

2025 EU Systemic Risk Assessment

- *Cover note* -

Executive Summary

Wikipedia’s 2025 EU Systemic Risk Assessment—in spreadsheet form—is appended to this cover note. It sets out the Wikimedia Foundation’s assessment of “systemic” risks, in the EU, linked to the use of Wikipedia. This is a legally-required document under the European Union Digital Services Act (DSA).

Out of 8 EU systemic risks that were assessed, the 2025 risk assessment designates three risks as having the highest salience for Wikipedia:

1. Disinformation around conflicts, and/or civic and electoral processes
2. Disinformation around historical/geographical narratives
3. Harassment among the volunteer community

The designation of these three risks as the most salient for Wikipedia reflects a high degree of continuity with the systemic risks identified in our 2023 and 2024 submissions, given the relative consistency of the external threat environment over the past year and the relatively static nature of our systems. Indeed, the systemic risks identified in this submission remain largely consistent with those reported in 2024 as a result. It is important to acknowledge that Wikipedia as a platform - that is, the technical aspects and user experience of the platform - is both low risk and relatively static. It does not undergo significant changes year over year, nor do any of the systemic risks we have identified appear to rise to a level that would require sudden and urgent changes to mitigate them. This feature is by design and stands in contrast to other online platforms, which seek to rapidly evolve and incorporate new technologies and features in order to pursue their profit-driven incentives. Any changes made to Wikipedia’s technical infrastructure or user experience occur incrementally over a number of years and are carried out through an intentional process that incorporates robust efforts to solicit the input and feedback of readers and members of our volunteer community from across the globe. For example, the Foundation introduced in 2023 the first update to Wikipedia’s interface in a decade. As this 2023 [article](#) in Fast Company observed, “For years, Wikipedia’s volunteers have been hashing out an update to the internet’s foremost encyclopedia. You might not even notice—which, for such a widely used service, is kind of the point.” Any such changes to Wikipedia are centered around the Foundation’s

commitment to upholding users' rights to privacy and access to information. Thus, the systemic risks we identify in this exercise are expected to be largely consistent year over year and remain at a level where new mitigations are attempts to further improve and increment technology and policies that are already working generally well to mitigate these risks.

Nonetheless, the Foundation continues to implement a range of existing and planned mitigations for all 8 risks identified in this risk assessment on an ongoing basis. An updated mitigations plan will also be submitted to regulators—as a follow up to this Risk Assessment—upon its completion without delay. Globally, the Foundation continues to deal with a much wider range of risks, such as growing threats of physical intimidation to Wikimedia volunteers in authoritarian countries outside of the EU, but these remain out of scope for the DSA.

[The Wikimedia Foundation's 2025-26 Annual Plan](#) provides a more holistic overview of the issues our volunteer editors and audience of readers are facing every day, and what is being done to protect and defend the rights of our audiences and editors to access and share free knowledge across borders. Furthermore, the Foundation's [Human Rights Policy](#) describes the organization's commitment to identify and mitigate human rights risks globally.

About the Wikimedia Foundation

The Wikimedia Foundation is the nonprofit organization that hosts [Wikipedia and other free knowledge projects](#). The vision of our free knowledge communities is a world in which every single human being can freely share in the sum of all knowledge. To this end, we support a vibrant community of more than 300,000 volunteers around the world, who contribute to the Wikimedia projects by adding, editing, and verifying content in over 55 million articles across more than 300 languages, all for free and without ads.

Contents

Executive Summary

About the Wikimedia Foundation

Contents

Background

About this document

About the Wikimedia model: advantages and challenges when it comes to EU DSA SRAM

How are relevant risks selected?

What risks are excluded by this approach?

How are the selected risks assessed?

Stakeholder engagement

What are the key risks highlighted in this 2025 exercise?

Conclusion

Background

About this document

Wikipedia's 2025 EU Systemic Risk Assessment sets out the Wikimedia Foundation's assessment of systemic risks linked to the use of Wikipedia in the EU. This is a legally-required document under the DSA. It has been completed by the Foundation's Legal department in consultation with a range of internal subject matter experts and stakeholders across the Foundation. The Risk Assessment is an intermediate step in a wider exercise. The DSA's overall Systemic Risk Assessment and Mitigation (SRAM) process requires the Foundation (as the hosting provider of Wikipedia) to not only assess risks, but also determine what risk mitigations are appropriate.

This process of assessment and mitigation is repeated at least annually. Interim updates may be appropriate, in response to significant changes in the risk landscape.

The DSA SRAM exercise should not be viewed in isolation - it is a smaller piece of broader ongoing risk and mitigation assessment work by the Foundation (see Figure 1), and a much broader overall program of work for the Foundation and in the wider Wikimedia ecosystem, to fulfil our mission to empower and engage people around the world to collect and develop educational content under a free license or in the public domain, and to disseminate it effectively and globally.

