

Your response

Question	Your response
<p>Question 1: How do you measure the number of users on your service?</p>	<p>Confidential? – N</p> <p>The Wikimedia Foundation is a non-profit organisation that hosts several free knowledge projects, the largest of which is Wikipedia. Wikipedia, the free encyclopaedia, is a collaborative project created and maintained in over 300 languages by volunteers across the globe. The community of volunteers, who comprise the global Wikimedia Movement, collaboratively write and edit the content of the encyclopaedia.</p> <p>Our policies and practices regarding user data and metrics prioritise security and privacy, by design. We do not require users to log in, and we do not prevent users from having multiple accounts. Accordingly, we cannot generate direct statistics about the number of users on our sites. Instead, we attempt to count unique <i>devices</i> visiting the sites. In some cases, one device = one user; but since a family can share a tablet PC, while single adults can use multiple devices at home (and more at work), that 1:1 equivalence is unlikely to hold true, generally.</p> <p>In line with our data minimization principles, we comply with the EU Digital Services Act (DSA)'s requirement to publish an average number of monthly “active,” “unique” human users by estimating it, using the existing data collected about visits to our sites—namely, unique device counts. We convert unique device data to an estimated unique user count by dividing it (currently) by 2.4, which a survey by Cisco suggests is an approximate number of devices per capita for internet users. In other words, we use third party survey data to estimate that an “average” human accessing Wikipedia, in a</p>

Question	Your response
	<p>given month, does so using 2.4 unique devices (phone, laptop, work devices, “Alexa”-type home assistants, their usual browser’s “private”/”anonymous” sessions, etc). As updated survey data emerges, this conversion factor can be updated, to match evolving trends in internet usage.</p> <p>For more information about our Unique Device counting methodology—including how we attempt to discount bots—see here, here; see also our Privacy Policy and our New User Welcome Survey Privacy Statement. Our conversion of that into usercount data, for DSA categorisation purposes, is described here.</p> <p>This approach, which preserves privacy while providing the relevant regulators with sufficiently reliable and detailed information to appropriately categorise our platform for DSA purposes, was vetted and approved by the European Commission. It is vital, both for privacy and resourcing sustainability purposes, that the UK accepts this approach.</p> <p>Our firm commitment to protect the privacy of our large international user base is necessary so that volunteers and readers alike can trust that they will not be tracked in their activities on Wikimedia platforms. This is further supported by international human rights standards, which stipulate that states have a duty to protect children’s right to form and express their opinions without interference from automated processes of information filtering and profiling (See Human Rights Policy).</p>

Question	Your response
<p>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)?</p>	<p>Confidential? – N</p> <p>With the exception of “special” pages (like the platform Terms of Use or Privacy Policy, or system-generated pages like an “Article history”, which just lists edits made to an article), and also with the exception of on-wiki messages posted by Wikimedia Foundation staff members (e.g. in response to questions from the public), the content on Wikipedia is <i>all</i> user-generated content. That content, and community discussions about it, as well as edits or changes to improve such encyclopaedic content, are publicly-viewable by anyone who visits the site. This is how the platform was designed from its inception more than twenty years ago.</p> <p>Every article has a “history” section, which indicates what changes have been made and who has made those changes, and a “discussion” section, where users can discuss changes they want to make before hitting “edit.” These basic safeguards build accountability into the editing process and put content moderation tools and processes in the hands of the entire community.</p> <p>Such “user-generated” and non “user-generated” content is highly intermingled. On a page like https://en.wikipedia.org/wiki/Wikipedia:Fundraising/2023_banners , for instance, there are messages — even, potentially, <i>parts of sentences</i> — that have been posted/modified by both Wikimedia Foundation staff (i.e. not “users”, but rather, staff of the web host), and ordinary users. Differentiating the two would be artificial, and practically impossible. The English language Wikipedia alone contains 6.4 million articles and approximately 555 are added <i>per day</i>.</p>

