




Who Has Your Back?

CENSORSHIP EDITION 2018



ELECTRONIC FRONTIER FOUNDATION

Authors: Nate Cardozo, Andrew Crocker, Gennie Gebhart, Jennifer Lynch, Kurt Opsahl, Jillian C. York

With assistance from: Michelle Chang, Hugh D'Andrade, Lena Gunn, Jason Kelley, Laura Schatzkin

A publication of the Electronic Frontier Foundation, 2018.

“Who Has Your Back? 2018” is released under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

View this report online: <https://www.eff.org/who-has-your-back-2018>


















Who Has Your Back?

CENSORSHIP EDITION 2018

Nate Cardozo, Senior Staff Attorney
Andrew Crocker, Staff Attorney
Gennie Gebhart, Researcher
Jennifer Lynch, Senior Staff Attorney
Kurt Opsahl, Deputy Executive Director and General Counsel
Jillian C. York, Director for International Freedom of Expression

MAY 2018

	Legal takedown requests	Platform policy takedown requests	Provides meaningful notice	Allows appeals	Limits geographic scope
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★
	★	★	★	★	★

Contents

Executive Summary	7
Introduction: Government Censorship in the Age of HTTPS	8
Scope	9
Major Findings and Trends	9
Overview of Criteria	11
Transparent About Legal Takedown Requests	11
Transparent About Platform Policy Takedown Requests	12
Provides Meaningful Notice	12
Allows Appeals	13
Limits Geographic Scope	13
Looking Ahead	13
Company Reports	14
Apple App Store	14
Dailymotion	15
Facebook	16
Google+	17
Google Play Store	18
Instagram	19
LinkedIn	20
Medium	21
Pinterest	23

Reddit.....	24
Snap.....	25
Tumblr.....	26
Twitter.....	27
Vimeo.....	29
WordPress.com.....	30
YouTube.....	31
References and helpful links.....	33

Executive Summary

We are at a critical moment for free expression online and for the role of Internet intermediaries in the fabric of democratic societies. In particular, governments around the world have been pushing companies to take down more speech than ever before. What responsibilities do the platforms that directly host our speech have to protect—or take down—certain types of expression when the government comes knocking?

The first step toward answering that question is transparency. How often are governments asking companies to remove speech, and how do the companies handle those demands? Furthermore, what processes do companies afford to users whose content is removed and whose accounts are suspended?

Given policymakers' and the public's intense focus on cracking down on speech they consider undesirable, this year's Who Has Your Back report features substantially redesigned categories and criteria. Since the Electronic Frontier Foundation began publishing Who Has Your Back in 2011, it has generally focused on the practices of major consumer-facing Internet companies regarding government requests to produce user data. This year, we shift our focus to companies' responses to government requests to take down user content and suspend user accounts.

For our 2018 report, we assess companies' policies against five all-new criteria:

- Transparency in reporting government takedown requests based on both **legal requests** and requests alleging **platform policy violations**
- Providing **meaningful notice** to users of every content takedown and account suspension
- Providing users with an **appeals process** to dispute takedowns and suspensions
- **Limiting the geographic scope** of takedowns when possible

Three platforms—the Apple App Store, Google Play Store, and YouTube—earned stars in all five of these categories. And three more—Medium, Reddit, and WordPress.com—earned stars in all but the notice category, which proved the most challenging category for the companies we assessed. Some companies fell notably short overall; Facebook's and Instagram's policies in particular lagged behind comparable tech companies and social networks. However, it's clear that public pressure is resulting in real change in corporate policy and practice. We look forward to more long-term improvements across the industry in future years as companies take steps to be more accountable to their users and those users' right to freedom of expression.

Introduction: Government Censorship in the Age of HTTPS

Private censorship hasn't always been governments' tool of choice for blocking information. But for governments interested in suppressing information online, the old methods of direct censorship are getting less and less effective. In the past, governments could pressure ISPs, hijack domain name queries, and otherwise directly take down online content they wanted to censor. So why are governments increasingly asking online platforms to do their censorship work for them?

The first part of the answer is HTTPS. When HTTPS encrypts your browsing, it doesn't just protect the contents of the communication between your browser and the websites you visit. It also protects the specific pages on those sites, preventing censors from seeing and blocking anything "after the slash" in a URL.

This means that if a critical video or controversial picture shows up on a website, government censors cannot identify and block just the pages on which it appears. In an HTTPS-only world that makes such granularized censorship impossible, the government's only direct censorship option is to block the site entirely.

That might still leave governments with tenable censorship options if critical speech and dissenting activity only happened on certain sites, like devoted blogs or message boards. A government could try to get away with blocking such sites wholesale without disrupting users outside a certain targeted political sphere.

This brings us to the second part of the answer as to why governments can't censor like they used to: **mixed-use social media sites**. All sorts of user-generated content—from calls to revolution to cat pictures—are converging on mixed-use websites, which people of all backgrounds use and rely on. When content is both HTTPS-encrypted and on a mixed-use social media site, it can be too politically expensive for a government to block the whole site.

Instead, the only option left is for the government to request that the platform engage in targeted removals at the government's request.

