Opportunities for successive validation*

**A.37.7** While interoperability across platforms remains in its infancy, a number of providers are exploring how age signals – particularly from inference and estimation – can be made portable, privacy-respecting and policy-aligned. This is a critical next step to enable users to avoid repeating intrusive checks and to allow services to build consistent assurance without persistent profiling.

**A.37.8** In summary, successive validation represents a mature and adaptable model of age assurance, well suited to the diverse risk environments encountered in Australia's digital ecosystem. It allows systems to calibrate their assurance level in response to both contextual risk and user characteristics. When governed transparently and implemented in accordance with privacy and security best practices, successive validation has the potential to support inclusive, proportionate and scalable age assurance across sectors.

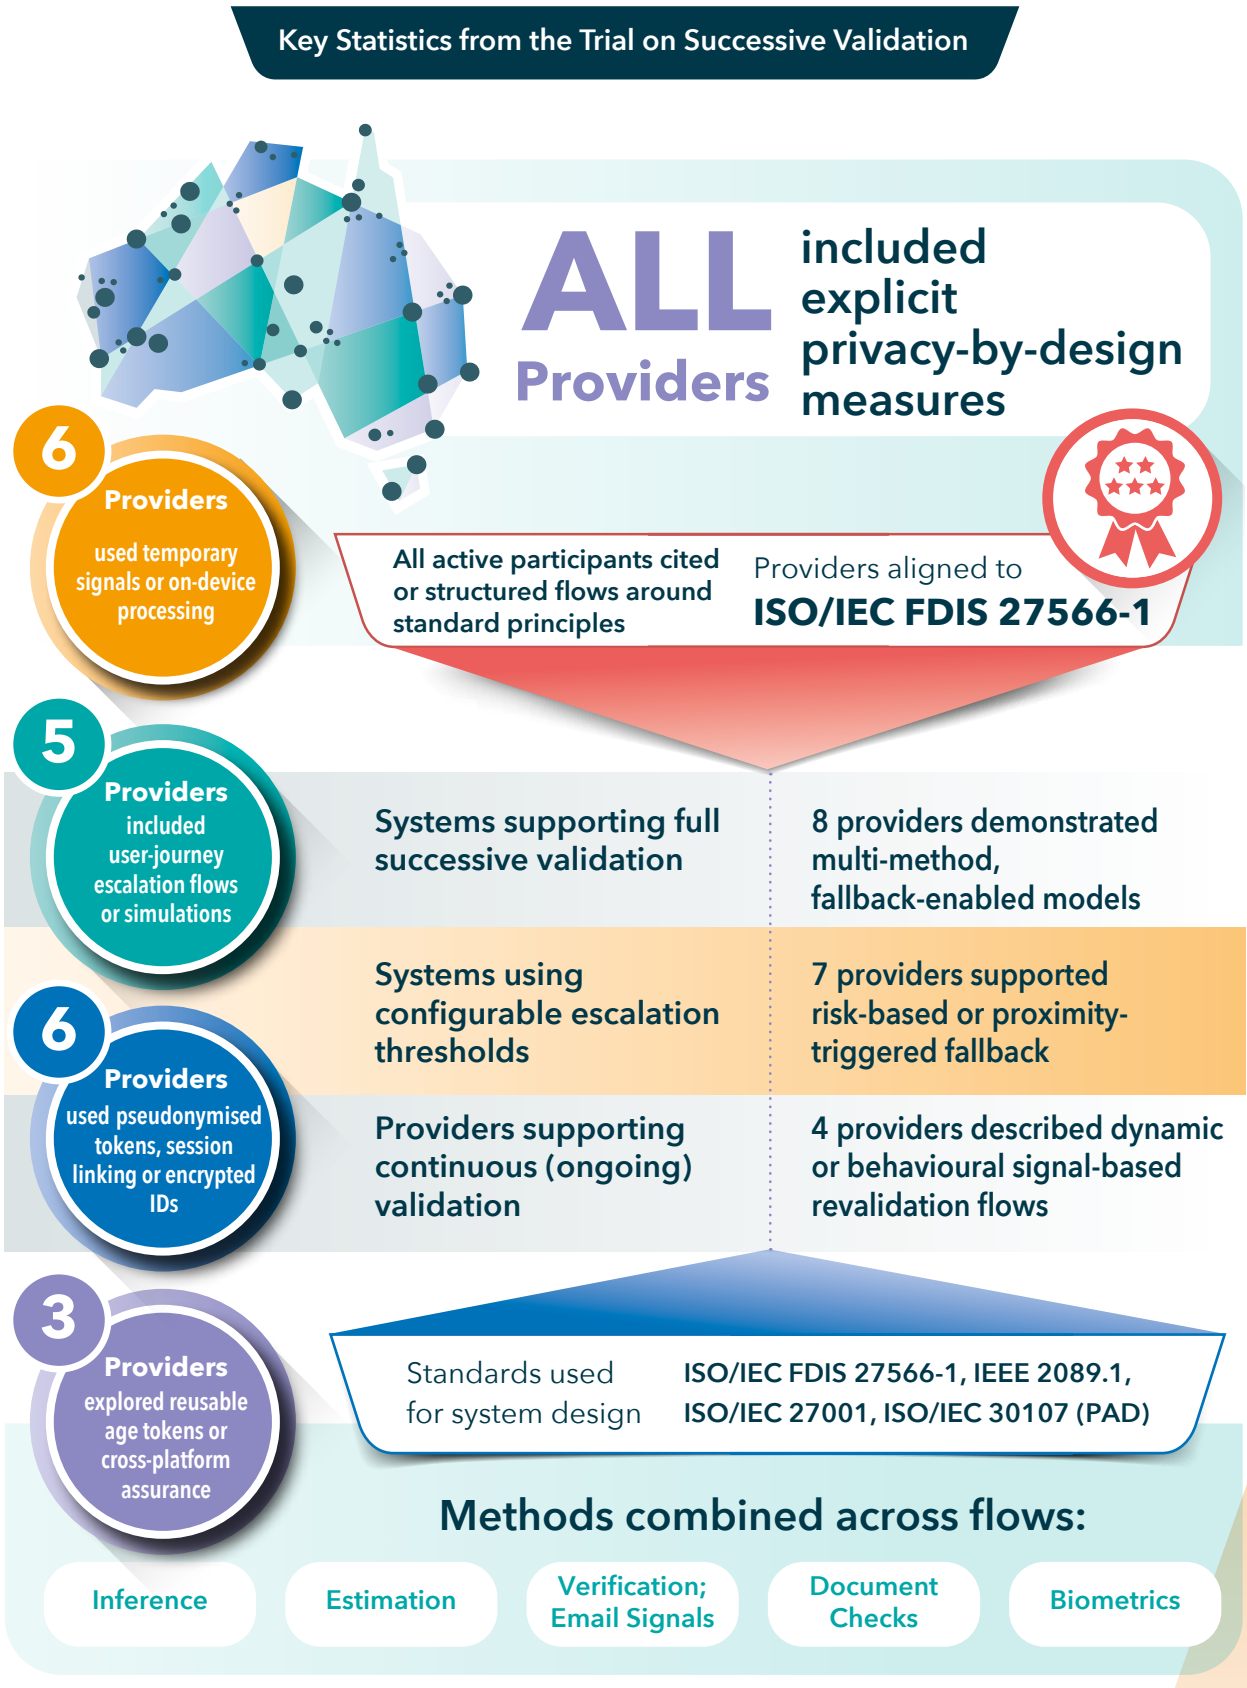


***Figure A.37.1*** *Key Statistics from the Trial on Successive Validation*

## A.38 Who Participated in the Trial of Successive Validation











## A.39 Observations About Successive Validation

A.39.1 **Successive validation is technically** feasible and aligns with emerging international standards in Australia. When applied proportionately, layering methods such as inference, estimation and verification enables services to provide scalable and risk-based age assurance. ISO/IEC FDIS 27566-1 and IEEE 2089.1 recognise successive validation as a recommended practice for improving confidence in age assessment.

A.39.2 There are **no substantial technological limitations** to implementing successive validation approaches. The Trial demonstrated that age assurance providers could integrate layered methods into service workflows using current technologies, supported by secure APIs orchestration logic and modular architecture.

A.39.3 Successive validation systems demonstrated **internal consistency and standards alignment.** Providers articulated well-defined escalation logic, fallback triggers and confidence thresholds, supported by privacy-preserving data handling and compliance with clauses from ISO/IEC FDIS 27566-1.

A.39.4 **There is no one-size-fits-all configuration,** but flexible models exist across services and sectors. Approaches varied by risk context and use case – from estimation-first models with fallback to document checks, to real-time platform-based escalation triggered by behavioural contra-indicators.

A.39.5 An evolving and innovative sector **is actively exploring layered age assurance models.** Providers demonstrated dynamic user journey flows, including real-time prompts, device-based checks and reuse of validated identities, reflecting a maturing industry focused on inclusion, compliance and user experience.

A.39.6 **Strong privacy-by-design principles** were observed across successive validation stages. Early-stage signals were typically anonymised or temporary and higher-assurance steps included safeguards such as pseudonymised tokens, strict data separation and one-time use of biometric data.

A.39.7 Successive validation can **enhance demographic inclusion and reduce bias.** By combining methods, systems can support users without formal ID, including young people near threshold ages and underrepresented communities. Some providers began exploring culturally grounded or context-aware assurance.

