IT

This document has not been adopted or endorsed by the European Commission. Any possible measures indicated in this paper are the preliminary elements being considered by the Commission services, they do not preclude the measures to be finally considered in the Impact Assessment and may not in any circumstances be regarded as stating an official position of the Commission. The information transmitted is intended only for the entity to which it is addressed for discussions and for the preparation of the Impact assessment and may contain confidential and/or privileged material.

IMPACT ASSESSMENT ON ILLEGAL CONTENT ONLINE
STAKEHOLDER CONSULTATION:
INTERNET COMPANIES

Introduction:

In the context of the ongoing work on the Impact Assessment on Illegal Content Online, the Commission would like to get your views on a number of issues set out below. These views will complement the Open Public Consultation (OPC, available here), as well as the data collection exercise based on the table of indicators.

The Commission started work on an impact assessment outlining potential problems, objectives and options in the attached Inception Impact Assessment (IIA). As part of the options to be considered, the Commission will analyse the current situation (baseline scenario) as well as actions to reinforce the voluntary measures as well as possible sector-specific legislation (including in particular on terrorism content online) as well as horizontal legislation applicable to all types of illegal content.

The measures presented in the Inception Impact Assessment[1] are initial ideas, and additional actions and options could be considered. The actions to be undertaken would be mainly addressed to online platforms, but could also require further action by Member States.

In addition to the requests for factual data as part of the reporting exercise within the EU Internet Forum and the possibility to contribute to the Open Public Consultation that closes on 25th June, we would like to offer you the possibility of providing further input to the Impact Assessment by replying to **the questions below and provide any additional considerations in writing by 15th of July.** We are also available on the week of 18-22 June to hold a meeting or videoconference, at a time to be arranged, in order to discuss your input, clarify any questions you may have and discuss additional elements which you consider should be taken into account.

[1] https://ec.europa.eu/info/law/better-regulation/initiatives/ares-2018-1183598_en

Questions

1. What are the main risks or concerns for your company as regards terrorist content online which could be hosted in your platform? Please indicate your agreement with the following statements, with a short justification to the extent possible.

Statement	Impact	Justification
Such content has a negative impact on our users	☐ **Very negative** ☐ Negative ☐ No impact ☐ I don't know	Research has shown that while the Internet does not single handedly lead to radicalisation, it is used by terrorist organisations to spread propaganda. Using YouTube to incite violence, spread violent extremist propaganda, recruit for terrorism, or celebrate or promote terrorist attacks is strictly and specifically prohibited by our Community Guidelines. We do not want the very platforms that have enabled open and free expression to be abused by those who wish to promote terrorism or extremism.
Such content damages the reputational image of the company	☐ **Very negatively** ☐ Negatively ☐ No impact ☐ I don't know	Displaying terrorist propaganda infringes the core values of the company, and fails our mission to serve the Internet

		user with useful information, when not done with the intent to explain, denounce or put into context. We have taken significant action to protect our community against violent or extremist content, testing new systems to combat emerging and evolving threats. We tightened our policies, increased our enforcement teams, and invested in powerful new machine learning technology to scale the efforts of our human reviewers to take down videos and comments that violate our policies.
Such content impacts on the company's business model (e.g. risks of losing advertising or users switching to other platforms)	☐ **Very negatively** ☐ Negatively ☐ No impact ☐ I don't know	It is critical for us to identify and remove terrorist content from our platforms to protect our community, including creators, viewers, and advertisers. In addition, we have taken several steps over the last year to protect advertisers to

		ensure that their ads run alongside content that reflects their values. Google ads may not be placed on any page that contains content that is illegal, promotes illegal activity or infringes on the legal rights of others.
Such content undermines the trust by users when using the Internet	☐ To a large extent ☐ To some extent ☐ To a limited extent ☐ Does not undermine trust ☐ **I don't know**	
Risks of litigation by hosting such content	☐ Is a serious concern ☐ **Is a concern** ☐ Is not a concern	As a company that operates on a global scale, we are always mindful of the risks of litigation.
Risks of diverging legislation in different countries to address such content posing excessive regulatory burden on companies	☐ **Is a serious concern** ☐ Is a concern ☐ Is not a concern	The E-Commerce Directive has offered an effective, EU-wide framework for dealing with illegal content, and has enabled companies like ours to develop voluntary systems that build on top of that. Where countries seek to diverge from that framework, such fragmentation indeed increases the burden on companies, and also risks diminishing the

