Age Assurance Technology Trial

PART E
Age Inference

August 2025

Funded by

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

Project by

Age Check
Certification Scheme

Findings on Age Inference

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.

© Copyright of the Commonwealth of Australia

This document is available for reproduction on royalty-free, perpetual, attributed source, non-commercial rights to reproduce. This permission allows for unlimited reproduction and distribution, provided that proper credit is given to the original author(s) and source. This grant applies to all formats and media worldwide. For queries about commercial use or the licence in general, please contact the publishers. All rights to materials on links are reserved to the author of those materials.

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 978-1-0681646-4-4

Table of contents

Introduction and Overview

I

E.1 Introduction to Part E: Age Inference 6
E.2 Executive Summary 8
E.3 Who Participated in the Trial of Age Inference Technology 12

Context, Standards and Methodology

II

E.4 What is Age Inference 16
E.5 Evaluation Approach for Age Inference Systems 22
E.6 International Standards for Age Inference Methods 24

Detailed Analysis of Age Inference Findings

III

E.7 Age Inference Can Be Done 30
E.8 No Substantial Technological Limitations to Age Inference in Australia 41
E.9 Privacy-Conscious Implementation of Age Inference Systems 46
E.10 Transparency and Standards Alignment in Age Inference Practice Statements 50
E.11 Configuration Management and Threshold Sensitivity in Age Inference Deployment 53

E.12 Innovation and Research in Age Inference: Insights From the Australian Context 60

E.13 Privacy in Age Inference: Protection, Proportionality and Digital Footprint Management 63

E.14 Inclusion and Cultural Grounding in Age Inference 67

E.15 Improving Technological Foundations for Reliable Age Inference 77

E.16 The Future of Embedded Age Inference in Digital Environments 83

E.17 Enhancing Age Inference Through Parental Control Signals 87

E.18 Age Inference and Verified Credentials in Digital Wallets 92

E.19 Security and Information Protection in Age Inference Systems 99

E.20 Mitigating the Risk of False Legends and Manipulated Digital Footprints 103

E.21 Data Minimisation and Avoidance of Over-Collection in Age Inference Systems 107

Age Assurance Technology Trial

PART E
Introduction and Overview

E.1 Introduction to Part E: Age Inference

E.1.1 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.

E.1.2 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[1], which defines the privacy, purpose limitation and effectiveness expectations for age assurance systems and IEEE 2089.1[2], which outlines performance and interoperability requirements for age-related signals.

[1] *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.*

[2] *All references to IEEE 2089.1 throughout the suite of reports are referring to IEEE 2089.1–2024 – IEEE Standard for Online Age Verification.*

E.1.3 This part of the report presents the Trial's findings on age inference technologies, including their accuracy in different contexts, their performance across diverse populations (including those with limited identity documentation or digital histories), their capacity to support privacy-first deployments and their potential integration into broader digital ecosystems, such as interoperable wallets or in-app safety features. The analysis supports future development of evidence-based guidance, certification frameworks and trustworthy implementation models within Australia's evolving age assurance landscape.

E.2 Executive Summary

E.2.1 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.

E.2.2 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.

E.2.3 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.

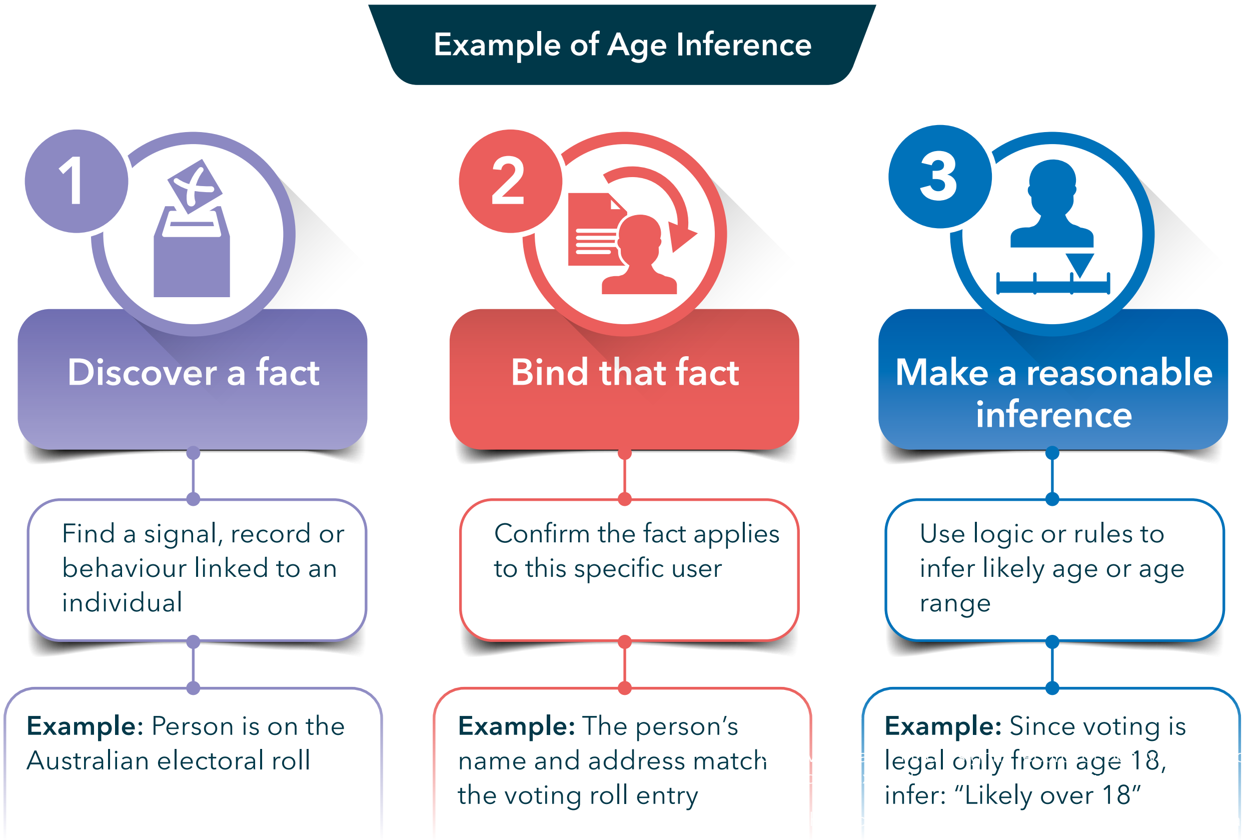

Example of Age Inference

1 Discover a fact

Find a signal, record or behaviour linked to an individual

Example: Person is on the Australian electoral roll

2 Bind that fact

Confirm the fact applies to this specific user

Example: The person's name and address match the voting roll entry

3 Make a reasonable inference

Use logic or rules to infer likely age or age range

Example: Since voting is legal only from age 18, infer: "Likely over 18"

Figure E.2.1 *Example of Age Inference*

E.2.4 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.

E.2.5 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

E.2.6 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.

E.2.7 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 E.2.2 *Key Statistics from the Trial on Age Inference*

E.3 Who Participated in the Trial of Age Inference Technology

E.3.1 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).

E.3.2 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

Age Assurance Technology Trial

II

PART E
Context, Standards and Methodology

E.4 What is Age Inference

E.4.1 Age inference is an age assurance method based on verified information which indirectly implies that an individual is over or under a certain age or within an age range.

E.4.2 It involves the following stages:

- Finding, locating, identifying or sourcing a fact or facts about an individual (other than their date of birth [DOB]) from a document, record, database, image, online activity or any other authoritative and reliable source.
- Binding that fact or facts to the correct individual (i.e. making sure that they relate to that actual individual).
- Analysing those facts or drawing an inference from a fact that can reasonably give an indication of the individual's age range.
- Communicating an age-related indication of that individual's age to a relying party (i.e. that they are in an age range, over or under a particular age threshold).

E.4.3 ISO/IEC FDIS 27566-1 defines age inference as:

> *"Age inference is an age assurance method based on verified information which indirectly implies that an individual is over or under a certain age or within an age range."*

E.4.4 It is particularly useful in low-friction or privacy-preserving contexts or where individuals do not have identity documents or decline to use biometrics. Its effectiveness depends on the quality and reliability of the facts and the logic of the inference.

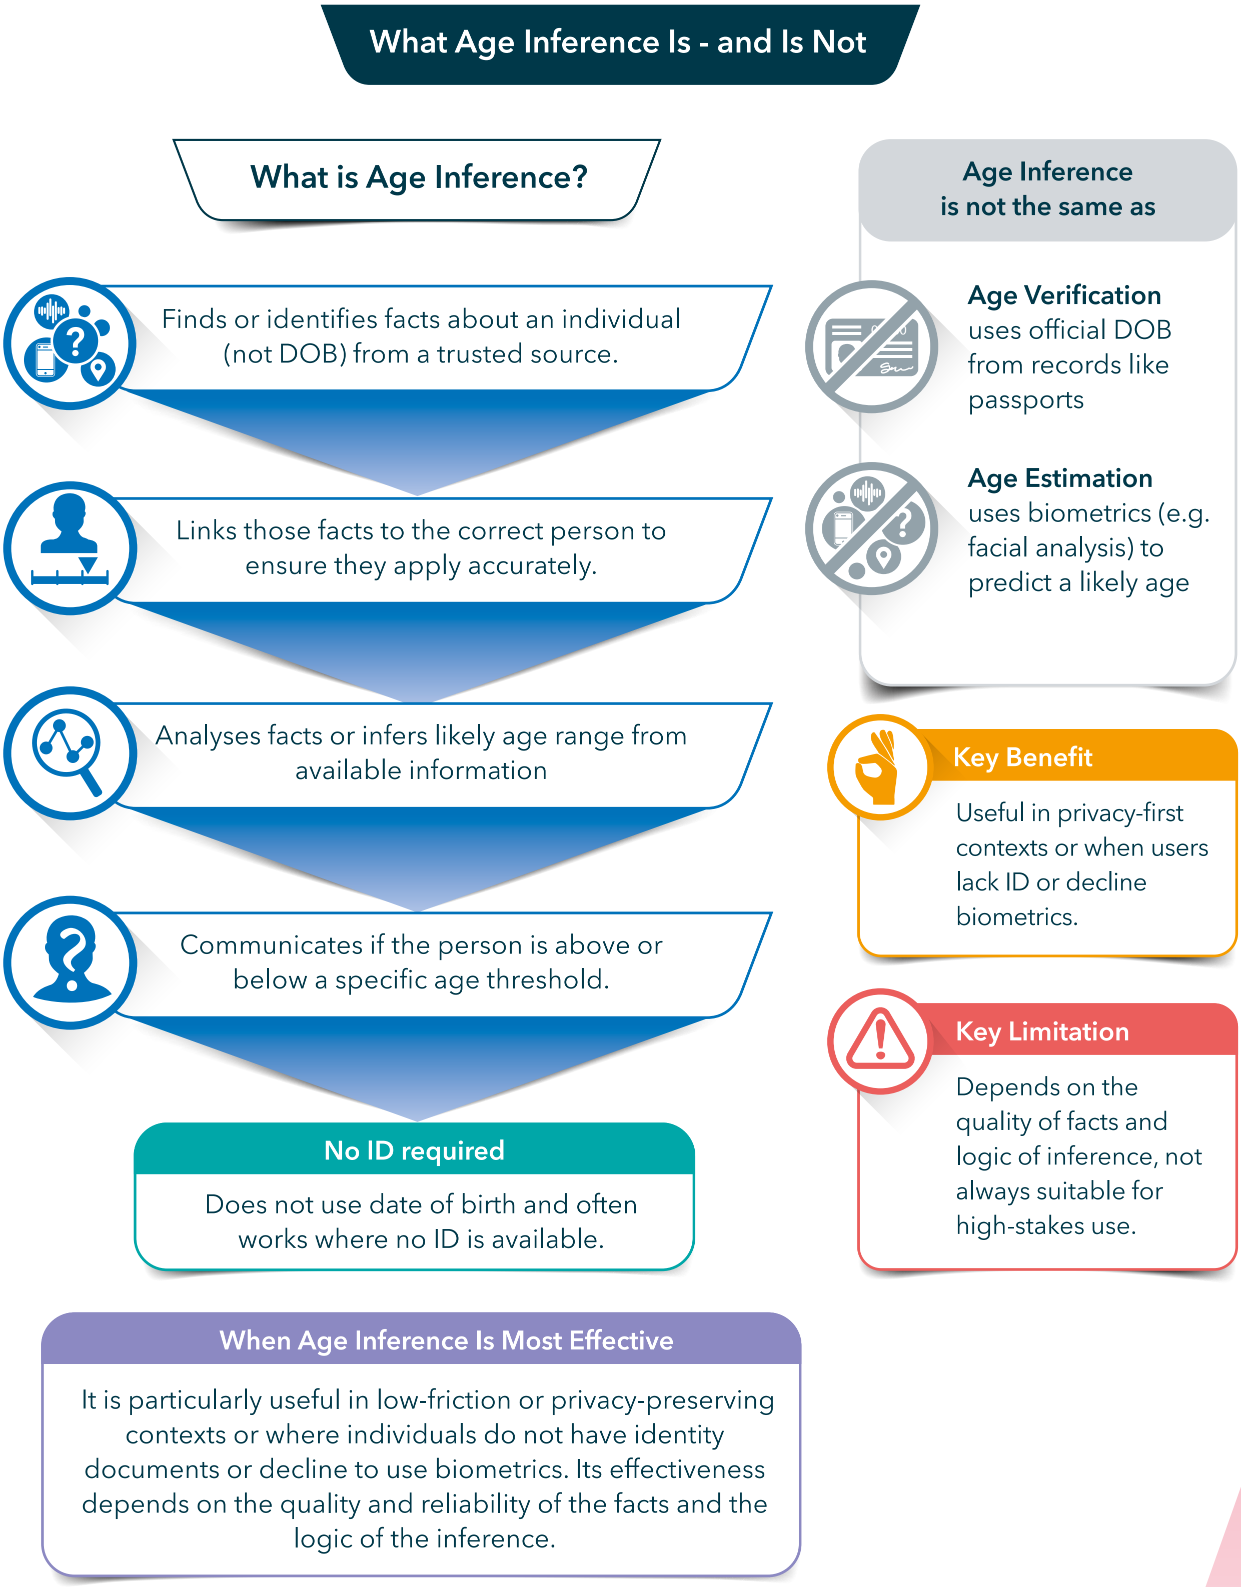

What Age Inference Is - and Is Not

What is Age Inference?

Finds or identifies facts about an individual (not DOB) from a trusted source.

Links those facts to the correct person to ensure they apply accurately.

Analyses facts or infers likely age range from available information

Communicates if the person is above or below a specific age threshold.

No ID required

Does not use date of birth and often works where no ID is available.

When Age Inference Is Most Effective

It is particularly useful in low-friction or privacy-preserving contexts or where individuals do not have identity documents or decline to use biometrics. Its effectiveness depends on the quality and reliability of the facts and the logic of the inference.

Age Inference is not the same as

Age Verification uses official DOB from records like passports

Age Estimation uses biometrics (e.g. facial analysis) to predict a likely age

Key Benefit

Useful in privacy-first contexts or when users lack ID or decline biometrics.

Key Limitation

Depends on the quality of facts and logic of inference, not always suitable for high-stakes use.