A.39.8 **Configuration and escalation logic** would benefit from clearer standardisation and guidance. Some implementations lacked transparency on fallback thresholds or policy triggers. Better tooling and policy support would help relying parties consistently apply risk-based successive validation.

A.39.9 **Security practices aligned with best practice** and addressed emerging attack surfaces. Providers employed defences such as rate-limiting, liveness detection, cryptographic token binding and spoofing mitigation to protect validation chains against manipulation.

Age Assurance Technology Trial

# PART G
# Parental Control

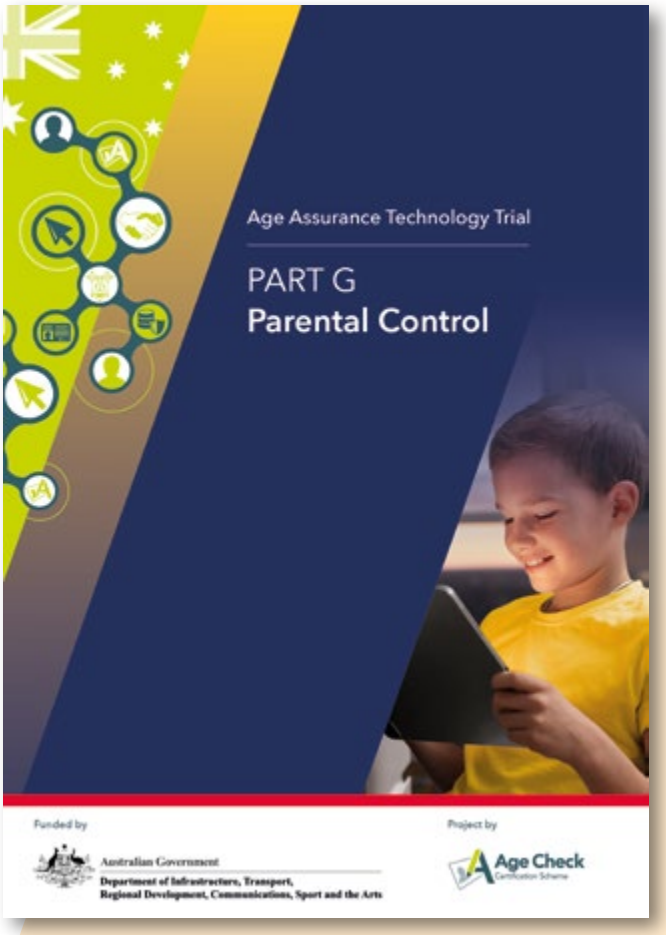


***See full report: Part G** – Parental Control*

## A.40 Findings on Parental Control

**A.40.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the topic of parental control.

1. Parental control **systems can be effectively applied** in Australia in many contexts.

2. Most systems **focus on restriction rather than participation;** there is limited accommodation for children's evolving capacity.

3. Parental control is a **proactive mechanism within layered assurance models,** supporting risk reduction in lower-risk, family-led environments.

4. Well-designed parental controls can generate **strong contextual age signals,** emitting useful indicators of a user's likely age range.

5. Effectiveness depends on **accurate and engaged setup by caregivers.** Configuration accuracy affects reliability of controls.

6 **Parental controls enable private forms of access management,** allowing restrictions without requiring direct age verification.

7 **Contextual signals should not be reused without consent;** this increases the risk of data misuse.

8 **Inclusivity and accessibility require ongoing attention,** though systems were broadly consistent across demographics.

9 Parental control signals **should not be treated as verified age data,** but rather as supplementary indicators.

10 Platforms seeking ways to **integrate parental control signals;** shared formats could support more consistent implementation across systems.

## A.41 What is Parental Control

**A.41.1** A parental control system is a set of tools or settings that allow parents or guardians to manage, restrict or monitor a child's access to digital content, services or device functions.

**A.41.2** Parental control systems are established in advance of the encounter by a child of age-restricted goods, content, services, venues or spaces. This differs from parental consent (which arises as a result of an age assurance process that then requires a decision by a parent or guardian about granting access or permission).

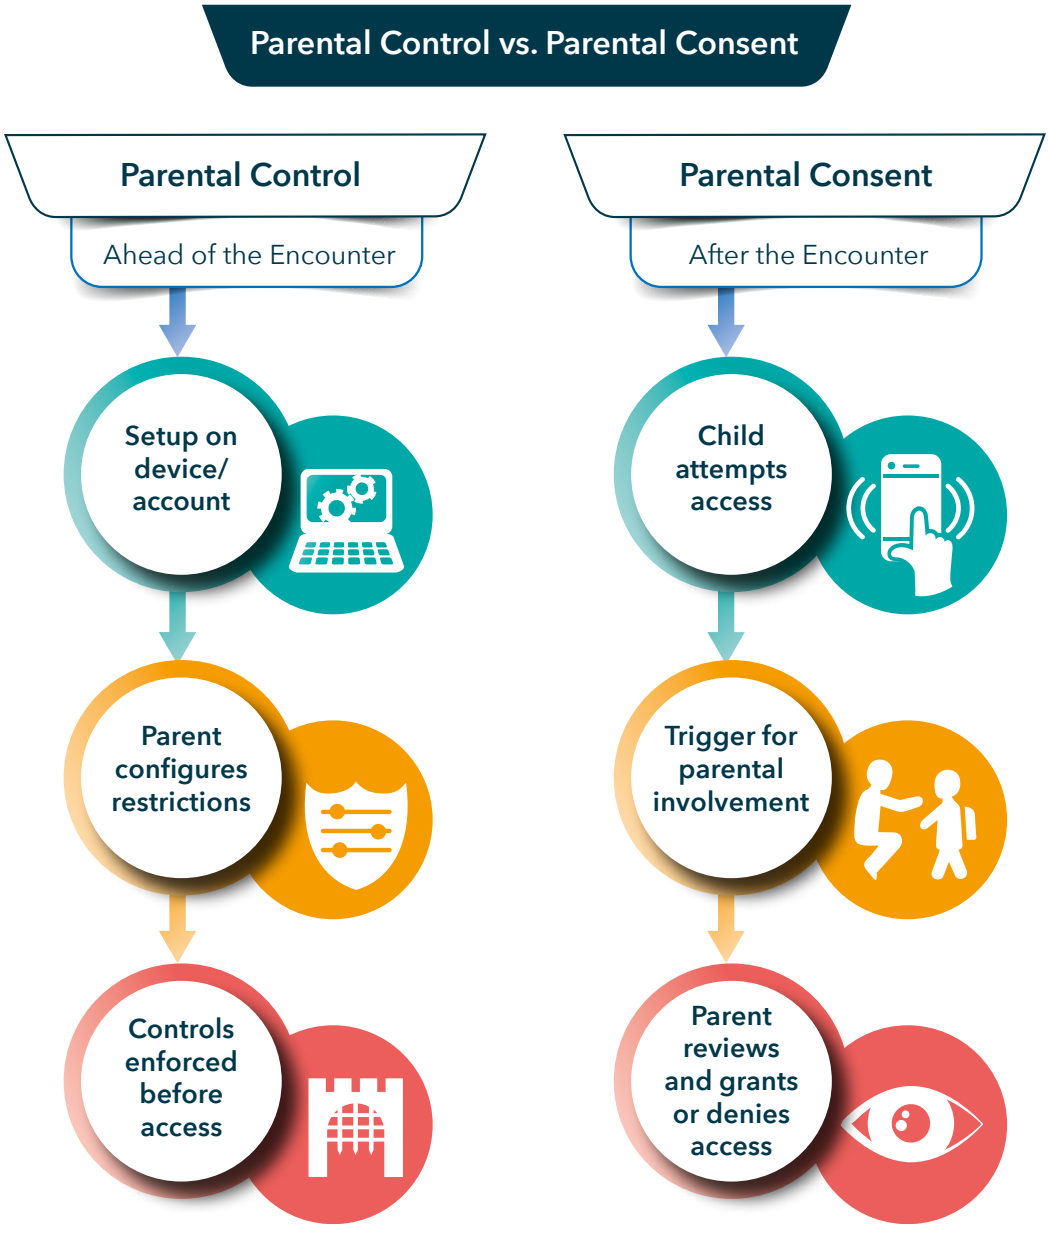


***Figure A.41.1*** *Parental Control vs. Parental Consent*

## A.42 Introduction to Part G: Parental Control

**A.42.1** Part G of the Age Assurance Technology Trial focuses specifically on parental control systems – the tools, configurations and supervisory features that allow parents or guardians to manage a child's access to digital content, services, devices or online functions. Parental controls play a significant role in digital safety ecosystems by providing families with the means to restrict or guide a child's exposure to age-inappropriate content, particularly in contexts where direct age verification or estimation may not be viable or proportionate.

**A.42.2** Unlike other age assurance methods – such as age verification, age estimation or age inference – parental control mechanisms are proactive and pre-configured. They are generally established by a parent or guardian before the child encounters age-restricted material and are managed either through platform tools (such as family account settings), device configurations (such as screen time or content filters) or network-level controls (such as home router restrictions). These systems may also provide indirect age signals to relying parties, contributing to layered or context-aware age assurance strategies.

**A.42.3** The evaluation undertaken in this part of the Trial assessed how parental control features operate in real-world deployments across platforms, devices and content services. The Trial explored how effectively these systems support safe digital engagement for children and adolescents, their limitations and the extent to which relying parties can trust or incorporate parental control status as part of their own compliance with age-related policies or regulations.

