Executive Manager, Industry Regulation and Legal Services
Office of the eSafety Commissioner
PO Box Q500 Queen Victoria Building
NSW 1230

January 20, 2024

The Canadian Centre for Child Protection (C3P) is very supportive of the Australian government and the eSafety Commissioner's work in establishing strong, forward-looking industry standards for the technology industry.

In particular, we are supportive of requirements that impose duties to act responsibly and to proactively block or mitigate harm before it occurs, as opposed to reactive-only frameworks. We wholeheartedly support requirements to deploy widely available and proven technologies to detect and block the distribution of known child sexual abuse material (CSAM) throughout a provider's various services.

These strategies are reasonable and are privacy preserving. They are also consistent with widely accepted technologies already in use for virus and malware detection/blocking across services millions of citizens use daily, including email, file storage and even end to end encrypted messaging services.

Our section-by-section feedback is available below. However, broadly, the core themes of our feedback include:

- The importance of central reporting by providers when breaches occur;
- The need to be more prescriptive in certain areas such as definitions, or response time expectations;
- Paying special attention to providers who allow users accessing their services via a Tor exit relay or VPNs to interact with their service, especially when establishing risk levels and expectations;

While the drafts standards are comprehensive, our view is that the documents omit at least one key overarching obligation we believe is crucial to preventing repeat breaches over the same (or similar) material. We believe providers ought to be required to bolster their deterrence and image detection tools, informed by their previously actioned breaches.

For example, if a user reports a legitimate breach of the provider's term of service due to the discovery of class 1A content, the draft standards should impose an obligation on the provider to ensure metadata about the violative content (e.g. hashes, filenames, etc) be integrated into their subsequent detection and mitigation strategies. Without such an obligation, providers may be satisfied with acting on complaints, but not taking steps to prevent a repeat of the same harm into the future.

C3P refers to this phenomenon as content “recidivism”. In a report published in 2021, we found that close to half of all violative images our content detection tools discovered on a provider’s service had previously been detected on that same service.

More information on media recidivism:
https://content.c3p.ca/pdfs/C3P_ProjectArachnidReport_en.pdf

Thank you for considering our feedback and experience in supporting victims and survivors of online harm as part of your public consultation.

Signy Arnason
Associate Executive Director
Canadian Centre for Child Protection
615 Academy Road
Winnipeg, MB R3N 0E7
Canada
Tel: [redacted]

- - - -

Section by section comments: **Relevant Electronic Services**

See full draft document.

Section	Comment
5(2)	We are concerned about the possibility for too much discretion being given to providers to determine or shape which code/standard their service will be governed by in cases where they may share characteristics that could arguably be covered under different codes. For example, the websites Pinterest and Tumblr are hybrid platforms that merge some elements of social media, while effectively acting as an image-hosting platform. Who will ultimately determine the appropriate code/standard that will apply? We are concerned that providers will strategically favour a specific code/standard that is more likely to alleviate the burden of complying with regulations, rather than be pressed to prioritise the overall objectives of the Act.
6(1)	The definition of “Child sexual exploitation material” does not capture still images taken from frames of a known child sexual abuse video, where, for example, the child is still clothed or semi clothed that is taken during the progression of the sexual abuse. In isolation these may not necessarily meet a criminal law threshold, but they are part of a larger sequence of illegal material. These images are also commonly used to promote the availability of CSAM to users who may recognize the victim or scenes. We recommend broadening the definition to capture this type of content. These images also pose a privacy risk for victims when used in training AI models designed to output abuse imagery with their likeness. The example included for the definition of “end user” may leave the impression that only those with an account (or using the account of someone else) can be considered an end-user, but there are services where one need not have an account to "use" the service, or to be able to access content displayed on the service.