Under the First Amendment, intermediaries generally have a right to decide what kinds of expression they will carry. But just because companies can act as judge and jury doesn't mean they should. If and when companies do comply with government requests to remove content or suspend accounts, these decisions must be transparent to their users and the general public.

Compared to governments directly blocking content, private censorship can effectively disappear information across the entire web rather than merely within a certain country's borders. The stakes are high, especially in unstable political environments and for those living under repressive regimes. The targets of takedown requests include

journalists, activists, and dissidents, and requests to take down their content or block their pages often serve as an ominous prelude to further government targeting.

In an HTTP environment where governments could directly filter Internet content, entire websites could disappear behind obscure and misleading error messages. But now, on today's increasingly HTTPS-secured, social media-dominated web, transparency from tech companies about the pressure they face may give us increased visibility into both corporate and government censorship decisions.

Scope

This report provides objective measurements for analyzing the policies of major technology companies when it comes to government-requested censorship. We focus on a handful of specific, measurable criteria that reflect attainable best practices.

We assess those criteria for some of the biggest online platforms that publicly host a large amount of user-generated content. The group of companies and platforms we assess does not include infrastructure providers (e.g. Cloudflare), storage (e.g. Dropbox, Google Drive), communications providers (e.g., Gmail, Outlook), or search engines (e.g. Bing, Google).

The scope of this report does not include several types of censorship. We do not cover removals of child exploitation imagery, or intellectual property removals, restrictions, and reporting. Further, for the two “app stores” we evaluate this year, we limit the scope of our review to developer accounts and the apps themselves.

We also distinguish between how companies censor at the request of governments from how they censor independently. The latter type of private censorship is not the focus of this report. However, many of this report's criteria—including transparent reporting, meaningful notice, and appeals—serve to encourage best corporate practices regardless of whether censorship is driven by governments, companies, or a combination of the two.

Major Findings and Trends

Our major findings include:

- Three companies—the **Apple App Store**, **Google Play Store**, and **YouTube**—are receiving credit in all five categories.
- The companies that got four out of five stars—**Medium**, **Reddit**, and **WordPress.com**—all missed the same category: meaningful notice.
- **Facebook's** and **Instagram's** policies fall short of those of other similar technology companies.

- Although many large tech players did not score well in our assessment, it is clear that public pressure and scrutiny around content moderation has paid off.

We are pleased to announce that three companies earned stars in every category we evaluated in this year's report: the Apple App Store, Google Play Store, and YouTube. And three more earned stars for all but one of this year's categories: Medium, Reddit, and WordPress.com. These companies serve to provide examples of strong policy language for others hoping to raise the bar on content moderation policy.

Notably, two of the highest-scoring companies are app stores. This means that the users submitting content are typically app developers with strong incentives to comply with guidelines, and that the companies have one primary type of content—apps—to moderate. The Apple App Store and Google Play Store also operate at a relatively smaller scale compared to large social media networks. These variables combine to make it easier for these companies to implement the requirements of this year's report, particularly providing users with meaningful notice and an appeals process.

Also notably, YouTube gets a perfect score as well. As a massive platform with an incredibly diverse user base, YouTube provides a good counterexample to the two app stores. If YouTube can achieve this level of transparency with regard to content takedowns and account suspensions, it's reasonable to expect other companies of its size and scope to also meet that standard.

However, other major companies fell short in more than one category. In particular, Facebook's and Instagram's policies lag behind other social networks and large tech companies. Facebook's transparency report does not meet our requirement to provide detail on all government-requested content removals. Further, Facebook does not commit to providing users with meaningful notice of, nor an appeals process to dispute, many types of content takedowns and account suspensions. Given that Facebook is the site of some of the most controversial content moderation decisions and has the largest number of users among platforms of its type, its decision to deny its users more comprehensive notice and appeals is particularly concerning.

Indeed, providing users with meaningful notice of content takedowns and account suspensions posed the greatest challenge for the companies we assessed. The three companies that earned four out of five stars all missed only the notice category. Notice is a critical component of accountable content moderation. With timely, informative, on-the-record notice, users are in a better position to appeal to companies directly as well as to draw public attention to government targeting and censorship.

Although many companies fell short on notice and other categories, it is clear that intense public scrutiny around content moderation is overall leading to concrete changes for the better in corporate policies. For example, Facebook recently made its internal moderator guidelines more public, and released a preliminary report on its community standards enforcement. In another first, Twitter began reporting recurring terms of service takedowns as well as specific details regarding terrorism-related takedowns.

Content moderation policy and ethics is a complex area that cannot fit neatly into just five criteria. This report is just a beginning snapshot as we look forward to more long-term improvements across the industry.

Overview of Criteria

Only publicly available statements can qualify for credit in this report. Positions, practices, or policies that are conveyed privately or internal corporate standards, regardless of how laudable, are not factored into our decisions to award companies credit in any category.

Requiring public documentation serves several purposes. First, it ensures that companies cannot quietly change an internal practice in the future in response to government pressure, but must also change their publicly posted policies—which observers can note and document. Second, by asking companies to put their positions in writing, we can examine each policy closely and prompt a larger public conversation about what standards tech companies should strive for. Third, it helps companies review one another’s policies around content moderation, which can serve as a guide for startups and others looking for examples of best practices.

