



Digital Trust
& Safety Partnership

Age Assurance

Guiding Principles and Best Practices

September 2023



Table of Contents

| | |
|---|----|
| Executive Summary | 2 |
| Introduction | 3 |
| Age Assurance Methods | 5 |
| Method Implementation Choices | 9 |
| Key Challenges & Trade-Offs | 10 |
| Guiding Principles and Best Practices | 12 |
| Definitions | 18 |
| Appendix: Mapping age assurance principles and practices against the DTSP Best Practices Framework | 19 |

Executive Summary

Digital services work to design safe, age-appropriate experiences, including by implementing “age assurance” practices to establish a user’s age.

A variety of age assurance approaches exist, including age verification based on review of identity documents or parental consent; age estimation based on inferences made from user data, physical characteristics, or other measures; and self-declaration by the user.

Every approach to age assurance presents trade-offs. Key challenges include the fact that more accurate methods may depend on collection of new personal data, and thus can be in tension with a service’s privacy commitments to users and legal obligations. Methods may also create inequities among users, unfairly discriminating against certain people, and may not be economically feasible for smaller companies.

There is no one-size-fits all solution in this area. Instead, different services opt for different approaches based on a variety of factors, including but not limited to users of the service, type of service offered, risk calculation, privacy expectations, and economic feasibility.

As part of the Digital Trust & Safety Partnership’s practicing companies’ commitment to its overarching [Trust & Safety Best Practices Framework](#), this document describes a range of current best practices for age assurance.¹ We identify five guiding principles and then note how companies have used these principles to develop example best practices for age assurance. Of course, the specific practices that services use may vary by digital product or feature and evolve with both the challenges faced and advances made in age assurance technologies.

The five guiding principles are:

1. Identify, evaluate and adjust for risks to youth to inform proportionate age assurance methods, as part of implementing safety-by-design.
2. Account for risks to user privacy and data protection as part of development, implementation, and ongoing assessment of age assurance approaches.
3. Ensure assurance approaches are broadly inclusive and accessible to all users, regardless of age, socioeconomic status, race, or other characteristics.
4. Conduct layered enforcement operations to implement age assurance approaches.
5. Ensure that relevant age assurance policies and practices are transparent to the public, and report periodically to the public and other stakeholders regarding actions taken.

¹ This publication was developed by the DTSP Working Group on Age-Related Practices, facilitated by Betsy Masiello and Derek Slater of Proteus Strategies.

Introduction

Young people² increasingly rely on digital services in every aspect of their lives, from educational experiences to social interactions with friends, to engagement with entertainment, games, news and other information, and much more. They also move across and through these services fluidly. Digital services work to design age-appropriate experiences, including by designing specific unique content experiences and safety features. Developing digital services that align with the age of individual users poses various challenges. These challenges include (but are not limited to) the complexity in defining age-appropriate content across culturally diverse communities across the world, accounting for the role of the parent in a teen’s life, and the inherent tension in accurately determining a user’s age while also respecting privacy. There is also a lack of clear universally agreed upon standards upon which to gauge a service’s efforts, although efforts to create such standards are ongoing.³ The practices delineated herein concentrate specifically on principled strategies to assure a user’s age with sufficient accuracy (“age assurance”) amidst the myriad difficulties associated with this particular objective. It does not address other aspects of delivering age appropriate experiences.

While digital services and all relevant stakeholders have aligned around the need to offer age appropriate experiences, they have not aligned on a solution that works for all services. A variety of age assurance approaches exist at present, and every approach presents trade-offs. Enhancing confidence in users’ age carries implications for safeguarding their privacy rights, ensuring their access to information, and preserving their freedom to engage in digital experiences without constraints. These implications hold significance not only for young individuals but for all users of the service. Moreover, the risks associated with youth accessing the service will vary depending on its nature. For instance, a service may be for adults-only, youth-directed, or purpose-built for a mixed audience; services might involve public or private interactions with third parties, or focus instead on engagement with particular types of content. Each service’s risk profile will vary, as do the expectations of its users based on the principles, values, and features of the service. Additionally, different age groups (such as teens) move fluidly across varying services. All of this requires careful consultation with different stakeholders, including with youth themselves, to develop best practices.

To help define an overall framework for what constitutes responsible approaches for digital services, the Digital Trust & Safety Partnership (DTSP) has developed a flexible set of best practices arising from its [Trust & Safety Best Practices Framework](#). As part of fulfilling the Commitments in the Best Practices Framework document, some services may employ age assurance practices as part of an overall approach to addressing risks to young users and developing age appropriate experiences.

