



October 2018

The Transparency Reporting Toolkit: Content Takedown Reporting

Spandana Singh & Kevin Bankston

Last edited on October 24, 2018 at 8:54 a.m. EDT

Acknowledgments

In addition to the many stakeholders across civil society and industry that have taken the time to talk to us over the years about our work on transparency reporting, we'd particularly like to thank Daphne Keller, the Director of Intermediary Liability at the Stanford Center for Internet and Society, Annemarie Bridy, a Professor of Law at the University of Idaho, and the Ranking Digital Rights team at New America for their help in refining the recommendations in this toolkit. We would also like to thank the MacArthur Foundation for its generous support of our work in this area.

About the Author(s)

Spandana Singh is a policy program associate in New America's Open Technology Institute.

Kevin Bankston is the director of New America's Open Technology Institute, where he works in the public interest to ensure that all communities have equitable access to an internet that is both open and secure. He previously served as OTI's policy director.

About New America

We are dedicated to renewing America by continuing the quest to realize our nation's highest ideals, honestly confronting the challenges caused by rapid technological and social change, and seizing the opportunities those changes create.

About Open Technology Institute

OTI works at the intersection of technology and policy to ensure that every community has equitable access to digital technology and its benefits. We promote universal access to communications technologies that are both open and secure, using a multidisciplinary approach that brings together advocates, researchers, organizers, and innovators.

Contents

Introduction and Executive Summary	6
Overview of Types of Content Takedowns and Reporting	11
1. Government and other legal content demands	11
2. Copyright requests	12
3. Trademark requests	12
4. Network shutdowns and service interruptions	13
5. "Right to be Forgotten" delisting requests	13
6. Community Guidelines-based content takedowns	14
General Best Practices for Content Takedown Reporting	16
1. Issuing regular reports on clearly and consistently delineated reporting periods	17
2. Issuing reports specific to the type of demand	17
3. Reporting on types of demands using specific numbers	18
4. Breaking down demands by country	19
5. Reporting on categories of objectionable content targeted by demands	19
6. Reporting on products targeted by demands	21
7. Reporting on specific government agencies/parties that submitted demands	22
8. Specifying which laws pertain to specific demands	23
9. Reporting on the number of accounts and items specified in demands	24
10. Reporting on the number of accounts and items impacted by demands	25
11. Reporting on how the company responded to demands	26
12. Additional best practices to consider	27

Contents Cont'd

Issue-Specific Reporting Best Practices	32
1. Reporting on government and other legal content demands	33
2. Reporting on copyright requests	33
3. Reporting on trademark requests	35
4. Reporting on network shutdowns and service interruptions	35
5. Reporting on Right to be Forgotten delisting requests	37
6. Reporting on Community Guidelines-based content takedowns	38
Appendix	44

Introduction and Executive Summary

In 2016, New America's Open Technology Institute and Harvard University's Berkman Klein Center for Internet & Society released the **Transparency Reporting Toolkit**, a joint project that aimed to make it easier for companies to create and improve their transparency reports around government demands for user data and to make transparency reporting more consistent, easier to understand and more effective.

As the internet has become an increasingly important tool for free expression for individuals around the world, so too have the platforms and networks that carry that expression become speech gatekeepers, often removing or blocking users' content for various legal or policy reasons. As questions and controversies around these internet companies' content takedown practices have multiplied, many have expanded their transparency reports to include data on content takedowns and restrictions as well as network shutdowns and service interruptions. However, as with reporting on government demands for user data, there is a great deal of variance in the approaches, styles, and scope of data that companies are using in their reporting, making it hard to combine or compare data in a meaningful way.

Transparency reporting on content takedowns is critically important because it helps hold companies accountable in their role as gatekeepers of our online speech, and helps the public identify where they think the companies are doing too much—or not enough—to address content issues on their platforms and networks. Transparency reporting also offers a number of benefits to the companies themselves, including helping to build trust with their users and policymakers. For companies who are consistently under pressure to act on problematic speech, this is an opportunity to highlight to users and policymakers the responses they have implemented, and communicate the size and complexity of the problems they are addressing, as well as the impact they have had thus far. Reporting can also help reveal where particular governments or laws are leading to a disproportionate amount of content takedowns, thus helping hold public authorities to account as well. For all these reasons, tools for promoting, improving, and standardizing the issuance of content takedown reporting are much-needed, hence this toolkit.

For this edition of the Transparency Reporting Toolkit, focused on content takedown reporting, we've relied not only on extensive consultations and convenings with a wide range of companies and civil society experts but also on the rigorous work of the **Ranking Digital Rights** project at New America, which has developed a broad set of indicators for judging how well companies perform when it comes to protecting human rights. We've also applied the lessons we've learned from our recent **Getting Companies to Do The Right Thing** project,

which included an **extensive history** of the practice of transparency reporting and how it has evolved over the years.

We surveyed 24 international and 46 domestic internet and telecommunications that issue transparency reports of some kind, and found 35 of them have reported on content-related demands and takedowns, mostly falling into six categories:

- Government and other legal content demands
- Copyright requests
- Trademark requests
- Network shutdowns and service interruptions
- Right to be Forgotten delisting requests
- Community Guidelines-based content takedowns

Charts **1a** and **1b** reflect which companies are issuing reports on each of these six categories.

Based on this survey, we identified 11 general best practices for content takedown reporting which are applicable across these different types of content reports:

- **Issuing regular reports on clearly and consistently delineated reporting periods:** Companies should consistently publish reports at regular intervals covering clearly defined reporting periods.
- **Issuing reports specific to the type of demand:** Companies should report separately on different types of content demands.
- **Reporting on types of demands using specific numbers:** Companies should report separately on the number of demands they have received for each category.
- **Breaking down demands by country:** Companies should break down the demands they have received for each category by country.
- **Reporting on categories of objectionable content targeted by demands:** Companies should break down demands received by the different types of objectionable content (e.g. nonconsensual pornography, extremist content) that the demands target.

- **Reporting on products targeted by demands:** Companies should break down demands by which of their products or services the demands target.
- **Reporting on specific government agencies/parties that submitted demands:** Companies should break down demands received by which government agencies or parties submitted them.
- **Specifying which laws pertain to specific demands:** Companies should specify which laws are associated with reported demands.
- **Reporting on the number of accounts and items specified in demands:** Companies should break down how many accounts and pieces of content were specified in demands they received.
- **Reporting on the number of accounts and items impacted by demands:** Companies should break down how many accounts and pieces of content were impacted (i.e., taken down or otherwise restricted) in response to demands.
- **Reporting on how the company responded to demands:** Companies should break down their different responses to demands (e.g., complied, partially complied, rejected, etc.).

Charts 2a through 6 in the [appendix](#) reflect, for the first five of the six different categories of reports, which companies satisfy these best practices. We did not include a chart for Community Guidelines-based content takedowns because that practice is still relatively rare and highly variable in how it is reported on.

Most of the above best practices are focused on quantitative transparency—counting incidents and items. Our survey also revealed several additional best practices that companies should follow that are mostly aimed at enhancing the qualitative transparency around takedowns by offering greater context and explanation, or features aimed at making the data more useful. These include:

- Defining terms clearly
- Providing meaningful explanations of internal policies
- Offering case studies to illustrate the company’s practices and the issues it faces
- Reporting on specific notices where reasonable and permitted by law

- Providing meaningful numbers that reflect how many pieces of content or accounts were taken down, blocked or otherwise restricted based on automated flagging or review
- Linking relevant reports to one another
- Publishing reports at static and functioning URLs
- Publishing data in a structured data format
- Publishing reports using a non-restrictive Creative Commons license
- Offering a Frequently Asked Questions section for the report

Finally, our survey identified a handful of additional best practices specific to reporting on certain categories of takedowns:

- Reporting on copyright requests should include
 - The number and impact of counter-notices to DMCA-based requests
 - The different format or medium of content (audio, video, photograph, text, etc.) targeted by copyright requests
 - Data on who is submitting copyright requests
- Reporting on network shutdowns and service interruptions should include
 - The date or date range of the network shutdown or service interruption
 - The duration of the network shutdown or service interruption
- Right to be Forgotten delisting requests should include
 - The categories of content requested to be delisted (e.g. personal information, professional information, information about crime and professional wrongdoing, etc.).
 - The categories of websites targeted (e.g., news, social network, directory, etc.)

These recommendations—which are not legal advice, but simply our survey of the features of current reports—reflects our understanding of current best practices in this space. Our survey was based on transparency reports published as of August 2018 (notably, South Korean internet company Daum Kakao has since issued a new transparency report that unfortunately is much less detailed than its previous reports; the data and examples regarding Daum Kakao in this toolkit are based on its prior report, for which we use links to [archived versions](#) of its report at Archive.org).

As companies continue to expand their transparency reporting on these issues, we hope to revisit these resources and update them. If you believe we missed out on anything, please feel free to reach out to us and let us know. We hope this edition of the Toolkit will help facilitate discussions around best practices for content takedown reporting as well as foster future discussions on best practices for other types of transparency reporting.

Overview of Types of Content Takedowns and Reporting

Below is an overview of the six most common types of content takedowns that internet and telecommunications companies can and have been reporting on. The first five are types of content takedowns or blocking done in response to certain types of legal requests or demands; the sixth category regards content takedowns based on a company's own content policies.