Question	Your response
<p>Question 3: Do you measure different segments of users on your service?</p> <ul style="list-style-type: none"> • 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? 	<p>Confidential? – N</p> <p>As previously mentioned, we do not directly measure different segments of <i>users</i> of the various language editions of Wikipedia. Anyone can read Wikipedia, and even edit some articles, without creating an account. We do not require that users create an account to use Wikipedia; even if they create one (or more than one), they do not have to login to them. As for unregistered or logged-out users, we cannot uniquely identify which one is a particular person. Accordingly, we cannot measure or segment “users”, i.e. individual people across platforms and devices. Instead, we base our measurements (and segmentation) on the <i>unique devices</i> that access the platform. To do so, since 2016, we use a Last-Access cookie to track the number of unique devices per domain (e.g. en.m.wikipedia.org) and per project (e.g. Wikipedia), and can separate devices by country (see, for example, our userbase statistics for EU-based users of the Wikimedia projects as required under the DSA.</p> <p>There are limitations to this approach: we are not able to identify users from the data passed in the cookie—the cookie contains only a year, month, and day. With unique device counts, if a user visits the desktop and mobile sites, that user’s devices would be counted twice. Moreover, device geolocation is generally inaccurate—for example, some users may be using Virtual Private Networks (VPNs) in order to avoid local surveillance and/or censorship. It is worth noting that the European Commission does not require us to increase the amount of data collected about our users in order to comply with the DSA.</p> <p>Do you segment user measurements by different parts of your service? For example, by website vs app, by product, business unit</p>

Question	Your response
	<p>As previously mentioned, we utilise the unique device counts to extrapolate the number of active users based on project and, when possible, country/region. We further segment by which app version a mobile device user accesses the sites through (namely, iOS and Android). For reference: in 2022, there were approximately 1.4M monthly active users on the Wikipedia iOS app and 4.5M on the Wikipedia Android app, globally.</p> <p>We measure pageviews and article edits, as well, and further segment those into different categories. For example, monthly pageview metrics are segmented by domain: mobile pageviews are counted using the mobile domain (e.g., en.m.wikipedia.org); Desktop pageviews are counted using the desktop domain (e.g., en.wikipedia.org; cy.wikipedia.org).</p> <p>We also segment pageviews, edits, and unique devices by which of the 300 language versions of Wikipedia is accessed. UK readers regularly access language versions of Wikipedia <i>other than English</i>, including: Chinese (6M UK pageviews monthly); French (3M UK pageviews monthly); Polish (3M UK pageviews monthly); Persian (2M UK pageviews monthly); Arabic (1M UK pageviews monthly); Welsh (90K UK pageviews monthly); Urdu (41K UK pageviews monthly); Irish (12K UK pageviews monthly); Scottish gaelic (10K UK pageviews monthly); Breton (9K UK pageviews monthly); Manx (4K UK pageviews monthly). Together, these figures represent over 15M pageviews <i>per month</i> by UK readers.</p> <p>As highlighted by Lord Moylan and others throughout the OSB's journey through the House of Lords, Welsh Wikipedia is the largest Welsh-language website in the world, and would be in danger of vanishing entirely if Wikipedia becomes inaccessible in the UK (See Lord Moylan's comments in</p>

Question	Your response
	<p>support of an amendment seeking to exempt Wikipedia from the OSB). Welsh Wikipedia is viewed over 7.5 million times every month, and is now part of the secondary curriculum in Wales, within a module of the Welsh Baccalaureate.</p> <p>Do you segment user measurement into different types of users? For example: creators, account holders, active users.</p> <p>As outlined above, we segment device/account (<i>not</i> user) measurement into different categories for different purposes, including editors (see active editors of English Wikipedia by country), anonymous editors, and anonymous users. We do the same to measure bot activity (e.g., bot editors; articles edited by bot editors — note: these are usually community moderator-created bots that perform simple tasks, e.g. typo corrections or undoing basic vandalism).</p> <p>We further measure certain categories of registered accounts with enhanced enforcement powers, including “administrators” and “bureaucrats”. These experienced volunteer editors, selected by the broader community of volunteer editors, have the ability to block or unblock accounts, temporarily protect pages from being edited, and delete pages entirely. These users have registered accounts and utilise those accounts to access as well as take actions on articles.</p> <p>How much flexibility does your user measurement system have to define new or custom segments?</p> <p>The only way to track and count unique users, as opposed to devices, would be if we required</p>