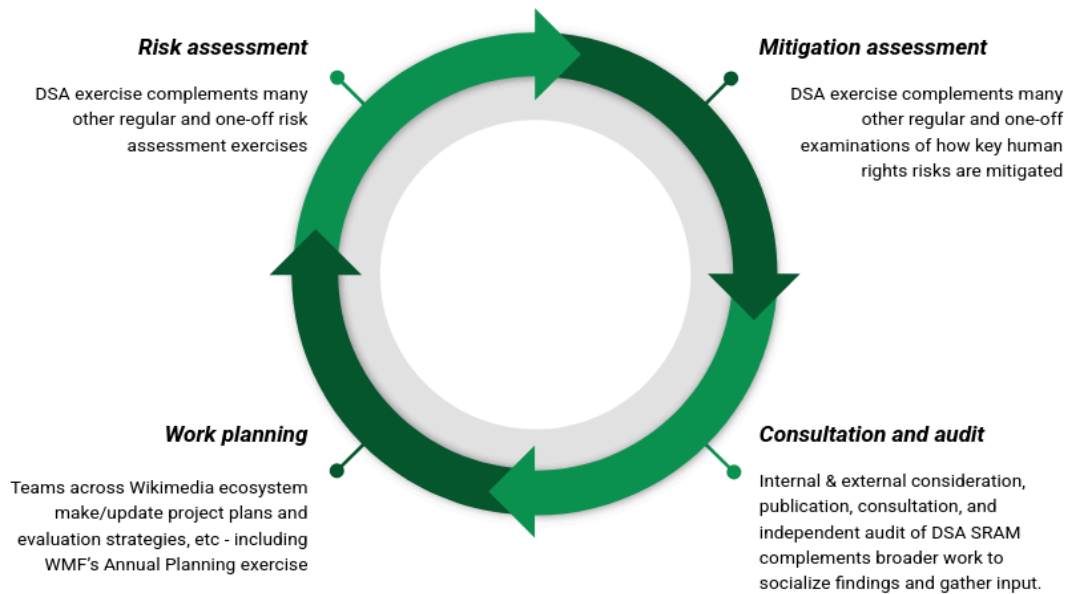


Figure 1: The EU DSA SRAM process adds to work that is already done in order to ensure that the Wikimedia projects are a force for good in society; the role of human rights impact assessment and mitigation at the Foundation is shown in this cycle diagram.

The Foundation is opting for a single, living SRAM document which can be updated on an ongoing basis. This is also sometimes referred to as the Foundation’s DSA Risk Register. Each year, an annual snapshot will be taken of that document, and will be filed with regulators.

The DSA sets out risk assessment and mitigation as a two-step process (requiring two successive filings: first the risk assessment, and then the mitigation plan—see Figure 1, above). Despite this, our DSA SRAM Register has been structured so that it can be used in a single-step strategy in future years if appropriate (see Figure 2, below).

The Foundation submitted its first plan to mitigate the risks identified in its inaugural Systemic Risk Assessment for 2023 on July 12, 2024. The Foundation has already begun implementing and will continue to implement these mitigations while it evaluates their efficacy and identifies any new strategies that could be implemented to further mitigate risk. An updated mitigation plan will be submitted in the coming months.

| Category | Risk | Risk Description | Article 34.1a (Illegal Content) | Article 34.1b (Fundamental Rights) | Article 34.1c (Disinfo / Civic & Political Discourse) | Article 34.1d (Gender / Child Welfare) | Article 34.2a (Algorithmic Systems) | Article 34.2b (Content Moderation) | Article 34.2c (Tou Enforcement) | Article 34.2d (Advertisements) | Article 34.2e (Data Practices) |
|-----------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|-------------------------------------|-------------------------------------------------------|----------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------------|--------------------------------|
| Harmful Content | Dangerous content | Users may attempt to place content on Wikipedia (e.g. articles, text, photos, videos, or other content) in breach of applicable Foundation and community rules and policies, such as the Foundation's Terms of Use. Existing Foundation and editor community processes may not always identify this content or identify it in a timely manner. Use of generative AI by malicious actors to create such content at scale could accelerate its introduction onto Wikipedia pages and, potentially, overwhelm existing content moderation practices. As a result, vulnerable users, including children, could access dangerous or harmful content on Wikipedia (such as content promoting self-harm or eating disorders, hateful or racist speech, pornography, sexually explicit images, and/or child sexual abuse material in a manner or context that fails to align with encyclopedic quality, Wikipedia editorial standards, and educational value), which in the case of children could then have the effect of diminishing children's rights to be free from physical or mental violence, injury or abuse, neglect or negligent treatment, maltreatment or exploitation (UDHR 3, 5, 25 / CFREU 1, 3, 4, 6, 24). | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