² There is no standard definition of young person, youth, child, teenager or adult when discussing age appropriate experiences online. For the purposes of this report, we will rely on the terms “youth” or “young user” and “adult” throughout this paper to represent a general bifurcation between “of-age” and “underage” users, noting that the specific determinant age will vary by jurisdiction around the world. As with the UN Convention on the Rights of the Child, we intend here to refer to “means every human being below the age of eighteen years unless under the law applicable to the child, majority is attained earlier;” although services certainly make further delineations based on a range of factors, including local law.

³ See, e.g., <https://www.iso.org/standard/80399.html>



Digital Trust & Safety Partnership

In this document, we outline the challenges service providers face in age assurance, and then highlight age assurance practices online. We identify five relevant guiding principles for developing and deploying best practices, and then share specific practices for how the framework may be put into action. Of course, the specific practices that services use may vary by digital product or feature and will evolve with both the challenges faced and advances made in age assurance technologies.

While we limit our focus here to age assurance, it is important to recognize that it is only one part of how services may address the challenge of building age appropriate experiences online.



Age Assurance Methods

There are a wide range of practices deployed to gain confidence in a user's age, and an accompanying range of language used to describe them. Age assurance is typically an umbrella term used to describe the full range of practices that services deploy to establish, determine, or confirm a user's age with some level of confidence. These methods may be used by themselves or in combination with one another within a given service.

Age Verification

At one end of the age assurance spectrum, there are a number of methods to verify a user's age that rely on a trusted authority. In these cases, digital service providers rely on third-party, authoritative credentials, or more recently, a government Application Programming Interface (API) to get confidence of a user's age; alternatively, providers rely on parental⁴ consent. All of these methods have limitations and drawbacks that must be taken into consideration when balancing against the benefit of having a highly trusted third party confirm a user's age.

Identity document verification is a mechanism for gathering age information and credentials. It requires access to third-party documentation, typically government-issued IDs or other forms of documentation that otherwise verify an individual's age.⁵ In the context of age verification, service providers typically retain only the information related to the person's age while discarding the actual identity documents provided. Nonetheless, the user must provide more information than what is needed to share their age. Digital service providers aim to address privacy concerns by minimizing the retention of sensitive personal information. In addition to these challenges, identity document verification can present equity challenges as access to government IDs or other necessary documentation to verify their age and identity, such as birth certificates, varies across socioeconomic backgrounds. In some cases, the social or political context makes it risky for certain users to obtain or carry identity documentation.

Implementing identity verification in-house is a significant operational challenge, and presents trade-offs. Much like other trust and safety functions, identity verification requires a combination of machine learning techniques and human review in order to detect fraudulent attempts at verifying an identity and age, and to provide users with adequate appeals processes. In addition to the operational complexity of implementing these reviews, collection of this additional identity data raises issues related to privacy and data protection.

There are a number of third-party vendors that offer identity-verification-as-a-service, which reduces the operational overhead for a digital service provider to stand up their own verification workflow, but still represents an added cost and may also create an opacity around data practices that is challenging for digital service providers to navigate.

⁴ For purposes of this paper, we intend "parent" to refer broadly to the legal guardian of the young user.

⁵ Documentation that might be accepted in an identity verification process may include, for example, a birth certificate, school IDs, utility bills in the individual's name, bank account or credit card verification, or real estate ownership documentation.

More recently, some governments are deploying ID systems⁶ that, in turn, support identity verification. Governments have been successful in common standards for Machine Readable Travel Documents⁷ and an equivalent effort could be made to establish a common format for digital government IDs for use online.

For businesses, these may reduce operational overhead, although that depends greatly on precisely how they're implemented; for instance, if every government develops its own approach rather than relying on a common standard, then it will still be cumbersome for many services to use different digital IDs in a consistent manner. Moreover, reliance on government ID systems necessarily raises concerns around government collection of data and monitoring of user's daily lives online.

Parental consent is a method relying on validation from a parent (or legal guardian) and is in some cases required in existing regulations. Some regulators have relied on parental consent when mandating digital service providers verify that a user is "of-age" to access a given service. For instance, the Children's Online Privacy Protection Act (COPPA) in the United States requires that services get parental consent if they know that a user is under the age of 13. COPPA has had the effect of bifurcating services into those that are directed at youth under 13 and implement parental consent, and those that do not allow users under 13. Similar rules exist in other jurisdictions when it comes to collection and analysis of personal data, albeit with varying ages of consent. For instance, in Europe, the General Data Protection Regulation's (GDPR) Article 8 allows member states to set the age between 13 to 16; Australia's guidelines presume that anyone under 15 does not have capacity to consent; and, in Korea, data privacy law sets the age at 14. Age of consent is also under active debate in many jurisdictions.

