



GitHub

88 Colin P Kelly Jr Street,
San Francisco, CA 94107
Tel: 415-448-6673 (main)

January 2018

To: United Nations Special Rapporteur David Kaye
Via email: freedex@ohchr.org
Re: **Submission to study on social media, search, and freedom of expression**

GitHub welcomes the opportunity to make a submission for your upcoming report on the effect of content regulation on free expression. More than 27 million people globally use GitHub to share and search for content. As a result, content regulation and its impact on free expression is a continuous concern for us, and we appreciate being able to contribute to a report that considers the role of content platforms.

GitHub is the world's largest software development platform, enabling users and businesses to collaboratively develop open-source and proprietary software projects. GitHub's global community includes individual developers, startups, small businesses, large companies, NGOs, and governments. GitHub-hosted software projects include applications designed for web or mobile devices, as well as the source code that powers entire businesses. Developers on GitHub work together, sharing code and knowledge, to build the future of software.

GitHub hosts content in over 77 million repositories. Users generate the content they store in repositories. They can share content they post and can search for other content on GitHub. GitHub hosts content beyond software code. People use GitHub to work on many types of content, particularly where they develop it collaboratively and when features like version control are useful.

As a platform that hosts user-generated content, GitHub promotes the freedom of expression in our policies and in our application of those policies to specific cases, consistent with international human rights law's articulation of the right to freedom of expression and its limitations in the International Covenant on Civil and Political Rights (ICCPR). In addition, GitHub promotes transparency by (i) directly engaging our users in developing those policies; (ii) explaining our reasons for making different policy decisions; (iii) notifying users when we need to restrict content, with our reasons; (iv) allowing users to appeal removal of their content; and (v) publicly posting takedown requests in real time in a dedicated repository.

In making this submission, GitHub seeks to highlight the development of our standards with visible, public engagement from our community and the transparency in our processes for content moderation. GitHub also notes that broad efforts to regulate offensive content or unauthorized distribution of audio and visual works often have unintended effects on platforms, like GitHub's, where users are typically posting different and more complex types of content, like source code. This creates



a chilling effect that increases the likelihood that content will be taken down by the platform, particularly in the case of smaller platforms that are not equipped to implement rigorous content review that would allow them to filter content narrowly.

Consistent with Article 19 of the ICCPR, GitHub promotes the right of our users to hold opinions without interference and the right to freedom of expression, including every user’s “freedom to seek, receive and impart information and ideas of all kinds, regardless of frontiers, either orally, in writing or in print, in the form of art, or through any other media of his choice.” Our limitations are consistent with Articles 19 and 20 of the ICCPR, that is, when provided by law and necessary, including for the respect of others’ rights or where the content constitutes harassment, abuse, threats, or incitement of violence towards any individual or group. To this end, GitHub’s Terms of Service establish rules we use for content moderation on our platform. In addition, GitHub’s Community Guidelines provide further guidance on the boundaries of acceptable speech on our platform.

Crucial to our credibility and legitimacy to our users, GitHub develops the Terms of Service and Community Guidelines with open, ongoing, public engagement that we invite and encourage from our users (see <https://github.com/github/site-policy>). Recognizing the importance of involving people in the development of standards that affect them, GitHub hopes our approach gets recognized as a best practice that is adopted by other platforms. To this end, GitHub also openly licenses our policies to encourage others to use them and modify them as they see fit (see <https://github.com/github/site-policy#license>).

Through our policies and their application, GitHub promotes free expression while protecting our users from abuse. As described in more detail below, when faced with illegal or abusive content, GitHub takes the utmost care to take the narrowest course of correction. For example, we provide users an opportunity to remove specific content, when possible, rather than blocking entire repositories, and we remove content only in the jurisdiction where it is shown to be illegal, rather than worldwide.

1. Company compliance with State laws

What processes have companies developed to deal with content regulation laws and measures imposed by governments?

There are two main categories of requests GitHub receives to remove or block content pursuant to law: (A) government takedown requests (based on violation of a particular government’s law) and (B) copyright infringement notices (based on the U.S. Digital Millennium Copyright Act (DMCA), section 512). By far, the most common kind of takedown request we receive is for alleged copyright infringement.