1. Government and other legal content demands

This category focuses on government-issued legal demands, as well as other legal demands for content takedowns issued by individuals asserting a violation of local law (e.g., defamation). Although copyright and trademark requests are also forms of legal demands, companies typically report on them separately. Therefore we also separately evaluate company performance on reporting about copyright and trademark requests, later in our survey.

Currently, most companies lump together government and non-government legal demands into one report, rather than counting them separately, and as a result our evaluation of their reporting in Charts 2a and 2b also assesses them collectively. However, a few companies, such as **Daum Kakao** and **Apple**, report separately on at least some forms of government and non-government requests. Going forward it is recommended that all companies do so, for greater clarity and transparency.

Reporting on government and other legal demands is currently the most prevalent form of content takedown reporting across both the internet and telecommunications sectors. Reporting on this category of demands generally includes data on the number of content demands a company received, how many accounts or items these demands impacted and how the company responded to such demands.

When discussing impact, it is important to note that some companies do not always respond to demands by completely removing content from their entire platforms or networks. Rather, they may otherwise restrict or geo-block content, like **Facebook** does, or block access to certain IPs and URLs in a country or region if the applicable law is regionally specific, like **Telefonica** does.

As a result, some internet platforms also report specifically on content restrictions short of complete takedown, and some telecommunications companies also report on IP/URL or content blocking.

2. Copyright requests

One of the most common types of legal requests seeking content takedown are those alleging copyright infringement. Among the companies we surveyed, the most commonly cited copyright law is the Digital Millennium Copyright Act (DMCA), a U.S. statute that enables copyright holders to file requests or notices for content takedown based on the premise that a certain piece of posted content infringes upon copyright. Most internet intermediaries that are based in the United States or that largely operate in the United States follow the DMCA's guidelines on copyright. However, international internet intermediaries whose primary hubs of operation are outside the United States tend to receive content requests based on their own local copyright protections and laws.

Reporting on copyright takedowns is currently the second most prevalent form of reporting on content takedowns. Reporting on this category of demands typically includes data on the number of related requests a company received and how they responded. However, this data often excludes vital points such as information on the number of counter-notices a company has received and how much content was restored as a result.

3. Trademark requests

Just as copyright holders issue requests demanding the takedown of allegedly copyright-infringing content, so do trademark holders who submit takedown requests alleging infringement of their marks. Among the companies we surveyed, the majority of internet platforms evaluate trademark requests based on U.S. trademark law. However, international internet intermediaries whose primary hubs of operation are outside the United States receive content requests based on their local trademark protections and laws.

Reporting on trademark takedowns typically includes data on the number of requests a company received, how a company responded to those requests and the number of accounts or items that were impacted by those requests. Compared to other forms of reporting on intellectual property-related takedowns such as copyright, reporting on trademark takedowns is not as prevalent on communications-centered websites such as social media platforms, although it is more prevalent on e-commerce platforms. Reporting on trademark takedowns is, however, gradually becoming a more widespread practice across internet platforms that host user-generated content such as blogs.

4. Network shutdowns and service interruptions

This category includes government-mandated network shutdowns of internet services as well as incidents in which telecommunications systems, online platforms, or other internet-based services have been disrupted by government agencies. This category does not include demands to block access to specific pieces of content, IP addresses, or URLs, but rather focuses on system- and service-wide shutdowns and interruptions.

Network shutdowns primarily impact telecommunications operators. However, internet disruptions may also be targeted at companies that offer services over the internet in addition to or instead of the telecoms themselves, especially during times of political turmoil and instability. For example, a government may force a shutdown of a telecom's internet service, or it may force that telecom to disrupt access to (e.g.) Twitter alone. Both shutdowns and disruptions aim to bar access to certain services or the internet in general.

Although network shutdowns are becoming increasingly prevalent across the world, internet and telecommunications platforms do not typically document or report on instances of shutdowns or disruptions in their transparency reports. Sometimes this is due to legal restrictions which prevent such disclosures in countries that companies operate in.

Companies that report on this category well typically disclose data on the location, timing, and reason behind the shutdown. A handful of companies also provide quantitative data on the number of shutdowns and the duration of these shutdowns.

5. "Right to be Forgotten" delisting requests

In May 2014, the Court of Justice of the European Union ruled that citizens of the EU could request search engines to delist search results tied to their names if the information in the result was "inadequate, irrelevant or excessive in relation to the purpose of processing." This right to erasure was also included in the EU's General Data Protection Regulation which came into effect in May 2018. A similar Right to Erasure law came into effect in Russia in January 2016.

Although reporting on delisting requests received as a result of the Right to be Forgotten laws are a form of company reporting on legal requests, companies have begun reporting on these requests separately. As a result, we have assigned this form of reporting its own category.

Currently, transparency reporting around Right to be Forgotten delisting requests is the least prevalent form of content takedown-related transparency

reporting—only **Microsoft** and **Google** currently issue such reports (for Bing and Google Web Search, respectively). This is most likely because these regulations are directed specifically at search engines, and as a result other types of services don't receive—or at least, shouldn't be receiving—these requests. However, to the extent services outside the scope of these laws are receiving such requests, that is something the public and policymakers should know, and the best practice for those companies would be to report on those requests—especially because the Right to be Forgotten may ultimately be expanded to cover them.

6. Community Guidelines-based content takedowns

A company's Community Guidelines (also often referred to as Community Standards or Terms of Service) outline its self-defined standards and norms for what type of content is acceptable and unacceptable on its service. Companies often remove or restrict content, or suspend or deactivate accounts, based on violations of such guidelines, whether that violation was flagged by users or discovered by the service through its own automated or human review. The need for transparency and accountability in this area is particularly strong because the company is essentially enforcing its own private law, rather than complying with public law subject to the typical checks and balances of the political process.

Reporting on Community Guidelines-based content takedowns is still a new and emerging practice. Until recently, the very few platforms that reported on Community Guidelines-based content takedowns at all only did so for one or another category of content (for example, **Microsoft** only reports on nonconsensual pornography takedowns, while **Automattic** only reports on terror content takedowns). However, over the past few months major internet platforms such as **Google (via YouTube)** and **Facebook** have issued more comprehensive transparency reports covering Community Guidelines-based takedowns across multiple content categories.

Each of these six types of reporting on content takedowns are integral to understanding the nature of and limitations on online freedom of expression today, and we recommend that companies issue reports providing data on all of the above content demands where relevant to their services and the types of requests they receive. However, there is currently a significant amount of variance between how and what internet and telecommunications companies report on.

Charts **1a** and **1b** summarize which of the six types of content demands the 35 companies we surveyed are (or are not!) reporting on. Two important notes for understanding these charts: First, note that for the Community Guidelines-based content takedowns section, a company earned a check mark if it published any

data on its Community Guidelines-based content takedown efforts. However, the quality and extent of the reported data varied widely, and as previously noted only two companies are currently offering anything close to comprehensive numbers. Second, these charts reflect the first of our 11 general best practices—regular issuance of a report over consistent periods—by highlighting each company’s last reported period. Therefore that best practice is not addressed in later charts that assess compliance with the overall best practices.

General Best Practices for Content Takedown Reporting

Through our survey of internet and telecommunications companies and their approaches to reporting on content demands, we have identified a set of general best practices for such reporting. Some of these best practices are applicable to all forms of content-and network-related reporting, regardless of the issue area they fall under. These include:

- Issuing regular reports on clearly and consistently delineated reporting periods
- Issuing reports specific to the type of demand
- Reporting on types of demands using specific numbers
- Breaking down demands by country
- Reporting on categories of objectionable content targeted by demands
- Reporting on products targeted by demands
- Reporting on specific government agencies/parties that submitted demands
- Specifying which laws pertain to specific demands.
- Reporting on the number of accounts and items specified in demands
- Reporting on the number of accounts and items impacted by demands
- Reporting on how the company responded to demands

Below are brief overviews of what these 11 best practices for reporting include and why they are important. These best practices also include examples of companies that we believe demonstrated these best practices in an outstanding manner.

1. Issuing regular reports on clearly and consistently delineated reporting periods

The first and most fundamental best practice in transparency reporting of any kind is to regularly issue such reports on a consistent timeline, and clearly and consistently delineate the reporting period for each issued report.

Currently there is no industry-wide standard for how often companies should publish transparency reports, and reports typically range from quarterly to annual publication schedules. Because reports issued more often and covering shorter periods can offer more granular information, the best practice is to publish quarterly, if practical. Regardless of the publication schedule a company adopts, it should clearly define the time period covered by each report, and subsequent reports should cover the same length of time so that reports may be easily compared. If a company fails to publish a transparency report for an expected period, they should explain why.

As visible in Charts **1a** and **1b**, a number of internet and telecommunications platforms have been inconsistent in publishing their reports. Below are some examples of companies that have issued regular reports, using the three most common time periods: quarterly, biannually, and annually.

CREDO Mobile: CREDO Mobile publishes a transparency report quarterly. Each report covers three months.

Verizon: Verizon publishes a transparency report bi-annually, covering the periods of January-June and July-December, respectively.

Reddit Reddit publishes a transparency report annually. The report covers the period of January 1 - December 31 of each year.