**A.42.4** Importantly, the Trial was designed as a technological evaluation and does not recommend or mandate any policy decision. It assessed whether technologies – including parental controls – are deployable, functional and privacy-preserving, but does not judge whether they should be adopted at a regulatory level. That distinction between capability and policy is fundamental.

**A.42.5** Through this section of the report, we examine the extent to which parental control systems can support risk-appropriate, low-friction and inclusive approaches to age assurance in Australia. We analyse their usability, configurability, consistency, privacy impact and technical maturity. The report also considers how parental controls interact with or supplement other forms of age assurance and explores their potential role in a broader, layered approach to child online safety.

## A.43 Summary of Parental Control

**A.43.1** Parental control systems are a well-established and widely available element of digital safety infrastructure. These tools enable parents and guardians to manage children's access to online content, services and devices – providing practical ways to guide digital engagement and reduce exposure to inappropriate material. While not a form of age assurance in themselves, parental controls can contribute meaningfully to broader, layered assurance models by emitting contextual signals that indicate a child is present.

**A.43.2** As part of the Trial, parental control systems were evaluated for their technical feasibility, usability, inclusivity, privacy impact and alignment with international standards and child rights principles. The assessment drew on vendor practice statements, interviews, system walkthroughs and alignment with ISO/IEC 29146 (framework for access management), IEEE 2089.1 and the UN Convention on the Rights of the Child (UNCRC)[15].

**A.43.3** The Trial found that parental control systems can be effectively and securely deployed across Australian platforms and contexts. Tools implemented at the device, network, platform and account levels are mature and already in widespread use, enabling families to configure access restrictions, time limits, content filters and supervision protocols. These systems provide a scalable and privacy-conscious foundation for access management in home environments and are particularly effective for younger children in lower-risk settings.

15. *The UNCRC is a legally binding agreement which outlines the fundamental rights of every child, regardless of their race, religion or abilities. Australia became a signatory to the UNCRC on 22 August 1990 and ratified it on 17 December 1990.*

**A.43.4** Parental controls operate pre-emptively – configured before a child attempts to access restricted content – which distinguishes them from reactive assurance mechanisms like age estimation or verification. When properly designed and implemented, parental control signals can indicate the presence of a supervised child profile or device, allowing relying parties to apply safeguards without needing to collect identity data.

**A.43.5** However, the evaluation also identified key areas for refinement. Most systems are static and do not adapt easily to children's evolving maturity, preferences or rights to participate in decisions about their digital lives. As children grow older, the absence of mechanisms for graduated autonomy or shared configuration can limit both the effectiveness and acceptability of these tools. In some cases, children may be subject to restrictions without visibility or recourse – raising important questions around dignity, fairness and transparency.

**A.43.6** Framed through the lens of the UNCRC, parental control systems should protect children from harm (Article 17), while also upholding their rights to privacy (Article 16), to express their views and be heard (Article 12) and to have increasing autonomy as they develop (Article 5). Striking this balance does not mean discarding parental control – it means evolving it to function as a guidance framework rather than a blunt restriction model and creating opportunities for children to participate in ways that are appropriate to their age and capacity.

**A.43.7** In addition, while many systems are designed with privacy in mind, concerns remain around persistent activity logging, data retention and the potential for over-surveillance – particularly if contextual signals are reused across services without consent. Configuration also presents usability challenges in some contexts, especially for families with limited digital literacy or those using shared devices. Cultural assumptions about caregiving models can limit the inclusivity of controls for First Nations families, multigenerational households and non-traditional guardians.

**A.43.8** Despite these limitations, the sector is evolving rapidly. Several providers are embedding more dynamic, context-aware features into real-time environments, improving flexibility and responsiveness. Platforms are increasingly interested in integrating parental control status into content gating or feature restriction workflows, signalling an opportunity to standardise control signals and promote interoperability, trust and consistency.

**A.43.9** Overall, parental control systems represent a valuable and effective part of the age assurance and digital safety toolkit. They can play a meaningful role in protecting children's online experiences – particularly when configured thoughtfully and deployed proportionately. But to realise their full potential, future development must embrace children's rights alongside safety goals, enabling systems that protect, include and empower.