Figure 2.1: Example extract from the Foundation’s EU DSA SRAM Register. Its components will be updated and filed at least once a year with EU regulators. The DSA foresees the Mitigations column (blue) being completed and filed separately, as a second step following Risk Assessment. In practice, our DSA SRAM Register will record both our assessment of risks and their mitigations; this example, which shows all columns, illustrates our planned end-state for the Register.

| Category | Risk | Comment on the influence of components listed under Article 34.2 on risk | Likelihood | Impact | Notes on Likelihood and Impact Scoring (see: Human Rights Risk Assessment Framework in cover note annex) | Salience (Derived from Likelihood x Impact) |
|-----------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| Harmful Content | Dangerous content | The manifestation of this risk could, though may not always, violate the Foundation's Terms of Use, which prohibits, "Posting or modifying content with the intention to seriously harm others, such as deliberate inducements to self-harm, or deliberate triggering of epilepsy" and "Posting or trafficking in obscene material," (34.2c). However, the manifestation of this risk primarily challenges the implementation of community governance practices and community-developed policies around the suitability of content for Wikipedia's encyclopedic mission, which together, form Wikipedia's main approach to content moderation. For the most part, since non-educational or otherwise inappropriate content is routinely removed by the community based on its own restrictive policies and processes, this is not typically a matter in which Foundation content moderation is necessary or appropriate (34.2b). | 2 | 2 | Likelihood was scored as a 2 due to the assessed possibility that dangerous as described here content could be found on Wikipedia within the six-month horizon. Impact was scored as a 2 due to the assessed impact of users, including children, facing moderate emotional hazards as a result of being exposed to such content. | Medium Salience |

Figure 2.2: Example extract from the Foundation’s EU DSA SRAM Register (continued).

About the Wikimedia model: advantages and challenges when it comes to EU DSA SRAM

Wikipedia and other Wikimedia projects provide free access to neutral, well-sourced information about science, culture, history, and other encyclopedic subjects. The information in question is added, organized, and edited by a decentralized community of volunteers who engage in open debate to reach consensus around content decisions and policies.

Volunteers address most everyday content issues on the Wikimedia projects, such as intentional vandalism or edits that do not meet Wikipedia’s reliability and neutrality standards, without interference from the Wikimedia Foundation.

This volunteer-led, self-governing model provides for several advantages that allow our movement to advance knowledge equity and to combat disinformation through collective contributions and open debate about content and content moderation. A global community of committed volunteers with specific expertise and local knowledge and language skills have built an unprecedented repository of encyclopedic information that is available and culturally relevant to more than 300 language communities around the globe. These same qualities allow these communities to monitor Wikipedia pages for disinformation and to rapidly remove content that may range from common vandalism to artifacts of coordinated disinformation campaigns. All of this work adding, editing, and verifying content takes place in the open and is governed by policies developed and implemented democratically by these volunteer communities themselves.

This unique model has allowed the Wikimedia projects to achieve the prominence and high quality they have now, and the volunteer-led, distributed decision-making that occurs on the Wikimedia projects is a manifestation of the ideals of online participation and freedom of expression.

Despite these advantages, some challenges naturally arise.

One challenge relevant to this work emanates from our community governance model. Beyond the Foundation's Terms of Use and other high-level policies, most policies and decision making related to platform governance and content moderation are developed and implemented by these volunteer communities. Thus, any significant changes to platform governance and content moderation policies or practices—such as possible recommendations stemming from human rights impact assessments—require consultation with, buy-in from, and leadership by these communities. This model of governance, therefore, can be slower to effect change than that which is feasible for platforms with top-down models.

Another challenge is that the Wikimedia Foundation operates as a non-profit organization. The Foundation's largest share of revenues comes from individual donations; it does not generate ad revenue or sell user data to generate income. As a result of this model, the Foundation has fewer financial and human resources available to tackle significant technical changes or upgrades than do other VLOPs. Any reallocation of resources to make significant changes to Wikipedia for risk mitigation requires pulling resources from other priority areas, such as requested feature development, fixing bugs or broken tools, enhancing the overall stability and reliability of the websites we host, or—more fundamentally—continuing to create an environment in which volunteers want to engage and improve the projects.

This presents a challenge with respect to the EU DSA SRAM process: the DSA expects greater intervention from platform operators (“accountability” / “responsibilization” of platforms). This is a sensible objective for more conventional, for-profit, social media platforms, whose content policies and moderation are imposed on a top-down basis.

However, when it comes to the Wikimedia projects, regulatory obligations should not be interpreted in a way that reduces the Wikimedia communities' autonomy, enthusiasm, and control.