In this report, we strive to offer ambitious but practical standards. To that end, we only include criteria that at least one major company has already adopted. This ensures that we are highlighting existing and achievable, rather than theoretical, best practices.

We analyzed five criteria for this report:

Transparent About Legal Takedown Requests

To earn a star in this category, the service provider must regularly publish records of government requests for takedowns based on claims of legal violations, for instance in its transparency report. This should include, at a minimum, the information necessary to determine:

- the number of requests received,
- the country from which the request originated, and
- the number of requests acted upon and/or the number of posts removed or restricted or accounts suspended.

Takedown requests include requests to restrict public access to the post, including geographic limitation, and account suspensions that limit access to posts for a period of time.

A request is categorized as a “government request” if it is provided through official channels (such as an order issued by a competent judicial authority); if the requestor identifies themselves as a government official or relies upon their governmental

position or authority; or if the provider otherwise is aware a government is being represented in the request.

If the service provider is restricted by applicable law from disclosing the request, it may delay including the request until that restriction is lifted and still get credit in this category.

Transparent About Platform Policy Takedown Requests

To earn a star in this category, the service provider must regularly publish records of content or account restrictions based upon identifiable government allegations of violations of the provider's policies, such as Terms of Service or Community Standards, regardless of whether the request came through channels for government requests or through customer service channels. This includes government requests alleging facts that lead to a content or account restriction based on a provider's policies.

The provider's reporting should include, at a minimum, the information necessary to determine:

- the number of requests received,
- the country from which the request originated, and
- the number of requests acted upon and/or the number of posts removed or restricted or the number of accounts suspended.

A request is identifiable from a government if it is provided through official channels (such as an order issued by a competent judicial authority); if the requestor identifies themselves as a government official or relies upon their governmental position or authority; or if the provider otherwise is aware a government is being represented in the request.

Provides Meaningful Notice

To earn a star in this category, the service provider must publicly commit to provide meaningful notice to users of every removal and suspension, unless prohibited by law, in very narrow and defined emergency situations,¹ or if doing so would be futile or ineffective.²

For legal takedowns, the notice must (1) identify the specific content that allegedly violates the law, and (2) inform the user that it was a legal takedown request. For policy takedowns, this notice must (1) identify the specific content that allegedly violates a provider policy, and (2) include the specific provider policy the content allegedly violates.

¹ The exceptions should not be broader than the emergency exceptions provided in the Electronic Communications Privacy Act, 18 USC § 2702 (b)(8).

² An example of a futile scenario would be if a user's account has been compromised or their mobile device stolen, and informing the "user" would concurrently—or only—inform the attacker.

This category excludes spam, phishing, and child exploitation imagery.

Allows Appeals

To earn a star in this category, the service provider must publicly commit to **respect due process** by providing users with an appeals process. This process must provide users with effective mechanisms to appeal all provider-policy based content and account restriction decisions, including *during* temporary suspensions. Upon a successful appeal, the account or material must be reinstated promptly.

This category excludes spam, phishing, and child exploitation imagery.

Limits Geographic Scope

To earn a star in this category, the service provider must publicly commit to reasonable efforts (such as country-specific domains or relying upon user-provided location information) to **limit legally ordered content restrictions** to jurisdictions where the provider has a good-faith belief that it is legally required to restrict the content.

Looking Ahead

We hope this report will encourage more tech companies and platforms to adopt principles of transparency, notice, appeal, and limited scope with regards to government-ordered censorship. Otherwise, tech giants can misuse their power not only to silence vulnerable speakers, but also to obscure how that censorship takes place and who demanded it.

Company Reports

Apple App Store



Transparent About Legal Takedown Requests. Apple has publicly committed to reporting government takedowns in its future transparency reports:

Starting with the Transparency Report period July 1 - December 31 2018, Apple will report on Government requests to take down Apps from the App Store in instances related to alleged violations of legal and/or policy provisions.

Transparent About Platform Policy Takedown Requests. Apple has publicly committed to reporting government takedowns in its future transparency reports:

Starting with the Transparency Report period July 1 - December 31 2018, Apple will report on Government requests to take down Apps from the App Store in instances related to alleged violations of legal and/or policy provisions.

Provides Meaningful Notice. The Apple App Store publicly commits to notifying users of every app removal with the reason for removal:

Apple sometimes receives notices that require us to remove content on the App Store. We may also remove content for the reasons set forth in the App Review Guidelines or any of our agreements with you. Apple will notify you when and why an app is removed from sale, with the exception of situations in which notification would be futile or ineffective, could cause potential danger of serious physical injury, could compromise Apple's ability to detect developer violations, or in instances related to violations for spam, phishing, and child exploitation imagery.

For account suspensions, Apple does not lock or disable accounts except in relation to security issues.

Allows Appeals. The Apple App Store allows users to appeal app removals, as well as app rejections.

Apple also provides a contact support form for locked or disabled accounts.

Limits Geographic Scope. The Apple App Store has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

Whenever possible, apps that are removed from the App Store will only be removed in countries and territories specific to the issue, and will remain available in locations that aren't impacted.

