



RESEARCH REPORT

Fixing the Age Assurance Conundrum

Privacy, Proof of Age, UX. Pick Two.

Table of Contents

03 Introduction

18 Assuring Age -
Tools and
Technologies

05 Why is Age
Assurance
Important

31 Challenges to
Age Assurance

10 Age Assurance
and the Regulatory
Landscape

37 Conclusion

Chief Editor:

Nick Holland, Director of Research
Jennifer Berry, President

Contributing Authors:

Will Charnley, Manager, Advisory Services
Travis Jarae, CEO
Yifan Li, Senior Associate, Advisory Services

Introduction

There is probably no issue considered more important to humanity than the protection of children. The desire to guard the most vulnerable members of society from harm is hardwired into us as a species.

Today, a third of internet users are children, and this population has a right to leverage technology for learning, communication and entertainment.¹ The internet is also an environment that is rich with content that is inappropriate for children and is increasingly offering goods and services that would be age restricted offline. At its nadir, the internet can be a conduit for the very worst of society to gain access to children.

Age assurance is rapidly becoming a defining issue for the future of the internet. How consumers access content and how service providers will monetize these consumers is at an inflection point. We are on the cusp of building ecosystems that will enable digital identity to become a significant enabler of high value digital services. However, for the internet to evolve further, it must develop mechanisms for ensuring that children are safe online.

One of the greatest challenges comes from the need to know definitively whether an individual is of age -- this is difficult to ascertain in person, yet alone remotely. Further compounding the issue, privacy is paramount; mechanisms must only know the age of the person, without capturing other forms of personally identifiable information (PII) in the process.

Defining Age Assurance

Age assurance is a collective term for the technologies and activities that can be conducted to assure to a reasonable degree the age or age range of an individual. This is primarily to ascertain whether the person is of a specific age that meets the requirements for access to age restricted services online, or to make purchases of age restricted products and services.

There are two subsets of age assurance:

- 01 Age verification** - proof of age using deterministic data derived from trusted and verified sources such as birth certificates, passports and driving licenses, via proxies such as credit cards, or via digitized tokens that confirm that an individual is over a certain age without divulging identity information. Biometrics also falls under age verification when matching an individual to an existing biometric template (1:1 matching).
- 02 Age estimation** - an approximation of age made via probabilistic data such as behavioral biometrics and capacity testing. Biometrics also fall under age estimation when used, for instance, to assess the age of an individual via their facial structure or the size of their fingerprints.

While there have been some cursory steps made by service providers to police themselves, regulatory entities are now stepping up to the challenge, which is likely to have far reaching consequences. New regulations for age assurance not only impact service providers that offer age restricted goods and content; the very fabric of how apps and platforms are designed must now be protective of children by default. This presents significant challenges to businesses, from eCommerce platforms that sell age-restricted goods like vape pens all the way to video streaming platforms, who may attract children with its content. It also provides a significant opportunity for age assurance solution providers.

It will be important that regulatory interventions are not overly heavy handed. The internet is a critical component for learning and development of children. With COVID-19 still very much impacting the ability for children to be learning in schools, remote access is vital for their education and communication needs.

If managed in a manner that is appropriate for the protection and education of children, while respecting the privacy rights of all internet users, age assurance will almost certainly be an important catalyst in the development of digital ecosystems of trust.



Why is Age Assurance Important?

There are numerous reasons why age assurance is becoming a critical topic for the identity management industry — an increasingly under-18 digital population has heightened pressure from regulatory and non regulatory entities on service providers to implement more stringent assurances that children are kept safe online. There are also macro trends pushing age assurance to the forefront — an increased desire for consumers (adult and children) to control where their personal data is shared and how it is monetized, together with increasing concerns that big tech is placing profits over protection when it comes to children online. Collectively, these forces are shaping the future of age assurance technologies and policies.

A significant segment of the digital population is under-18.



According to UNICEF, one in three internet users globally are children.

UK regulator Ofcom reports that 79% of 12-15 year olds state that they have had at least one harmful experience online in the past 12 months.²

With the COVID-19 pandemic still a significant mitigating factor in children returning to in-person education, the reliance on online access is ever more critical. Yet the internet was never designed with children in mind and has some significant risk areas defined as the “4Cs.”³

01 Content

Children can discover content that is potentially harmful. This can include violent or gory content, hateful or extremist content, and pornographic content that may be illegal or harmful.

02 Contact

The child can be contacted, or targeted, in a potentially harmful adult-initiated interaction.

03 Conduct

Children can discover content that is potentially harmful. This can include violent or gory content, hateful or extremist content, and pornographic content that may be illegal or harmful.

04 Contract

The child is party to or exploited by potentially harmful commercial interests (e.g., - gambling, age inappropriate marketing). There is also increasing commercial activity on the internet that allows the purchase and delivery of goods and services that are age restricted in the physical world, such as alcohol, tobacco, marijuana and firearms.

The internet is a reflection of the physical world, where there are abundant opportunities for both positive and negative experiences. The challenge is in ensuring that the same checks and balances applicable in the physical environment can be layered upon the digital environment for the protection of minors, while continuing to respect the privacy rights of all digital citizens.

Age assurance regulations are gaining momentum, and teeth.

The requirement to protect children on the internet has been a pervasive issue since its genesis. Having been designed without any real thought of the consequences of where it would evolve to today, the online world has been often compared to the “Wild West,” where any rules that are applied are hard to maintain and to police. Nonetheless, the protection of children online has been a significant motivator for internet regulation. In 2000, the Children’s Online Privacy Protection Act (COPPA) became effective, requiring that website operators seek verifiable consent from a parent or guardian for collection of data pertaining to marketing to under 13 year olds. This has been foundational in shaping regulation internationally around the digital age of consent.

More recent regulations such as the EU Audio Visual Media Services Directive (AVMSD) and the UK Age Appropriate Design Code (AADC) are already proving to be significant in reshaping the way that content is delivered to children across new and emerging media platforms. AVMSD was originally designed to moderate network broadcast TV, but in 2017 was updated to include online streaming media platforms, widely extending the reach and impact of the regulation. AADC is a 2021 amendment to an EU General Data Protection Regulation (GDPR) provision, which includes compliance requirements for not only services aimed at children, but those likely to be used by them. AADC non-compliance falls under the same regulatory framework as GDPR, meaning potential fines of up to EU20 million, or 4% of annual revenues, whichever is highest.

Non-regulatory entities are also clamping down

While regulators are tightening requirements for age assurance, non-governmental associations are also demonstrating a less tolerant approach to companies that are not putting the necessary age assurance controls in place. Financial institutions and payment networks have significant power to demonetize service providers that are not compliant with their terms of service.

In December 2020, an investigation by the New York Times accused adult content site Pornhub.com of harboring illegal content, including child pornography. Shortly thereafter, card networks Visa, Mastercard, and Discover imposed a block on all credit card transactions being made on the site (American Express does not allow any card transactions on adult content sites).

In April 2021, Mastercard announced that it was extending existing “Specialty Merchant Registration” requirements, meaning that banks that connect merchants to the network would need to certify that the seller of adult content has effective controls in place to monitor, block and, where necessary, take down all illegal content.⁴ Other requirements included:

- Documented age and identity verification for all people depicted and those uploading the content
- A content review process prior to publication
- A complaint resolution process that addresses illegal or non consensual content within seven business days
- An appeals process allowing for any person depicted to request their content be removed

Compliance for these requirements went into effect on October, 16 2021.⁵

It's not just card networks that have forced increased stringency on age-restricted content. In August 2021, the popular user generated content site OnlyFans.com abruptly announced that it would no longer be allowing adult content on its platform. The reason, according to CEO Tim Stokely — “banks.” Stokely named three major banks that refused service because of “reputational risk” associated with the UK-based OnlyFans’ sexual material — BNY Mellon, Metro Bank, and JPMorgan Chase. BNY Mellon specifically had “flagged and rejected” every wire transaction involving OnlyFans, threatening its ability to pay creators.

Reasons for payment networks to be this stringent are numerous — the volume of chargebacks from adult content sites is significantly higher than other forms of online content, represented by payment networks charging transaction fees 10x or more to cover this risk.⁶ A greater risk may be that these networks are also culpable if users are paying for access to

content that is of children. Unless providers can provide reasonable assurances that content is of adults, and for adults, they can find themselves cut off from revenue generation and without a viable business model.


Digitally astute internet citizens want greater controls

While demographic and cultural shifts are pushing more children online for longer periods of time and regulatory forces are pushing for greater protection of these constituents, other broader trends are also bringing age assurance to the forefront.

Consumers (both adults and children) are becoming increasingly conscious of the value of their data and privacy, and are likely to enact more controls over where their data is shared and with whom if presented with the opportunity to do so. One metric that is indicative of this trend is the percentage of consumers that have opted not to be tracked across third-party apps in recent versions of Apple's iOS. In a July 2021 survey conducted by Liminal, over three quarters of consumers opted not to be tracked when given the option. Similarly, a quarter of consumers in the same survey stated that they actively avoided using fingerprint or facial biometrics on smartphones due to privacy concerns. As further evidence of consumers wanting greater controls over their digital destinies, over three quarters of consumers were interested or very interested in having the ability to control which companies get to see and share their digital identity, and to revoke access at any time.

Consumers demonstrate interest in having more autonomy over their digital identity

Question: How interested are you in having the ability to control which companies get to see and share your digital identity, including the ability to revoke access at any time?