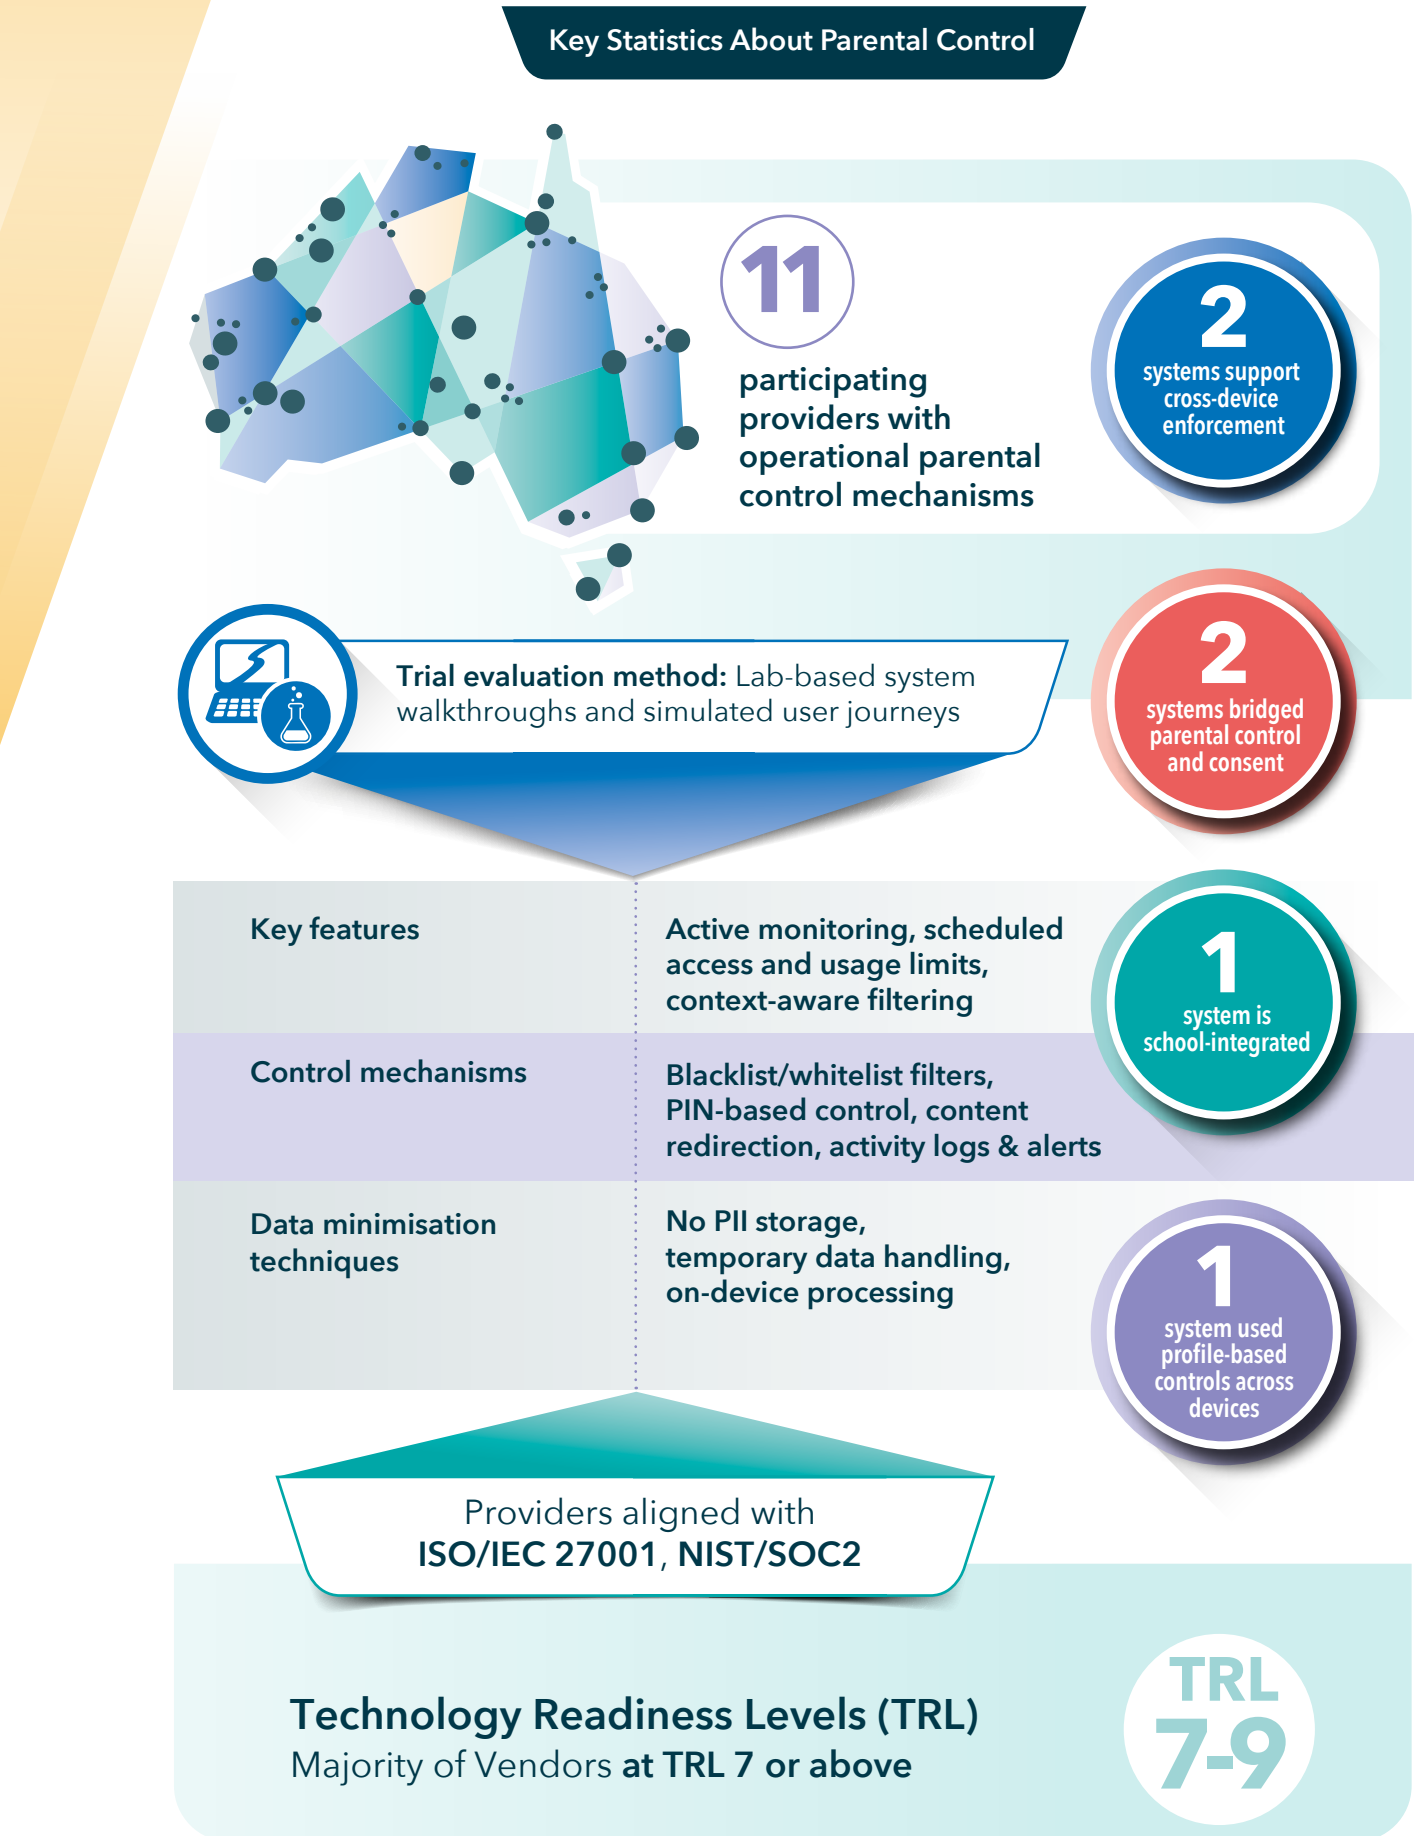


***Figure A.43.1*** *Key Statistics from the Trial on Parental Control*

## A.44 Who Participated in the Trial of Parental Control Technology










## A.45 Observations About Parental Control

A.45.1 **Parental control systems can be effectively applied in Australia**. Trial participants demonstrated the capability to configure and enforce age-appropriate restrictions using family control centres, device settings, platform tools and account-linked supervision. These systems are mature, usable and well-suited to managing children's access in many contexts.

A.45.2 Most parental control systems **focus on restriction rather than participation**. While effective at limiting access, current tools offer limited accommodation for children's evolving capacity, privacy or ability to be heard – key rights relevant to digital engagement and autonomy.

A.45.3 Parental control is a **proactive mechanism within layered assurance models**. Unlike estimation or verification, parental control is configured before a child interacts with restricted content. It supports risk reduction through parental oversight and is especially useful in lower-risk, family-led environments.

A.45.4 Well-designed parental controls can generate strong contextual age signals. When tied to a managed device or child profile, parental controls **can emit useful indicators of a user's likely age range**, supporting content moderation and gating in a privacy-respecting manner.

A.45.5 **Effectiveness depends on accurate and engaged setup by caregivers**. The reliability of controls as a proxy for age depends on how accurately they are configured. Misstatements, lack of understanding or social pressure can weaken the value of the emitted signals.

A.45.6 Parental controls enable **private forms of access management**. These systems allow services to apply restrictions without requiring direct age or identity verification. Signals like "child-supervised account" can enable safeguards while minimising personal data collection.

A.45.7 **Contextual signals should not be reused without consent**. Signals generated by parental controls are relevant to the specific platform and use case. Reuse across services without clear consent increases the risk of unintended profiling or data misuse.

A.45.8 **Inclusivity and accessibility require ongoing attention**. While systems were generally consistent across demographics, some families – particularly those with limited digital literacy or different caregiving models – may face challenges in setup and maintenance.

A.45.9 Parental control signals **should not be treated as verified age data**. These tools provide useful contextual input but do not meet the assurance standards required for regulatory compliance. International standards reinforce that they serve as supplementary, not standalone, indicators.

A.45.10 **Platforms are seeking ways to integrate parental control signals**. Trial participants reported demand from service providers to use parental control status in access logic. Developing shared formats could support safer and more consistent implementation across systems.

Age Assurance Technology Trial

H

# PART H
# Parental Consent

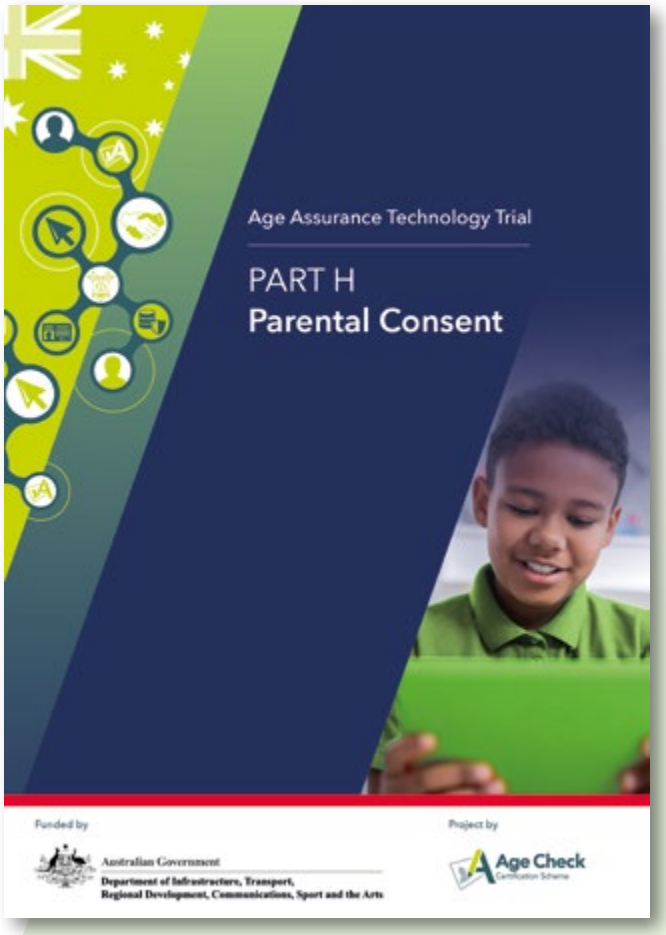


***See full report: Part H** – Parental Consent*

## A.46 Findings on Parental Consent

**A.46.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the topic of parental consent.

1. Parental consent systems **can be effectively applied** in Australia across different services and platforms.

2. Consent mechanisms offered **private, event-driven models** flowing from age assurance outputs; typically triggered at point of access.

3. **Design approaches varied significantly** across providers; evaluated systems ranged from lightweight verification to more formalised models involving ID checks.

4. Most **systems assumed conventional family structures**; did not routinely account for more complex guardianship arrangements.

**5** Long-term consent logging practices varied, with **implications for privacy and transparency**.

**6** Emerging innovations showed **potential to support more dynamic consent workflows**; may facilitate more responsive consent experiences.

**7** **Alignment with international standards** was evident, though implementation maturity differed.

**8** Consent was generally positioned as a **one-time event**, with limited ongoing interaction and designs focused on single transactions.

## A.47 What is Parental Consent

**A.47.1** A parental consent mechanism is a process that enables a parent or legal guardian to provide or revoke permission for a child to access digital goods, content, services, venues or spaces.

**A.47.2** Unlike parental control, which is configured in advance and operates continuously, parental consent arises in response to an age assurance trigger – typically when a child attempts to access something that requires age verification or compliance with legal or policy restrictions.

**A.47.3** Parental consent mechanisms typically involve five stages:

1. Identifying the parent or guardian, usually via account credentials, digital ID or other verified identity tools
2. Binding the parent or guardian to the correct child, confirming their legal relationship
3. Capturing informed consent for a specific action, such as joining a service, purchasing digital goods or engaging with age-restricted content
4. Communicating consent status to the relying party or service provider, often through a verifiable token or signal that the child has parental permission for the requested access
5. Providing a facility for consent to be revoked


***Figure A.47.1*** *Five Stages of Parental Consent*

**A.47.4** Parental consent mechanisms can be found in many online and offline services, such as:

- Online platforms: social media networks, multiplayer games, educational portals and content platforms often request guardian permission for children under a certain age
- Mobile and app ecosystems: app stores and in-app purchase systems may require verified consent before allowing downloads or transactions
- Offline environments: schools, healthcare providers or recreational venues (such as trampoline parks or soft play centres) may require guardian signatures or digital forms to authorise child participation in services or activities

## A.48 Introduction to Part H: Parental Consent

**A.48.1** Part H of the Age Assurance Technology Trial focuses on parental consent – a form of age assurance where a parent or guardian confirms a child's access to age-restricted goods, services or content, typically in digital environments. Unlike age estimation, inference or verification, parental consent does not seek to determine a user's age directly. Instead, it relies on the intervention of a responsible adult, who attests to the child's eligibility, often in response to an age-related trigger.