In both cases, we post the requests in a public repository in real time. GitHub does this because we are concerned about Internet censorship, and we believe that transparency on a specific and ongoing level is essential to good governance. By posting the notices publicly, we can better inform the public about what content is being withheld from GitHub, and why.



As we explain in both our government takedowns repository (<https://github.com/github/gov-takedowns>) and our DMCA takedowns repository (<https://github.com/github/dmca>), if we post a notice or request, it only means that we received it on the indicated date. It does not mean that we deem the takedown unlawful or otherwise wrong, nor does it imply the opposite – that the notice is necessarily meritorious. These notices and requests are posted only for informational purposes.

Given that GitHub is a platform where users contribute layers of content in collaboration with one another and where content often results from contributions from multiple users, we need to be especially judicious in removing content.

In addition, GitHub tracks both government and DMCA requests to remove or block user content in our transparency reports so that our users know when we have taken down content and why. We do this because viewing aggregate takedown information, as opposed to case-specific takedown information in the repositories, helps complete the picture about public and private takedown activity.

A. Government takedowns

From time to time, GitHub receives requests from governments to remove content that they judge to be unlawful in their local jurisdiction. Although GitHub may not always agree with those laws, we may need to block content if we receive a valid request from a government official. Given that not complying with a government request often means a wholesale ban of our platform in that jurisdiction, we may need to comply so that our users in that jurisdiction may continue to have access to GitHub to collaborate and build software. Whenever we block content at the request of a government, we post the official request that led to the block in a publicly accessible repository (<https://github.com/github/gov-takedowns>).

Regarding our process, when we receive a request, we confirm:

- that the request came from an official government agency;
- that the official sent an actual notice identifying the content; and
- that the official specified the source of illegality in that country.

If we believe the answer is yes to all three, we block the content in the narrowest way as we see possible, for example by restricting the removal only to the jurisdictions where the content is illegal. We then post the notice in our government takedowns repository, creating a public record where people can see that a government asked GitHub to take down content.

B. Copyright takedowns

Section 512 of the Digital Millennium Copyright Act (DMCA) sets out a process for user-generated-content platform providers to follow in order to avoid liability on their platforms when copyright owners put them on notice of hosted copyright infringement. When GitHub receives a DMCA takedown notice that is sufficiently detailed to meet the DMCA's requirements, the law requires us to take down that content. The DMCA also provides a counter-notice procedure that allows a platform to reinstate content when the alleged infringer states that the takedown was erroneous. As we explain on our website at <https://help.github.com/articles/dmca-takedown-policy/>, GitHub follows the DMCA takedown



processes for notice-and-takedown and counter-notice procedures while respecting the right of Fair Use.

To elaborate on our process, when we receive a valid DMCA takedown notice, we:

- post the notice in a publicly accessible repository (<https://github.com/github/dmca>);
- contact the user to request that s/he modify or delete the content that is identified as infringing the copyright of another within 24 hours (or, if the notice alleges that the contents of an entire repository infringes copyright, we will disable the repository); and
- notify the copyright owner that we have taken this step. (Or, if the notice alleges that the contents of an entire repository infringes copyright, we will disable the repository.)

If the user decides to make changes, s/he must tell us within the 24-hour period. If not, we will disable the repository. If the user notifies us that s/he made changes, we will verify that the changes were made and then notify the copyright owner.

If a user believes that their content was disabled as a result of a mistake or misidentification, they may send us a counter notice. As with the original notice, we will make sure that the counter notice is sufficiently detailed according to the DMCA and, if so, will post it to our public repository and pass the notice back to the copyright owner by sending them the link.

Thus, GitHub will only disable a user's content if: (i) the copyright owner has identified the entire repository as infringing; (ii) the user has not made any changes after being given an opportunity to do so; or (iii) the copyright owner has renewed their takedown notice after the user had a chance to make changes. If the copyright owner chooses instead to revise the notice, we will repeat the process as if the revised notice were a new notice.

To highlight how our approach minimizes restricting speech and promotes transparency:

- we only act on a takedown notice when it is sufficiently detailed according to the DMCA;
- we post the notice in a public repository;
- we provide the affected user with the opportunity to modify or delete the content rather than removing it directly ourselves (unless an entire repository is infringing);
- when we give the user a chance to make changes, we notify the copyright owner and send a link to the revised content; and
- if an affected user wishes to appeal the decision, s/he can send us a counter notice, which we will process according to the same steps as an original notice (provided it is sufficiently detailed according to the DMCA) and post to our public repository.