Accordingly, in stark contrast to other commercial platforms, the Foundation intends to generally refrain from dictating changes to content policy, or displacing effective community mechanisms (e.g. for complaint handling, or efforts to tackle certain categories of problematic content). Our focus is instead on creating the right conditions for success. This is achieved by continuing the Foundation's catalytic, framework role, for instance by creating useful tools, supporting community structures (such as key committees), offering training, fostering discussion, offering guidance and support, and engaging in the co-creation of new policies and systems, such as the Universal Code of Conduct ("UCoC") and its enforcement processes. The Foundation continues to assess that this approach is effective, while monitoring for possible risks that may develop over time, particularly in smaller communities that may be less resilient due to the lower number of available editors to address issues. The UCoC itself contemplates this issue, with more global support available for languages that may have fewer volunteers to provide local governance and editorial processes.

The focus is therefore on ensuring that the Wikipedia communities have the means to tackle these important problems—rather than usurping their leading role.

How are relevant risks selected?

Our starting point is the Foundation's global, non-DSA-specific view of the risk landscape for Wikimedia projects. In particular, our DSA exercise builds on the earlier results of a detailed, third-party-led, enterprise-wide [Human Rights Impact Assessment \(HRIA\)](#), last conducted in 2020, which laid out a detailed analysis of risks relating to the Wikimedia projects generally, and how they could be better mitigated. This body of human rights due diligence is complemented by ongoing efforts to identify and mitigate human rights-related risks on Wikimedia projects, such as the Foundation's [Child Rights Impact Assessment](#), which was published in January 2024, and a more recent human rights impact assessment evaluating the opportunities and risks associated with artificial intelligence and machine learning on Wikimedia projects, which will be published in the coming weeks. Following a recommendation provided in the Foundation's first DSA audit report, an updated enterprise-wide HRIA will be commissioned in the coming months to evaluate how previously identified risks have evolved over the past five years and whether new risks have emerged. The findings of this HRIA will be incorporated into subsequent DSA systemic risk assessments.

The DSA is clear that the Foundation's global, non-DSA-specific work should not be used as-is for this specific exercise. For one thing, the HRIA (and related work) are not EU-specific. For another, not all risks in the HRIA have the EU-wide scope and scale to be

“systemic”, as the DSA requires. And not all of them are also Wikipedia-specific.¹ For example, the use of algorithms to recommend content on commercial platforms has resulted in significant concern among the general public around the privacy of users’ data and how that data feeds into those algorithms, the internal workings of which are rarely visible to the public. On Wikipedia, however, the few algorithms used to recommend encyclopedic articles do not rely heavily on individuals’ user data because Wikipedia only collects the minimal amount of user data feasible. Furthermore, the internal workings of such algorithms are documented transparently using model cards for the public to scrutinize.

These differences represent an important distinction between how commercial platforms and Wikipedia use these technologies – a difference that is often underappreciated by the general public. We are therefore being careful to assess the seriousness of each risk identified in our risk register based on two factors: 1) how that risk could manifest on Wikipedia given how the platform actually operates; and 2) the likelihood of that risk arising within the EU.

Based on the wording and structure of the law, our first step for this exercise was therefore to select risks by taking the risks that were salient at a more global level (identified through large scale but possibly infrequent exercises such as our HRIA, and from other sources such as reporting, ad hoc studies, etc), and then applying DSA-specific filters:

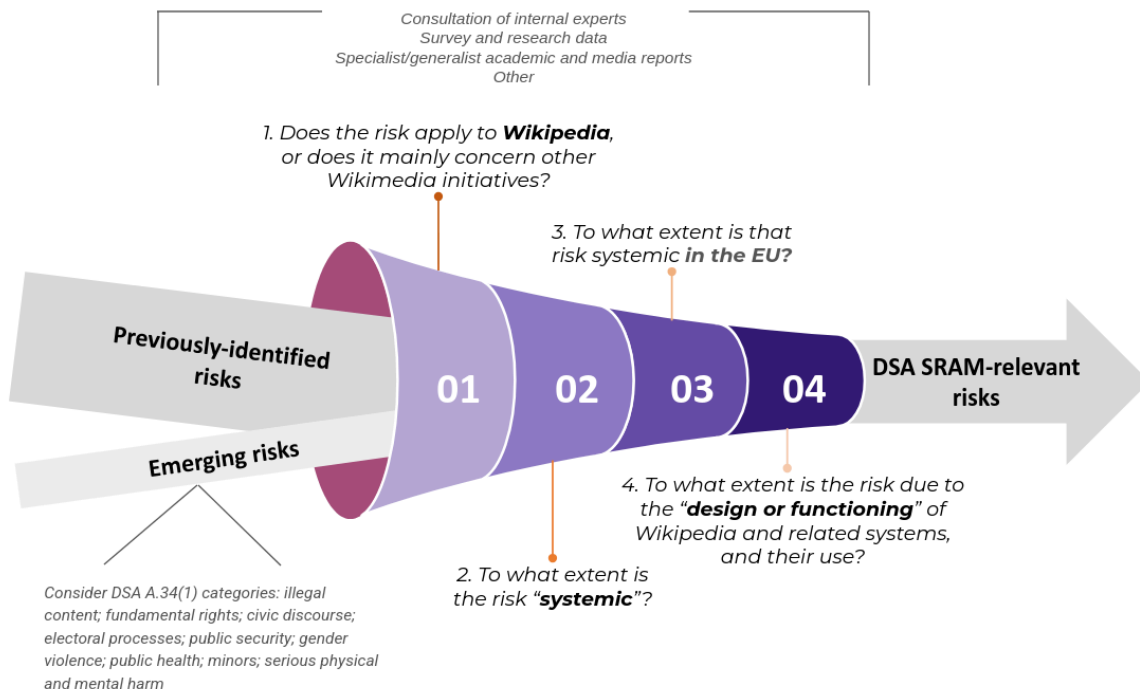


Figure 3: Diagram showing how the Wikimedia Foundation selects risks for inclusion in the DSA SRAM process.

¹ Of all the Wikimedia projects, only Wikipedia is in-scope for the DSA SRAM exercise—it is the only project large enough to have been designated as a Very Large Online Platform (VLOP); the other Wikimedia projects are subject to several other DSA obligations, but not SRAM.

What risks are excluded by this approach?

DSA-specific filtering means that important global risks, which are the focus of extensive Foundation and wider community activity, *are not reflected in the DSA SRAM Register*.

For example, country-wide blocking of the Wikimedia projects is a major risk for affected populations and for the Wikimedia projects themselves, which then have reduced participation from affected citizens. We do not currently consider this an EU systemic risk, nor is it clearly linked to Wikipedia's design/functioning/use, so it has not been included in this DSA-specific exercise, despite its wider importance.

In addition, it should be noted that Wikipedia was designated as a VLOP based on its number of monthly active visitors from the EU. However, certain of the risks included in our DSA SRAM exercise are likely to be experienced only by actively-engaged volunteers, such as those making and discussing changes to Wikipedia articles. This is a much smaller number of people—far below the 45 million user VLOP threshold. By way of example, English-language Wikipedia only sees (approximately) 59,000 active editors a month; of these, *at least* 34,000 come from non-EU countries.² For French Wikipedia, there are only approximately 5,700 active editors from France, and another 420 from Belgium. Accordingly, even though the Foundation is focused on tackling risks such as volunteer harassment, globally, it was unclear to us to what extent that should be selected as a DSA-specific EU systemic risk. We have not made any adjustments based on these criteria this year, but plan to continue to monitor the EU-specific extent and impact of the risks we identify in the EU context. Some risks were sunsetted last year based on their lower likelihood of occurrence and the Foundation's judgment that they have been effectively mitigated.

How are the selected risks assessed?

Based on the four factors shown in Figure 3, the global risks and their assessments are—or are not—transposed (modified, if appropriate) into the DSA SRAM Register. For the most part, the Foundation's broader assessments of global risks, e.g. the analysis contained in our HRIA, will remain somewhat relevant to the risks transposed into the DSA SRAM Register. Those documents therefore provide important further reading for those with an interest in the various risks logged in the DSA SRAM Register.

Starting from that existing (wider) assessment of risks in exercises such as our growing body of human rights impact assessments, we then conduct a second pass assessment from a DSA-specific perspective. This includes consideration of the matters listed in DSA Article

²

[https://stats.wikimedia.org/#/en.wikipedia.org/contributing/active-editors-by-country/normal|table|last-month|\(activity-level\)~5..99-edits|monthly](https://stats.wikimedia.org/#/en.wikipedia.org/contributing/active-editors-by-country/normal|table|last-month|(activity-level)~5..99-edits|monthly)

34(2) - such as what our Terms of Use permit or prohibit, and how that is enforced), although in practice, a wider range of considerations are relevant than just those; consistent with Wikimedia principles, user community policies and moderation practices tend to be more restrictive than our basic Terms of Use, and are therefore an important consideration.

In terms of inputs, this second-pass exercise relies heavily on existing assessments (those done at the more “global”, non-DSA-specific level), bringing in additional knowhow, studies and data (often EU-specific, country-specific or European language-specific), if available - for example the latest [reports](#) and [results of studies](#) into EU election-related disinformation.