	In the definition of known child sexual abuse material, the reference to the database managed by the National Center for Missing and Exploited Children needs more context. There are several databases managed by that organisation, including ones that are industry driven. Also, consideration should be given to international CSAM databases such as those provided by the Internet Watch Foundation and C3P (Project Arachnid). These expanded examples would have implications for section 15(3) as well.
6(2)	We recommend the addition to (d) the following: "including, in the case of material depicting children, the best interests of the child."
7	As a general statement we believe it is crucial for the government to normalise adherence to basic online safety standards as an expected "cost of doing business" in the technology industry. In the same way that we establish, for example, clear and uniform food-handling rules for food establishments whether a small kiosk or a national fast-food chain, we hope that the threshold for what constitutes non-feasibility due to financial cost — even for small providers — should not be set too low.
8(8)	We are concerned about the possibility for too much discretion being given to providers to determine their own risk profile independently as part of the initial step. A process ought to be established to review and approve risk profiles, and if that is not feasible, then at minimum the risk profile should have to be shared by the provider with the government once completed rather than simply recorded. Section 10 may be where it is anticipated the system will intercept inappropriate self-risk assessment (along with section 33 and 37) but we want to ensure this is flagged as a potential gap.
9(5)	With regards to matters to be taken into account when establishing the methodology for the conduct of a risk assessment, we believe the list of items ought to be more specific and include such matters as: 1. Whether the service permits user traffic via a Tor relay to access or post content on the service; 2. Whether the service permits user traffic via a VPN to access or post content on the service; 3. Whether the service is end-to-end encrypted; 4. Whether the service accepts cryptocurrencies as a form of payment for paid services; 5. The extent to which the service allows for user discoverability (i.e. ease with which other users not previously connected to the user can be discovered, contacted); 6. Whether the service collects sufficient information to verify users (i.e. know-your-client practices). Given the very high-risk factors associated with the above points, it is our view than any of those risk factors ought to automatically merit a Tier 1 risk profile.
10	The risk assessments referred to in this section should be required to be dated, and someone specific like an officer of the corporation or the owner should be required to sign or certify it before a lawyer or someone similar.
12	We believe matters to be taken into account for establishing appropriate action ought to also include patterns of behaviour/track records for the ESP. For example, if the ESP acts appropriately each time a breach occurs, would it truly

	be considered appropriate action if the same breach occurs repeatedly? Should there be consideration for escalated actions? In our experience, several providers will comply reactively with bare minimum compliance requirements, however, will often subsequently not seek to prevent the same harm from reoccurring on their service into the future, because generally speaking, regulations do not require them to.
15(2) + 15(3)	For the purpose of section 15(2)(b), we recommend not limiting the good faith belief to be about a person in Australia. We understand there may be a jurisdictional issue underpinning this clause, but it must be kept in mind that the nature of online crimes is that a person in one country can pose a serious threat to a person in another country, and all countries must do their part in ensuring that such threats make it to a policing agency that can take action. Police in most countries do cooperate and share information as many online investigations start in one country but then end up impacting another. By limiting the obligation to "a person in Australia", when a serious threat is about a person in another country (or it is not clear in which country the threatened person is located), there may be no report and therefore no opportunity for law enforcement in any country to take action. Regarding 15(2) and (3), through lessons learned in Canada, we recommend prescribing a specific entity for both the enforcement entity and the organisation to which reports must be sent. The main reason is to ensure these reports can be tracked, measured, analysed and be used to hold non-reporters accountable when information suggests they have not complied with their reporting obligations. In the U.S., mandatory reporting laws require providers to report to NCMEC. This central reporting ensures policy makers have a complete picture of reporting. In Canada, under the *Act respecting the mandatory reporting of Internet child pornography by persons who provide an Internet service,* a provider of an internet service to the public that has reasonable grounds to believe their services is being or has been used to commit a CSAM offence must report to police (and there are retention requirements to follow), however there is no specific police agency specified. Providers of an internet service to the public who are advised of an IP address or URL where CSAM may be being made available to the public are obligated to report it to the Canadian Centre for Child Protection (C3P). The decentralised reporting structure when it comes to the police notification requirement makes it a challenge to know how often or how compliant providers are in Canada with regards to their reporting obligations. Act: https://laws-lois.justice.gc.ca/eng/acts/I-20.7/page-1.html Regulation: https://laws-lois.justice.gc.ca/eng/regulations/SOR-2011-292/page-1.html#h-775160
16(2) + 16(3) + 17	Without a response time expectation tied to this section, we are concerned long delays, and wilful understaffing could occur. We recommend setting expectations for response times, rather than using terminology such as "as soon as practicable". Some regulations in jurisdictions such as Germany and France