Question	Your response
	<p>everyone to login to use wiki projects, and to provide information about their true identities (presently, editors who wish to protect their identities for whatever reason, including concerns about harassment and other physical threats can create user profiles using pseudonyms, and are not required to offer information that could be used to identify them offline). As previously stated, collecting more information about user and editor identities would not only go against our long standing commitments to our volunteer community and the prioritisation of user privacy and security; the community would resist such changes.</p> <p>If the Foundation is forced to collect new types of data about users, this would likely result in the loss of some of Wikipedia's most important and active contributors, editors and admins, including those from the UK who cannot see the sense of, for example, proving their age and putting up with Foundation-imposed measures, just to carry on doing what they have successfully done - as volunteers, and through no duty to anyone - for <i>over two decades</i>.</p> <p><u>Our stance on privacy, when it comes to analytics, also links to some of our much wider concerns about this Bill, notably around privacy, and barriers to learning and exploration.</u></p> <p>For instance, forcing Wikipedia to institute age verification and/or requiring every individual who wants to read or otherwise access the information on Wikipedia to create and log-into an account would jeopardise the functionality of these services and, more broadly, the availability and quality of neutral, well-sourced information across the internet ecosystem.</p> <p>There are many services that rely on knowledge from Wikipedia to serve vital information to their users. This includes search services like</p>

Question	Your response
	<p>Google, Bing, and DuckDuckGo; virtual assistants like Apple's Siri and Amazon's Alexa; and ChatGPT.</p> <p>We can reasonably expect that most people as well as information aggregation services will instead turn to other sources of information that are not age-gated. Their information diet will be — for instance — a tabloid website, or a free (and likely tampered-with, or advertising-laden) Wikipedia clone hosted by an unknown entity. (A website that re-publishes Wikipedia content — and worse still, tampers with it — but cannot be edited by users, falls outside the scope of this Bill.)</p>
<p>Question 4: Do you publish any information about the number of users on your service?</p>	<p>Confidential? – N</p> <p>Yes. Many of these metrics are published on publicly-available sites like Wikimedia Statistics (aka WikiStats). We also publish data related to requests to alter or takedown content and user data requests from authorities, divided by country, in our bi-annual Transparency Reports; in line with DSA requirements, we will in future be adding (for Wikipedia, only - not our smaller projects) our estimated active monthly user counts, at national level (the DSA requires the data to be broken down by EU country, for Very Large Online Platforms). This will complement our existing DSA (estimated) userdata publication, here: https://foundation.wikimedia.org/wiki/Legal:EU_DSA_Userbase_Statistics</p>
<p>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?</p>	<p>Confidential? – N</p> <p>No. We are a nonprofit that lacks the resources and staffing necessary to build out systems to contribute to further databases.</p>

Question	Your response
<p>Question 6: Do you have evidence of functionalities that may affect how easily, quickly and widely content is disseminated on U2U services?</p> <ul style="list-style-type: none"> • 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? 	<p>Confidential? – N</p> <p>Are there particular functionalities that enable content to be disseminated widely on U2U services?</p> <p><u>Algorithmic Amplification</u></p> <p>The Wikimedia projects are structured and governed in a way that does not allow content to spread virally on the projects, limiting the threat of such content being widely viewed. In the case of Wikipedia, since Wikipedia is organised around a singular goal, the construction and maintenance of an online encyclopaedia, the factual information posted by users is not the type that goes viral.</p> <p><u>Recommendation Systems</u></p> <p>Wikipedia <i>does</i> have limited recommendation features including, for example, featured (community selected) “article of the day” and “on this day in history.”</p> <p>There are recommendation/sharing-like features that are specific to the mobile browser and app interfaces, namely:</p> <ul style="list-style-type: none"> • “Related Articles” links at the bottom of an article; • “Explore” feed on the mobile apps is constantly updated and can be customised to the user’s interests; • “Most read” lists on the mobile apps; • “Share this article” widget on the mobile apps so on-app links can be shared easily on other applications. <p>Features such as these should not automatically be Category 1 or 2B triggers; they are extremely commonplace. Instead, functionality-based categorization criteria should also ask:</p>