Source: Liminal Consumer Identity Survey

This desire to exert control may be symptomatic of the current state of dissatisfaction with identity management. In a Liminal interview with Dr. Margaret Cunningham, a human factors engineer and behavioral psychologist, she describes the consumer rejection of tracking and biometrics as indicative of the broader frustrations of internet users.

“It’s exercising that last bit of control that we feel we have over our data, over who we are and how we’re represented in a space that we don’t understand,” says Cunningham. “I don’t know who owns all of my data points and I think that it’s something that we can say no to still, and that’s very powerful for people.”



Age Assurance and the Regulatory Landscape

Age assurance is becoming a lightning rod for the evolution of the internet today, with recent regulation telegraphing that companies doing business on the internet will need to design their platforms to ensure that children are protected, or face significant consequences. Regulations have been most influential in shaping the requirements of today's service providers to enact controls to ensure that children accessing their products or services are identified and protected. These trends also highlight the challenges of regulating age assurance while still providing for the privacy rights of the individual.

UN Convention on the Rights of the Child (CRC)

The 1989 UN Convention on the Rights of the Child (CRC) can be considered an overarching template for how children should be protected, both online and offline.⁷ While significantly pre-dating the current internet, it stands today as a framework for how children's rights regulation is enacted, including defining the age of a child, parental guidance requirements, respect for freedom of expression, education and information, rights to privacy, and protection from negligence and violence. These rights form the backbone of much of the age assurance regulation that has developed since.

The Children's Online Privacy Protection Act (COPPA)

In 2000, the Children's Online Privacy Protection Act (COPPA) became effective. The act applies to the online collection of personal information by persons or entities under U.S. jurisdiction about children under 13 years of age including children outside the U.S., if the company is U.S.-based. It details what a website operator must include in a privacy policy, when and how to seek verifiable consent from a parent or guardian, and what responsibilities an operator has to protect children's privacy and safety online including restrictions on the marketing of those under 13.

In December 2012, the Federal Trade Commission (FTC) issued revisions which created additional parental notice and consent requirements, amended definitions, and added other obligations for organizations that operate a website or online service that is "directed to children" under 13 and that collects "personal information" from users, or knowingly collects personal information from people under 13 through a website or online service.

After July 1, 2013, COPPA required operators to:

- Post a clear and comprehensive online privacy policy describing their information practices for personal information collected online from persons under age 13,
- Make reasonable efforts (taking into account available technology) to provide direct notice to parents of the operator's practices with regard to the collection, use, or disclosure of personal information from persons under 13, including notice of any material change to such practices to which the parents have previously consented,
- Obtain verifiable parental consent, with limited exceptions, prior to any collection, use, and/or disclosure of personal information from persons under age 13,
- Provide a reasonable means for a parent to review the personal information collected from their child and to refuse to permit its further use or maintenance,
- Establish and maintain reasonable procedures to protect the confidentiality, security, and integrity of the personal information collected from children under age 13, including by taking reasonable steps to disclose/release such personal information only to parties capable of maintaining its confidentiality and security,
- Retain personal information collected online from a child for only as long as is necessary to fulfill the purpose for which it was collected and delete the information using reasonable measures to protect against its unauthorized access or use; and
- Prohibit operators from conditioning a child's participation in an online activity on the child providing more information than is reasonably necessary to participate in that activity.

The law has been foundational in defining the way that commercial operators have been able to interact with children digitally. Although children under 13 can legally give out personal information with their parents' permission, many websites disallow children under 13 from using their services altogether due to the cost and work involved in complying with the law.

COPPA non-compliance has been enforced on numerous occasions. In February 2019, the FTC issued a fine of \$5.7 million to ByteDance for failing to comply with COPPA with their TikTok app.⁸ The infraction — collecting personal information from children without their consent. According to the FTC, TikTok (then named Musical.ly) met COPPA's definition of a site “directed to children”, based on audience composition statistics, as well as additional factors like subject matter, visual content, music, and the presence of child celebrities or celebrities who appeal to children. The complaint also alleged that Musical.ly had actual knowledge the company was collecting personal information from children; users' profiles revealed that many of them gave their date of birth or grade in school. And since at least 2014, Musical.ly had received thousands of complaints from parents of kids under 13 who were registered users. ByteDance agreed to pay the largest COPPA fine since the bill's enactment and to add a kids-only mode to TikTok. The message, according to the FTC — “whether a company intends – or doesn't intend – to have a site directed to kids isn't what controls the analysis. Instead, the FTC will look at the site's look and feel, as well as evidence that the company had actual knowledge that users are under 13.”

On September 4, 2019, the FTC issued another fine of \$170 million to YouTube for COPPA violations, including tracking viewing history of minors in order to facilitate targeted advertising.⁹ In the complaint the FTC and New York Attorney General alleged that YouTube, and its parent company, Google, violated COPPA by collecting personal information, in the form of “cookies”, from viewers of child-directed channels, without first notifying parents and getting their consent. YouTube earned millions of dollars by using the cookies to deliver targeted ads to viewers of these channels, according to the complaint. The FTC and New York Attorney General alleged that while YouTube claimed to be a general-audience site, some of YouTube's individual channels, such as those operated by toy companies, were child-directed and therefore required to comply with COPPA. As a result, YouTube announced that as part of the settlement, it would require channel operators to mark videos that are “child-oriented” and would use machine learning to automatically mark them as “child-oriented” if not marked already.

The Digital Economy Act 2017

The UK's Digital Economy Act 2017 is representative of the challenges faced in implementing privacy preserving solutions for age assurance and what can go wrong. One of the provisions of the act was the creation of a UK age-verification regulator to publish guidelines about how pornographic websites which operate “on a commercial basis” should ensure their users are aged 18 or older. The regulator would be empowered to fine those who fail to comply up to

£250,000 (or up to 5% of their revenue), to order the blocking of non-compliant websites, and to require those providing financial or advertising services to non-compliant websites to cease doing so.

The introduction of the scheme was subject to multiple delays. It was expected to begin in 2018 but was delayed until spring 2019, then to July 2019, and then for a further period in the region of six months. In October 2019, the government abandoned the mandate altogether, in favor of replacing it with a forthcoming wider scheme of internet regulation.

The reasons for these delays were related to the inability to verify age without relying upon bulk data ingestion that otherwise violated privacy rights. The provisions for the age verification of pornographic website users raised concerns about the privacy implications of collecting user data, and the possible ineffectiveness of a method focused on restricting payments to pornographic websites.¹⁰ Further, the regulation would impinge on existing legal rights for freedom of expression in its existing specification.

The inability for the act to be workable demonstrates the significant difficulties faced in implementing age assurance that is also privacy preserving for all users. Sites offering products and services that may be legal, but provide content or services that may be culturally sensitive or even taboo (pornography, marijuana, etc.), may experience decreased traffic due to an aversion due to PII capture concerns and the potential for de-anonymization through the reselling of this information, or through data breaches where the sensitive nature of accessing age restricted goods and services could be embarrassing, or worse. There is precedent for these concerns. The 2015 hack of extramarital affair dating site Ashley Madison exposed data of 32 million users, some of which was used for “sextortion” scams where users were blackmailed in return for their usage on the site being kept secret from partners.¹¹

General Data Protection Regulation (GDPR)