Dailymotion



Note: Despite attempts to contact Dailymotion via their contact form, Twitter, and LinkedIn, we did not receive responses.

Transparent About Legal Takedown Requests. Dailymotion does not publish a transparency report.

Transparent About Platform Policy Takedown Requests. Dailymotion does not publish a transparency report.

Provides Meaningful Notice. Dailymotion does not publicly commit to providing meaningful notice to users of every removal and suspension.

Allows Appeals. Dailymotion does not have a published policy or process for users to appeal takedowns and suspensions.

Limits Geographic Scope. Dailymotion does not have a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required.

Facebook



Transparent About Legal Takedown Requests. While Facebook produces a transparency report on government takedown requests that breaks requests down by country and reports the number of pieces of content removed, it does not report the total number of government takedown requests received.

Transparent About Platform Policy Takedown Requests. While Facebook produces a transparency report on government takedown requests, that report does not include platform policy-based requests.

And while Facebook has produced a preliminary report of its own Community Standards enforcement, that report does not specify instances in which Community Standards violations are reported by governments.

Provides Meaningful Notice. While Facebook has a published policy of notifying users of content takedowns, it limits that notice only to certain categories of policy-based restrictions: "posts that were removed for nudity/sexual activity, hate speech or graphic violence."

Allows Appeals. While Facebook allows users to appeal takedowns, and further commits to timely human review of those appeals, it limits the option to appeal only to certain categories of policy-based restrictions: "posts that were removed for nudity/sexual activity, hate speech or graphic violence."

Limits Geographic Scope. Facebook has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

When we restrict content based on local law, we do so only in the country or region where it is alleged to be illegal.

Google+



Transparent About Legal Takedown Requests. Google publishes a transparency report that includes all government takedown requests. The transparency report states:

Governments contact Google with content removal requests for a number of reasons. Government bodies may claim that content violates a local law, and include court orders that are often not directed at Google with their requests. Both types of requests are counted in this report.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests.

Transparent About Platform Policy Takedown Requests. Google publishes a transparency report that includes all government takedown requests. The transparency report states:

We also include government requests to review content to determine if it violates our own product community guidelines and content policies.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests

Provides Meaningful Notice. Google+ does not publicly commit to providing meaningful notice to users of every removal and suspension.

Allows Appeals. While Google+ allows appeals for account suspensions, it does not have a published policy or process for users to appeal content takedowns.

Limits Geographic Scope. Google has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required. Its government requests FAQ states:

Where possible, we remove or restrict access to the content in the country where it is deemed to be illegal.

Google Play Store



Transparent About Legal Takedown Requests. Google publishes a transparency report that includes all government takedown requests. The transparency report states:

Governments contact Google with content removal requests for a number of reasons. Government bodies may claim that content violates a local law, and include court orders that are often not directed at Google with their requests. Both types of requests are counted in this report.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests

Transparent About Platform Policy Takedown Requests. Google publishes a transparency report that includes all government takedown requests. The transparency report states:

We also include government requests to review content to determine if it violates our own product community guidelines and content policies.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests

Provides Meaningful Notice. Google Play publicly commits to notifying users of every app takedown and account termination with the reason for removal.

For takedowns:

If your app violates any of our policies [which include policy against illegal activity], it will be removed from Google Play, and you will receive an email notification with the specific reason for removal.

For account termination, the Google Play Store publicly states it will include a “reason for termination” in the “email sent to your registered developer email address.”

Allows Appeals. Google Play allows users to appeal app rejections, removals, and suspensions, as well as account termination.

Limits Geographic Scope. Google has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required. Its government requests FAQ states:

Where possible, we remove or restrict access to the content in the country where it is deemed to be illegal.

Instagram



Instagram

Transparent About Legal Takedown Requests. While Instagram’s parent company Facebook produces a transparency report on government takedown requests that breaks requests down by country and reports the number of pieces of content removed, it does not report the total number of government takedown requests received.

Transparent About Platform Policy Takedown Requests. While Instagram’s parent company Facebook produces a transparency report on government takedown requests, that report does not include platform policy-based requests.

And while Facebook has produced a preliminary report of its own Community Standards enforcement, that report does not specify instances in which Community Standards violations are reported by governments.

Provides Meaningful Notice. While Instagram’s parent company Facebook has a published policy of notifying users of content takedowns, it limits that notice only to certain categories of policy-based restrictions: “posts that were removed for nudity/sexual activity, hate speech or graphic violence.”

Allows Appeals. While Instagram’s parent company Facebook allows users to appeal takedowns, and further commits to timely human review of those appeals, it limits the option to appeal only to certain categories of policy-based restrictions: “posts that were removed for nudity/sexual activity, hate speech or graphic violence.”

Limits Geographic Scope. Instagram’s parent company Facebook has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

When we restrict content based on local law, we do so only in the country or region where it is alleged to be illegal.

LinkedIn



Transparent About Legal Takedown Requests. LinkedIn’s transparency report does not include government takedown requests.

Transparent About Platform Policy Takedown Requests. LinkedIn’s transparency report does not include government takedown requests.

