



Age Assurance Technology Trial

PART J Tech Stack

August 2025



Funded by



Australian Government

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

Project by



Findings on the Tech Stack

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

1

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

2

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

3

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

4

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

5

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

6

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

7

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

8

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

9

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

© 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 978-1-0681646-8-2



Table of contents

Introduction and Overview

I

J.1	Introduction to Part J: Technology Stack	6
J.2	Executive Summary	8
J.3	Who Participated in the Trial of the Tech Stack	13
J.4	What is the "Tech Stack"	14

Context, Standards and Methodology

II

J.5	Exploring the Tech Stack	22
J.6	Platform vs App Store	35

Detailed Analysis of Findings



J.7	Deployment Opportunities in the Tech Stack	44
J.8	App Store Model Analysis	50
J.9	Privacy and Control Considerations	56
J.10	Interoperability	70
J.11	Technology Readiness Assessment	80
J.12	Critical Implementation Challenges	85
J.13	Responsibility and Liability	90
J.14	Proximity to Risk	97
J.15	Virtual Private Networks and Geolocation	101
J.16	Emerging Challenges and Future Considerations for the Tech Stack	113



Age Assurance Technology Trial

PART J Introduction and Overview



J.1 Introduction to Part J: Technology Stack

J.1.1 Part J of the Age Assurance Technology Trial focuses on the technology stack and services that can assist age assurance but are not themselves age assurance. Traditional regulatory models, such as those defined by COPPA¹ in the US or the EU's Audio-Visual Media Services Directive², placed the burden of age restriction directly on the service provider. However, this approach can result in inconsistent or ineffective deployment, as each service implements its own system for determining and enforcing age-appropriate access. Increasingly, alternative models are being proposed – ones that shift responsibility to other layers of the technology stack, including the device, the network or even app stores.

J.1.2 In this part of the report, we examine these emerging approaches to technology stack deployment, with the goal of informing future policy and technical development. We explore:

Emerging approaches

Defining the Technology Stack for Age Assurance and Parental Control or Consent

What constitutes the stack and how its layers – from the user device to the internet backbone – can contribute to or hinder effective age-based access controls.

Deployment Models and Interoperability

The theoretical and practical ways in which age assurance mechanisms can be integrated at different levels of the stack, including how they may interoperate across services and providers.

1. *The Children's Online Privacy Protection Act, or COPPA, is a law that was passed by the United States Congress in 1998 with the aim of protecting the privacy and personally identifying information of children under the age of 13 who use online services.*
2. *The EU's Audiovisual Media Services Directive (AVMSD) governs EU-wide coordination of national legislation on all audiovisual media, both traditional TV broadcasts and on-demand audiovisual media services.*

Emerging approaches

Tools and Services Available Across the Ecosystem	A review of current tools and services offered by app stores, mobile network operators, ISPs ³ , platforms, browsers, geolocation services and website monitoring tools, evaluating how each might contribute to or support age assurance.
Technology Readiness Levels (TRLs)	An assessment of the maturity of individual technologies and of the tech stack as a whole in the context of age assurance. Drawing from submissions to the trial, we highlight discrepancies between perceived and actual readiness.
Risks and Threats	Analysis of the core threats posed by tech stack implementations, focusing on functionality, performance, privacy, security and acceptability – critical considerations in safeguarding children without unduly compromising user rights.
Liability in the Tech Stack Ecosystem	Exploration of accountability across stakeholders – whether platform, developer, device maker or network provider – and whether distributed responsibility leads to fragmented or nefarious outcomes, along with strategies for resolving such liability gaps.
Proximity to Risk	How different placement models for age assurance mechanisms affect their proximity to online harms and the implications of this in designing effective protections for children.

3. This refers to an Internet Service Provider. ISP and other abbreviations can be found in:



Cross Reference: *Part K - Glossary Section*

J.2 Executive Summary

J.2.1 This section of the Trial examined how age assurance, parental consent and control mechanisms could be embedded more systematically across the digital ecosystem by leveraging the technology stack – ranging from user devices and browsers to networks, app stores and backend services. The aim was to explore whether stack-level deployment could move beyond fragmented, service-by-service implementation and support more consistent, interoperable and privacy-conscious approaches to protecting children online.

J.2.2 The evidence gathered through submissions, interviews and analysis suggests that while stack-based models offer real potential, their practical maturity is still limited. Most approaches remain conceptual or at early development stages and few are ready for scalable, real-world deployment. App store-based models were the most fully conceptualised, with companies like Meta and Snap proposing frameworks in which platforms collect and securely share age-related attributes. However, existing implementations by operators such as Apple and Google still rely primarily on self-declared or parent-entered information and do not incorporate independent age verification or support for open, cross-platform interoperability.

J.2.3 Alternative models at the device or network level offer broader enforcement possibilities – especially for browser-based or unauthenticated services – but raise complex questions around privacy, data minimisation and user autonomy. These models, while promising in theory, must overcome significant compliance and usability barriers, particularly in environments where devices are shared or controls are imposed without user awareness.

J.2.4 Several participants proposed approaches to interoperability, including reusable age credentials, digital wallets and orchestration layers that could work across services and jurisdictions. Although diverse in architecture, these models shared a common ambition: enabling a user to verify their age once and reuse that assurance in a privacy-preserving way. However, implementations remain fragmented, technically incompatible and often reliant on proprietary interfaces or ecosystem buy-in that has not yet materialised.

J.2.5 A clear theme across the Trial was the mismatch between participants claimed Technology Readiness Levels (TRLs) and the actual state of deployment or integration. Many systems were rated as mid-to-high TRL despite lacking demonstrated interoperability, system-level testing or platform integration. This suggests that the field is still in an innovation phase, with most solutions yet to be validated in operational environments.

J.2.6 In addition to technical challenges, the Trial identified a range of implementation issues relating to latency, reliability, transparency and user acceptability. Systems operating deep in the stack – such as at the network or browser level – may offer coverage, but risk alienating users through opaque controls or poor alignment with household realities. Similarly, systems that rely on parents to configure or enforce protections must account for digital literacy, language barriers and socio-economic context.

Snapshot of Technology Stack Trial Inputs

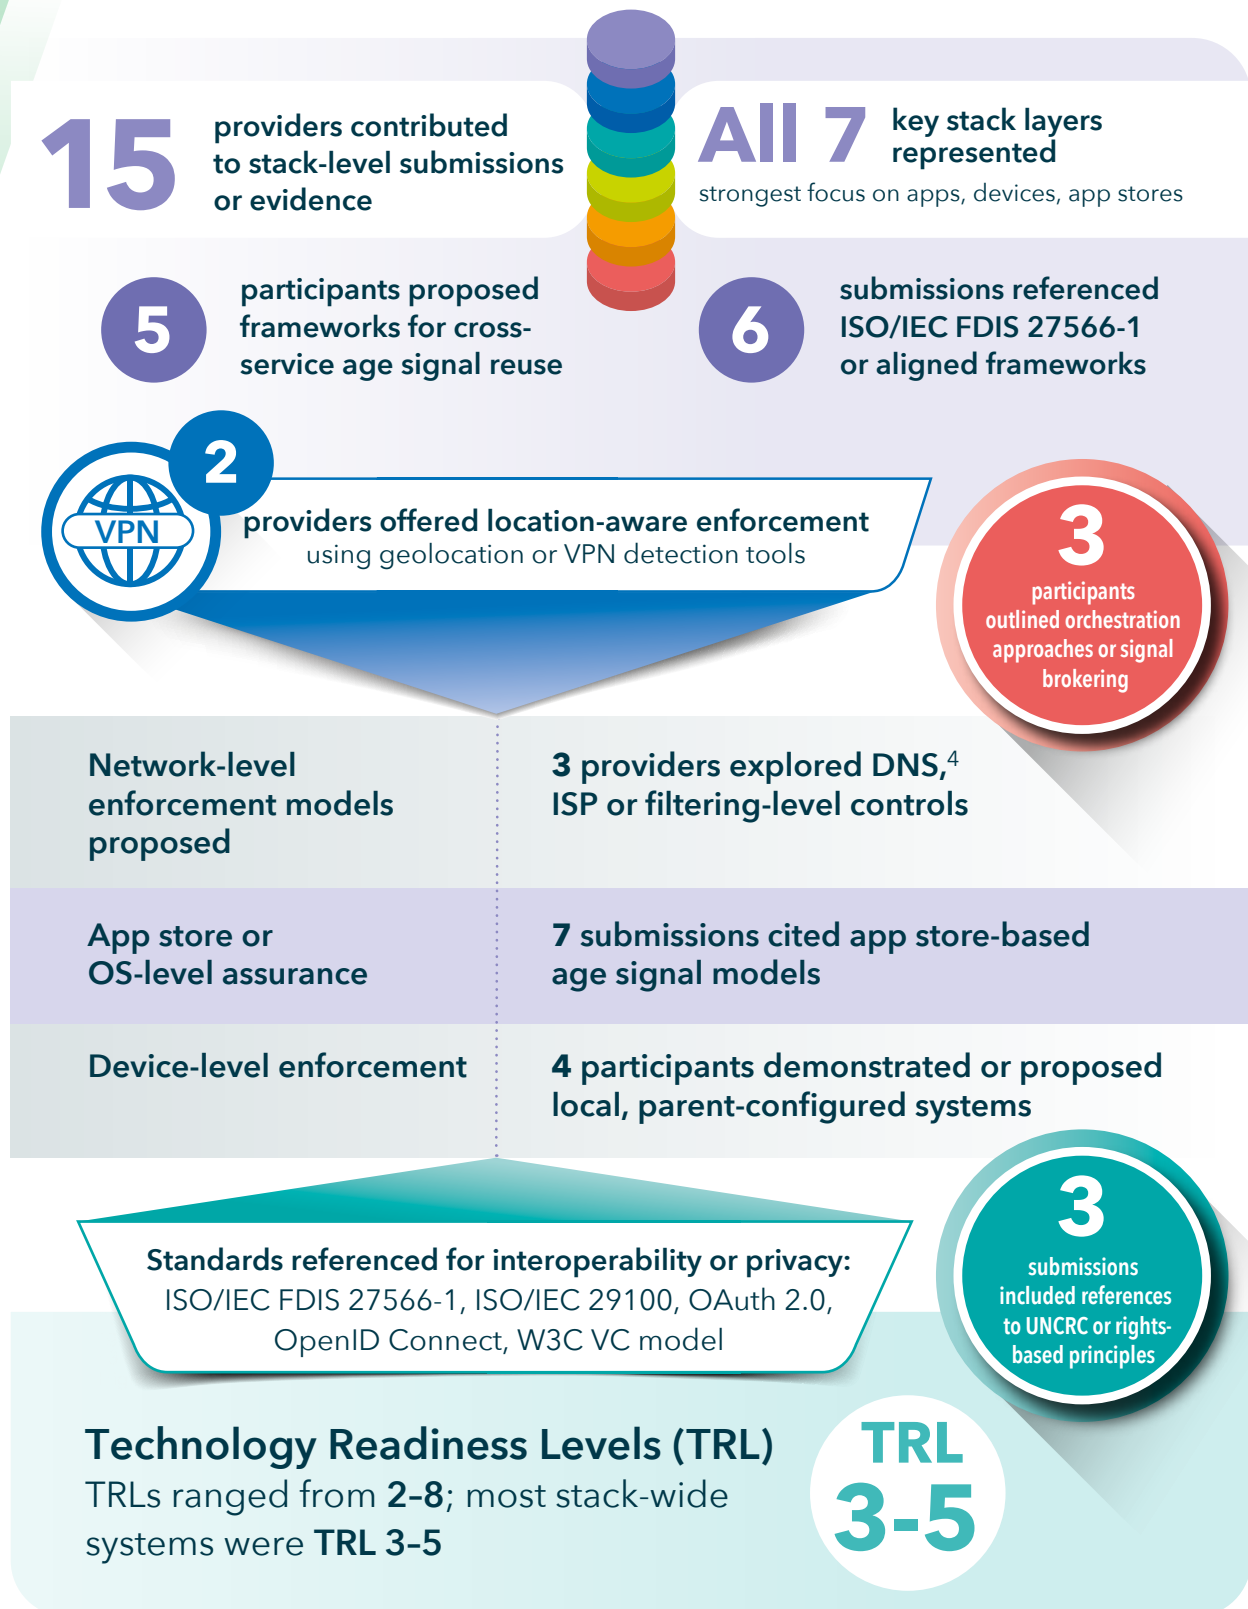


Figure J.2.1 Snapshot of Technology Stack Trial Inputs

4. Domain Name System, or DNS, matches domain names to IP addresses. It is often used in the context of filtering, whereby it is the process of using the DNS to block malicious websites or filter out harmful or inappropriate content..

J.2.7 One of the most pressing gaps exposed through the Trial was the lack of clarity around responsibility and liability. In a distributed system where verification might be performed by one party, credentials issued by another and enforcement triggered by a third, it is not clear who is accountable when protections fail. Participants often defined their role narrowly, disclaiming responsibility for broader outcomes. Without clear legal, regulatory or contractual frameworks, enforcement and redress will be challenging.

J.2.8 The effectiveness of stack-based assurance also depends heavily on its proximity to risk. The closer the mechanism is to the point of potential harm – such as within an app or on a child’s device – the more accurately and responsively it can protect users. However, these mechanisms tend to have narrower reach and depend more heavily on developer or user action. Conversely, distant mechanisms such as app store filters or network-level blocks may provide broader coverage but are less capable of responding to dynamic risks or nuanced behaviours.

J.2.9 The Trial also explored how geolocation and VPN detection services – such as those offered by GeoComply – can support enforcement by identifying circumvention attempts. While not a substitute for age verification, these tools can enhance jurisdictional compliance and help close gaps exploited through spoofing or routing manipulation. The idea that VPN use invalidates compliance obligations is not supported by law or practice. Tools for VPN detection are well established in other regulated sectors and can be effectively applied to age assurance contexts.

J.2.10 Looking ahead, the Trial identified several emerging areas likely to shape the future of stack-based age assurance. These include edge-based enforcement that sits between device and ISP layers; AI-powered moderation linked to age signals; auditability and logging standards for enforcement transparency; and post-quantum credential resilience. Design improvements that centre children and guardians – such as intuitive interfaces, scalable consent tools and support for multi-child households – will be essential. Equally, stack-based systems must function in low-connectivity and shared-access environments to ensure equity and inclusion.

J.2.11 In conclusion, the technology stack holds real promise as an infrastructure layer for systemic, privacy-aware age assurance. However, the current ecosystem is still fragmented, immature and dependent on cooperation from a small number of dominant actors. Realising the full potential of the stack will require sustained development, regulatory clarity and shared standards to support trust, scalability and user protection. The Trial provides a foundational map of what is possible – and what work still lies ahead.

J.3 Who Participated in the Trial of the Tech Stack



euCONSENT



GeoComply[®]



J.4 What is the “Tech Stack”

| Understanding the Term

J.4.1 The term “technology stack” (or “tech stack”) is commonly used in software development and digital infrastructure. It refers to the layers of technologies that work together to deliver a digital service. Think of it like a layered cake – each layer depends on the one beneath it and together they form a whole.

J.4.2 In practical terms, the tech stack includes everything from the device you’re using (like a phone or laptop), to the apps and browsers you access, to the networks that connect you to the internet and the platforms that host the content or services you use. Each layer performs a different role – but they are all interconnected.

| How Does the Tech Stack Work?

J.4.3 When a person opens an app or visits a website, many parts of the tech stack are involved:

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

J.4.4 Each layer has its own role, but they can also be used strategically to apply controls or protections. This is particularly relevant in areas like **age assurance** and **parental consent**.