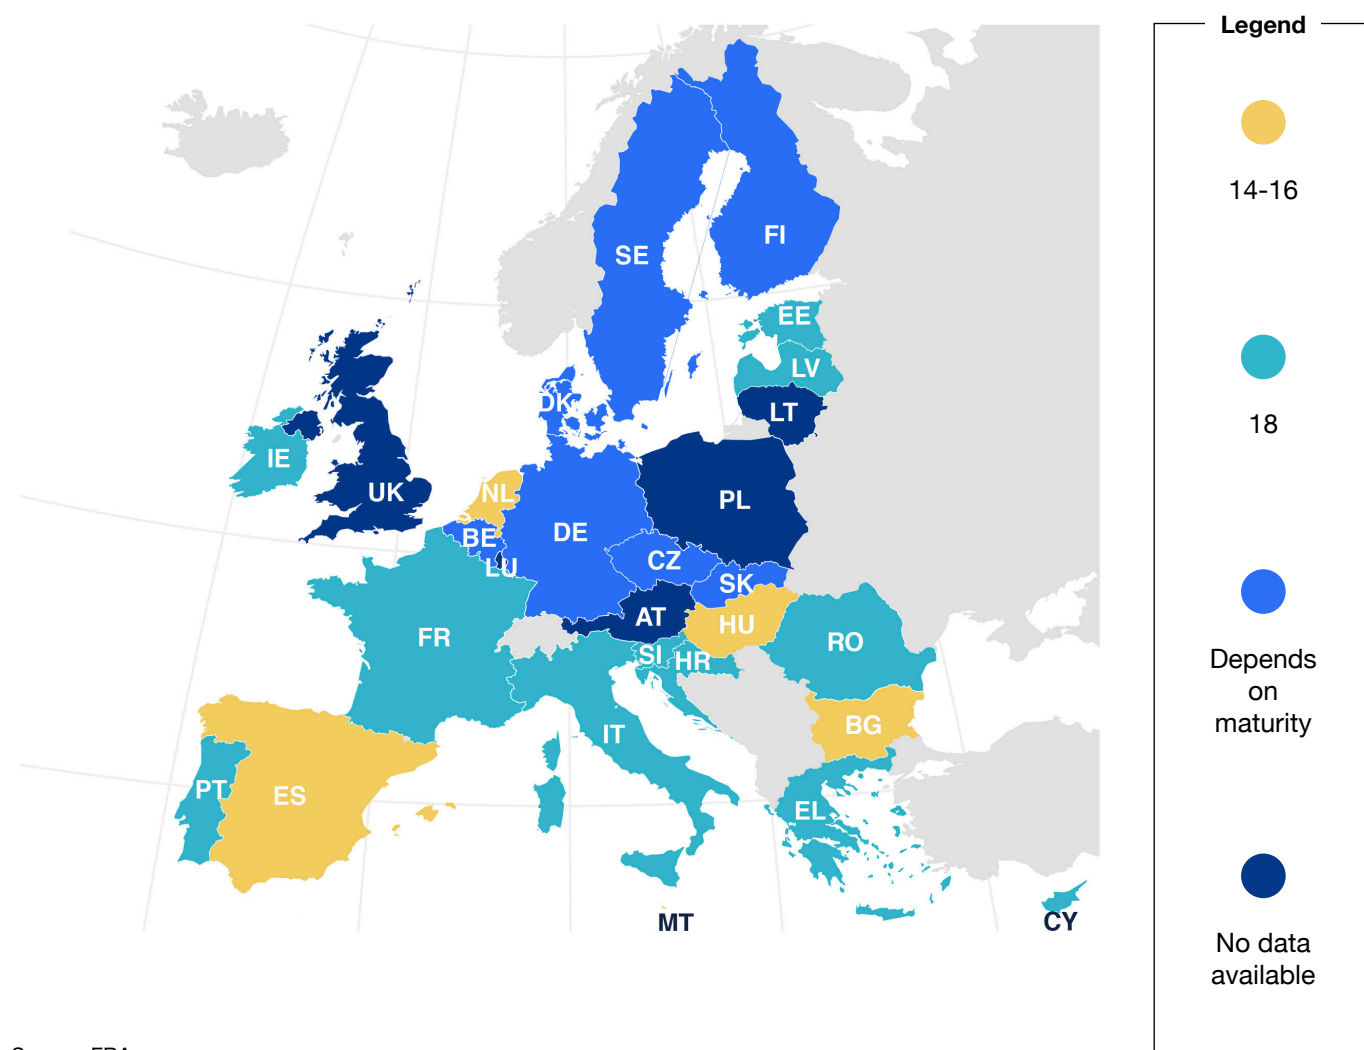
GDPR went into effect in Europe in May 2018, providing sweeping provisions for privacy protection. According to Article 8 of the GDPR, when processing personal data in relation to information society services offered directly to children under 16, consent shall be given or authorised by the holder of parental responsibility. Member States, however, may provide for a lower age, not below 13 years.¹² Non-compliance is punitive; if an organization is not compliant with GDPR, it is exposed to fines of up to €20 million, or 4% of annual worldwide turnover, whichever is higher.

While a laudable attempt to transpose similar standards as COPPA upon the European Union, implementation lacked consistency across member states in terms of the digital age of consent:

- **13 years:** the Czech Republic, Denmark, Ireland, Latvia, Poland, Spain, Sweden, the United Kingdom
- **14 years:** Austria, Italy
- **15 years:** Germany, Hungary, Lithuania, Luxembourg, the Netherlands, Slovakia

This patchwork of regional compliance requirements made child consent management for organizations working across the EU highly complicated, highlighting the significant challenges of policing a globally accessible platform across geographically nuanced regions.

Age at which children can provide consent for the use of their personal data



Audiovisual Media Services Directive (AVMSD)

The EU's Audiovisual Media Services Directive (AVMSD) governs EU-wide coordination of national legislation on all audiovisual media including traditional TV broadcasts and on-demand services. Originally implemented in 2007, the directive was updated in 2017 to reflect changes in digital media consumption patterns since its inception.¹³

The update included specific rules to protect minors from inappropriate on-demand media audiovisual services, segmenting content into two areas:

Content which might seriously impair minors

- Content that might seriously impair minors must not be included in any program, signifying a total ban.
- Programs which “might seriously impair” the development of minors containing pornography or gratuitous violence are prohibited.
- Programs which might simply be “harmful” to minors can only be transmitted when it is ensured that minors will not normally hear or see them. This can be done by selecting the time of the broadcast or by any technical measure such as e.g. encryption.
- When harmful programs are not encrypted, they must be preceded by an acoustic warning or by means of a clearly identifiable visual symbol throughout their duration.

Content which is likely to impair minors

- Such content must be ensured, by selecting the time of the broadcast or by any technical measure (e.g. encryption), that minors in the area of transmission will not normally hear or see such broadcasts.



Programs which “might seriously impair” the development of minors are allowed in on-demand services, but they may only be made available in such a way that minors will not normally hear or see them via age assurance mechanisms. There are no restrictions for programs which might simply be “harmful.”

AVMSD is noteworthy in how it has expanded the definition of broadcaster from an entity that was traditionally terrestrial (broadcast or cable TV), or regionally specific (such as satellite TV), to now encompass global streaming media services such as Netflix and Amazon Prime, as well as social media networks capable of delivering content such as Twitter, Facebook, and YouTube. This made these entities responsible for all user generated content, placing a staggering onus on them to monitor and take down content that is deemed inappropriate, and to provide the requisite age assurance tools to keep minors separated from programs that “might seriously impair them.”

China Audio-Video and Digital Publishing Association (CADPA) Age Ratings

Regulation body China Audio-Video and Digital Publishing Association (CADPA) introduced an age-based rating system for games in the country in December 2020. Under the standard, games are divided into three age categories, each represented by a color: 8+ (green), 12+ (blue), and 16+ (yellow). Games must display these labels on their website and other relevant materials.

The policy expanded upon measures taken in 2019 to protect minors, and specifically their eyesight from an over-exposure to screens. According to research by the Chinese authorities conducted in 2015, over half of the population over 5 years of age experienced myopia. While this could be the result of a confluence of factors, online gaming was singled out as one of the core areas for regulatory intervention.¹⁴ These measures included an obligation for players to register in a game under their real name to be able to verify their ages, as well as spending and playtime limits depending on age. The China cyber-curfew imposes a 10pm - 8am restriction on access to online gaming and also restricts access to 90 minutes of gaming on weekdays and three hours on weekends and holidays for under-18s.¹⁵

What separates CADPA from other age assurance regulations is the use of biometrics by gaming companies as a means of enforcement of age specific time restrictions. Tencent Games introduced a facial biometric check called “Midnight Patrol” as its response to the CADPA requirements. In June 2021, Tencent prompted an average of 5.8 million users a day to use facial biometrics for age assurance, blocking more than 90 percent of those who rejected or failed facial verification from access to their accounts.

The example of CADPA, and Tencent’s solution to age assurance, demonstrates some of the

cultural and societal friction that can occur with the introduction of emerging technologies as a means of proving age. Even in China, a nation state known for its tight controls of citizen privileges, the deployment of biometrics for age assurance has been contentious, with The China Security and Protection Industry Association stating that despite benefits, the wide use of facial recognition makes it ripe for abuse.¹⁶

Age Appropriate Design Code (AADC)

Age Appropriate Design Code (AADC) is not a new law, but a clarification on how GDPR is to be enforced and extends the regulators' reach beyond the basic protection of children's personal data into how children's digital experiences are designed. AADC makes clear that companies can't be in compliance with GDPR if they use personal data in ways that don't have the best interests of their potential child users in mind.¹⁷ It extends privacy protections to:

- Children up to 18 years old
- Not only services aimed at children, but those likely to be used by them
- The passive collection of data by connected devices (such as voice assistants or toys)
- 'Inferred data' such as that created by ad targeting platforms
- Any company with operations in the UK, any non-EEA company with users in the UK, and, post-Brexit – any EEA company with users in the UK

While AADC is regional, it is expected to be received as a template for other developing regulatory standards around age verification, as was the case with Europe's General Data Protection Regulation (GDPR) and data privacy.

AADC had only recently gone into effect at the time of publication, but its implications are likely to be far reaching, as well as something of an unwelcome surprise for many internet service providers that have yet to hear of the regulation. As a GDPR amendment, non-compliance carries the same punitive measures — fines of up to EU20 million, or 4% of annual revenues, whichever is higher. Further, AADC is likely to be applied to a wide swathe of service providers that are unaware of its reach — those that are offering services “likely to be used” by children. This ambiguity, along with the rather nebulous nature of “inferred data,” and the geographically challenging scope of “any company with operations in the UK, any non-EEA company with users in the UK, and, post-Brexit – any EEA company with users in the UK” could result in some significant fines being levied against companies that had no idea that they were out of compliance. It is probably not coincidental that the day AADC regulation went into effect (9/2/2021), WhatsApp was fined EU225 million, the second highest ever fine for GDPR non-compliance.¹⁸

Assuring Age - Tools and Technologies

Returning to the “4C’s” of risk for children online — content, contact, conduct and contract, there are various technologies and policies designed to provide assurances that children are protected:

01 Content

Age assurance tools can be used to require users to demonstrate that they are 18 or older in order to block content that is illegal or deemed harmful.

02 Contact

Age assurance tools can be used to flag adult users that are interacting with children in an all ages platform, to keep adults out of online environments designed specifically for children, and to keep children out of online environments designed for adults.

03 Conduct

Age assurance can be used to reduce the risks of younger children being influenced or coerced by older children or adults, but cannot prevent harm that can occur between users of the same age groups.