	have determined 24 or 48 hours to be considered an acceptable response time. These are the kinds of regulation that should prompt the use of proactive detection and other design to prevent the reports in the first place. Germany: https://www.loc.gov/item/global-legal-monitor/2021-07-06/germany-network-enforcement-act-amended-to-better-fight-online-hate-speech/ France: https://www.assemblee-nationale.fr/dyn/15/textes/l15t0388_texte-adopte-seance#B2298414350
17(2)(c)	We recommend defining what constitutes "remove material". Does it mean, to make inaccessible to users? To delete permanently from servers? To preserve for record keeping or perhaps industry collaboration on flagging mechanisms?
17(4)	This section does not specify preservation requirements.
18(3)	What constitutes "sufficient" [trust and safety] personnel? We recommend establishing guidelines and/or factors that must be taken into consideration, including, but not limited to: ● User volume; ● User engagement (e.g. volume of user-generated content contributed); ● Risk level; ● Staffing commensurate with abuse report response time response.
19(4)(f)	We recommend expanding on what constitutes "private by default". Our understanding/expectations of privacy are somewhat context dependent. What an end user understands as "private by default" is almost certainly not shared by many providers whose business models often depend on users behaving in more "open" ways. Examples of private by default ought to include settings such as, but not limited to: ● No geo-location; ● No public friend/network listing; ● No direct-message capability without being an existing contact; ● Strong limits discoverability through search functions for non-contacts;
19(4)	We recommend adding a requirement that providers give users the ability to restrict interactions with users outside of a geo-gated location set by the user. For example, a young-child account could be set to limit the user's network to a narrow geographic area such as a city or state/province (as opposed to it being open to the global internet).
19(6)	We recommend implementing user verification requirements for dating services to minimise the risk of adult end-users connecting with the children and to prevent children from accessing the service.
19(7)	We believe that closed communication relevant electronic services that are end-to-end encrypted represent a significantly increased risk of breaches of the Act. It is our view that there ought to be enhanced safety features required when the messaging service uses end-to-end encryption, for example preventing illegal material from entering the end-to-end pipeline. In addition, the service should be required to take some measure to verify the validity of the phone number, email address or other identifier supplied by the end user.

19(9)	We recommend the addition of a condition that if users of the service are or may be children, the information must use language and terminology that is age-appropriate for that audience.
20(2)	We recommend being more specific about the expectations surrounding the use of hashing algorithms used for detection, as not all are equally effective. For example, detecting known content using exact cryptographic hashes (e.g. sha1 or md5) will miss images that have been resized, reformatted, or reuploaded from a screenshot. These may be detected using perceptual (i.e. approximate or fuzzy) hashes (e.g. PhotoDNA or PDQ).
20(3), (5) and (6)	While these sections are focused on technical feasibility, it is also worth noting in terms of financial feasibility that Project Arachnid, which is operated by the Canadian Centre for Child Protection, offers a no-cost known CSAM detection API that is currently used by several ESPs internationally. The eSafety Commissioner may want to familiarise itself with available known-CSAM detection tools and recommend them as appropriate to ESPs. More info: https://www.projectarachnid.ca/en/#shield. Also, some parameters, guidelines or examples around what may be legitimately considered not technically feasible would be helpful.
20(4)	Based on the issuance of more than 35 million CSAM removal notices since 2017, our data shows significantly more than half of all notices are typically acted on within 24 hours, with the remainder often taking days, weeks or even months before being actioned. We recommend adopting removal time standards, rather than use language such as “as soon practicable”. Respectfully, ESPs likely do not share the same sense of urgency as the victims depicted in the images.
20(6)	In cases where the deployment of tools for detecting CSAM is deemed to be not technically feasible, we strongly recommend establishing that “appropriate alternative action” mandatorily include the blocking of Tor and VPN traffic, as these are highly correlated risk factor for CSAM upload and distribution as observed by C3P, but also noted by Dr. Neal Krawetz who holds a PhD in Computer Science from Texas A&M University. Source: https://www.hackerfactor.com/blog/index.php?/archives/720-This-is-what-a-TOR-supporter-looks-like.html
22	As noted previously, Tor and VPN user traffic are highly correlated with illegal or higher-risk activities by users. We recommend the blocking of user traffic, or at least the blocking of user uploads, from these two sources as one of the key deterrence strategies, as these are the primary methods used for masking one’s true IP address/identity.
23(1)(b)	We believe the threshold of 1 million **average** monthly active end-users to be far too high and lacking clarity. For example, does an end-user not registered with an account count as an end user? Perhaps the metric should be “an average of 1 million unique visitors”? Also, given that a Tier 1 service by definition is considered a high risk service for the solicitation, access and distribution of class 1A material or class 1B material, the 1 million threshold is far too high of a threshold before the requirements of section 23 are triggered, particularly considering the long term harm and damage to victims when this type of material is not promptly addressed.