Question	Your response
	<p>1. What purpose does it serve, and what sort of content does it share (in Wikipedia’s case, it is merely to spark intellectual curiosity, and aid discoverability of encyclopaedic content);</p> <p>2. Is it neutral/contextual (based on relevance to the context the user is in), or is biased in some fashion, e.g.</p> <ul style="list-style-type: none"> a) Intimately tailored to the user/usergroup demographics; and/or b) Manipulated in order to push certain types of content (paid placement, political messaging, highly sophisticated addiction/engagement tailoring)? <p>3. Has it been made core to the user experience, design, etc, or merely ancillary (as can be seen on a page like https://en.m.wikipedia.org/wiki/2023_Mar-rakesh-Safi_earthquake , someone would have to scroll extensively before seeing the “Related Articles” box).</p> <p>Are there particular functionalities that prevent content from being easily, quickly and widely disseminated on U2U services?</p> <p><u>Human Editors</u></p> <p>The Wikimedia Movement’s approach to addressing potentially harmful or illegal content has been tailored over years of community and organisational practice to promote fairness and minimise harm. This necessarily involves close collaboration between volunteer moderators and professional trust and safety staff. The Wikimedia volunteer community also enforces project-specific policies which address illegal content, like these from English Wikipedia.</p> <p>The Wikimedia community is already highly effective at removing illegal and harmful content on the projects. Researchers at the Berkman Klein</p>

Question	Your response
	<p>Center for Internet and Society at Harvard University found that the median amount of time harmful content remained on English language Wikipedia was 61 seconds. They found that Wikipedia's system of identifying and removing harmful content is largely effective, despite Wikipedia's large scale, the variety of content, and different interpretations of the Wikimedia Foundation's guidelines and policies.</p> <p>There are certain situations which cannot be handled by volunteers and are escalated to the Wikimedia Foundation Trust & Safety emergency response team to address. This includes situations where there is a threat of serious harm to someone's physical safety, as well as some higher level conduct issues which require a full, confidential investigation. This type of escalation is possible because of the trusted relationship between the Foundation and the volunteer administrators who maintain the Wikimedia projects.</p> <p><u>Automated Tools to Support Human Editors</u></p> <p>Editors on Wikipedia employ a multi-layered approach to discovering and removing harmful speech on the projects. The Foundation seeks to empower users to participate in content moderation processes by, for example, providing them access to machine learning tools which they can use to improve or quickly remove content. While the Foundation may assist developers with building tools, they are used and maintained by community members.</p> <p>One of the tools editors can use is ClueBot NG, an automated tool which uses a combination of different machine learning detection methods and requires a high confidence level to automatically remove vandalism on the projects. Another tool is a machine learning tool called Objective Revision Evaluation Service (ORES) which assigns scores to</p>

Question	Your response
	<p>edits and articles in order to help human editors improve articles. Additionally, users with special privileges have access to the AbuseFilter extensions, which allows them to set specific controls and create automated reactions for certain behaviours. While automated tools are used to support existing community moderation processes, the bulk of the work is still done manually.</p>
<p>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?</p>	<p>Confidential? – N</p> <p>When considering which functionalities are relevant to setting categorisation thresholds, Ofcom should take into account that the types of content hosted by a platform for particular purposes—like a collaboratively-maintained online encyclopaedia with policies that require information to be neutrally-presented and well-sourced.</p> <p>The functionality and safety features that are employed on Wikipedia and the other Wikimedia projects are designed to serve the specific user communities that build and maintain them. The projects attract different user bases and are centred around different types of content (images on Wikimedia Commons, for example).</p> <p>We have found that the larger the userbase, the safer the projects usually are from proliferation of content that is not welcome or does not belong on the platform. The maintenance work required for Wikipedia or Commons, for example, relies on multiple volunteer users. The larger volunteer userbase contributes to a greater diversity of voices in the all-important discussions regarding particular content and seeking consensus regarding what actions are to be taken, and related project policies.</p> <p>Forcing changes on collaboratively-maintained projects will necessarily dissuade volunteer</p>