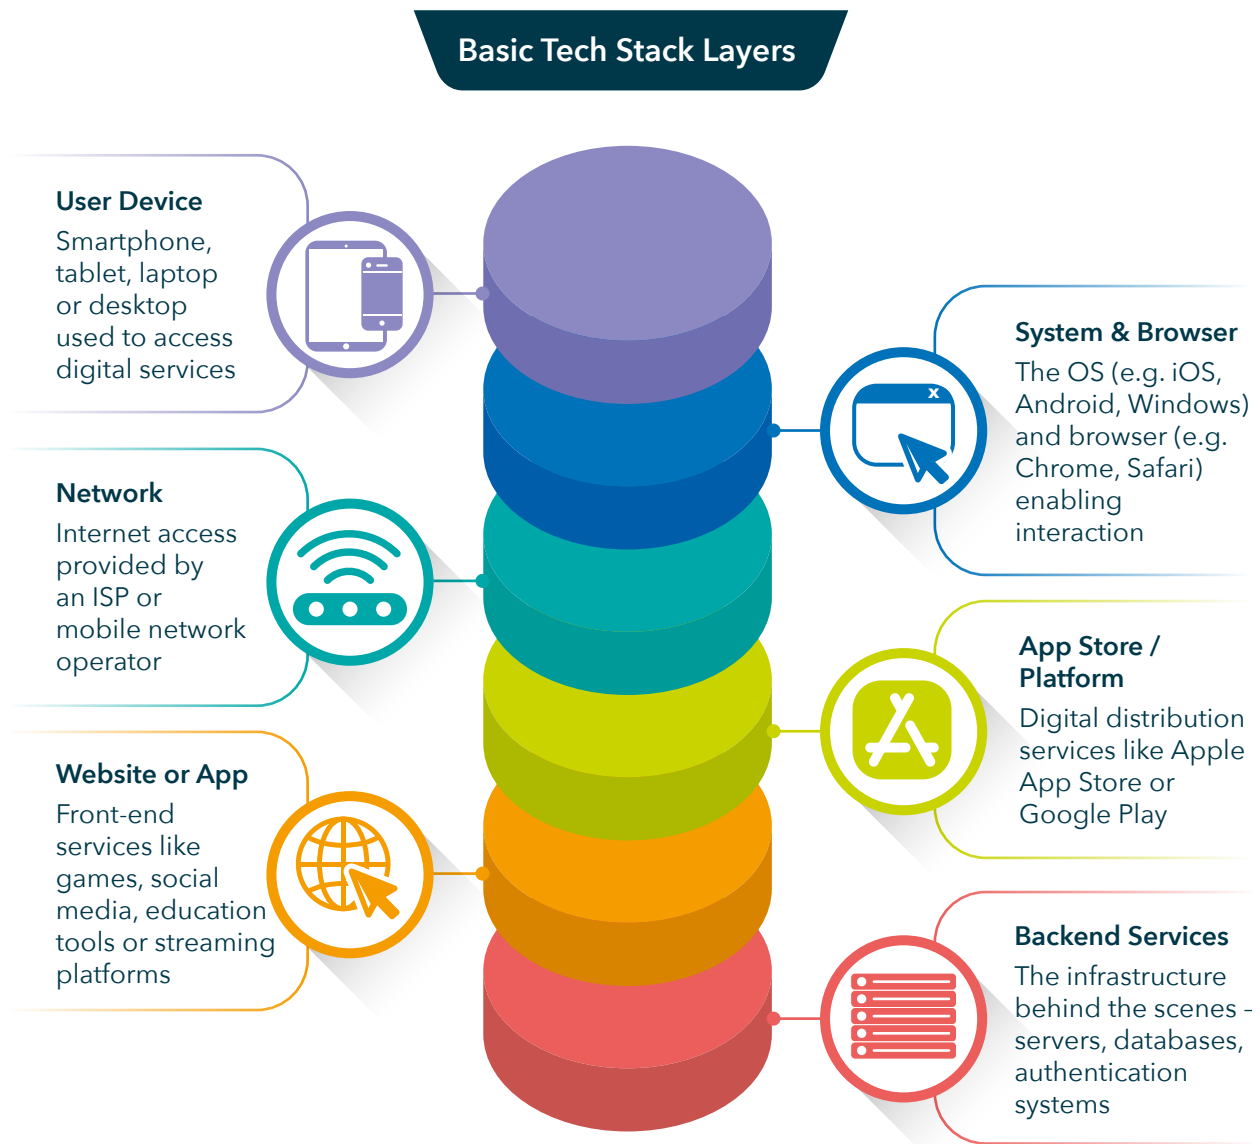


Figure J.4.1 Basic Tech Stack

| Why is the tech stack important for age assurance and parental control?

J.4.5 Historically, most age restrictions have been enforced at the application layer – the service itself asks a user their age or offers a parental control function. This approach relies heavily on self-declaration and has proven easy to bypass, particularly for motivated children and teens.

J.4.6 A growing body of thinking – and several submissions to the Trial – has suggested that placing age controls at other layers of the stack could offer a more reliable, systemic approach. For example:

- **Device-level controls** can help parents manage what apps or websites their child can access.
- **App store settings** may restrict downloads of apps based on age classifications.
- **Network-level solutions** might block or filter content for underage users, regardless of which app or browser is used.
- **Geolocation and device fingerprinting** can help detect attempts to bypass controls via VPNs or fake profiles.

J.4.7 In this way, the tech stack is not just about how services function – it's also a powerful toolset for enforcing age-based protections more consistently and pervasively across the digital ecosystem.

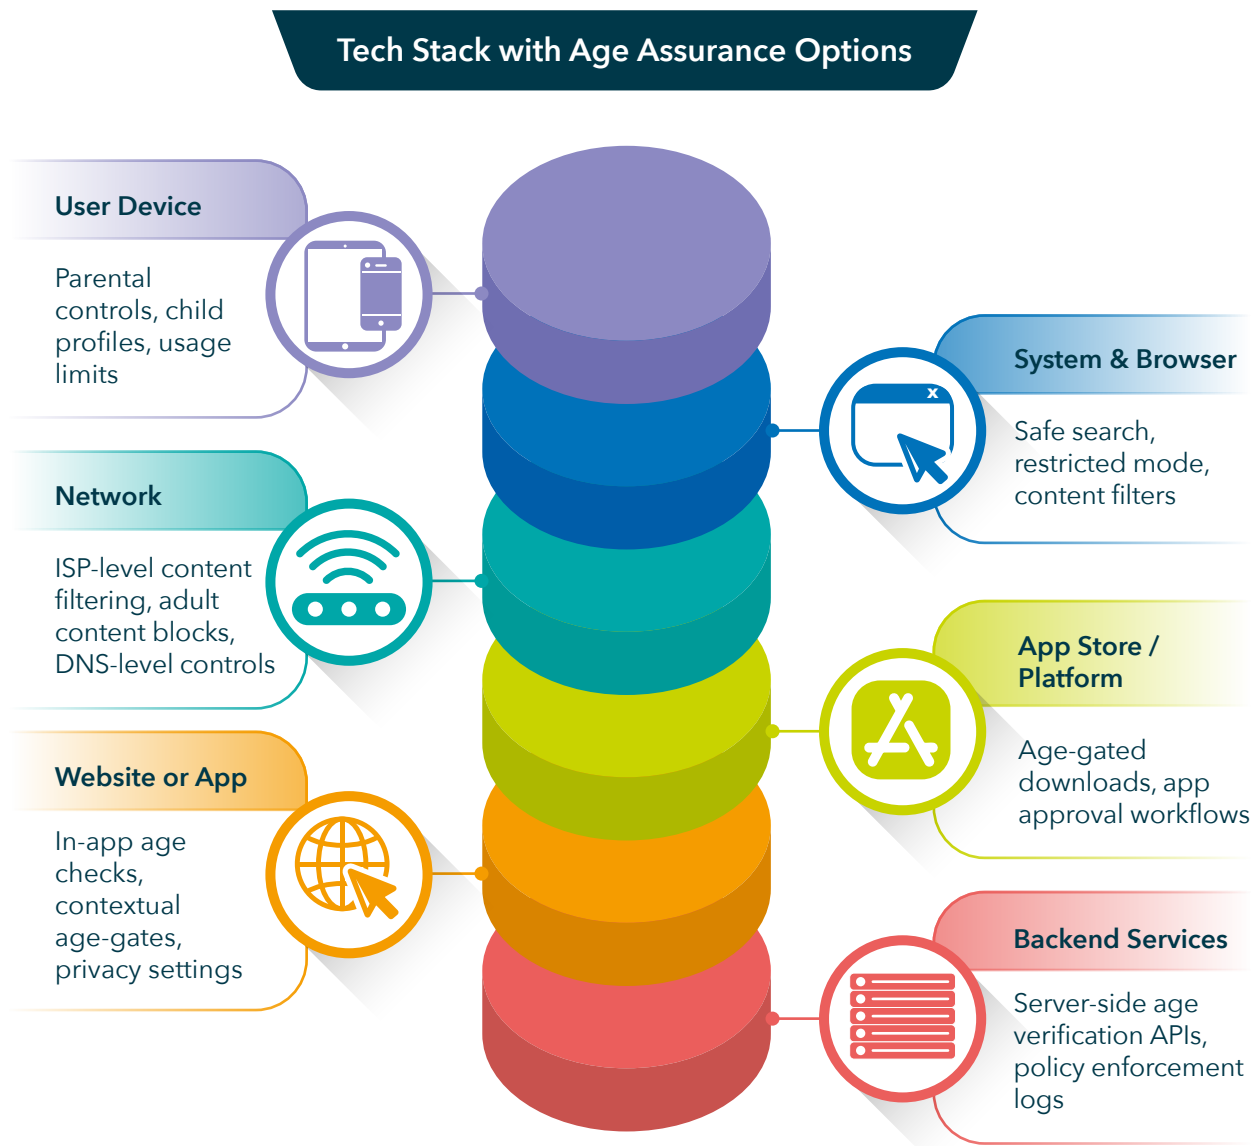


Figure J.4.2 Tech Stack with Age Assurance Options



| Challenges and limitations

J.4.8 Here is a comparison table of age assurance by stack:

Stack Layer	Potential Role in Age Assurance	Strengths	Risk / Limitations	Current Adoption
User Device	Local parental controls, age-based profiles	Device-level enforcement, user-specific settings	Easily bypassed, varies by device/platform	Widely available on phones & tablets
Operating System / Browser	Safe search, restricted modes, default content filters	Broad content filtering, pre-installed tools	Limited granularity, may miss app-level data	Medium to high adoption
Network (ISP)	DNS filtering, content blocking, identity-linked services	Network-wide coverage, no need for device access	Shared devices complicate enforcement, weak identity binding	Moderate adoption in households
App Store / Platform	Age-gated downloads, parental approval flows	Centralised control over app access	Can be bypassed (e.g. sideloading), age self-declaration	High (Apple/Google ecosystems)
Website or App	Onboarding age checks, contextual gating, child-safe modes	Real-time, service-specific control	Inconsistent UX, regulatory fragmentation	Variable (depends on service)
Backend Services	Server-side verification APIs, audit logs, token validation	Centralised logging	Privacy concerns, latency, integration complexity	Emerging (especially for VC tokens)



Age Assurance Technology Trial



PART J

Context, Standards and Methodology



J.5 Exploring the Tech Stack

J.5.1 Having introduced the concept of the technology stack as a layered model underpinning digital services, this section explores each of its key layers in more detail – highlighting their distinct functions, the role they can play in age assurance and parental controls and the relevance to the Australian market. Understanding these components is essential to identifying both opportunities and limitations for systemic online protections.

| User devices

J.5.2 User devices are the primary access point to the internet for most Australians – whether through smartphones, tablets, laptops, desktops, smart TVs or gaming consoles. These devices often include embedded parental control settings, screen time limits and the ability to set up child user profiles.

J.5.3 In Australia, Apple, Samsung and Google dominate the smartphone market, while Microsoft and Apple lead in personal computers. Gaming devices like the Nintendo Switch, Sony PlayStation and Xbox are widely used by children and teenagers.

J.5.4 Many of these platforms provide native parental tools (e.g. Apple's Screen Time, Google's Family Link), but their adoption and correct configuration vary significantly between households.

apple.com/uk/family-sharing/

families.google/familylink/

J.5.5 The Trial noted that several proposed solutions for device-level enforcement depend on these local user configurations, which can vary widely in effectiveness and enforceability.



Implication: Devices offer a direct point of intervention and high proximity to risk, but effectiveness depends on parental engagement and platform compatibility.

| Operating systems and browsers

J.5.6 Operating systems (OS) manage device functions and serve as the environment within which apps and browsers run. Browsers (e.g. Safari, Chrome, Edge) are particularly relevant for web-based services, many of which are not accessed through dedicated apps.

J.5.7 In Australia, iOS (Apple) and Android (Google) dominate mobile OS usage, while Windows is prevalent on desktops and laptops.

J.5.8 Browsers can implement safe search defaults, incognito mode restrictions or integration with child user accounts, but few currently enforce robust age-based filtering or assurance by default.

J.5.9 Some submissions to the Trial referenced browser-level restrictions or plug-ins that could facilitate age-based gating, but most solutions focused on other layers.



Implication: While browsers and OS environments are widely used, their ability to enforce age controls independently is limited and often user dependent.

| Network infrastructure (ISPs and MNOs)

J.5.10 Network providers – both fixed-line and mobile – serve as intermediaries between the user and online content. In theory, they are well positioned to offer systemic protections because they see all traffic flowing to and from a user's device.

J.5.11 Major providers in Australia include Telstra, Optus, TPG Telecom (Vodafone/iiNet) and Aussie Broadband.

J.5.12 Some offer optional family filtering services or allow customers to block adult content. However, such filters are usually broad and not age-dynamic (i.e., they do not differentiate between a 6-year-old and a 15-year-old).

J.5.13 The Trial received limited proposals from network-level providers but did identify conceptual models for ISP-based or DNS-layer filtering integrated with verified user age attributes.



Implication: The network layer has high coverage potential, but raises serious privacy, scalability and accuracy concerns if not designed with proportionality and user transparency.

| App stores and platforms

J.5.14 App stores are often the gatekeepers to the digital ecosystem. They can restrict what users download based on device settings, age declarations or parental permissions.

J.5.15 The Australian market is dominated by the Apple App Store and Google Play Store, with additional niche platforms like Steam, Epic Games Store and console-specific stores (e.g. PlayStation Store, Nintendo eShop).

J.5.16 Some parental features are in place (e.g. Family Sharing, ask-to-buy features), but none of the major platforms independently verify age or expose standardised age attributes to apps.

J.5.17 Trial participants identified app-store based assurance as the most conceptually developed model, proposing that verified ages could be stored at the app-store level and shared downstream to services. However, this remains largely theoretical and lacks support from major platform providers.



Implication: App stores are strategically positioned for systemic age control but face governance and adoption barriers, particularly due to platform control by a small number of global players.

| Apps and websites

J.5.18 This is the layer where most age assurance currently occurs – or fails to. Individual services (e.g. YouTube, TikTok, Roblox) implement their own systems for account creation, age gates, content filtering and parental features.

J.5.19 Many services accessed by Australian children rely on self-declared age with no verification. Others offer parent dashboards or teen account modes, but these differ widely in approach and effectiveness.

J.5.20 The Trial included multiple service providers whose solutions included onboarding age checks, in-app parental controls or external age verification integrations.

J.5.21 However, these solutions often operate in isolation and are not interoperable across platforms.



Implication: Service-level controls are essential but fragmented. Reliance on voluntary or proprietary measures leaves many children unprotected or inconsistently treated.

| Backend services

J.5.22 Behind every visible digital service are backend systems – servers, databases and APIs⁵ that store data, manage content and perform key logic. This is where age attributes may be processed, stored or shared.

J.5.23 Major backend providers include Amazon Web Services (AWS), Google Cloud and Microsoft Azure, as well as specialised providers in identity verification, content moderation or parental oversight platforms.

J.5.24 The Trial received several submissions focused on age verification, age estimation (e.g. via AI models) and interoperable attribute sharing, all of which operate largely at the backend.

J.5.25 These solutions face data protection and interoperability challenges, especially when attempting to share verified age data across services and jurisdictions.



Implication: Backend services are critical for enabling interoperable, privacy-preserving age assurance, but require robust trust frameworks and architectural alignment to scale.

5. This refers to an Application Programming Interface. API and other abbreviations can be found in:



Cross Reference: Part K - Glossary Section

| Geolocation services

J.5.26 Geolocation services determine a user's physical location based on their IP address, GPS data, Wi-Fi signals or other metadata. These services are already widely used across industries for purposes such as regional content restrictions, fraud detection and targeted advertising. In the context of age assurance, geolocation has growing relevance, particularly in enforcing jurisdiction-specific age restrictions and detecting attempts to circumvent controls via VPNs or proxy services.

J.5.27 In Australia, providers such as GeoComply, MaxMind, IP2Location and Google's Geolocation API are used to estimate a user's region or city with varying degrees of precision.

J.5.28 Submissions to the Trial highlighted that VPN use among under-18s is widespread, particularly to access adult content or evade restrictions on gaming or social media platforms.

J.5.29 Geolocation tools can be used to flag when a user's IP address does not match expected regional patterns, suggesting possible use of a VPN or anonymiser.

J.5.30 These tools may also support geo-fencing, where services adjust access or features based on a user's location (e.g., enabling different levels of verification in line with local legislation).