Figure E.4.1 *What Age Inference Is and Is Not*

Signature

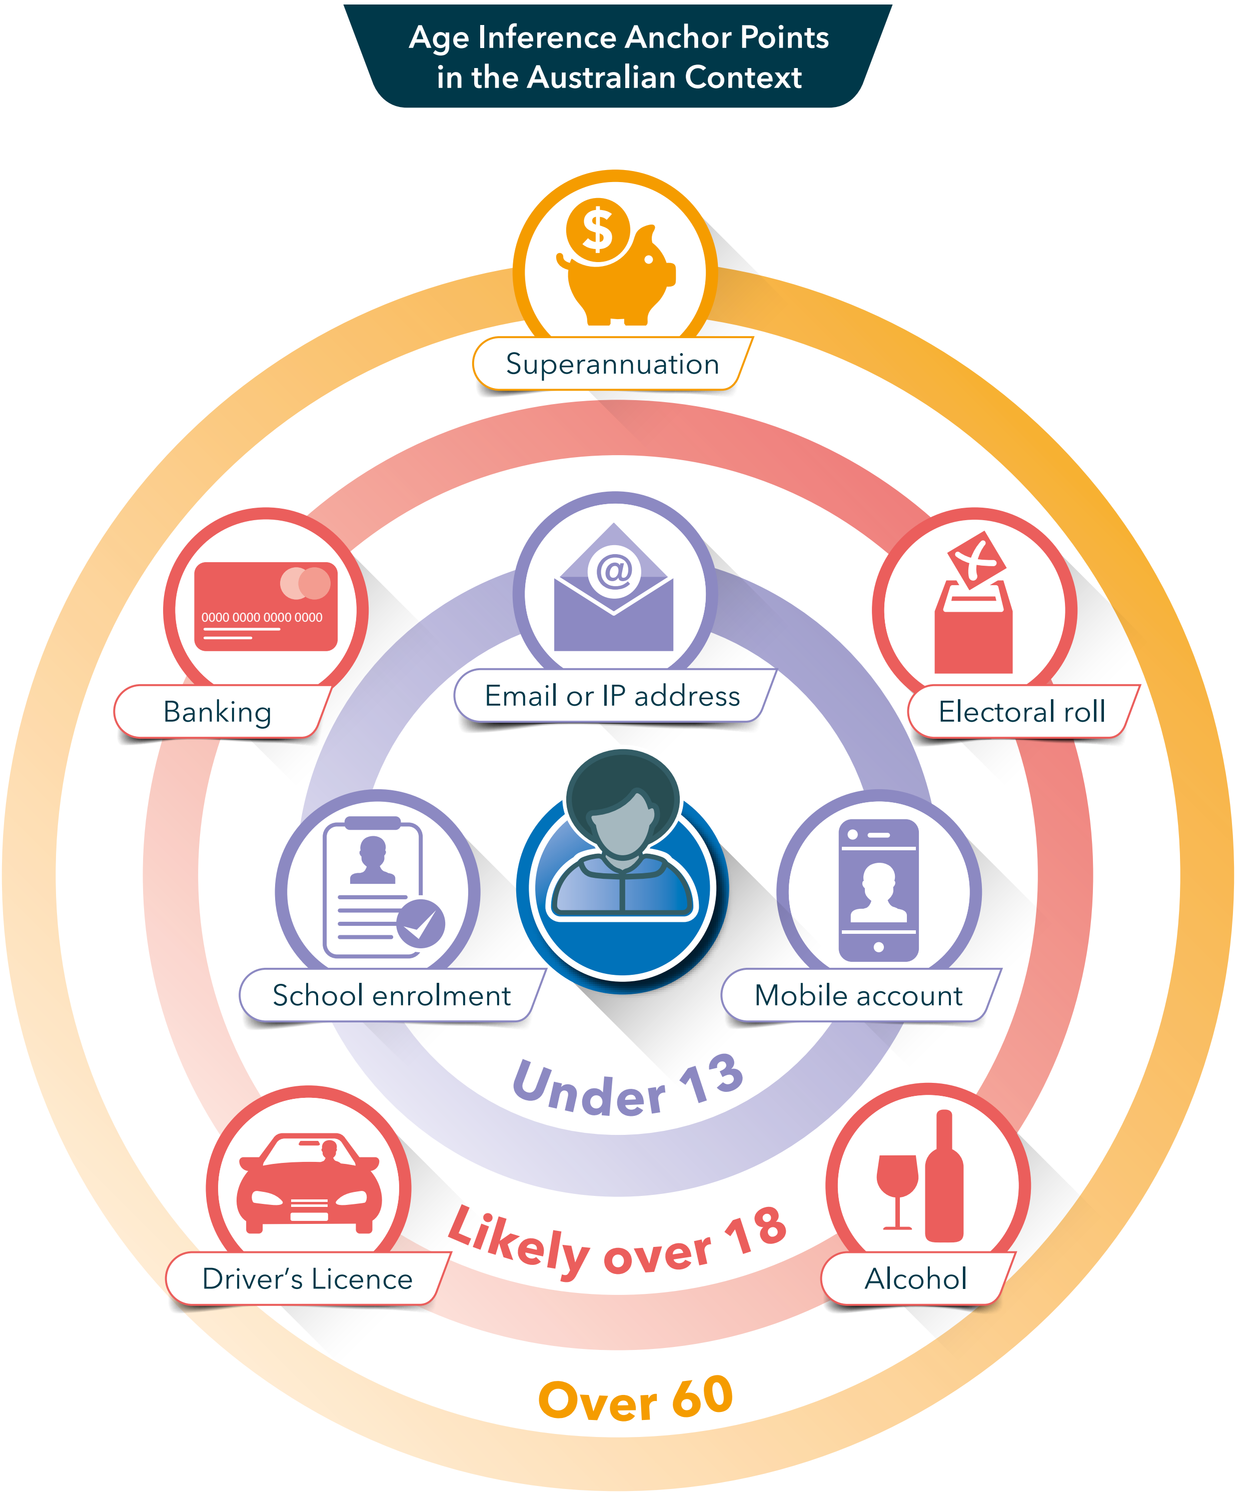

Figure E.4.2 Age Inference Anchor Points in the Australian Context

E.4.5 Age inference in Australia is highly viable due to the depth and accessibility of verifiable signals linked to individuals' life stages. When those signals are appropriately bound and interpreted with conservative, explainable logic, age inference becomes a powerful tool for scalable, privacy-sensitive age assurance, particularly in digital contexts where users may not wish – or be able – to provide formal ID.

Practical Australian examples of age inference

Inferred Fact	Source or Context	Legal Threshold	Age Inference
On the Australian electoral roll	Confirmed via name/address match	Must be 18+ to enrol and vote	Inferred: Likely over 18
Holds a commercial pilot licence	Verified aviation credential	Must be 21+ in Australia	Inferred: Likely over 21
Is enrolled in Year 9 in school	From education database or school access	Year 9 students are typically 13–15 years old	Inferred: Likely 13-15
Member of a seniors' card scheme	Based on card use or eligibility	Varies by state; generally 60+	Inferred: Likely over 60
Purchased a regulated product (e.g., online alcohol order with proof)	Past purchase record	Must be 18+ to buy alcohol	Inferred: Likely over 18
Has an account with a child-specific online platform (e.g., with COPPA-like registration)	Account metadata	Registered as under 13	Inferred: Likely under 13
Uses vocabulary or slang highly associated with preteens (via NLP*)	Language analysis	Model-based	Inferred: Likely under 12, low confidence

Notes and Clarifications

- *Voting age in Australia: 18+ (compulsory enrolment from age 18)*
- *Commercial airline pilot minimum age: 21+ (for ATPL – Air Transport Pilot Licence) as per CASA regulations*
- *Year 9 students in Australia: Typically 14–15 years old (some younger depending on school entry age)*

** This refers to Natural Language Processing.*

E.4.6 Each of these examples relies on an indirect fact or activity to generate a probabilistic age signal, not a certainty. As such, the strength of the inference depends on:

- The quality and the currency of the underlying information
- The strength of binding to the individual
- The logic applied in drawing the age-related conclusion

E.4.7 Age inference is a powerful and flexible technique that allows platforms or relying parties to make reasonable, risk-aware assumptions about a user's age, without requiring full identity verification or biometric analysis. When based on high-quality signals, applied proportionately and designed with explainability and fairness, age inference can be a key part of scalable, privacy-preserving age assurance.

E.5 Evaluation Approach for Age Inference Systems

E.5.1 *Core methodology*

Age inference systems were assessed through a structured desktop review, using a multi-dimensional criteria framework aligned to relevant international standards. The review was based on provider-submitted practice statements, privacy policies, vendor interviews and publicly available information, including material from provider websites.

E.5.2 *Evaluation criteria*

Age inference systems were assessed against a structured set of criteria:

Key Attributes	Criteria
Accuracy (as claimed)	The provider's stated ability to infer whether a user is likely above or below key age thresholds (e.g. 13+, 16+, 18+), based on available input signals.
Reasonableness of Inference	Whether the inputs described (e.g. school enrolment, account metadata, payment activity) logically support the age-related conclusions drawn.
Reliability (as described)	The degree to which the systems are presented as consistent across varying use cases, data contexts and operational environments.
Privacy and Security	Evidence of data minimisation, non-persistent data handling and separation of identity from inference logic, as described in documentation and interviews.
Demographic Inclusion	Claimed support for users with limited documentation or digital history, including attention to First Nations and other underrepresented groups.
Transparency and Explainability	Clarity in the provider's description of inference logic, confidence handling and the rationale behind automated decisions.
Technology Readiness Level (TRL)	An indicative assessment of maturity based on documentation, deployment claims and alignment with standards such as ISO/IEC FDIS 27566-1.

E.6 International Standards for Age Inference Methods

E.6.1 ISO/IEC FDIS 27566-1 formally recognises age inference as one of the methods available within an age assurance framework.

E.6.2 Key provisions from the standard include:

- Valid Method of Age Assurance: Age inference is acknowledged as a suitable approach and its effectiveness is seen as contingent on the quality of the underlying data and the reasonableness of the inference drawn from it. The standard does not classify it as inherently low assurance but stresses that fitness for purpose should be demonstrated through evidence and context.
- Risk-Based Application: The standard encourages age inference to be used proportionately, considering the level of risk associated with the service or decision. Where inference is relied on, it should be justified by the quality and appropriateness of the signals used and the outcomes should be statistically sound, explainable and fair.
- Transparency and Accountability: ISO/IEC FDIS 27566-1 requires that systems using inference be transparent about their use and provide users or guardians with appropriate information and challenge mechanisms where feasible. This includes describing the types of signals used and the rationale for inferring age-related eligibility.
- Bias, Fairness and Explainability: The standard stresses that any age assurance method – including inference – should be subject to evaluation for bias and should operate in a manner that is demonstrably fair across demographic groups. Inference systems should have documented reasoning and, where relevant,

demonstrate calibration and performance consistency.

- Layered Assurance Support: While inference may be used on its own where justified, the standard also supports its use as part of a layered approach, where it complements other methods (e.g. estimation or verification) or is used to prompt further checks based on confidence levels or contextual risk.

Cross reference: *Part F – Successive Validation*

E.6.3 In short, ISO/IEC FDIS 27566-1 treats age inference as a versatile and powerful age assurance tool, capable of supporting high-quality age-related decisions – provided it is implemented with clear logic, accountable governance and appropriate safeguards. It does not limit the method's assurance level by default, but instead emphasises that its effectiveness depends on context, configuration and clarity of implementation.

E.6.4 *Review methods*

Key Attributes	Criteria
Document and Interview Review	All evaluations were based on providers' practice statements, privacy policies, interviews and public domain information. No hands-on testing of systems was conducted.
Policy and Design Analysis	Evaluators analysed privacy protections, configuration options (e.g. thresholds, fallback logic) and how inference results are intended to be used by relying parties.
Inclusion and Accessibility Claims	Providers' responses were reviewed for alignment with inclusivity principles, including any stated efforts to support cultural or contextual age inference methods.
Standards Alignment	Systems were reviewed for how closely their documentation and described practices adhered to the principles and requirements of ISO/IEC FDIS 27566-1, IEEE 2089.1 and other applicable standards.

This assessment approach allowed the Trial to examine the readiness, privacy posture and conceptual integrity of age inference systems as described by participating vendors – without performing direct functional testing or independent technical validation.

E.6.5 *Demographic inclusion and signal validity*

- Systems were reviewed based on provider documentation and interviews for how they propose to support populations with limited digital or documentation footprints, including First Nations and Torres Strait Islander communities.
- Particular attention was given to any stated or planned use of culturally grounded or alternative inference pathways (e.g. school enrolment, community indicators or non-biometric life-stage signals).

E.6.6 *Usability and integration*

- Provider materials were reviewed to assess how clearly age inference outputs are described (e.g. "Likely Over 18") and whether outputs are designed to be easily integrated by relying parties.
- The Trial also examined claims regarding configuration flexibility, including fallback to higher-assurance methods when inference confidence is low.

E.6.7 *Static and policy reviews*

- The Trial conducted a desktop review of each provider's privacy and security posture, including safeguards described for data minimisation, purpose limitation and user privacy – against relevant ISO/IEC FDIS 27566-1 clauses.
- Policies on data retention, digital footprint management and user rights were analysed through submitted documents and interviews.

E.6.8 *Limitations*

- The evaluation did not involve functional testing, penetration testing or adversarial robustness trials.
- No validation of system accuracy was conducted using synthetic or real-world data inputs.
- The review focused on threshold-based inference use cases (e.g. under/over 18), rather than continuous or longitudinal age modelling.

Age Assurance Technology Trial

PART E
Detailed Analysis of Age Inference Findings

E.7 Age Inference Can Be Done

| Summary finding

E.7.1 Age inference can be done in Australia, as there are a wide range of verifiable facts about individuals that exist in documents, records, civil infrastructure or day-to-day activities of Australians – such as electoral registration, financial footprint data, school enrolment records, driver's licences, medical data and participation in age-specific services or transactions using verifiable anchor points, such as e-mail addresses, mobile telephone numbers or address data.

| Detailed analysis

E.7.2 Age inference is a feasible and increasingly practical method of age assurance in Australia, thanks to the widespread availability of verifiable, life-stage-linked facts embedded in daily life, government systems and commercial services. These facts – when accurately bound to an individual – can be used to draw strong, reasoned conclusions about a person's likely age or age range, often without requiring access to full identity documents or biometric analysis.

E.7.3 Australia's robust digital and civil infrastructure provides a rich ecosystem of indirect signals that can support inference systems. These include both structured datasets and behavioural footprints that are unique, contextual and legally regulated.

Functional characteristics (ISO/IEC FDIS 27566-1 Reference)

E.7.4 Age inference technologies exhibit the core functional characteristics described in ISO/IEC FDIS 27566-1:

ISO/IEC FDIS 27566-1	Criteria
Friction Minimisation (Clause 5.3.2)	Age inference enables services to assess a user's likely age without requiring identity documents, biometric capture or complex user journeys. By analysing contextual or behavioural signals already present in the interaction – such as email metadata, account activity or device settings – age inference supports low-friction, seamless user experiences.
Selective Disclosure (Clause 8.3.2)	Inference systems typically output minimal information, such as a binary or threshold-based indicator (e.g. "Likely Over 18: Yes"), without revealing the underlying data or exact inferred age. This supports privacy-by-design and ensures that users are not exposed to unnecessary data sharing.
Appropriateness to Risk (Clause 4.3.4)	Age inference is particularly well suited to low- and moderate-risk environments where formal identification is disproportionate or unavailable – such as access to age-gated content, app store downloads or purchase eligibility checks. Its use helps balance user autonomy with protective access controls in these settings.
Adaptability and Context Awareness (Clause 9.6)	Providers demonstrated configurable systems that can adapt to the context and needs of the relying party – for example, by adjusting confidence thresholds, layering with fallback methods (like age estimation or verification) or tailoring signal combinations to suit sector-specific risks.

Deployment contexts in Australia

E.7.5 **Availability:** Most Australians engage regularly with structured services (e.g. banking, education, taxation, licensing), leaving behind verifiable, timestamped records.

E.7.6 **Regulatory backing:** Many signals – such as voting eligibility or driver's licences – are governed by law and linked to age-specific thresholds.

E.7.7 **Reusability:** Common identifiers such as email addresses, mobile numbers or residential addresses can act as binding anchors across systems when appropriately matched.

E.7.8 **Temporal value:** Participation in life-stage-linked activities (e.g. superannuation or immunisation) often correlates with known age brackets.

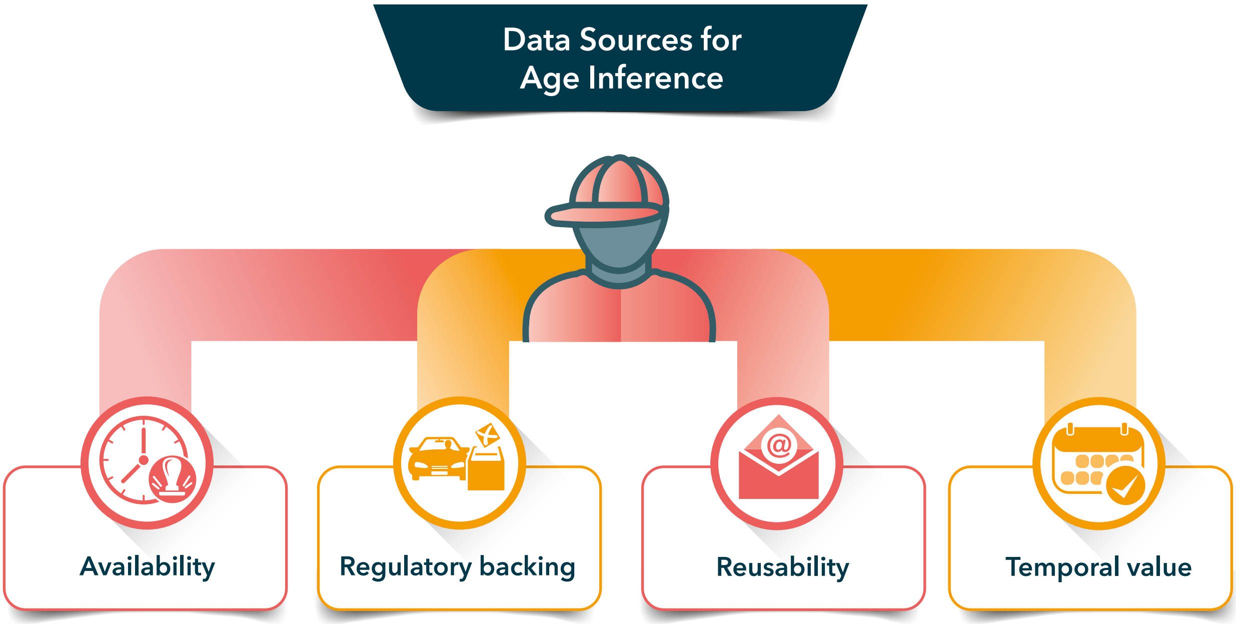

Figure E.7.1 Data Sources for Age Inference

E.7.9 When inference systems use reliable signals, apply conservative logic and operate with transparency, they can support statistically robust and context-appropriate age-related decisions – especially when paired with confidence thresholds and fallback mechanisms.

Reasonableness of inference: Evaluation approach and supporting evidence

E.7.10 In evaluating age inference methods, the Trial applied a structured framework to assess the reasonableness of the conclusions drawn from non-identity-based facts. Unlike age verification (which uses a known date of birth) or biometric estimation, age inference draws on indirect indicators – such as school enrolment, credit card ownership, account tenure or usage patterns – to suggest a likely age or age range.

E.7.11 Given the desktop-review nature of the Trial, reasonableness was assessed through the documentation and claims provided by vendors, focusing on whether the logic of their systems appeared statistically sound, contextually appropriate and ethically responsible. The following three factors guided this evaluation:

Relevance of the underlying fact

E.7.12 The Trial considered whether the input signals used had a plausible, evidence-based correlation with age. Strong signals included:

- **Electoral roll registration**, which legally requires individuals to be aged 18 or over in Australia.
- **Credit card ownership**, which similarly implies adult status under Australian financial regulations.
- **School year enrolment**, such as Year 9 attendance, which typically corresponds to the 13–15 age bracket.

E.7.13 Conversely, signals that lacked clear age correlation – such as app usage preferences or general browsing behaviour – were treated with caution unless accompanied by additional contextual indicators or fallback logic.

| Strength and consistency of the correlation

E.7.14 Inferences based on strong, binary signals (e.g. verified credit card transaction) were considered more reliable than those based on broader behavioural trends. Providers submitted documentation indicating:

- **High-confidence use of binary triggers**, such as credit card verification or electoral enrolment matching.
- **Moderate-strength signals** for less distinct thresholds, such as school attendance or military service.
- **Weaker, contextual signals**, like content engagement or account activity patterns, were typically used in combination or as early screening layers.