When parental consent can be obtained, it provides some assurance that the child's parents are comfortable with the child accessing a given service or feature. Parental consent can be an important tool for involving parents in a young user's online experience, but it does not guarantee that user is the required minimum age, nor does it offer service providers knowledge of the child's actual age. At best it offers the parent an opportunity to provide the user's stated age, rather than the user doing it themselves. Parental consent also comes with the burden of verifying parental responsibility and the relationship between two user accounts; it may be possible to verify one account holder is an adult with a credit card, for example, but that may not verify in any reliable way that the second account is that adult's child. Furthermore, some youth lack parents who can be relied upon to provide verification of the user's age and parents may verify age inaccurately. The Irish Data Protection Commissioner has noted that "there aren't yet many ways of checking parental consent which are accurate, proportionate and that actually work in practice."⁸

⁶ For example, the European Union is developing a European Digital Identity, see https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en; a number of US states have either implemented or are in the process of developing digital IDs or otherwise add a state ID card to a mobile device - see e.g., Colorado's existing digital ID program at <https://mycolorado.state.co.us/> and California's in-development digital ID at <https://cdt.ca.gov/digitalid/>

⁷ See <https://www.icao.int/publications/pages/publication.aspx?docnum=9303>

⁸ See https://www.dataprotection.ie/sites/default/files/uploads/2023-04/DPC_ChildrensData_ParentalConsent.pdf

Age Estimation

As part of a comprehensive risk management system, digital service providers sometimes deploy age estimation techniques to mitigate the risk of underage users gaining access to inappropriate content or experiences. These techniques do not verify a user's age, but rather provide the service provider with some degree of confidence that a user is above or below a given age, based on assessment of some inherent features or behaviors of the user. These techniques may be used on top of asking for age or verifying age via documentation.

Capacity testing is a method that deploys a test to determine a user's capacity for knowledge or analytical thought as a proxy for age. For example, a user may be asked to solve a puzzle or complete a math test. While capacity testing requires very little data to be collected about a user, it has limitations. As a practical matter, it can lack precision, making it unhelpful in determining age-appropriateness except in the most general of senses. It also can generate inequitable outcomes based on users' exposure to educational opportunities and developmental age. Finally, it becomes increasingly ineffective at differentiating ages as a person grows older and gains greater literacy and overall capacity, and users may circumvent the tool by, for instance, getting help from another person.

Age inference is a more advanced estimation method that analyzes data related to a user in order to estimate a user's age. This technique can be deployed on data that is already collected by a digital service provider as a function of using the service. While the data may be produced by the user's interaction with the service itself, it's also possible to incorporate data collected through other services operated by a given provider or third-party sources.

For example, a social networking site may evaluate behavioral signals based on how a user engages with the service, or a media platform may evaluate signals based on the type of content a user is consuming. If a user seeks to consume mostly youth-directed cartoons, then a service might infer that they are a young user themselves; by contrast, if a user primarily seeks out topics like mortgage information, then a service might infer that they are an adult. In turn, if it has a certain level of confidence that the user is young, a service might then prompt the user to produce other evidence of their age; similarly if a user appears to be posing as a young person to gain access to other young users, a service might also prompt the user to produce other evidence of their age.

Because it is necessarily dependent on signals collected during use of a digital service, age inference is a technique that cannot easily be deployed upon account creation. Instead, it is typically deployed as part of a broader age assurance program that incorporates multiple methods, which may be built and operated in-house or by third-party vendors.

Though inference may not depend on additional data collection (at least prior to prompting for further information), it can still raise privacy and data protection concerns. For instance, use of a user's web browsing history might raise concerns when used to do age assurance on a separate service operated by the same provider.

Social vouching is the use of a user's social connections to add additional verification layers to a user's age. This technique relies on asking other users, whose ages and/or identities have already been verified with some degree of confidence, to "vouch" for the stated age of a given user.

Physical age estimation is the use of physical characteristics, most notably facial images, to estimate a user's age. This method relies on collection of data at account creation, or during a subsequent age assurance process after a user is suspected of being underage. Typically this data would not otherwise be collected in the provisioning of the service. This additional data collection heightens the associated privacy risks for the user, and demands strict compliance with data protection standards on the part of the service provider or vendor. In part as a result of these trade-offs, age estimation using facial images is a method that has, to date, been most widely deployed by vendors as opposed to built in-house by service providers.