04 Contract

Data protection legislation in many countries prevents companies from collecting data from children under the age of 13, however there is room for greater granularity in terms of age determination and this limit does not apply to children from 13-18. Age assurance tools could provide greater specificity in the age that children become eligible for data collection and targeted marketing.

Forms of Age Assurance

There are a spectrum of technologies for age assurance that vary widely in their ability to assess the age of an individual, the accuracy of the findings, and the amount of private data that is captured in the process. While there is yet to be some form of equivalent to a “Turing Test” for age assurance, this field is experiencing a renaissance in innovation and investment as the digital domain seeks to reliably answer the question, “how old are you?”

Self Declaration

Self declaration is the weakest form of age assurance, offering in some cases nothing more than a checkbox or button for verification that the user is of an appropriate age to access the subsequent content or services. For instance, some websites offering the sale and delivery of alcoholic beverages have a pop up that appears on the homepage asking the visitor “are you 21 or older?” before allowing access, depending on the choice of “yes” or “no.” While self declaration is inherently weak, it can be designed in a way that makes the assurance slightly more robust.

- When a child enters a date of birth that is below the minimum age, the service can deny access on any repeated attempts from the same IP address, even if an appropriate age is entered.
- Users can be requested to enter a specific date of birth, which is more likely to be input incorrectly if a child is trying to calculate an age that is appropriate.
- Where a child has entered an age that is above the minimum age, the provided age can be checked again later in the process. Children who provided a false date of birth may not remember the date that they originally entered on the second prompt.

Strengths:

- Self declaration is easy for children to use.
- Self declaration is easy for parents to provide consent for, presuming children would actually ask for content.
- Self declaration captures no private information.

Weaknesses:

- Self declaration offers a very low level of assurance -- children can, and do lie about age.
- From a regulatory perspective, COPPA and others have strict guidelines around parental consent for each time a child accesses self declarative age gating. It is unreasonable to expect that a child will engage their parents for permission on every single occasion that they open a site.

Database Checks of Identity Attributes

Database checks of identity attributes, or “hard identifiers”, are verified sources of identification to prove age, such as a driver’s license, passport or other forms of public or private entity issued identification. Some identifiers display a date of birth explicitly, while others, such as a credit card, act as a proxy since a cardholder must be over 18 to have an account. These identifiers are considered a high level of age assurance since the documents or credentials have already passed a verification process and typically have proof of authenticity such as holograms and watermarks that would be challenging to counterfeit.

Strengths:

- Hard identifiers are relatively ubiquitous among the adult population.
- Hard identifiers are issued by reputable authorities that have designed proof of authenticity into the credentials themselves.
- Hard identifiers are relatively easy to procure, given the maturity and abundance of organizations such as credit bureaus and data aggregators.

Weaknesses:

- Hard identifiers often contain a lot more personal information than just age, requiring more intensive privacy preserving measures to be taken when they are being used purely for age assurance purposes. This presents a challenging “catch-22” situation for companies relying on hard identifiers for age verification -- they can be used to provide age assurance, but storage of these credentials would not be compliant with regulations such as COPPA and GDPR.
- Hard identifiers are robust credentials, but there is no proof that the individual using them is the legitimate owner in an online environment without augmentation with identity proofing such as facial biometrics to match the user to the hard identifier.
- Hard identifiers may disadvantage children and adults that do not have access to official documentation.



Biometrics

Biometrics measure a human's physical characteristics to verify and authenticate an identity. Biometrics encompass several physiological modalities, including fingerprint, face, iris, palm vein, and voice. Biometric data can be used for both age verification and age estimation. For example, facial recognition can be used to secure access to devices and content by locking them to the specific features of an individual, whereas facial analysis can estimate the age of an individual without requiring any further comparison to information.

Biometric data can be processed in near real-time and, in conjunction with contextual data such as length of time using a service, it can be used to build an age profile. It can also be rendered more robust with the addition of other forms of age assurance such as hard identifiers or self declaration.

Strengths:

- Biometrics vary in strength and can be used from everything from low to high risk age assurance scenarios.
- Biometrics can be used passively for age estimation and, combined with other signals such as device intelligence, can provide robust assurance that an individual is who they claim to be.
- Biometrics are simple to operate and relatively ubiquitous in the current generation of smartphones.
- In age estimation scenarios, there is no need for personal data capture beyond access to the age restricted goods or content.

Weaknesses:

- Some people express a strong degree of reluctance to use biometrics. A Liminal survey conducted in July 2021 found that a quarter of US consumers never use biometrics on their smartphones specifically due to privacy concerns. Tencent's usage of facial recognition for age verification has flagged privacy violation concerns, even in China where civilian surveillance is a fact of life.
- Biometrics are a form of deterministic data, meaning that in most circumstances, a biometric is compared to a stored record. This could be in contravention of regulations relating to the storage of data on children.
- Certain forms of biometrics such as facial recognition have proven to have significant inaccuracies in recognizing characteristics of dark skinned people.¹⁹ Further, the efficacy of using biometrics for people of disabilities is also unclear.²⁰
- The accuracy of facial analysis is lower on younger children. Also, the margin of error with age estimation may mean that children close to an age boundary such as 18 could be falsely allowed access.

Behavioral Biometrics

Behavioral biometrics are a class of authentication solutions that use dynamic identifiers based on human behavioral patterns. Distinct from ‘traditional’ biometrics, which uses absolute identifiers such as fingerprints and facial features, behavioral biometrics can be kinesthetic or device-based. Behavioral biometric data (e.g., typing speed, signature, gait analysis) is more probabilistic than biometric data, inferring that a pattern of behavior belongs to a specific individual. For age assurance purposes, its usage is in its infancy.

As with biometric data, behavioral biometrics can be processed in near real-time and, in conjunction with contextual data such as length of time using a service, can be used to build an age profile. It can also be rendered more robust with the addition of other forms of age assurance such as hard identifiers or self declaration.

Strengths:

- Behavioral biometrics can be used passively for age estimation and verification, therefore not triggering end user privacy concerns.
- When behavioral biometric data is being used specifically for age estimation, no PII data is exposed in the process.

Weaknesses:

- Many of the same issues relating to biometrics are also applicable to behavioral biometrics; there is significant scope for false positives, particularly with some ethnographic groups.
- For age assurance, behavioral biometrics may just be too vague for anything other than “guesstimates” of a specific age range.
- Biometric collection and storage without user notification and consent is currently a legally ambiguous area, but is likely to become increasingly regulated in the wake of malpractice by big tech companies.²¹ This may seriously impact the ability to use biometrics, and specifically behavioral biometrics passively for age assurance purposes.

Profiling

Profiling refers to the processing of data to analyze and infer information about a user, or to predict aspects of user behavior. Data used for profiling is derived from information that users choose to share about themselves, as well as information that is inferred about them, such as time spent on a webpage, where their cursor hovers, time of access, interests, friends, and location.

Profiling service providers have developed predominantly for commercial advertising purposes, but the granularity of data collected can often be sufficient for general age estimation. For example, a service can determine a child's height, daytime location, interests, and ages of best friends from information that is shared or inferred.

While the data is probably not sufficient for accurate age assurance alone, combined with other sources such as self declaration, it can be an important tool for building a solid profile of the individual's age to a reasonable degree.

Strengths:

- Profiling can build on existing deterministic data to provide a more accurate estimation of age without explicitly requesting user input.
- The profiling industry is relatively mature and is rapidly evolving due to developments in AI and ML.
- Data is increasingly accurate, leveraging a broad range of data sources.

Weaknesses:

- Profiling creates a significant tension between data collection, processing, and privacy rights. For children, this may be directly in conflict with regulation designed to protect their privacy. AADC, for instance, specifically prohibits collection of inferred data on children.
- Profiling is highly likely to capture information far beyond that needed for age assurance purposes, which can be exploited for purposes such as targeted advertising.
- Companies that have been most active in profiling have business models that are based on the sharing of information with third parties, further increasing the risk of children being exposed to or targeted with age inappropriate content.

Capacity Testing

Capacity testing is analogous to the now familiar CAPTCHA tests for proof of human rather than machine intelligence, but instead estimating a person's age based on their aptitude or capacity to complete a test. For example, a child may be requested to complete a language test, solve a puzzle, or undertake a task that gives an indication of their age range.

Strengths:

- Capacity testing is privacy preserving.
- Capacity testing, combined with behavioral biometrics and profiling information can provide a fairly accurate measure of age estimation while not requiring a significant amount of user interaction.
- Capacity testing can be designed in a way that gamifies the age assurance process, making it “fun” for end users.

Weaknesses:

- Capacity and problem solving skills are not linear with age; many children of the same age have widely varying levels of problem solving and language skills. This may exclude children of an eligible age from accessing products and services.



eIDs

An electronic identity (eID) is a digital representation of an individual's real world identity, made up from attributes such as their name, date of birth, address, etc. These attributes are derived from hard identifiers such as a passport or social security number, as well as biometric attributes such as a facial scan or fingerprint. Once the eID is established, this can be stored in a digital identity wallet (typically a mobile device) and these attributes can be shared with service providers as required to establish identity and other characteristics such as age. In theory, the user can choose which attributes are shared and which are not, meaning that privacy control can be self sovereign. eIDs can also leverage tokenization technology that has long been established in the card payment industry as a means of sharing credentials without exposing any valuable personal identity information (PII). For example, PAS1296:2018 is a UK standard for an age checking token exchange, providing assurance that an individual is between the age of 13 and 17, while sharing no other data.²²

Age assurance is likely to be one of the most important use cases for eIDs, but there is something of a “Catch-22” situation in that this use case is dependent on eIDs to be developed and to proliferate - one cannot occur without the other. There are also significant challenges with interoperability between digital wallets, with some developers taking an exclusive approach to the use of their platforms. It will be incumbent on the identity industry to ensure that interoperability is designed into ecosystems to assure that digital identities and their usage can evolve without the potential for balkanization.