**A.48.2** Parental consent operates downstream of other age assurance methods. A user is typically flagged as a child (or possible child) through inference, estimation or declared age, after which a parent or guardian is asked to approve access or authorise an account. Parental consent thus acts as a decision point – not a measurement tool – and must be implemented with clear evidence of adult identity, informed consent and safeguards to prevent coercion, misrepresentation or circumvention.

**A.48.3** This part of the report examines how parental consent systems are designed, how they operate in real-world deployments and the extent to which they meet the requirements of emerging international standards – particularly ISO/IEC FDIS 27566-1 and IEEE 2089.1. These standards set functional expectations for parental involvement, identity binding, consent logging, data minimisation and the appropriate use of parental permissions across different risk contexts.

**A.48.4** The Trial was established to evaluate the technical feasibility and privacy implications of a wide range of age assurance methods in the Australian context. It does not make policy recommendations, nor does it seek to determine whether parental consent should be mandated for any particular use case. Rather, it addresses whether parental consent technologies are practically implementable, user-friendly, secure and reliable in supporting age assurance – particularly for children under regulatory thresholds such as 13, 15 or 18.

**A.48.5** This report explores the strengths and challenges of parental consent as a method of age assurance, including:

- How the consent process is initiated and verified
- How parent-child relationships are authenticated
- What safeguards are in place to protect the child's and guardian's data
- How systems prevent misuse or false assertions of parental status

**A.48.6** Importantly, we examine how well current technologies can balance the rights of children, responsibilities of guardians and expectations of relying parties, while ensuring the experience is accessible, inclusive and meaningful across diverse communities and service contexts.

## A.49 Summary of Parental Consent

**A.49.1** Parental consent mechanisms represent a widely recognised component of age assurance strategies, particularly in digital environments where children seek access to services or content subject to regulatory thresholds, such as those below the ages of 13 or 16. These mechanisms enable a responsible adult – typically a parent or legal guardian – to authorise a child's participation in age-restricted environments, usually after an initial trigger such as self-declared age or inferred risk. Unlike parental controls, which are configured in advance and applied continuously, parental consent is typically event-based, requiring an affirmative, verifiable action by an adult at a specific point in the child's user journey.

**A.49.2** The Trial evaluated a range of parental consent systems currently deployed or in development across Australian and international contexts. It found that these systems are technically feasible and can be effectively deployed using existing infrastructure. Participating providers demonstrated varied approaches to capturing consent, including email-based verification, credit card micro-payments, digital identity checks and token-based authorisation. While many of these mechanisms were already operational, their implementation styles varied in rigour, user experience and alignment with international frameworks.

**A.49.3** Across the evaluated systems, most implementations were designed around conventional, binary parent–child relationships. As a result, few consent models explicitly accommodated non-traditional caregiving arrangements, such as those involving foster carers, kinship care or shared parental responsibility. Similarly, most mechanisms were static in design, offering limited support for consent renewal, expiry or adaptation as the child matures. This often left little scope for recognising the evolving capacities of children or involving them meaningfully in the consent process.

## Key Statistics About Parental Consent

**12** participating providers with operational parental consent mechanisms

**3** systems offering consent revocation or expiry features

**Trial evaluation method:** Lab-based system walkthroughs and simulated user journeys

**2** systems supporting non-parental carers (e.g. foster, kinship)

| | |
|---|---|
| Common service contexts | Account creation, feature unlocking, digital purchases, content access |
| Consent mechanisms | 4 providers tested with scoped, risk-based decision-making |
| Data minimisation techniques | 4 providers showed mechanisms using pseudonymisation or data minimisation techniques |

**1** (exploratory) system enabling child participation or co-signing

**5** systems using implementations with indefinite consent logging

3 providers referencing international standards **ISO/IEC 29184, IEEE 2089.1**

**Technology Readiness Levels (TRL)**
Majority of Vendors **at TRL 7 or above**

**TRL 7-9**

***Figure A.49.1*** *Key Statistics from the Trial on Parental Consent*

**A.49.4** Verification of the adult's identity and legal authority varied in strength. Some systems relied primarily on self-declaration or account-based continuity, while others integrated more robust checks. Although several providers referenced alignment with standards such as ISO/IEC 29184 (Online privacy notices and consent) and IEEE 2089.1, practical application of principles like informed consent, accessibility and revocability differed across implementations.

**A.49.5** Emerging innovations – such as scoped, time-bound consent signals and privacy-preserving credential frameworks – indicated a growing maturity in the field, but also revealed challenges related to interoperability and ecosystem fragmentation. The retention and use of consent logs also varied, with some systems demonstrating strong privacy-by-design features (such as pseudonymisation or limited signal exposure), while others retained long-term records without clear boundaries on scope or re-use. This variability has implications for the privacy and data protection of both children and guardians.

**A.49.6** Overall, the Trial found that parental consent technologies are functionally mature and capable of supporting access governance where age-related restrictions apply. However, the consistency, inclusiveness and contextual adaptability of these mechanisms remains uneven. The findings suggest that while the technology underpinning parental consent is largely in place, further evolution in design, scope and implementation may be necessary to ensure these systems work equitably, proportionately and in support of both children's rights and service provider obligations.

## A.50 Who Participated in the Trial of Parental Consent Technology










## A.51 Observations About Parental Consent

A.51.1 **Parental consent systems demonstrated technical viability in the Australian context**. The Trial found that a variety of parental consent mechanisms were functional and implementable across services and platforms. Core elements – such as parent identification, child linkage and consent logging – were supported using existing infrastructure.

A.51.2 **Consent mechanisms offered private, event-driven models** flowing from age assurance outputs. In contrast to proactive parental controls, these systems were typically triggered at the point of access and enabled parents or guardians to make access decisions without the need for direct age verification of the child.

A.51.3 Design approaches varied significantly across participating providers. The systems evaluated ranged from lightweight verification (e.g. email loops) to more formalised models involving ID checks or credentialed consent tokens. **This variation affected consistency in how consent authority, accountability and revocation were handled**.

A.51.4 Most systems **assumed conventional family structures** and static relationships. The majority of implementations were structured around a single parent-child interaction and did not routinely account for more complex guardianship arrangements or evolving relationships. Systems also generally lacked features enabling child input into the process.

A.51.5 **Long-term consent logging practices varied**, with implications for privacy and transparency. While some systems used minimal audit trails, others retained detailed consent records over extended periods. In several cases, retention timelines and reuse of consent signals were not clearly bounded by time or context.

A.51.6 Emerging innovations showed potential to support **more dynamic consent workflows**. Some providers demonstrated credential-based or tokenised models of consent that included features such as time-bounded approval, scope limitation or revocation. These approaches may support more responsive or flexible consent experiences as services evolve.

A.51.7 **Alignment with international standards was evident**, though implementation maturity differed. Most participants referenced frameworks such as ISO/IEC 29184 and IEEE 2089.1. However, the extent to which consent mechanisms incorporated core elements – such as verifiability, informed action and accessibility – varied between implementations.

A.51.8 **Consent was generally positioned as a one-time event**, with limited ongoing interaction. Few systems enabled consent to be updated, refreshed or adapted over time as children grew older or circumstances changed. Most designs focused on a single transaction rather than a continuing parent–child–service relationship.

Age Assurance Technology Trial

J

# PART J
# Tech Stack

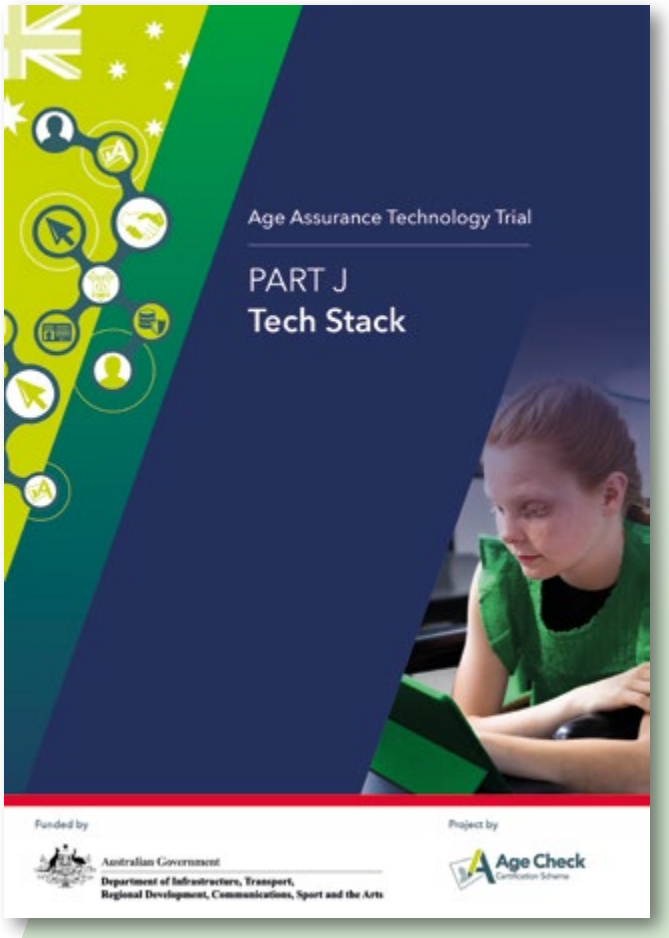


***See full report: Part J** – Tech Stack*

## A.52 Findings on the Tech Stack

**A.52.1** These are our headline findings. In line with the overall findings of the Trial, these findings relate specifically to the technology stack.