2. Issuing reports specific to the type of demand

By reporting separately on different types of demands or takedowns—breaking down numbers as between, e.g., government and other legal demands vs. copyright requests vs. trademark requests vs. Right To Be Forgotten delisting requests vs. community guidelines violations, rather than lumping them all together—a company is able to highlight the breadth of demands they have received and the volume and impact of each of these takedown categories. Some examples of companies that exhibit this best practice in their transparency reports are:

Automattic: In its latest transparency report, which covers the first half of 2018, Automattic separately reports that they received 9,166 total copyright notices and 318 trademark notices.

Microsoft: In its latest transparency report, Microsoft separately reports on government requests for content removal, copyright removal requests, and Right to be Forgotten requests.

Snap Inc.: In its latest transparency report, which covers the the second half of 2017, Snap Inc. separately reports on the number of government-issued content removal requests and copyright requests they received.

Some companies report on the number of requests per issue area in unique ways and do not always explain the rationale behind such approaches. **Telefonica** is one example of this. In some cases, Telefonica separately reports the number of requests received in a country for the blocking and filtering of contents and for geographical or temporary suspensions of the service. However, in some countries, such as Argentina, they lump these figures together in conjunction with the number of requests they received for access to metadata. Ideally, companies will report consistently, or at least explain when they do not.

3. Reporting on types of demands using specific numbers

By reporting separately on the number of demands a company receives over a given time period for each different type of takedown, companies can highlight which types of demands are most common. This best practice requires the publication of specific numbers; percentages alone are not sufficient (though they are a helpful supplement to specific numbers), nor are numeric ranges. Some examples of companies that exhibit this best practice in their transparency reports are:

Apple: Apple specifies that between July and December 2017 it received 7 requests for account restriction/account deletion globally.

Oath: In its latest transparency report, which covers the second half of 2017, Oath reports that it received 77 total government-issued removal requests from around the world.

Pinterest: In its latest transparency report, which covers the period of January - March 2018, Pinterest reports that it received 26 government-issued content removal requests.

4. Breaking down demands by country

In order to demonstrate the geographic scope of demands a company is receiving, and to highlight which countries' governments or laws are most actively restricting online free expression, companies should specify how many requests originate from each specific country. The most effective way to provide this information is to create a list or map of all countries relevant to a company's operations and indicate the number of demands received from each. Some examples of companies that exhibit this best practice are:

Mapbox: Mapbox reports on the number of DMCA takedown notices and the number of government requests to withhold content it has received using a map of the world that shows a specific number for each country. (So far, Mapbox has not received any such requests, so all of the figures on its map are 0.)

Microsoft: In its latest transparency report, which covers the period July-December 2017, Microsoft breaks down data on Right to be Forgotten delisting requests by country. Some of the countries included in this breakdown are Austria, Finland, Germany, Romania, Russia and the United Kingdom.

Telenor: In its March 2017 Authority Requests Disclosure Report, Telenor breaks down the number of requests it has received by the countries it operates in. The countries covered in this report are Norway, Sweden, Denmark, Hungary, Serbia, Montenegro, Bulgaria, Pakistan, India, Bangladesh, Myanmar, Thailand and Malaysia.

5. Reporting on categories of objectionable content targeted by demands

By reporting on the categories of objectionable content targeted by different types of content demands, a company can highlight the varying reasons that parties are asking for content to come down, and also indirectly highlight the relative prevalence of different types of problematic content on their services. This data can be particularly revealing when combined with the prior best practice of breaking down demands by country, allowing readers to spot (e.g.) which countries are most aggressively seeking to censor which types of content. Some examples of categories of objectionable content include extremist content, obscenity and child pornography, content that violates privacy, and defamatory content.

In order to receive credit for this best practice, companies need to provide specific numbers breaking down how many requests pertained to what kinds of objectionable content.

In the context of government and other legal content demands, satisfying this best practice means highlighting the types of objectionable content specified in requests.

Google: In its latest transparency report on government requests to remove content, which covers the second half of 2017, Google enables users to filter the number of requests received and the number of items specified within those requests based on the reasons behind them. Users can also select individual countries and view quantitative breakdowns on the number of requests received for each category of objectionable content relevant to that country. Common reasons cited in government requests to remove content include national security, defamation, privacy and security, hate speech, drug abuse, religious offense, impersonation, obscenity and nudity. Google also has a separate dedicated section that explains the most common reasons cited for content removal.

In the context of copyright requests, because the category of objectionable content—content that allegedly infringes on a copyright holders' rights—is self-evident, satisfying this best practice means reporting on the format or medium of content being targeted by requests.

Tumblr: In its latest transparency report on copyright and trademark requests, which covers the second half of 2017, Tumblr highlights the kinds of content DMCA requests targeted, including content formats (some Tumblr-specific) such as images, texts, audio, video, links, asks, quotes and chats. This is also applicable to trademark requests. However, currently no companies report on the format or medium of content being targeted by trademark requests.

In the context of network shutdowns and service interruptions, satisfying this best practice means reporting on the reasons behind those shutdowns and interruptions. This does not always require specific numbers, and can be qualitative instead.

Millicom In its latest Law Enforcement Disclosure Report, which covers the year 2017, Millicom reports that since 2014, authorities in El Salvador and Honduras have required telecommunications companies

to shut down services or reduce signal capacity in and around prisons in order to prevent criminal gangs from using smuggled cell phones.

In the context of Right to be Forgotten delisting requests, satisfying this best practice means reporting on the categories of content targeted by delisting requests.

Google: In its latest transparency report on search removals under European privacy law, Google reports on the “categories of content requested for delisting”. The categories of content include insufficient information, personal information, professional information, self authored, crime and professional wrongdoing.

6. Reporting on products targeted by demands

Some companies maintain and support multiple products. For example, Facebook offers its core product, the Facebook social network service, as well as Instagram, WhatsApp, and its Facebook Messenger app/service. By specifying which of a company’s products are being targeted by which demands, a transparency report can better reflect how those demands are impacting the range of its offerings, highlight differences in impact between its services, and better enable comparison of that impact with other companies’ comparable services. Some examples of companies that exhibit this best practice are:

Google: In its latest transparency report on government requests to remove content, which covers the second half of 2017, Google enables users to filter the number of requests received and the number of content items specified within these requests based on the product that was targeted by requests, broken down into four categories: Web Search, YouTube, Google AdWords, and All Others. Users can also select individual countries and explore which of Google’s products have been targeted by how many content removal requests there. Additionally, Google’s report has a separate section that breaks down, in percentages, the top three products and services cited in government requests during the selected reporting period (currently YouTube, Web Search, and Blogger).

Wikimedia Foundation: In its latest transparency report, which covers the first half of 2018, the Wikimedia Foundation outlines which specific Wikimedia projects were targeted by requests for content alteration and takedown. Examples of Wikimedia projects that have been targeted by such requests are English Wikipedia, Wikimedia Commons, German Wikipedia and Wikispecies.

Ideally, this best practice would be implemented by a company across all of the types of takedowns it reports on. Some companies do this, but some do not.

Facebook, for example, only breaks down its reporting by product in its copyright and trademark requests report. These reports provide separate data for both the Facebook and Instagram platforms. Facebook does not, however, provide similar separately reported data for Instagram in its reports on other issue areas such as government and other legal content demands, nor do they provide any specific data about content takedowns or blocking on other products and services such as WhatsApp and Facebook Messenger.

This best practice is generally not applicable to telecommunications companies who only offer connectivity services as they typically do not have more than one product or service. However, it is applicable to telecommunications companies who offer multiple services such as email.

In the context of Right to be Forgotten delisting requests, satisfying this best practice means reporting on the categories of websites that the delisting requests target. This is because Right to be Forgotten delisting requests currently target web search products only, and so a breakdown by product for this category of demands is not applicable.

Google: In its latest transparency report on search removals under European privacy law, Google reports on the “categories of websites hosting requested for delisting”. The categories reported on include news, directory and social media sites.

7. Reporting on specific government agencies/parties that submitted demands

By reporting on which specific government agencies or entities submit content-related demands, a company can highlight which elements of government in which countries are the most active in seeking to police online content, which can in turn help identify misuse or overuse of authority or actions outside of a particular part of government’s jurisdiction, as well as overall trends in what content which parts of government are targeting. Some examples of companies that exhibit this best practice are:

Daum Kakao In the last Daum Kakao report that we surveyed, which covers the first half of 2018, Daum Kakao reports on the number of requests it received from specific government agencies. In this reporting, Daum Kakao breaks down requests received from South Korean government agencies such as the Korea Communications Standards Commission, the Ministry of Food and Drug Safety, the

Korea Internet and Security Agency, the National Policy Agency and the Military Manpower Administration. These requests are also broken down by requests received by Daum and requests received by Kakao (Daum Kakao is a company that formed as a result of a merger between Daum Communications and Kakao).

Google: In its latest transparency report on government requests to remove content, which covers the second half of 2017, you can filter the number of requests received and the number of items specified within these requests based on the government branch that submitted requests. In most countries, Google divides up requests based on whether they were received from the judicial or executive branches.

Reporting specific information about the private parties that submitted content demands may raise privacy and safety issues that government requests do not, and so we are not urging that companies publish the identity of every private party that makes a content request of any kind. However, it is worth highlighting that we have seen exceptional transparency from a few companies in regard to those who submit copyright requests:

Google: In its latest transparency report on removals of URLs from its search results due to copyright, Google highlights the identity of the six copyright-holding companies that submitted the most requests, as well as highlighting the top six reporting organizations that have submitted the most requests, and goes on to list the more than 15,000 other copyright holders and the more than 13,000 reporting organizations that have submitted DMCA takedown requests.