While this form of age estimation can provide a great degree of accuracy in estimating a user's age, there are limitations associated with demanding this additional collection of personal data from users. Some users may be concerned about the additional collection, and use of a specialized age assurance vendor does not necessarily address this concern. In fact, when a vendor is providing the service for the consumer-facing digital service provider, it may be that much more challenging for users to easily understand how their data will be used and protected. Furthermore, there are potential issues to consider with regard to accuracy and bias, including potential racial bias, of physical age estimation techniques.

Self-declaration

Self-declaration (sometimes referred to as "asserted age assurance") is the practice of asking a user to declare their age. Some self-declaration implementations merely ask a user to affirm they are "of age", asking a yes or no question as to whether a user is above a given age; others are considered "neutral" methods that simply ask the user to state their age. This is the most widely deployed method to date of age assurance, but is also highly susceptible to users providing false information. It is often deployed as one method within a broader risk management system; for example, a service might use self-declaration and subsequently deploy a behavioral age estimation method against all users.

Method Implementation Choices

There are also a range of different ways to implement age assurance methods. For instance:

- **Platforms and devices age assurance:** Some platforms and device manufacturers may implement an age assurance method that results in an age for the user which is then made available to apps and other services on the device.
- **Third-party validation:** Third-party validation, in the context of identity verification and age assurance, refers to the involvement of an external entity or service provider to verify and validate the identity or age of an individual. Instead of relying solely on internal processes or self-declaration, organizations use third-party services or systems to independently confirm the accuracy and legitimacy of the information provided by the individual. For instance, a bank, utility, or credit agencies might provide information to a third-party service that the person is a customer and thus is of a given age. This can also take the form of government validation, such as through digital identity programs, as discussed above.
- **Vendor-provided age assurance:** Some services will hire vendors that perform age assurance methods, using proprietary technology or by aggregating data from a variety of sources. A user's personal data is passed to these vendors, who are then responsible for performing the actual age assurance test, as well as implementing data privacy and security practices.
- **In-house services:** Some service providers will implement age assurance methods entirely in-house, using a combination of proprietary technologies, open source software, and operational enforcement mechanisms. A provider may also make these services available to third parties.

Services might rely on a variety of implementations, and may apply different methods in different geographies based on local requirements.

Key Challenges & Trade-Offs

Several challenges emerge in developing age assurance approaches. Key characteristics that digital service providers look to incorporate in developing these approaches include:

- **Effective:** Having confidence that a user is a given age allows a digital service provider to provide them with an age-appropriate experience, in a way that is accurate and hard to circumvent.
- **Accessible, inclusive and equitable:** Age assurance should not result in inequitable outcomes for a given user, and the complexity of processes should not overly burden users in ways that discourage appropriate use of a service.
- **Privacy-preserving and data-protecting:** Protecting a user's privacy, especially a young user's privacy, demands adherence to key privacy principles including data minimization, as well as implementation of security measures to protect data.
- **Affordable:** Implementation costs must be reasonable and proportionate.
- **Risk-appropriate:** The approaches deployed should be proportionate to the risks associated with underage access to a given service, as well as the risks of inaccurate determinations regarding age.

Incorporating each characteristic comes with trade-offs, and there is no one-size-fits-all solution. Highly accurate age assurance methods may depend on collection of new personal data such as facial imagery or government-issued ID. Some methods that may be economical may have the consequence of creating inequities among the user base. And each service and even feature may present a different risk profile for younger users; for example, features that are designed to facilitate users meeting in real life pose a very different set of risks than services that provide access to different types of content.

Services will face different challenges based on their level of maturity and size, as well as the nature of risks associated with their service. Some age assurance methods may not be economical for smaller companies, and if these methods are not risk-appropriate may thus unduly burden these companies. While larger service providers may find use of these vendors a viable investment to make, particularly if their service is oriented toward higher risk activities, smaller companies may not be able to sustain their business if these costs are imposed. It is therefore important that companies are able to mitigate age-associated risks in different ways, for example by mitigating the risk of youth-centric harms occurring on a service before implementing a costly age verification method.

Similarly, age assurance methods may have tensions with and need to be balanced against a service's privacy commitments to users, which may be entirely appropriate based on the risk profile and purpose of a service. It is important that services retain flexibility to make these decisions for implementing age assurance in the broader context of mitigating age-associated risks, including based on the nature of the service provided.