Strengths:

- eIDs solve age assurance to a high degree of accuracy; it allows for an individual to prove digitally that they are who they claim to be, including proof of exact age.
- eIDs could be privacy preserving, divulging only the information required to answer the question, “are you old enough?”

Weaknesses:

- eIDs are nascent and while there is much development occurring in this space, standards and interoperability remain unclear.
- eIDs could be privacy preserving, but invariably are not. Additional information is routinely shared unnecessarily at the request of service providers that have business models built around data harvesting and reselling to third parties.
- Several eID schemes rely upon the verification of a mobile driver's license (mDL) or verification against a civil registry. eIDs schemes will require the development of a voluntary ID intended for use by children. As an example, Belgium has rolled out Kids-ID to facilitate children's access to youth-only sites.²³

Self Sovereign Verifiable Credentials

Self-sovereign identity (SSI) is an emerging framework based on decentralizing all digital identity attributes of the user. Individuals retain management of their digital identity (i.e., such that the user is ‘sovereign’ over the digital identity of their own self). SSI is often distinguished from current methods/practices by using a secure distributed ledger technology and shifting more power and decision authority to the user.

SSI can be considered a philosophical perspective on an approach to digital identity. This philosophy has many different manifestations, but almost all are rooted in the ten SSI principles defined by Christopher Allen in “The Path to Self-Sovereign Identity”²⁴:

1. **Existence:** Users must have an independent existence.
2. **Control:** Users must control their identities.
3. **Access:** Users must have access to their own data.
4. **Transparency:** Systems and algorithms must be transparent.
5. **Persistence:** Identities must be long-lived.
6. **Portability:** Information and services about identity must be transportable.
7. **Interoperability:** Identities should be as widely usable as possible.
8. **Consent:** Users must agree to the use of their identity.
9. **Minimization:** Disclosure of claims must be minimized.
10. **Protection:** The rights of users must be protected.

Self Sovereign Verifiable Credentials are SSI tokens of identification that can be distributed and retracted as needed by an end user to confirm, for instance, they can confirm that the credential is provided by an individual who is over a certain age, but do not need to divulge their date of birth or any other form of identifying characteristics.

Strengths:

- SSI and Self Sovereign Verifiable Credentials are the ultimate in user controlled identity consent management, enabling the end user to determine what about them is shared with who, and for how long.

Weaknesses:

- SSI could be considered “self-attested,” and therefore unreliable for age assurance in the same way that capacity testing is seen today.
- SSI and Self Sovereign Verifiable Credentials are in their infancy and may not come to widespread adoption due to innovation focused on eIDs / mDLs as a “good enough” solution to the age assurance problem, particularly combined with tokenization designed to protect PII.

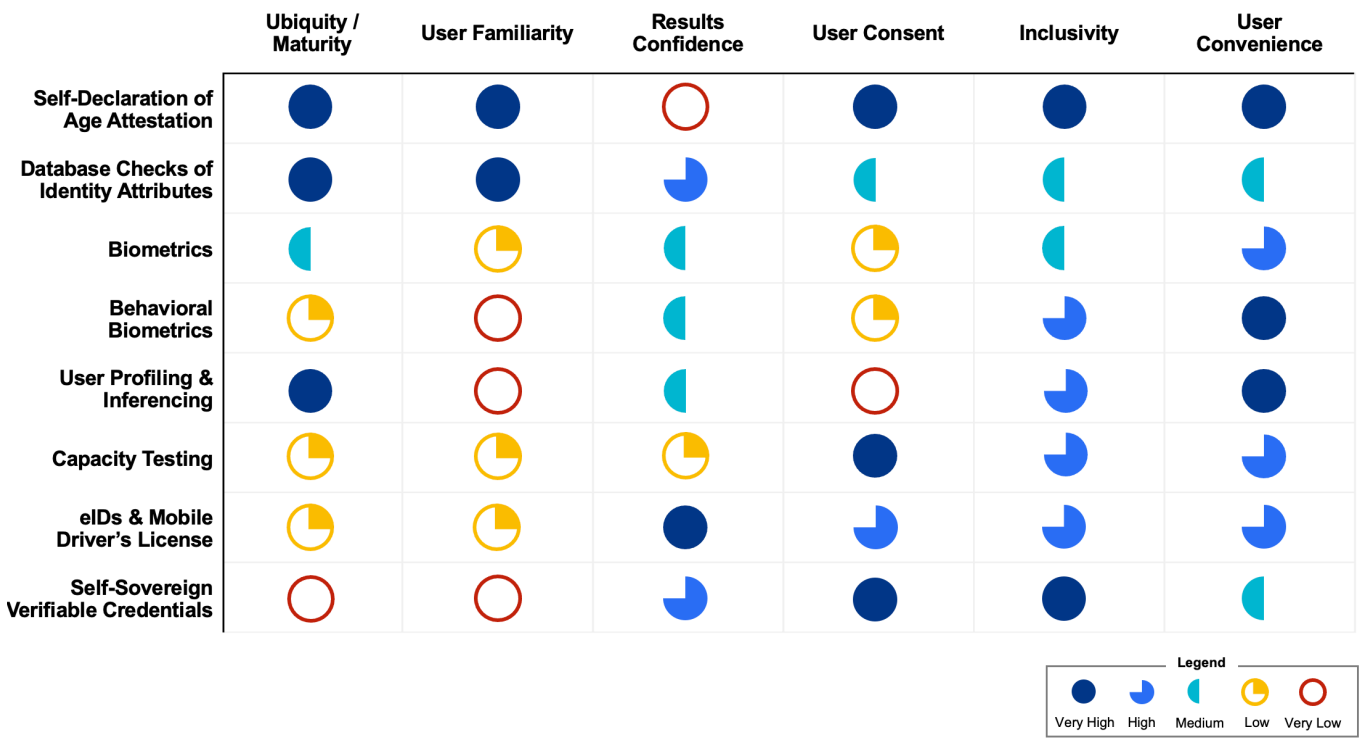
Comparison of Age Assurance Modalities

Assessing the spectrum of age assurance solutions that are available, we consider a number of criteria that will be significant in the adoption of each of these:

- **Ubiquity / Maturity:** how commonplace the technology is in the general population
- **User Familiarity:** how accustomed / comfortable end users are with using this technology
- **Results Confidence:** how reliable the findings of this technology are in assuring age
- **User consent:** how much control the end user has in the data provided during or after performing the activity
- **Inclusivity:** the extent to which the use of the technology is unbiased and non discriminatory for all members of society
- **User convenience:** the degree of effort required for the end user to perform the age assurance task

Age Assurance Modalities Offer Varying Degrees of User Acceptance

Today’s market is defined by fragmented solutions. In isolation, no current age assurance modality is uniquely fit for purpose



Comparing the modalities, it becomes apparent that there is no single solution that is widely available today that fulfils the needs of the criteria above.

- Self declaration, being self attested, is inherently weak for age assurance by itself.
- Hard identifiers are inconvenient, exclusionary and overshare PII.
- Biometrics and behavioral biometrics are emerging, unfamiliar, and come with consumer privacy and regulatory concerns.
- User profiling is not consensual and possibly illegal in its data trawling of information on children for age assurance purposes.
- Capacity testing is insufficiently accurate to determine age, or even an age range.

This leaves eIDs and Self Sovereign Verifiable Tokens as the best candidates to deal with the conundrum of age assurance technologically. But, neither of these have yet gained the ubiquity, maturity and familiarity needed for widespread adoption.