26(2)	We recommended providing the exact text (or a series of options deemed suitable) required to be displayed to ensure completeness and accuracy of information. It is also recommended that certain display parameters be prescribed such as minimum font size and font colour.
27(2)	We recommend considering working with industry to build toward a standardised user reporting mechanism, menu and language. The eSafety Commissioner may consider reviewing the study titled "A comparative analysis of platform reporting flows" by researcher Alex Leavitt (UC Berkeley/Roblox) for a better understanding of this issue. For information on gaps in online reporting flows for select providers as of 2020, please review C3P's report here: https://content.c3p.ca/pdfs/C3P_ReviewingCSAMMaterialReporting_en.pdf
28	While we agree that the tools, processes or technologies must be available "in service", consideration must be given to mandating a mechanism for individuals who are not active users or formal subscribers to the service to make a report or complaint. See comment in 29(1) below.
29(1)	This list should also indicate that a tool, process or technology for reporting breaches or making a complaint should not require a user to have a registered account, especially in cases where the violative material is visible without an account.
32	Same comment as 28 and 29(1)(b).
35	While it is somewhat implied that the provider would be assessing risk when contemplating adding a new feature or function, there should actually be an express obligation for the provider to revisit the risk assessment and risk profile as part of the process . The result of their assessment ought to be part of the information provided when notifying the Commissioner, and where the provider believes it will not significantly increase the risk, they should be required to document why and on what basis that determination was made, with the records being dated and certified by an officer/owner of the provider at the time the decision was made.
37(3)	This section should also include an obligation to provide the number of reports made to an enforcement authority (class 1A) and to an organisation (class 1B).
40	An additional requirement (4) should require the provider to retain a record of all complaints, whether they were handled pursuant to 40(2) or 40(3), and such information should be required as part of the reporting requirements in section 37.

Section by section comments: **Designated Internet Services**

See full draft document.

Section	**Comment**
5(2)	We are concerned about the possibility for too much discretion being given to providers to determine or shape which code/standard their service will be governed by in cases where they may share characteristics that could arguably be covered under different codes. For example, the websites Pinterest and Tumblr are hybrid platforms that merge some elements of social media, while effectively acting as an image-hosting platform. Who will ultimately determine the appropriate code/standard that will apply? We are concerned