How should companies respond to State content regulation laws and measures that may be inconsistent with international human rights standards?

Regarding what companies should do when they confront this situation, while we recognize the challenges posed and do not claim to know how to fully address them, our view is companies should look at exactly what they are asked to do and exactly what that State's law requires, including to whom it applies and its jurisdictional limitations. In some cases, a State may be asking a company to do something that is beyond the scope of that law or that State's jurisdiction.



Where a State's content regulation laws and measures may be inconsistent with international human rights standards, we believe companies should note this inconsistency and look for ways to comply that are the least restrictive on human rights (for example by removing only the minimum amount of content in question or by geo-blocking content only in the jurisdiction(s) in question).

This is the approach we take at GitHub. In addition, given that so many of our users rely on GitHub to do their work and keep widely used services running, we also consider the social and economic importance that our services represent, and what taking them away would mean to our users. We believe that other similarly situated platforms should take that into account as well.

2. Other State Requests

Do companies handle State requests for content removals under their terms of service differently from those made by non-State actors? Do companies receive any kind of content-related requests from States other than those based on law or the company's terms of service (for example, requests for collaboration with counter speech measures)?

When States request content removals, they invariably claim that the content violates a State law, rather than one of GitHub's Terms of Service. Our Terms of Service prohibit unlawful content, so if a State actor were to report unlawful content as a violation of our Terms of Service, we would process that as a government takedown request. That includes confirming that the request is coming from a genuine State official and posting the government's request. Conversely, when a non-State actor reports unlawful content as violating our Terms of Service, we do not feel compelled to take it down in the same way that we do when the takedown is demanded by a State official. GitHub does not receive other kinds of content-related requests from States.

3. Global removals

How do / should companies deal with demands in one jurisdiction to take down content so that it is inaccessible in other jurisdictions (e.g., globally)?

When we determine a need to remove content, GitHub geo-blocks content so that we only restrict access in the jurisdiction that reported the content as illegal. We do this in an effort to balance a country's sovereignty in its jurisdiction with the right to free speech outside of that jurisdiction.

4. Individuals at risk

Do company standards adequately reflect the interests of users who face particular risks on the basis of religious, racial, ethnic, national, gender, sexual orientation or other forms discrimination?

GitHub's Community Guidelines explicitly call out the bases for the particular risks mentioned in this question, in addition to others: "age, body size, disability, ethnicity, gender identity and expression, level of experience, nationality, personal appearance, race, religion, or sexual identity and orientation."



GitHub designed the Community Guidelines to guard the interests of marginalized groups and to encourage members of communities to respect one another.

We are committed to the continuous guarantee that GitHub is a welcoming environment for all the different voices, perspectives, and code contributions in our community, while maintaining a space where people are free to express themselves. In addition to the thoughtful review of our Terms of Service Support and Legal Support teams, GitHub has a dedicated Community and Safety team committed to making GitHub safer and more welcoming to users of diverse backgrounds. GitHub's Community Guidelines elaborate on our approach (see <https://help.github.com/articles/github-community-guidelines/>).

We developed our Community Guidelines with feedback from our community through a review and notice period (see <https://github.com/blog/2267-introducing-github-community-guidelines>) and open-sourced them along with the other policies that govern use of our site. Open-sourcing in this context means we posted them online (on GitHub) to allow our users and customers to review, comment, and suggest changes, and we gave them rights to adapt and reuse the policies so that together we can promote new and better standards in platforms' site policies, including terms of service.

In addition to encouraging our community to engage in developing the standards we and others use, we offer recommendations on how to encourage inclusivity and at the same time protect users from abuse. For example, we post recommendations on how to minimize abuse (<https://githubengineering.com/community-and-safety-feature-reviews/>), develop community standards and codes of conduct for a project (<https://github.com/blog/2380-new-community-tools>), and encourage community participation (<https://github.com/blog/2397-making-it-easier-to-grow-communities-on-github>) for their projects.

5. Content regulation processes

What processes are employed by companies in their implementation of content restrictions and takedowns, or suspension of accounts?