		effectiveness of their measures.
Other; please elaborate: Risk of rigid legislation	**Is a serious concern**	The methods as well as the platforms used by terrorist organisations keep changing. At the same time, hosting providers continuously invest in technology and new ways to tackle the challenge of terrorist content. We need an open framework to be able to adapt our answer to a constant evolving threat in the most efficient and balanced way possible.
Risk to free expression		The E-Commerce Directive maintains a needed balance adapted to the reality of the Web ; any change to the liability regime would harm the entire ecosystem. Increased pressure through liability risks moving the needle towards extremely-aggressi ve enforcement by companies, to avoid any legal and financial risks, without taking the time to fully understand a piece of content in order to take the right decision. We value

		the quality of a decision rather than its speed, as some tough calls require an in-depth reflection and analysis of the intent as well as underlying implication of a material.

2. What measures could be developed to **reinforce the voluntary approach** (e.g. a Memorandum of Understanding or a Code of Conduct between the EU and the industry including specific commitments building upon the Recommendation[2])?

We are committed to fighting illegal content online. We participate voluntarily in important fora like the EU Internet Forum and are leaders of the industry's Global Internet Forum to Counter Terrorism (GIFCT). We are members of the shared industry database of hashes, using digital fingerprints to help other companies match and remove terrorist content. We continue to speed up removals and we are committed to working together to find new ways to quickly get this content off our platforms, through our own measures and by working closely with law enforcement and governments. We do this because it is the right thing to do.

We are also committed to expanding and regularly releasing a report on how we're enforcing our Community Guidelines on YouTube. This quarterly update will help show the progress we're making in removing violative content from our platform. We published our first report in April, and by the end of the year, we plan to refine our reporting systems and add additional data, including data on comments, speed of removal, and policy removal reasons.

We meet regularly with EU IRU and other EU law enforcement authorities to improve our referral system, exchange feedback, and conduct refreshed training on our Trusted Flagger tools and Community Guidelines. We appreciate the opportunity to have open and honest exchanges on content referrals. However, the Recommendation's strict requirement for 1-hour turn-around times from time of referral does not provide space to have discussions and exchange feedback where there are areas of dispute or where more information is required before we can make an informed and responsible decision. As noted elsewhere, we have serious concerns about a requirement for 1-hour turn around times. An improved system would ask companies to make decisions as quickly as possible, in line with capacity and concern for free expression. It would allow significantly more time time where a referral requires additional information or exchange.

[2] https://ec.europa.eu/digital-single-market/en/news/commission-recommendation-measures-effectively-tackle-illegal-content-online

3. Which actions could be taken to **support small companies** and start-ups in tackling terrorist content online effectively? Should these be taken by larger companies, public authorities or both?

Through the GIFCT, we are committed to helping smaller companies by engaging in shared learning. We aim at helping them develop the technology and processes necessary to tackle terrorist and extremist content online as well as developing best practices. We've begun discussing ways to share machine learning technology with smaller companies, and are working to expand the hash-sharing consortium's membership, and continuing to organize workshops and trainings for smaller technology companies in Europe and around the world.

4. What are your views on **regulating at EU level in the following areas** and how would you qualify the **impact on your business (positive or negative)**? Please provide a short justification of your assessment.

Definition of terrorist content	As our Community Guidelines already state, "We do not permit terrorist organizations to use YouTube for any purpose, including recruitment. YouTube also strictly prohibits content related to terrorism, such as content that promotes terrorist acts, incites violence, or celebrates terrorist attacks." Our Terms of Service require users to comply with the Community Guidelines and all applicable laws.
Requirements regarding the companies' terms of service	Our Terms of Service already require users to comply with the Community Guidelines and all applicable laws.
General requirement for companies to put the necessary measures in place to ensure that they do not host terrorist content (complemented by self regulation)	Companies already have many incentives to put in place necessary measures, in line with their capabilities. Under the E-Commerce Directive, hosting providers must act expeditiously to remove content upon knowledge that it consists of illegal activity or information if they wish to avail themselves of the safe harbor from liability for that content. Beyond that, companies like ours have implemented many voluntary measures that go above and beyond what is required by the law, including by developing innovative new technology and collaborating across industry. Further, the recently-finalized update to the Audiovisual Media Services Directive (VSPs) already calls for video sharing platforms to implement necessary measures to respond to terrorist content specifically, in line with the

