Your response

Question	Your response
Question 1: How do you measure the number of users on your service?	N/A
Question 2: If your service comprises a part on which user-generated content is present and a part on which such content is not present, are you able to distinguish between users of these different parts of the service? If so, how do you make that distinction (including over a given period of time)?	N/A
Question 3: Do you measure different segments of users on your service? • **Do you segment user measurement by different parts of your service? For example, by website vs app, by product, business unit.** • **Do you segment user measurement into different types of users? For example: creators, accounts holders, active users.** • **How much flexibility does your user measurement system have to define new or custom segments?**	N/A
Question 4: Do you publish any information about the number of users on your service?	N/A

Question	Your response
Question 5: Do you contribute any user number data to external sources/databases, or help industry measurements systems by tagging or sharing user measurement data? If not, what prevents you from doing so?	N/A
Question 6: Do you have evidence of functionalities that may affect how easily, quickly and widely content is disseminated on U2U services? • **Are there particular functionalities that enable content to be disseminated easily on U2U services?** • **Are there particular functionalities that enable content to be disseminated quickly on U2U services?** • **Are there particular functionalities that enable content to be disseminated widely on U2U services?** • **Are there particular functionalities that prevent content from being easily, quickly and widely disseminated on U2U services?**	The following responses will cover the spread of terrorist content on user to user (U2U) and search services, based on the expertise of Tech Against Terrorism. You can access all of our reports via our Public Resources page on our Knowledge Sharing Platform. There are several functionalities that enable the easy, quick and wide spread of content. The ability to forward or repost content allows users to share other users' content and increase its reach. This is maximised by the ability to share posts across different platforms. The ability to livestream and share recordings of livestreams also increases the reach and spread of content allowing large numbers of users to view and comment on content at the same time. For a detailed summary of functionalities that enable the sharing of terrorist content, see the response to question 9. The use of automated content moderation systems is one example of a functionality which limits the spread of content. These include automated detection systems, such as hashing, and hiding and deprioritizing schemes, such as filtered searching and Geo-blocking. Tech Against Terrorism's Terrorist Content Analytics Platform (TCAP) can also support with the targeted alerting of terrorist content on a platform, speeding up content moderation workflows and limiting the spread of content. As per our upcoming TCAP Transparency Report, our open-source intelligence experts submitted a total of 18,995 URLs containing terrorist content and the TCAP sent 10,174 alerts to 57 tech companies, 82% of which is now offline. In total, 150 tech companies are registered and able to receive alerts as soon as we detect terrorist content on their platforms. Please get in touch if you have any additional questions regarding the TCAP. There are also some platform features that limit the dissemination of content. Easily accessible on-

Question	Your response
	platform user reporting mechanisms support the moderation of content which violates a platforms' terms of service. Other platform feature limitations can also limit the spread of content. Restrictive features such as limiting platform group sizes, only allowing users to access and post content via an account, and only allowing users to share content in-platform, can reduce the accessibility and quick spread of content.
Question 7: Do you have evidence relating to the relationship between user numbers, functionalities and how easily, quickly and widely content is disseminated on U2U services?	N/A
Question 8: Do you have evidence of other objective and measurable factors or characteristics that may be relevant to category 1 threshold conditions?	Platform resources: Platform resources should be a characteristic considered for Category 1, 2A an 2B thresholds. Platform resources refer to both the size of the platforms workforce (number of employees) for human content moderation and its access to technical resources for automated content moderation. Under platform resources, Tech Against Terrorism recommends that Ofcom considers the size of platforms' workforce when developing the thresholds for all categories. Recent analysis conducted by Tech Against Terrorism examined the relationship between tech platform size by number of employees [1] and the amount of terrorist content identified and alerted via the TCAP. It is important to note that information about the size of a tech platforms' trust and safety team is not usually publicly available, therefore this analysis focused on platforms' total number of employees, where this information was available.

[1] Tech Against Terrorism's classification of platform size by number of employees: Very Early: 0-10 employees; Early: 11- 49 employees; Mid: 50 - 249 employees; Enterprise: 250 + employees.