Instead of a single approach, we acknowledge that appropriate age assurance will vary among services, based on an assessment of the risks and benefits of a given context. A single service may also use different approaches for different aspects or features of the service, taking a multi-layered approach.

The following principles and practices represent how Practicing Companies may put into action age assurance, in the context of a broader trust and safety program and the overall DTSP Best Practices Framework.

Guiding Principles and Best Practices

Principle 1

Identify, evaluate and adjust for risks to youth to inform proportionate age assurance methods, as part of implementing safety-by-design.

Aim: Ensure that companies engage in adequate forethought to the specific risks to youth, and incorporate insights into age assurance methods.

Commentary: Practicing Companies care deeply about the development of age-appropriate experiences that both protect the safety and respect the rights and interests of youth, including their rights to express and access information, and their right to privacy.⁹

To develop a proportionate approach for age assurance, Practicing Companies evaluate risks that are specific to youth – for instance, content that may be inappropriate for younger users, or contact from adults for purposes of grooming. Such risks and impacts are evaluated as products and features are developed, and they may evolve over time after a product or feature is launched. In turn, ongoing assessment and evaluation of risks and impacts is critical. They also assess how age assurance methods may impact both youth and other users.

Examples of specific practices to assess and analyze risks to youth when deploying age assurance include:

- Identify and categorize risks to youth related to content, conduct, contact with third parties, and commercial relationships made possible by a product or feature.
- Develop specialized expertise, insight, and analysis capabilities related to high-impact risks to youth and appropriate mitigation measures.
- Develop and implement frameworks and best practices for risk and impact assessment, which take into account, for instance, the likelihood of youth engaging with a service or feature; the audience the service is designed for; the evolving capacities of young people as they age; foreseeable risk; and the best interests of the young people.
- Consult third parties, including youth and families, in assessing risks and impacts, and selecting age assurance methods.
- Include experts in young people and their safety throughout the product development process, and provide for ongoing feedback both pre-launch and post-launch on risks to youth as well as upholding their rights.

⁹ See UN Convention on the Rights of the Child, UN Office of the High Commissioner, November 1989; General Comment No. 25 (2021) on children's rights in relation to the digital environment, OHCHR, March 2021.



Principle 2

Account for risks to user privacy and data protection as part of development, implementation, and ongoing assessment of age assurance approaches.

Aim: Ensure that age assurance approaches respect data protection and privacy rights, including and especially the privacy rights of young people.

Commentary: Privacy and data protection must be taken into account when evaluating what is most appropriate for a given feature or product. Each method of age assurance has different impacts on user privacy, and different services have different baseline privacy practices that impact user expectations. Along with considering how best to respect privacy when implementing age assurance approaches themselves, a particular challenge for digital service providers is gaining insight into and confidence in the privacy practices of third-party age assurance vendors, which is necessary not only to meet their obligations to users but also to meet regulatory and compliance responsibilities appropriately.

Examples of specific practices to protect user privacy rights include:

- Minimize collection of personal data for age assurance, in a manner proportionate to assessed risk, and design tailored practices regarding the retention, deletion, and use of data.
- Use sensitive data collected solely for age assurance only for that purpose, and delete such data expeditiously once a particular method is complete.
- Analyze personal data exclusively on-device wherever possible, to prevent its transmission to servers (including the digital service provider's servers) that are beyond the individual's control.
- Use age estimation methods on data that is already collected as a function of providing the service, to prevent collection of new personal data.
- Implement age assurance through a vendor such that any new personal data that is collected (e.g., a selfie photo) is only sent to the vendor that performs the age assurance test, not first to the digital service provider.
- Require that vendors apply high privacy and security standards, ensure appropriate third-party review and confirmation that those standards are met.
- Provide transparency to end-users about how their data is collected, used, and retained.
- Complete a Data Protection Impact Assessment before implementing any new age assurance method.
- Where possible, rely on interoperable age assurance solutions that minimize the burden on the user to provide additional information to new services, and mitigate the data protection risks the user must bear.

Principle 3

Ensure assurance approaches are broadly inclusive and accessible to all users, regardless of age, socioeconomic status, race, or other characteristics.

Aim: Ensure age assurance does not unduly impede access to the service, taking into account the disparate impact that age assurance methods may have.