VPN and Geolocation Mismatch Map

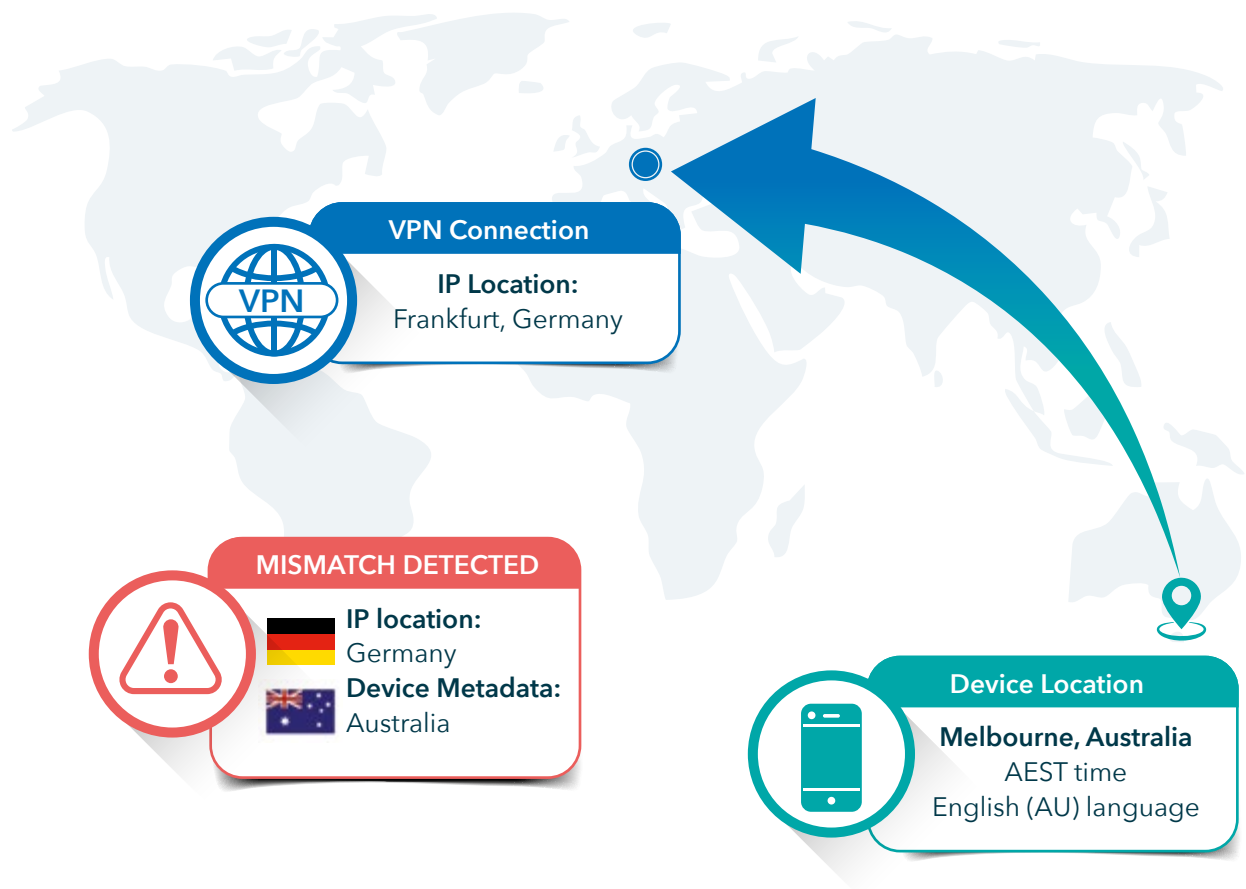


Figure J.5.1 VPN and Geolocation Mismatch



Implication: Geolocation services offer a valuable enforcement tool, particularly for cross-jurisdictional compliance and VPN detection. However, they are not foolproof – location spoofing is possible and aggressive implementation may impact user privacy or accessibility for legitimate users (e.g., those travelling).

| Web crawlers and sweeper services

J.5.31 Web crawler or “sweeper” services systematically scan and index internet content. Originally developed for search engines and compliance monitoring, these tools are now increasingly used for content classification, risk analysis and policy enforcement – including in the field of age assurance.

J.5.32 These services can be used to detect the presence of age-inappropriate content, identify non-compliant websites or monitor whether child protection measures are being properly implemented.

J.5.33 In the context, web sweeper tools could play a role in identifying websites that:

- Do not apply proper age gates
- Feature harmful or adult content accessible without verification
- Provide circumvention tools (e.g., guides for spoofing age or using VPNs)

J.5.34 Examples of relevant tools include Netsweeper, NetClean, SafeToNet and enterprise-grade web monitoring platforms, many of which are used by schools, regulators and telecom providers in Australia for safeguarding and filtering.

J.5.35 These tools can also help generate and maintain blocklists or allowlists used by ISPs, app stores or parental control systems.



Implication: Sweeper services provide important visibility across the broader online ecosystem, enabling enforcement agencies, platforms or parents to monitor and respond to risks at scale. However, they typically operate retrospectively and are most effective when combined with proactive gatekeeping mechanisms.

Website Mapping Using Automated Crawlers

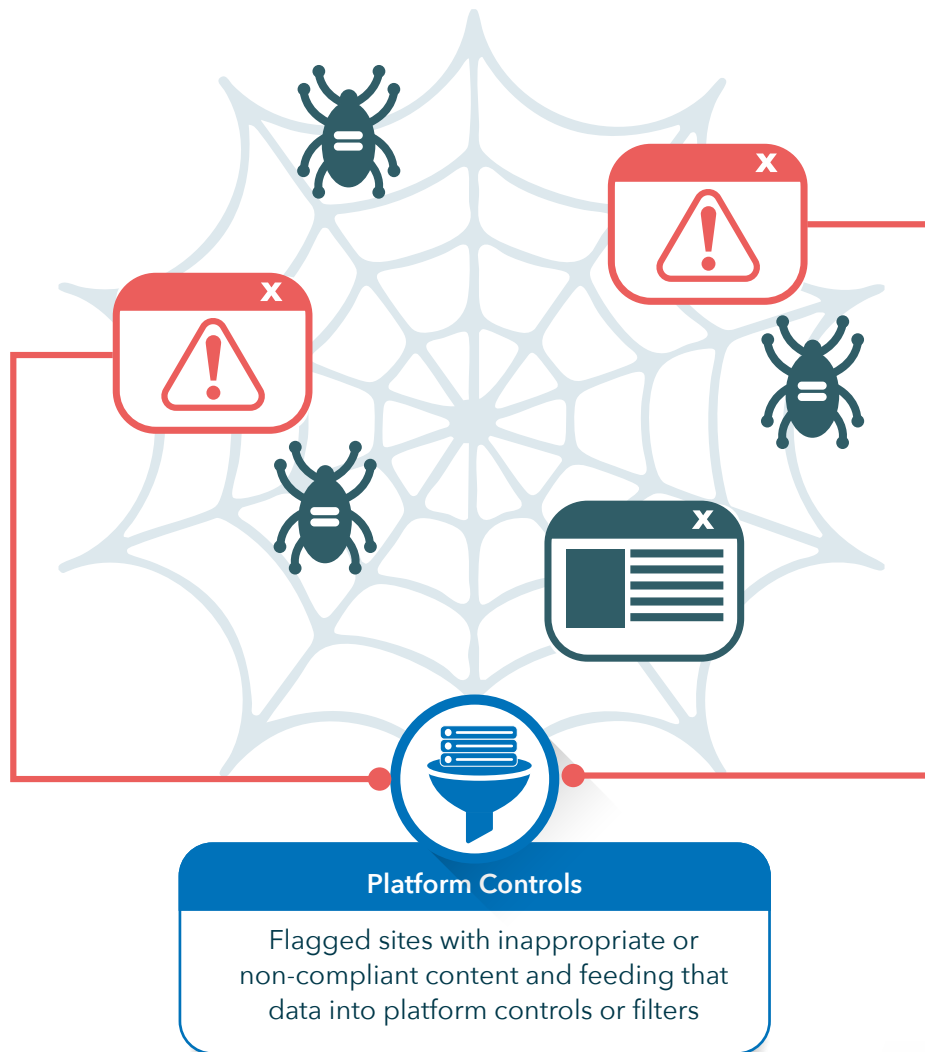


Figure J.5.2 Website Mapping Using Automated Crawlers

| Connecting the layers: how the Trial mapped the stack

J.5.36 Throughout the Trial, submissions and interviews revealed that many participants were implicitly working across different layers of the tech stack – even if their solution focused primarily on one. For example:

- Some apps relied on device-level controls or OS settings to enable parental oversight.
- Others integrated with backend identity providers to verify age or share attributes downstream.
- A few explored browser extensions or network-level signals to restrict access or detect circumvention.

J.5.37 While no participant deployed a full tech stack solution, several theoretical models were submitted that outlined how controls at different layers could interoperate. The Trial highlighted that no single layer could provide complete protection – but that by combining efforts across the stack, a more robust and scalable solution may emerge.

| Applying standards to the tech stack: observations from the Trial

J.5.38 During the Trial, it became clear that many proposed solutions – while well-intentioned – were at varying stages of alignment with relevant standards. Key observations included:

- Few participants referenced technical standards beyond ISO/IEC FDIS 27566-1, indicating a gap between solution design and underlying internet or interoperability standards.
- Some solutions that claimed interoperability lacked implementation of recognised attribute sharing protocols (e.g. OpenID Connect or Verifiable Credentials).
- Accessibility and user-centred design considerations – central to W3C’s and IEEE’s guidance – were often overlooked, particularly in child-facing age verification flows.
- Cross-stack implementations (e.g. app-store to in-app signal sharing) require standards for data hand-off, trust models and attribute governance – yet no universal framework currently exists.



Conclusion: Applying international standards consistently across the tech stack is essential to achieving scalable, interoperable and trustworthy age assurance systems.

This requires not just compliance with age assurance frameworks like ISO/IEC FDIS 27566-1, but also technical adherence to identity, privacy, accessibility and data exchange standards developed by global bodies such as IETF, W3C, IEEE and ISO.

J.5.39 The following visual (*Figure J.5.3*) outlines how various global standards bodies contribute to defining best practices, protocols and safeguards at each layer of the digital infrastructure relevant to age assurance and parental control.

Standards Across the Age Assurance Tech Stack

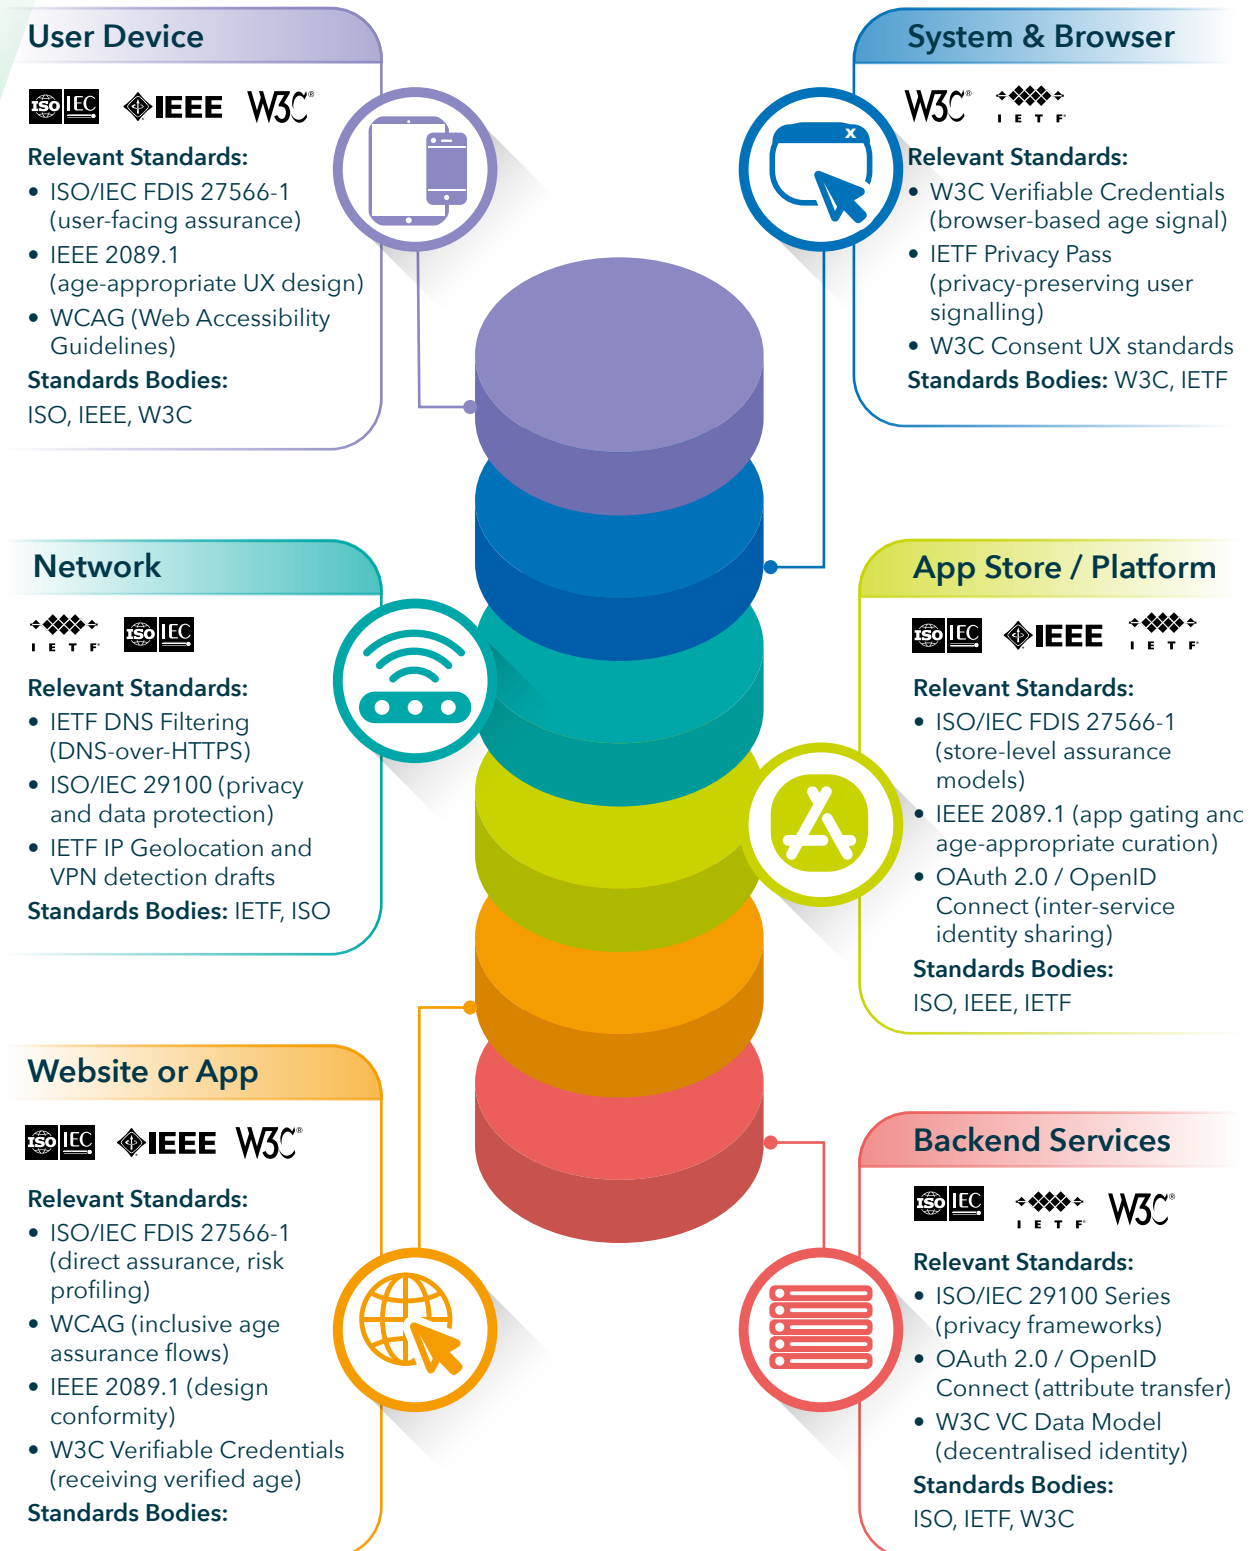


Figure J.5.3 Mapping International Standards Across the Age Assurance Technology Stack

J.6 Platform vs App Store

J.6.1 A key theme across the submissions and interviews from major social media platforms – including Snap and Meta – is a strong preference for age assurance to be implemented at the app store or operating system (OS) level, rather than on a service-by-service basis. This approach is often described as a “whole-of-ecosystem” or centralised model and contrasts with the traditional expectation that individual services verify users’ ages themselves.

J.6.2 Under this proposed model, age assurance – including age verification, parental consent and age-based content gating – would occur when a user first sets up their device or attempts to download an app. App stores like those run by Apple and Google would act as the central point for collecting and verifying a user’s age and then share age signals (such as an age range) with apps via secure APIs. This would allow services to apply age-appropriate settings or restrict access without having to re-verify each user themselves.

J.6.3 The platforms argue that this model would:

- Reduce friction and duplication for users and parents;
- Enable consistent age enforcement across services;
- Simplify parental oversight via built-in OS tools; and
- Support privacy by limiting the sharing of sensitive data to basic age-range signals.

J.6.4 For example:

1. **Meta** proposed an OS-level API that allows apps to receive the user's age range and parental approval status, rather than collecting this information independently for every app.
2. **Snap** proposed a digital wallet-based model in which assured age attributes are stored centrally with the OS or platform account (e.g. Google or Apple ID) and accessed as needed by services.
3. While **Apple** and **Google** - as the operators of the two dominant global app stores - are central to any model that embeds age assurance at the app store or OS level, neither company has formally endorsed the proposal as put forward by social media platforms such as Meta and Snap.



| Apple's position

J.6.5 Apple's submission to the Trial (dated June 2025) outlines a comprehensive suite of new and enhanced features designed to support age-appropriate experiences across its ecosystem. These include:

- Child Accounts integrated into Family Sharing, requiring parental linkage and default safety settings.
- The Declared Age Range API, allowing apps to request a child's age range without accessing precise birth dates.
- Updated App Store age ratings, expanded to include new granular bands (13+, 16+, 18+).
- Tools like PermissionKit, Screen Time and FamilyControls, which allow fine-grained parental control and user experience restrictions.

J.6.6 Apple's approach supports the spirit of centralised age assurance but retains tight control over how age data is shared and used, emphasising privacy and user choice. Importantly, while Apple has created APIs that allow apps to access age range information, it does so only with explicit parental permission and does not provide an open, interoperable assurance framework as proposed by Meta or Snap.