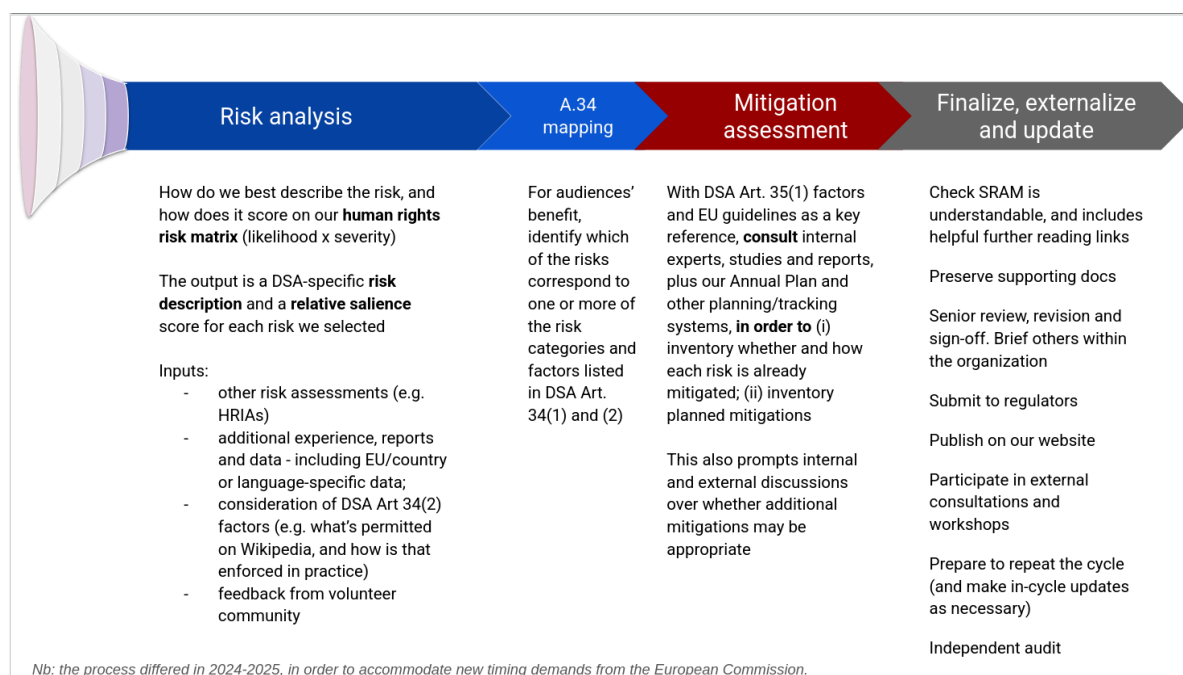


Figure 4: Process chart showing how the Wikimedia Foundation assesses risks and their mitigations (and circulates the findings to ensure scrutiny and real-world impact) as part of the DSA SRAM process.

Our primary consideration is (1) the likelihood of a risk manifesting, and (2) the impact a risk is likely to have, if it manifests. Among other factors, this takes into account the mitigations which are already in place. These components of risks are evaluated according to the Foundation’s Human Rights Risk Assessment Framework, which consists of a set of heuristic criteria and norms that are intended to reduce bias and facilitate discussions about ranking and prioritization that lead to consensus. These criteria can be found in Annex: Human Rights Risk Assessment Framework.

The distribution of Child Sexual Abuse Material (CSAM), for instance, can be damaging, but its prevalence on Wikipedia (as an EU systemic risk) appears to be low. Unlike other websites, the openness of Wikipedia, i.e. that all contributions and edits are logged publicly, and volunteer editors can identify and remove such content, makes it a highly unattractive platform for distributing such material. That transparency similarly makes Wikipedia a

relatively unattractive environment for grooming, or other types of conduct that thrive in obscurity.

Finally, our risk prioritization (including the ability to devote resources to a selected risk's mitigation, for DSA purposes) also takes into account the Foundation and wider editor and reader communities' global (non-DSA-specific) priorities. As noted above, the Wikimedia projects, globally, are exposed to important risks, such as mass censorship, or impediments to global knowledge equity, that demand significant prioritization outside of this EU, DSA-specific exercise.

For all these reasons, it is important to also note that the results of the Foundation's DSA SRAM exercises cannot be compared on a like-for-like basis with those produced by other platforms and search engines that have been designated as VLOPs and VLOSEs. For instance, the number of EU-based users exposed to some of these risks on large social media platforms could be orders of magnitude greater than on Wikipedia.

Stakeholder engagement

A significant number, though unknown precisely, of the more than 300,000 volunteers around the world that contribute to Wikimedia projects reside in the EU, participate in EU-based organizations affiliated with the Wikimedia movement, and may be directly impacted by the systemic risks identified in this exercise. The Wikimedia Foundation is committed to substantively engaging such affected stakeholders in order to meet the commitments articulated in our Human Rights Policy, particularly with regards to conducting ongoing human rights due diligence and assessing systemic risks. Indeed, such stakeholder engagement is strongly encouraged under the Digital Services Act.