Twitter: In its latest transparency report on copyright notices, which covers the second half of 2017, Twitter highlights the top five copyright reporters on the platform within each reporting period. This reporting includes information on the number of copyright takedown notices they have filed, the percentage of total takedown notices their notices comprise and the number of materials that were withheld as a result of their copyright requests. Notably, the five top requesters during the last reporting period were responsible for over half of the requests.

8. Specifying which laws pertain to specific demands

Because most major internet and telecommunications companies operate in multiple countries, it is important to understand which laws and legal frameworks govern user speech and communications. Some examples of companies that do a good job in their transparency reports of explaining which laws are responsible for which types of takedowns include:

Daum Kakao: In the last Daum Kakao report that we surveyed, which covers the first half of 2018, Daum Kakao provides an overview of the different types of requests users and governments can submit. These include requests for content removal or moderation based on copyright infringement, trademark and portrait rights infringement, the leaking of personal information, defamation and restrictions on obscene materials and materials that purport to sell drugs and medical supplies. For each of these reasons, Daum Kakao provides the applicable law that governs these requests including the Act on Promotion of Information and Communications Network Utilization and Information Protection, etc. and the Pharmaceutical Affairs Act.

Telefonica: In its latest transparency report, which covers the year 2016, Telefonica specifies the relevant legal framework in each country for each category of demands (if there is one). For example, In El Salvador, Telefonica specifies that geographical or temporary suspensions of service can be requested under the Special Law Against the Crime of Extortion (Art. 13 and 14). Similarly, in Spain, Telefonica specifies that the blocking and filtering of certain contents is possible under three legal frameworks: the Royal Decree 1889/2011 of 30 December, regulating the functioning of the Intellectual Property Commission (articles 22 and 23), the Revised Text of the Intellectual Property Law, approved by Royal Legislative Decree Law 1/1996 of 12 April (Article 138) and Law 34/2002 of 11 July on services of the information society and electronic commerce (article 8).

Trade Me: Government and other legal content demands on Trade Me center around New Zealand's Harmful Digital Communications Act. In its 2017 transparency report, Trade Me explains the Act, how it works and reports on the number of relevant demands received under it.

9. Reporting on the number of accounts and items specified in demands

In addition to reporting on the number of notices or demands a company has received for each relevant content category (highlighted in best practice #2), companies should also report on the number of accounts and items specified in demands as this enables a better understanding of the full breadth of those demands. By reporting on the number of users/accounts and items/pieces of content that are specified in content demands, companies can highlight the scope and extent of requests, how many users and items within their network these requests could potentially impact, and allow comparison with the number of users or pieces of content actually impacted (best practice #10).

Currently, companies typically report on *either* the number of accounts specified in demands *or* the number of items specified in demands. Examples of these companies are offered below. At this point we believe the best practice should be to publish data about *both* specified accounts and specified content. However, because no companies currently do so, we gave credit in our charts to companies reporting on one or the other.

Dropbox: In its latest transparency report on government removal requests, which covers the second half of 2017, Dropbox highlights the number of accounts specified in requests it received in each relevant country. For example, Dropbox reports that in the Netherlands, 88 accounts were targeted by 45 requests.

Microsoft: In its latest transparency reporting on Right to Be Forgotten delisting requests, which covers the second half of 2017, Microsoft specifies the total number of URLs that were requested for delisting in each relevant country. For example, Microsoft reports that in Croatia, 20 URLs were requested to be delisted by 16 requests.

Twitter: In its latest transparency report on legal requests for content removal, which covers the second half of 2017, Twitter discloses the number of accounts specified by requests received in relevant countries. For example, In India, the company received 2 court order removal requests and 142 government removal requests (called “other legal demands” in the report). These requests specified 800 accounts.

10. Reporting on the number of accounts and items impacted by demands

In addition to reporting on the number of accounts and items specified in demands (highlighted in best practice #9), companies should also report on the number of accounts and items impacted by those demands. Such reporting offers the most direct measure of how many speakers and how much free expression is being silenced as a result of such demands (and how many/how much is being effectively defended by the company). It also enables a comparison of the requested impact versus the actual impact, which in turn offers a greater understanding of both the quality and legality of the requests being made and the company’s rates of compliance with those requests.

Currently, and similarly to best practice #9, companies typically report on either the number of accounts impacted by demands or the number of items impacted by demands. Although companies received credit for this best practice in the charts if they did either, the best practice would be to publish data on both, and here—unlike with best practice #9—there are companies that actually satisfy this

best practice. Some examples of companies that report on both the number of accounts and the number of items of content impacted are include:

Tumblr: In its latest transparency report on copyright and trademark requests, which covers the second half of 2017, Tumblr separately reports on the number of accounts, Tumblr posts, and individual pieces of content affected by DMCA notices on a monthly basis. For trademark complaints, also on a monthly basis, Tumblr specifies the number of URLs affected by requests, indicates the percentage of blog content that was removed compared to the amount of blog content that was complained about, and the percentage of URLs that were found to be misleading and were required to be changed, compared to the number of URLs complained of.

Twitter: In its latest transparency report on legal requests for content removal, which covers the second half of 2017, Twitter specifies the number of accounts and tweets “withheld” or restricted on a country by country basis as a result of legal requests.

11. Reporting on how the company responded to demands

Reporting on how a company responds to requests across different issue areas is vital for understanding how companies comply with legal frameworks, government demands and user requests. In addition, reporting on how a company responds to requests also highlights the role companies play in protecting or censoring speech. Currently, companies demonstrate this best practice in a variety of ways including reporting on the number or percentage of requests complied with or removal percentages associated with demands for each category of content. Companies that disclosed this data received credit for satisfying this best practice in the charts. However, the preferred method of satisfying this best practice is for companies to clearly break down responses to demands using categories that highlight the range of possible responses a company can employ when they receive a request. These categories could highlight, for example, the number or percentage of requests complied with in whole, complied with in part, or rejected. This method of reporting goes beyond simply reporting on the number of impacted users or pieces of content as specified in best practice #10, or stating a compliance rate that does not adequately summarize how a company could have responded to demands. Some examples of companies that exhibit this best practice in their transparency reports are:

Automattic In its latest transparency report on copyright requests, which covers the first half of 2018, Automattic breaks down its responses to requests using the following categories: “Percentage of

notices where some or all content was removed”, “Percentage of notices rejected as incomplete”, and “Percentage of notices rejected as abusive”.

Dropbox: In its latest transparency report on government removal requests, which covers the second half of 2017, Dropbox categorizes its responses to requests into three categories: content blocked in response to a request, content blocked pursuant to Acceptable Use Policy, and no action taken.

Kickstarter In its latest transparency report on copyright and trademark requests, which covers the year 2015, Kickstarter reports on the number of claims they rejected, as well as the possible reasons behind such rejections. For example, in its reporting on copyright claims, Kickstarter states “we reject claims when they are incomplete, when they involve material that can’t be protected under copyright, or when they target fair use.” In addition to reporting on the number of projects that were “hidden” (taken down) as a result of DMCA or trademark claims, Kickstarter also reports on the number of projects that were targeted by copyright and trademark requests, but that they were able to avoid hiding by “helping creators make project modifications to address a copyright claim” and “by encouraging claimants to resolve the dispute directly with the project creator or by helping the creators make a modification to their projects to address the claim”. For copyright and trademark claims, Kickstarter also reports on the number of projects that have been returned to public view and that remain hidden.

12. Additional best practices to consider

In addition to the 11 general best practices for content takedown reporting outlined above, companies should also aim to follow these additional best practices focused on enhancing the reports useability and supplementing its quantitative data with qualitative data and additional context:

Defining terms clearly: By clearly defining terms such as “content takedowns”, “content restrictions”, “network shutdowns” and “Community Guidelines-based takedowns”, a company can clarify what specific data points are being reported on in its transparency reports, and how reported data points differ from one another. Clearly defined terms also make cross-company comparisons easier. Companies that demonstrate this best practice often do so in the form of a glossary. Some examples of companies that exhibit this best practice in their transparency reports are:

Dropbox: In its latest transparency report, which covers the second half of 2017, Dropbox provides a definition and overview of the scope of information reported on for each request type. For example, regarding government removal requests, Dropbox highlights that its reporting on this area includes “court orders and written requests from law enforcement and government agencies seeking the removal of content based on the local laws of their respective jurisdictions”.

Telefonica: In its latest transparency report, which covers the year 2016, Telefonica has a specific glossary-like section dedicated to defining the different types of indicators and categories of requests they report on. For example, Telefonica defines “geographical and temporary suspensions of the service” as “A request from the competent authorities to temporarily or geographically limit the provision of a service. These requests are usually connected with circumstances involving situations of force majeure, such as natural catastrophes, acts of terrorism, etc.”

Telenor: In its March 2017 Authority Requests Disclosure Report, Telenor has a specific glossary-like section dedicated to defining the different types of authority requests they receive and engage with. In this section, it clearly defines “network shutdowns” as “requiring shut down of the operators network in part or in full” and “content restrictions” as “impos[ing] restriction on electronic content distributed through its network, such as blocking of URLs.”