GitHub's Terms of Service and Community Guidelines establish for our users what content we permit on GitHub. Section 512 of the DMCA (discussed in question 1) sets out a process GitHub can follow to avoid liability when users notify us of allegedly copyright-infringing content on our platform. Based on our processes implementing those standards and consistent with Articles 19 and 20 of the ICCPR, we only remove content that violates our Terms of Service, our Community Guidelines, or the law. All content review is done by GitHub employees and we do not use automated filtering.

To promote users' free expression rights while balancing them against illegality, abuse, and copyright infringement, we prefer to avoid destructive actions like permanently deleting content, and to hide content rather than permanently removing it, until we have had sufficient time for the relevant teams to review the content in question.

We consider the nature of the violation in determining whether to hide content, disable a repository, or suspend an account. We aim to hide or disable *only* the infringing content. Given the nature of our



platform, we need to be especially judicious with respect to takedowns because the content in question is often built of layers of contributions by multiple users.

We track the reports we receive and our responses to them.

6. Bias and non-discrimination

How do companies take into account cultural particularities, social norms, artistic value, and other relevant interests when evaluating compliance with terms of service? Is there variation across jurisdictions? What safeguards have companies adopted to prevent or redress the takedown of permissible content?

In evaluating compliance with our Terms of Service, recognizing the global nature of our platform, we take great care to consider a broad range of factors and circumstances that might influence a user's understanding of those terms and our understanding of how that content may or may not violate our terms. We also bring in additional reviewers and teams as necessary to broaden the perspectives in the review. In continuously developing our Community Guidelines with collaboration from our community, we engage directly with people who use our platform, and they can help us approach our terms from varying viewpoints and incorporate them into our terms. As for safeguards, beyond analyzing reported violations carefully and thoughtfully, we provide users the ability to appeal takedown of content.

7. Appeals and remedies

How should companies enable users to appeal mistaken or inappropriate restrictions, takedowns or account suspensions? What grievance mechanisms or remedies do companies provide?

Regarding enabling user appeals, our view is that companies should provide a means for a user to contest a decision to remove content, with reasons for the decision and an opportunity to discuss it. Companies should also provide more than one level of review of the decision as appropriate.

GitHub notifies users when we determine content needs to be removed. Users can contact us to appeal the decision, and in the case of the DMCA, can file a counter notice to reinstate the content. We provide multiple levels of review of a decision as appropriate. We post notices leading to takedowns in a public repository and track them in our transparency report.

8. Automation and content moderation

What role does automation or algorithmic filtering play in regulating content? How should technology as well as human and other resources be employed to standardize content regulation on platforms?

As noted above, GitHub does not use automation or algorithmic filtering for content moderation. All content review is done by GitHub employees. As a broader point, using automation or algorithmic filtering for content moderation is not necessarily something we even *could* do effectively, at least with existing technology. For instance, for some kinds of content, filtering technologies might not



necessarily be effective because they sweep up permissible content as false positives. Thus, when making rules that apply to various kinds of content, it is necessary to think about all kinds of content affected and to recognize that certain approaches might not work for all types of content.

With this in mind, in response to what role technology, as well as human and other resources, should play in standardizing content regulation on platforms, our view is that companies should aim to strike a balance between protecting their users in a scalable way while not leaving adjudication to potentially discriminatory or otherwise ineffective algorithms. In addition, we feel companies should have well-trained support staff who analyze and respond to reports.

Furthermore, we believe companies' approaches should correlate with the nature of their platforms and the risks posed by activities on their platforms. For instance, GitHub does not host advertisements, so moderation practices that focus on the use of algorithmic filtering regarding online advertisements are not relevant to GitHub.

Similarly, requirements on companies and Internet users should correlate with the nature of the activities and associated risks. Yet, presently, we see examples where this is not the case. Some proposed requirements are overly broad in their scope and, as applied to GitHub and our user community, could be so cumbersome as to prevent developers from being able to launch their work. (For example, Article 13 of the proposed EU Directive on Copyright in the Digital Single Market is intended to address mass unauthorized distribution of audio and visual works yet its requirements would require platforms used to share code to proactively monitor content using upload filters.)

9. Transparency

Are users notified about content restrictions, takedowns, and account suspensions? Are they notified of the reasons for such action? Are they notified about the procedure they must follow to seek reversal of such action?