Question	Your response
	<p>participation, thereby reducing the number of users monitoring pages, negatively impacting the community's ability to effectively moderate content on the platform, and decreasing the diversity of voices among the volunteer userbase.</p>
<p>Question 8: Do you have evidence of other objective and measurable factors or characteristics that may be relevant to category 1 threshold conditions?</p>	<p>Confidential? – N</p> <p><u>Summary:</u> In addition to user numbers and functionalities, Ofcom <i>must</i> consider the following in all of its categorisation criteria: the platform's mission (public interest v.s. profit); governance model (bottom-up, community governance v.s. top-down, centralised); resourcing and sustainability.</p> <p><u>User base</u></p> <p>A one-size-fits-all approach for setting Category 1 threshold conditions based on a platform's user numbers or reach <i>without regard to its specific purpose and operations</i> would risk over-regulation of legal and less harmful content. This approach creates a substantial risk that platforms with high visitor numbers and well-established policies <i>and</i> practices that govern user behaviour and content moderation—like Wikipedia—would be placed in an inappropriate category and face outsized regulation relative to their risk level.</p> <p>As previously stated, altering our policies and the functionality of the Wikimedia projects by forcing all users to create and log-into accounts would not only go against our longstanding commitments to and prioritisation of user privacy and security, but would risk destroying our governance model via decimation of our user communities. This would likely mean the loss of some of Wikipedia's most important and active contributors, editors and admins, who are critical to the operation and continuation of the projects.</p>

Question	Your response
	<p>What is worse, if a legislative precedent is established that forces us to collect such data about UK users, we can expect that many other governments around the world will impose similar requirements in ways that will seriously expose our community to security and human rights risks. Such increased personal data collection and processing would present an unacceptable security risk, but particularly those in parts of the world that <i>truly</i> need Wikipedia and rely on safe and secure ways to access it.</p> <p><u>Sustainability and mission</u></p> <p>The obligations imposed under the OSB on nonprofit, public interest platforms with decentralised, volunteer-run content moderation models like Wikipedia <i>should be</i> different from those required of for-profit platforms that have top-down, centrally-directed content moderation systems that support advertising-driven business models.</p> <p>Platforms like Wikipedia, which are not-for-profit, face unique regulatory challenges, which are only exacerbated when categorisation thresholds lack nuance and set the same requirements and expectations for all platforms with a certain number of users.</p> <p>Unlike profit-oriented platforms, Wikimedia projects provide information to individuals <i>without</i> exploiting their data, attention, or targeting them with ads. Platforms that do not derive revenue from advertising or based on the number of 'clicks' on their content do not have the same incentives as commercial platforms to push potentially harmful content to users, whether through algorithmic recommendations or otherwise.</p>

Question	Your response
	<p>Also, the Wikimedia projects pursue the public interest; more specifically, they empower everyday citizens to determine the inclusion, presentation and prominence of content. There is no commercial or political incentive for the Wikimedia Foundation to promote some of that content (e.g. to accept paid political advertising in order to maximise investor returns).</p> <p>A categorisation scheme that does not recognize and account for these factors would improperly place Wikipedia in the same tier as Facebook, YouTube, and Twitter, based almost entirely on user numbers, and risk destroying Wikipedia's sustainability.</p> <p>Categorisation criteria could for instance ask for evidence of nonprofit status, for example (in our case) proof of "501(c)(3)" (non-profit) status under US law. It could also look at advertising revenue (as a proxy for attention/virality-based business models which could be at higher risk of problematic content amplification, misplaced incentives around content moderation policy, etc.).</p> <p>Alternatively, or in addition, criteria could examine an organisation's legally-enshrined objectives/purpose, as documented (usually) in its founding documents (Articles of Incorporation, charter, etc.).</p> <p>It has previously been suggested, including during the Bill's Lords Report stage and by a vast diversity of signatories to this Open Letter, that organisations with a <i>public interest</i> mission, whose services are "<i>provided for the purpose of indexing, curating, adapting, analysing, discussing or making available content in the public interest, including but not limited to historical, academic, artistic, educational, encyclopaedic, journalistic, or statistical content</i>", deserve protection from the Bill's excesses. We concur, and therefore suggest that if</p>