J.6.7 Apple has not signalled support for a federated or third-party interoperable system, nor has it suggested that it would verify or vouch for users' actual ages beyond declared information. In short, Apple is moving toward enhanced app store governance and parental empowerment, but within a closed, Apple-defined framework, rather than an ecosystem-wide model.

| Google's position

J.6.8 While Google did not submit a formal statement for publication in the Trial documents, its position can be inferred from:

- The Google Play parental controls that allow age-based app filtering and require parental approval for downloads.
- Family Link, which enables parents to manage screen time, content filters and approvals for apps and purchases.
- The Google Credential Manager API, mentioned in Snap's proposal, which could be extended to support interoperable age signals in the future.

J.6.9 Unlike Apple, Google has not publicly introduced an equivalent to the Declared Age Range API. However, its infrastructure (e.g. Google accounts linked to birth dates and integration with Android OS) could technically support similar models.

J.6.10 Critically, neither Apple nor Google currently perform independent age verification at the app store level. They rely on self-declared information during account creation and app-based age ratings as a signal for parents, but not as a regulatory mechanism.

| Summary

J.6.11 Both Apple and Google have developed tools that align partially with the app store-level age assurance model:

- They provide app ratings, parental controls and child account structures.
- They allow for developer-accessible APIs to tailor age-appropriate experiences.

J.6.12 Despite these features, they stop short of endorsing or implementing a cross-platform, interoperable age assurance infrastructure as envisioned by platforms like Meta and Snap.

J.6.13 This strategic hesitation reflects deeper issues around platform governance, liability and control over user data. It also underscores a core challenge: while social media platforms are pushing for systemic age assurance at the app store level, actual implementation depends on the willingness of a small number of dominant tech companies to enable or share control of these processes.

J.6.14 This model is central to current industry thinking about scalable and interoperable age assurance, particularly in light of growing international regulatory pressure. The Trial has treated this proposal as a conceptual deployment model within the technology stack and has explored both its potential benefits and implementation challenges, which are examined in detail in the following sections.

J.6.15 This debate about App Store-level age assurance connects closely to several of the key findings from the Trial. It illustrates broader tensions around where responsibility sits within the tech ecosystem, how systemic solutions might be implemented and what trade-offs exist between effectiveness, privacy and feasibility. Here's how the debate supports and intersects with the Trial's findings:

1. **Systemic, interoperable solutions are needed**

The app store-level model proposed by Meta and Snap directly responds to this by offering a centralised solution that could apply consistently across thousands of apps. Instead of asking each app to perform its own age check, a shared age signal issued at the platform level would enable interoperability and reduce duplication.

2. **Responsibility across the tech stack is diffuse and needs clarification**

This debate exemplifies the problem. Social media platforms argue that Apple and Google should take responsibility as central gatekeepers of the digital ecosystem, while Apple and Google position themselves as enablers, not enforcers. This back-and-forth highlight the lack of clarity over who is ultimately accountable for ensuring underage users are protected – a central challenge the Trial identified.

3. **Technology readiness and adoption are not aligned**

The app store-level model is a case in point: while platforms like Meta have clearly outlined how it could work, the actual operators of app stores have not adopted it and in some cases (e.g. Apple), have built alternative tools that stop short of full interoperability or external data sharing. The gap between concept and deployment reflects the broader TRL discrepancies observed during the Trial.

4. **Privacy, usability and acceptability must be balanced**

App store-level age assurance is often positioned by its proponents as more privacy-preserving and easier for parents, since it would only require age verification once, rather than on every app. This supports the Trial's finding that solutions must be both effective and minimally invasive, especially when dealing with sensitive data like children's ages or government IDs.

5. **Proximity to risk matters**

Proponents of app store-level assurance argue that it achieves both proximity and scale – by acting at the “golden moment” when a child first gets a device or tries to download an app. This aligns with the Trial's emphasis on interventions that occur early and are enforced broadly, rather than relying on back-end mechanisms that only trigger after harm has occurred.

6. **Innovation is emerging, but fragmented**

The proposals from Snap (digital wallets), Meta (age signal APIs) and Apple (Declared Age Range API) show different and unaligned approaches to solving the same problem. The lack of standardisation or mutual recognition highlights the fragmentation of innovation.

J.6.16 The debate over app store-level age assurance encapsulates many of the core themes of the Trial:

- It illustrates the push toward systemic, scalable solutions.
- It raises real questions about accountability in a layered digital ecosystem.
- It reflects both promise and limitation in terms of technological readiness.
- It underscores that coordination among dominant infrastructure providers is essential if any truly ecosystem-wide age assurance model is to succeed.



Age Assurance Technology Trial



PART J

Detailed Analysis of Tech Stack Findings



J.7 Deployment Opportunities in the Tech Stack

| Summary finding

J.7.1 Technology stack deployment offers potential for systemic and interoperable age assurance. Theoretical models indicate that placing age assurance mechanisms at different layers of the technology stack – such as on the user’s device, within the network infrastructure or at the app-store level – could provide consistent and cross-cutting protections across services. Interoperability across components will be essential to realise this potential.

| Detailed analysis

J.7.2 Technology stack deployment refers to the strategic placement of age assurance mechanisms at different layers of the digital infrastructure – ranging from user devices and operating systems, to networks, app stores and backend services. In the context of age assurance, this approach moves beyond isolated, service-level enforcement, enabling broader, systemic protections that can be applied across multiple digital environments.

J.7.3 The Trial explored this concept both theoretically and through participant submissions. While few end-to-end stack-based implementations currently exist, several proposals presented early models or architectural designs in which age signals, parental controls or content gating are managed at foundational levels of the stack – such as during device setup, app store downloads or network-level filtering.

J.7.4 This concept aligns with the framework set out in ISO/IEC FDIS 27566-1 (Age Assurance Systems), which is technology-agnostic but supports flexible implementation of age assurance methods across a digital ecosystem. It also draws on adjacent technical standards for identity, security, accessibility and privacy, including OpenID Connect, ISO/IEC 29100 and IEEE 2089.1.

| Observed models from the Trial






J.7.5 Participants proposed several models for stack-level deployment of age assurance:

- Device-level approaches relied on establishing a verified age or parental control configuration during device setup (e.g. Apple Child Accounts), with age-based policies enforced through operating system features.
- App store-level models, particularly championed by Meta and Snap, envisaged app stores collecting verified age and parental approval and securely passing this to apps via APIs.
- Network-layer concepts proposed that Internet Service Providers or Mobile Network Operators apply content filtering or access restrictions based on user age or account type.
- Backend and federated identity models included use of digital wallets or credential frameworks to store and share verified age attributes across services, offering potential for ecosystem-wide assurance without re-verification.

J.7.6 These models suggest that integrated, cross-service age assurance could be achievable if mechanisms across the stack are aligned, interoperable and governed by shared standards.

Interoperability and standards landscape

J.7.7 The effectiveness of stack-based deployment depends on the ability of components across the stack to communicate securely and consistently. Relevant standards and frameworks include:

International Standards		
	ISO/IEC FDIS 27566-1	Functional specification for age assurance systems
	OAuth 2.0 / OpenID Connect	Protocols for secure identity and attribute sharing
	W3C Verifiable Credentials	Enabling privacy-preserving, decentralised age proofing
	ISO/IEC 29100	Privacy frameworks for handling sensitive user data
	WCAG and IEEE 2089.1	Ensuring age assurance systems are accessible and developmentally appropriate

J.7.8 However, the Trial highlighted that integration of these standards into working systems is still at an early stage. Many solutions proposed proprietary approaches or required adaptation of platform policies (e.g. app store data sharing) that have not yet been standardised.

Opportunities for technological improvement

J.7.9 The analysis of Trial submissions, supported by vendor interviews and industry input, points to several opportunities for improving the technological underpinnings of stack-based age assurance:

- Developing interoperable APIs and protocols that allow secure, consent-based transmission of age signals across stack layers (e.g. from app store to app or OS to browser), reducing the need for repeated verification.
- Enhancing alignment between platform-level age controls and service-level enforcement, such that settings applied at the OS or account level can inform access control or user experience in downstream apps.
- Expanding support for parental configuration tools and controls at the device and OS level, which may reduce friction for families and improve the usability of age assurance systems.
- Implementing shared data formats and governance structures to facilitate consistent interpretation of age attributes (e.g. age ranges vs. precise age) while preserving user privacy.

J.7.10 Exploring wider application of existing digital identity frameworks, including those used in government, banking or telecommunications sectors, to support trusted issuance and verification of age without exposing sensitive personal information.

J.7.11 Each of these areas represents an ongoing area of technical development and innovation. The Trial findings suggest that while individual components may already exist in the market, their integration into coherent, scalable and interoperable stack-level solutions remains an emerging challenge.

Benefits and considerations

J.7.12 If implemented effectively, technology stack deployment could support:

- Consistent user experiences across apps and services.
- Reduced burden on individual app providers.
- Centralised parental control with less need for repetitive configuration.
- Privacy-preserving approaches that avoid excessive data collection.
- Improved access control for younger users, especially where enforcement is linked to verified accounts or profiles.

J.7.13 However, the evaluation also identified important considerations:

- Adoption by platform operators remains a critical dependency. For example, the app store-level model requires cooperation from Apple and Google, who have taken different paths toward age-related feature development but have not formally embraced full-scale interoperable assurance.
- Proximity to risk varies by deployment layer. While network-level controls may be broad, device- or service-level controls may be more accurate and timelier.
- Public trust and transparency are vital, especially where enforcement occurs invisibly or is perceived as surveillance.

J.7.14 These considerations reflect the broader findings of the Trial: that a layered, interoperable and privacy-conscious approach is essential to advancing robust age assurance – but that real-world implementation will require coordination across actors, standardisation of mechanisms and thoughtful user engagement.

| What tech stack deployment will not provide

J.7.15 While technology stack deployments offer the potential for more systemic and consistent age assurance, they do not provide a complete solution in isolation. Most notably, app store or OS-level models cannot verify or control what happens within an app or service after download. Even if a child's age is accurately verified at the point of device setup or app installation, the app itself may still need to apply age-appropriate experiences, restrict certain features or detect attempts at circumvention. In this way, stack-level controls set the foundation, but enforcement still relies on downstream actors to uphold protections within their own environments.

J.7.16 Furthermore, deployments at the app store or OS level do not detect or prevent all forms of misuse or evasion. For example, a verified age might be linked to an account, but that account could be shared among multiple users or used on a shared family device without supervision. Stack-level models also do not typically address real-time risks such as grooming, bullying or exposure to harmful content, which often require dynamic monitoring, moderation or user reporting systems within the service itself. Nor do they inherently provide ongoing behavioural insights, which some platforms use to detect underage use or risky interactions after initial verification. As such, while stack deployments can improve baseline confidence in a user's age, they must be part of a broader ecosystem of protections, not a substitute for service-level responsibility.

J.8 App Store Model Analysis

| Summary finding

J.8.1 App-store based models are being developed but lack critical adoption and verification features. While app-store based models were the most fully conceptualised, they currently rely on self-declared or parent-set age information. Without independent age verification and without adoption by key operators such as Apple and Google, these models do not currently meet the criteria for robust age assurance.

| Detailed analysis

J.8.2 App store-based age assurance models represent one of the most fully conceptualised strategies for embedding systemic age controls into the digital ecosystem. In this model, app stores – operated globally by a small number of dominant providers, primarily Apple and Google – act as central gatekeepers, capturing or inferring age-related information during account setup or device onboarding and sharing this with apps through dedicated application programming interfaces (APIs).

J.8.3 This model has been advocated by major social media platforms including Meta and Snap as a means to reduce duplication, increase consistency and shift age assurance closer to the point where users first access the digital environment. The Trial received detailed submissions reflecting these views, as well as public documentation and input from the platform providers themselves. While both Apple and Google are moving toward expanded support for age-based configuration and parental oversight, their approaches differ significantly in architecture, openness and policy posture.

| Apple's approach: controlled, privacy-centric implementation

J.8.4 Apple's submission to the Trial set out an ecosystem-centred model focused on parent-verified child accounts, age-based content filtering and minimal data sharing. The key components include:

- Child Accounts created via Family Sharing, which link children's profiles to parental oversight.
- Declared Age Range API, which allows third-party apps to access a user's age band (e.g. under 13, 13-15, etc.) if parental permission is granted.
- Integration with frameworks like PermissionKit and FamilyControls, which manage user permissions, app gating and screen time.
- App Store metadata and ratings to restrict access to age-inappropriate content at the point of download.

J.8.5 This model emphasises strong parental control, user privacy and Apple's system-level governance over device and ecosystem interactions.

J.8.6 There are opportunities for improvement:

- Apple's approach currently relies on parent-declared age information without independent verification. This presents an opportunity to explore whether secure, privacy-protective verification mechanisms could be embedded in setup flows.
- The model is closed to external assurance providers, which may limit interoperability across platforms or services. Opportunities exist for developing standards-based interfaces that enable recognised third-party signals to integrate with Apple's framework.

| Google's approach: flexible and decentralised architecture

J.8.7 Google, while not providing a formal submission, has publicly implemented a suite of age-related and parental control features through:

- Google Play Store controls that allow download restrictions by age rating.
- Family Link, a parental management system for supervising child accounts, usage and permissions.
- The Credential Manager API, which could in the future support age attribute sharing or consent signals across applications.
- The Android ecosystem's general openness, which allows third-party app stores and extensive developer customisation.

J.8.8 This approach places greater emphasis on user flexibility and developer extensibility, allowing services and families to configure controls based on their needs and technical capacity.

J.8.9 Opportunities for improvement:

- Google's more decentralised architecture can lead to fragmentation and inconsistent enforcement, especially across devices or custom Android builds. This may present opportunities for designing reference implementations or encouraging cross-vendor consistency in applying age controls.
- Parental controls are available but not always activated by default. There is scope to explore mechanisms that could guide or prompt setup at key "onboarding moments" (e.g., first login or device activation) to improve adoption.

| What these approaches aim to achieve

J.8.10 Both Apple and Google appear focused on providing age-appropriate user experiences, reducing the risk of inappropriate content exposure and empowering parents with tools to oversee their children's digital lives. These approaches are designed to:

- Establish a consistent age context across services accessed through the app store.
- Simplify compliance for app developers by enabling system-level gating.
- Reduce the need for services to collect or process sensitive age data directly.
- Embed parental involvement into the setup and permission process for younger users.

| What these approaches do not provide

J.8.11 Despite their ambitions, neither approach currently provides independently verified age assurance as defined by ISO/IEC FDIS 27566-1. Both Apple and Google rely on self-declared or parent-declared age information and do not (as of the time of the Trial) offer age verification based on identity documents, biometric estimation or third-party credentials.

J.8.12 They also do not prevent children from using shared devices, creating new accounts with falsified age or accessing services via web browsers or sideloaded apps. In practice, these methods may not prevent circumvention without additional, service-level controls.

J.8.13 Moreover, once a user enters an app, enforcement relies on the app provider to apply age-appropriate functionality. The app store can gate access, but it does not directly enforce safe design, moderation or behaviour-based protections within individual services.

| Liability and market power considerations

J.8.14 The centralisation of age assurance within app stores raises significant questions about liability, governance and competition:

- **Liability:** If app stores become the primary source of age signals, the extent to which they could (or should) be held responsible for downstream failures in content moderation or child protection remains ambiguous. Conversely, developers relying on app store-provided signals may limit their own obligations or capabilities to verify age.
- **Market dominance:** Entrusting age assurance to two global platform operators raises concerns about barriers to entry for smaller age assurance providers, app developers or new technologies. The concentration of power may risk stifling innovation, particularly where platform rules, APIs or commercial policies limit integration by external services.
- **Innovation and diversity of approach:** The richness of the age assurance ecosystem – including biometrics, identity wallets, school-issued credentials and parental confirmation methods – could be diminished if platform gatekeepers define a single approach. This raises longer-term concerns about choice, flexibility and contextual appropriateness for different users or jurisdictions.

J.8.15 There are opportunities for improvement:

- To preserve a diverse and innovative assurance ecosystem, there is scope to explore models in which app stores recognise verified age attributes from multiple sources, rather than enforcing a single method or exclusive API.
- Liability frameworks could evolve to provide clarity around shared responsibilities, particularly in cross-stack scenarios where multiple actors contribute to the age assurance process.
- Privacy safeguards can be further strengthened by implementing granular, user-consented data sharing, where apps receive only the minimum information required to enforce protections (e.g. an age range, not a date of birth).

Privacy and user visibility

J.8.16 App store-level models may reduce the number of times age data is collected – but they also introduce new privacy risks. Sharing a user’s age range across multiple services, even with limitations, increases the visibility of children online. If these signals are not adequately protected, they could be misused for targeting, profiling or discrimination.

J.8.17 There is also potential for unintended identification of child users, especially if age signals are combined with other identifiers or used by services in ways that are not fully transparent. These concerns highlight the importance of data minimisation, transparency and usage controls – areas where further technical development may offer valuable safeguards.

J.9 Privacy and Control Considerations

| Summary finding

J.9.1 Deployment at the network or device level raises privacy and control considerations. Implementing age assurance at the device or ISP level could enable broader coverage, including services accessed through browsers or third-party platforms. However, these approaches raise significant concerns regarding user privacy, autonomy and data protection compliance.

| Detailed analysis

J.9.2 Deployment of age assurance mechanisms at the device or network (ISP) level is often seen as a way to provide broad, cross-service enforcement, particularly in contexts where age controls applied by individual apps or websites can be bypassed. These approaches operate at foundational points in the digital stack – either at the user’s hardware or at the gateway to internet services – and therefore offer the potential to restrict or shape access regardless of which service, browser or app is used.