Provides Meaningful Notice. LinkedIn does not publicly commit to providing meaningful notice to users of every removal and suspension.

LinkedIn’s commitment to transparency regarding content blocked from the site states (emphasis added):

When we block content that you have authored due to the local legal requirements of the country from which you access LinkedIn, we will attempt to provide you with a notification that your content has been blocked. LinkedIn would provide this notice to the primary email address that you gave to LinkedIn

or through a message on the site. In some cases, local legal requirements may prevent us from providing you with a notification that your content has been blocked.

However, that same page also states (emphasis added):

If your content or the content you attempt to access has been blocked by LinkedIn in all locations because we believe the content is illegal or violates the terms of our [User Agreement](#) and/or [Professional Community Guidelines](#), you may not receive a notification that this content was removed.

Allows Appeals. LinkedIn allows users to appeal takedowns and suspensions:

If your account has been restricted or content removed and you believe the action was in error, you can appeal your case and we'll review your account.

Limits Geographic Scope. While LinkedIn's notice policies suggest that it aims to limit the geographic scope of content restrictions when complying with legal takedown requests, it does not have a published policy that explicitly states it will limit legally ordered content restrictions to jurisdictions where such restriction is required.

Medium



Medium

Transparent About Legal Takedown Requests. Medium sends all takedown requests it receives to Lumen (formerly Chilling Effects), a database for collecting and documenting legal complaints and takedown requests for online content. Its rules state:

Medium submits to the Lumen database government requests to restrict access to content (redacted where appropriate to protect privacy or prevent harm to a person), regardless of what or whether action is taken on the request. This includes government requests to review content to determine if it violates these Rules or other Medium content policies.

Each record specifies the country from which the request originated and the URL in question, as well as an explanation for the request, the law or regulation that motivated it, and the government agency that made the request.

Transparent About Platform Policy Takedown Requests. Medium sends all takedown requests it receives to Lumen (formerly Chilling Effects), a database for collecting and documenting legal complaints and takedown requests for online content. Its rules state:

Medium submits to the Lumen database government requests to restrict access to content (redacted where appropriate to protect privacy or prevent harm to a person), regardless of what or whether action is taken on the request. This includes government requests to review content to determine if it violates these Rules or other Medium content policies.

Each record specifies the country from which the request originated and the URL in question, as well as an explanation for the request, the law or regulation that motivated it, and the government agency that made the request.

Provides Meaningful Notice. While Medium has a policy of advance notice before taking down content, as well as a policy of notice specifically for government takedown requests, it does not publicly commit to specifying in that notice the specific content in question and the legal or policy reason for taking it down. Medium also does not commit to providing notice for account suspensions.

Regarding notice before disabling content, Medium's rules state:

Before disabling content associated with your account, we will give you advance notice, unless we believe your account is automated or operating in bad faith, or that notifying you is likely to cause, maintain or exacerbate harm to someone.

Regarding notice for government takedown requests, Medium's rules state:

If Medium receives a request from a government actor to restrict access to content associated with your account, we will notify you unless we are prohibited by law or believe doing so may endanger others.

Allows Appeals. Medium allows users to appeal takedowns and suspensions:

If you believe your content or account have been restricted or disabled in error, or believe there is relevant context we were not aware of in reaching our determination, you can write to us at yourfriends@medium.com. We will consider all good faith efforts to appeal.

Limits Geographic Scope. Medium has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required. Its rules state:

Where applicable, we will work to limit legally-ordered content restrictions to jurisdictions where we have a good faith belief that we are legally required to restrict the content.

Pinterest



Transparent About Legal Takedown Requests. While Pinterest publishes a transparency report that includes the total number of government takedown requests and the number complied with, it does not break down those requests by country.

Transparent About Platform Policy Takedown Requests. While Pinterest publishes a transparency report that includes the total number of government takedown requests and the number complied with, it does not break down those requests by country.

Provides Meaningful Notice. Pinterest does not publicly commit to providing meaningful notice to users of every removal and suspension.

Allows Appeals. Pinterest allows users to appeal takedowns and suspensions through its contact form, which is linked from its help center. To appeal content takedowns, users can select “Reporting something” and “Appeal a policy violation removal” from the drop-down menus. To appeal account suspensions, users can select “Getting into my account” and “Appeal account suspensions” from the drop-down menus.

Limits Geographic Scope. Pinterest has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

When Pinterest complies with a government request to remove content, we restrict that content from appearing only in the country where the request originated. That content will still be available to all other users.

Reddit



Transparent About Legal Takedown Requests. Reddit publishes a transparency report that breaks down all government takedown requests by country, as well as noting the company’s compliance rate, the type and number of posts affected, and the reason for their removal. Reasons for removal include “legal reasons.”

Transparent About Platform Policy Takedown Requests. Reddit publishes a transparency report that breaks down all takedown requests from government by country, as well as noting the company’s compliance rate, the type and number of posts affected, and the reason for their removal. Reasons for removal include “violation of the Content Policy.”