	ECD and with a bias for co-regulatory frameworks. For that reason, we believe a legislative mandate to this effect is not necessary.
Specific requirements in terms of action upon referral (including time limit of one hour)	Imposing a time limit for a decision is often unworkable, and it creates the wrong incentive by encouraging deletions of content and speedy decisions over a qualitative and in-depth analysis. Numerous tough calls that have to regularly be taken by Internet platforms require a full analysis of the content, requiring cultural and linguistic expertise. The same symbol or word can have a different meaning in a different context, and it is important to take the right decision to respect the freedom of expression, as well as ensure there is enough material online putting into context terrorism to help understand, remember as well as counteract. One hour is often not enough time to make a responsible assessment in these instances. A regulation in that sense would have very negative consequences.
More explicit and detailed obligations to deploy specific proactive measures (including automatic detection)	While we have developed and implemented various measures to facilitate the detection and removal of certain kinds of content, we are concerned that an obligation to deploy specific automatic measures – such as automatic filtering – would violate Article 15 of the E-Commerce Directive, which was carefully written to balance the need to tackle illegal content with respect for fundamental rights. As the threat of terrorist propaganda as well as technological capabilities to tackle it keep evolving, it would be premature as well as counter productive to anchor down detailed obligations on how to tackle it. There is no silver bullet, and automatic detection is still imperfect. If not used thoughtfully, it can lead to false negatives and deleting legitimate content. We are working faster and working smarter, but we still don't always get this right. For example, recently YouTube mistakenly removed violent and disturbing videos from the war in Syria. Unfortunately, these videos weren't terrorist propaganda - they were documentary evidence of atrocities uploaded by human rights activists. When we make mistakes, we quickly reverse them, and we are working with the human rights community to encourage them to include context so we understand their intent. But we face challenges

	on both sides of the equation every day, even as we push the edge of technology.
Specific requirements to cooperate with other hosting service providers to avoid the dissemination across platforms	Platforms in the GIFCT's hash-sharing coalition are already sharing hashes; such efforts exist and do not not require any legislation. However it requires a common understanding between platforms of what needs to be shared and what safeguards are required, so it cannot be available to platforms that have not met the criteria for membership.
Sanctions in case of non-compliance	1) Creating sanctions put the regime of the e-commerce directive at risk, by creating a precedent where the mere hosting of content could create a liability. 2) It creates a huge burden on companies, due to the breadth of existing content, and the limits of automatic detection. Experience has shown that automatic detection of controversial content can get it wrong, as when YouTube deleted a parliamentary debate on torture or activists videos documenting what was happening in Syria.
Exchanges of information with law enforcement to limit any interference with investigations and to feed into the analysis of terrorist material	We need to take into account the broader context of MLAT as well as the recent U.S. CLOUD Act. Any reform should only happen in a concerted way, and not lead to a conflict of jurisdictions and law for Internet companies. The e-evidence proposal needs first to be fully assessed, debated and improved, before pursuing any further legislation.
Clarify that companies engaged in proactive measures benefit from the liability exemption (Good Samaritan clause)	We welcome the principle of a Good Samaritan clause that would enable companies to take voluntarily proactive measures, without fear of liability if their decisions to review, remove or keep up content are challenged or such monitoring is imperfect. It must be unquestionable that a hosting service does not become liable for any of the information that it hosts simply by virtue of the fact that it has taken voluntary action in good faith, whether of an automated or a non-automated nature. It should also be understood that such actions do not imply that the service provider has knowledge of or control over the information which it transmits or stores. Clarity in this baseline is crucial to encouraging the most efficient and innovative methods to tackle illegal content online.
Requirement to Member States to increase referral capabilities, quality criteria	It is vital the EU law enforcement authorities join YouTube's Trusted Flagger program and use their