Question	Your response
	<p>those services are not added to Schedule 1 (and thus spared the Bill's general burdens and risks), then at the very least such a purpose/mission should be recognised in questions relating to categorisation. (And as noted here, "public interest" is a recognised UK legal concept, allowing Ofcom to rely on that existing caselaw, if there were questions about an organisation's true nature).</p> <p><u>Volunteer-driven content governance model</u></p> <p>The communities who build Wikipedia, for instance, collaborate to effectively and swiftly remove content that runs counter to the purpose of writing a fact-based, well-sourced encyclopaedia article, or does not otherwise meet high quality standards. Wikipedia's successful model of community collaboration and deliberation empowers volunteers to consider the context and sourcing of every sentence, data point, or image. This allows them to make nuanced and thoughtful decisions, and to avoid the mistakes and over-censorship common to the automated flagging and removal processes used for content moderation by many commercial platforms.</p> <p>New obligations <i>on the platform host</i> (the Wikimedia Foundation) to monitor, automatically remove, block or filter certain content, or to respond to complaints within timeframes so short that they prevent meaningful community decision-making are not compatible with community governance models like Wikipedia's. A duty to reliably shield users (adults or children) from accessing certain content could be seen - particularly in the face of occasional moral outrages, political pressure, and/or ill-considered judicial precedent) as a duty to monitor every piece of information that is uploaded to a platform, or even</p>

Question	Your response
	<p>changes to specific sentences in an article, and review the content with regard to its potentially illegal nature (under highly specific UK legal criteria) or harmful effects. As such it is detrimental to the way that volunteers on Wikipedia can make responsible decisions together without the Foundation interfering in those established processes.</p> <p>Category 1 status would also require a site like Wikipedia to allow UK users to “prevent non-verified users from interacting with content which that user generates, uploads or shares on the service” (Clause 15(9) and 15(10)(b)). There does not appear to be any carveout for “interaction” consisting of community moderation. In essence, this would mean that all users - even administrators, and the “arbitration committee” members who can review their actions - would <i>have</i> to submit to identity verification before they can do <i>anything</i> - even <i>read</i> - content posted or edited by those UK users. The vast majority of Wikipedia users, admins, arbitrators, etc, could either refuse or be unable to reliably verify their identity to continue doing something they have reliably done for over two decades. If assigned to Category 1 status, Wikipedia would be made essentially powerless to continue as a community-moderated endeavour, <i>worldwide</i>, in order to serve a UK audience. It is critical, therefore, that Category 1 criteria exclude sites where moderation is substantially dependent on content “interactions” by other users (in the sense of clause 15(10)(b)).</p> <p>The existence, or not, of meaningful community-based moderation is generally objectively visible in and of itself. If it <i>also</i> had to be <i>measured</i>, then categorization criteria could for instance compare (i.e. measure a ratio of) moderation actions taken by the website host (in our case, the Wikimedia Foundation) versus moderation-type actions taken by the userbase.</p>

Question	Your response
	<p data-bbox="699 327 1070 360"><u>Other Systems & Processes</u></p> <p data-bbox="699 434 1369 512"><u>Existing policies and practices to evaluate, mitigate risks and harms</u></p> <p data-bbox="699 526 1377 882">Platforms that have effective policies and practices to mitigate risks and harms potentially posed by U2U content, and to protect the fundamental rights of their users—implemented <i>prior</i> to the enactment of any legal mandate to do so—should be recognized as inherently lower-risk, and thus more disproportionately affected by Category 1 status.</p> <p data-bbox="699 898 1377 1254">For example, the Wikimedia Foundation has conducted multiple impact assessments to identify and evaluate human rights risks related to the Wikimedia projects, as well as opportunities to address and mitigate those risks (see Human Rights Impact Assessment (HRIA); HRIA Report (2020); Child Rights Impact Assessment (CRIA) to be published later this year).</p> <p data-bbox="699 1270 1369 1859">A further objective and measurable criterion could be whether the platforms are already regulated under the EU DSA <i>and</i> are generally implementing the requisite online safety measures in a geographically universal way. Provided that the UK also benefits from those compliance actions, the platform’s risk (from a UK perspective) can also generally be assumed to be lower, and thus less meriting of Category 1 status than (for instance) a UK-only site that has never had to consider the DSA, or a global site that excludes UK users from the protective changes it has made to its services for EU DSA compliance purposes.</p> <p data-bbox="699 1874 1345 2000">Note that EU DSA VLOP status applies based purely on usercount data, so it would be wrong and potentially damaging to assume that a</p>