Providing meaningful explanations of internal policies: By explaining the internal policies and guidelines companies use to process and act upon requests in a meaningful manner, a company can clearly communicate to users what types of requests they process, what factors are considered when processing requests, and what types of content requests fall out of scope. Some examples of companies that exhibit this best practice in their transparency reports are:

Daum Kakao: In the last Daum Kakao report that we surveyed, which covers the first half of 2018, Daum Kakao uses a **detailed flow-chart style infographic** to explain to users what process they use for “handling rights of infringement reports”.

Facebook: In the first half of 2018, Facebook published a comprehensive overview of its Community Standards, which delineate the kinds of content that it removes and the kinds of content they do not. For example, under the section “Violence and Criminal Behavior”, Facebook specifies that they “remove content, disable accounts, and

work with law enforcement” when they assess a “genuine risk of physical harm or direct threats to public safety.”

Google: In its latest YouTube Community Guidelines enforcement report, which covers the second quarter of 2018, Google has a video called “**The Life of a Flag**” which explains how flagged content is assessed and if needed, moderated.

Tumblr: In its latest transparency report on copyright and trademark requests, which covers the second half of 2017, Tumblr uses a **detailed flow-chart style infographic** to explain to users how they “handle copyright infringement notifications under the Digital Millennium Copyright Act (DMCA)”.

Vodafone: Although Vodafone does not publish quantitative information on content demands, and is therefore not otherwise included in this report, it does provide a significant amount of qualitative information on how it responds to legal and government demands for content blocking and network shutdowns, and what legal frameworks exist in each country it operates in that enable this. This includes outlining a set of “**Vodafone Freedom of Expression Principles**” which delineate what Vodafone does in response to requests, what it does not do and what it believes governments should do.

Offering case studies to illustrate the company's practices and the issues

it faces: By publishing case study examples a company can provide greater qualitative information on how they have processed past cases. This generates greater understanding on their policies and processes, as well as on the types of cases they typically receive and how they respond to them. Some examples of companies that exhibit this best practice in their transparency reports are:

Google: In its traffic and disruptions transparency report, Google provides links to external news sources that have reported on traffic and disruption events the company has experienced. In addition, in its reporting on Right to be Forgotten delisting requests (under “Search removals under European privacy law”), Google provides examples of requests it has received. Each example provides details on the country in which the request was received, the details of the request, and details on how Google responded.

Millicom: In its latest Law Enforcement Disclosure Report, which covers the year 2017, Millicom categorizes shutdown and disruption events as “major events”. For some of these cases, Millicom provides a contextual case study. For example, they explain that since 2014 there

has been an “ongoing shutdown of services in prisons in Central America” and provide relevant contextual information.

Reporting on specific notices where reasonable and permitted by law:

When possible and legal, and consistent with privacy and safety considerations, companies should report on the specific legal demands and content removal requests they have received. This enables for greater understanding on the prevalence and scope of content takedowns online. One way companies can do this is by reporting specific requests to the [Lumen Database](#), which is a project of the Berkman Klein Center for Internet & Society at Harvard University. Google, for example, does this when it receives and complies with a DMCA notice to remove content from its web search product. Instead of being able to see the copyrighted content, a user is instead directed to the [Lumen Database](#) where they can see the DMCA takedown notice. Similarly, GitHub hosts its own [repository](#) with text of DMCA takedown notices it has received.

Providing meaningful numbers that reflect how many pieces of content or accounts were blocked or otherwise restricted based on automated flagging or review: Companies utilize automated tools to seek to identify content that is illegal or that violates their Community Guidelines. Because such automated flags and even automated takedowns are prevalent, and raise unique due process and accuracy concerns, specific reporting on how often they are used—and how often they make mistakes—is desperately needed. Currently, [Facebook](#), [Google](#), and [Etsy](#) report on the use of automated tools, in relation to their Community Guidelines-based takedowns. This best practice should not only spread to other companies but also to other categories of reporting, because automated tools are increasingly being used to seek out and block a variety of different types of illegal and infringing content, beyond content that merely violates Community Guidelines.

Linking relevant reports to one another: All of a company’s various transparency reports ideally should be available through a single convenient portal. Companies should have a central webpage, list or drop-down menu where all of their current and past transparency reports can be easily accessed. Parent companies that own subsidiaries or products that publish independent transparency reports should similarly collect all of their available transparency reports in one central location.

Publishing reports at static and functioning URLs: Transparency reports are useless if people can’t find them. Maintaining static and functioning URLs is especially important for older versions of a company’s transparency report. If a company is acquired or re-branded and the links to its transparency reports subsequently change, they should clearly note this and provide directions to updated links.

Publishing data in a structured data format: Companies should make all report data available in a CSV (comma separated values) format, rather than or in addition to (e.g.) a flat PDF file. This format is most helpful to researchers, journalists, and others who want to make use of the report data, as it simplifies the data extraction process and makes reports more accessible.

Publishing reports using a non-restrictive Creative Commons license: Companies should use a non-restrictive Creative Commons license for their reports, such as the “ShareAlike” license, which “lets others remix, tweak, and build upon your work even for commercial purposes, as long as they credit you and license their new creations under the identical terms.” For more information on Creative Commons licensing, visit <https://creativecommons.org/licenses/>, and for information about choosing a Creative Commons license, visit <https://creativecommons.org/choose/>.

Offering a Frequently Asked Questions section: Although topics such as a company’s values, definitions of key terms and an explanation of why the company publishes a transparency report will likely be included in the narrative of the report, companies should also include a Frequently Asked Questions (FAQ) section. This goes above and beyond to inform readers about company practices and how content requests are handled.

Issue-Specific Reporting Best Practices

This section and the corresponding charts in the [appendix](#) will look at how well different companies are doing at applying the above best practices to the six different categories of content takedown reporting, while also identifying a few individual best practices that only apply to specific categories of reporting.

→ HOW TO READ THE CHARTS

Charts 2a through 6 in the [appendix](#) below outline how well the companies we surveyed complied with the 11 core best practices identified above. There is a chart for each category of content takedowns/reporting discussed above (i.e., government and other legal content demands, copyright requests, trademark requests, network shutdowns and service interruptions, and Right To Be Forgotten Requests). The one category for which there is not a chart is Community Guidelines-based takedowns; that area is so new and the practice so varied that a chart would have been of limited usefulness.

Each chart only includes data on companies who have issued reports on the type of content takedown in question. For example, Millicom is not listed in Chart [3b](#), which surveys international company copyright reporting as they do not report on copyright requests.

Designations of N/A in the tables indicate that a best practice is not applicable to a particular company because, for example, they only have one product, and therefore cannot report on requests by product, or only operate in one country, and therefore cannot break down demands by country.

Designations of N/R in the tables indicate that a company has received no requests across any time period for the specified content category, and therefore it is not possible to gauge their performance on certain best practices. Note that if a company has not received any demands for the specified content category during their most recent reporting period, we refer back to the most recent report which includes data on such demands. Instances where this data is referenced in the charts is indicated in footnotes. N/R is only utilized when no data on the specified demands are available from any reports.

For each category of takedown and reporting, we will also address any category-specific best practices that should be considered.

1. Reporting on government and other legal content demands

Charts [2a](#) and [2b](#) highlight how companies have been reporting on government and other legal content demands, including instances of content and IP restrictions and blocking. Generally, most companies report on the number of requests they received, and break down requests by country. They also generally report on how they responded to requests.

We don't have any category-specific best practices to offer here. However, it is worth noting that demands to DNS intermediaries to block content via the Domain Name System (DNS) are becoming increasingly prevalent. However, reporting on this type of content blocking is not yet common. We urge domain name registrars and registries to adopt the practice of transparency reporting going forward.

2. Reporting on copyright requests

Charts [3a](#) and [3b](#) highlight how companies have been reporting on copyright requests.

In addition to following the general best practices previously described, the best transparency reports on copyright requests also report on the following aspects:

Counter-Notices and Impact of Counter-Notices for DMCA-Based

Requests: A DMCA counter-notice is a notice submitted by a party who has been accused of infringing copyright in response to a DMCA takedown notice as a mechanism to counter the takedown and reinstate the content. There is debate over whether DMCA counter-notices are necessary or useful when it comes to protecting against improper removals, but there is little data to inform this debate. More reporting on the number of counter-notices that companies receive as well as the impact of those counter-notices would provide some much-needed perspective on this issue. Some companies that currently exhibit this best practice are:

Medium: In its latest transparency report, which covers the year 2015, Medium reports on the number of counter-notices received as well as the number of posts reinstated as a result.

Tumblr: In its latest transparency report on copyright and trademark requests, which covers the second half of 2017, Tumblr has a dedicated

section for DMCA counter-notices which highlights the number of counter-notices received, the number of counter-notices which were deemed valid, the number of accounts whose content was restored, the number of posts restored and the number of pieces of content restored per month during this reporting period.

Twitter: In its latest transparency report on copyright notices, which covers the second half of 2017, Twitter has a dedicated section for DMCA counter-notices which highlights the number of counter-notices received, the percentage of material restored as a result of the counter-notices, as well as the number of tweets restored and the number of pieces of media restored. This data is broken down on a monthly basis within each reporting period.