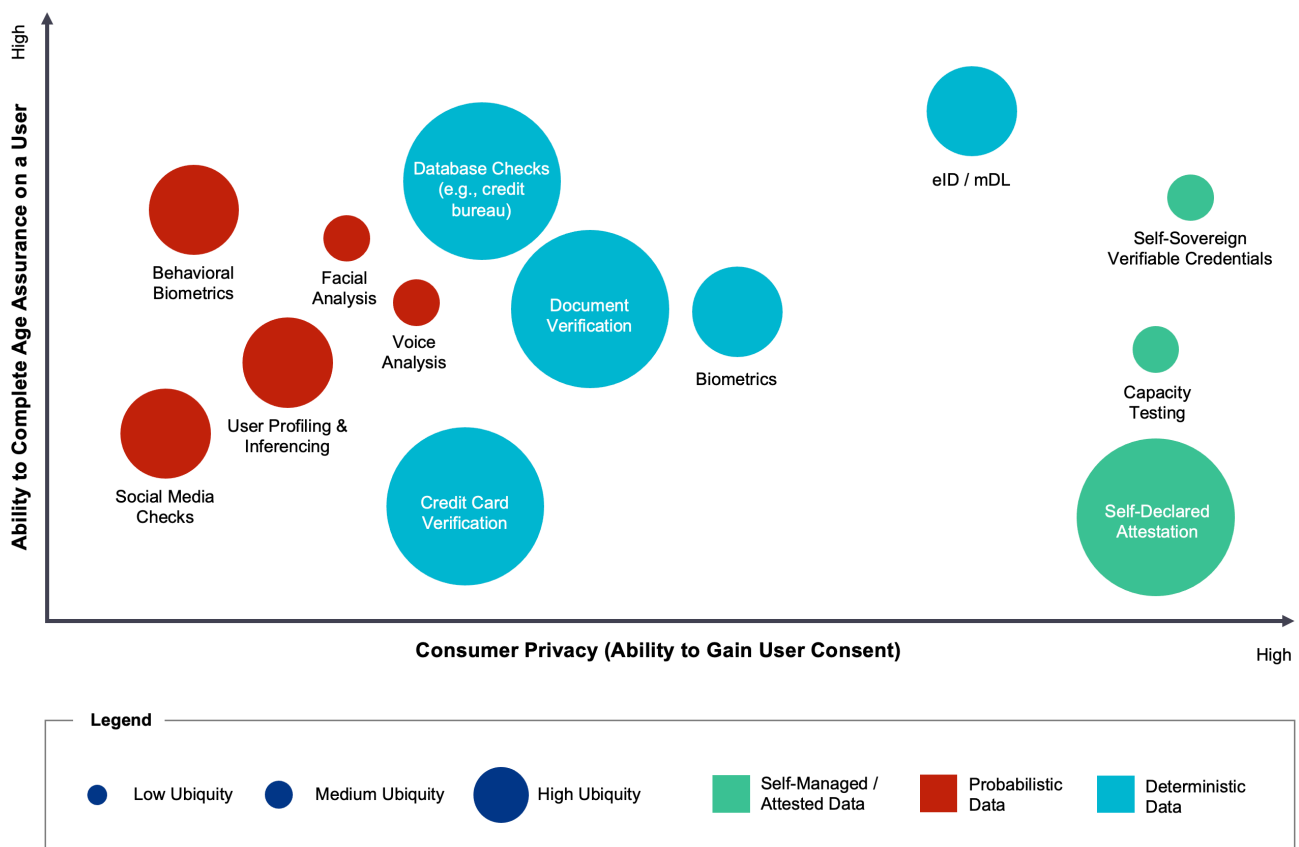
Mapping the various age assurance modalities to the attributes of user consent (privacy) and the ability to complete age assurance on a user, we can divide these into three categories — self managed / attested data, probabilistic data, and deterministic data.

- Self managed / self attested data has thus far been weak for age assurance, as the attestations of children can be considered unreliable with something as simple to bypass as a button or checkbox for entry. Capacity testing shows some promise as it develops as a new innovative field for age assurance, but a definitive Turing Test for age proves elusive at this time. With this in mind, this and self attestation should only be used for low risk age restrictions where access to goods and services are unlikely to be seriously detrimental to the welfare of a child. Self sovereign verifiable credentials would potentially meet the sweet spot of a high degree of end user control combined with highly reliable age assurance. However, as highlighted, the technology currently lacks meaningful levels of adoption.
- Probabilistic data (behavioral biometrics, social media checks, profiling, and facial and voice analysis) have some promising capabilities, particularly in the ability to passively infer or predict the age of a person. Alone, however, they can be considered relatively weak in their ability to accurately assure age. Further, the ability for passive collection of data on anyone, but particularly children, is likely to invite more regulatory scrutiny and the likely outcome that unwarranted collection will become outlawed in many regions.

- Deterministic data derived from hard identifiers (document verification, database checks, and credit card verification, for instance), have varying degrees of assurance, based on the issuing authority. They do however have the advantage of being relatively easy to acquire, particularly via data aggregators such as credit bureaus. However, where these identifiers are problematic for age assurance is the likelihood that extraneous PII is also captured and stored. This may be in direct conflict with privacy protecting regulations such as GDPR and CPPA, as well as child protection regulations such as COPPA and AADC. Further, some of the agencies most well known for providing data checks have been subject to massive data breaches in recent years, including the current high watermark for exposed PII caused by the 2017 Equifax data breach of over 147 million records. Can they be trusted with something as valuable as children's identities?

Self-Managed Data Offers Most Opportunity for Privacy-Preserving Age Assurance

Self-attested data that offers high assurance still offers the least ubiquity, making deterministic and probabilistic data sets a more realistic option for organizations rapidly attempting to meet evolving compliance requirements



eIDs show some significant promise in their potential to protect privacy with data sharing that is controlled by the identity owner, or the use of tokenization to obscure non essential information. It may also be considered more reliable than self managed / attested data in the form of self sovereign verifiable credentials in that the eID is issued by a state entity or other trusted third party. However, as with self sovereign verifiable credentials, the deployment of eIDs remains relatively nascent.

The best solution therefore for today's organizations that need to meet age assurance compliance standards is a combination of deterministic and probabilistic sources that collectively provide enough information to assure age to what would be considered "reasonable" from a regulatory perspective.

eIDs show some significant promise in their potential to protect privacy with data sharing that is controlled by the identity owner, or the use of tokenization to obscure non essential information. It may also be considered more reliable than self managed / attested data in the form of self sovereign verifiable credentials in that the eID is issued by a state entity or other trusted third party. However, as with self sovereign verifiable credentials, the deployment of eIDs remains relatively nascent.



Challenges to Age Assurance

While requirements for age assurance are clearly a significant area of innovation at the moment, there remain a number of speed bumps in the road before solutions match requirements.

Some can and probably will be resolved via technological evolution. Nation states are collectively gaining momentum to digitize identities for their citizens, and private entities such as payment networks and telcos see new opportunities in leveraging their trusted reputations and globally prolific reach to develop new initiatives, solving problems such as internationally recognized vaccine passports and digital voting.

Other challenges are however less predictable. Companies that have grown immensely lucrative via online advertising revenues are being politely requested to behave more ethically in how they capture data and serve content. Their subsequent behavior is likely to be catalytic in defining the regulatory response and ultimately determining the degree of age assurance lockdown that is implemented.

Regulations vs. Today's Age Assurance Capabilities

Finding the right balance between sufficient child protection and overreach is a fine art. The UK's Digital Economy Act 2017 is a good example of how a badly implemented age assurance regulation can be unworkable in its ability to preserve privacy rights for those old enough to access age restricted content and services.

Regulations can also expand to include industries that have evolved without the need for age assurance to the same levels of compliance that have developed over decades in others. One such example is AVMSD, which forced online service providers to be responsible for content available to minors across their platforms, including user generated content which can quickly be virally distributed across networks. Compounding the problem, the more contentious / offensive the content, the more likely it is to be a viral sensation and amplified by reshares. Technologically, AI and ML are not sufficiently sophisticated to parse video and photographic content for obscene or damaging content, leaving the task primarily to human judgement — a truly herculean task. Consequently, mistakes are made — according to research by NYU Stern, one in ten Facebook moderating decisions is inaccurate, lending itself to 300,000 errors per day.²⁵ It is inevitable therefore that children will be exposed to content on this and other social media sites that is not age appropriate.

While AVMSD has had a seismic impact on the way that social media has been required

to add age assurance controls to their portals, AADC has the potential to be even more disruptive, expanding GDPR child assurance regulations to cover:

- Children up to 18 years old
- Not only services aimed at children, but those likely to be used by them
- The passive collection of data by connected devices (such as voice assistants or toys)
- ‘Inferred data’, such as that created by ad targeting platforms
- Any company with operations in the UK, any non-EEA company with users in the UK, and, post-Brexit – any EEA company with users in the UK

The most obvious problem here is the onus placed on the service provider to ensure that they are compliant when they may be unaware of their under 18 audience, or that they are passively collecting information on children, or that they are collecting and inferring information on an individual that may or may not be a child. There is also a high likelihood that they are completely unaware of AADC in regions outside of the UK and EU.

One further area of note relating to regulations -- there are often contradictory regulatory requirements for age assurance and privacy preservation, which are technologically challenging to meet today.

How do I confirm age to a reasonable degree of certainty using today's deterministic and probabilistic solutions, yet ensure that no information on that individual is stored, accidentally captured, or even inferred without their consent?

This regulatory request can not easily be met with today's age assurance technologies.



The Age Assurance Conundrum


There is a significant way to go in the development and implementation of solutions that can reliably assure age without unnecessarily capturing private information on both adults and children without their consent. Beyond the choice between age assurance and privacy invasion, there is however a third factor for consideration — user experience (UX).

UX friction is a real problem, particularly as consumers become increasingly accustomed to a digital, on demand experience in all aspects of their lives. As is demonstrated in Liminal consumer research, 42% of consumers have given up on a digital onboarding process for a bank account or loan, with the most common reason being that they were requested to provide too many personal questions. In a physical world scenario, this may have been acceptable. In today's digitally transformed environment, a few clicks is too many.

This leads to the age assurance conundrum.