Commentary: Age assurance would be counterproductive if it had the effect of eliminating access to digital services for wide swaths of users for whom those services are appropriate. Deploying an age assurance method that relies on, for example, a government-issued ID may have the effect of discriminating against younger users and users who've had no need to obtain a government-issued ID in their locale. Similarly, certain types of age estimation such as those based on facial images may have greater or lesser confidence levels for populations of a certain ethnicity or age demographic. Practicing Companies take these different effects on inclusivity into account when designing an age assurance approach.

Examples of specific practices to ensure inclusivity and accessibility for all users include:

- To the extent feasible and aligned with legal requirements, select age verification vendors that provide options beyond government-issued IDs, such as birth certificates and school IDs, or a combined unique account identifier and picture of the user, to ensure users without access to government-issued IDs are not discriminated against.
- Provide an accessible, easy-to-navigate appeals mechanism for those users who failed an age estimation test.
- Perform an impact analysis before deploying age estimation techniques to specific populations of users to understand any discriminatory outcomes and mitigate them.
- Provide a process for users to flag that a user is underage, and in concert with an enforcement action give that user in question access to an appeals process to prove their age.
- Conduct reviews of age assurance implementations, including with credentialed third-party vendors or service providers as appropriate.
- Consult with third parties to evaluate the impacts of age assurance methods under consideration.

Principle 4

Conduct layered enforcement operations to implement age assurance approaches.

Aim: Ensure operational capacity exists to prevent users from accessing services or features that are inappropriate to the level of risk, limit access for those who are discovered to have accessed risk-inappropriate services or features, and to provide an appeals process for users whose access is impacted because of age assurance processes.

Commentary: Services define and train an enforcement function within the company that is equipped to implement policies on age appropriate access based on the output of age assurance methods. Based on evaluations of risks, companies invest in a range of technologies and personnel to both select appropriate methods for age assurance and ensure their enforcement on an ongoing basis. These operations are “layered” in the sense that different approaches may be combined and different approaches may apply to different parts of a service, content, or features, based on levels of risk. Services also reevaluate and adjust these operations based on evolving technologies and best practices.

Examples of specific practices of layered enforcement operations to implement age assurance approaches include:

- Set default limits on access to and discovery of the service, certain features, or particular content, subject to in-product notifications about age appropriateness of that content and/or some other form of age assurance.
- Label and, where appropriate, classify services as appropriate for only certain ages, and coordinate with distribution platforms to apply relevant limits for downloads and access by underage users.
- Deploy an age assurance check if a user changes their self-declared age from one that was <18 to an age >18.
- Analyze behavior on a service for all users who self-declare an appropriate age at account creation, identifying users who may have inputted false information and are underage, and proactively putting these users through an additional age assurance test. Behavior indicative of a self-declaration being inaccurate might be, for example, a message celebrating a user’s own 11th birthday or repeated engagement with predominantly youth-directed content.
- Train enforcement teams to identify indicia of a user who has mis-reported their age (for instance, their appearance signals they are actually younger), and a method for triggering further age assurance.
- Allow users to report users who may have mis-reported their age and thus should be limited from the service or certain features.

- Implement technical methods that can help prevent users who have failed an age assurance test and been deemed ineligible from circumventing the controls (e.g., by immediately signing up with a different account).
- Apply new age assurance tests to existing users as a company improves or changes its assurance processes or makes relevant changes to a product or feature.
- Offer family accounts in connection with parental verification by attaching a young person's account to the parent's.
- Empower users or community moderators to set age requirements for engagement with particular content or communities within a service.

Principle 5

Ensure that relevant age assurance policies and practices are transparent to the public, and report periodically to the public and other stakeholders regarding actions taken.

Aim: Ensure users and the public have insight into a service's age assurance methods.

Commentary: Transparency serves a key function in informing the public and educating various stakeholders about a Practicing Company's age assurance practices, while also building trust over time in the sufficiency of an industry's standard of care. At the same time, transparency needs to be considered alongside the risk of users figuring out how to game age assurance systems.

Examples of specific practices to provide transparency about age assurance practices include:

- Explanations why age or birth date is collected as part of account sign-up.
- Implementation of open source age assurance solutions, such that the implementing code can be easily inspected by external stakeholders and experts.
- Providing third-party researchers access to implementation details and data on age assurance effectiveness such that evaluations of a given method's appropriateness can be made by external actors.
- Publishing data that explains the cost burden of different age assurance methods for providers of varying scale.
- Help center articles explaining a service provider's partnership with an age assurance vendor, including what data is shared and an overview of the vendor's data practices.
- Providing quantitative and qualitative information about enforcement of age assurance policies and practices.