Format or Medium of Content Targeted by Copyright Requests: For copyright requests, the category of objectionable content targeted by demands is already evident: it is content that allegedly infringes on a copyright holder's rights. Therefore, when it comes to specifying the category of objectionable content in transparency reports on copyright requests, we instead encourage companies to specify the format or medium of the targeted content: for example, whether the targeted content is a video, an image, specific text, a specific link, etc. Copyright notices can target a range of content types including text, images, audio and video. This can be seen in practice in the example below:

Tumblr: In its latest transparency report on copyright and trademark requests, which covers the second half of 2017, Tumblr offers numbers of copyright notices received, broken down by the format or medium of the content targeted by the notice, including content types such as images, texts, audio, video, links, asks, quotes and chats.

Data on Who is Submitting Copyright Requests: As already described above, some companies have begun reporting on the top copyright reporters they receive notices from or otherwise highlighting the types of copyright holders who have submitted requests to the company, and we encourage this as a best practice. We also encourage companies, especially American companies, to report on non-DMCA copyright takedown requests they receive. Currently, **Medium** does this. However, little detail is provided on what the applicable laws and regulations around these requests were, and what the impact of these requests was.

3. Reporting on trademark requests

Charts 4a and 4b in the appendix highlight how companies have been reporting on trademark requests. In our survey we did not identify any best practices for reporting that were specific to trademark requests. As a result, Charts 4a and 4b only reflect performance against the general best practices for content takedown reporting mentioned previously. There was little standardization across the companies that were reporting on this issue area.

Currently, some companies, such as Facebook and Etsy, also report on other forms of intellectual property-related requests, such as those related to counterfeit and patent takedowns. Although these reports are not surveyed in this toolkit, we encourage platforms to report on these intellectual property-related takedowns if appropriate for their service.

4. Reporting on network shutdowns and service interruptions

Charts 5a and 5b highlight how companies have been reporting on network shutdowns and service interruptions. For telecommunications companies, this is usually in the form of a demand from a government to shut down or otherwise disrupt its offering of internet service, often to a specific geographic area. For internet platforms, these events are typically implemented by government agencies directly pressuring telecommunications to block the service for them, and as a result the affected companies themselves do not receive requests to execute such actions. Therefore, internet platform-related data in these charts primarily pertains to internet companies reporting on externally-implemented shutdown and disruption events that have impacted their services. Because network shutdowns and service interruptions are system- and service-wide disruptions, the best practice of reporting on the number of accounts and items specified in demands is not applicable to this content category.

Currently, only a handful of internet and telecommunications companies report on network disruptions and shutdowns, and there is little standardization across these forms of reporting. However, many of the companies that report on network disruptions and shutdowns do so in a detailed manner, and provide a large amount of qualitative information on these events.

In addition to reporting on the general best practices previously described, the best reporting on network shutdowns and service interruptions also includes the following information.

The date or date range of the network shutdown or service interruption:

Highlighting the date or date range of each network disruption or shutdown

event a company has experienced enables a company to demonstrate the extent of the impact on their business and their users. It also enables companies to highlight any ongoing social, political, or economic events that may have contributed to these disruption or shutdown events.

Facebook: In its latest transparency report, which covers the second half of 2017, Facebook reports that there were 2 nationwide network disruptions in Togo, the first between September 5, 2017 and September 10, 2017 and the second between September 19, 2017 - September 22, 2017, although it offers no other context for these disruptions.

Google: In its latest traffic and disruptions transparency report, Google specifies the date a network disruption event took place. For example, Google reports that on February 20, 2014, Google's web search experienced a disruption in Syria due to a internet shutdown there, and links to external coverage of that shutdown. In addition, in its reporting Google also specifies whether the disruption event is ongoing. For example, Google reports that there are ongoing disruptions to its web search feature in Congo-Brazzaville and China.

The duration of the network shutdown or service interruption:

Highlighting the specific duration of a network disruption or shutdown event enables a company to demonstrate the impact each of these events has had on their services.

Facebook: In its latest transparency report, which covers the second half of 2017, Facebook reports that they experienced 14 internet disruptions in India. These disruptions lasted a total duration of 5 weeks, 2 days, 21 hours. In its reporting, Facebook also provides the duration of each of these individual 14 internet disruptions. For example, Facebook reports that between July 6, 2017 and July 7, 2017, Facebook's services experienced a disruption in the Kashmir Valley for 13 hours.

Google: In its latest traffic and disruptions report, Google specifies the duration of network disruptions in days, hours and minutes. For example, on September 5, 2017, Google's web search experienced a disruption in Togo that lasted 4 days, 19 hours and 55 minutes.

The reasons behind network shutdowns and service interruptions:

Because the best practice of reporting on categories of objectionable content targeted by demands is not directly applicable during network shutdowns and service interruptions, companies who report on this category should instead report on the reasons behind their shutdowns and interruptions, if known.

Millicom: In its latest Law Enforcement Disclosure Report, which covers the year 2017, Millicom reports that since a specific date in 2014, authorities in El Salvador and Honduras have required telecommunications companies to shut down services or reduce signal capacity in and around prisons in order to prevent criminal gangs from using smuggled cell phones. Their discussion of this shutdown highlights that it is ongoing and because it is legally mandated will likely continue indefinitely.

5. Reporting on Right to be Forgotten delisting requests

Chart 6 highlights how companies have been reporting on Right to be Forgotten delisting requests. It reflects company performance against the general best practices for content takedown reporting outline above.

In regard to the best practice of reporting on the categories of objectionable content targeted, because Right to be Forgotten delisting requests only apply to search engines and their results, and only to particular types of data, the categorizations will likely be different than with other types of reporting. The most detailed and specific breakdown of reporting on this best practice we've seen so far is in **Google's** latest transparency report on search removals under European privacy law, where it provides numbers for a variety of different "categories of content requested for delisting". The categories of content include insufficient information, personal information, professional information, self-authored, crime and professional wrongdoing. **Google** also provides numbers broken down by "categories of websites hosting requested for delisting," including categories like news, directory and social media sites. We encourage other services to follow these best practices.

We also encourage all platforms that receive Right to be Forgotten delisting requests to report on those requests, even if they don't offer search services and don't believe the current regulations apply to them. If complainants are trying to use the Right to be Forgotten outside of the bounds of the law, that information should be publicized.

Notably, search engines themselves aren't the only ones that might have Right to be Forgotten-relevant data worth reporting. For example, the **Wikimedia Foundation** has a dedicated page where it posts "notices of delisted project pages" it has received from search engines that have delisted Wikimedia content as a result of Right to be Forgotten requests. However, because it does not report on requests it has directly received related to the Right to be Forgotten, it does not receive credit for doing so in **Chart 6**.

6. Reporting on Community Guidelines-based content takedowns

Reporting on content takedowns based on a company's Community Guidelines (also often known as Terms of Service or Community Standards) is one of the newer and emerging forms of transparency reporting. Until recently, the few companies that shared any data at all about their Community Guidelines-based content takedowns did so in an inconsistent manner, either failing to publish regularly or focusing only on takedowns of a particular kind of guidelines-violating content rather than comprehensively looking at the breadth of guidelines-based takedowns. The following companies offer examples of such limited reporting.

Automatic: Beginning in July 2017, Automatic began reporting on the number of reports of terrorist and extremist propaganda in violation of its terms that they received from government Internet Referral Units (IRU's) around the world. (We treat these as Community Guidelines-based removals because they are not legal demands.) Automatic also reports on the "number of notices where sites/content were removed as a result" and the "percentage of notices where sites/content were removed as a result". Automatic does not, however, report on the amount of extremist content they identify through other means like user flagging or automated systems and then remove based on their Community Guidelines (which they refer to as "User Guidelines" or "Terms of Service"), and it does not report on any other content aside from terrorist or extremist content.

Etsy: In its latest transparency report, which covers the year 2017, Etsy reports on the total number of accounts closed because of violations of its Seller Policy other than intellectual property violations. According to the report, "violations may range from reselling commercial items to listing prohibited items to repeatedly providing poor customer service." The Etsy team uses a combination of "automated systems, community flags, and proactive review" to enforce this policy. In the report, Etsy highlights that they received 394,564 flags regarding "potentially noncompliant listings, including both internal and community generated flags." Etsy also automatically processed 110,037 of those flags (28%) and the Etsy team manually reviewed 369,694 of those flags (94%), including some automatically flagged items to ensure accuracy. The report breaks down flags by source (community or Etsy) and breaks down accounts reviewed by region (Americas, EMEA and Asia Pacific). However, the report does not offer any granularity in terms of the types of policy violations at issue.

Medium: In its latest transparency report, which covers the year of 2015, Medium reports on takedown requests received “based on violations of Medium’s Terms of Service or Rules, including harassment and privacy” under the category “other takedown requests”. Medium’s reporting includes quantitative data on “total other takedown requests”, “posts removed in whole or part” and “account suspensions”. However, it does not provide a breakdown based on content category, nor does it indicate how much content was flagged by users, by governments, or by automated systems.

Microsoft: In its latest transparency report, which covers the second half of 2017, Microsoft reports on the number of “requests for removal of nonconsensual pornography (‘revenge porn’)” it has received. This reporting includes quantitative data on the total number of “requests reported”, “requests accepted” and the “percentage of requests accepted”. Microsoft however does not provide such reporting for other content categories and also does not indicate how much content was flagged by users, by governments, or by automated systems.