In April 2025, the Foundation facilitated a stakeholder engagement session to help these stakeholders understand the Foundation's regulatory obligations under the DSA, the process by which we identify and assess systemic risks under the DSA, and the nature of the risks identified. This session occurred as a part of the 2025 [Wikimedia Europe General Assembly](#) and brought together approximately 27 representatives of affiliate organizations from across Europe. Participants were leaders in their organizations and generally possessed a thorough understanding of the issues, risks, and challenges Wikimedians experience in their countries. The session had two primary objectives: to educate participants how the Wikimedia Foundation identified and assessed systemic risks relating to Wikipedia in the European Union as required under the DSA, and to learn from participants' reflections on how they experience these risks and others in their own work.

In general, participants concluded that the systemic risks the Foundation had identified generally align with their experiences and shared some ideas for future systemic risk assessment reporting. Participants discussed harassment on Wikimedia projects and how women and LGBTQIA+ people often experience this threat disproportionately. Foundation staff also engaged participants on the subtle difference between disinformation and

misinformation and the risks they pose to European society, as well as the emerging challenges posed by generative AI. This exercise helped to inform and validate the content presented in this systemic risks assessment. In the future, the Foundation will undertake efforts to continue this kind of stakeholder engagement within the EU.

The Wikimedia Foundation also engaged stakeholders in broader civil society at the European Rights & Risks Forum hosted by the Global Network Initiative and the Digital Trust & Safety Partnership in Brussels in June 2025. This forum, which was the second of its kind, gathers civil society, academics, investors, and very large online platforms and search engines (respectively, VLOPs and VLOPSEs) subject to the DSA, to stakeholders' experiences with producing risk assessments and lessons learned. Representatives of the Wikimedia Foundation shared its unique experiences supporting Wikipedia—the only nonprofit platform designated as a VLOP— in complying with the additional requirements those platforms face under the law. Representatives also shared the above-mentioned experience working with the volunteer community to understand and provide feedback around Wikipedia's risk assessments. These exchanges provided unique opportunities for Foundation staff leading our DSA compliance work to engage directly with these stakeholders to answer questions, share experiences, and benefit from mutual learning. A summary report of this forum is available [here](#).

What are the key risks highlighted in this 2025 exercise?

Our assessment for 2025 identifies three risks of highest salience:

1. Disinformation regarding civic and electoral processes, and conflicts
2. Disinformation regarding historical/geographical narratives
3. Harassment among the volunteer community

The designation of these three risks as the most salient for Wikipedia reflects a high degree of continuity with regard to the external threat environment and our systems since 2023 and 2024. Work to address these risks has been underway across the whole communities of Wikipedia contributors and allies for some time and will continue in the coming months and years. More generally, Wikipedia is a broadly welcoming environment for contributors that results in reliable, neutrally-presented information on a vast range of topics. This is the result of longstanding efforts by the community, the Foundation, and others. But as is clear not just from our human rights impact assessments, but also community conversations and outside research, more can be done.

These most salient risks have been long-term focuses of the Wikimedia Foundation and its volunteer communities. In general terms, the Foundation's approach to mitigating these risks is as follows, though more granular elements of these approaches will be provided in our upcoming mitigation inventory for 2025:

1. **Anti-harassment efforts** are being supported by rollout of the [Universal Code of Conduct](#) and its related guidance and enforcement processes; and the Foundation is working on [a new, supplemental way for UCoC-violating incidents to be reported](#);
2. **Anti-disinformation efforts** are sufficiently diverse that the Foundation has developed an overarching Anti-Disinformation Strategy. Specific disinformation risks—for example ahead of certain elections—are responded to by supporting the formation of temporary task forces (by the volunteer community, or the Foundation, or both). This approach is backed by a program of work to improve technical tools to support moderation and content curation carried out by volunteer editors, and a [wider research program to support Knowledge Integrity](#) on Wikipedia and its sister projects. An e-learning module has also been developed and deployed to assist volunteer editors in identifying and combatting disinformation. These same efforts also assist in combating disinformation or bias that may be introduced as a result of users adding content developed through generative AI.

Other risks that were assessed and are being mitigated, as part of this exercise, include the following. More details on these risks and mitigations can be found in the Risk Register itself, and the underlying materials it is based on.

Medium Salience:

- Dangerous content: Exposure of vulnerable individuals to content that may be dangerous in that context
- Child sexual abuse material (CSAM)
- Terrorist or violent extremist content (TVEC)

Lower Salience:

- Propagation of disinformation or bias through generative AI
- Privacy risks for young contributors

Conclusion

The Wikimedia Foundation appreciates the opportunity to discuss these risks with the European Commission, members of the volunteer community that make free knowledge projects like Wikipedia possible, and, indeed, the broader public. The Foundation looks forward to engaging in sustained dialogue with these stakeholders around these risks, how it is already working to mitigate them now and in the future, and how our work to identify and mitigate risks will continue in the coming years as the digital ecosystem evolves.