J.9.3 Participants in the Trial explored theoretical and early-stage models for such deployment. While these were not tested in operational form during the Trial, the conceptual analysis highlighted both the coverage potential and substantial risks associated with device-level and network-level interventions. These include challenges related to privacy, autonomy, technical feasibility and compliance with data protection laws.

| Device-level deployment

J.9.4 At the device level, age assurance can be implemented by configuring user profiles with age-related attributes, enforced through the operating system and associated apps. This may occur during initial setup, as part of a child account configuration or via parental control tools embedded in the OS.

J.9.5 Here are some example from the Trial and market:

- Apple's Child Accounts and Screen Time allow parents to establish age-linked usage limits, restrict access to certain content and approve or deny app downloads.
- Android's Family Link enables similar parental oversight, including app approvals and time management tools.
- Some device-level proposals in the Trial explored extending age signals into web browsers or using trusted device identity to enforce age-gated access across services.



J.9.6 These are some potential benefits:

- Broad enforcement across installed apps, web browsers and services accessed via the device.
- Proximity to user behaviour, enabling more context-aware and responsive controls.
- Parental usability, with a single point of configuration and oversight for a child's entire digital experience.

J.9.7 These are some privacy and control considerations:

- Device-level approaches often involve persistent identifiers and profile-level metadata that could be used to infer or disclose the user's age. If not securely handled, this may increase exposure to profiling or tracking.
- These models risk limiting user autonomy, particularly in households where devices are shared or where a child matures beyond the default configuration. The risk of "over-blocking" or age-based discrimination may arise.
- Devices may be configured with inaccurate age data, either through error or deliberate evasion, with no independent verification to support assurance claims.

J.9.8 These are some privacy and control considerations:

- Implementing device-level controls in an aligned way requires that data minimisation, transparency and user control principles are applied. Where controls are imposed without clear user consent, legality under privacy laws (such as the Australian Privacy Act or GDPR) may be questioned.
- Children's rights to privacy and evolving autonomy must be balanced against protective measures – particularly as devices become more personalised over time.

| Examples of device-based solutions from the Trial

J.9.9 Several participants in the Trial explored device-based age assurance models – solutions that operate entirely or primarily on the user’s local device, without requiring persistent connectivity to a backend server or the integration of third-party SDKs⁶. These approaches typically leverage on-device processing, cryptographic storage or parent-driven configuration to support privacy-preserving, user-centric control over age gating.

J.9.10 Device-based solutions differ from SDK or API-led models by prioritising local enforcement, data minimisation and user control. In many cases, they aim to avoid the transmission of personal or biometric data across networks, instead using the device as the enforcement boundary. This makes them particularly well suited to contexts with limited connectivity, privacy sensitivity or low trust in external processors.

6. SDK refers to Software Development Kits. This and other abbreviations can be found in:



Cross Reference: Part K - Glossary Section

Vendor Case Study



PRIVATELY

Website

privately.eu

Privately offers a lightweight, on-device facial age estimation system designed for privacy-by-design deployments, especially in youth-focused contexts such as education and family settings.

Three Key Facts

1

Age checks run directly on user devices, ensuring privacy without sending data to servers.

2

Real-time processing on edge devices provides instant results without internet dependence.

3

Demonstrates efficient, secure AI by leveraging local hardware for biometric age estimation.

Observation

Strong privacy by design: zero biometric retention, on-device only architecture, no cloud contact – aligns with ISO/IEC FDIS 27566-1 expectations.

Flexible integration options with local fallback mechanisms and threshold tuning make Privately adaptable across different use cases and risk levels.

Practice Statement

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

Privacy Policy

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

Technology Trial Test Report

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

Technology Trial Interview

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

Summary of Results

Privately's edge-based solution is a strong example of a privacy-preserving, operationally ready technology that eliminates the need for identity documents or server infrastructure – demonstrating no significant technical limitations to deployment in Australian contexts.

| Characteristics of device-based approaches

J.9.11 Device-based models in the Trial shared a number of common features:

- No persistent cloud-based identity or tracking: User data, including age, is processed and retained locally, with minimal or no outbound data transmission.
- Parent-managed control models: These systems often empower parents to configure access via one-time setup processes, including PINs or physical device pairing.
- Low integration complexity for third-party apps: By working independently of app-level SDKs, these solutions reduce implementation overhead for service providers.
- Friction minimisation for compliant users: Once configured, access is seamless for users, avoiding repeated prompts or verification events.



J.9.12 However, these benefits are offset by challenges:

- Limited ecosystem reach: Without integration into broader identity or app store systems, device-based models may not be recognised by third-party services or interoperable across platforms.
- Difficulty with shared devices: When multiple users share the same hardware (e.g. tablets or home computers), device-level enforcement may either over-restrict or under-protect individual users.
- Parent-device dependency: Solutions like SafeGen rely on the parent taking initial action and may not cover children who use unmanaged or hand-me-down devices.

J.9.13 These examples demonstrate that on-device age assurance can offer a powerful, privacy-preserving path forward – particularly when high assurance is needed without sacrificing anonymity or user experience. As a category, these tools may be particularly valuable in schools, families or jurisdictions that prioritise data sovereignty or that lack national digital identity infrastructure. Their design also presents opportunities for further development around interoperability, credential caching and hybrid assurance models that combine local processing with optional network-based validation.

| Risks of device sharing

J.9.14 Device sharing among young people presents a risk to confidence in age assurance because it undermines the core assumption that digital interactions are linked to a specific, known user. When multiple individuals – especially children of different ages or a mix of children and adults – share a device, age-based protections can become inaccurate, ineffective or easy to bypass.

J.9.15 Key risks associated with device sharing include:

- **Misapplication of age signals:** If one child verifies their age (or has a profile set up by a parent) and another child uses the same device, the second user may inherit age-appropriate or inappropriate access based on the first user's credentials or settings. For example, a 14-year-old could access content intended for a 17-year-old sibling.
- **Inaccurate enforcement of restrictions:** Systems that rely on device-level attributes – such as browser settings, app store configurations or declared age on the OS – assume a 1:1 relationship between user and device. Shared use violates this assumption and can either over-restrict or under-protect depending on who is using the device at the time.
- **Bypassing in-app controls:** Even when apps implement robust in-service age gates, switching between user accounts or using guest mode on shared devices may allow children to circumvent restrictions, particularly if devices are shared informally without parental oversight.
- **Lack of persistent identity:** In shared environments like households, schools or public access points (e.g., libraries), users may not log in to unique profiles. This makes it difficult to apply persistent, user-specific controls and reduces the reliability of behavioural or session-based verification tools.

J.9.16 These are the implications for age assurance:

- Higher false positive/negative rates: Device sharing increases the risk of both failing to protect children (false negatives) and unnecessarily blocking legitimate access (false positives).
- Challenges for interoperability: Shared devices create confusion when age credentials are stored locally (e.g. in a digital wallet or app) but used by multiple people – raising privacy, consent and trust issues.
- Design pressure for child-specific profiles: This risk highlights the importance of implementing systems that can differentiate users on shared devices – for example, through separate OS-level accounts, biometric login or parent-managed access profiles.

J.9.17 In short, unless age assurance solutions can distinguish between users on a shared device, they risk being easily circumvented or inappropriately applied. This is particularly important in multi-child households, lower-income families with limited devices or settings like after-school programs where shared access is the norm.



| Network-level deployment (ISP/MNO)

J.9.18 Network-level approaches propose that Internet Service Providers (ISPs) or Mobile Network Operators (MNOs) implement content filtering, age-based access controls or user profile restrictions at the point where internet traffic leaves or enters a household or device.

J.9.19 These are some examples from the Trial and market:

- Some conceptual models submitted to the Trial included DNS-level filtering, IP-based age signal detection or mechanisms that block certain domains for underage profiles.
- Australian ISPs such as Telstra, Optus and TPG already offer optional parental filtering services, though these are typically based on general content categories, not verified age.

J.9.20 These are some potential benefits:

- Service-agnostic coverage: These controls can apply regardless of which app, browser or device is used.
- Early interception of harmful content, especially for non-authenticated or non-profiled users.
- Low setup burden for families, if deployed as default or “opt-in” features by ISPs.

J.9.21 These are some privacy and control considerations:

- Network-level age assurance may involve traffic inspection, device profiling or user metadata analysis, all of which raise significant privacy risks.
- Controls are often invisible to users, leading to concerns about transparency, consent and the ability to challenge or modify restrictions.
- If user identities or attributes are inferred from device behaviour or IP addresses, the risk of false positives or misclassification increases.

J.9.22 Australia - Mobile Network Operators (as at July 2025)

Provider	Parental / Content Filtering	Age Assurance / Signal Sharing	Notes / Relevant Programs
Telstra	Offers Broadband Protect and Mobile Protect with category-based filtering (e.g., adult, gambling).	No verified age signals or sharing with apps or services.	Parental controls configurable via customer portal or app.
Optus	Provides Smart Controls for website or keyword blocking via modem-level settings; requires opt-in.	No age verification or interoperable age assurance mechanisms.	Offers digital safety education resources for families.
TPG Telecom	Includes Family Friendly Filter options; content filtering configurable per device.	No support for verified age signals or cross-platform enforcement.	Encourages use of optional controls; aligned to voluntary ISP codes.
Vodafone (TPG group)	Basic content filtering and family safety tips; relies on user-set restrictions.	No system-level age assurance or age signal API integration.	Offers a Digital Parenting hub with advice but no enforcement infrastructure.
Aussie Broadband	Offers guidance for setting up content filtering via modem or third-party tools; no default age-based restrictions.	No age verification capabilities or user-profile enforcement.	Parental controls optional and require manual configuration.

**This table is based on open source searching because the mobile network operators did not participate in the Trial.*

J.9.23 These are some compliance considerations:

- Use of network data for age determination may require specific legal authority or user consent, especially where personal identifiers are involved.
- Proportionality becomes a key issue: applying broad content filters or restrictions based on inferred age must be weighed against users' rights to access and freedom of expression.
- Shared connections (e.g., public Wi-Fi or family households) complicate enforcement and may result in over- or under-restriction.

Limitations and risks across both models

J.9.24 Despite their reach, both device- and network-level models share limitations:

- Lack of independent verification: Without validated age data, these models often rely on inferred or declared attributes that may not meet assurance thresholds.
- Risk of circumvention: Users may switch devices, change networks or use VPNs to bypass restrictions.
- Potential for profiling: Persistent identification of devices or traffic can lead to user tracking or surveillance concerns, especially for minors.
- Technical rigidity: Once configured, controls may be hard to adjust in real time or may fail to respond to the user's evolving maturity and autonomy.

Opportunities for technological improvement

J.9.25 The Trial surfaced several areas where further development could improve the utility and trustworthiness of device- and network-level deployments:

- Privacy-preserving architectures: There is potential to explore models that enable enforcement without persistent identifiers or intrusive data processing – such as through local enforcement agents, tokenised access or edge computing approaches.
- User-centric control interfaces: Designing tools that provide clear visibility and manageability for users (including parents and older children) could improve adoption and trust.
- Context-aware enforcement: Integration of dynamic policy engines could enable restrictions that adjust based on time, content type or other contextual signals, reducing the risk of over-blocking.
- Alignment with identity standards: Device and network controls could interoperate with external identity or assurance systems, supporting layered verification without centralising control in a single actor.

J.9.26 These opportunities highlight the potential of these models to provide wider coverage, especially for services accessed outside logged-in environments – but also the critical need for safeguards to protect privacy, maintain flexibility and support user rights.

Vendor Case Study



Website

netsweeper.com

Netsweeper participated in the Trial as a provider of network-based content filtering technology used by ISPs and governments. Originally developed for education settings, its platform now supports country-scale enforcement by applying content and access policies at the network level - without requiring software on user devices.

Three Key Facts

1

Minimises data collection; no age inference; AI-based filtering; supports privacy compliance without storing identifiable user information.

2

Netsweeper operates at the ISP gateway, where it can filter websites based on content type and presence of age gates.

3

Privacy risks exist; ethical use requires transparency, opt-outs, and adherence to child rights and data minimisation.

Practice Statement

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

Technology Trial Test Report

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

Privacy Policy

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

Technology Trial Interview

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

Summary of Results

Netsweeper illustrates how network-level enforcement can support national age assurance goals by complementing - rather than replacing - device or service-level controls. Its utility lies in broad, systemic content governance, especially in environments where individual services lack authentication or oversight.

J.10 Interoperability

| Summary Finding

J.10.1 Interoperability solutions are emerging but remain early-stage and non-standardised. Several Trial participants proposed mechanisms to support interoperability across different age assurance systems. These are still nascent and varied in design, making it difficult to generalise about their functionality or maturity.

| Detailed analysis

J.10.2 Interoperability refers to the capacity of different systems, services or providers to work together seamlessly – transferring trusted age signals or credentials across platforms, services or devices without re-verifying the user. This concept has gained increasing attention as the age assurance ecosystem grows more complex and fragmented. Within the Trial, several participants proposed early-stage mechanisms aimed at supporting interoperable age assurance, reflecting growing recognition that single-provider or siloed solutions are unlikely to meet emerging regulatory and user needs at scale.

J.10.3 While none of these systems are yet mature or standardised, the Trial surfaced a clear direction of travel: toward reusable, privacy-preserving age credentials that can operate across services and jurisdictions. This section explores three such approaches – euConsent’s Age Aware infrastructure, General Identity’s GIP protocol and interoperable digital wallets proposed by participants like Snap and Opale.io.

| euConsent: the age aware infrastructure

J.10.4 The euConsent Age Aware project, supported by European regulators and industry stakeholders, aims to develop a pan-European, standards-based framework for interoperable age assurance. While not a formal Trial participant, its influence was acknowledged in submissions from several vendors.

J.10.5 Key features include:

- Federated identity architecture: Users can verify age once with a trusted provider (e.g., government eID, AV provider), then use that verified age to access multiple regulated services without repeating the process.
- Double-anonymity protocols: The system allows websites to confirm that an age check has occurred without learning the user's identity or verification method.
- Cross-border operability: Built to comply with the eIDAS framework and GDPR, supporting recognition of credentials across EU Member States.

J.10.6 This model highlights the opportunity for interoperable credentials to minimise friction and enhance privacy, but also underscores the need for consistent standards, trust frameworks and aligned governance.

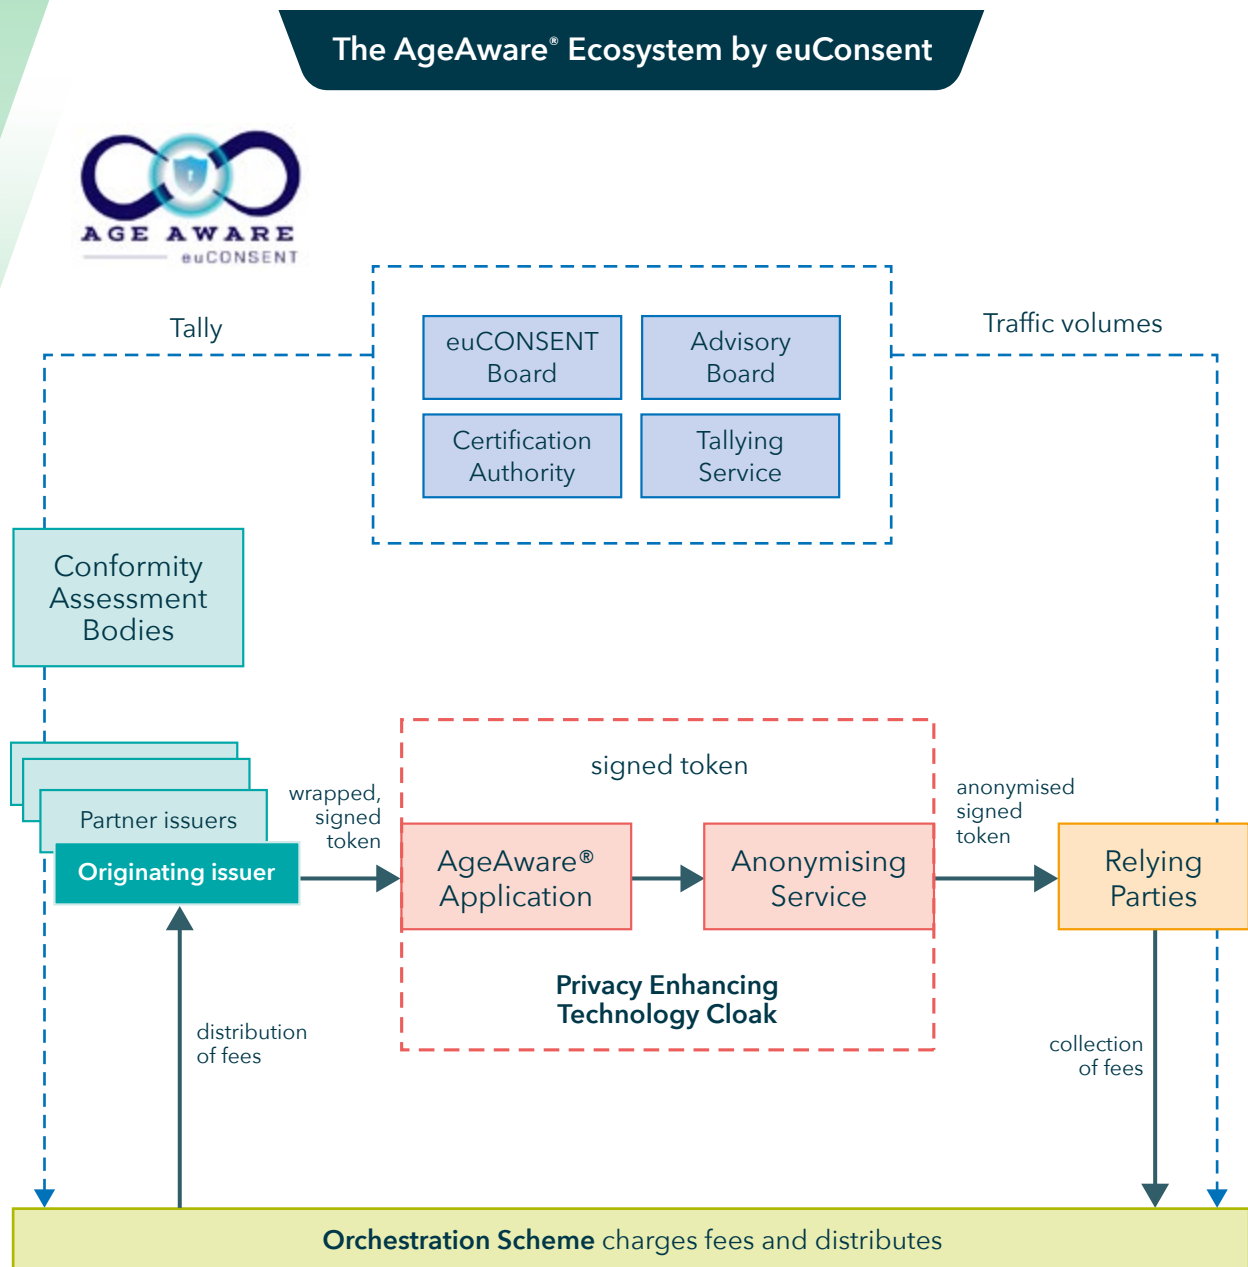


Figure J.10.1 The AgeAware Ecosystem by euConsent

| General identity protocol (GIP)