Provides Meaningful Notice. While Reddit has a published policy of providing notice to users whose accounts have been suspended, it does not publicly commit to providing notice to users whose posts have been taken down. In neither case does Reddit publicly commit to notifying users of the reason for the suspension or takedown.

Allows Appeals. Reddit allows users to appeal takedowns, suspensions, and any other decisions by contacting the admins.

Limits Geographic Scope. Reddit has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required. Its transparency report states:

Where appropriate, rather than removing a post outright, Reddit may make the post inaccessible in a particular country (“Geoblock”).

Snap



Snap Inc.

Transparent About Legal Takedown Requests. Snap publishes a transparency report that breaks down all government takedown requests by country, as well as noting the company's compliance rate.

Transparent About Platform Policy Takedown Requests. Snap does not track this kind of takedown request in its transparency report:

Although we do not formally track when we remove content that violates our policies when a request has been made by a governmental entity, we believe it's an extremely rare occurrence.

Provides Meaningful Notice. Snap does not publicly commit to providing meaningful notice to users of every removal and suspension.

Allows Appeals. Snap does not have a published policy or process for users to appeal takedowns and suspensions.

Limits Geographic Scope. Snap has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

When we believe it's necessary to restrict content that is deemed unlawful in a particular country, but does not otherwise violate our policies, we seek to restrict access to it geographically when possible, rather than remove it globally.

Tumblr



tumblr.

Transparent About Legal Takedown Requests. Tumblr’s parent company Oath publishes a transparency report that breaks down all takedown requests from governments by country, as well as reporting the number of requests, the number of “items specified” in those requests, and the company’s compliance rate. This includes legal requests.

Transparent About Platform Policy Takedown Requests. Tumblr’s parent company Oath publishes a transparency report that breaks down all takedown requests from governments by country, as well as reporting the number of requests, the number of “items specified” in those requests, and the company’s compliance rate. This includes content that violates Oath’s Terms of Service and Community Guidelines.

Provides Meaningful Notice. While Tumblr has a published policy of providing notice to users, it does not specify when users may or may not receive notice of government-ordered takedowns or suspensions. Its community guidelines state:

If we conclude that you are violating these guidelines, you may receive a notice via email. If you don't explain or correct your behavior, we may take action against your account.

Allows Appeals. Tumblr allows users to appeal content takedowns and account suspensions through Tumblr’s support interface. This interface allows users to choose a category in which their problem fits, including “Terminated blog” and “Blog incorrectly marked as explicit.”

Limits Geographic Scope. Tumblr does not have a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required.

Twitter



Transparent About Legal Takedown Requests. Twitter reports legal takedown requests from governments in its transparency report, specifying:

The removal requests reflected in this section of the Transparency Report only include official legal process, such as court orders served on Twitter, and other legal requests that are specifically directed to our intake channels for law enforcement and other authorized reporters (“Legal Requests”). This section does not include requests, including those submitted by government officials, that are directed to our customer support team through our online support forms.

Twitter reports the number of legal requests per country, the type of legal request, its compliance rate, the number of accounts specified in requests, and the number of tweets and accounts ultimately withheld.

Transparent About Platform Policy Takedown Requests. While Twitter reports government requests related to platform policy in its transparency report, it does not break them down by country.

Provides Meaningful Notice. While Twitter has a published policy of providing notice to users, it does not provide notice in cases related to “terrorism”:

Twitter may notify you of the existence of a legal request pertaining to your account unless we are prohibited or the request falls into one of the exceptions to our user notice policy (e.g., emergencies regarding imminent threat to life, child sexual exploitation, terrorism).

Because Twitter does not commit to providing notice in cases related to “terrorism,” a class of content that is difficult to accurately identify and prone to mistakes, it does not earn a star in this category.

Allows Appeals. Twitter allows users to appeal tweet takedowns and account suspensions.

For tweet takedowns:

When we determine that a Tweet violated the Twitter Rules, we require the violator to delete it before they can Tweet again. We send an email notification to the violator identifying the Tweet(s) in violation and which policies have been violated. They will then need to go through the process of deleting the violating Tweet or appealing our review if they believe we made an error.

For permanent account suspensions:

Violators can appeal permanent suspensions if they believe we made an error. They can do this through the platform interface or by filing a report. Upon appeal, if we find that a suspension is valid, we respond to the appeal with information on the policy that the account has violated.

For other types of account locks and suspensions:

If you are unable to unsuspend your own account using the instructions above and you think that we made a mistake suspending or locking your account, you can appeal.

Further, Twitter provides step-by-step instructions for users whose accounts have been temporarily locked or limited, and allows appeals via a specific form for locked or suspended accounts.

Limits Geographic Scope. Twitter has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

In our continuing effort to make our services available to people everywhere, if we receive a valid and properly scoped request from an authorized entity, it may be necessary to withhold access to certain content in a particular country from time to time. Such withholdings will be limited to the specific jurisdiction that has issued the valid legal demand or where the content has been found to violate local law(s)."

Vimeo



Transparent About Legal Takedown Requests. Vimeo does not publish a transparency report.

Transparent About Platform Policy Takedown Requests. Vimeo does not publish a transparency report.

Provides Meaningful Notice. Vimeo does not publicly commit to providing meaningful notice to users of every removal and suspension.