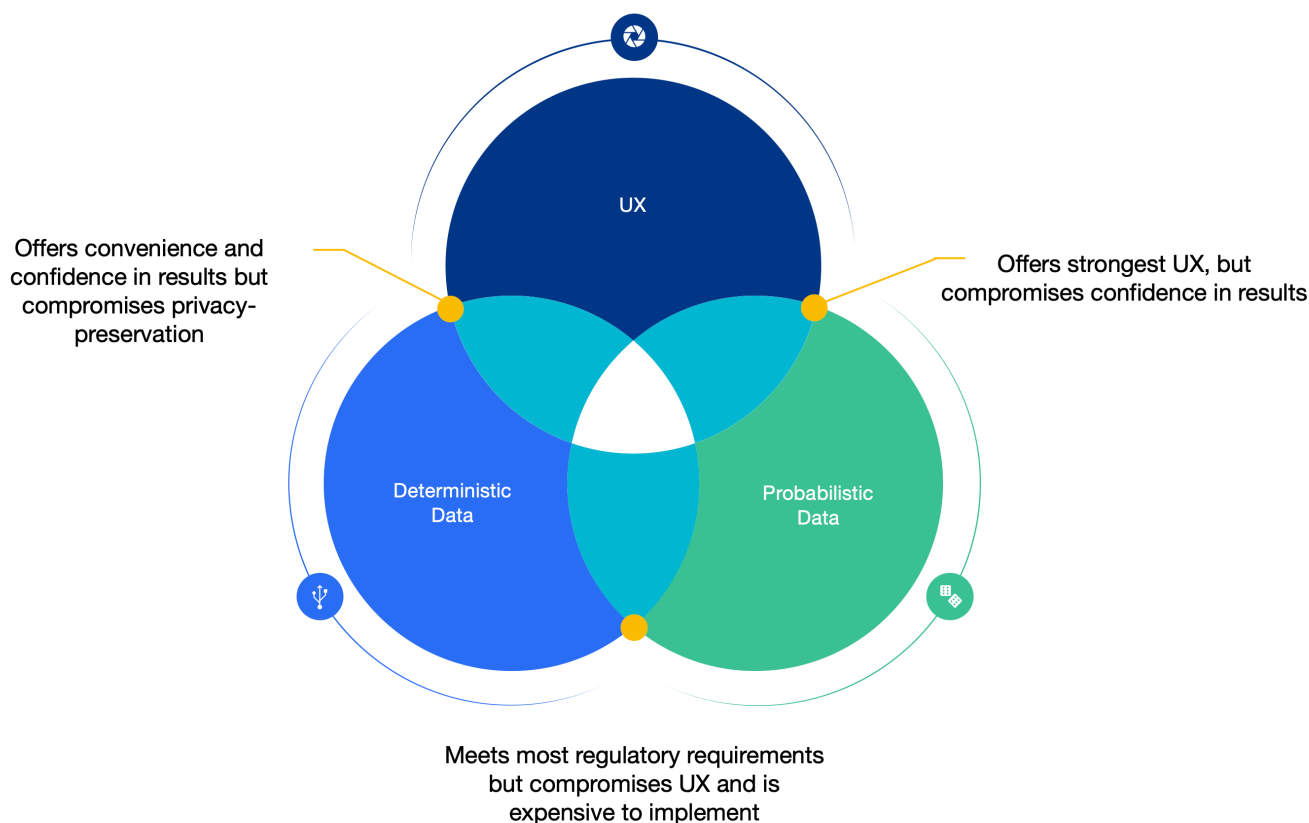
You can have user convenience, age assurance and privacy protection via consent management, but based on available technologies today, you can realistically only pick two.

- **A UX focus + probabilistic data** = convenient and privacy preserving (no need for explicit consent), but **not age assuring**. This scenario allows access for under age users and regulatory non compliance.
- **A UX focus + deterministic data** = age assuring and convenient, but potentially **not privacy protecting / facilitating user consent**. This scenario allows for the capture and storage of extraneous PII on children, and regulatory non compliance.
- **Probabilistic + deterministic data** = age assuring and privacy protecting via consent, but **not convenient**. This scenario meets regulatory requirements, but may present too much end user friction, onerous process implementation for the service provider, and not be operationally viable from a business perspective.



The Age Assurance Conundrum

You can have user convenience, age assurance, and privacy protection, but based on available technologies today, you can only pick two



Source: Liminal Advisory Services

Service providers are therefore forced to make a trade off to meet the regulatory requirements with the tools at their disposal today, and hope that their age assurance / privacy preservation activities are considered “reasonable” from a legal perspective, or choose to meet and even exceed compliance requirements from both an age assurance and privacy perspective, but risk significant customer drop-off due to a less than optimal user experience. There is no easy answer other than for companies to choose the risk tolerance level that they are most comfortable with, and be prepared to pay the price if this leads to regulatory enforcement.

Compliance vs. Capitalism

One of the greatest challenges with age assurance and the use of data pertaining to children is that the platforms and service providers that are the most popular with children have built business models primarily around the harvesting and selling of data on individuals to third parties. Facebook, for example, generates 98% of its revenue from advertising. In 2020,

Facebook's revenue from advertising exceeded \$84 billion – a 21% rise from 2019.²⁶

Age assurance and privacy regulations place a significant challenge on Facebook and other social networks (Instagram, TikTok, Snap, YouTube, etc.) to restrict the sharing of personal information with third parties, particularly information relating to children. But do the punishments fit the crime? Arguably, yes, so far. YouTube's \$170 million COPPA fine in 2019 did force the company to require channel operators to mark videos that are "child-oriented" as such, as did the fine of \$5.7 million to ByteDance for failing to comply with COPPA with their TikTok app, forcing the company to add a kids-only mode to TikTok.

It remains to be seen whether these companies will continue to comply with the evolving age assurance landscape, or whether the temptation to flout regulations for financial gain is worth the risk. TikTok's parent company, ByteDance, more than doubled its annual revenue to \$34.3 billion in 2020, a 111% YoY increase — a \$5.7 million fine could be considered a drop in the ocean.²⁷

Big Tech Ethics

Signs are not good that "Big Tech" will do its job in protecting children from harm given that, as detailed previously, revenue generation through advertising is the primary mode of business and children present a tempting swathe of the internet population that is ripe for monetization.

There are bigger issues with the toxic effects of current digital media consumption than targeted advertising. A September 2021 investigation by the Wall Street Journal (WSJ) uncovered internal documents from Facebook that the company knew that its Instagram service was damaging to the mental health of teenage girls.²⁸ "We make body image issues worse for one in three teen girls," stated one slide from a 2019 presentation by researchers that was posted to Facebook's internal message board. Among teens who reported thinking about suicide, 13 percent of British users and 6 percent of American users traced the desire to take their own lives to Instagram. Backlash from regulators and parents groups was swift, forcing Facebook to pause development of its Instagram Kids project as a result of the WSJ disclosure and calls from regulators for greater accountability and transparency from Facebook and its affiliates for how its algorithms are designed. The situation remains ongoing at the time of publication.

YouTube has also come under criticism for algorithms that lead to increasingly dangerous and controversial content being delivered in order to increase viewer stickiness. To gather data on specific recommendations being made to YouTube users, a Mozilla research project used a browser extension that let users self-report YouTube videos they "regretted" watching.²⁹

The research reported a wide variety of "regrets," including videos spreading COVID-19 fear-

mongering, political misinformation and “wildly inappropriate” children’s cartoons, per the report — with the most frequently reported content categories being misinformation, violent/graphic content, hate speech and spam/scams. A substantial majority (71%) of the regret reports came from videos that had been recommended by YouTube’s algorithm itself. The research also found that recommended videos were 40% more likely to be reported by the volunteers than videos they had searched for themselves. Mozilla even found several instances when the recommender algorithm put content in front of users that violated YouTube’s own community guidelines and/or was unrelated to the previous video watched.

Some companies have taken a more privacy centric approach to their business. Following the 2014 leak of hundreds of nude celebrity photographs from Apple iCloud accounts, the company embarked on a rebranding exercise, articulating that it was taking a privacy first approach.³⁰ However, more recently, cracks in Apple’s data privacy image have occurred -- the company reversed plans to include a child sexual abuse material (CSAM) scan of content on individual’s Apple devices in the recent iOS 15 update due to pushback from over 90 privacy advocate organizations.³¹ As the company continues to diversify from selling devices to selling services, Apple’s privacy focus may be a casualty of the transition.



Conclusion

Age assurance is becoming one of the defining issues of the internet today. Implemented well, age assurance could usher in a new era of use cases and societal advantages derived from the fusion of digital and physical domains around more reliable digital identity management... a safer and more trustworthy internet, greater inclusion of disenfranchised groups in civic initiatives, and a significant reduction in all forms of fraud, to name just a few.

Implemented poorly, however, age assurance could make the internet a place that is considered largely unsafe for children, stifling opportunities for education, entertainment, communication and personal growth, and holding back the evolution of digital identity ecosystems that would greatly benefit society.

There are greater forces at play in enforcing age assurance protections for minors than regulators or payment networks — notably, parents. As a maturing, digitally savvy society that is spending more and more time connected online either actively or passively, we are becoming more attuned to where children are going and what they are experiencing digitally. The internet should develop in a way that is analogous to the physical world, where there are the necessary checks and balances in place to ensure that it is safe for all, but not restricting access to experiences that can be considered formative.

For all digital citizens, the online protection of children cannot come at the sacrifice of their digital privacy. Further, any technology / methods implemented must not encroach upon data privacy rights for all.

The above statement is a laudable goal for how the internet should be. Unfortunately, at the crux of where age assurance stands today is the age assurance conundrum — user convenience, age assurance and privacy protection, pick two. There simply isn't a privacy preserving, age assuring solution that is easy to use or prolific enough. Yet.

This situation is likely to be resolved over time, primarily in the short term as eID initiatives gain traction. One significant advantage of state sponsored digital identity solutions such as eIDs is that they can be applied with more stick than carrot to enforce rather than encourage adoption, attributes that are less available to private sector initiatives that have to persuade usage with value added services, such as digital payments, loyalty programs, and convenience over other form factors (cards, and cash, for instance).