1. Technology stack deployment offers potential for **systemic and interoperable age assurance**, with potential for cross-cutting protections across services.

2. **App-store based models** are being developed but lack critical adoption and verification features.

3. Deployment at the **network or device level raises significant privacy and control considerations**.

4. **Interoperability solutions** are emerging but remain early-stage and non-standardised, resisting generalisation on functionality and maturity.

5 **Technology Readiness Levels (TRLs) vary widely**, with many solutions overstating maturity.

6 Functionality, performance, privacy and acceptability present **critical implementation challenges;** concerns include latency and public trust.

7 **Responsibility and liability** in a distributed tech stack are unclear and require further definition.

8 **Proximity to risk** is key to assessing effectiveness; location within the stack affects response to harmful content.

9 **Geolocation services** can play a role in detecting and preventing circumvention via VPNs.

## A.53 Summary of the Tech Stack

**A.53.1** This section of the Trial examined how age assurance, parental consent and control mechanisms could be embedded more systematically across the digital ecosystem by leveraging the technology stack – ranging from user devices and browsers to networks, app stores and backend services. The aim was to explore whether stack-level deployment could move beyond fragmented, service-by-service implementation and support more consistent, interoperable and privacy-conscious approaches to protecting children online.

**A.53.2** The evidence gathered through submissions, interviews and analysis suggests that while stack-based models offer real potential, their practical maturity is still limited. Most approaches remain conceptual or at early development stages and few are ready for scalable, real-world deployment. App store-based models were the most fully conceptualised, with companies like Meta and Snap proposing frameworks in which platforms collect and securely share age-related attributes. However, existing implementations by operators such as Apple and Google still rely primarily on self-declared or parent-entered information and do not incorporate independent age verification or support for open, cross-platform interoperability.

**A.53.3** Alternative models at the device or network level offer broader enforcement possibilities – especially for browser-based or unauthenticated services – but raise complex questions around privacy, data minimisation and user autonomy. These models, while promising in theory, must overcome significant compliance and usability barriers, particularly in environments where devices are shared or controls are imposed without user awareness.

**A.53.4** Several participants proposed approaches to interoperability, including reusable age credentials, digital wallets and orchestration layers that could work across services and jurisdictions. Although diverse in architecture, these models shared a common ambition: enabling a user to verify their age once and reuse that assurance in a privacy-preserving way. However, implementations remain fragmented, technically incompatible and often reliant on proprietary interfaces or ecosystem buy-in that has not yet materialised.

**A.53.5** A clear theme across the Trial was the mismatch between participants claimed Technology Readiness Levels (TRLs) and the actual state of deployment or integration. Many systems were rated as mid-to-high TRL despite lacking demonstrated interoperability, system-level testing or platform integration. This suggests that the field is still in an innovation phase, with most solutions yet to be validated in operational environments.

**A.53.6** In addition to technical challenges, the Trial identified a range of implementation issues relating to latency, reliability, transparency and user acceptability. Systems operating deep in the stack – such as at the network or browser level – may offer coverage, but risk alienating users through opaque controls or poor alignment with household realities. Similarly, systems that rely on parents to configure or enforce protections must account for digital literacy, language barriers and socio-economic context.

## A.54 What is the Tech Stack

### | Understanding the Term

**A.54.1** The term "technology stack" (or "tech stack") is commonly used in software development and digital infrastructure. It refers to the layers of technologies that work together to deliver a digital service. Think of it like a layered cake – each layer depends on the one beneath it and together they form a whole.

**A.54.2** In practical terms, the tech stack includes everything from the device you're using (like a phone or laptop), to the apps and browsers you access, to the networks that connect you to the internet and the platforms that host the content or services you use. Each layer performs a different role – but they are all interconnected.

### | How Does the Tech Stack Work?

**A.54.3** When a person opens an app or visits a website, many parts of the tech stack are involved:

| | |
|---|---|
| **The device** | Provides the interface (e.g. your phone screen and keyboard). |
| **The browser or app** | Interprets your input and displays content. |
| **The operating system** | Manages communication between your apps and hardware. |
| **The network** | Connects your device to remote services. |
| **The platform or app store** | Distributes the app and may apply rules or restrictions. |
| **The service provider** | Delivers the content or functionality – be it social media, gaming, shopping or streaming. |

**A.54.4** Each layer has its own role, but they can also be used strategically to apply controls or protections. This is particularly relevant in areas like **age assurance** and **parental consent**.

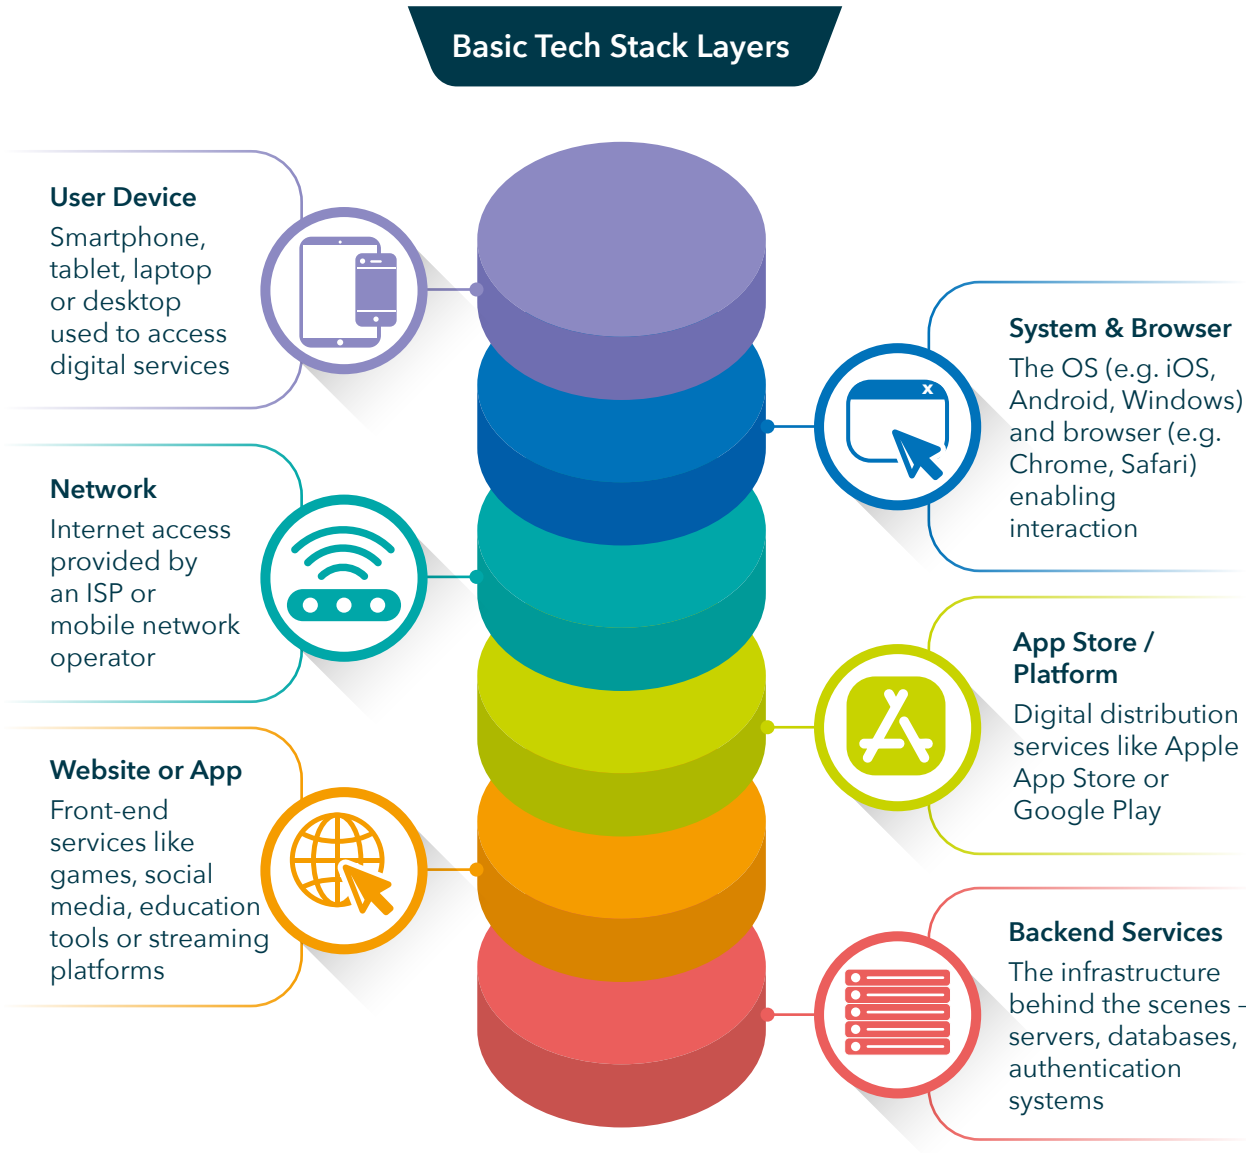


***Figure A.54.1*** *Basic Tech Stack*

## A.55 Who Participated in the Trial of the Tech Stack

euCONSENT

FastAs

GIP GENERAL IDENTITY PROTOCOL

GeoComply

Google

IDexchange

netsweeper

Opale.io

PRIVATELY

R2LABS

SHAYYPE

TikTok

## A.56 Observations About the Tech Stack

**A.56.1** Technology stack deployment offers potential for **systemic and interoperable age assurance**. Theoretical models indicate that placing age assurance mechanisms at different layers of the technology stack – such as on the user's device, within the network infrastructure or at the app-store level – could provide consistent and cross-cutting protections across services. Interoperability across components will be essential to realise this potential.