E.7.15 A comparative table of signal strengths was included in the Trial's findings, highlighting which indicators were most suitable for specific age thresholds (e.g. 13+, 16+, 18+).

| Degree of certainty expressed

E.7.16 The Trial reviewed how vendors communicated confidence in their inferences. High-quality submissions:

- Applied **buffer thresholds** around critical ages (e.g. disregarding inferences near 17.5–18.5 yrs, escalating to biometric estimation).
- Presented outputs as **categorical or probabilistic signals** (e.g. "Likely Under 13"), rather than absolute classifications.
- Demonstrated alignment with ISO/IEC FDIS 27566-1 Clause 6.4, which encourages proportionate use of inference and the availability of fallback or escalation mechanisms.

E.7.17 Some providers disclosed detailed **confidence scoring**, fallback logic or decision rules in their practice statements, often using flowcharts or pseudocode to clarify how inferences were derived and interpreted.

Observed good practice across the Trial

E.7.18 In addition to structured logic, the Trial found strong alignment with good practice in how providers addressed uncertainty and edge cases:

- **Single-variable inference** was discouraged unless based on a high-certainty, legally backed signal.
- **Fallback mechanisms** – such as escalation to age estimation or additional signals – were commonly described for ambiguous results.
- Vendors showed awareness of **demographic fairness**, including early-stage exploration of culturally grounded inference signals relevant to First Nations users.
- Providers also showed evidence of **data minimisation**, avoiding persistent profiling or over-collection when drawing inferences.

E.7.19 This evaluative framework enabled the Trial to assess age inference methods on their conceptual robustness and standards alignment, even in the absence of direct functional testing. It also reinforced the need for age inference to be implemented transparently, proportionately and with strong governance – particularly where automated decisions may affect access to age-restricted services or content.

Vendor Case Study

verifymy

Website

verifymy.io

Verifymy participated in the Trial as an age inference provider using metadata, email domains and app interaction context. Security practices include internal ISO 27001 frameworks and session-based logic evaluation.

Three Key Facts

1. Email domain (e.g.,ac.uk,.edu.au). Device and usage context. Declared user preferences and enrolment flows.

2. Threshold-based inference with "Likely Over" outputs. Used fallback to age estimation when email metadata was insufficient.

3. Signals anonymised and discarded post-session. Ad hoc data handling audits implemented. No identity or PII collection used.

Strengths

- Suitable for use cases with email registration
- Modular API structure
- Fast and low-cost to implement for relying parties

Practice Statement

ageassurance.com.au/v/vmy/#PS

Privacy Policy

ageassurance.com.au/v/vmy/#PP

Technology Trial Test Report

ageassurance.com.au/v/vmy/#TR

Technology Trial Interview

ageassurance.com.au/v/vmy/#VI

Summary of Results

Verifymy provided a lightweight, efficient inference service grounded in contextual metadata. Its non-reliance on identity documents and session-only logic makes it suitable for low-assurance age thresholding applications.

Caveats and considerations

E.7.20 While Australia's infrastructure supports high-quality inference, the strength of the inference depends on:

- Binding accuracy: Ensuring that the fact or data point truly relates to the user in question (e.g. not just a shared household)
- Timeliness: Outdated records (e.g. school enrolment from 5 years ago) may not reflect current age
- Fairness and inclusion: Some groups (e.g. new migrants, remote First Nations communities) may not have consistent participation in certain systems

E.7.21 These challenges highlight the need for clear documentation, fallback paths and optionality for users, as supported by ISO/IEC FDIS 27566-1 Clause 8.2 (Privacy by Design) and Clause 7.2 (Indicators of Confidence).

E.7.22 Age inference in Australia is highly viable due to the depth and accessibility of verifiable signals linked to individuals' life stages. When those signals are appropriately bound and interpreted with conservative, explainable logic, age inference becomes a powerful tool for scalable, privacy-sensitive age assurance, particularly in digital contexts where users may not wish – or be able – to provide formal ID.

| Why age inference is important

E.7.23 Age inference is a versatile, low-friction method of age assurance that offers a practical and privacy-conscious way to assess whether an individual is likely to meet a given age threshold – without requiring formal identity documents, biometric data or declared dates of birth. Instead, it draws on existing contextual or behavioural signals to support age-related decisions. This makes age inference especially valuable in moderate-risk or high-volume environments, such as social media platforms, streaming services, online marketplaces and educational or community services.

E.7.24 While it does not deliver the binary precision of date-of-birth verification, age inference can still provide highly effective and reasonable age judgments, particularly were strong, verifiable indicators – such as financial activity, school enrolment or long-term service usage – support confident conclusions. When supported by clear thresholds, fallback logic and robust integration practices, age inference offers scalable, context-aware solutions for platforms that must balance access, safety and inclusion.

E.7.25 When implemented with privacy-first, standards-aligned and bias-aware principles, age inference offers a practical and proportional method for age assurance that supports:

- Risk-appropriate access control
- User privacy and autonomy
- Operational efficiency and real-world deployment

E.7.26 Its adaptability means that age inference can be embedded into real-time digital interactions and may be layered with other assurance methods – such as estimation or verification – depending on context. As interoperability increases across identity frameworks and digital wallets, inference-derived age signals (e.g. "Likely Over 18") may be incorporated into broader ecosystems, helping to expand access to age-restricted services without unnecessary intrusion or identification.

E.7.27 As confidence grows in the reliability and ethical use of inference-based systems, age inference is emerging as a critical tool in the age assurance landscape, offering inclusivity, proportionality and practicality – particularly for users who may lack formal ID or are excluded by more intrusive methods.

Provider	Description
PRIVO® Privacy • Permission • TRUST	Mature platform (PRIVO iD) actively used in regulated services (e.g. COPPA), with multiple age inference methods (contextual data, parental consent metadata). TRL 8 reflects mature system with strong documentation, but limited clarity on direct deployment in Australian context.

TRL 8

Website
privo.com

Practice Statement
ageassurance.com.au/v/pvo/#PS

Technology Trial Test Report
ageassurance.com.au/v/pvo/#TR

Privacy Policy
ageassurance.com.au/v/pvo/#PP

Technology Trial Interview
ageassurance.com.au/v/pvo/#VI

verifymy	System incorporates email domain analysis and metadata inference and is deployed in moderated content settings. While operational, its age inference capability is narrower and more targeted (e.g. “likely under 13”) and evidence suggests it is part of layered access control.

TRL 7

Website
verifymy.io

Practice Statement
ageassurance.com.au/v/vmy/#PS

Technology Trial Test Report
ageassurance.com.au/v/vmy/#TR

Privacy Policy
ageassurance.com.au/v/vmy/#PP

Technology Trial Interview
ageassurance.com.au/v/vmy/#VI

E.8 No Substantial Technological Limitations to Age Inference in Australia

E.8.1 Age inference technologies in Australia face no substantial technological limitations to implementation.

E.8.2 The Trial found that age inference technologies can be implemented effectively in Australia, with no substantial technological limitations preventing their deployment or operation in real-world conditions. Trial participants were able to demonstrate operational systems – some already in production – that used verifiable behavioural, contextual or account-based signals to derive a user's likely age or age range.

E.8.3 Age inference solutions were shown to be technically feasible, computationally efficient and readily deployable using standard infrastructure such as cloud-based machine learning models, edge-device analytics or API integration with relying party platforms.

| Worked examples from age inference providers

E.8.4 Example: In-platform behaviour analysis

IDverse™
A LexisNexis® Risk Solutions Company

- One Trial participant implemented age inference based on real-time user interactions – including text input, navigation speed, help request patterns and engagement duration.
- The system detected patterns strongly associated with pre-teen behaviour and flagged users as likely under 13, triggering escalation to a more robust verification step.
- The underlying model operated in-browser with no significant latency or bandwidth impact, demonstrating that inference can be deployed without complex infrastructure changes.

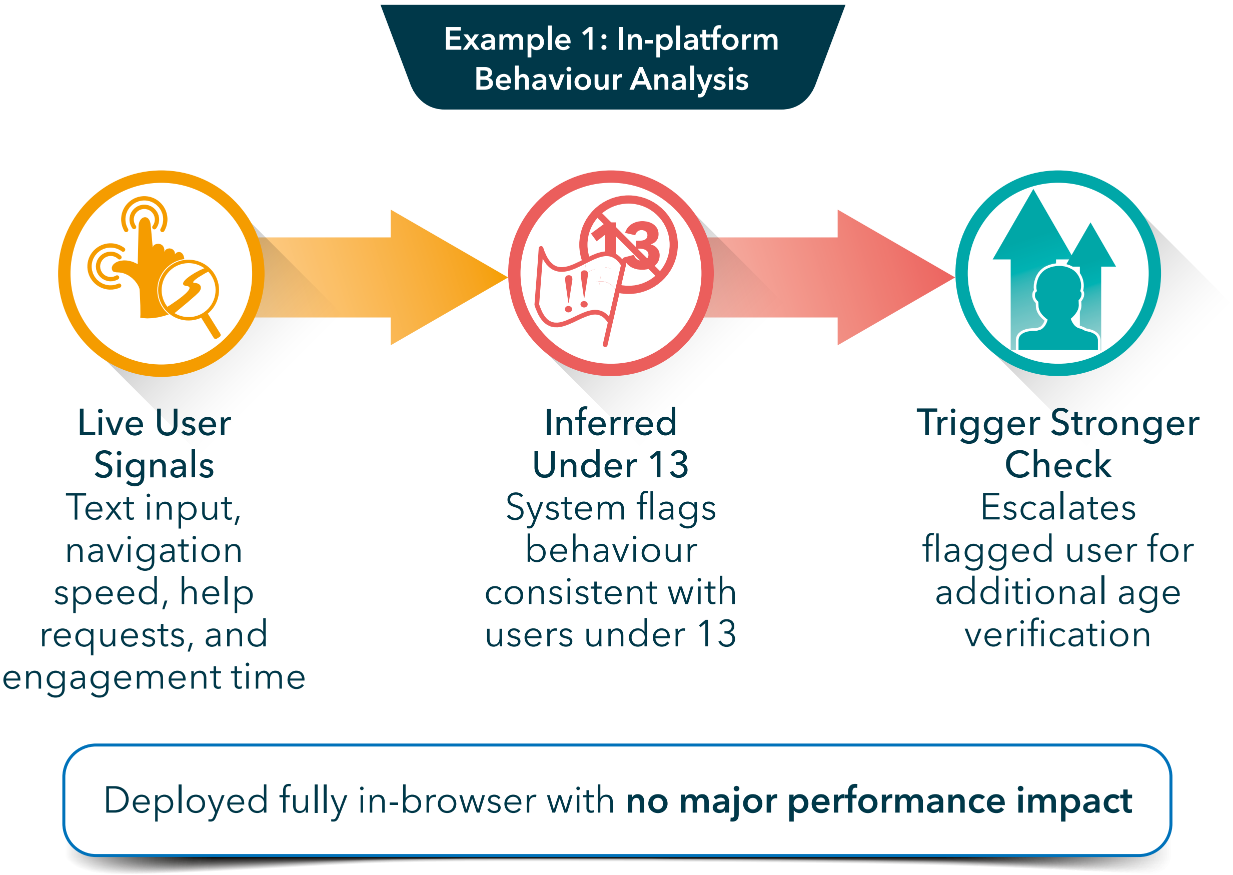

Example 1: In-platform Behaviour Analysis

Live User Signals
Text input, navigation speed, help requests, and engagement time

Inferred Under 13
System flags behaviour consistent with users under 13

Trigger Stronger Check
Escalates flagged user for additional age verification

Deployed fully in-browser with **no major performance impact**

Figure E.8.1 *In-platform Behaviour Analysis – IDVerse*

E.8.5 Example: Account metadata and self-asserted data validation

verifymy

- Verifymy analysed contextual metadata such as email domain type, account creation age, device usage and user-declared input sequences to infer likely age bands.
- Accounts registered using school-based email domains or patterns of shared/family device use were flagged as likely to belong to minors, particularly those under 13.
- The approach required no identity documents or biometric data and operated within Verifymy's existing moderation and access-control framework.

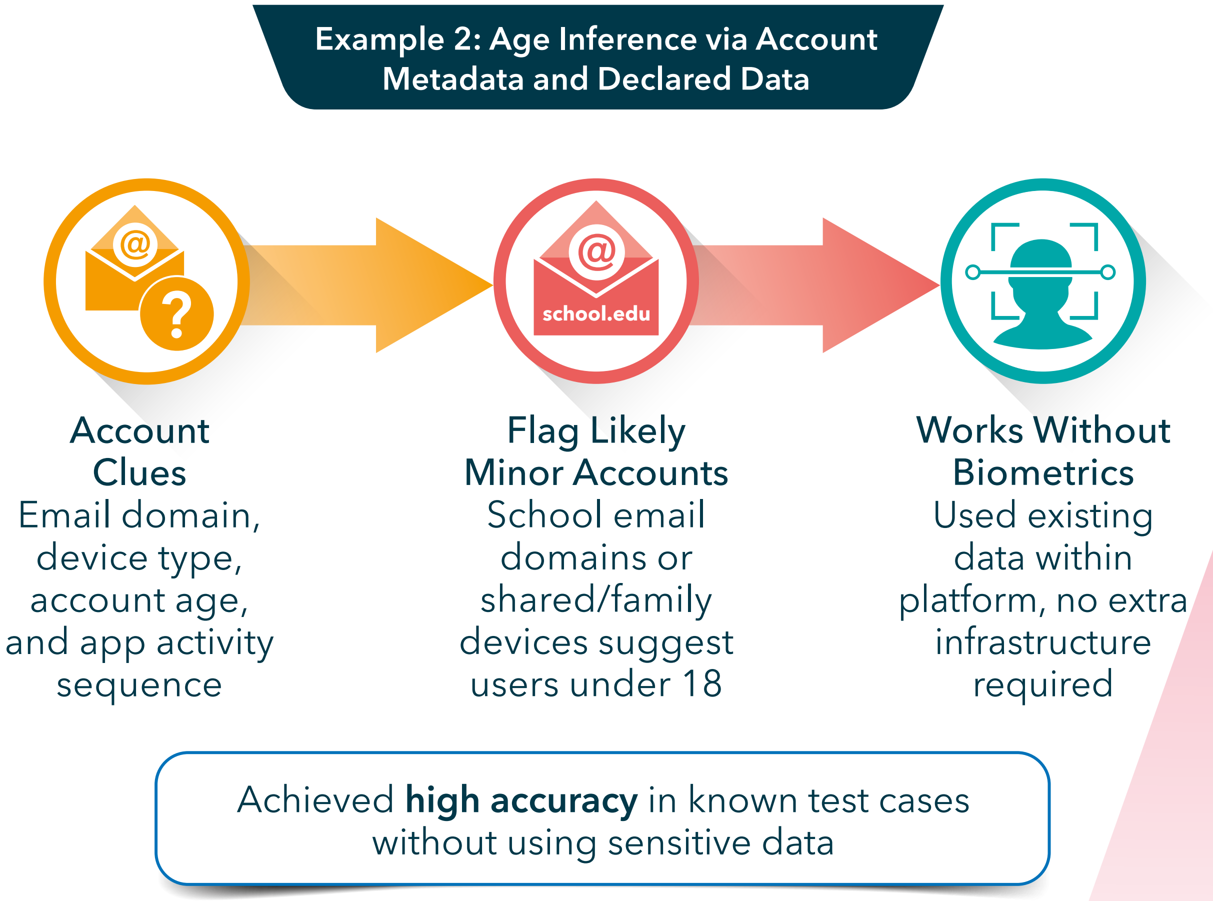

Example 2: Age Inference via Account Metadata and Declared Data

Account Clues
Email domain, device type, account age, and app activity sequence

Flag Likely Minor Accounts
School email domains or shared/family devices suggest users under 18

Works Without Biometrics
Used existing data within platform, no extra infrastructure required

Achieved **high accuracy** in known test cases without using sensitive data

***Figure E.8.2** Age Inference via Account Metadata and Declared Data*

E.8.6 **Example:** Transactional history and service usage

EQUIFAX

- Equifax used historical transaction data, purchase behaviour and regulatory service use (e.g. alcohol delivery, age-restricted subscriptions) to infer whether users were above age thresholds such as 16 or 18.
- For example, repeated use of services legally limited to adults supported a “Likely Over 18” inference when combined with other contextual signals.
- The system applied simple, auditable logic using pseudonymised identifiers and was integrated with retail partners through secure API interfaces.