This process of systemic risk assessment continues to serve as a timely incentive for the Foundation to take stock of the broader societal impacts of Wikipedia within the EU. We welcome the opportunity to describe to the Commission, and the EU citizens it represents, how our organization plans to mitigate risks that the DSA is intended to address. Indeed, we

welcome efforts by governments to work with platforms on processes that foster greater transparency and accountability for online platforms.

As discussed above, many important actions to mitigate the risks identified in this submission will require buy-in and action from Wikimedia affiliates and volunteers throughout the EU. We hope our publication of material about the risks and mitigations will ultimately spark important discussions with the Wikimedia communities about their important and essential roles in mitigating risks to themselves and European society more broadly, thereby empowering Wikimedians across Europe to exercise leadership and initiative in our collective endeavor.

Annex: Human Rights Risk Assessment Framework

| Risk | Impact | | Likelihood | |
|-----------------------------------------------------------------|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|----------------------------------------------------------------------------------------------------------------|
| Freedom of Expression and Government Censorship | 3 | Wikimedians face heightened threats for creating and editing content; Wikimedia projects are blocked; Wikimedians are charged and/or jailed for conveying information; Governments actively engage in disinformation campaigns on Wikimedia projects, including instances of project capture | 3 | Ongoing, Probable, or Imminent in the EU OR Probable and requires a long lead time for treatment |
| Online Privacy and Government Surveillance | | Wikimedians are broadly and actively surveilled by government entities, including targeted surveillance in violation of human rights norms; Personal data and private communications of Wikimedians are intercepted by government entities. | | |
| Conflict | | Severe impact to Wikimedia platforms, affiliates, people, and events, including (but not limited to): physical security threats to individuals, critical damage to physical assets, blocking of access to platforms | | |
| Diversity, Equity, & Inclusion | | Women, racial and ethnic minorities, persons with disabilities, and LGBTQIA+ individuals face pervasive legal, social, and/or cultural discrimination in online fora or in-person gatherings, including in public accommodation | | |
| Technology, Products, & Tools | | Users' access, freedom of expression, and/or | | |

| | | | | |
|-----------------------------------------------------------------|---|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|-------------------------------------|
| | | privacy online are severely impacted | | |
| Child Rights | | Minors (under the age of 18) face significant physical and significant emotional dangers, including experiencing or being exposed to sexual abuse or exploitation. | | |
| Freedom of Expression and Government Censorship | 2 | Wikimedians face limitations on their freedom of expression and may be harassed by government entities for sharing information on particular topics; Governments seek to exploit legal mechanisms to remove or alter content it finds objectionable | 2 | Possible in next 6 months in the EU |
| Online Privacy and Government Surveillance | | Some Wikimedians are surveilled by government entities for reasons beyond their involvement in the Movement; Government entities seek to use legal requests to obtain user data | | |
| Conflict | | Minor impact to platforms, affiliates, people, and events, including (but not limited to): minor damage to physical assets, temporary disturbance to individuals' everyday life, etc. | | |
| Diversity, Equity, & Inclusion | | Women, racial and ethnic minorities, persons with disabilities, and LGBTQIA+ individuals face occasional legal, social, and/or cultural discrimination in online fora or in-person gatherings, including in public accommodation | | |
| Technology, Products, & Tools | | Users' access, freedom of expression, and/or privacy online are moderately impacted | | |
| Child Rights | | Minors (under the age of 18) face emotional hazards, such as bullying, harassment, or teasing. | | |

| | | | | |
|-----------------------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|------------------------------------------------------|
| Freedom of Expression and Government Censorship | 1 | Wikimedians face minimal risks to freedom of expression; Wikimedians are not subjects of clearly rights-violating government interest | 1 | Unlikely to occur or deteriorate this year in the EU |
| Online Privacy and Government Surveillance | | Wikimedians are not generally surveilled online by government entities; Wikimedia projects are not censored; Government entities follow and respect legal processes for removing or altering content. | | |
| Conflict | | Wikimedia platforms, affiliates, people, and events are not seriously affected | | |
| Diversity, Equity, & Inclusion | | Women, racial and ethnic minorities, persons with disabilities, and LGBTQIA+ individuals face minimal legal, social, and/or cultural discrimination in online fora or in-person gatherings, including in public accommodation | | |
| Technology, Products, & Tools | | Users' access, freedom of expression, and/or privacy online are not meaningfully affected | | |
| Child Rights | | Children (minors under the age of 18) face minimal emotional hazards. | | |

*