Twitter: In its latest reporting on “Government TOS reports”, which covers the second half of 2017, Twitter reports on “worldwide reports” the company received from “official government representatives through our standard customer support intake channels to review content against our terms of service. This section covers government requests to remove content that may violate Twitter’s Terms of Service (TOS) under the following Twitter Rules categories: abusive behavior, copyright, promotion of terrorism, and trademark. It does not include legal demands regardless of whether they resulted in a TOS violation.” For each of these categories, Twitter provides data on the number of “reports received” from governments, the number of “accounts reported” and the “percentage of accounts actioned.” However, Twitter does not report on the amount of content it or other non-government actors have flagged for removal based on the company’s Terms of Service or Community Guidelines. And although Twitter provides reporting on government requests regarding the above-specified content categories, it does not report on the full complement of content categories the platform prohibits.

Especially because the vast majority of takedowns that occur on major platforms are based on those platforms’ Community Guidelines, such limited reporting is not nearly sufficient. The need for in-depth reporting is made even greater by the fact that much of that takedown activity is done unilaterally by the companies—rather than in response to legal process subject to political checks and balances—and because that activity is often automated, raising unique due process, accuracy, and bias concerns.. Therefore, much greater transparency is needed on

how companies identify and remove (or otherwise restrict) content based on their own Community Guidelines.

In April 2018, both Facebook and Google's video service YouTube released transparency reports on the enforcement of its community standards and related content takedown efforts. Thus far these are the most comprehensive transparency reports on Community Guidelines-based content takedowns, and represent the latest best practices in this fast-evolving area.

Facebook: Facebook's Community Standards Enforcement Report covers three key metrics: how prevalent are Community Standards violations on Facebook (measured by estimating the percentage of views violating content receives), how much content the company takes action on, and how much violating content the company identifies before users do. The company also reported that it is working on instituting a fourth metric — how quickly it takes action on violations. The report covers six content categories: graphic violence, adult nudity and sexual activity, terrorist propaganda (ISIS, al-Qaeda and affiliates), hate speech, spam and fake accounts. These are the biggest categories of guidelines-violating content, but not every category; we hope that future reports are even more comprehensive.

Prior to the release of its Community Standards Enforcement Report, Facebook also announced it was broadening its appeals process, which enables users to appeal the removal of their content. Greater transparency on the number of appeals the company receives and the impact of these appeals is a logical next step in the transparency reporting process.

Google: Google's Community Guidelines enforcement report focuses on Community Guidelines-based content takedowns on YouTube. Google issues this report on a quarterly basis as opposed to bi-annually like the rest of Google's content-related transparency reports. The report highlights the number of videos removed based on the type of flagger (human or automated), the percentage of videos that were flagged by automated flagging tools and were removed before and after they were viewed, the percentage of human flaggers by type (government agency, NGO, user and individual trusted flagger), the top 10 countries by human flagging volume, and examples that provide further insight into the flagged video management process. The report also breaks down by category of violation the content removed as a result of human flagging. Unfortunately, it does not offer the same breakdown by category of violation for content originally flagged by automated means, which is especially concerning since automated systems flagged the vast majority of objectionable content.

The release of both Google and Facebook’s Community Guidelines content takedown reports are vital first steps towards increased transparency around content removals online. However, civil society will continue to push companies for greater transparency and accountability in this area. Currently, the best recommendations for reporting on Community Guidelines-based content takedowns are the [Santa Clara Principles on Transparency and Accountability Around Online Content](#). These principles were developed by New America’s Open Technology Institute and a coalition of organizations, advocates and academic experts who support the right to free expression online, with the intention of sparking further in-depth dialogue on Community Guidelines-based transparency reporting in the future. These principles outline minimum levels of transparency and accountability that companies should be providing around enforcement of their Community Guidelines. However, companies should also continue to develop their own unique metrics and forms of reporting, such as Facebook’s prevalence metric, as these also provide valuable insights into how content is removed and how content policies are enforced.

Below are the three Santa Clara Principles—Numbers, Notice, and Appeal—that were intended by the document’s signers “as initial steps that companies engaged in content moderation should take to provide meaningful due process to impacted speakers and better ensure that the enforcement of their content guidelines is fair, unbiased, proportional, and respectful of users’ rights.” As these principles represent some of the latest thinking from civil society and academic experts on the issue of content takedown transparency and accountability, and are mostly applicable to all categories of takedowns and not just Community Guidelines-based takedowns, they are an appropriate place for us to conclude our report. Go forth and be more transparent!

→ **BOX 1**

The Santa Clara Principles on Transparency and Accountability Around Online Content Moderation:

NUMBERS: Companies should publish the numbers of posts removed and accounts permanently or temporarily suspended due to violations of their content guidelines.

At a minimum, this information should be broken down along each of these dimensions:

- Total number of discrete posts and accounts flagged.

- Total number of discrete posts removed and accounts suspended.
- Number of discrete posts and accounts flagged, and number of discrete posts removed and accounts suspended, by category of rule violated.
- Number of discrete posts and accounts flagged, and number of discrete posts removed and accounts suspended, by format of content at issue (e.g., text, audio, image, video, live stream).
- Number of discrete posts and accounts flagged, and number of discrete posts removed and accounts suspended, by source of flag (e.g., governments, trusted flaggers, users, different types of automated detection).
- Number of discrete posts and accounts flagged, and number of discrete posts removed and accounts suspended, by locations of flaggers and impacted users (where apparent).

This data should be provided in a regular report, ideally quarterly, in an openly licensed, machine-readable format.

NOTICE: Companies should provide notice to each user whose content is taken down or account is suspended about the reason for the removal or suspension.

In general, companies should provide detailed guidance to the community about what content is prohibited, including examples of permissible and impermissible content and the guidelines used by reviewers. Companies should also provide an explanation of how automated detection is used across each category of content.

When providing a user with notice about why a post has been removed or an account has been suspended, a minimum level of detail includes:

- URL, content excerpt, and/or other information sufficient to allow identification of the content removed.
- The specific clause of the guidelines that the content was found to violate. How the content was detected and removed (flagged by other users, governments, trusted flaggers, automated detection, or external legal or other complaint). The identity of individual flaggers should generally not be revealed, however, content flagged by governments should be identified as such, unless prohibited by law.

- Explanation of the process through which the user can appeal the decision.

Notices should be available in a durable form that is accessible even if a user's account is suspended or terminated. Users who flag content should also be presented with a log of content they have reported and the outcomes of moderation processes.

APPEAL: Companies should provide a meaningful opportunity for timely appeal of any content removal or account suspension.

Minimum standards for a meaningful appeal include:

- Human review by a person or panel of persons that was not involved in the initial decision.
- An opportunity to present additional information that will be considered in the review.
- Notification of the results of the review, and a statement of the reasoning sufficient to allow the user to understand the decision.

In the long term, independent external review processes may also be an important component for users to be able to seek redress.

Appendix

How to Read the Charts

Charts 2a through 6 in the appendix below outline how well the companies we surveyed complied with the 11 core best practices identified above. There is a chart for each category of content takedowns/reporting discussed above (i.e., government and other legal content demands, copyright requests, trademark requests, network shutdowns and service interruptions, and Right To Be Forgotten Requests). The one category for which there is not a chart is Community Guidelines-based takedowns; that area is so new and the practice so varied that a chart would have been of limited usefulness.

Each chart only includes data on companies who have issued reports on the type of content takedown in question. For example, Millicom is not listed in Chart 3b, which surveys international company copyright reporting as they do not report on copyright requests.

Designations of N/A in the tables indicate that a best practice is not applicable to a particular company because, for example, they only have one product, and therefore cannot report on requests by product, or only operate in one country, and therefore cannot break down demands by country.

Designations of N/R in the tables indicate that a company has received no requests across any time period for the specified content category, and therefore it is not possible to gauge their performance on certain best practices. Note that if a company has not received any demands for the specified content category during their most recent reporting period, we refer back to the most recent report which includes data on such demands. Instances where this data is referenced in the charts is indicated in footnotes. N/R is only utilized when no data on the specified demands are available from any reports.

Chart 1a: Company Reporting Across the Six Issue Areas - U.S.-Based Companies

Company	Country of HQ	Last Reported Date	Government and Other Legal Content Demands	Copyright Requests	Trademark Requests	Network Shutdowns and Service Interruptions	Right to Be Forgotten Delisting Requests	Community Guidelines-Based Content Takedowns
Amazon	USA	June 2018	✓					
Apple	USA	December 2017	✓ ¹					
Automattic	USA	June 2018	✓	✓	✓			✓
Cheezburger	USA	December 2014		✓	✓			
DreamHost	USA	December 2014	✓	✓	✓			
Dropbox	USA	June 2017	✓					
Etsy	USA	December 2017		✓	✓			✓
Facebook	USA	December 2017	✓	✓	✓	✓		✓
GitHub	USA	December 2017	✓	✓				
Google	USA	December 2017 ²	✓	✓		✓	✓	✓ ³
Kickstarter	USA	December 2015		✓	✓			
Mapbox	USA	N/A	✓	✓				
Medium	USA	December 2015	✓	✓	✓			✓
Microsoft	USA	December 2017	✓	✓			✓	✓
Oath ⁴	USA	December 2017	✓					
Pinterest	USA	June 2018	✓					
Reddit	USA	December 2017	✓	✓				
Snap Inc.	USA	December 2017	✓	✓				
Tumblr	USA	December 2017		✓	✓			
Twitter	USA	December 2017	✓	✓	✓			✓
Wikimedia Foundation	USA	June 2018	✓	✓				
AT&T	USA	June 2018	✓					
CREDO Mobile	USA	June 2018	✓					
Sonic	USA	December 2016	✓					
Verizon	USA	June 2018	✓					

¹ Apple reports on "Worldwide Government Account Deletion/Restriction Requests," and on United States Private Party Requests for Account Restriction/Deletion. However, they do not report on requests regarding the removal of specific content (as opposed to accounts), such as requests to remove specific apps from the Apple Store.