Question	Your response
	<p>service that has “Very Large Online Platform (VLOP) status under the DSA is automatically so risky that it should get Category 1 status under the UK OSB. This is also because UK OSB Category 1 obligations go beyond those imposed by DSA VLOP status, and there must therefore be tighter criteria around Category 1 designation, in order to maintain proportionality.</p> <p><u>Collaboratively-edited community content</u></p> <p>Category 1 status for a site like Wikipedia would have enormous consequences, because of the obligations that would then flow from that status. As discussed above, Category 1 status should not apply to sites (or parts thereof) where content moderation <i>depends</i> on interaction by non-verified users. It also should not apply if Category 1-specific content visibility duties would cover content that has multiple authors.</p> <p>For example, the Clause 15 user empowerment “duty to include in a service features which adult users may use or apply if they wish to filter out non-verified users” would be extraordinarily problematic for something like Wikipedia. Any given sentence of any given Wikipedia article might have been edited (in whole, or in parts as small as a single letter), by dozens, hundreds or even thousands of users around over the last 20+ years. They could but would therefore be senseless for Category 1 status to apply in respect of parts of a U2U platform whose content has multiple different authors; the platform cannot reasonably show only the parts of a sentence or word that happened to be posted by a “verified” user.</p> <p>Categorisation criteria should therefore prevent Category 1 status applying to parts of platforms</p>

Question	Your response
	<p>hosting collaboratively-edited content.</p>
<p>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?</p> <ul style="list-style-type: none"> • 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? 	<p>Confidential? – N</p> <p>As peers from across the political spectrum have pointed out, Wikipedia poses little risk to children compared to the commercial platforms that the OSB was meant to target. To be sure, Baroness Kidron identified Wikipedia as a service that is “inherently in a child’s best interest.”</p> <p>Wikimedia free knowledge projects are important resources through which children across the UK exercise and access their rights by sharing knowledge and gaining media literacy in any number of languages, not only in English (see discussion of Welsh Wikipedia in the response to Question 3).</p> <p>Additionally, Wikimedia UK regularly works with schools and universities to put on classroom education activities, teaching students how to contribute to Wikipedia and educating them about how information is shared and spread online. These programs were designed with digital literacy skills development in mind, and help students to better exercise their writing, research, and critical thinking skills while navigating content online.</p> <p><u>Limited Purpose and Scope</u></p> <p>Wikipedia and the other Wikimedia projects are structured in a way that does not allow users to post whatever they want—they are designed for specific purposes, and the guidelines that govern the projects enforce the projects’ primary purposes. In Wikipedia’s case, it’s a collaboratively-edited</p>

Question	Your response
	<p>online encyclopaedia. User-generated content that is not encyclopaedic in nature, and not neutral and reliably sourced according to rules set by the volunteer editor community does not belong, and will be removed as being out of scope and inappropriate for the website.</p> <p>Since Wikipedia is organised around a singular goal, the construction and maintenance of an online encyclopaedia, the types of potential harm on the platform are different than on most social media platforms. Further, they present less risk of potential harm to users and to the internet ecosystem more generally. The Wikimedia projects are totally free to use, and we aim for it to stay that way in perpetuity. Unlike some commercial platforms, there are no subscription models or multi-tiered memberships on our projects. Further, users cannot monetise the informational content they post or their use of the service, thereby reducing the incentives that exist on other platforms to share click-bait or other viral content.</p> <p>The ethos of being a free culture project also often disincentivizes posting illegal content on the Wikimedia projects. For example, Wikimedia Commons, our free image repository, often removes copyrighted content even if there may be other legal exceptions or justifications for hosting the content.</p> <p><u>Effective moderation policies & practices</u></p> <p>The Wikimedia Movement's approach to addressing potentially harmful or illegal content has been tailored over years of community and organisational practice to promote fairness and minimise harm. This necessarily involves close collaboration between volunteer moderators and professional trust and safety staff. The Wikimedia volunteer community also enforce project-specific policies which</p>