**A.56.2** **App-store based models** are being developed but lack critical adoption and verification features. While app-store based models were the most fully conceptualised, they currently rely on self-declared or parent-set age information. Without independent age verification and without adoption by key operators such as Apple and Google, these models do not currently meet the criteria for robust age assurance.

**A.56.3** Deployment at the **network or device level raises privacy and control considerations**. Implementing age assurance at the device or ISP level could enable broader coverage, including services accessed through browsers or third-party platforms. However, these approaches raise significant concerns regarding user privacy, autonomy and data protection compliance.

**A.56.4** **Interoperability solutions** are emerging but remain early-stage and non-standardised. Several Trial participants proposed mechanisms to support interoperability across different age assurance systems. These are still nascent and varied in design, making it difficult to generalise about their functionality or maturity.

**A.56.5** **Technology Readiness Levels (TRLs) vary widely**, with many solutions overstating maturity. A significant number of Trial participants reported higher TRLs than could be substantiated. Some conceptual solutions were rated as TRL 3 or higher without evidence of analytical validation or testing. Most interoperable tech stack models remain at a conceptual or early prototyping stage.

**A.56.6** Functionality, performance, privacy and acceptability present **critical implementation challenges**. Even theoretically promising models face practical threats to performance and adoption. Key concerns include latency, data handling practices, user transparency and public trust – particularly where technologies operate beyond the user's immediate control.

**A.56.7** **Responsibility and liability** in a distributed tech stack are unclear and require further definition. Where age assurance functions are spread across multiple technical layers and actors, accountability becomes diffuse. Without clear regulatory or contractual frameworks, there is a risk of ambiguity in liability, weakening enforcement and redress mechanisms.

**A.56.8** **Proximity to risk** is an important factor in assessing effectiveness. The location of the age assurance mechanism within the stack affects its ability to respond to harmful content or risky interactions. Solutions closer to the user or service (e.g., device-level or in-app) may offer more accurate contextual control but may also have narrower coverage.

**A.56.9** **Geolocation services** can play a role in detecting and preventing circumvention via VPNs. Some participants highlighted the potential for software to support age assurance systems by identifying when users attempt to mask their true location – such as through Virtual Private Networks (VPNs) – to bypass regional age restrictions and can then be required to use geolocation software to prove their real location. While promising, this approach raises its own challenges related to accuracy, evasion tactics and the implications for user privacy and cross-border service access.


Age Assurance
KJR

Age Assurance Technology Trial

# PART K
# Glossary, Bibliography & Literature Review

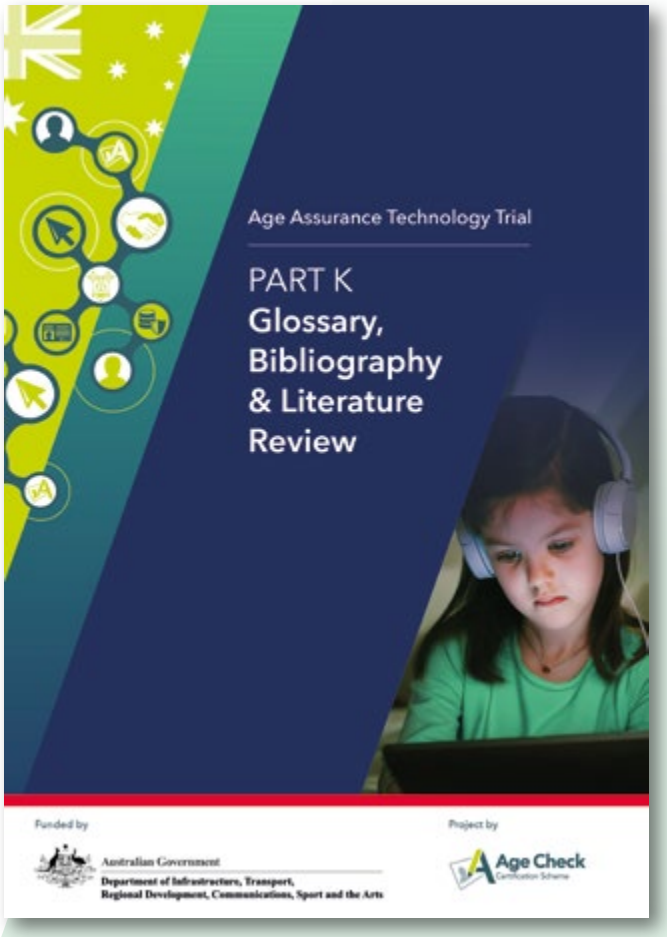


***See full report: Part K*** *– Glossary, Bibliography & Literature Review*

## A.57 Introduction to the Supporting Materials

**A.57.1** Part K of the Report provides the core reference materials that support the evaluation findings across all parts of the report suite. It serves as a foundational resource for readers seeking clarity on key terms, source materials and the broader evidence base that informed the Trial's design, execution and analysis.

**A.57.2** The Glossary within Part K standardises terminology used throughout the reports. It includes definitions for age assurance concepts (e.g. age verification, age estimation, age inference), key actors and roles (e.g. relying party, age assurance provider) and important terms from the Australian regulatory context, including references to the Privacy Act 1988 and the Online Safety Act 2021. Definitions have been drawn from international standards such as ISO/IEC FDIS 27566-1, as well as the Trial's own evaluative documentation.

**A.57.3** The Bibliography outlines the full range of materials consulted in the development of the Trial, structured hierarchically from legal sources to academic literature. It includes relevant legislation, regulatory frameworks, standards, research studies, white papers, industry submissions and commentary. This section provides a consolidated view of the global and Australian-specific knowledge landscape regarding age assurance technologies.

**A.57.4** Part K also contains the Literature Review, which summarises the current state of understanding around age assurance, including international practices, public attitudes, technical challenges and ethical considerations. It identifies evidence gaps and outlines how the Trial was designed to address them, with particular emphasis on Australian community needs and the inclusion of First Nations perspectives.

## | Trial website and data repository

Trial website
**www.ageassurance.com.au**

**A.57.5** The Trial's public webiste hosts key outputs and supporting materials, including the Trial overview, methodology summaries, published FAQs, stakeholder engagement updates and downloadable versions of the Final Reports. It also includes links to relevant standards, policy references and media coverage of the Trial. The site was designed to ensure transparency and facilitate public understanding of the technologies under evaluation.

**A.57.6** For those interested in accessing the Trial's underlying data and test documentation, structured outputs have been published to the Open Science Framework (OSF) at Age Assurance Technology Trial – OSF:


**A.57.7** This repository includes anonymised test datasets (where ethical and privacy requirements permit), evaluation templates, data dictionaries and audit records of methodology alignment with ISO and IEEE standards. This supports reproducibility, independent review and future research.

**A.57.8** Together, Part K and the associated digital resources offer a complete reference framework, ensuring that readers, researchers and policymakers can trace the basis of the Trial's conclusions and explore the wider landscape of age assurance technology evaluation.