Definitions

For purposes of the DTSP Age-Related Best Practices Framework and the accompanying Commentary, the following definitions apply:

Age Assurance: an umbrella term used to describe the full range of practices that services deploy to gain confidence in a user’s likely age, including age verification and age estimation solutions. The word ‘assurance’ refers to the varying levels of certainty that different solutions offer in establishing an age or age range.

Age Estimation: refers to a range of methods aimed to develop a degree of confidence that a user is above or below a given threshold, in order to allow or deny access to age-restricted online content or services.

Age Verification: measures which determine a person’s age to a high level of certainty, usually through relying on third-party documentation or parental consent to verify a user’s age.

Commitment: The actions committed to by Practicing Companies to identify and address Content- and Conduct-Related Risk as part of the overall DTSP Best Practices Framework. This document does not include new Commitments, but rather Practices that may be used to implement the overall Commitments.

Content- and Conduct-Related Risk(s): refers to the possibility of certain illegal, dangerous, or otherwise harmful content or behavior, including risks to human rights, which are prohibited by relevant policies and terms of service.

Contact-Related Risk(s): refers to the possibility of illegal, dangerous, or otherwise harmful contact from a third party to a young user, which are prohibited by relevant policies and terms of service.

Commercial Relationship-Risk(s): refers to the possibility of illegal, dangerous or otherwise harmful commercial and contractual relationships or pressures that a young user may experience in using a service.

(References to “Risks” shall be understood to refer cumulatively to **Content-, Conduct-, Contact-, and Commercial Relationship-Related Risk(s)**.)

Practicing Companies: Providers of products or services that have adopted the Commitments described in the overall DTSP Best Practices Framework.

Trust and Safety: refers to the field and practices that manage challenges related to Content- and Conduct-Related Risk(s), including but not limited to consideration of safety-by-design, Product Governance, risk assessment, detection, and response, quality assurance, and transparency.

Youth and Young User: For the purposes of this report, we will rely on the terms “youth” or “young user” and “adult” throughout this paper to represent a general bifurcation between “of-age” and “underage” users, noting that the specific determinant age will vary by jurisdiction around the world. As with the UN Convention on the Rights of the Child, we intend here to refer to “means every human being below the age of eighteen years unless under the law applicable to the child, majority is attained earlier,” although services certainly make further delineations based on a range of factors, including local law.

Appendix: Mapping age assurance principles and practices against the DTSP Best Practices Framework

| Guiding Principle 1: Identify, evaluate and adjust for risks to youth to inform proportionate age assurance methods, as part of implementing safety-by-design. | | | | | |
|---|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| Identify and categorize risks to youth related to content, conduct, contact with third parties, and commercial relationships made possible by a product or feature. | ✓ | | | ✓ | |
| Develop specialized expertise, insight, and analysis capabilities related to high-impact risks to youth and appropriate mitigation measures. | ✓ | | | ✓ | |
| Develop and implement frameworks and best practices for risk and impact assessment, which take into account, for instance, the likelihood of youth engaging with a service or feature; the audience the service is designed for; the evolving capacities of young people as they age; foreseeable risk; and the best interests of the young people. | ✓ | | | ✓ | |
| Consult third parties, including youth and families, in assessing risks and impacts, and selecting age assurance methods. | ✓ | | | ✓ | |
| Include experts in young people and their safety throughout the product development process, and provide for ongoing feedback both pre-launch and post-launch on risks to youth as well as upholding their rights. | ✓ | | | ✓ | |

Guiding Principle 2: Account for risks to user privacy as part of development, implementation, and ongoing assessment of age assurance approaches.

| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
|---|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| Minimize collection of personal data for age assurance, in a manner proportionate to assessed risk, and design tailored practices regarding the retention, deletion, and use of data. | ✓ | ✓ | ✓ | | |
| Use sensitive data collected solely for age assurance only for that purpose, and delete such data expeditiously once a particular method is complete. | | ✓ | ✓ | | |
| Analyze personal data exclusively on-device wherever possible, to prevent its transmission to servers (including the digital service provider's servers) that are beyond the individual's control. | | ✓ | ✓ | | |
| Use age estimation methods on data that is already collected as a function of providing the service, to prevent collection of new personal data. | | ✓ | ✓ | | |
| Implement age assurance through a vendor such that any new personal data that is collected (e.g., a selfie photo) is only sent to the vendor that performs the age assurance test, not first to the digital service provider. | | ✓ | ✓ | | |
| Require that vendors apply high privacy and security standards, ensure appropriate third-party review and confirmation that those standards are met. | | ✓ | ✓ | | |
| Complete a Data Protection Impact Assessment before implementing any new age assurance method. | ✓ | ✓ | | | |
| Where possible, rely on interoperable age assurance solutions that minimize the burden on the user to provide additional information to new services, and mitigate the data protection risks the user must bear. | | ✓ | ✓ | | |