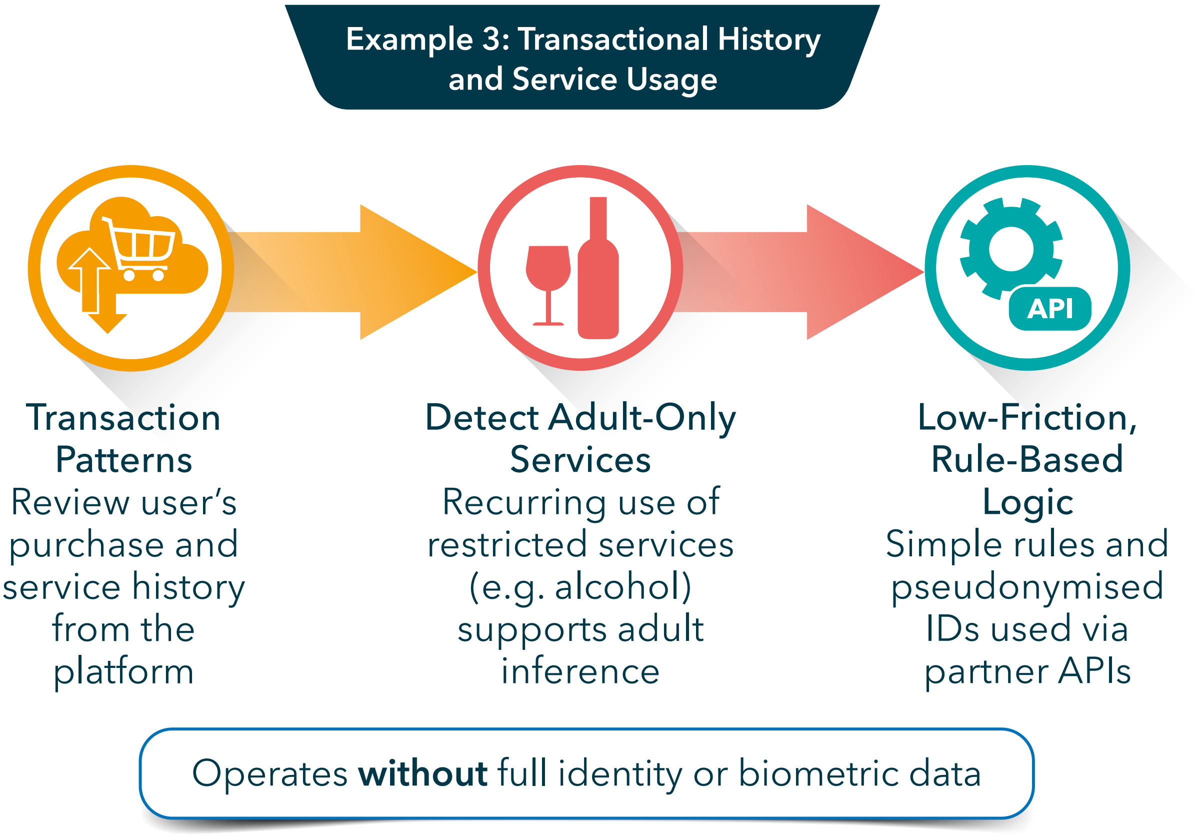

Example 3: Transactional History and Service Usage

Transaction Patterns
Review user’s purchase and service history from the platform

Detect Adult-Only Services
Recurring use of restricted services (e.g. alcohol) supports adult inference

Low-Friction, Rule-Based Logic
Simple rules and pseudonymised IDs used via partner APIs

Operates **without** full identity or biometric data

__Figure E.8.3__ Transactional History and Service Usage

| Why implementation was technologically straightforward

E.8.7 Consistency across age inference providers:

- **Data Availability:** In Australia, users typically engage with digital services using email addresses, mobile numbers, accounts and payment methods – providing ready access to inference signals.
- **System Compatibility**: Age inference systems operated effectively across web, mobile and in-store platforms, without the need for specialised hardware or user-side installation.
- **Deployment Flexibility:** Providers offered SDKs[3], APIs[4] or fully hosted inference platforms, enabling fast deployment for relying parties.
- **Standards Alignment:** Most systems demonstrated alignment with relevant standards such as ISO/IEC FDIS 27566-1, particularly in terms of transparency, explainability and data minimisation.

| Considerations and observations

E.8.8 While there were no significant technical constraints, some Trial participants noted implementation considerations that require attention, including:

- The need to document inference logic to ensure transparency
- Ensuring that signals used are culturally and demographically inclusive
- Avoiding over-reliance on inference alone for high-risk or high-assurance use cases

[3] *SDKs are Software Development Kits. This and other abbreviations used throughout the reports can be found in Part K's Glossary section.*

[4] *API means Application Programming Interface. This and other abbreviations used throughout the reports can be found in Part K's Glossary section.*

E.8.9 The Trial found that age inference technologies are technically mature, adaptable and well-suited to the Australian infrastructure and digital landscape. Trial participants demonstrated that systems could be integrated quickly, ran efficiently and made effective use of existing data and platform signals. As such, technology is not a barrier to adoption – successful deployment hinges on good design, governance and ethical implementation.

E.9 Privacy-Conscious Implementation of Age Inference Systems

E.9.1 Providers and relying parties have approached age inference with careful, responsible deployment, emphasising privacy, data protection and proportionality. The effectiveness of these systems depends on policy and algorithmic decisions that ensure inferences are reasonable and context specific. When well-calibrated, age inference can support age-based requirements without revealing full identity or sensitive information.

E.9.2 The Trial found that participating age inference providers consistently prioritised privacy, data protection and proportionality in the design and deployment of their systems. Unlike verification systems, which may rely on identity documents or estimation systems, which use biometric data, age inference techniques can be implemented in a way that avoids the collection or processing of directly identifying information – making them particularly well-suited to privacy-preserving deployments.

E.9.3 Critically, this aligns with the privacy requirements of ISO/IEC FDIS 27566, which provides clear expectations for data protection in age assurance systems, including:

ISO/IEC FDIS 27566-1	Criteria
Privacy and Data Protection (Clause 8.3)	Requires that personal data be minimised, anonymised where possible and only processed to the extent necessary for the age assurance purpose.
Use of Identity and Personal Data (Clause 8.2)	Emphasises that systems should avoid the use of identity data unless it is essential and supports the implementation of systems that do not reveal the full identity of the user.
Expression of Confidence and Uncertainty (Clause 7.2)	Supports proportional responses, depending on the strength of the inference and discourages overconfident or unjustified conclusions.
Avoidance of Digital Footprint Expansion (Clause 8.4)	Specifically warns against systems that add to or accumulate new persistent data about individuals, particularly when performing repeated or passive inferences.

| Responsible approaches observed in the Trial

E.9.4 Participating age inference providers demonstrated responsible privacy practices, including:

- Using in-platform signals only, such as browsing behaviour or user activity logs, without collecting new identifiers or personal information
- Designing inference systems that immediately discard raw signals after the inference is made (e.g. transient session data)
- Providing binary or categorical outputs (e.g. "likely under 13", "possibly over 18") without retaining full decision trails or probability scores
- Avoiding persistent user profiles unless specifically justified, reducing the risk of profiling or behavioural surveillance
- Implementing confidence thresholds and fallback mechanisms, ensuring that low-confidence inferences were not used for hard access decisions

E.9.5 One provider, for instance, deployed a system that inferred age likelihood from session behaviour and content preferences, but stored only a temporary tag for moderation purposes. No persistent identifier was kept and the user was not profiled across services – demonstrating the kind of ephemeral, privacy-centric design promoted by ISO/IEC FDIS 27566-1.

| Policy and algorithmic decision-making

E.9.6 The effectiveness and ethical acceptability of age inference depends not just on the data used, but on the policies and design choices that shape how the inference is made. During the Trial, systems performed best when:

- The input signals had strong and justifiable links to age (e.g. enrolment status, voting eligibility, regulated product use)
- Algorithms were calibrated with clear logic and explainable thresholds
- Inferences were contextual – e.g. flagging likely underage users for further checks, rather than outright blocking access

E.9.7 This reflects a core principle of ISO/IEC FDIS 27566-1: the need for proportionality and transparency in all age assurance activities, especially when using inferential techniques that may be invisible to users.

E.9.8 Age inference can be a highly privacy-preserving method of meeting age assurance requirements, particularly when built on minimal, ephemeral and context-specific data. The Trial confirmed that providers are increasingly aware of the risks of overreach and profiling and have designed their systems to comply with the data minimisation, purpose limitation and footprint avoidance principles enshrined in ISO/IEC FDIS 27566-1.

E.9.9 By calibrating inference logic carefully and implementing clear governance frameworks, providers can deliver effective age assurance without identity exposure, enabling a more inclusive and privacy-aligned digital environment for all users.

E.10 Transparency and Standards Alignment in Age Inference Practice Statements

E.10.1 Our analysis of practice statements provided by age inference technology providers whose systems had a TRL of 7 or above accurately and fairly reflected the technological capabilities of their products, processes and services (to the extent applicable to the Trial's evaluation criteria). These statements explained how age inference was conducted and how internal policy decisions of the age inference technology providers informed the interpretation and use of inferred age outputs.

E.10.2 One of the key findings of the Trial was that participating age inference providers were able to clearly and consistently describe their systems in Practice Statements developed in accordance with ISO/IEC FDIS 27566-1. These statements offered structured, transparent insight into how each provider's age inference logic was constructed, the nature of the signals used and the safeguards applied to manage accuracy, bias, privacy and fairness.

E.10.3 ISO/IEC FDIS 27566-1 encourages the use of practice statements (similar to certification declarations or conformance documents) as a means of communicating the design intent, functionality and limitations of an age assurance system. These documents are not only useful for evaluators and regulators but also support downstream relying parties in understanding how inference methods operate in practice.

Examples of high-quality practice statements

E.10.4 Several providers submitted detailed documentation that included:

- Logic diagrams showing how signals were weighted and processed
- Pseudocode or flowcharts illustrating how inference results were reached
- Tables of accepted and rejected signals, e.g., stating that social media followers were excluded due to low reliability
- Threshold models, indicating where binary classifications (e.g. "likely under 13") were issued vs where fallback to estimation or verification occurred

E.10.5 These disclosures reflected a mature and self-aware approach, indicating that providers were not only respecting of ISO/IEC FDIS 27566-1 expectations, but had integrated them into internal design governance.

| Implications for trust and accountability

E.10.6 TThe practice statements played a critical role in enabling the Trial evaluation team to:

- Assess the reasonableness of each inference method
- Understand the risk model and confidence logic underpinning system outputs
- Compare data minimisation and user privacy strategies across providers
- Determine the technology readiness level (TRL) and maturity of each solution

E.10.7 Moreover, by using the ISO/IEC FDIS 27566-1 structure, providers supported interoperability and comparability – a key benefit for relying parties seeking to integrate age assurance technologies responsibly.

E.10.8 The age inference systems evaluated during the Trial were accompanied by clear, standards-aligned practice statements, offering valuable insight into system operation, limitations and privacy features. These practice statements were consistent with ISO/IEC FDIS 27566-1 Clause 11 and demonstrated a high level of transparency, maturity and accountability. This level of documentation will be critical in supporting future certification, procurement and regulatory review of age inference systems as they become more widely adopted.

E.11 Configuration Management and Threshold Sensitivity in Age Inference Deployment

E.11.1 We identified opportunities to improve clarity around configuration management during implementation by relying parties – particularly regarding buffer thresholds. For example, relying parties may choose to disregard inferred age results close to critical age boundaries (e.g. 18–21) and instead apply alternative age assurance methods.

E.11.2 Some age inference methods may be more directly applicable to a given age threshold. So, the possession of a credit card and the ability to perform a zero-dollar authenticated transaction with one-time passcode to the credit card holder's device may give moderate confidence that the holder is over 16 (the minimum age for holding a credit card (as an additional user) in Australia). However, other data inferences (such as married status, military services, school attendance, etc) may provide less clear age boundaries at older child/young adult age categories.

E.11.3 There are a wide range of age inference technologies that are technologically advanced and have been used for many years (such as credit card validation, electoral registration and school attendance). There are also newer technologies (such as the algorithmic evaluation of identity anchors like email address) that have demonstrated to the Trial their effectiveness.

E.11.4 Most age inference approaches only delineate adults from children, but as an adult creates a larger digital footprint, it is possible in the future that activity history over a sustained period could start to see age inference approaches for older adults. Conversely, the use of age inference to map out age groups for children and young people is challenging. At that age, they are unlikely to have developed a sufficient and accessible digital footprint to make age inference effective (save perhaps for inference drawn from school of attendance).

E.11.5 During the Trial, the evaluation team observed that the effectiveness and appropriateness of age inference systems were often shaped as much by how relying parties implemented and configured them as by the core technology itself. One important area where this surfaced was in the management of buffer thresholds – zones near critical age boundaries where inferred age results may be considered too uncertain for automated decision-making.

Buffer thresholds and confidence calibration

E.11.6 Many relying parties and technology providers used a concept of buffering to account for the probabilistic nature of inference. For example, even when an inference engine indicated a high probability that a user was over 18, if the output fell within a buffer zone (e.g. between 17.5 and 18.5), the relying party would often choose to:

- Discard the result,
- Trigger a fallback mechanism, such as age estimation or verification or
- Prompt the user for additional signals or self-assertion.

E.11.7 This conservative approach reflects good practice in accordance with ISO/IEC FDIS 27566-1 Clause 9.6, which states:

> *"Where confidence is insufficient, age assurance should default to a more accurate method or provide the user with a means of escalation."*

| Worked examples from the Trial

E.11.8 **Example:** Email Domain Inference

verifymy

- Verifymy used email domain analysis (e.g. .edu.au versus public/free domains), registration metadata and login patterns to infer likely age bands.
- This was particularly effective for identifying school-aged users, but for individuals in the 16–19 age range, results often clustered near the 18+ threshold.
- To address uncertainty, relying parties typically discarded borderline results and triggered facial age estimation as a fallback.

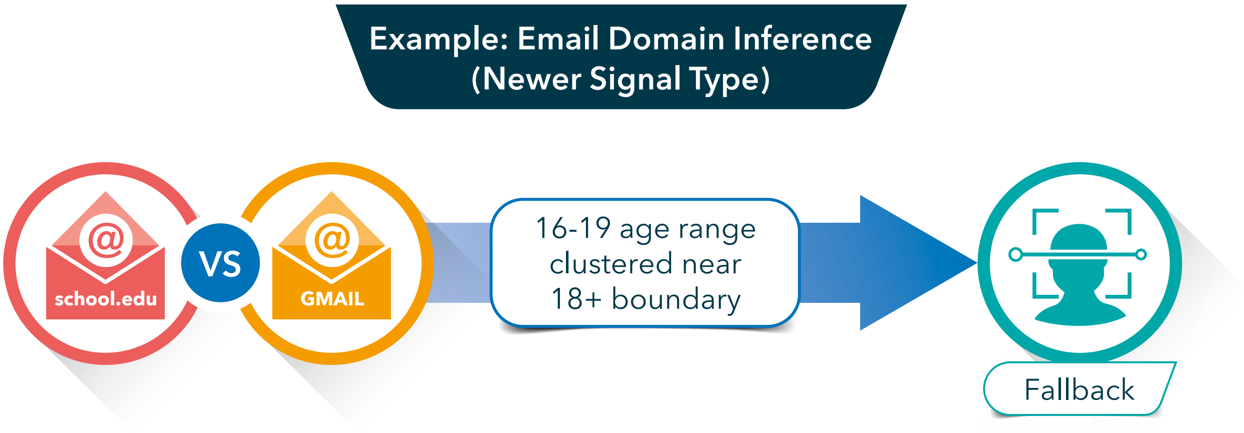
Example: Email Domain Inference
(Newer Signal Type)
school.edu
VS
GMAIL
16-19 age range clustered near 18+ boundary
Fallback

Figure E.11.1 *Email Domain Inference (Newer Signal Type)*

E.11.9 **Example:** Credit Card Ownership as a Proxy for Adulthood

agechecked

- AgeChecked integrated with payment processors to test users' ability to initiate a zero-dollar authenticated credit card transaction, with verification via one-time passcode.
- Since credit cards are legally restricted to those 16+ in Australia (as additional cardholders), a successful result offered a medium-confidence, binary signal of being over 16 at least.
- Relying parties did not apply a buffer threshold in these cases, treating the signal as definitive.

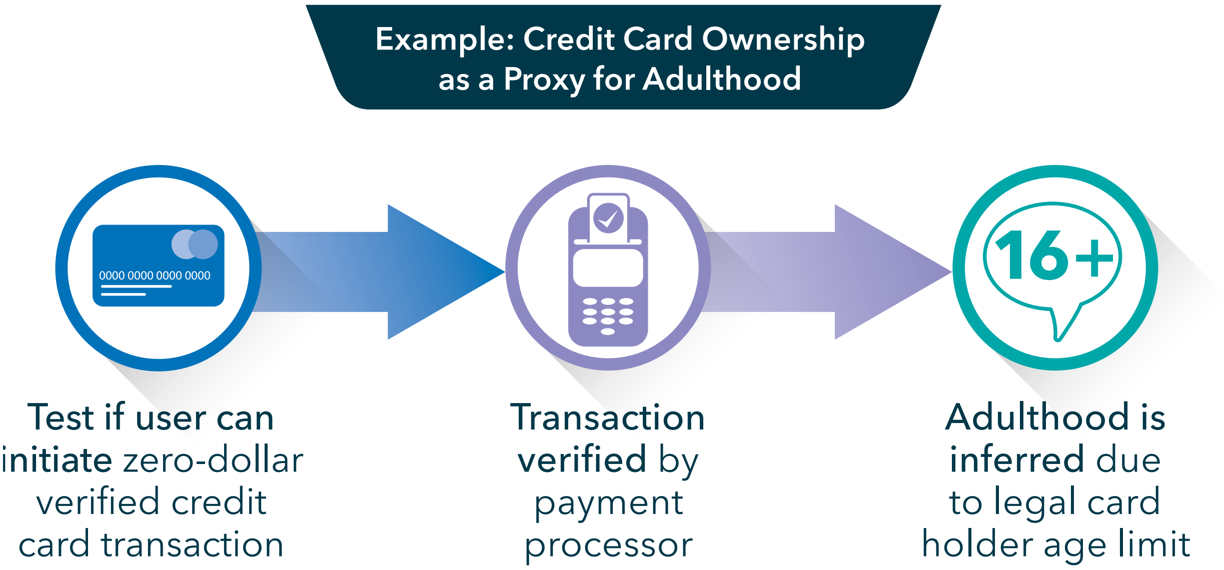

Example: Credit Card Ownership as a Proxy for Adulthood

Test if user can initiate zero-dollar verified credit card transaction

Transaction verified by payment processor

Adulthood is inferred due to legal card holder age limit

Figure E.11.2 *Credit Card Ownership as a Proxy for Adulthood*

E.11.10 **Example:** Electoral Roll Matching

EQUIFAX

- Equifax also enabled matching against electoral roll records, linking names and addresses to verified enrolment.
- As enrolment is mandatory from age 18, a successful match provided strong justification to infer that the user was an adult.
- Some relying parties introduced a ±6-month buffer to account for possible delays in enrolment, especially near birthdays.