## A.58 Project Team and Structure

**Tony Allen**
*Project Director*

**Andrew Hammond**
*Deputy Project Director*

WP1

**George Billinge**

WP2

**Dr Asad Ali &**
**Dr Dinindu Koliya Wedenage**

WP3

**Iain Corby**

WP4

**Dr Mark Pedersen**

WP5

**Rhianne Kiddle**

WP6

**Keith Robinson**
**& Nicola Elkin**

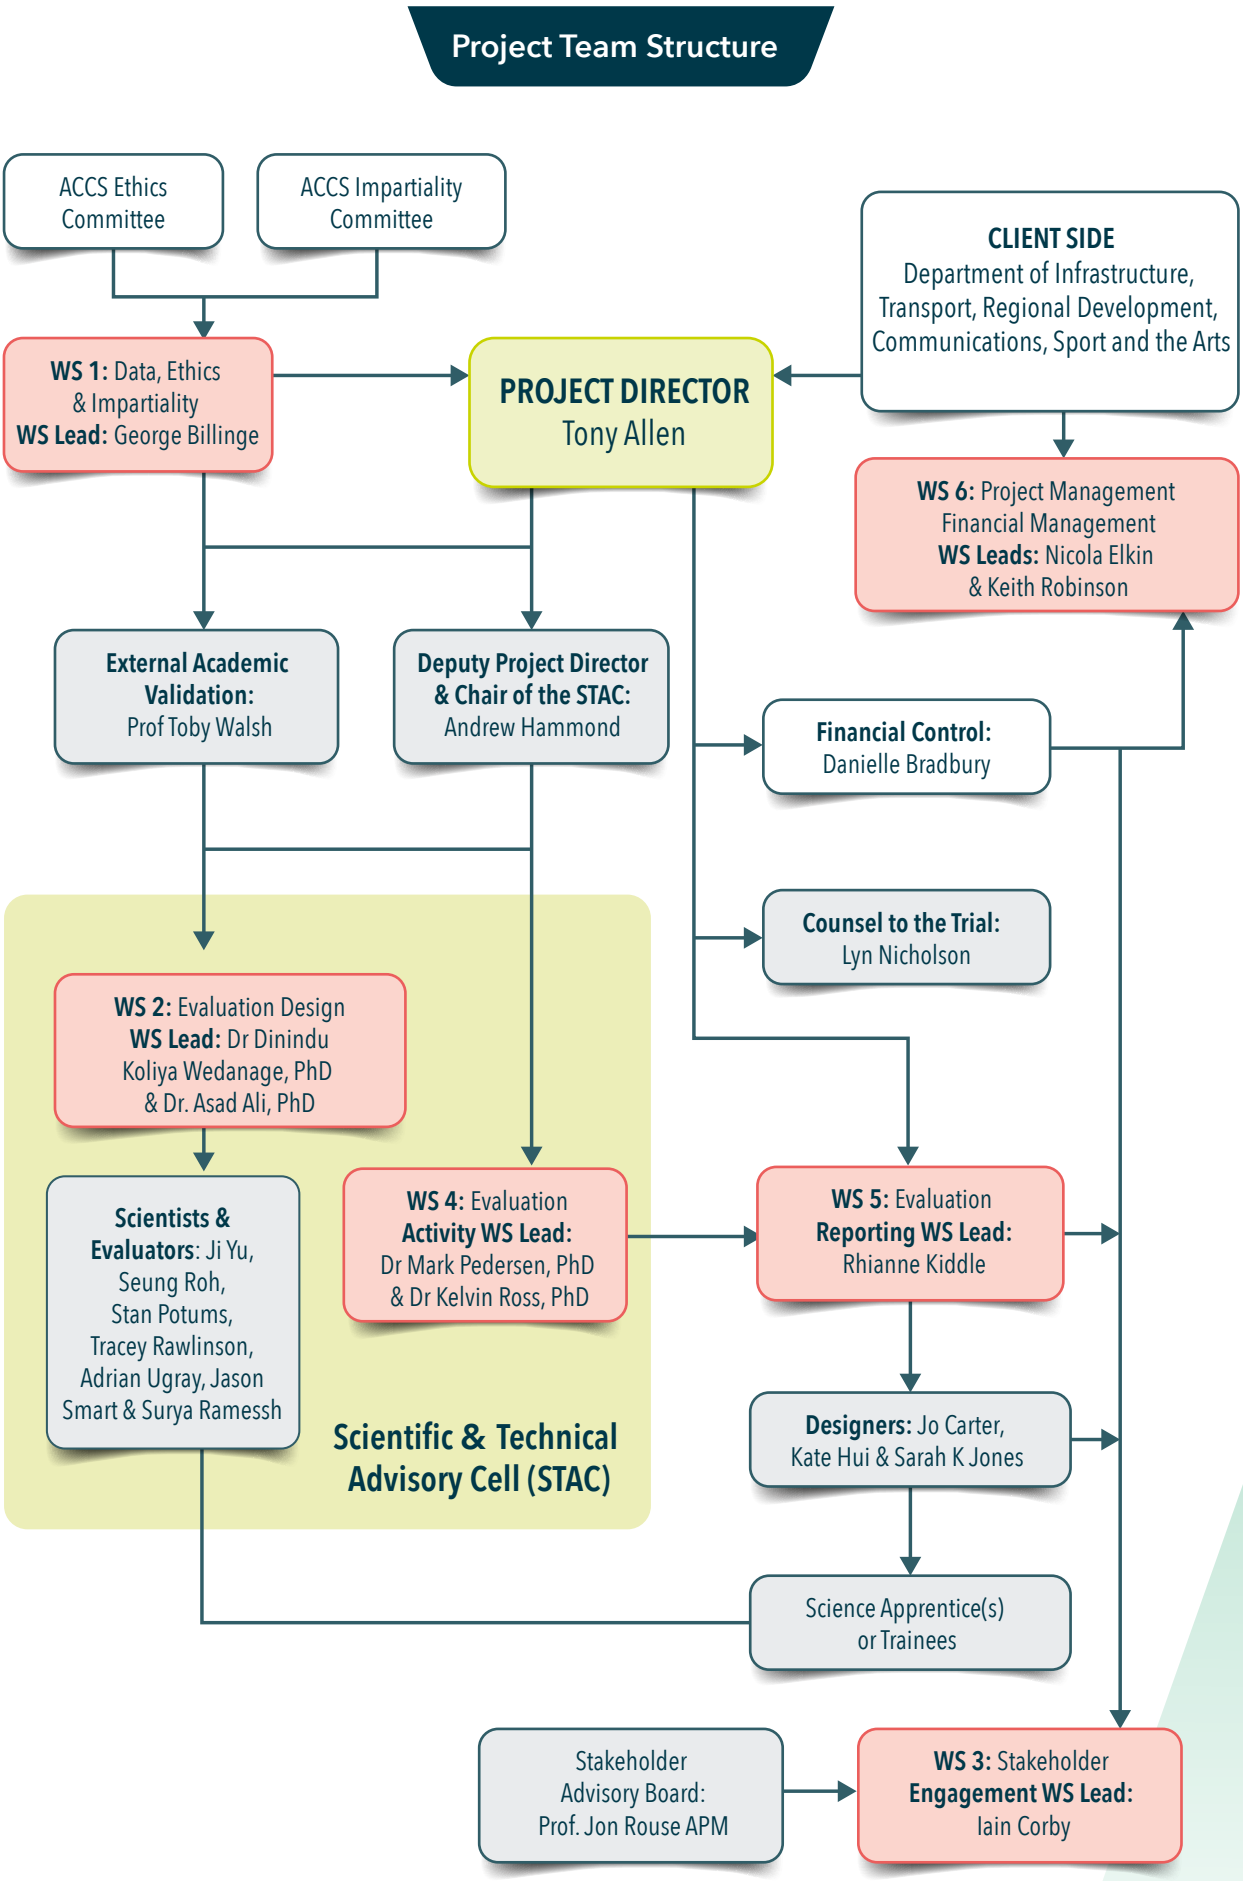


***Figure A.58.1*** *Project Team Structure*

**The world's leading independent conformity assessment body for age assurance technologies.** ACCS test that ID and age check systems work. ACCS are headquartered in the UK, but operate globally, including clients in Australia.

An Australian **software quality engineering consultancy**, that specialises in software testing and AI implementation.

**Leading data science, ethics, age assurance tech and scientists** from Koliya Group (AU) and Illuminate Tech (UK).

**Specialist Legal and Data Privacy Advisors**

**Creative and web development communications** from Heartburst (AU) and SoJo Creative (UK), industry engagement specialists, SafetyTech Limited (UK) and additional freelance graphic designers, mystery shopping providers and user experience analysts as needed.

## A.59 Statement of Impartiality

**A.59.1** The Age Check Certification Scheme (ACCS) is an independent, third-party conformity assessment body committed to the highest standards of integrity, objectivity and impartiality. Our involvement in the Australian Age Assurance Technology Trial was conducted without direction or influence from any Australian Government department, agency or regulatory body, although we sought their views and took them into consideration throughout the Trial. The Trial was funded following an open tender exercise carried out by the Department of Infrastructure, Transport, Regional Development, Communications and the Arts, which has been acknowledged on the front page of all reports from the Trial.

**A.59.2** In particular, ACCS affirms its full independence from the Australian Government's Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, as well as from the eSafety Commissioner, the Office of the Australian Information Commissioner (OAIC) and all other statutory or regulatory authorities in Australia.

**A.59.3** Our assessment, evaluation and reporting were guided solely by evidence gathered during the Trial and aligned with internationally recognised methodologies for conformity assessment, data ethics and digital safety. We did not seek or receive approval for the findings or language used in this report from any government body or external organisation.

**A.59.4** This independence is fundamental to our role. It enables ACCS to provide stakeholders with a transparent, objective and trustworthy account of the technologies evaluated, free from political, commercial or institutional bias. Our goal was to support an open dialogue about the role of age assurance technologies, grounded in factual evidence, robust analysis and a commitment to the rights and safety of users – particularly children and young people.

The **Age Assurance Technology Trial** is a landmark national initiative evaluating the real-world performance, privacy, usability and security of age assurance technologies. Commissioned by the **Australian Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts,** the Trial assessed 48 vendors and over 60 distinct technologies across various sectors, including social media, gaming, adult content and online retail. Through lab-based testing, interviews, analysis, school-based trials and mystery shopper evaluations, the Trial investigated how well different solutions could confirm, estimate or imply a user's age in ways that are secure, privacy-preserving and inclusive.

At the heart of the trial was one fundamental question: **Can age assurance be done?** The answer – based on thousands of data points, stakeholder interviews and international standards – is **yes, it can.** While no single solution fits all contexts, the Trial found that a wide variety of technologies already meet meaningful thresholds for accuracy, security and privacy when carefully selected and implemented. The report offers a comprehensive evidence base to support regulators, industry leaders and the broader public in shaping a safer, age-appropriate digital environment for all Australians.

Visit us on social media...

@AgeCheckCert

Age Check Certification Scheme

www.ageassurance.com.au

AVID Certification Services Ltd t/a Age Check Certification Scheme, registered in England 14865982 • Unit 321 Broadstone Mill, Broadstone Road, Stockport, SK5 7DL, United Kingdom • ABN 76 211 462 157

ISBN 978-1-0681646-0-6