Question	Your response
	We found that platforms with fewer employees are on average the most heavily exploited by terrorist actors, based on the amount of terrorist content collected via the TCAP. Platforms with the smallest workforces (very-early stage and early stage) also had on average the lowest removal rate of URLs alerted to them via the TCAP. The lower average removal rates of terrorist content for platforms with fewer employees is unsurprising given the direct correlation with fewer resources and capacity for human moderation. However, insights from Tech Against Terrorism's Mentorship Programme have indicated that the size of a platforms workforce does not always equate to the size of a platforms trust and safety team. For example, one platform with a small workforce (11-49 employees) was found to have dedicated approximately half of its workforce to trust and safety. Another platform with a much larger workforce (250+ employees) had a significantly smaller trust and safety team. As such, consideration of a platforms workforce should specifically focus on the number of allocated trust and safety employees. TAT remains available to share more information about platforms it works with. Under platform resources, Tech Against Terrorism recommends that Ofcom also considers platforms' capacity to implement technical tools. While some platforms with limited resources can access certain tools, Ofcom should look beyond the availability of technical tools to a platform when developing its thresholds for categorisation. While many small platforms may have access to an automated tool it will still need to effectively implement it into its content moderation workflow. For instance, the possibility to access a database of hashed terrorist content does not mean that it can be easily implemented into content moderation workflows as to workflows, as to utilise the database platforms would need to have already hashed content on their services. Please get in touch if you have any questions concerning the inclusion of platform resources as an additional threshold characteristic.

Question	Your response
	Risk assessments and the prevalence of terrorist content: Tech Against Terrorism recommends that Ofcom considers platforms' illegal content risk assessments when developing thresholds for all categories to better understand the prevalence of terrorist content on each platform in scope and to inform the overall boundaries of categories of regulated services.
Question 9: Do you have evidence of factors that may affect how content that is illegal or harmful to children is disseminated on U2U services? • **Are there particular functionalities that play a key role in enabling content that is illegal or harmful to children to be disseminated on U2U services?** • **Do you have evidence relating to the relationship between user numbers, functionalities and how content that is illegal or harmful to children is disseminated on U2U services?**	Platform functionalities: A recent analysis of the relationship between platform type and the amount of terrorist content identified and alerted via the TCAP found that the most at-risk platforms were file sharing, archiving sites and messaging platforms, for different reasons. Over half of the platforms on which we identified terrorist content were file-sharing in their functionality (106 out of 187 platforms). File-sharing sites are used by terrorist actors to host content such as text, images, and videos, which can then be accessed through aggregated outlinks on beacon platforms. A large volume of terrorist content was also identified on a small number of archiving and pasting sites. Archiving and pasting platforms are likely to be popular with terrorist actors due to their multifunctional nature. Archiving sites are used to aggregate outlinks to content stores as well as providing access to historic content stores following removal by content moderators. Meanwhile, pasting sites are used to store content and aggregate information, such as lists of URLs which link to further content and are not immediately identifiable as necessarily terrorist in nature. Both can be used to evade content moderation. Messaging platforms were found to host the most far-right terrorist content and the average removal rate for messaging platforms was one of the lowest of all platform types. On terrorist exploitation of features, terrorist and violent extremist actors search for three main characteristics in a tech platform: security, stability, and audience reach. Tech Against Terrorism has added an additional fourth characteristic: usability. The following is outlined in our 'Terrorist Use of E2EE' report: • Security: Enhanced security and privacy features. • Stability: Limited capacity, or in some cases a limited willingness, to remove content or

Question	Your response
	ban accounts, resulting in a more stable online presence for terrorist accounts or groups. Open-source software is also appealing as it offers terrorists the opportunity to develop their own platforms. • Audience reach: Features that increase their ability to reach a wide audience, such as large-capacity groups or channels with unlimited audience. • Usability: Encompassing the different features that make an app user-friendly, usability includes those that make the platform attractive to a wider audience and prove useful for organisational and idea-sharing purposes. Tech Against Terrorism also outlined tech platform features commonly exploited by terrorists for different purposes, based on their known threat - how the feature has been exploited in the past and potential threat – how the feature may be exploited in the future. Terrorist and violent extremists can utilise features such as the ability to create and post content without the need to create an account. Similarly, the ability to post content via anonymous accounts which are not tied to email addresses, phone numbers or any other PII, make platforms more attractive for exploitation. Other highly exploited features are for content hosting. File mirroring is a popular propaganda dissemination technique of uploading the same content across a range of platforms, then providing outlinks to the various platforms on a centralised channel. This allows users to open all the URLs in sequence and find one that is still online. A platform's password protection feature is also often used to ensure that content is only accessible to specific users. Search functions on a platform also make it simple for terrorist and violent extremists to identify content and users. They are also able to circumvent content moderation efforts by search for deliberately misspelled code words which pertain to specific types of terrorist content. Targeted communication features such as End-to-end Encryption and private, password-protected or vetted channels, are considered high risk as they ensure enhanced security. They also reduce the