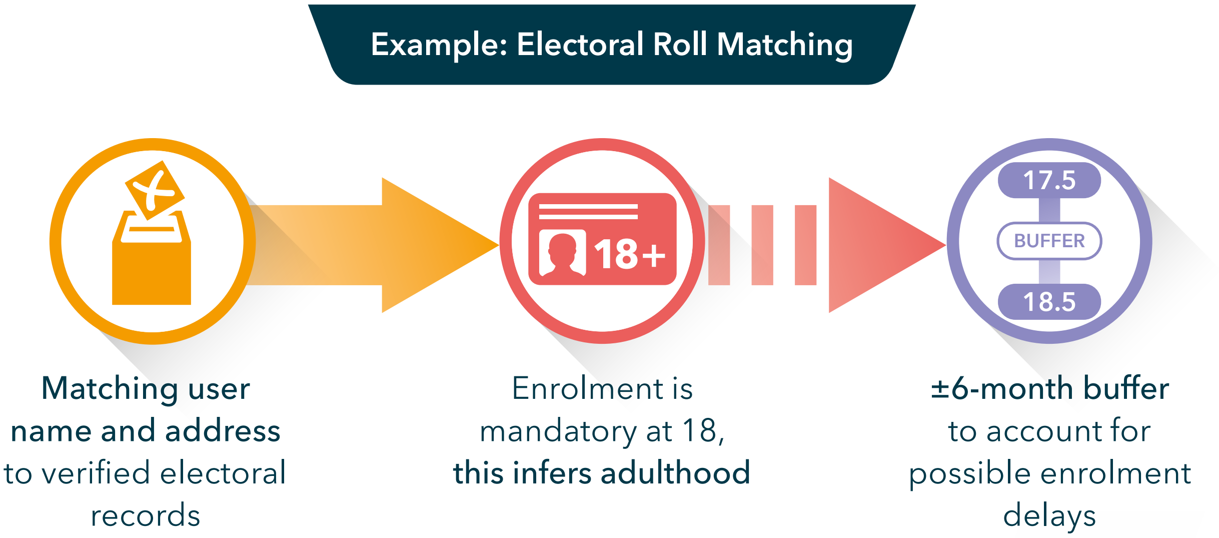
Example: Electoral Roll Matching
18+
17.5
BUFFER
18.5
Matching user name and address to verified electoral records
Enrolment is mandatory at 18, **this infers adulthood**
±6-month buffer to account for possible enrolment delays

Figure E.11.3 *Electoral Roll Matching*

E.11.11 Example: School Attendance Inference

MyMahi

- MyMahi inferred age from school enrolment data, identifying users in Year 9 or Year 10, typically aged 14–16.
- This was effective for flagging likely underage users, but near the 16+ boundary, the data lacked granularity.
- Relying parties treated these cases with caution, especially when access to sensitive or regulated content was involved.

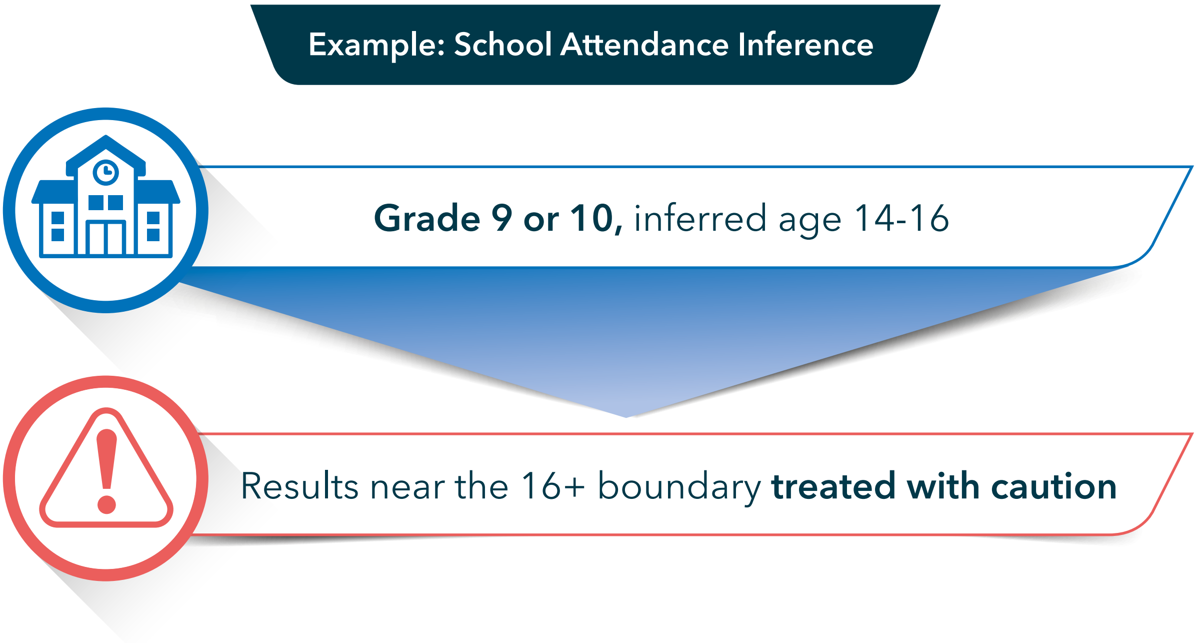

Example: School Attendance Inference

Grade 9 or 10, inferred age 14-16

Results near the 16+ boundary **treated with caution**

Figure E.11.4 School Attendance Inference

Observations on future maturity and usage

E.11.12 While most age inference implementations during the Trial were focused on adult vs child distinctions, the Trial observed that as adults generate deeper digital footprints, it may become feasible to apply inference for older age banding (e.g. 30-45, 60+). Examples might include:

- Retirement planning activity
- Long-term health service engagement
- Superannuation or pension access history

E.11.13 Conversely, for younger users, inference remains more difficult. Adolescents often have limited public records, payment credentials or distinct online habits, making fine-grained inference across age bands (e.g. distinguishing 13 from 16) inherently less reliable. In these cases, inference may still play a valuable role in triggering layered assurance steps rather than forming the basis for a final decision.

E.11.14 The Trial found that age inference systems can be powerfully effective when aligned to the right thresholds and signals – but their deployment depends heavily on how they are configured and interpreted by relying parties. Well-defined buffer zones, fallback logic and clear signal-strength assessments are essential to ensuring that age inferences are applied ethically, accurately and proportionately.

E.11.15 Future guidance – potentially embedded within ISO/IEC FDIS 27566-1 practice statements – could help standardise the configuration of thresholds, fallback decisions and confidence interpretation, supporting more consistent and trustworthy implementations across diverse platforms.

E.12 Innovation and Research in Age Inference: Insights From the Australian Context

E.12.1 We found a vibrant, creative and innovative age inference service sector, with mature and actively deployed solutions. Providers demonstrated ongoing commitment to research and development, refining methods to infer age from behavioural, contextual and transactional data. Continuous improvements, independent testing and user-focused design have enhanced both accuracy and accessibility across age-restricted services.

E.12.2 The Trial highlighted a dynamic and innovative age inference sector in Australia. Providers showcased mature solutions actively deployed across various platforms, demonstrating a commitment to continuous improvement through research and development. These efforts focus on refining methods to infer age from behavioural, contextual and transactional data, enhancing both accuracy and accessibility in age-restricted services.

Academic research informing age inference

E.12.3 Several Australian academic studies provide insights into the feasibility and ethical considerations of inferring age from various data sources:

- **Longitudinal Study of Australian Children (LSAC[5]):** This study tracks the development of children over time, collecting data on behaviours, environments and outcomes. While not directly used for commercial age inference, LSAC provides a rich dataset that can inform models predicting age-related behaviours and transitions. *Australian Institute of Family Studies+2Australian Institute of Family Studies+2Oxford Academic+2*
- **Understanding the Digital Behaviours of Older Australians:** This research by the eSafety Commissioner examines how older Australians engage with digital technologies, highlighting patterns that could inform age inference models, particularly in distinguishing older age groups based on digital literacy and online behaviours. *eSafety Commissioner*
- **Unexplored Territory: Information Behaviour in the Fourth Age:** This study explores the information needs and behaviours of individuals in advanced age, providing insights into how age-related changes affect information-seeking behaviours, which could be relevant for refining age inference techniques for older populations. *Charles Sturt University Research Output+1PMC+1*

[5] *All references to the following academic research can be found in Part E's Bibliography section within the Part K Report.*

Examples from the Trial Demonstrating Innovation

E.12.4 Participants in the Trial implemented various innovative approaches:

- **Behavioural Analysis:** Some providers analysed user interaction patterns, such as typing speed, navigation habits and content preferences, to infer age groups. These methods align with findings from academic studies on digital behaviours across age groups.
- **Transactional Data Utilisation:** Providers leveraged transactional data, like purchase histories and payment methods, to infer age, considering legal age requirements for certain transactions. This approach reflects the practical application of understanding age-related behaviours in commerce.
- **Contextual Signals:** Use of contextual information, such as time of activity and device usage patterns, helped in refining age inference models, demonstrating the sector's commitment to incorporating diverse data points for more accurate age estimation.

E.12.5 The age inference sector in Australia is characterised by its innovative spirit and commitment to integrating research findings into practical applications. While certain academic studies provide foundational knowledge, ongoing collaboration between researchers and industry practitioners is essential to develop robust, ethical and effective age inference methods. This synergy ensures that age-restricted services can be both accessible and aligned to privacy and data protection standards.

E.13 Privacy in Age Inference: Protection, Proportionality and Digital Footprint Management

E.13.1 We found robust understanding of and internal policy decisions regarding the handling of personal information in age inference – particularly where inference is based on signals drawn from an individual's digital footprint, such as behavioural patterns, service usage or contextual data. Providers demonstrated clear separation between operational use, training and validation data, with strong privacy safeguards in place.

E.13.2 Securely processed and not retained in a form that could identify individuals. Most providers retained only non-identifying transaction codes. We also observed early-stage development of methods that rely solely on indirect digital indicators – such as browsing habits or interaction histories – to estimate age without revealing identity.

E.13.3 While privacy protections vary by deployment context, these developments reflect a strong commitment among age inference providers to respecting user expectations and enabling privacy-preserving age assurance.

E.13.4 Across the Trial, age inference providers demonstrated a high level of privacy awareness, integrating strong safeguards into both the design and operational deployment of their systems. While age inference draws on indirect data signals – such as behavioural patterns, service usage and contextual metadata – rather than identity documents or biometrics, these systems still intersect with personal data. As such, adherence to privacy-by-design principles and careful data governance is essential.

E.13.5 Providers consistently showed deep understanding and deliberate policy choices to minimise risk, particularly when working with data drawn from users' digital footprints. This aligns directly with the privacy and data protection provisions in ISO/IEC FDIS 27566-1, which outlines specific requirements for age assurance technologies to protect users from unnecessary or disproportionate intrusion.

Emerging Zero-Identity Inference Techniques

E.13.6 Some providers demonstrated early-stage research into techniques relying solely on:

- Browsing habits,
- App interaction frequency,
- Interface navigation speed,
- Or content engagement style.

E.13.7 These methods promise extremely privacy-preserving inference models that can detect likely age bands without any identifiable input from the user. Such developments represent a meaningful advance towards zero-knowledge age assurance, particularly for low-risk environments.

Balancing risk, utility and expectations

E.13.8 Privacy risks in age inference are context dependent. A low-risk platform (e.g. a gaming app with optional chat) may apply passive inference techniques with minimal disclosure. By contrast, a platform offering access to sensitive or regulated services (e.g. online alcohol sales) may need tighter controls, clearer user notice and fallback mechanisms.

E.13.9 Providers were found to be attuned to user expectations – designing systems that respected the principle of proportionality, avoided tracking across services and did not build persistent user profiles unless explicitly required by the relying party.

E.13.10 Age inference systems trialled through the Trial demonstrated strong alignment with ISO/IEC FDIS 27566-1 privacy requirements, with consistent application of:

- Data minimisation
- Non-identifiability
- Separation of functional and training uses
- Digital footprint reduction

E.13.11 This reflects a maturing field in which privacy and performance are not mutually exclusive and where responsible design choices allow for age assurance to be both effective and ethical. As age inference technology continues to evolve, this foundation provides a solid basis for future trust, certification and regulatory acceptance.

Vendor Case Study

frankieone

Website

frankieone.com

FrankieOne aggregates age assurance services, enabling inference through third-party integrations; supports orchestration logic for fallback, confidence thresholds and multi-vendor identity decisioning.

Practice Statement

ageassurance.com.au/v/fra/#PS

Technology Trial Test Report

ageassurance.com.au/v/fra/#TR

Privacy Policy

ageassurance.com.au/v/fra/#PP

Technology Trial Interview

ageassurance.com.au/v/fra/#VI

Summary of Results

Aggregator platform integrating third-party inference and IDV solutions. Supports inference via partners and orchestration logic but does not offer native age inference capability. Strong on configuration and delivery; TRL reflects reliance on external providers.

E.14 Inclusion and Cultural Grounding in Age Inference

E.14.1 Age inference provides an important opportunity for individuals who do not have formal ID documents or a strong digital history – to demonstrate age-related eligibility for access to goods, services, content, venues or spaces. Unlike age estimation, which relies on analysing biological features, age inference can draw on a broader range of contextual signals.

E.14.2 For First Nations People, limited representation in conventional datasets can present challenges. However, we observed efforts to explore culturally appropriate inference methods. In particular, the knowledge younger community members possess about their land, animals, plants and tribal features can provide strong, culturally grounded indicators of age or life stage. These insights offer promising pathways to more inclusive, respectful and accurate age inference systems, tailored to the lived realities and traditions of Indigenous communities.

E.14.3 Age inference plays an important role in improving accessibility and inclusion within age assurance frameworks – particularly for individuals who may lack formal identification documents or who are underrepresented in digital or biometric datasets. Unlike age verification (which requires authoritative documents) or age estimation (which often relies on facial imagery), age inference can draw on contextual, behavioural and culturally grounded signals, allowing for more adaptable and community-sensitive deployments.

E.14.4 This flexibility is central to the acceptability and inclusion principles defined in ISO/IEC FDIS 27566-1, which states in Clause 5.5:

> *"Age assurance systems should be designed and implemented in a way that is accessible, inclusive and culturally appropriate, particularly for underserved or marginalised populations."*

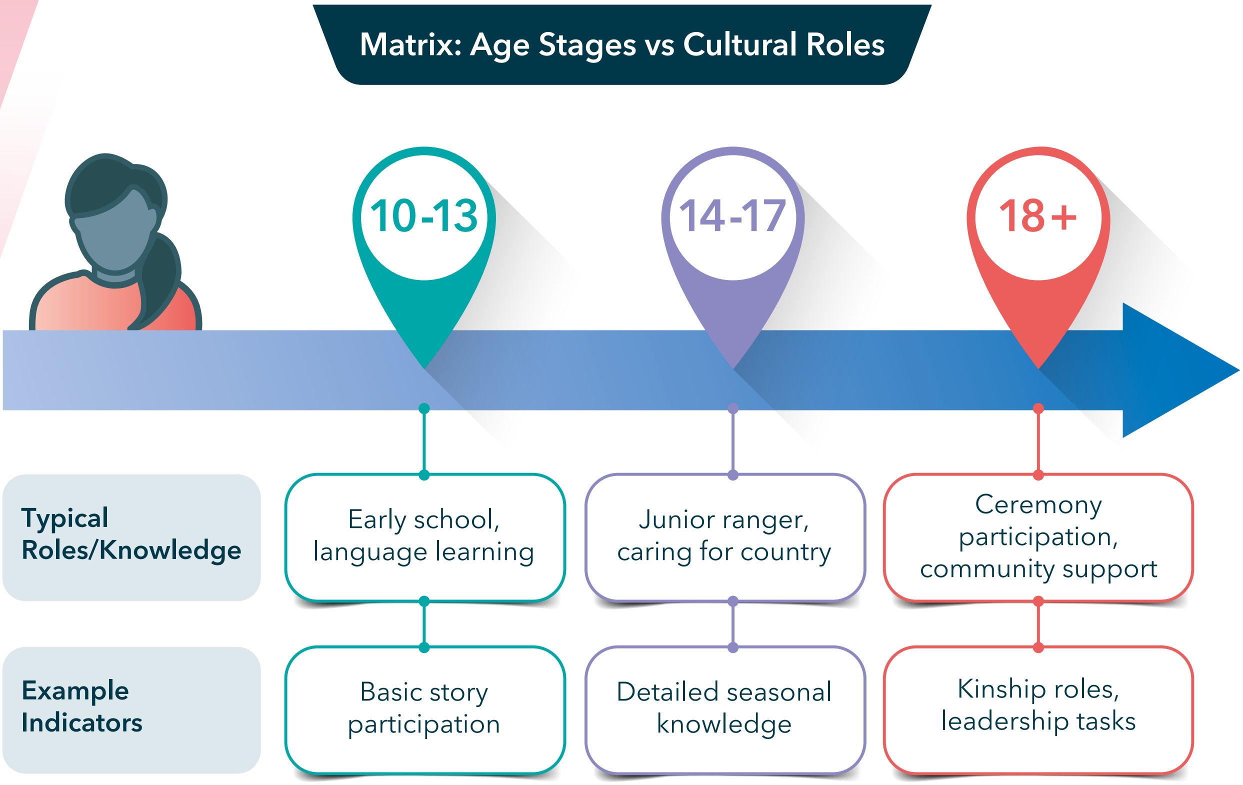

Matrix: Age Stages vs Cultural Roles

	10-13	14-17	18+
Typical Roles/Knowledge	Early school, language learning	Junior ranger, caring for country	Ceremony participation, community support
Example Indicators	Basic story participation	Detailed seasonal knowledge	Kinship roles, leadership tasks

Figure E.14.1 *Matrix: Age Stages vs Cultural Roles*

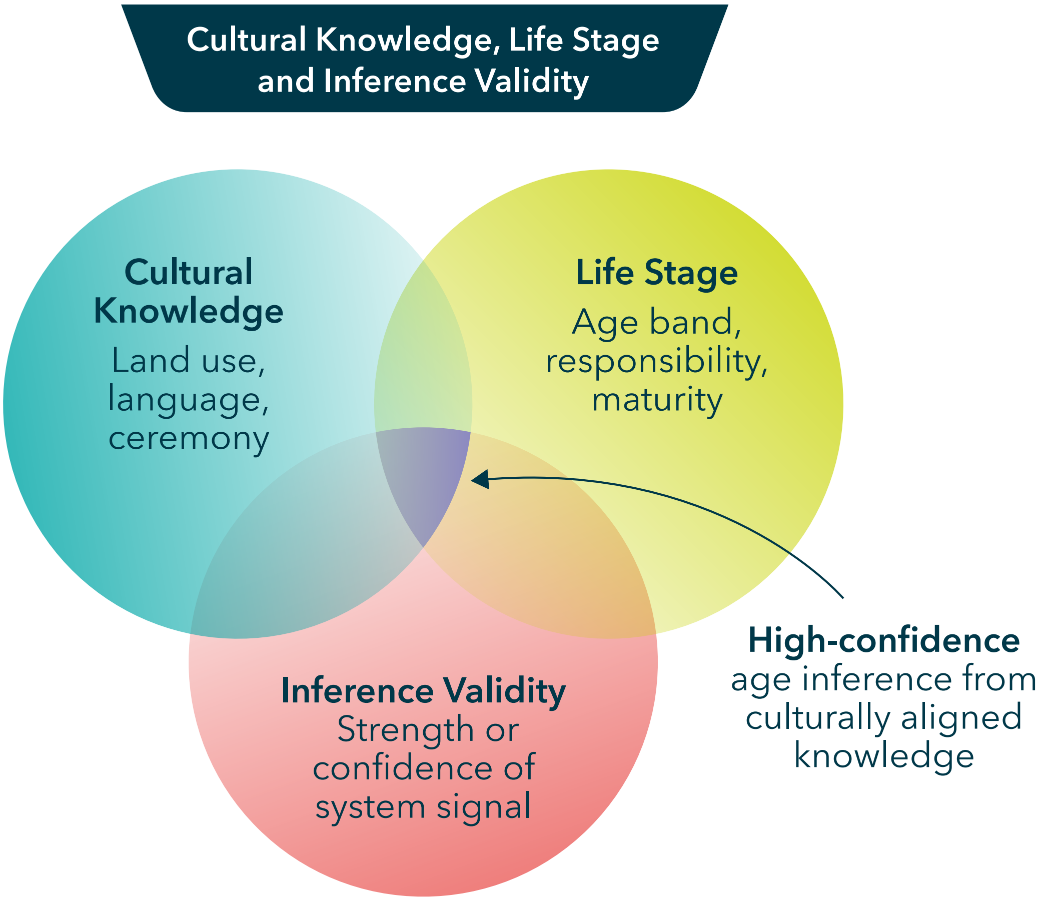

Cultural Knowledge, Life Stage and Inference Validity

Cultural Knowledge
Land use, language, ceremony
Life Stage
Age band, responsibility, maturity
Inference Validity
Strength or confidence of system signal
High-confidence age inference from culturally aligned knowledge

Figure E.14.2 *Overlap of Cultural Knowledge, Life Stage and Inference Validity*

Cultural inclusion in the Australian context

E.14.5 In the context of the Trial, several participants acknowledged the challenges of applying conventional inference methods to Indigenous communities, including Aboriginal and Torres Strait Islander Peoples. These populations are frequently underrepresented in structured datasets, face lower levels of digital footprint visibility and may use shared or communal devices, making standard inference models less reliable.