	that providers will strategically favour a specific code/standard that is more likely to alleviate the burden of complying with regulations, rather than be pressed to prioritise the overall objectives of the Act.
6(1)	The definition of “Child sexual exploitation material” does not capture still images taken from frames of a known child sexual abuse video, where, for example, the child is still clothed or semi clothed that is taken during the progression of the sexual abuse. In isolation these may not necessarily meet a criminal law threshold, but they are part of a larger sequence of illegal material. These images are also commonly used on social media to promote the availability of CSAM to users who may recognize the victim or scenes. We recommend broadening the definition to capture this type of content. These images also pose a risk for use in training AI models designed to output abuse imagery in the likeness of known victims of CSAM. The example included for the definition of “end user” may leave the impression that only those with an account (or using the account of someone else) can be considered an end-user, but there are services where one need not have an account to "use" the service, or to be able to access content displayed on the service. In the definition of known child sexual abuse material, the reference to the database managed by the National Center for Missing and Exploited Children needs more context. There are several databases managed by that organisation, including ones that are industry driven. Also, consideration should be given to referencing the database managed by Interpol and used by Project Arachnid. These expanded examples would have implications for section 15(3) as well. We recommend expanding the definition of “machine learning model platform service” to include services that collect, label and otherwise process media to create pools of images to be used by services that develop AI models. For example, the company behind the public LAION-5B dataset is used by companies such as Stability AI to create AI models. It is important to ensure training data providers are explicitly covered by this definition as the inputs used for training models affect the outputs AI models generate. Recently, researchers from Stanford University discovered that the LAION-5B dataset included at least 1,679 CSAM images that had been collected from social media posts and popular adult websites. This report can be found here: https://purl.stanford.edu/kh752sm9123
6(4)	We recommend the addition to (d) the following: “including, in the case of material depicting children, the best interests of the child.”
7	As a general statement we believe it is crucial for the government to normalise adherence to basic online safety standards as an expected “cost of doing business” in the technology industry. In the same way that we establish, for example, clear and uniform food-handling rules for food establishments

	whether a small kiosk or a national fast-food chain, we hope that the threshold for what constitutes non-feasibility due to financial cost — even for small providers — should not be set too low.
8(8)	We are concerned about the possibility for too much discretion being given to providers to determine their own risk profile independently as part of the initial step. A process ought to be established to review and approve risk profiles, and if that is not feasible, then at minimum the risk profile should have to be shared by the provider with the government once completed rather than simply recorded. Section 10 may be where the system intercepts inappropriate self-risk assessment (along with section 34 and 38) but we want to ensure this is flagged as a potential gap.
9(5)	With regards to matters to be taken into account when establishing the methodology for the conduct of a risk assessment, we believe the list of items ought to be more specific (less broad) and include such matters as: 1. Whether the service permits user traffic via a Tor relay to access the service; 2. Whether the service permits user traffic via a VPN to access the service; 3. Whether the service is end-to-end encrypted; 4. Whether the service accepts cryptocurrencies as a form of payment for paid services; 5. Whether the service collects sufficient information to verify users (i.e. know-your-client practices). Given the very high-risk factors associated with the above points, it is our view that any of those items ought to automatically merit a Tier 1 risk profile.
10	The risk assessments referred to in this section should be required to be dated, and someone specific like an officer of the corporation or the owner should be required to sign or certify it before a lawyer or someone similar.
12	We believe matters to be taken into account for establishing appropriate action ought to also include patterns of behaviour/track records for the provider. For example, if the provider acts appropriately each time a breach occurs, would it truly be considered appropriate action if the same breach occurs repeatedly? Should there be consideration for escalated actions? In our experience, several providers will comply reactively with bare minimum compliance requirements, however, will often subsequently not seek to prevent the same harm from reoccurring on their service into the future, because generally speaking, regulations do not require them to.
15(2) + 15(3)	For the purpose of section 15(2)(b), we recommend not limiting the good faith belief to be about a person in Australia. We understand there may be a jurisdictional issue underpinning this clause, but it must be kept in mind that the nature of online crimes is that a person in one country can pose a serious threat to a person in another country, and all countries must do their part in ensuring that such threats make it to a policing agency that can take action. Police in most countries do cooperate and share information as many online investigations start in one country but then end up impacting another. By