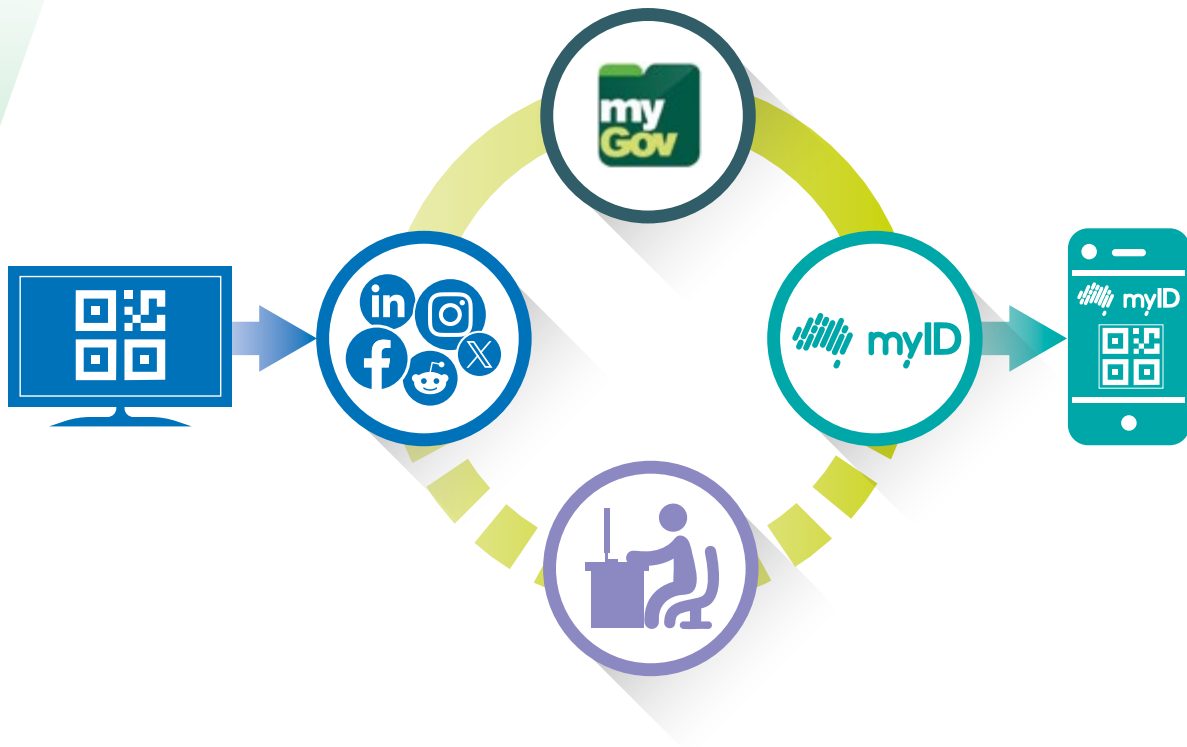
J.10.7 New Zealand-based startup GIP Ltd submitted its General Identity Protocol (GIP) to the Trial as a proposed architecture for networked identity and interoperable age assurance. The GIP model envisions a three-party ecosystem involving:

- Identity Providers (IdPs), who hold verified personal information.
- Credential Brokers, who route and verify attribute requests.
- Relying Parties, such as social media platforms or age-restricted services.

J.10.8 A key feature of the GIP model is the use of dynamic authentication codes and minimal client-side integration. The system supports multiple assurance pathways – including facial estimation, government eID or verified parental approval – and is designed to be scalable, cost-effective and privacy-preserving. It allows the user to present age claims without revealing identity and includes built-in audit and trust controls to prevent unauthorised or unverified claims.

J.10.9 GIP's strength lies in its architectural flexibility: it can operate via browser, app or call-centre scenarios and connect to national infrastructure such as myGov, bank records or third-party wallets. Its current status is conceptual with prototype demonstrations, but it presents a credible blueprint for future interoperable deployments.

Technical Demonstration Online



- 1 Age check instigated
- 2 SM site requests and R18 code
- 3 GIP sends R18 code
- 4 SM site displays code
- 5 Individual scans the code
- 6 Individual authenticates
- 7 myID passes code to GIP
- 8 GIP checks code matches
- 9 GIP accesses myGov for age check
- 10 GIP confirms success
- 11 SM site and myID show success

Figure J.10.2 GIP Technical Demonstration Online

| Snap and Meta: digital wallet-based interoperability

J.10.10 Submissions from Snap Inc. and Meta both outlined ecosystem-level approaches that aim to enable reusable age signals across apps. These models envision:

- A verified age credential stored at the OS, app store or digital wallet level.
- APIs that allow third-party apps to query and apply that age credential, gated by user or parental consent.
- Support for multiple age assurance inputs, including ID-based verification, facial age estimation or parent-supplied attributes

J.10.11 Snap's model is notable for its technology neutrality, supporting various forms of verification and a low-data footprint. Meta's proposal emphasises ease of use for parents, with a "golden moment" approach that captures the child's age at the time of device onboarding. Both envision a future in which apps trust the ecosystem – rather than perform age checks in isolation.

J.10.12 However, both models currently depend on cooperation from OS and app store providers, which may limit interoperability across vendors or jurisdictions if closed ecosystems are maintained.

| Opale.io and privacy-preserving interoperability

J.10.13 Another relevant submission came from Opale.io, which offers AgeKey, a reusable proof-of-age system built around privacy and double-anonymity principles. Users verify their age once with an accredited provider; the resulting AgeKey is stored locally and bound cryptographically to their device, enabling subsequent access without further data disclosure.

J.10.14 The model is designed to support interoperability across services, including adult content, gaming and social media platforms. It enables the reuse of age credentials while minimising tracking risk, supporting granular, revocable and consent-based sharing.

| Observations and opportunities

J.10.15 Although still early-stage, these approaches point to a future in which interoperable frameworks allow users to verify their age once and reuse that signal across the digital environment, reducing burden and improving consistency. However, current implementations differ in:

- Architectural model (federated, centralised, decentralised).
- Credential type (asserted age range, verified claim, age token).
- Data sharing protocol (APIs, QR codes, digital wallets).
- Trust management (issuer whitelisting, public key infrastructure, third-party audits).

J.10.16 This diversity offers a rich landscape for innovation, but also highlights the need for common standards, which could emerge from W3C, ISO or eIDAS-aligned frameworks. Without shared schemas and governance, the risk remains that interoperability will be partial or fragmented, particularly across borders or platform boundaries.

J.10.17 Participants also raised concerns about:

- Privacy and control: Interoperability must not mean increased traceability. Zero-knowledge and consent-based models remain important.
- Trust models: Questions remain about how relying parties evaluate and accept third-party credentials.
- Vendor lock-in: Without open APIs or credential standards, dominant platforms may consolidate control, limiting competition and innovation.

| Orchestration service providers

J.10.18 Orchestration service providers have the potential to play a pivotal role in enabling interoperable, scalable and user-friendly age assurance systems. These providers act as intermediaries that coordinate multiple verification methods, credential issuers or user interactions – allowing services to integrate with one platform while supporting a variety of assurance pathways. In doing so orchestration layers abstract the complexity of integrating with multiple assurance providers, reduce redundancy for users (who might otherwise need to re-verify across services) and create a flexible foundation for adopting future technologies. This makes them particularly valuable in environments where a single standard or verification method may not be appropriate across all use cases, jurisdictions or user contexts.

J.10.19 During the Trial, Opale.io emerged as a clear example of an orchestration-enabled architecture. Their system, AgeKey, is designed to accept credentials from multiple age verification providers, offering a consistent, reusable token to relying parties – while ensuring that no personally identifiable information is shared unnecessarily. Similarly, the General Identity Protocol (GIP) proposed by GII can be understood as an orchestration framework: it allows relying parties to request specific attributes (e.g., “Is this user over 16?”) and coordinates the routing of those requests to the appropriate identity or age credential providers. These models suggest that orchestration can support faster adoption by service providers, greater user privacy and broader ecosystem compatibility, especially if aligned with open standards and consent-based data sharing practices.

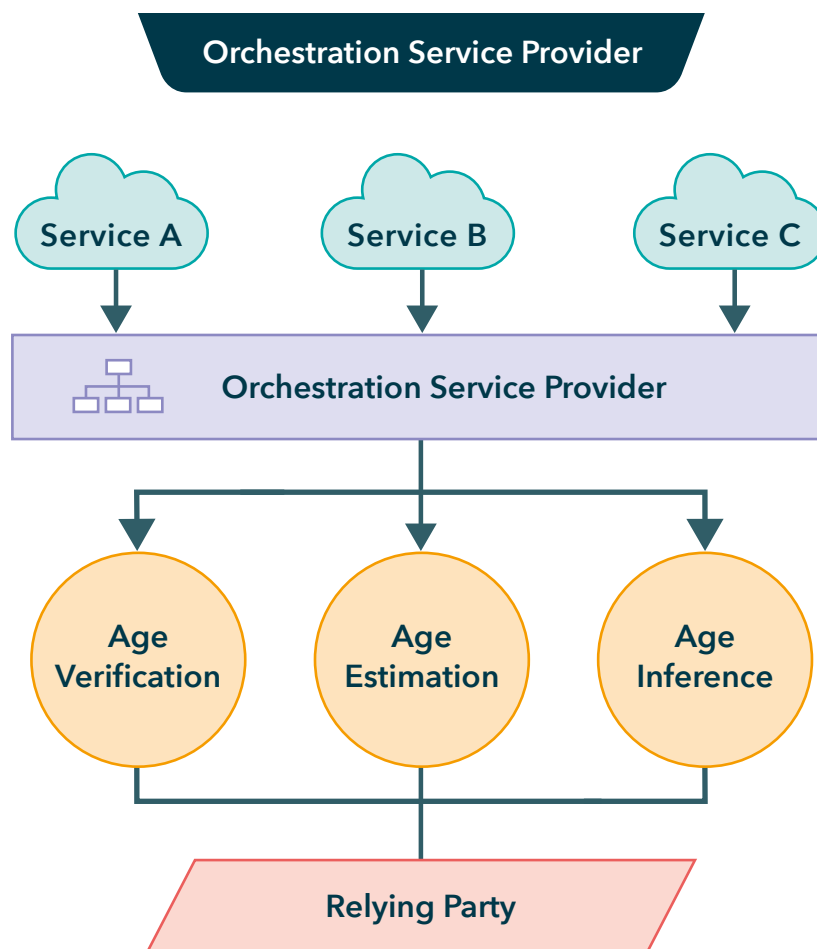


Figure J.10.3 *Orchestrated Service Provider*

| Summary

J.10.20 The Trial identified growing momentum toward interoperable age assurance systems, with multiple participants offering conceptually aligned – but technically distinct – approaches. Whether through federated networks, reusable tokens or system-level APIs, these models aim to balance user privacy, scalability and ease of integration.

J.10.21 While no dominant standard has yet emerged, the early work of participants such as GIP, Opale.io, Snap and others presents clear opportunities for technical convergence and regulatory harmonisation. The future development of interoperable systems will likely depend on trust frameworks, standards alignment and the ability to protect users' privacy, choice and agency across a decentralised ecosystem.

Vendor Case Study



Website

opale.io

Opale.io's AgeKey system proposes a reusable, on-device age credential that allows users to verify their age once with an accredited provider and then use a cryptographically bound token to access age-gated services. This architecture prioritises data minimisation and local trust enforcement, which aligns strongly with privacy-by-design principles.

Three Key Facts

1

AgeKey offers privacy-preserving, reusable proof-of-age tied securely to user devices.

2

AgeKey supports interoperability across platforms while minimising tracking through consent-based, revocable credential sharing.

3

Opale.io exemplifies orchestration-enabled architecture, integrating multiple providers without exposing personally identifiable information.

Practice Statement

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

Technology Trial Test Report

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

Privacy Policy

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

Technology Trial Interview

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

Summary of Results

AgeKey showed technical promise and privacy alignment but lacked broad testing, integration, and adoption, limiting its system-level readiness and maturity.

J.11 Technology Readiness Assessment

| Summary finding

J.11.1 Technology Readiness Levels (TRLs) vary widely, with many solutions overstating maturity. A significant number of Trial participants reported higher TRLs than could be substantiated. Some conceptual solutions were rated as TRL 3 or higher without evidence of analytical validation or testing. Most interoperable tech stack models remain at a conceptual or early prototyping stage.

| Detailed analysis

J.11.2 The Trial evaluated proposed solutions not only in terms of their technical function (e.g. age verification, consent flows), but also their ability to operate within or across the digital technology stack – including at the device, browser, network, app store, platform and backend levels. In this context, Technology Readiness Levels (TRLs) were used to assess the maturity of deployment models, including whether theoretical frameworks had been translated into demonstrable system integration, user pathways or technical feasibility assessments.

J.11.3 While TRL self-assessments provided by participants were helpful in understanding the perceived maturity of their solutions, the Trial revealed a pattern of overstatement – particularly in relation to interoperable, stack-based deployment models. Many proposals were rated at TRL 3 or above (indicating proof-of-concept or early validation), yet lacked concrete evidence of component integration, lab testing or real-world trials within the technology stack itself. In some cases, proposed implementations had not progressed beyond conceptual architecture diagrams or slideware and could not be demonstrated in operation.

Challenges in stack-level deployment

J.11.4 This discrepancy highlights several important dynamics:

- Age assurance methods themselves (e.g. facial estimation, ID-based verification) may have higher TRLs, but their deployment across stack layers – such as integration with operating systems, ISP networks, app stores or federated backends – often remains in the conceptual or early prototyping phase.
- Parental consent mechanisms, while sometimes implemented in apps or web environments, were rarely shown to operate systemically across devices or services. Proposals for embedding consent flows in device setup or network-level onboarding lacked demonstrable architecture or user testing.
- Cross-layer orchestration, particularly involving trusted signal handoff (e.g. from app store to app or device to browser), was described by some participants but not proven through lab environments or interoperability trials. Many solutions had not yet been tested with live services, under operational constraints or with representative user groups.

J.11.5 This illustrates that while participants were innovating in their respective domains, most had not yet bridged the gap between single-layer functionality and stack-wide implementation. TRL assessments at the system level were often inflated by the maturity of underlying verification technologies, without equivalent validation of their integration into the broader digital ecosystem.

Structural limitations contributing to lower TRLs

J.11.6 Several structural and technical factors appear to limit TRL advancement for tech stack-based models:

- **Dependency on platform operators:** Many concepts require cooperation from Apple, Google or ISPs to integrate age assurance mechanisms into app stores, operating systems or network gateways. Without platform-level support, even well-designed concepts remain hypothetical.
- **Lack of common standards:** Inconsistent APIs, trust frameworks and data formats hinder the ability of solutions to interoperate across layers, lowering the system-level readiness even when individual components are mature.
- **Absence of controlled integration environments:** Few participants had access to testbeds, sandboxes or pilot environments that would allow staged deployment and verification of stack-wide functionality under controlled conditions.
- **User and regulatory feedback loops:** Many systems had not been tested with real users – particularly parents, children or edge-case households – or reviewed against local legal frameworks. This limited the ability to validate operational effectiveness and compliance readiness.

Opportunities for Improvement

J.11.7 The Trial findings suggest that while conceptual models for stack-based deployment are developing, there is significant scope to improve system integration maturity through:

- Simulation and sandbox environments where multi-layer interoperability can be tested in realistic conditions.
- Collaborative engagement with ecosystem operators, including app stores, ISPs and OS vendors, to enable demonstrators or plug-in pathways.
- Shared development of open APIs and protocols that enable trustable age attributes to pass securely and verifiably between layers.
- Reference implementations of age assurance deployment within common user flows, such as device setup, app install or network subscription onboarding.