E.14.6 However, the Trial also surfaced early-stage exploration of culturally respectful inference models, including approaches that draw on culturally embedded life-stage markers and local knowledge systems. These approaches would require community ownership, consent and oversight to ensure alignment with Indigenous values and self-determination.

Worked examples: culturally grounded indicators of age

E.14.7 Knowledge-based life stage markers

- In many Indigenous communities, certain cultural responsibilities or knowledge domains are learned or passed down at specific ages or rites of passage.
- For example, deep knowledge of seasonal land use, plant life or hunting practices may correlate with teenage or early adult roles, while more advanced knowledge of kinship structures or ceremonial practice may indicate older life stages.

E.14.8 Language and story-telling participation

- Participation in particular forms of storytelling (e.g. Dreaming stories appropriate to certain age groups) or the ability to recount specific Country-based knowledge may reflect age-linked educational exposure or community positioning.
- Some inference systems could be adapted to analyse language complexity, story themes or questionnaire responses in culturally sensitive ways.

E.14.9 School and community role correlation

- Rather than relying solely on enrolment records, some providers are exploring the use of roles within school or community events (e.g. junior ranger programs, youth leadership roles) as proxy indicators of age range or maturity level.

E.14.10 Community-attested roles

- With consent, community attestation or role confirmation (e.g. confirmation from a community leader that a person is of "senior schooling age") could be used as a fallback or complementary signal in layered inference systems.

E.14.11 These approaches are still largely in research or conceptual stages, but they reflect a growing awareness of the need for culturally grounded data practices, as promoted by organisations such as the Lowitja Institute, which advocates for health and data models grounded in Indigenous values and realities.

lowitja.org.au

Systemising culturally grounded age inference in practice

E.14.12 Culturally grounded age inference offers the potential for respectful, community-aligned methods of establishing age eligibility without relying on biometric data or formal ID. For First Nations and Torres Strait Islander Peoples – who may not possess conventional identity documents or digital histories – such systems can provide accessible, non-intrusive alternatives aligned with their knowledge systems and community structures.

E.14.13 To be effective and resistant to spoofing, these inference methods must:

- Be based on genuinely learned knowledge that reflects life stage or cultural progression,
- Be community contextual, not generic or easily searchable online, and
- Involve validation processes rooted in local authority, engagement or dynamic questioning.

Trial Cultural Advisor, John Fejo and his family.

| Practical examples of culturally grounded age inference

E.14.14 Here are three contextualised examples that could inform pilot programs:

E.14.15 *Dynamic Knowledge Test Based on Country-Specific Seasons*

Use Case: Establishing whether an individual is likely over age 14 for access to teen social platforms in a remote community.

Mechanism:

- Present the user with a dynamic, locally specific question such as:

  "In your community, what animals or plants appear in the dry season and how do people prepare for them?"

- Responses are free-text or voice-based, analysed for:
  - Use of local terminology (e.g., names of native species in language),
  - Sequencing or reasoning that matches expected lived knowledge, and
  - Cross-validation with local environmental data (e.g., fire season, blooming cycles).

Spoof resistance:

- Questions rotate and are not standardised.
- Google-style searches are unlikely to yield answers with cultural depth or specificity.
- Can be cross-verified by local elders or used in layered inference (not as sole input).

E.14.16 *Ceremonial or Role-Based Knowledge Markers*

Use Case: Determining if a user is likely to be 16+ to access training services.

Mechanism:

- The user is prompted to indicate:
  - Their recent involvement in ceremonial roles, youth camps or community responsibilities (e.g. caring for country, junior ranger schemes).
  - Their understanding of what is expected at those life stages.

Systemisation:

- Community-developed role matrices define which activities typically correlate to which age ranges.
- Responses are evaluated via:
  - Role-based inference models, updated by the community.
  - Integration with school or program participation data.

Spoof resistance:

- Responses are validated against known role patterns or confirmed by a community-nominated verifier (e.g. youth worker or teacher).

E.14.17 *Language and Story Participation*

Use Case: Filtering out users under 13 from chat-based environments.

Mechanism:

- System engages the user in narrative-based prompts involving Dreaming stories or place-based knowledge:

  "Tell us about a story from your area about the hills and the animals that live there."

- Responses are assessed for:
  - Cognitive and narrative complexity associated with older age.
  - Use of place-based references known primarily to locals.
  - Engagement with intergenerational stories unlikely to be memorised without lived experience.

Spoof resistance:

- Real-time story structuring is difficult to fake.
- Evaluated using trained AI models fine-tuned on language complexity and regional relevance.

E.14.18 While no provider implemented culturally grounded age inference in production during the Trial, several providers demonstrated awareness of inclusion challenges, particularly for Indigenous communities and expressed a commitment to community-led exploration of culturally respectful methods. Conceptual models – such as dynamic cultural knowledge prompts or role-based indicators – were discussed and may form the basis for future pilot programs.

Vendor Case Study

MyMahi

Website

mymahi.com

MyMahi provides verifiable credentials based on school-sourced identity and age data, enabling inference of student life-stage through structured education signals and contextual indicators.

Practice Statement

ageassurance.com.au/v/mym/#PS

Technology Trial Test Report

ageassurance.com.au/v/mym/#TR

Privacy Policy

ageassurance.com.au/v/mym/#PP

Technology Trial Interview

ageassurance.com.au/v/mym/#VI

Summary of Results

Conceptually strong use of verified school data for age-bound credential issuance. However, practical deployment of age inference (e.g. age range detection without explicit credentials) is still emerging. TRL 5 reflects a validated design with future implementation potential.

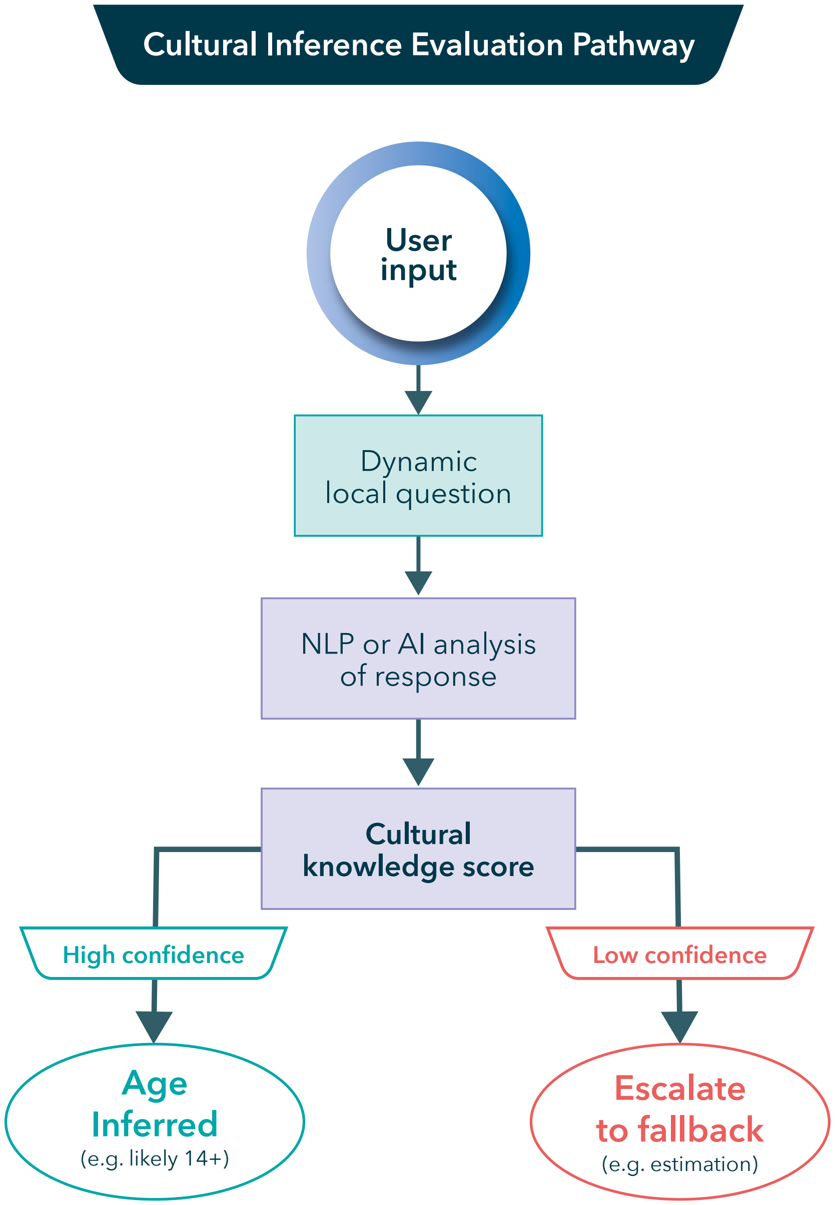

Figure E.14.3 *Cultural Inference Evaluation Pathway*

E.15 Improving Technological Foundations for Reliable Age Inference

E.15.1 We found opportunities for technological improvement of age inference systems, particularly in developing reliable, ground-truthed data sources to support accurate and context-aware inference. This includes behavioural and contextual signals linked to known age indicators, rather than formal identity documents.

E.15.2 While large-scale models improve inference accuracy, collecting high-quality, representative data – especially for children – remains challenging, costly and raises ethical concerns. Many promising inference methods, such as those based on user interaction patterns, interface behaviours or engagement with age-typical content, are still in the academic or early development stages.

E.15.3 The Trial found that while age inference systems are increasingly viable and privacy-conscious, there remain significant opportunities for technological improvement, particularly in relation to the quality and representativeness of the data used to train, calibrate and validate these systems.

E.15.4 Most age inference models rely on behavioural, contextual or interactional signals rather than identity-linked documents. These inputs offer privacy advantages and support inclusion, but their effectiveness ultimately depends on:

- The validity of the signals in relation to age,
- The consistency and quality of the underlying data and
- The robustness of models trained on these data to diverse user populations.

E.15.5 This aligns with the emphasis in ISO/IEC FDIS 27566-1 Clause 6.1.2 (Measurement of Effectiveness), which requires that:

> *"The age assurance provider shall define and document the performance characteristics of the system and identify how these are measured and validated using representative datasets."*

| Challenges in developing ground-truth data

E.15.6 Unlike age verification systems that use date of birth from authoritative sources, age inference relies on indirect signals – which are only useful if they have been tested and statistically linked to known age outcomes.

E.15.7 However, The Trial identified persistent challenges in this area:

- Lack of public, ethically collected datasets of behavioural signals tied to confirmed ages
- Difficulties in collecting ground-truth data for children, especially under age 13, due to heightened ethical and legal protections
- Bias risks where training data is skewed toward adult or high-income populations, reducing accuracy for marginalised or underrepresented groups

Vendor Case Study

EQUIFAX

Website
equifax.co.uk

Equifax used its access to financial signals and account metadata to support inference of age likelihood. Equifax is certified to ISO/IEC 27001 and operates under extensive regulatory compliance obligations.

Three Key Facts

1. Credit product age and transaction metadata. Device and session characteristics. Zero-dollar card verification for age 18+.
2. Deterministic rules (e.g., credit card ownership adult). Inference triggered by transaction context.
3. Inference highly accurate for adult classification, no direct PII used in inference logic. Used known payment infrastructure for confidence boosting.

Strengths

- Trusted signals with legal age constraints
- High confidence for adult gating
- Well-suited to e-commerce and digital transaction contexts

Practice Statement
ageassurance.com.au/v/equ/#PS

Privacy Policy
ageassurance.com.au/v/equ/#PP

Technology Trial Test Report
ageassurance.com.au/v/equ/#TR

Technology Trial Interview
ageassurance.com.au/v/equ/#VI

Summary of Results

Equifax provided a robust, high-assurance inference model for 18+ classification. Its deterministic, rule-based approach based on credit eligibility demonstrated reliability and legal clarity.

| Early-stage and promising inference techniques

E.15.8 Several methods trialled or explored during the Trial showed promise but remain in academic or pre-commercial phases. These include:

Inference Signal Type	Example Indicators	Development Status
User interface interaction patterns	Tap speed, scrolling style, menu navigation	Pilot-tested
Language complexity in chat or input	Syntax, vocabulary, phrase repetition	Used in research
App engagement and feature access	Use of settings, time spent, feature switching	Limited deployment
Content consumption patterns	Types of videos, games or apps engaged with	Privacy-sensitive
Environmental cues	Time-of-day usage, device type, geolocation	Experimental

E.15.9 While these signals may correlate with developmental stages or age-related maturity, validating their reliability requires systematic study using datasets anchored to known age, i.e., ground-truthed data.

Vendor Case Study

LUCIDITI®

Website

luciditi.co.uk

Luciditi offers age inference through digital identity orchestration using document checks, behavioural signals and open data sources, supporting verification and estimation across multiple age thresholds.

Practice Statement

ageassurance.com.au/v/luc/#PS

Technology Trial Test Report

ageassurance.com.au/v/luc/#TR

Privacy Policy

ageassurance.com.au/v/luc/#PP

Technology Trial Interview

ageassurance.com.au/v/luc/#VI

Summary of Results

Luciditi platform supports inference, estimation and verification. Claims of interoperability and real-world integration; however, specific details of age inference deployment were limited to interview insights rather than practice statement-level granularity.

| Ethical and practical limitations

E.15.10 Developing such datasets presents real constraints:

- **Ethical concerns:** Collecting detailed behavioural data from children requires rigorous consent, data minimisation and oversight protocols.
- **Cost and complexity:** Curating representative datasets is resource-intensive and may not be commercially viable for all providers.
- **Data governance:** Ensuring long-term control, auditability and privacy-preserving retention of ground-truth datasets is non-trivial.

E.15.11 ISO/IEC FDIS 27566-1 reinforces that effectiveness should not be assumed – it must be measured, disclosed and validated, particularly when systems are being used to enforce access to regulated services (Clause 5.4.1 and 5.4.2).

E.15.12 While age inference shows clear promise, particularly in low-friction and privacy-first deployments, its long-term success depends on robust, ground-truthed and ethically sourced datasets. These are needed to:

- Improve model performance across age groups
- Minimise bias and demographic variability
- Support explainable and auditable assurance outcomes

E.15.13 The Trial highlighted several early-stage techniques – such as interaction-pattern analysis and contextual signal matching – that may yield reliable results in the future but require further data development and validation to meet the effectiveness measurement expectations of ISO/IEC FDIS 27566-1.

E.16 The Future of Embedded Age Inference in Digital Environments

E.16.1 We see strong future potential for age inference technologies to become more seamlessly integrated into in-app, in-game and in-purchase experiences, using natural digital behaviours to infer age in privacy-preserving and context-sensitive ways.

E.16.2 The Trial identified strong future potential for age inference technologies to be seamlessly embedded into digital environments – particularly in-app, in-game and in-purchase experiences – where user engagement provides a natural opportunity to analyse behavioural and contextual signals to infer age. These environments offer a rich, low-friction context for privacy-preserving, unobtrusive age assurance, aligning with the growing expectation that digital safety mechanisms should integrate natively into user experiences.

| In-app and in-game integration

E.16.3 In real-time digital ecosystems such as mobile apps and games, age inference could:

- Evaluate patterns of interaction (e.g. gesture frequency, session length, feature access)
- Detect language usage or decision-making maturity through gameplay or chat
- Adapt access to content dynamically based on inferred age bands

E.16.4 **For example:** A mobile game designed for general audiences might use embedded inference models to detect if a user is likely under 13 and automatically disable monetised features, in-app chat or links to external content. This could occur without requiring any personal data entry or interruptive age verification screen, preserving both engagement and compliance.

| In-purchase and E-commerce journeys

E.16.5 In online retail or digital marketplaces, age inference may operate within:

- Browsing and cart behaviour (e.g. frequency and type of age-restricted items)
- Interaction with terms and conditions (e.g. skipping fine print, request for help)
- Payment method indicators (e.g. inferred credit card ownership, billing details)

E.16.6 For example: An alcohol delivery app may infer the likelihood of a user being underage based on their device profile, delivery address history and interaction patterns. If the system detects a low-confidence age band, it could escalate to real-time facial estimation or verified ID – ensuring age assurance remains proportionate and frictionless for most users.