Guiding Principle 3: Ensure assurance approaches are broadly inclusive and accessible to all users, regardless of age, socioeconomic status, race, or other characteristics.

| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
|---|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| To the extent feasible and aligned with legal requirements, select age verification vendors that provide options beyond government issued IDs, such as birth certificates and school IDs, or a combined unique account identifier and picture of the user, to ensure users without access to government issued IDs are not discriminated against. | | ✓ | ✓ | | |
| Provide an accessible, easy-to-navigate appeals mechanism for those users who failed an age estimation test. | | | ✓ | | |
| Perform an impact analysis before deploying age estimation techniques to specific populations of users to understand any discriminatory outcomes and mitigate them. | ✓ | | | | |
| Provide a process for users to flag that a user is under age, and in concert with an enforcement action give that user in question access to an appeals process to prove their age. | | | ✓ | | |
| Conduct reviews of age assurance implementations, including with credentialed third-party vendors or service providers as appropriate. | | | | ✓ | |
| Consult with third parties to evaluate the impacts of age assurance methods under consideration. | ✓ | ✓ | | | |



Guiding Principle 4: Conduct layered enforcement operations to implement age assurance approaches.

| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
|---|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| Set default limits on access to and discovery of the service, certain features, or particular content, subject to in-product notifications about age appropriateness of that content and/or some other form of age assurance. | | | ✓ | ✓ | |
| Label and, where appropriate, classify services as appropriate for only certain ages, and coordinate with distribution platforms to apply relevant limits for downloads and access by underage users. | | ✓ | ✓ | ✓ | |
| Deploy an age assurance check if a user changes their self-declared age from one that was <18 to an age >18. | | | ✓ | ✓ | |
| Analyze behavior on a service for all users who self-declare an appropriate age at account creation, identifying users who may have inputted false information and are underage, and proactively putting these users through an additional age assurance test. Behavior indicative of a self-declaration being inaccurate might be, for example, a message celebrating a user’s own 11th birthday or repeated engagement with predominantly youth-directed content. | | | ✓ | ✓ | |
| Train enforcement teams to identify indicia of a user who has mis-reported their age (for instance, their appearance signals they are actually younger), and a method for triggering further age assurance. | | | ✓ | ✓ | |
| Allow users to report users who may have mis-reported their age and thus should be limited from the service or certain features. | | | ✓ | ✓ | |



| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
|--|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| Implement technical methods that can help prevent users who have failed an age assurance test and been deemed ineligible from circumventing the controls (e.g., by immediately signing up with a different account). | | | ✓ | ✓ | |
| Apply new age assurance tests to existing users as a company improves or changes its assurance processes or makes relevant changes to a product or feature. | | | ✓ | ✓ | |
| Offer family accounts in connection with parental verification by attaching a child's account to the parent's. | | ✓ | ✓ | ✓ | |
| Empower users or community moderators to set age-requirements for engagement with particular content or communities within a service. | | ✓ | ✓ | ✓ | |



Guiding Principle 5: Ensure that relevant age assurance policies and practices are transparent to the public, and report periodically to the public and other stakeholders regarding actions taken.

| Age Assurance Principles and Practices | DTSP Best Practices Framework | | | | |
|--|-------------------------------|--------------------|---------------------|---------------------|----------------------|
| Best Practices | Product Development | Product Governance | Product Enforcement | Product Improvement | Product Transparency |
| Explanations why age or birth date is collected as part of account sign-up. | | | | | ✓ |
| Implementation of open source age assurance solutions, such that the implementing code can be easily inspected by external stakeholders and experts. | ✓ | ✓ | | | ✓ |
| Providing third party researchers access to implementation details and data on age assurance effectiveness such that evaluations of a given method's appropriateness can be made by external actors. | | | | | ✓ |
| Publishing data that explains the cost burden of different age assurance methods for providers of varying scale. | | | | | ✓ |
| Help center articles explaining a service provider's partnership with an age assurance vendor, including what data is shared and an overview of the vendor's data practices. | | ✓ | | | ✓ |
| Providing quantitative and qualitative information about enforcement of age assurance policies and practices. | | | | | ✓ |