Question	Your response
	likelihood of content moderation, user reporting. Meanwhile, features that allow one-way messaging to an audience, or Beacon channels, allow terrorists and violent extremist communities to post leadership and organisational messages. A type of potentially high-risk feature is in-app content editing. In-app content editing allows terrorist and violent extremists to manipulate content to avoid automated content moderation. Platforms which allow easy content editing are likely to be attractive content hosting options for terrorist and violent extremist entities, and especially for hosting content that is regularly removed, such as officially branded content or attacker-produced crisis content. If content is not identified and moderated at the point of editing, then it is likely that the now-edited content will circulate on the target platform more successful. As content moderation efforts become more sophisticated in the detection and removal content, terrorist and violent extremist entities will almost certainly seek to exploit in-app content editing functions in the future. Please get in touch if you have any additional questions concerning the exploitation of tech platform features. User numbers: Tech Against Terrorism's recent analysis of the relationships between a platforms average user base[2] and the amount of terrorist content identified and alerted via the TCAP, found that that small platforms (between 100,000 and 10 million average monthly users) were the most at risk of terrorist exploitation, based on the volume of content identified on their services. Small platforms also averaged a lower removal rate of alerted terrorist content than Large and Medium-sized platforms. However, an important caveat is that the size of a platforms user-base does not equate to the resources it has, such as the size of its workforce and its capacity for technical tools to address the prevalence of terrorist content.

[2] Tech Against Terrorism's classification of platform size by average user base: Micro - < 100,000 average users per month; Small: > 100,000 average users per monthly; Medium: > 10 million average user per month; Large: > 1 billion average users per month

Question	Your response
	Please get in touch if you have any additional questions concerning platform size – both user numbers and platform resources.
Question 10: Do you have evidence of other objective and measurable characteristics that may be relevant to category 2B threshold conditions?	Please refer to the response to Question 8 on tech platform resources.
Question 11: Do you have evidence of matters that affect the prevalence of content that (once the Bill takes effect) will count as search content that is illegal or harmful to children on particular search services or types of search service? For example, prevalence could refer to the proportion of content surfaced against each search term 16 that is illegal or harmful to children, but we welcome suggestions on additional definitions. • **Do you have evidence relating to the measurement of the prevalence of content that is illegal or harmful to children on search services?**	Tech Against Terrorism recommends that easy access to terrorist and violent content via indexed search results should be considered when developing thresholds for category 2A search services. Terrorist content is easily discoverable through search services. Search engines risk facilitating the discovery or promotion of terrorist and violent extremist networks and their material, on terrorist operated websites (TOWs) different platforms, on the indexed web via search results. Indexed search results almost certainly assist in the discoverability of TOWs, especially when the indexed result is available within the first page of results. Users are also able find content on specific online platforms, including social media and video-sharing services, despite the content being "hidden" via on-platform search. For instance, certain social media or content sharing platforms may have blocked specific keywords searches and hashtags or may have acted on certain violent extremist content and accounts by hiding it from users on the platform. An indexed search result can also help a user find terrorist content that has been removed through accessible metadata, such as the username of the original poster, keywords and phraseology. Beyond searching for content and material produced by terrorist and violent extremist actors, Branislav Todorovic and Darko Trifunovic notably highlighted the use of online mapping tools and of search engines to "plan attacks, monitor news, and identify potential recruits."

Question	Your response
	Language disparities and lack of coordination in search engines' content moderation practices make terrorist content in non-English languages (such as Arabic) easily accessible via search services on the surface web. This means that even if terrorist content is inaccessible through one search engine, they will still be discoverable through others. Tech Against Terrorism has previously outlined the exploitation of search engines in a state of play report, however our ongoing disruption and analysis of terrorist operated websites indicates that more analysis on the role of search services in the accessibility of terrorist content online, specifically terrorist operated websites is needed. For more information on terrorist operated websites, read our threat report and mitigation strategies report. URL shortening services can also be used to support the exploitation of search engines. URL services can play a significant role in facilitating the promotion and diffusion of terrorist and violent extremist networks and material by acting as key nodes between different platforms. URL services can be crucial bridges in online terrorist and violent extremist networks, as they can redirect supporters to different platforms making-up an online network and to where content is hosted. Please get in touch if you have any additional questions concerning the accessibility of terrorist content via search services.
Question 12: Do you have evidence relating to the number of users on search services and the level of risk of harm to individuals from search content that is illegal or harmful to children? • **Do you have evidence regarding the relationship between user numbers on search services and the prevalence of search content that is illegal or harmful to children?**	N/A

Question	Your response
Question 13: Do you have evidence of other objective and measurable characteristics that may be relevant to category 2A threshold conditions?	N/A

Please complete this form in full and return to os-cfe@ofcom.org.uk.