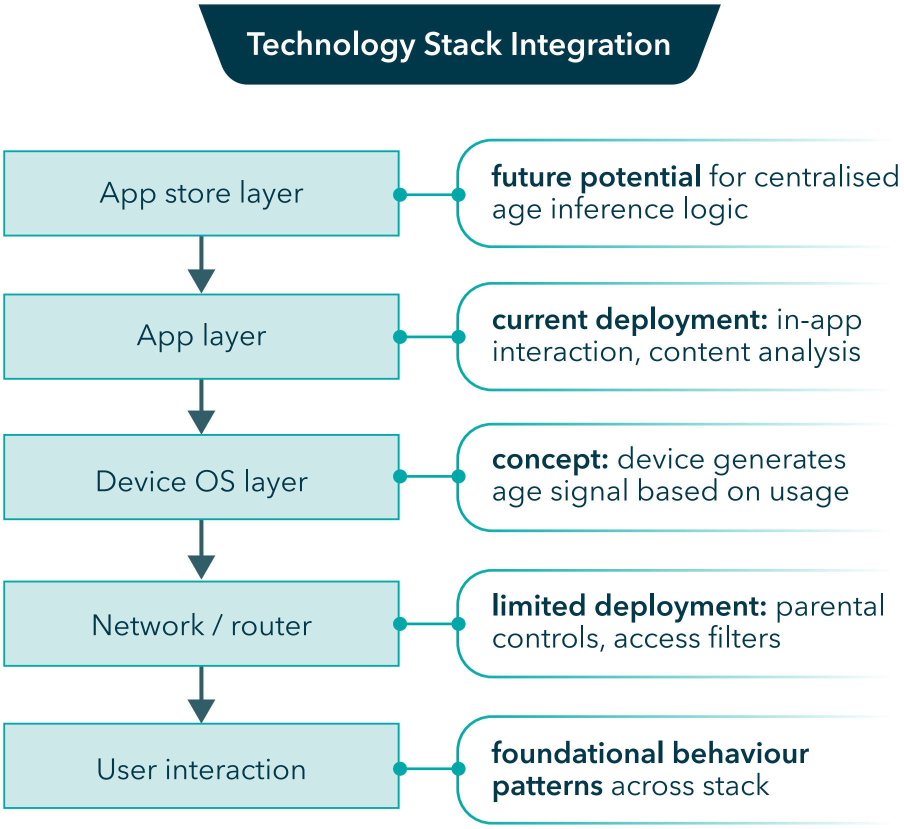
Technology Stack Integration
App store layer
future potential for centralised age inference logic
App layer
current deployment: in-app interaction, content analysis
Device OS layer
concept: device generates age signal based on usage
Network / router
limited deployment: parental controls, access filters
User interaction
foundational behaviour patterns across stack

Figure E.16.1 Technology Stack Integration

| Alignment with the technology stack report

E.16.7 As noted in the Technology Stack Report, future-ready age inference solutions are expected to operate at the application layer, with optional support for:

- Platform-wide integration (e.g. app store level)
- Device-level age signals (e.g. parental controls, device ownership)

E.16.8 While the Trial did not identify current deployments at platform or device level using inference alone, participants expressed strong interest in exploring such capabilities, particularly for edge device processing where inference models could operate locally without sending raw data to external servers.

E.16.9 This reflects the emphasis ISO/IEC FDIS 27566-1 places on:

ISO/IEC FDIS 27566-1	Criteria
Clause 6.5	Minimising data transfer
Clause 6.6	Avoiding digital footprint expansion
Clause 7.2	Maintaining functional transparency in embedded environments

***Cross Reference: Part J -** Tech Stack*

| Cross-reference to age verification and estimation

E.16.10 In contrast to age verification, which often requires document submission or external identity checks and age estimation, which typically uses a biometric capture moment (e.g. facial analysis), age inference allows systems to evaluate user age characteristics passively and in context.

E.16.11 The embedded nature of inference allows for:

- Early-stage screening, especially useful for progressive assurance journeys
- Real-time adaptation of content access or purchasing controls
- Avoidance of intrusive prompts or data collection for most users

E.16.12 This makes age inference particularly valuable for user-centric designs, where UX continuity is as important as regulatory compliance.

E.16.13 Age inference technologies are on track to become key components of embedded age assurance ecosystems, offering lightweight, privacy-sensitive mechanisms that support compliance and safety without compromising user flow. As platforms evolve and digital experiences become more adaptive, inference-based tools will likely become standard components of in-app safety stacks, complementing verification and estimation in layered frameworks.

E.16.14 To realise this potential at scale, ongoing work is needed on:

- Confidence calibration
- Device-side processing
- Industry standards for interoperable inference modules

E.16.15 These developments, aligned with ISO/IEC FDIS 27566-1 and insights from the broader Technology Stack Report, represent the future of seamless, respectful and scalable age assurance.

E.17 Enhancing Age Inference Through Parental Control Signals

E.17.1 Parental controls can enhance age inference systems by providing contextual signals – such as device settings, usage restrictions or guardian-managed profiles – that suggest likely age ranges. While not definitive proof of age, these indicators contribute to a broader picture, supporting more accurate inferences without serving as direct, ground-truth evidence of an individual's age.

E.17.2 While age inference systems generally rely on behavioural and contextual indicators associated with the user's interaction with a platform, these systems can be augmented by data from parental control settings, which provide an additional layer of environmental context. Such controls – when established at the device, platform or network level – can offer strong contextual indicators of likely age ranges, particularly in shared or family-managed digital ecosystems.

| Types of signals from parental controls

E.17.3 When integrated responsibly, the following types of parental control signals can contribute to inference models:

Signal Type	Typical Inference Value
Age set in parental control profile	May directly indicate a declared age band
Restricted access to content/apps	Suggests user is likely under platform's default age (e.g. under 13)
Time-of-day usage restrictions	Often associated with child protection settings
Shared or supervised device profile	Indicates the device is likely used by a minor
App store or browser filter level	Correlated with maturity setting of the user

E.17.4 While these signals do not constitute evidence of actual age, they can enhance the confidence level of a layered inference system – particularly when the platform wishes to apply graduated responses (e.g. flagging for further age assurance only if other signals corroborate youth).

Vendor Case Study

PRIVO®
Privacy • Permission • TRUST

Website
privo.com

PRIVO specialises in child protection and underage user assurance, especially in US COPPA-aligned contexts. Its inference services focus on determining likelihood of under-13 status.

Three Key Facts

1. Parental consent flows. App behaviour and declared age. Shared device patterns.

2. Probabilistic underage detection with platform context. Used to trigger consent or block access.

3. Effective in sandboxed and child-facing app environments. Used only indirect, non-PII signals. High reliability in US-style app gating.

Strengths

- Purpose-built for child protection
- Strong fallback to verifiable consent
- Easily embedded in education and kids' apps

Practice Statement
ageassurance.com.au/v/pvo/#PS

Privacy Policy
ageassurance.com.au/v/pvo/#PP

Technology Trial Test Report
ageassurance.com.au/v/pvo/#TR

Technology Trial Interview
ageassurance.com.au/v/pvo/#VI

Summary of Results

PRIVO delivered a context-specific, privacy-sensitive model for identifying underage users. Its strength lies in early-stage filtering and integration into family-oriented digital ecosystems.

| Role in layered age assurance

E.17.5 This aligns with the endorsement in ISO/IEC FDIS 27566-1 of layered age assurance strategies, where indirect indicators (Clause 6.2.2) are used to:

- Inform decisions without needing biometric or identity data
- Respect user privacy by drawing only on ambient, contextual information
- Avoid unnecessary digital friction in low-risk or early-stage use cases

E.17.6 In most age inference deployments observed during the Trial, parental control signals were treated as self-asserted or indirect indicators – useful in:

- Flagging users likely under 13 or under 18, prompting further verification
- Supporting profile pre-qualification before offering content or features
- Providing device-level context in shared environments

E.17.7 In several Trial deployments:

- Systems queried device-level metadata (where permitted) to detect presence of child profiles or guardian-set age restrictions
- Some platforms used multi-device inference, recognising that a household tablet with strict controls likely belonged to a child or adolescent, even if login credentials were ambiguous
- Inferences drawn from parental control settings were often cross-checked with usage behaviours and content engagement before determining any access restrictions

E.17.8 This approach mirrors the principle in ISO/IEC FDIS 27566-1 that contextual and probabilistic indicators should be used in a way that is proportionate, transparent and explainable, especially when not independently verified.

E.17.9 Parental control signals offer a valuable source of context to age inference systems. While they do not serve as standalone proof of age, their presence can meaningfully boost the accuracy of age classification in a non-intrusive, privacy-respecting manner. When used as part of a layered assurance approach, these signals help reduce false positives and unnecessary escalations – supporting smoother user journeys and greater inclusivity, especially for children using family-managed devices.

E.17.10 To ensure responsible implementation, providers must:

- Transparently disclose the role of such signals in practice statements
- Avoid over-reliance or hard decisions based solely on parental controls
- Calibrate thresholds with awareness of demographic and platform variance

E.17.11 As inference technologies mature, privacy-aware, context-rich indicators like parental controls will become increasingly important in delivering proportionate, user-sensitive age assurance across digital environments.

E.18 Age Inference and Verified Credentials in Digital Wallets

E.18.1 The rise of holder services, such as digital wallets, may enhance age inference capabilities – particularly when used to store and present verified attributes like "Likely Over 18" based on behavioural or contextual signals, rather than a declared date of birth. These credentials could enable low-friction access to age-restricted services without disclosing sensitive personal information.

E.18.2 However, user-centric holder services can also introduce risks. If behavioural or inferred data points are linked across contexts, they could create a persistent digital footprint that raises privacy concerns. For example, an age-inferred credential used for accessing gambling platforms could, if not properly protected, influence unrelated applications such as loan or mortgage assessments diminishing user privacy.

E.18.3 As the adoption of holder services (such as digital wallets) accelerates in Australia and globally, the opportunity to store and present age-related credentials in a privacy-preserving way is expanding. Traditionally, such wallets have stored identity-derived attributes (e.g. date of birth or driver's licence), but increasingly, systems are experimenting with inferred attributes, such as "Likely Over 18," based on behavioural or contextual signals.

E.18.4 These inference-based credentials could be issued by age assurance providers and held in the user's digital wallet to enable low-friction access to age-restricted services – without requiring the user to disclose identity documents or biometric data.

From age inference to verified credential

E.18.5 For an age inference output to be eligible for use as a verified credential within a digital wallet, it must meet key quality and trust requirements. These are articulated in part in ISO/IEC FDIS 27566-1 (particularly Clauses 6.5, 6.6 and 7.1), which emphasise:

Requirement	How Age Inference Can Satisfy It
Confidence in Result	Inference system must demonstrate statistically valid outputs
Binding to Individual	Inference must be bound to a session or device with authentication
Minimised Disclosure	Only age band or binary attribute ("Over 18") is shared
Auditability and Traceability	Credential must indicate method used, source of inference
Contextual Appropriateness	Credential validity may be scoped to domains (e.g. "valid for online content access")

E.18.6 Once these requirements are met, the inference result can be cryptographically signed by the issuing provider and presented to the holder service as a selectively disclosable credential.

Vendor Case Study

YOTI

Website

yoti.com

Yoti participated in the Trial offering behavioural and contextual age inference services. Certified to ISO/IEC 27001 and SOC 2, Yoti aligns with ISO/IEC FDIS 27566-1 and is a recognised leader in privacy-first assurance models.

Three Key Facts

1. Account metadata (e.g., age of account, frequency of use). Content engagement patterns. Device and browser context.

2. Real-time inference at point-of-interaction. Confidence scoring with thresholds. Output as binary decision: Likely Over/ Underage threshold.

3. Used minimum viable signal sets to preserve privacy. Triggered fallback to facial estimation where confidence was low.

Strengths

- Clear documentation and transparency in logic
- Fully operational SDKs and browser-based integrations
- High relevance for low-friction, privacy-focused age gates

Practice Statement

ageassurance.com.au/v/yot/#PS

Privacy Policy

ageassurance.com.au/v/yot/#PP

Technology Trial Test Report

ageassurance.com.au/v/yot/#TR

Technology Trial Interview

ageassurance.com.au/v/yot/#VI

Summary of Results

Yoti's system was mature, privacy-centric and well-aligned to ISO/IEC FDIS 27566-1. Its session-based design, confidence management and fallback logic make it suitable for dynamic and scalable digital deployments.

Worked Example: issuing an inference-based credential

E.18.7 This scenario illustrates how age can be verified without revealing a user's full identity. By analysing behavioural signals and issuing a temporary credential, users can access age-restricted services while maintaining privacy and control over their personal information.

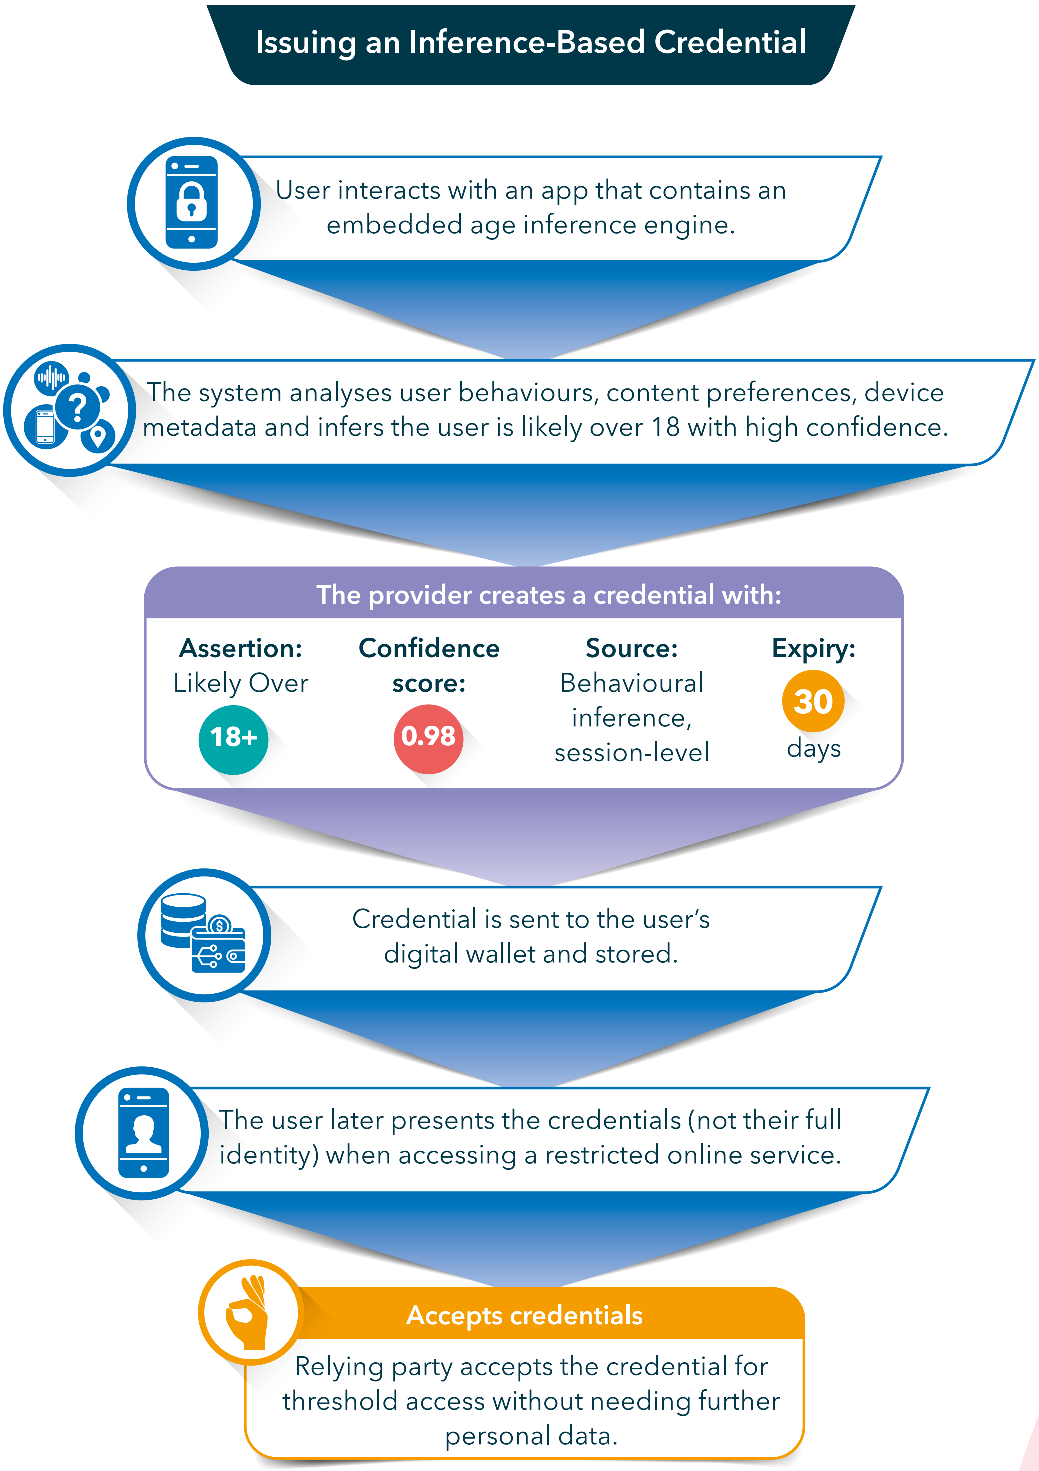

Issuing an Inference-Based Credential

User interacts with an app that contains an embedded age inference engine.

The system analyses user behaviours, content preferences, device metadata and infers the user is likely over 18 with high confidence.

The provider creates a credential with:

Assertion: Likely Over 18+

Confidence score: 0.98

Source: Behavioural inference, session-level

Expiry: 30 days

Credential is sent to the user's digital wallet and stored.

The user later presents the credentials (not their full identity) when accessing a restricted online service.

Accepts credentials

Relying party accepts the credential for threshold access without needing further personal data.

***Figure E.18.1** Issuing an Inference-Based Credential*

LIFEGUA

Privacy and ethical considerations

E.18.8 Despite its promise, this model introduces privacy risks if not carefully governed.