	limiting the obligation to "a person in Australia", when a serious threat is about a person in another country (or it is not clear which country the person threatened is in), there may be no report and therefore no opportunity for law enforcement in any country to take action. Regarding 15(2) and (3), through lessons learned in Canada, we recommend prescribing a specific entity for both the enforcement entity and the organisation to which reports must be sent. The main reason is to ensure these reports can be tracked, measured, analysed and be used to hold non-reporters accountable when information suggests they have not complied with their reporting obligations. In the U.S., mandatory reporting laws require ESPs to report to NCMEC. This central reporting ensures policy makers have a complete picture of reporting. In Canada, under the *Act respecting the mandatory reporting of Internet child pornography by persons who provide an Internet service,* a provider of an internet service to the public that has reasonable grounds to believe their services is being or has been used to commit a CSAM offence must report to police (and there are retention requirements to follow), however there is no specific police agency specified. Providers of an internet service to the public who are advised of an IP address or URL where CSAM may be being made available to the public are obligated to report it to the Canadian Centre for Child Protection (C3P). The decentralised reporting structure when it comes to the police notification requirement makes it a challenge to know how often or how compliant ESPs are in Canada with regards to their reporting obligations. Act: https://laws-lois.justice.gc.ca/eng/acts/I-20.7/page-1.html Regulation: https://laws-lois.justice.gc.ca/eng/regulations/SOR-2011-292/page-1.html#h-775160
16(2) + 16(3) + 17	Without a response time expectation tied to this section, we are concerned long delays, and wilful understaffing could occur. We recommend setting expectations for response times, rather than use terminology such as "as soon as practicable". Some regulations in jurisdictions such as Germany and France have determined 24 or 48 hours to be considered an acceptable response time. Germany: https://www.loc.gov/item/global-legal-monitor/2021-07-06/germany-network-enforcement-act-amended-to-better-fight-online-hate-speech/ France: https://www.assemblee-nationale.fr/dyn/15/textes/l15t0388_texte-adopte-seance#B2298414350
17(2)(c)	We recommend defining what constitutes "remove material". Does it mean, to make inaccessible to users? To delete permanently from servers? To

	preserve for record keeping or perhaps industry collaboration on flagging mechanisms?
17(2)	It would be prudent to add an obligation to preserve information for some specified period of time.
19	Considering that extreme crime and violence material and Class 1B material includes bestiality content and depictions of sexual violence, some more detailed information about requirements other than simply "appropriate action" is warranted for inclusion in section 19.
20(3)	What constitutes "sufficient personnel? We recommend establishing guidelines based on, but not limited to: ● User volume; ● User engagement (e.g. volume of user-generated content contributed); ● Risk level; ● Staffing commensurate with abuse report response time response.
21(2)	We recommend being more specific about the expectations surrounding the use of hashing algorithms used for detection, as not all are equally effective. For example, detecting known content using exact cryptographic hashes (e.g. sha1 or md5) will miss images that have been resized, reformatted, or reuploaded from a screenshot. These may be detected using perceptual (i.e. approximate or fuzzy) hashes (e.g. PhotoDNA or PDQ).
21(3), (5) and (6)	While these sections are focused on technical feasibility, it is worth noting in terms of financial feasibility that Project Arachnid, which is operated by the Canadian Centre for Child Protection, offers a no-cost known CSAM detection API that is currently used by several ESPs internationally. The eSafety Commissioner may want to familiarise itself with available known-CSAM detection tools and recommend them as appropriate to ESPs. More info: https://www.projectarachnid.ca/en/#shield Also, some parameters, guidelines or even examples around what may be legitimately considered not technically feasible would be helpful.
21(4)	Based on the issuance of more than 35 million CSAM removal notices since 2017, our data shows significantly more than half of all notices are typically acted on within 24 hours, with the remainder often taking days, weeks or even months before being actioned. We recommend adopting removal time standards, rather than use language such as "as soon practicable". Respectfully, ESPs likely do not share the same sense of urgency as the victims depicted in the images.
21(5) + 21(6)	In cases where the deployment of tools for detecting CSAM is deemed to be not technically feasible, we strongly recommend establishing that "appropriate alternative action" mandatorily include the blocking of Tor and VPN traffic, as these are highly correlated risk factor for CSAM upload and distribution as observed by C3P, but also noted by Dr. Neal Krawetz, who holds a PhD in Computer Science from Texas A&M University.