Vendor Case Study**R2LABS***Website*r2-labs.io

R2 Labs provides consent tokens tied to child identity using cryptographic proofs; supports revocation but not ongoing behavioural oversight or content filtering.

Practice Statementageassurance.com.au/v/r2l/#PS*Technology Trial Test Report*ageassurance.com.au/v/r2l/#TR*Privacy Policy*ageassurance.com.au/v/r2l/#PP*Technology Trial Interview*ageassurance.com.au/v/r2l/#VI**Summary of Results**

The system balances auditability (via immutable logging) with user rights (via revocable mechanisms), offering a compelling model for future adoption.

J.12 Critical Implementation Challenges

| Summary finding

J.12.1 Functionality, performance, privacy and acceptability present critical implementation challenges. Even theoretically promising models face practical threats to performance and adoption. Key concerns include latency, data handling practices, user transparency and public trust – particularly where technologies operate beyond the user’s immediate control.

| Detailed analysis

J.12.2 While many models for age assurance, parental consent and child protection within the technology stack appear promising in theory, the Trial revealed a range of critical challenges that threaten their performance and adoption in practice. These challenges extend beyond technical feasibility to encompass real-world constraints related to latency, data handling, user agency, social acceptability and trust. In many cases, systems that are architecturally sound or well-intentioned in design encounter barriers when implemented at scale, particularly when embedded at layers of the digital ecosystem that are less visible or controllable by users.

J.12.3 These issues are particularly acute for solutions implemented at foundational levels of the stack – such as at the network, browser or operating system layer – where controls may be applied without the user’s explicit awareness or real-time input. Even models that prioritise privacy or interoperability must contend with the human and social dimensions of technology deployment, which are often more decisive than technical sophistication in determining real-world outcomes.

| Functionality and performance limitations

J.12.4 Several Trial participants acknowledged or demonstrated performance concerns, particularly in relation to latency, accuracy and reliability:

- Latency can be introduced by remote age verification processes, especially where checks must occur in real time before a user can access content or services. This may be tolerable in isolated cases, but when embedded into system-level actions (e.g. during device setup, app download or page load), delays of even a few seconds can degrade user experience and increase dropout rates.
- Inconsistent behaviour across platforms or devices was also observed or anticipated. For example, solutions that work smoothly on Android may face restrictions on iOS or fail to function at all in privacy-enhanced browser environments. This inconsistency undermines reliability and increases integration burdens for developers.
- Some systems faced challenges in handling edge cases, such as shared devices, non-standard user journeys (e.g. kiosk or library access) or multilingual environments – limiting their practical coverage despite theoretical suitability.

| Privacy and data handling risks

J.12.5 Privacy is a central concern in all age assurance models, but especially in those operating beyond the user's direct control. The Trial identified several areas where data handling practices could present risks:

- **Opaque data flows:** Users (especially children and parents) often lack visibility into where their data goes, how long it is retained or who has access to it. This is particularly problematic in network- or device-level models, where data may be logged or inferred in ways that users cannot see or challenge.
- **Excessive data collection:** Some approaches proposed the use of government IDs, biometric images or persistent behavioural tracking - raising questions about proportionality and long-term privacy implications. Even where data minimisation was claimed, documentation of actual practices was inconsistent across participants.
- **Cross-service profiling risks:** Interoperable systems that share age attributes across services risk inadvertently creating digital breadcrumbs that could be re-identified or used for profiling. These risks increase when age signals are combined with device fingerprints, browsing behaviour or metadata.

| User acceptability and public trust

J.12.6 Even where systems perform well and protect data in principle, they may still face low adoption or resistance if perceived as intrusive, unfair or overly complex:

- Lack of user agency: Systems embedded deep in the stack – such as network filtering or automatic age gating at the app store level – can apply restrictions without offering meaningful user control or context. This can lead to frustration or workarounds, especially among teenagers who may feel over-restricted.
- Parental confusion or disengagement: Many models assume active and informed participation by parents, yet digital literacy, language barriers or socio-economic factors may prevent consistent engagement. Trial interviews highlighted that if parental controls are too complex or obscure, they are unlikely to be configured properly.
- Cultural and community mistrust: In some communities, particularly where governments or tech companies are not widely trusted, identity or age verification processes can be seen as surveillance or social control. This perception can reduce uptake, even of well-designed solutions.



| System-level trade-offs

J.12.7 A recurring theme in the Trial was the tension between centralisation and control. Models that operate high in the stack (e.g. app-store-based systems) may offer broad coverage and interoperability but reduce transparency and decentralisation. Conversely, user-managed or app-specific solutions offer more granularity but risk inconsistency or underuse.

J.12.8 Participants also noted challenges in balancing security with usability. For instance, systems that require frequent re-verification may be secure, but also cumbersome. Others that rely on local device storage may be user-friendly, but vulnerable to circumvention or tampering if not well-protected.

| Summary

J.12.9 The Trial demonstrated that technical feasibility is only one dimension of successful age assurance deployment. Real-world implementation of these models – especially those embedded across the tech stack – must contend with performance bottlenecks, data governance risks and user acceptance barriers that can undermine even the most sophisticated designs.

J.12.10 These findings highlight the importance of designing age assurance systems that are not only interoperable and standards-aligned, but also responsive to the social, behavioural and experiential realities of the users they aim to protect. Without addressing these broader dimensions, even the most promising models may fall short in practice.

J.13 Responsibility and Liability

| Summary finding

J.13.1 Responsibility and liability in a distributed tech stack are unclear and require further definition. Where age assurance functions are spread across multiple technical layers and actors, accountability becomes diffuse. Without clear regulatory or contractual frameworks, there is a risk of ambiguity in liability, weakening enforcement and redress mechanisms.

| Detailed analysis

J.13.2 As age assurance technologies evolve and are increasingly embedded across the digital technology stack, the question of who is responsible when things go wrong becomes more complex. In traditional regulatory models, liability for unlawful access by children to age-restricted goods or services lies clearly with the provider. But in a multi-layered digital ecosystem – spanning devices, networks, platforms, verification providers and parental controls – accountability becomes fragmented. This section explores the legal, contractual and practical implications of distributing age assurance functions across the tech stack and examines the challenges this creates for enforcement, redress and risk management.

| Legal foundations of liability in age-restricted access in Australia

J.13.3 In Australian law, responsibility for preventing minors from accessing age-restricted goods, content or services is typically assigned to the entity offering the product or service, particularly where access is prohibited by legislation. This applies across various domains:

- Alcohol, tobacco and gambling are governed by state and territory legislation, with clear statutory offences for supplying to minors.
- Media content restricted to adults is regulated under the Classification (Publications, Films and Computer Games) Act 1995.
- The Online Safety Amendment (Social Media Minimum Age) Act 2024 puts a legal obligation on 'age-restricted social media platforms' to take reasonable steps to prevent Australians under 16 years old from having an account on their service.

J.13.4 The general principle is that the provider of the regulated service bears responsibility for compliance. However, Australian consumer and common law also recognise circumstances where liability can be shared, outsourced or disclaimed, subject to tests of reasonableness, foreseeability and statutory consistency.



Views from Age Verification Providers Association on Liability:



The AVPA emphasises that, under most digital age assurance frameworks, final responsibility for preventing underage access remains with service providers – not platform operators or verification providers. In their commentary on EU digital wallets, AVPA notes that while governments and major platforms may facilitate age verification infrastructure, citizens and services must ultimately rely on trusted age assurance providers to achieve compliance. They also caution that over-reliance on platform-controlled systems (such as app stores or wallets) risks creating a “single-point failure” where liability becomes concentrated in the hands of a few large corporations, potentially limiting choice and increasing exposure for downstream services.

On shared liability, AVPA advocates for privacy-preserving, standards-based interoperability, including certification schemes and open protocols – like those being developed under euCONSENT and IEEE 2089.1. They express concern that liability should not be passed through exclusionary clauses or opaque delegation models, as these leave gaps in enforcement and do not align with the principle that the provider of age-restricted content must “stop underage access” – regardless of the stack layer. AVPA’s stance underscores the need for clear accountability frameworks that balance technical innovation with regulatory integrity.

| Diffuse accountability in a distributed tech stack

J.13.5 When age assurance is implemented via the technology stack – spanning multiple layers including devices, networks, app stores, identity brokers and third-party verifiers – responsibility is fragmented. No single party may have full visibility or control over the entire assurance pathway, raising significant legal and practical challenges:

- Service Providers (e.g. apps, websites) may claim that age assurance was delegated upstream to the app store or platform.
- App Stores or OS Operators may argue they provide infrastructure only, not content or service access decisions.
- Verification Providers (e.g. ID or biometric services) may disclaim responsibility for misuse or misapplication of their signals by relying parties.
- Parents may be positioned as gatekeepers under child account models but lack enforceable obligations.

J.13.6 This creates a liability vacuum in which enforcement agencies and courts must disentangle a complex web of contracts, configurations and system logs to determine who failed to act and whether the failure was material to the breach.

| Agency, principal and the role of contractual instruments

J.13.7 From a legal theory perspective, the complexity of distributed systems can be analysed through the lens of principal-agent relationships and contractual delegation:

- If a platform (e.g. an app store) engages an age verification provider to act on its behalf, it may be held liable for the agent's failure under the doctrine of vicarious liability, unless clear contractual limitations apply.
- Conversely, if a service contracts out its age assurance responsibilities to an upstream layer (e.g. via a platform API) and that upstream entity only offers age signals without warranties, the relying party may retain residual responsibility if the assurance fails.

J.13.8 In practice, these relationships are governed by terms of service, licence agreements, APIs terms and integration contracts – many of which include disclaimers, waivers, exclusions of liability and limitation clauses. For example:

- Verification providers may include clauses that state their age estimates are indicative only and not suitable for legal or compliance purposes.
- Platforms may offer services “as is,” with no guarantee that the age signal received reflects a verified or accurate user age.
- App developers may disclaim responsibility if relying on platform defaults or parental controls, especially if not modified by the user.

J.13.9 This contractual patchwork makes regulatory enforcement and legal redress difficult, as it may be unclear who assumed which risk and which assurances were promised to whom. The lack of a shared duty of care or statutory allocation of responsibility across the stack further complicates matters, especially where harm to a child occurs due to a failure in one part of the system.

| Implications for enforcement, redress and risk allocation

J.13.10 This distributed responsibility model creates substantial challenges for regulators and courts:

- Evidence chains must trace user access through multiple technical layers to determine where the failure occurred.
- Jurisdictional issues may arise, especially where different actors (e.g. an app in one jurisdiction, a verification provider in a second one and a child in a third jurisdiction) are involved.
- Inconsistent regulatory coverage across different actors within a supply chain of infrastructure, industry and consumers makes unified enforcement difficult.
- Without clear statutory guidance, disputes may be determined on a case-by-case basis, producing inconsistent outcomes and low legal certainty.

J.13.11 There is also the risk that liability will be “orphaned” – with each actor in the chain pointing to another and none accepting responsibility – undermining the deterrent effect of legal frameworks and frustrating the ability of affected families to seek redress.

| Insurance and mitigation strategies

J.13.12 Given this complexity, some participants and market observers are exploring how commercial risk mitigation mechanisms – particularly insurance, risk pooling and contractual indemnities – could support adoption of age assurance systems within the stack.

- Cyber liability insurance may cover certain failures (e.g. unauthorised access or data leakage) but rarely extends to failures of regulatory compliance in content gating or service restriction.
- Professional indemnity policies for verification providers or orchestration services could be adapted to cover false negatives or system failures, though risk pricing remains immature in this space.
- Mutual assurance frameworks (e.g. age assurance consortiums or trust frameworks) could provide collective liability cover for services using approved providers, subject to technical and governance standards.

J.13.13 Other potential mitigations include:

- Adoption of indemnification clauses in service contracts, where one actor (e.g. the verifier) agrees to compensate another (e.g. the app) in case of failure.
- Use of conformance frameworks or certification schemes (e.g. ISO/IEC FDIS 27566-1) to establish industry best practices that limit liability where followed.
- Development of third-party audit regimes or neutral attestation services, which may support regulatory confidence and reduce the burden of individual enforcement.

J.14 Proximity to Risk

| Summary finding

J.14.1 Proximity to risk is an important factor in assessing effectiveness. The location of the age assurance mechanism within the stack affects its ability to respond to harmful content or risky interactions. Solutions closer to the user or service (e.g., device-level or in-app) may offer more accurate contextual control but may also have narrower coverage.

| Detailed analysis

J.14.2 “Proximity to risk” refers to the distance – both technical and contextual – between the location where age assurance is applied and the place where a child is exposed to potential harm. In the context of a distributed technology stack, this principle recognises that the effectiveness of an age assurance mechanism is shaped not only by how accurate it is, but also by how close it is to the risk event itself.

J.14.3 Risk in a digital environment can arise at various points – during account registration, while consuming content, when interacting with strangers or when using specific device features (e.g., location sharing, live chat or in-app purchases). The closer the assurance mechanism is to the point where harm might occur, the more precisely it can respond to real-time user behaviour, apply context-specific controls or involve relevant safeguards (such as parent mediation or content warnings).

J.14.4 By contrast, when assurance is applied further away from the risk – such as during device setup, at the app store level or through network-based filtering – the system may be able to enforce broad access restrictions but lacks the nuance to address what the child is doing within a permitted environment or how they might encounter emerging harms.

| Risks of distant assurance

J.14.5 Conducting age assurance too far from the risk introduces a number of problems:

- **Blunt enforcement:** Mechanisms applied at the network or ISP level may lack the ability to distinguish between harmful and beneficial content, leading to overblocking of material that children have a right to access (e.g., sexual health information, youth mental health resources or trusted news services).
- **Static classification:** App store-based models may gate access at the time of download, but do not adapt to dynamic risks, such as changes in an app's features, in-app exposure to inappropriate content or peer-to-peer interactions within the service.
- **Collateral intrusion:** If stack-wide controls are applied without sufficient granularity, they may restrict older children from accessing content appropriate for their age group or inadvertently apply controls to adults using the same device or network. This undermines both utility and trust.
- **Limited real-time responsiveness:** Age assurance mechanisms operating far from the user or application context (e.g. in backend systems or external platforms) are less able to detect harmful user behaviour, grooming attempts or rapid shifts in content exposure. They may also fail to activate when children access services through anonymised browsers, VPNs or shared devices.

| Benefits of close proximity to risk

J.14.6 In contrast, mechanisms positioned close to the point of risk – such as within an app, on a child’s device or at the interface layer – can offer more precise and responsive protections:

- **Context-aware enforcement:** Device- or app-level systems can adjust controls based on usage patterns, time of day, content types or behavioural signals. For example, a messaging app might restrict contact with unknown users if the account is flagged as underage.
- **Granular user experiences:** In-service age checks enable differentiation between content types and features, allowing access to general information while still blocking harmful or exploitative material. This reduces the risk of “all-or-nothing” content blocking.
- **Adaptive safety measures:** Systems embedded in the application or interface can provide real-time warnings, pausing or escalation pathways – for instance, flagging unsafe interactions, prompting reflection or activating parental oversight.

J.14.7 However, proximity also comes with trade-offs: closer systems often require deeper integration, user trust and developer commitment. They may also struggle with coverage, as not all services implement age-specific controls or offer parity across platforms.

| Children's rights and proportionality

J.14.8 The principle of proximity to risk aligns closely with the UN Convention on the Rights of the Child (UNCRC) and its 2021 General Comment No. 25 on children's rights in the digital environment. In particular:

- Article 17 (Access to information): Children have a right to access age-appropriate information from diverse sources, including digital media.
- Article 13 (Freedom of expression): Age assurance should not unduly limit a child's ability to seek, receive and impart information.
- Article 16 (Right to privacy): Measures taken to protect children must not involve disproportionate surveillance or profiling.
- Article 3 (Best interests of the child): Any system must balance protection with developmental needs, agency and participation.

J.14.9 Applying age assurance at a distance from the risk may violate these principles if it bluntly limits access, applies static assumptions about age and risk or impedes developmental rights without sufficient justification. Proximity to risk ensures that protections are tailored, proportionate and justifiable within the specific context of harm.

J.15 Virtual Private Networks and Geolocation

| Summary Finding

J.15.1 Geolocation services can play a role in detecting and preventing circumvention via VPNs. Some participants highlighted the potential for geolocation software to support age assurance systems by identifying when users attempt to mask their true location – such as through Virtual Private Networks (VPNs) – to bypass regional age restrictions. While promising, this approach raises its own challenges related to accuracy, evasion tactics and the implications for user privacy and cross-border service access.

| What are Virtual Private Networks (VPN's)

J.15.2 A Virtual Private Network (VPN) is a technology that allows users to create a secure, encrypted connection between their device and a remote server on the internet. From the perspective of the websites or services being accessed, the user's internet traffic appears to originate from the location of the VPN server, not from their actual device or physical location.