- Cross-context linkage: If the same credential is used across unrelated domains or can be correlated, it may build a persistent profile of the user.
- Secondary use: For instance, if a gambling-access credential is used or visible during a loan or mortgage assessment, it could result in unfair bias or discrimination.
- Digital footprint expansion: Even when no identity is disclosed, repeating behaviour-based credentials in different contexts can lead to unintended re-identification.

E.18.9 ISO/IEC FDIS 27566-1 Clause 6.6 specifically warns against systems that expand a user's digital footprint through unintended reuse of assurance outputs.

Safe use of inference-based credentials

Design Principle	Implementation Strategy
Contextual Scoping	Ensure credentials are tagged for use in specific domains only
Short-lived Credentials	Expiry after 24–72 hours or use-once tokens
User Consent and Control	Let users approve each use and revoke credentials
Zero-Knowledge Proof Options	Use cryptographic proofs to share "Over 18" without disclosing source data
Wallet Governance Standards	Align with emerging wallet certification frameworks (e.g. W3C[6] VC, eIDAS[7])

[6] *This means a W3C (World Wide Web Consortium) standard for digital, cryptographically verifiable credentials.*

[7] *eIDAS (Electronic Identification, Authentication and Trust Services) is an EU regulation that governs digital identification and trust services throughout Europe.*

Cross-reference to Trial observations

E.18.10 During the Trial:

- Some participants proposed offline-compatible credentials issued after inference-based age checks that could be used in later sessions
- Others explored temporary access tokens issued to a wallet with a scoped TTL (time-to-live)
- No participants had yet deployed persistent, reusable inference credentials in wallets at scale, but several expressed interest in further R&D.

E.18.11 Inference-based verified credentials represent a novel and privacy-forward evolution of age assurance. They allow users to prove age eligibility without revealing identity and support seamless digital access across services. However, the potential for data misuse or cross-context leakage highlights the importance of clear policy, technical safeguards, and user governance in maintaining trust, safety, and alignment with the core privacy and security principles as set out in ISO/IEC FDIS 27566-1.

E.19 Security and Information Protection in Age Inference Systems

E.19.1 We found that the age inference systems were generally secure and consistent with information security standards. Most of the providers were able to demonstrate ISO/IEC 27001:2022 certified information security management and some had other supplementary security protocols (such as SoC2 or Fintech-level security).

E.19.2 Security is a foundational requirement of any age assurance technology and the systems evaluated under the Trial showed a strong commitment to secure design and operational integrity. Despite age inference systems typically handling less sensitive data than identity-based verification systems, the behavioural and contextual signals used in inference – particularly if aggregated – still constitute personal data and must be protected against misuse, exposure and unauthorised access.

| Alignment with ISO/IEC FDIS 27566-1 security characteristics

E.19.3 Clause 6.8 of ISO/IEC FDIS 27566-1 defines key security requirements for age assurance systems, which include:

ISO/IEC FDIS 27566-1 Security Characteristic	Relevance to Age Inference Systems Strategy
Integrity (9.3.1)	Ensuring inference logic and model outputs are not altered or spoofed
Confidentiality (5.3.3)	Protecting behavioural signals and context data during processing
Availability (6.5)	Systems must operate consistently during access control decisions
Resistance to Adversarial Attacks (8.3)	Preventing manipulation of inferred results via synthetic behaviours or data injection
Traceability (9.2)	Ensuring auditability of when, how and why an age inference occurred
Minimisation of Residual Risk (11.3)	Implementing fallback mechanisms and layered security for inference logic

Trial observations and provider practices

E.19.4 Most age inference providers demonstrated a high standard of security practice, typically including:

- ISO/IEC 27001:2022 certification for information security management systems (ISMS), indicating well-established frameworks for risk assessment, access control, monitoring and incident response.

E.19.5 Supplementary certifications, such as:

- SOC 2[8] Type II for cloud-based service integrity and availability,
- Fintech-grade security protocols (e.g. end-to-end encryption, secure enclaves for data processing),
- OWASP[9] Top 10 mitigation within frontend interfaces handling sensitive inference inputs.

E.19.6 In many cases, behavioural data was:

- Processed locally on the user's device (e.g. mobile or browser),
- Encrypted in transit and at rest and
- Anonymised or pseudonymised before any storage or aggregation.

[8] *SOC 2 stands for System and Organization Controls 2. SOC 2 is a security framework that specifies how organisations should protect customer data from unauthorized access, security incidents and other vulnerabilities.*

[9] *OWASP means the Open Worldwide Application Security Project. This and other abbreviations used throughout the reports can be found in the Glossary section of Part K.*

| Specific risks in age inference security

E.19.7 Even though age inference avoids directly handling identity credentials or biometrics, it introduces its own risk surface, including:

- Signal spoofing: Users attempting to mimic adult-like behaviours to trigger "Over 18" inference outputs.
- Model tampering: In environments where inference logic is client-side, attackers may attempt to reverse-engineer or modify scoring thresholds
- Inference leakage: If results are stored or reused across services, attackers could deduce behavioural profiles

E.19.8 These risks are addressed in part by:

- Implementing robust input validation and anomaly detection, especially in real-time inference models.
- Applying zero-knowledge proofs where inference output is shared without exposing source logic or input signals.
- Using signed credentials for any inference-based results presented to third-party services (e.g. digital wallets or APIs).

E.20 Mitigating the Risk of False Legends and Manipulated Digital Footprints

E.20.1 Age inference providers were acutely aware of the risks posed by individuals attempting to manipulate their digital footprint to create a false age narrative – also known as establishing a false legend. This includes deliberately generating misleading behavioural or contextual signals to appear older or younger than they are.

E.20.2 We evaluated the resilience of systems to such manipulation, including scenarios where users might simulate age-typical behaviours or inject false data points into the inference workflow. While international standards in this area are still evolving, we found providers actively developing safeguards against data injection and manipulation attacks.

E.20.3 The sophistication, effort and coordination required to construct and maintain a convincing false digital footprint were generally found to be high and unlikely to present a scalable threat. However, as inference systems become more integrated across services and contexts, ongoing vigilance and robust validation frameworks will be essential to maintaining system integrity and trust.

| False legends

E.20.4 One of the unique challenges of age inference systems – particularly those based on behavioural or contextual signals – is the potential for users to intentionally manipulate their digital behaviour to misrepresent their age.

E.20.5 The Trial evaluation recognised this risk and assessed the resilience of inference systems to such manipulation, especially as these systems become more widely integrated into in-app, in-browser and network-level deployments.

| What Is a false legend?

E.20.6 A false legend occurs when a user deliberately engineers their digital presence – such as:

- Simulating adult-like browsing or purchasing behaviour,
- Mimicking child-typical interaction speeds or content consumption patterns,
- Injecting or spoofing metadata (e.g. falsified device signals, fake session histories),
- Repeatedly interacting with certain features to game inference algorithms.

E.20.7 While more difficult to execute than simple misstatement in traditional verification, these techniques pose potential risks in inference-based systems due to their reliance on non-identity-linked data.

| Trial observations and testing

E.20.8 During the Trial:

- Several age inference providers were subjected to test cases involving simulated manipulation (e.g. controlled test accounts exhibiting misleading patterns).
- Systems were also evaluated for data injection vulnerabilities – where external scripts or tampered inputs attempted to feed false signals into the inference model .
- Providers applied safeguards such as:
    - o Anomaly detection algorithms to flag unnatural or inconsistent usage patterns,
    - o Minimum data thresholds before generating an inference,
    - o Use of session entropy checks to detect engineered behaviour,
    - o Restricting inference to real-time inputs, reducing replay attacks.

E.20.9 These approaches reflect a strong security posture, even in the absence of detailed international standards specifically addressing age inference manipulation. Providers referenced adjacent standards such as ISO/IEC 30107 (Biometric Presentation Attack Detection) and emerging practices in adversarial AI mitigation.

| Cross-reference: ISO/IEC FDIS 27566-1 on digital footprint minimisation

E.20.10 Importantly, Clause 6.6 of ISO/IEC FDIS 27566-1 cautions against age assurance systems that:

> *"Expand or accumulate a persistent digital footprint that may be used beyond the original context of assurance."*

E.20.11 Attempting to pre-empt false legends by tracking users over time or across contexts may itself create unintended privacy risks, by:

- Violating data minimisation principles,
- Building cumulative behavioural profiles,
- Introducing long-term identifiers that could enable re-identification.

E.20.12 As such, robust systems must strike a balance between:

- Validating behavioural signals for authenticity and
- Avoiding persistent surveillance or digital profiling.

E.20.13 The systems evaluated generally adhered to this balance by:

- Conducting inference in-session (ephemeral),
- Not storing raw signal data post-inference,
- Issuing age classification decisions without retaining identifiable metadata.

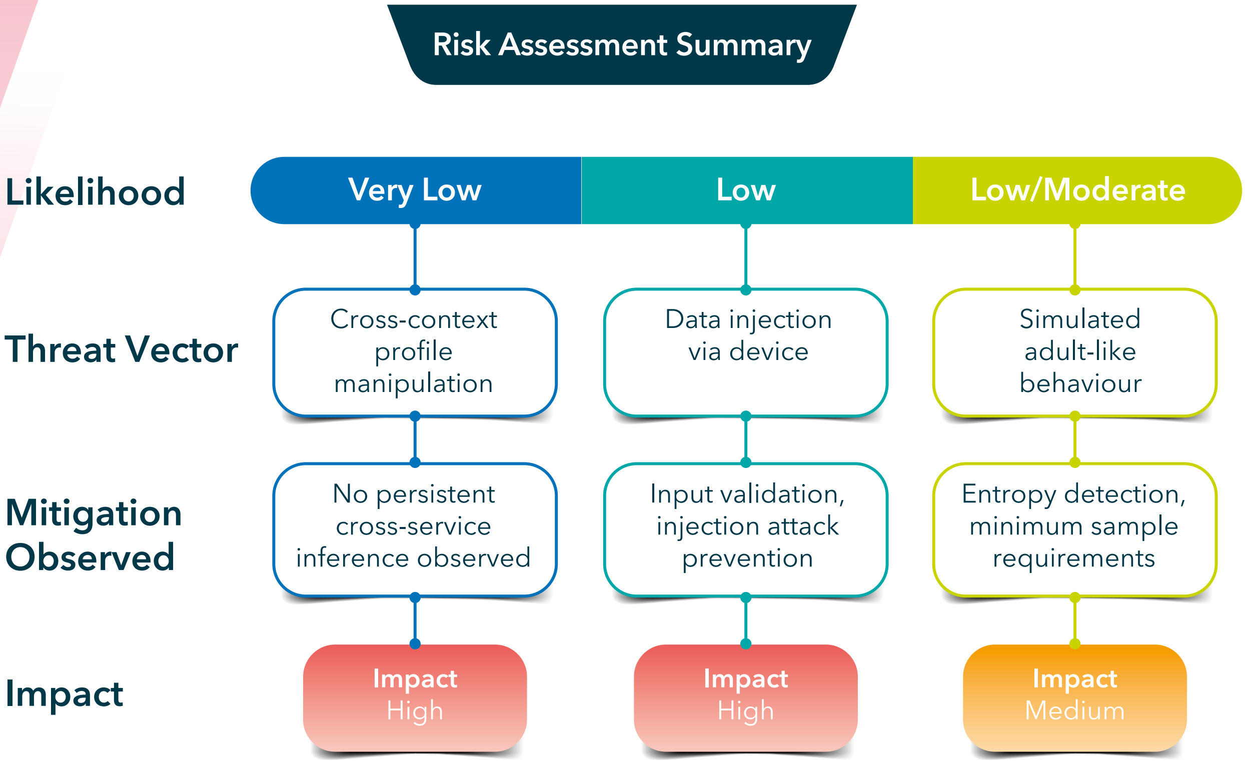

Risk Assessment Summary

Likelihood	Very Low	Low	Low/Moderate
Threat Vector	Cross-context profile manipulation	Data injection via device	Simulated adult-like behaviour
Mitigation Observed	No persistent cross-service inference observed	Input validation, injection attack prevention	Entropy detection, minimum sample requirements
Impact	Impact High	Impact High	Impact Medium

***Figure E.20.1** Risk Assessment Summary*

E.20.14 The age inference systems evaluated in the Trial demonstrated a strong understanding of the false legend risk and were actively implementing technical safeguards to detect and mitigate attempted manipulation. However, as these systems evolve and begin to integrate across apps, platforms and devices, ongoing vigilance will be required to:

- Strengthen resilience to adversarial behaviour,
- Avoid invasive profiling in the name of security and
- Maintain alignment with the guidance on digital footprint minimisation as set out in ISO/IEC FDIS 27566-1.

E.20.15 Further work is also needed to inform the development of specific international standards around inference security, spoofing detection and adversarial signal handling.

E.21 Data Minimisation and Avoidance of Over-Collection in Age Inference Systems

E.21.1 The data-minimising approach taken by age inference providers – particularly independent, third-party services – helped mitigate the risk of over-collection or retention of personal information in anticipation of future regulatory or investigative requests. These providers typically applied digital footprint analytics on a one-time, context-specific basis, inferring age from a discrete set of behavioural or contextual signals and discarding them immediately after use. This reduced the potential for profiling or surveillance and upheld strong privacy standards.

E.21.2 However, the sector could benefit from clearer regulatory guidance to ensure that this practice remains the norm and to prevent gradual drift towards persistent data retention or cumulative behavioural tracking. Independent age assurance providers would benefit from explicit frameworks that balance investigatory needs with privacy-preserving design.

E.21.3 In contrast, account-based online services – such as social media platforms – often employ continuous age inference techniques. These systems monitor user behaviour over time to validate the declared age and detect inconsistencies (or "contra indicators," as defined in ISO/IEC FDIS 27566-1). This allows platforms to intervene when signs suggest a user's actual age may differ from what is recorded, helping maintain compliance with age-related policies. While more effective in ongoing risk mitigation, this model raises greater concerns about long-term data profiling and highlights the need for transparent governance and user safeguards.

E.21.4 A core tenet of responsible age assurance – across verification, estimation and inference – is the principle of data minimisation: collecting only what is necessary, using it for a specific and defined purpose and discarding it as soon as it is no longer required. This principle is critical in the context of age inference, where behavioural and contextual signals – though often low in sensitivity individually – can be privacy-invasive when accumulated or stored over time.

| Strong practice among independent age inference providers

E.21.5 The independent, third-party providers participating in the Trial consistently demonstrated a data-minimising approach to age inference. In these systems:

- Age was inferred using a discrete, session-specific set of signals (e.g. interaction patterns, device context or content preferences),
- Inference occurred in real time or at the point of user action, such as content access or purchase attempt,
- Behavioural signals were immediately discarded after classification,
- Only the resulting age classification (e.g. “Likely Over 18”) – not the raw data – was passed to relying parties or stored.

E.21.6 This model reflects best practice under ISO/IEC FDIS 27566-1 Clause 7.3, which cautions against:

> *“Persistent collection or reuse of personal data that leads to the expansion of a user’s digital footprint.”*

E.21.7 This approach:

- Reduces the risk of surveillance or profiling.
- Prevents reuse of inference signals beyond the original context.
- Supports user trust and regulatory compliance.

| Need for regulatory clarity to prevent drift

E.21.8 Despite this strong starting point, there is a risk of gradual drift towards broader data retention, particularly as:

- Regulators or investigators begin to request traceability of assurance decisions,
- Market pressures push toward persistent credentials or longitudinal assurance tracking,
- Platforms seek to use inference for cross-contextual risk assessment or user analytics.

E.21.9 Without clear legal and policy frameworks, providers may feel compelled to retain behavioural data "just in case," undermining privacy and contradicting the guidance on purpose limitation and proportionality as set out in ISO/IEC FDIS 27566-1.

E.21.10 There is a pressing need for:

- Regulatory guidance on what, if any, audit data is appropriate for retention in inference contexts,
- Clear standards for just-in-time vs persistent inference logic and
- Shared expectations for investigatory cooperation that do not require long-term user profiling

| Comparison: Continuous inference in account-based services

E.21.11 By contrast, large account-based services (e.g. social media platforms, streaming services) often use ongoing, continuous age inference as part of their safety frameworks. These systems:

- Monitor behaviour over time (e.g. likes, follows, language use).
- Apply inference models to detect "contra indicators" – signals that a user's claimed age may be inaccurate (see ISO/IEC FDIS 27566-1, Clause 9.5).
- Use inference to flag accounts for review, suspend access or trigger escalated assurance.

E.21.12 While this model is potentially more effective at identifying age misrepresentation over time, it carries a higher risk of cumulative behavioural tracking:

- Users may be unaware that their interactions are being continuously analysed,
- Inference outputs may be repurposed for other features (e.g. personalisation, marketing),
- Digital footprints expand across sessions, apps and sometimes devices.

E.21.13 This illustrates the privacy trade-offs between contextual, point-in-time inference and persistent age tracking, reinforcing the need for:

- Transparency, so users understand what data is used and how long it is kept,
- Governance mechanisms, such as data minimisation audits,
- Scoped data use, ensuring inference signals aren't misapplied across domains.

Summary and alignment with ISO/IEC FDIS 27566-1

Practice	Independent Providers	Account-Based Platforms
Data Collection	Context-specific, one-time	Continuous, behavioural profiling
Retention of Signals	No (signals discarded post-inference)	Yes (ongoing log of user behaviour)
User Transparency	High (in-session, minimal storage)	Variable (may be opaque)
Alignment with ISO/IEC FDIS 27566-1	Strong (Clauses 6.5, 7.2, 7.3)	Mixed (requires stronger governance)

E.21.14 The Trial found that most age inference providers – particularly independent services – demonstrated commendable privacy-first practices, minimising data collection and avoiding persistent tracking. However, without clear regulatory guardrails, there is a risk of future overreach as inference systems mature and become more widely adopted.

E.21.15 To preserve trust, maintain compliance with ISO/IEC FDIS 27566-1 and balance safety with privacy, future development of inference systems should:

- Uphold just-in-time data use;
- Avoid unnecessary digital footprint expansion;
- Define clear limits on reuse; and
- Implement mechanisms for user oversight and revocation of inference-based credentials.

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 978-1-0681646-4-4

9 781068 164644 >