	Source: https://www.hackerfactor.com/blog/index.php?/archives/720-This-is-what-a-TOR-supporter-looks-like.html
21(9)(a) + 21(9)(b)	Image detection is one form of prevention, however other highly effective strategies also exist. As noted previously, Tor and VPN user traffic are highly correlated with illegal or higher-risk activities by users. We recommend the blocking of user traffic from these two sources as one of the key deterrence strategies, as these are the primary methods used for masking one's true IP address/identity. Generally speaking individuals will not exploit a provider's service for the distribution of CSAM if they cannot mask their IP address. Our experience in reviewing offender strategies is that they actively seek out providers that permit the use of Tor and VPNs to access their services.
23	See comment for 21(9)(a) + 21(9)(b)
24(1)(a)	We believe the threshold of 1 million **average** monthly users to be far too high and lacking clarity. For example, does a user not registered with an account count as a user? Perhaps the metric should be "an average of 1 million unique visitors"? A Tier 1 service by definition is considered a high-risk service for CSAM distribution so a 1 million threshold is far too high before the obligations in 24 are triggered. This should just be considered a cost of doing business once you are in that Tier 1 category.
24(1)(b)	We believe that given the elevated risks associated with end-user managed hosting services for the distribution of CSAM and illegal material, a threshold of 500,000 monthly active Australian end-users is far too great. Looking at the most common file-hosting services used to distribute CSAM based on Project Arachnid data, it is unlikely that any of these services — most of which would not be recognized by average citizens — would ever meet that threshold. And yet, they are also among the most damaging for CSAM distribution. We recommend considering either a lower user threshold or a threshold based on total bandwidth.
25(4)	We recommend strengthening these obligations with age or user verification guidelines. Service providers have massive commercial incentive to attract new users and build brand loyalty with users as soon as they begin to engage in digital spaces. Special attention should be paid to services that must gatekeep users based on age to ensure there are no gaps in this obligation. There has been recent media coverage, for example, reporting that Meta seldomly took steps to prevent underage users from accessing their services, despite having knowledge of their age. Source: https://www.nytimes.com/2023/11/25/technology/instagram-meta-children-privacy.html
28(2)	We recommended providing the exact text (or a series of options deemed suitable) required to be displayed to ensure completeness and accuracy of

	information. It is also recommended that certain display parameters be prescribed such as minimum font size and font colour.
29(2)	We recommend considering working with industry to build toward a standardised user reporting mechanism, menu and language. The eSafety Commissioner may consider reviewing the study titled "A comparative analysis of platform reporting flows" by researcher Alex Leavitt (UC Berkeley/Roblox) for a better understanding of this issue. For information on gaps in online reporting flows for select providers as of 2020, please review C3P research here: https://content.c3p.ca/pdfs/C3P_ReviewingCSAMMaterialReporting_en.pdf This list should also indicate that a tool, process or technology for reporting breaches or making a complaint should not require a user to have a registered account, especially in cases where the violative material is visible without an account.
30(2)(a)+30(2)(b)	Recommend establishing response time expectations for what constitutes "promptly". This comment applies to all references throughout this document that references response times.
31(2)	More detail surrounding what constitutes a "response" is needed.
36	While it is somewhat implied that the provider would be assessing risk when contemplating adding a new feature or function, there should actually be an express obligation for the provider to revisit the risk assessment and risk profile as part of the process. The result of their assessment ought to be part of the information provided when notifying the Commissioner, and where the provider believes it will not significantly increase the risk, they should be required to document why and on what basis that determination was made, with the records being dated and certified by an officer/owner of the provider at the time the decision was made.
38(8)	Section (a) should not just be child sexual exploitation material identified "by" the provider, but also "reported to" the provider by end users. This section should also include an obligation to provide the number of reports made by the provider to an enforcement authority (class 1A) and to an organisation (class 1B).
40	An additional requirement (4) should require the provider to retain a record of all complaints, whether they were handled pursuant to 40(2) or 40(3), and such information should be required as part of the reporting requirements in section 38.