for referrals and for referral entities in Member States to provide relevant support to companies in case of doubt about qualification as terrorist content (e.g. through points of contact)	dedicated Trusted Flagger account so their referrals can be routed through the dedicated queues. Our preference is for national authorities to submit terrorist content for removal to the Europol IRU, which could leverage its expertise and serve as a central source of notifications. This would help ensure coordinated and efficient communication with appropriate industry legal removals teams.
Nomination of point of contact within Companies	Rather than create single points of contact, we have built dedicated review queues for Europol's EU IRU and EU law enforcement authorities that are members of the Trusted Flagger program. That ensures that we can achieve 24/7 coverage for referrals, rather than depend on a single point of contact. Our Trusted Flagger dashboard allows participating members to submit comments about their referrals, and we maintain regular and open lines of communication with Europol and other active IRUs.
Reporting obligations for companies[3]	We strive to provide transparency on the way our platforms operate and on the decisions taken, as exemplified by our latest transparency report. Transparency disclosure should be encouraged but its format should not be mandated to allow for enough flexibility to design the best way to provide transparency on unique platforms.
Transparency requirements for companies vis a vis their users[4]	We already release a quarterly report on how we enforce our Community Guidelines on YouTube. We plan to refine our reporting systems and add additional data, However, we need to ensure that our systems to monitor performance in any area meets the high quality level appropriate for external public reporting.
Compulsory safeguards, such as the ones in the general chapter of the Recommendation	YouTube works hard to maintain a safe and vibrant community. We have Community Guidelines that set the rules of the road for what we don't allow on YouTube. We work hard to ensure we effectively enforce them. We rely on a combination of people and technology to flag inappropriate content and enforce these guidelines. Machine learning is allowing us to identify and remove content faster than ever before. Our investment in

[3] See point 41 of the Recommendation.
[4] See points 16 and 17 of the Recommendation.

	technology enables us to address enforcement of our content policies at scale. We rely on teams from around the world to review flagged videos and remove content that violates our terms; restrict videos; or leave the content live when it doesn't violate our guidelines. We support expeditious and transparent processing of notices and measures to prevent their abuse and discourage mistakes. We would welcome guidance on the minimum information required to constitute a valid notice. It should include the precise location of the allegedly infringing content, such as a URL. It should also include reasonable information to contact the notifier – common practice in notice procedures used around the world – unless in very specific circumstances (e.g., for safety). For some legal grounds, such information is indispensable to assessing the complaint. It also enables us to substantiate the notice; to notify the author of its outcome; to give an uploader the opportunity to ask a complainant to retract their complaint; and to process counter-notifications. And it helps prevent bad faith notices (examples at https://goo.gl/B24Eam). We also believe it is important to offer users means to dispute complaints that are made against them. On YouTube, a user's account may be terminated due to: (i) repeated violations of the Community Guidelines or Terms of Service; (ii) a single case of severe abuse (such as predatory behaviour or spam); or (iii) the account being dedicated to a policy violation (hate speech, harassment, impersonation, etc.). For violations falling within (i), there is a "3 strikes" system. If a strike is issued, the user will be sent an email and see an alert in their account's channel settings with information about why the strike was issued. If a user believes that the strike was issued without just cause, they can appeal it. YouTube also terminates user accounts that have received multiple takedown notices for copyright infringement. The user is notified and, if they believe the copyright claim against their account is improper or invalid, they can file a counter-notification. YouTube has various Help Centre articles (including guidance videos) designed to educate users on these issues and processes and, when an account is terminated, the account owner will receive an email detailing the reason. If a user believes that their account

	has been terminated in error, they may appeal. This process is still available to terminated users - access to the account is not required.
The establishment of an external audit/monitoring mechanism for assessing compliance of companies.	We believe that mandating an external audit or monitoring mechanism is neither necessary nor advisable. Companies are best suited to identify the reporting mechanisms that are most appropriate for them, and that do not put at risk the privacy of their users. Google has been a leader in transparency. We launched our first Google Transparency Report in 2010 and our latest just a couple weeks ago. We share data that sheds light on the scale and scope of government requests to remove content, and any user can sort that data by time, country, and product. We share data on Search removal requests under European privacy law and for results that may infringe on copyright. We also launched a quarterly YouTube Community Guidelines report to report on flags and removals. Our systems were designed to find the flagged content most likely to violate our guidelines and escalate that for action quickly. We are currently in the process of rebuilding our systems so they track this data and can report out on it with trusted levels of precision.