However, that doesn't solve for today, where big tech, and specifically social networks, are close to pariah status and appear on the cusp of heavy handed regulatory intervention in order to protect the safety and welfare of their underage users. While there is precedent for citizen

outrage at social media, such as the March 2018 Facebook / Cambridge Analytica scandal which largely abated after a few months, recent revelations pertaining to the impact of social media on children may be something that is considered less forgivable.

Age assurance will be a core component of the future internet, and one that will be fundamental in paving the way for a whole host of services that enable the seamless fusion of digital and physical domains. However, it will be incumbent upon all stakeholders to make the internet a child safe environment.

- **Develop solutions that fix the age assurance conundrum.** Collectively, the industry understands the problem and paradoxes of age assurance and privacy preservation. It will be essential to continue working towards this end goal via technology that meets these somewhat contradictory requirements, and via regulations that are not unworkable.
- **Design and Build Towards Identity Ecosystems.** Interoperability will be key to the success of age assurance solutions, particularly as digital identity ecosystems are in their infancy. Design of platforms locally should be considered in their global context, with an underlying ethos that collaboration will be key to ubiquity and adoption.
- **Protection before Profits.** It will be incumbent on big tech and other ecosystem players to self-police, or risk the consequences of regulatory intervention or worse, parental wrath — arguably nothing could galvanize support for the boycotting of service providers more than child risk. For companies that live by mantras such as “don’t be evil” and “build social value,” they can either walk the walk or face the consequences.
- **Regulate Proactively.** The UN Convention on the Rights of the Child remains a north star for how children have the right not just to protection, but to flourish. The internet is a powerful medium for education and enrichment and regulations should not overreach and impinge on the rights of children, as well as adults. Digital privacy rights for all need to be built into new regulations, not just protecting the child. Regulators need to stay ahead of digital developments and anticipate issues before they escalate. This will be especially challenging given the scope and velocity of new internet technologies, platforms and trends, but this is the only way to assure that online experiences remain protective of the most vulnerable.





About Us

LIMINAL

Liminal is a boutique strategy advisory firm serving digital identity, fintech, and cybersecurity clients, and the private equity and venture capital community. Since 2016, we have offered objective, high impact strategic advice, and analytical services, helping to support clients in crucial business decisions at all stages of the product and business lifecycle. We've advised many of the world's most innovative business leaders, investors, and government officials on building, buying, and investing in the next generation of integrated digital identity platforms and technologies. As a result, our clients trust us to set strategic direction in light of radically evolving ecosystem dynamics, pursue new growth strategies, capitalize on M&A opportunities, and optimize deal flow. We see the solutions to these complex digital challenges not as a 'what' but as a 'how.' We don't just tell you about the destination, we show you how to get there.

Learn More at www.liminal.co

Endnotes

1. Livingstone, S., Byrne, J., and Carr, J. "One in Three: Internet Governance and Children's Rights." UNICEF, January 2016. <https://www.unicef-irc.org/publications/795-one-in-three-internet-governance-and-childrens-rights.html>
2. ICO. "Internet Users' Concerns About and Experience of Potential Online Harms." Ofcom, March 2019. https://www.ofcom.org.uk/_data/assets/pdf_file/0028/149068/online-harms-chart-pack.pdf
3. Livingstone, S. and Stoilova, M. "The 4 Cs, Classifying Online Risk to Children." SSOAR, 2021. <https://doi.org/10.21241/ssoar.71817>
4. Verdeschi, John. "Protecting our network, protecting you: Preventing illegal adult content on our network." Mastercard, April 14, 2021. <https://www.mastercard.com/news/perspectives/2021/protecting-our-network-protecting-you-preventing-illegal-adult-content-on-our-network/>
5. "Mastercard issued an updated set of rules and requirements for adult merchants." Commercegate, 2021. <https://www.commercegate.com/blog/mastercard-issued-an-updated-set-of-rules-and-requirements-for-adult-merchants/>
6. Dugan, Kevin. "OnlyFans dustup gives a peek into the huge role banks play in the porn industry." Fortune, September 1, 2021. <https://fortune.com/2021/09/01/onlyfans-bank-fees-adult-performers-porn-industry/>
7. United Nations. "Convention on Rights of the Child." November 20, 1989. <https://www.ohchr.org/EN/professionalinterest/pages/crc.aspx>
8. Fair, Lesley. "Largest FTC COPPA settlement requires Musical.ly to change its tune." FTC, February 27, 2019. <https://www.ftc.gov/news-events/blogs/business-blog/2019/02/largest-ftc-coppa-settlement-requires-musically-change-its>
9. "Google and YouTube Will Pay Record \$170 Million for Alleged Violations of Children's Privacy Law." FTC, September 4, 2019. <https://www.ftc.gov/news-events/press-releases/2019/09/google-youtube-will-pay-record-170-million-alleged-violations>
10. Bielikova, Slavka. "Digital Economy Bill: Briefing to the House of Commons." Open Rights Group, September 12, 2016. <https://www.openrightsgroup.org/publications/digital-economy-bill-briefing-to-the-house-of-commons-on-second-reading/>
11. KrebsOnSecurity.com "Online Cheating Site AshleyMadison Hacked." July 19, 2015. <https://krebsonsecurity.com/2015/07/online-cheating-site-ashleymadison-hacked/>
12. Europa.eu. "Consent to Use Data on Children." FRA, April 24, 2018. <https://fra.europa.eu/en/publication/2017/mapping-minimum-age-requirements/use-consent>
13. Europa. EU. "Revision of the AVMSD." Policies, 2017. <https://digital-strategy.ec.europa.eu/en/policies/revision-avmsd>
14. BBC "China targets video gaming to tackle myopia in children." August 31, 2018. <https://www.bbc.com/news/world-asia-china-45366468>
15. BBC. "Video game addiction: China imposes gaming curfew for minors." November 6, 2019. <https://www.bbc.com/news/world-asia-50315960>
16. Borak, Masha. "Facial recognition in video games comes with security risks, Chinese industry group warns." SCMP, October 11, 2020. <https://www.scmp.com/abacus/tech/article/3105702/facial-recognition-video-games-comes-security-risks-chinese-industry>
17. Information Commissioner's Office. "Introduction to the Age Appropriate Design Code." 2021. <https://ico.org.uk/for-organisations/guide-to-data-protection/ico-codes-of-practice/age-appropriate-design-code/>
18. BBC. "WhatsApp issued second-largest GDPR fine of €225m." September 2, 2021. <https://www.bbc.com/news/technology-58422465>
19. Simonite, Tom. "The Best Algorithms Struggle to Recognize Black Faces Equally." Wired, July 22, 2019. <https://www.wired.com/story-best-algorithms-struggle-recognize-black-faces-equally/>
20. Whitaker, Meredith et al. "Disability, Bias, and AI." AI NOW Institute, New York University, November 2019. <https://ainowinstitute.org/disabilitybiasai-2019.pdf>
21. Larose, C. and Prescott, N. "Consent Is Key to Avoiding Child Privacy Class Actions." Bloomberg Law, June 2020. <https://www.mintz.com/sites/default/files/media/documents/2020-06-15/ChildPrivacyClassActionsECO-54017.pdf>
22. "PAS 1296:2108." Age Check Certification Scheme, 2021. <https://accscheme.wordpress.com/home/pas-12962018/>
23. Kingdom of Belgium. "The identity card for children younger than 12 (Kids-ID). Foreign Affairs, Foreign Trade and Development Cooperation, 2021, https://diplomatie.belgium.be/en/services/services_abroad/identity_card_for_belgians/kids_id
24. Allen, Christopher. "The Path to Self-Sovereign Identity." April 2016. <http://www.lifewithalacrity.com/2016/04/the-path-to-self-sovereign-identity.html>
25. Barrett, Paul M. "Who Moderates the Social Media Giants?" NYU Stern, June 2020. https://issuu.com/nyusterncenterforbusinessandhumanri/docs/nyu_content_moderation_report_final_version?fr=sZWZmZjI1NjI1Ng
26. Facebook. "Form 10-K." U.S. SEC, December 31, 2020. <http://d18rn0p25nwr6d.cloudfront.net/CIK-0001326801/4dd7fa7f-1a51-4ed9-b9df-7f42cc3321eb.pdf>
27. Lin, Liza. "TikTok Owner Bytedance's Annual Revenue Jumps to \$34.3 Billion." Wall Street Journal, June 17, 2021. <https://www.wsj.com/articles/tiktok-owner-bytedances-annual-revenue-jumps-to-34-3-billion-11623903622>
28. Wells, G., Horwitz, J. and Seetharaman, D. "Facebook Knows Instagram is Toxic for Teen Girls, Company Documents Show." Wall Street Journal, September 14, 2021. <https://www.wsj.com/articles/facebook-knows-instagram-is-toxic-for-teen-girls-company-documents-show-11631620739>
29. Mozilla. "Regrets Reporter." 2021: <https://foundation.mozilla.org/en/campaigns/regrets-reporter/>
30. Arthur, Charles. "Naked celebrity hack: security experts focus on iCloud backup theory." The Guardian, September 1, 2014: <https://www.theguardian.com/technology/2014/sep/01/naked-celebrity-hack-icloud-backup-jennifer-lawrence>
31. Bradford Franklin, S. and Nojeim, G. "International coalition calls on Apple to abandon plan to build surveillance capabilities into iPhones, iPads, and other products." Center for Democracy and Technology, August 19, 2021. <https://cdt.org/insights/international-coalition-calls-on-apple-to-abandon-plan-to-build-surveillance-capabilities-into-iphones-ipads-and-other-products/>