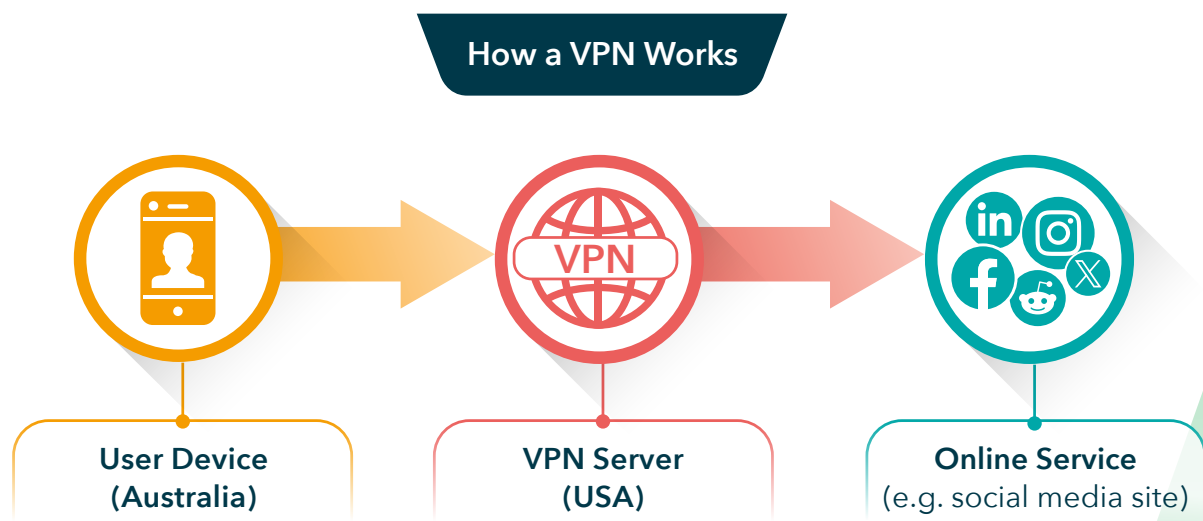


Figure J.15.1 How a VPN Works

J.15.3 VPNs are commonly used for privacy, security or to bypass geographic restrictions. For example, a user in Australia might use a VPN to appear as though they are browsing from the United States. While this can be useful for legitimate purposes (e.g. protecting data on public Wi-Fi), it can also be used to evade regional rules or content restrictions – including age assurance or parental controls that depend on a user’s real location.

J.15.4 In the context of age assurance, VPNs present a challenge: even if a child’s access is restricted based on local regulations or age thresholds, VPN use may allow them to route around those restrictions by appearing to access the internet from a jurisdiction with looser rules or different enforcement mechanisms.

| What is Geolocation Software?

J.15.5 Geolocation software determines the real-world location of a user’s device based on a combination of digital signals. It is commonly used for a wide range of purposes – such as ensuring legal compliance (e.g., in gambling or video streaming), offering location-specific services or detecting attempts to circumvent digital restrictions via VPNs or spoofing tools.

J.15.6 To determine a user’s location, geolocation systems typically aggregate multiple data points, including:

- IP address – identifies approximate geographic region.
- GPS signals – provide precise location when enabled.
- Wi-Fi network information – triangulates position based on nearby access points.
- Cell tower connections – gives coarse but useful data in mobile environments.
- Device and browser fingerprinting – detects consistency and possible manipulation.

J.15.7 In the context of age assurance, geolocation software can help determine whether a user is physically present in a jurisdiction where access to certain content or services is regulated – and whether they are attempting to bypass those controls via a VPN or spoofed identity.

How Geolocation Works to Validate User Location

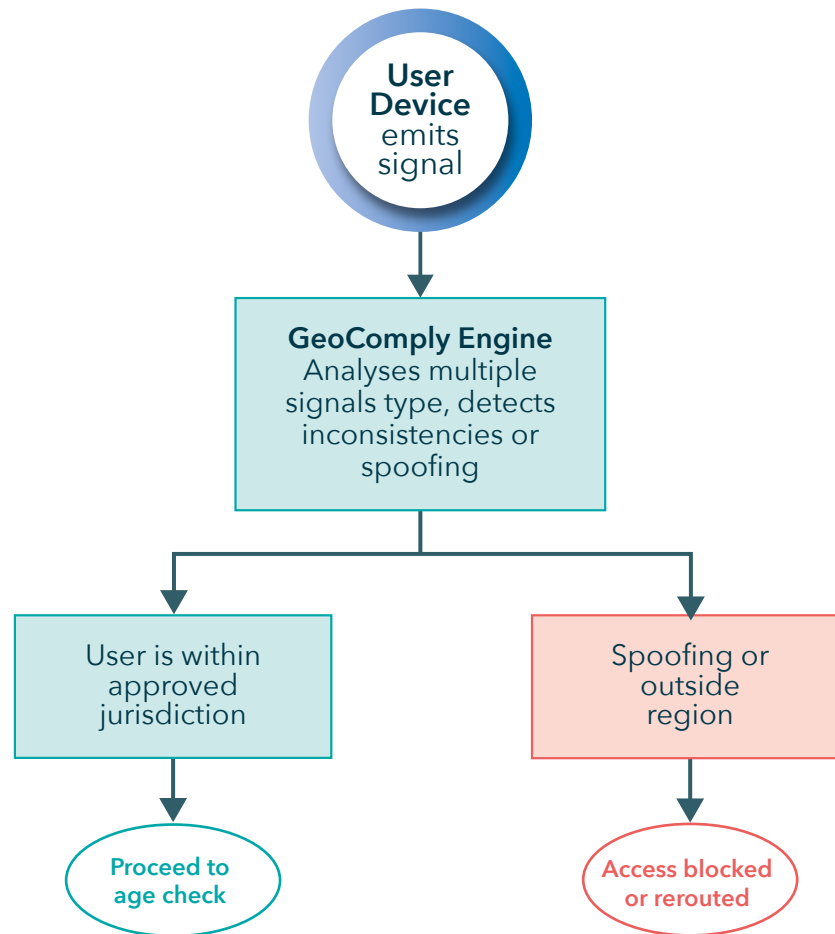


Figure J.15.2 How Geolocation Works to Validate User Location

Vendor Case Study

GeoComply[®]

Website

geocomply.com

GeoComply provides precise location and anti-VPN tools that help prevent circumvention of regional age restrictions and verify access legitimacy based on jurisdiction.

Three Key Facts

1

Supports high-accuracy applications such as online gambling and crypto compliance and is now expanding into social media and child safety contexts.

2

Processes anonymised data only, avoiding the collection of names, birthdates or biometric identifiers.

3

Acts as a gatekeeper - ensuring users are in a permitted location before they are routed to an age verification flow.

Practice Statement

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

Technology Trial Test Report

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

Privacy Policy

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

Technology Trial Interview

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

Summary of Results

GeoComply's approach has its roots in the highly regulated U.S. gambling sector, where strict state-level boundaries demand precise enforcement. Its application in age assurance allows platforms to enforce territory-based restrictions - such as limiting age-restricted services to jurisdictions with verified legal frameworks or obligations.

| Using geolocation to tackle the VPN fallacy

J.15.8 One of the more persistent myths in digital regulation – often referred to as the VPN fallacy – suggests that compliance with age-based access restrictions can be undermined or excused simply because a user might be using a Virtual Private Network (VPN). In reality, the law does not say “you must deny access to minors unless they use a VPN” – it says that regulated entities must prevent unlawful access, full stop. The use of a VPN by a minor does not exempt platforms or providers from responsibility under Australian or international law.

J.15.9 In this context, geolocation services offer a practical and established method to detect and respond to VPN usage, ensuring that digital compliance systems are not circumvented by relatively unsophisticated methods. This is not a new challenge: VPN detection and geolocation enforcement are well-established in other regulated sectors, particularly online gambling, financial services and content licensing.



| Global use cases of geolocation to counter VPN evasion

J.15.10 In Australia, geolocation tools are increasingly used to enforce content access restrictions (e.g., on licensed streaming platforms) and are emerging as part of age assurance infrastructure, including in proposals considered by the eSafety Commissioner.

J.15.11 In the United States, companies like GeoComply are used extensively by state gaming commissions to ensure online gambling only occurs within legally authorised state boundaries. VPN or location spoofing is detected and blocked in real-time, ensuring compliance with state-specific laws – e.g. gambling in New Jersey cannot be accessed from New York via a VPN.

J.15.12 In cross-border law enforcement, such as the sale of liquor or firearms across U.S. state lines, digital sales systems routinely use geolocation tools to identify prohibited interstate transactions. Even where users attempt to mask their location, regulatory enforcement uses IP traceback, timestamp logs and spoofing detection to confirm real location and intent.

| Why VPNs are detectable

J.15.13 While VPNs are designed to obscure a user's IP address, they are not invisible. Leading geolocation software detects VPN usage using a variety of indicators:

- IP range checks: Most commercial VPN providers use well-known IP address blocks, which are listed and updated regularly.
- DNS leakage and routing anomalies: VPNs often introduce detectable inconsistencies in how DNS requests or traffic routing behaves.
- Latency and traffic fingerprinting: VPN traffic often exhibits patterns (e.g. higher latency, constant tunnelling) that are recognisable.
- Multiple device or user anomalies: Repeated use of the same VPN IP across thousands of users in short intervals is a strong signal.

J.15.14 These indicators allow geolocation systems to flag suspicious access attempts with high accuracy, prompting services to take appropriate steps.

| VPNs are not inherently harmful - but require guardrails

J.15.15 It is important to acknowledge that VPNs are not inherently malicious. They serve important roles in digital privacy, cybersecurity, remote work and protection against censorship. However, in regulated environments – such as age-restricted services – they must be treated as risk signals.

J.15.16 The presence of a VPN should not automatically block access but can trigger additional safeguards, such as:

- Secondary verification prompts (e.g. requiring age proof via trusted method)
- Suspension of anonymous browsing (e.g. limiting session until a valid location is confirmed)
- Parent or administrator review (e.g. for underage users accessing via child accounts)
- Logging and audit flags for compliance reporting

J.15.17 This layered response allows providers to respect legitimate use of privacy tools while maintaining compliance and protecting children from unlawful access.

| Legal and regulatory response

J.15.18 Foxtel and Village Roadshow were among the first to push for legislative support against piracy and unauthorised streaming. Their advocacy contributed to the Copyright Amendment (Online Infringement) Act 2015, which allows Australian courts to issue site-blocking orders against offshore streaming services that violate copyright.

J.15.19 When users accessed these services via VPNs to circumvent geoblocking, broadcasters used court orders to require ISPs to block entire domains or IP ranges, even where VPNs were involved. These rulings reinforced the principle that use of a VPN does not override legal access restrictions – paralleling the regulatory position now emerging in age assurance.

| Use of geolocation and VPN detection technology

J.15.20 Streaming platforms such as Stan and Kayo Sports have implemented real-time VPN detection using commercial geolocation tools. These systems check the origin of IP addresses, DNS requests and device fingerprints to identify when a user is masking their location. If a VPN is detected:

- Access to the service may be temporarily suspended.
- Users are prompted to disable the VPN and confirm their location.
- In some cases, accounts may be flagged for review or cancellation.

J.15.21 This mirrors enforcement models used internationally by Netflix and BBC iPlayer, which also rely on blacklists of known VPN IP ranges and behavioural anomaly detection to limit circumvention.

| Education and subscription incentives

J.15.22 To reduce VPN-based circumvention and piracy, broadcasters also invested in making their content more easily accessible and competitively priced. By offering low-friction subscriptions, bundled offers (e.g. with mobile plans) and exclusive local content, they have reduced the incentive for users to seek VPN workarounds.

J.15.23 This multi-pronged approach – legal, technical and market-based – has allowed Australian broadcasters to significantly reduce unauthorised VPN use while still respecting legitimate user privacy in cases such as travel or workplace access.



| Using private browsing to evade age restrictions

J.15.24 VPNs and incognito mode are not the same and incognito mode does not provide an effective way to bypass age restrictions – especially when those controls are properly designed and implemented across the technology stack.

| What's the difference?

J.15.25 A VPN (Virtual Private Network) reroutes a user's internet connection through a remote server, masking their IP address and making it appear as though they are accessing the internet from a different location. This can be used to circumvent regional restrictions, including age-based access in some systems.

J.15.26 Incognito mode (or private browsing), by contrast, is a local browser feature. It prevents the browser from storing:

- Browsing history,
- Cookies,
- Site data,
- Form inputs.

J.15.27 However, incognito mode does not hide the user's IP address, device identity, location or login credentials. It does not prevent services from recognising the user, especially when logged into an account (e.g., YouTube, Google, TikTok) or when device-level or network-based age controls are active.

| Can kids use incognito mode to bypass age restrictions?

J.15.28 In some limited, poorly designed systems, incognito mode might temporarily reset or hide stored cookies used to determine age – but this is generally not a reliable or scalable method for circumvention.

J.15.29 Modern age assurance systems rely on persistent signals, such as:

- Device settings,
- Logged-in identity,
- Network-based filtering,
- App-level age gating,
- Parental controls,
- External age verification.

J.15.30 These do not rely solely on cookies or browser history and are unaffected by private browsing. Furthermore, many parental control systems and content filters are designed to block access to incognito mode altogether or provide parents with alerts when it is used.

J.15.31 While incognito mode might temporarily obscure browsing activity from a parent on a shared device, it is not a meaningful tool for bypassing well-implemented age assurance controls. It lacks the functionality of a VPN and provides no meaningful protection against location-based or account-based restrictions. In short, incognito mode is limited in scope and easily accounted for in system design.

| Relevance to age assurance

J.15.32 The idea that children will simply use VPNs to evade age assurance controls is a common misconception – but it does not reflect the practical or technical realities of sustained VPN use. While it is true that some digitally literate teens may attempt to bypass restrictions using VPNs, doing so consistently across multiple apps, services and devices is complex, inconvenient and often short-lived. VPNs can interfere with app functionality, content delivery, payment systems and regional settings – making them frustrating to maintain on a daily basis. Many commercial platforms now actively detect and restrict access when VPNs are in use and geolocation services can identify spoofing attempts in real time. Moreover, younger children typically lack the awareness, motivation or technical ability to use VPNs at all. The practical barrier, combined with detection technologies and enforcement policies, makes VPN-based evasion a limited and manageable risk – not a reason to avoid implementing meaningful age assurance.

J.15.33 There is proven precedent in the Australian digital environment: VPN usage is both detectable and manageable and does not present an insurmountable barrier to enforcement. As with content licensing, the use of geolocation and detection tools in age-restricted environments can help service providers meet their legal obligations – even when users attempt to evade them through VPNs.

| Summary

J.15.34 The idea that VPN use voids compliance obligations – the VPN fallacy – is not supported by law or practice. Regulators and courts expect service providers to use available technologies to detect circumvention attempts and respond appropriately. Geolocation services are proven tools that allow providers to spot and manage VPN use without over-blocking or compromising user rights. Their adoption within age assurance systems helps close enforcement gaps, uphold user trust and demonstrate a mature, rights-conscious approach to digital compliance.

J.16 Emerging Challenges and Future Considerations for the Tech Stack

| Summary finding

J.16.1 Several emerging topics – while not the primary focus of the Trial – have significant relevance to the evolution of the technology stack in age assurance. These include edge computing, auditability, child-centred design, and the resilience of assurance systems in diverse user contexts. They highlight areas for future technical development and regulatory engagement to support a scalable, trustworthy, and inclusive digital safety infrastructure.



| Detailed analysis

J.16.2 The Trial surfaced a wide range of models for age assurance within and across the technology stack. While this report has focused on those approaches actively submitted, trialled or assessed during the evaluation, there are several emerging areas of interest that are relevant to the future development of interoperable, privacy-preserving and risk-responsive age assurance infrastructure. These include:

- **Edge-based processing:** Age assurance capabilities deployed at the network edge (e.g. within home routers, content delivery networks or local gateways) could offer privacy-preserving, low-latency enforcement options. These models may be especially useful in education, enterprise or shared-use environments and offer a middle ground between device-based and ISP-level solutions.
- **Real-time, AI-powered content moderation:** As part of the in-service stack, some platforms are increasingly using machine learning to classify and adapt content exposure dynamically. When paired with stack-based age signals, this could support context-aware protections that go beyond static age gating.
- **Auditability and system logging:** Distributed assurance models introduce a need for clear, privacy-preserving records of enforcement decisions. Future development of logging standards – showing when and how age signals were validated, and by whom – could support accountability, regulatory compliance and user redress.
- **Child- and guardian-centric design:** While technical components of age assurance have received substantial attention, there is a growing need to align interface and user experience (UX) design with accessibility and child rights principles. This includes clear parental control dashboards, intuitive consent flows, and prevention of manipulative interface patterns.



- **Post-quantum and future-proofing credentials:** As cryptographic standards evolve, there may be long-term opportunities to align age assurance credentials with emerging quantum-resistant technologies, particularly in the context of national digital ID ecosystems and interoperable wallets.
- **Support for multi-child households:** Shared devices used by children of different ages pose specific implementation challenges. Future stack solutions could include family-linked identity models, per-child permissions, and scalable parental consent management systems across devices and services.
- **Offline and low-connectivity resilience:** In remote, mobile, or economically disadvantaged environments, systems that rely on live credential checks or continuous connectivity may exclude or fail. Device-bound, cacheable, or asynchronous validation models can support equity and functional robustness.

J.16.3 Together, these topics suggest directions for technical innovation, policy alignment and cross-sector engagement, ensuring the tech stack continues to support a safe and inclusive digital environment for children – even as technologies, standards and risks evolve. These emerging challenges also offer potential areas for international cooperation, industry standardisation and future-focused regulatory guidance.



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



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 978-1-0681646-8-2