Question	Your response
	<p>address illegal content, like these from English Wikipedia.</p> <p>As previously mentioned, researchers found that the median amount of time harmful content remained on English language Wikipedia was 61 seconds, and that Wikipedia’s system of identifying and removing harmful content is largely effective. There are certain situations which cannot be handled by volunteers and are escalated to the Wikimedia Foundation trust & safety emergency response team to address (i.e., CSAM or TVEC). This type of escalation is possible because of the trusted relationship between the Foundation and the volunteer administrators who maintain the Wikimedia projects.</p> <p>Outside of the circumstances described above, the Foundation believes that the open, participatory content governance on sites like Wikipedia guarantees that what is on the project serves socially useful purposes. Changing that balance, by forcing the platform to dictate policy, then on a day-to-day basis monitor, assess, categorise, and selectively or wholly deny access to content, fundamentally changes that dynamic, leads to editor attrition, and thus harms the very thing that makes these projects functional, relevant, and socially useful. And further undermines the primary purpose and core function of Wikipedia: a freely available and widely-accessible online encyclopaedia that is not age-gated or censored based on the age of the person holding the volume.</p>
<p>Question 10: Do you have evidence of other objective and measurable characteristics that may be relevant to category 2B threshold conditions?</p>	<p>Confidential? – N</p> <p>The criteria for 2B categorisation should be very clear and tightly-scoped, particularly because it will require increased resources expenditures by both the covered platforms and the regulator. The latest amendments to Category 1 thresholds suggest that tiny-but-risky sites can now be Category 1, which logically means that Category 2B services are</p>

Question	Your response
	<p>not expected to be particularly risky—the OSB generally does not require 2B services to take extra protective measures.</p> <p>As stated in our response to Question 8, in addition to user numbers and functionalities, Ofcom must consider the following in its categorisation criteria: the platform’s mission (public interest v.s. profit); governance model (bottom-up, community governance v.s. top-down, centralised); resourcing and sustainability.</p> <p>Public interest platforms (PIPs) like online encyclopaedias and libraries should be exempted from categorization—regardless of the number of users that access them, because their mission and core function is to provide the public with access to diverse and reliable sources of educational content and information. Wikipedia and other Wikimedia projects are designed to make information easily accessible and <i>freely</i> available. As previously mentioned, unlike some commercial platforms, there are no payment interfaces, subscription or tiered membership models, or monetization features of any kind, either in-app or on the websites on any of the Wikimedia projects.</p> <p>The reality is that the Wikimedia Foundation, as a global nonprofit supported by charitable donations, has a very small legal team that is already stretched thin working on: complying with the UK and EU GDPR (and UK Data Protection Act, plus all its related legislation), and the EU DSA; navigating growing regulatory uncertainty in the United States, the jurisdiction where we are headquartered; and, trying to defend the volunteer community and their efforts from attacks by a range of private actors and governments hostile to free knowledge — to mention but a few major concerns. We do not have a single UK-based staff member whose job is specifi-</p>

Question	Your response
	<p>cally devoted to regulatory compliance or government relations there. Even keeping up with OFCOM consultations and developments in this area, will be a major stretch.</p> <p>Other entities who run PIPs, from FixMyStreet to the Heritage Alliance, are equally worried about the threat posed by the OSB to PIPs, and to UK society, more broadly, and the Lords heard them when they expressed their concerns. We urge OFCOM and the Secretary of State to use the powers and discretion afforded to them to the fullest possible extent, to preserve and protect this important sector, and allow regulatory supervision to focus on the clearest sources of risk online. Categorisation criteria are one area for discretion, but less effective than the use of Secretary of State powers, as advertised by Lord Parkinson himself, to add service categories to Schedule 1 of the Bill.</p>

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