Transparency is highly important to GitHub. We notify users about content restrictions, takedowns, and account suspensions, and, when we determine a need to remove content, we provide reasons for our decision, with the ability for users to contact us to appeal the decision.

In addition, we engage our community in developing the standards themselves. GitHub's Community Guidelines and Terms of Service, which explain what is acceptable (and not acceptable) content on GitHub, are two of many GitHub site policies that we open-sourced to enable ongoing community review, as well as adoption by other companies. Users are free to suggest changes to the policies themselves, and to ask questions about the policies and their interpretation that the legal team addresses publicly.

Our transparency report tracks the number of takedown requests (requests to remove or block user content) that we have received in a given year and how we have handled them. We track both government takedown requests and DMCA takedown notices. (We also track requests to disclose user information, including subpoenas, court orders, and search warrants, as well as national security orders.) As explained in question 1, we also post the takedown requests we receive so that users



know that content was removed and why. We believe presenting the information in those two ways (real-time specific, aggregate) is very important in order to accomplish effective transparency.

10. Examples

Please share any examples of content regulation that raise freedom of expression concerns (e.g., account suspension or deactivation, post or video takedown, etc.), including as much detail as possible.

A. Government takedowns

In 2014, GitHub received requests from the Russian Federal Service for Supervision in the Sphere of Telecom, Information Technologies and Mass Communications (Roskomnadzor) to censor certain content on GitHub.com. A user had uploaded a file called “suicide.txt” that listed, in Russian, a number of ways one could kill themselves. Roskomnadzor demanded that we block access to the content under the laws prohibiting “information distributed over the Internet about the methods of suicide, as well as appeals to commit[] suicide” (Decree #1101 of the Government of the Russian Federation dated October 26, 2012).

After we were put on Roskomnadzor’s blacklist, the user deleted the file on his or her own, and GitHub was removed from the blacklist, without any involvement from GitHub. Soon after, GitHub received a number of new complaints from Roskomnadzor. We contacted the users and asked them to remove the files. After receiving reports that Github.com was being blocked in parts of Russia, we responded by geo-blocking the requested content in Russia only. Roskomnadzor took us off of the blacklist as soon as we responded.

B. Copyright takedowns

While government takedown requests are frequently known to raise freedom of speech concerns, requests to take down information for alleged copyright infringement also raise these concerns. In the U.S. DMCA’s notice-and-takedown process, the inclusion of a counter-notice procedure is an important way to check attempts to unduly restrict freedom of speech. The example below illustrates these aspects of the DMCA process.

In 2017, a company sent GitHub a DMCA notice to take down a case study that was publicly shared on GitHub but that the company had created for its internal use in interviews. GitHub verified that the notice came from the copyright owner or a person authorized to act on their behalf, clearly described what work was allegedly infringed, clearly identified the specific files or repository to be taken down, noted whether there were any options other than removal, included contact information, stated a good faith belief that the use was not authorized, taking into account fair use, and was signed, attesting to accuracy of the information in the notice. GitHub then took down the content and notified both the user and the claimant of the takedown.

The alleged infringer then filed a counter notice disputing the takedown as a violation of a proprietary license because the content was under an open source license. GitHub verified that the person



submitting the counter notice answered all questions in sufficient detail, including that the person identified the disabled content and where it was located; included their contact information; said that they had a good faith belief that the material was removed or disabled as a result of mistake or misidentification of the material; and provided a physical or e-signature. GitHub then notified the claimant of the counter notice and that the content would be re-enabled unless GitHub received a notice of a court action within ten business days. After ten business days, GitHub restored access.

* * *

GitHub is a place where millions of people share and search for content. In moderating content on our platform, GitHub espouses free expression and transparency while recognizing our obligations to comply with the law and to protect users from abuse. Thus, as explained in this submission, GitHub promotes free expression consistent with international human rights law, and, when we need to remove content, we aim to do so in the narrowest way possible to comply with the law and fend off abuse. In the interest of transparency, we involve users in the creation of the standards we use for content moderation; we notify users and provide reasons and ability to appeal removal of content; and we post takedown notices in a public repository in real time and track them in our transparency report.

We thank you again for the opportunity to contribute to your report on the effect of content regulation on free expression and to inform discussion about the role of platforms in this context.