Allows Appeals. Vimeo does not have a published policy or process for users to appeal takedowns and suspensions. Vimeo's terms of service explicitly state that users whose accounts are terminated may not re-register:

Vimeo may suspend, disable, or delete your account (or any part thereof) or block or remove any content you submitted if Vimeo determines that you have violated any provision of this Agreement or that your conduct or content would tend to damage Vimeo's reputation and goodwill. If Vimeo deletes your account for the foregoing reasons, you may not re-register for the Vimeo Service. Vimeo may block your email address and Internet protocol address to prevent further registration.

Limits Geographic Scope. Vimeo does not have a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required.

WordPress.com



Transparent About Legal Takedown Requests. WordPress.com’s parent company Automattic publishes a transparency report in which it details both overall and per-country numbers for total takedown requests, whether they are court orders or requests from law enforcement, compliance rate (“percentage of requests where content was removed solely in response to the demand”), and the number of WordPress.com sites specified in requests.

Transparent About Platform Policy Takedown Requests. WordPress.com’s parent company Automattic publishes a transparency report in which it details both overall and per-country numbers for total takedown requests, whether they are court orders or requests from law enforcement, compliance rate (“percentage of requests where content was removed due to a violation of our policies”), and the number of WordPress.com sites specified in requests.

Provides Meaningful Notice. While WordPress.com has a published policy of providing notice to users, it does not specify when users may or may not receive notice of government-ordered takedowns or suspensions:

In some cases, we may add a warning note in your dashboard. It will contain a link that you can use to contact us regarding the issue. We might also disable posting on your site, or discontinue other features on your account.

Allows Appeals. WordPress.com allows users to appeal takedowns, suspensions, or other errors:

We do make mistakes from time to time. If you feel that we’ve done anything in error, please contact us via the link on your dashboard or by using the form below. A real person will review your request and reply with our decision as soon as possible.

Limits Geographic Scope. WordPress.com’s parent company Automattic has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required:

We aim to promote freedom of expression around the world, and are also mindful of local laws that might impact that expression. When we receive an order to remove content, we may block it in only those jurisdictions where it violates local law (aka “geoblock”), so that it remains accessible in areas where it may not be illegal.

Further, Automattic documents in its transparency report the countries for which it exercises geoblocking, the websites blocked in those countries, and the error page users see when they attempt to access a geoblocked page.

YouTube



Transparent About Legal Takedown Requests. YouTube’s parent company Google publishes a transparency report that includes all government takedown requests. The transparency report states:

Governments contact Google with content removal requests for a number of reasons. Government bodies may claim that content violates a local law, and include court orders that are often not directed at Google with their requests. Both types of requests are counted in this report.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests.

Transparent About Platform Policy Takedown Requests. YouTube’s parent company Google publishes a transparency report that includes all government takedown requests. The transparency report states:

We also include government requests to review content to determine if it violates our own product community guidelines and content policies.

Information for each country includes the total number of takedown requests, the total number of items requested for removal, and the percentage of requests in which some content was removed. Country-level reports also categorize the reasons behind requests, as well as describing the details and outcomes of individual requests.

Provides Meaningful Notice. YouTube provides notice—or “strikes”—for account terminations as well as content takedowns due to Community Guidelines notifications or legal requests.

For account terminations:

When an account is terminated, the account owner receives an email detailing the reason for the termination.

For content takedowns due to Community Guidelines violations:

Community Guidelines strikes are issued when our reviewers are notified of a violation of the Community Guidelines. ... If a strike is issued, you'll get an email and see an alert in your account's Channel Settings with information about why your content was removed (e.g. for sexual content or violence).

For content takedowns due to legal requests:

YouTube makes reasonable efforts to notify creators when their content is restricted due to a legal request.

Allows Appeals. YouTube allows users to appeal takedowns and suspensions.

For content takedowns, users follow the process to appeal Community Guidelines actions.

For account suspensions, users can appeal through a dedicated form.

Limits Geographic Scope. Google has a published policy of limiting legally ordered content restrictions to jurisdictions where such restriction is required: Its government requests FAQ states:

Where possible, we remove or restrict access to the content in the country where it is deemed to be illegal.

References and helpful links

Apple App Store

Transparency report on government and private party requests:

<https://www.apple.com/legal/privacy/transparency/requests-2017-H2-en.pdf>

App store information for developers:

<https://developer.apple.com/support/app-store/>

If your Apple ID is locked or disabled:

<https://support.apple.com/en-ca/HT204106>

App review information:

<https://developer.apple.com/support/app-review/>

Dailymotion

Terms of Use:

<https://www.dailymotion.com/legal/termsOfSales>

Facebook

Content Restrictions Based on Local Law transparency report:

<https://transparency.facebook.com/content-restrictions>

Publishing Our Internal Enforcement Guidelines and Expanding Our Appeals Process:

<https://newsroom.fb.com/news/2018/04/comprehensive-community-standards/>

Community Standards Enforcement Preliminary Report:

<https://transparency.facebook.com/community-standards-enforcement>

Google+

Government requests to remove content transparency report:

<https://transparencyreport.google.com/government-removals/overview>

Government requests to remove content FAQs:

<https://support.google.com/transparencyreport/answer/7347744>

Form to appeal profile suspension:

https://support.google.com/plus/contact/suspension_appeal

Google Play Store

Government requests to remove content transparency report:

<https://transparencyreport.google.com/government-removals/overview>

Enforcement:

<https://play.google.com/about/enforcement/enforcement-process/>

Developer Program Policies:

https://play.google.com/about/developer-content-policy/#!?modal_active=none

Contact form for account termination or app removal:

https://support.google.com/googleplay/android-developer/troubleshooter/2993242?visit_id=1-6366166754-96211013-209119645&rd=1

and

https://support.google.com/googleplay/android-developer/troubleshooter/2993242?visit_id=1-6366166754-96211013-209119645&rd=1#ts=2993244

App removal information:

https://support.google.com/googleplay/android-developer/troubleshooter/2993242?visit_id=1-6366166754-96211013-209119645&rd=1#ts=2993244

Instagram

Content Restrictions Based on Local Law transparency report:

<https://transparency.facebook.com/content-restrictions>

“Publishing Our Internal Enforcement Guidelines and Expanding Our Appeals Process” Newsroom post:

<https://newsroom.fb.com/news/2018/04/comprehensive-community-standards/>

Community Standards Enforcement Preliminary Report:

<https://transparency.facebook.com/community-standards-enforcement>

LinkedIn

Transparency report:

<https://www.linkedin.com/legal/transparency>

About restricted accounts:

<https://www.linkedin.com/help/linkedin/answer/82934>

Account/content restricted or removed:

<https://www.linkedin.com/help/linkedin/answer/82934?query=account%20restriction>

Account suspensions appeal form:

<https://www.linkedin.com/help/linkedin/ask/hr>

Commitment to Transparency Regarding Content Blocked From Our Site:

<https://www.linkedin.com/help/linkedin/answer/46925/linkedin-s-commitment-to-transparency-regarding-content-blocked-from-our-site>

Medium

Medium Rules:

<https://medium.com/policy/medium-rules-30e5502c4eb4>

Lumen:

<https://lumendatabase.org>

Pinterest

Transparency report:

<https://help.pinterest.com/en/articles/transparency-report>

Reporting something on Pinterest:

<https://help.pinterest.com/en/articles/report-something-pinterest#Web>

Contact form for appeals (must be logged in):

<https://help.pinterest.com/en/contact>

Reddit

Transparency report:

<https://www.redditinc.com/policies/transparency-report>

Content policy:

<https://www.redditinc.com/policies/content-policy>

About suspensions:

<https://www.reddithelp.com/en/categories/rules-reporting/account-and-community-restrictions/suspensions>

Snap

Transparency report:

<https://www.snap.com/en-US/privacy/transparency/>

Tumblr

Oath government removal requests report:

<https://transparency.oath.com/government-removal-requests.html?gucounter=1>

Oath transparency report FAQs and glossary:

<https://static.tumblr.com/zyubucd/gmnopeat/combinedreport.pdf>

Community guidelines:

<https://www.tumblr.com/policy/en/community>

About reporting offensive content:

<https://tumblr.zendesk.com/hc/en-us/articles/226270628-Reporting-offensive-content>

Support form:

<https://www.tumblr.com/support>

Twitter

Removal requests report:

<https://transparency.twitter.com/en/gov-tos-reports.html>

Government TOS report:

<https://transparency.twitter.com/en/gov-tos-reports.html>

About country withheld content:

<https://help.twitter.com/en/rules-and-policies/tweet-withheld-by-country>

About suspended accounts:

<https://help.twitter.com/en/managing-your-account/suspended-twitter-accounts>

Form to appeal an account suspension or locked account:

<https://help.twitter.com/forms/general?subtopic=suspended>

Help with locked or limited accounts:

<https://help.twitter.com/en/managing-your-account/locked-and-limited-accounts>

Legal request FAQs:

<https://help.twitter.com/en/rules-and-policies/twitter-legal-faqs>

Range of enforcement options:

<https://help.twitter.com/en/rules-and-policies/enforcement-options>

Vimeo

Terms of service:

<https://vimeo.com/terms>

WordPress.com

Takedown demands report:

<https://transparency.automattic.com/takedown-demands/>

Country block list:

<https://transparency.automattic.com/country-block-list/>

About suspended content and sites:

<https://en.support.wordpress.com/suspended-blogs/>

YouTube

Government requests to remove content transparency report:

<https://transparencyreport.google.com/government-removals/overview>

About account terminations:

<https://support.google.com/youtube/answer/2802168?hl=en>

Community Guidelines strike basics:

https://support.google.com/youtube/answer/2802032?hl=en&ref_topic=2803138

About legal requests:

<https://support.google.com/youtube/answer/3001497?hl=en>

Appeal Community Guidelines Actions:

https://support.google.com/youtube/answer/185111?hl=en&ref_topic=2803138

“Unable to access a Google product” appeal form:

https://support.google.com/accounts/contact/suspended?p=youtube&visit_id=1-636610084672726027-1458380864&rd=1