² Google generally issues transparency reports on a bi-annual basis. However, its YouTube Community Guidelines enforcement report is issued on a quarterly basis, and its Search Removals Under European Privacy Law (Right to be Forgotten) report is updated daily.

³ Google publishes a comprehensive transparency report on community guidelines-based takedowns for YouTube only.

⁴ Oath is a subsidiary of Verizon Communications that serves as the umbrella company of its digital content subdivisions, including AOL and Yahoo!.

Chart 1b: Company Reporting Across the Six Issue Areas - International Companies

Company	Country of HQ	Last Reported Date	Government and Other Legal Content Demands	Copyright Requests	Trademark Requests	Network Shutdowns and Service Interruptions	Right to Be Forgotten Delisting Requests	Community Guidelines-Based Content Takedowns
Atlassian	Australia	December 2017	✓					
CyberGhost	Romania	December 2015		✓				
Daum Kakao ¹	South Korea	June 2018	✓	✓	✓			
LeaseWeb	Netherlands	December 2013	✓					
LINE	Japan	December 2017	✓					
Trade Me	New Zealand	June 2018 ²	✓					
Liquid Telecom Kenya	South Africa	December 2016				✓		
Millicom	Luxembourg	December 2017	✓			✓		
Telefonica	Spain	December 2016	✓			✓		
Telenor	Norway	March 2017	✓			✓		

¹ Since this survey of companies was completed, Daum Kakao has issued a new transparency report that unfortunately is much less detailed than its previous reports. The data and examples regarding Daum Kakao in this toolkit are based on its prior report, found at [Archive.org](#).

² Trade Me's most recent transparency report was published for the period of July 1 2017 - June 30 2018. However, reporting on the Harmful Digital Communications Act for this period was less detailed than their prior report. As a result, Trade Me's best practice reporting in this chart is based on their last transparency report, which covers the reporting period of July 1 2016 - June 30 2017.

Chart 2a: Reporting on Government and Other Legal Content Demands - U.S.-Based Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on categories of objectionable content targeted by demands	Reporting on products targeted by demands	Reporting on specific government agencies that submitted demands	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in demands ¹	Reporting on the number of accounts and items impacted by demands ²	Reporting on how the company responded to demands
Amazon ³	✓	✓								✓
Apple	✓	✓	✓					✓		✓
Automattic	✓	✓	✓		✓			✓		✓
Cheezburger	✓	✓	N/R	N/R	N/R	N/R	N/R	N/R	N/R	N/R
DreamHost	✓	✓	✓ ⁴						✓ ⁵	✓
Dropbox	✓	✓	✓					✓		✓
Facebook	✓		✓						✓	
GitHub	✓	✓								✓
Google	✓	✓	✓	✓	✓	✓		✓		✓
Mapbox ⁶	✓	✓	✓	N/R	N/A	N/R	N/R	N/R	N/R	N/R
Medium	✓	✓	N/R	N/R	N/A	N/R	N/R	N/R	N/R	N/R
Microsoft	✓	✓	✓							✓
Oath	✓	✓	✓					✓		✓
Pinterest	✓	✓			N/A					✓
Reddit	✓	✓	✓		N/A				✓	✓
Snap Inc.	✓	✓	✓							✓
Twitter	✓	✓	✓					✓	✓	✓
Wikimedia Foundation	✓	✓	✓		✓					✓
AT&T	✓	✓	✓							
CREDO Mobile ⁷	✓	✓	N/R	N/R	N/A	N/R	N/R	N/R	N/R	N/R
Sonic ⁸	✓	✓	N/R	N/R	N/A	N/R	N/R	N/R	N/R	N/R
Verizon ⁹	✓		✓						✓	

¹ Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

² Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

³ Amazon's most recent transparency report was published for the period of January-June 2018. However, in this period and in both reporting periods in 2017 and 2016, there were no content removal demands. As a result, Amazon's best practice reporting in this chart is based on the last available report indicating removal demands, which was during the reporting period of January-May 2015.

⁴ Although DreamHost reports the specific total numbers of different types of demands it receives, it then combines all of those different types of requests when reporting numbers broken down by country. The best practice would be to provide specific numbers in regard to each type of request both in total and when broken down by country.

⁵ Although DreamHost reports the specific total numbers of different types of demands it receives, it then combines all of those different types of requests when reporting the number of accounts impacted by demands. The best practice would be to provide specific numbers in regard to each type of request both in total and when broken down by number of accounts impacted by demands.

⁶ Mapbox does not specify the date of its transparency report and as a result it is not possible to refer back to older transparency reports as was done in the case of Amazon. The designation N/R is therefore used here. In the current report, Mapbox states they have received 0 "government requests to withhold content."

⁷ Beginning in 2015, CREDO Mobile began reporting on the "number of governmental requests to remove user content or accounts". However, they have not yet received such a request. As a result, it is not possible to determine their performance in reporting on these best practices and the designation N/R is used here.

⁸ Thus far, Sonic has not received any "government take down requests." As a result, it is not possible to determine their performance in reporting on these best practices and the designation N/R is used here.

⁹ In its latest, available transparency report, Verizon notes that they "did not receive a request from any government to remove user content in this reporting period." The data reflected in this chart is based on their reporting on "blocking demands".

Chart 2b: Reporting on Government and Other Legal Content Demands - International Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on categories of objectionable content targeted by demands	Reporting on products targeted by demands	Reporting on specific government agencies that submitted demands	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in demands ¹	Reporting on the number of accounts and items impacted by demands ²	Reporting on how the company responded to demands
Atlassian	✓	✓	³					✓	✓	✓
Daum Kakao	✓	✓	N/A ⁴		✓	✓	✓			✓
LeaseWeb	✓	✓								
LINE ⁵	✓	✓		N/R	N/R	N/R	N/R	N/R	N/R	N/R
Trade Me	✓	✓	N/A ⁶		N/A		✓			✓
Millicom			✓	✓	N/A					
Telefonica	✓	✓	✓				✓		✓	✓
Telenor	✓	✓	✓		N/A					

¹ Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

² Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

³ Atlassian did not receive credit for breaking down demands by country as they only report on "US government requests for account removal/content takedown", despite being an Australian company that operates in multiple nations.

⁴ Daum Kakao primarily operates in South Korea under South Korea laws. As a result, all reporting is focused on South Korea rather than being broken down by country.

⁵ Between July 2016 to December 2017 (the periods LINE has thus far issued transparency reports for) LINE reports receiving "no effective deletion requests" from the authorities. As a result, it is not possible to determine their performance in reporting on these best practices and the designation N/R is used here.

⁶ Trade Me primarily operates in New Zealand under New Zealand laws. As a result, all reporting is focused on New Zealand rather than being broken down by country.

Chart 3b: Reporting on Copyright Requests - International Companie

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on the format or medium of content being targeted by requests	Reporting on products targeted by requests	Reporting on specific parties that submitted requests ¹	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in requests ²	Reporting on the number of accounts and items impacted by requests ³	Reporting on how the company responded to requests	Specifying number of counter-notices received	Specifying outcome of counter-notices
CyberGhost	✓	✓	✓		N/A		✓					
Daum Kakao	✓	✓	N/A ⁴		✓	✓	✓			✓	N/A ⁵	N/A

¹ This column represents performance against the best practice of highlighting who submits copyright requests to companies. Rather than disclose the names of all individual copyright holders who file copyright notices to companies, companies can highlight the top copyright reporters they receive notices from or the type of copyright holders who have submitted requests to the company.

² Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

³ Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

⁴ Daum Kakao primarily operates in South Korea under South Korea laws. As a result, all reporting is focused on South Korea rather than being broken down by country.

⁵ Because Daum Kakao is a South Korean company that reports on copyright requests received under South Korean law, DMCA based processes such as counter-notices are not relevant.

Chart 3b: Reporting on Copyright Requests - International Companie

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on the format or medium of content being targeted by requests	Reporting on products targeted by requests	Reporting on specific parties that submitted requests ¹	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in requests ²	Reporting on the number of accounts and items impacted by requests ³	Reporting on how the company responded to requests	Specifying number of counter-notices received	Specifying outcome of counter-notices
CyberGhost	✓	✓	✓		N/A		✓					
Daum Kakao	✓	✓	N/A ⁴		✓	✓	✓			✓	N/A ⁵	N/A

¹ This column represents performance against the best practice of highlighting who submits copyright requests to companies. Rather than disclose the names of all individual copyright holders who file copyright notices to companies, companies can highlight the top copyright reporters they receive notices from or the type of copyright holders who have submitted requests to the company.

² Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

³ Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

⁴ Daum Kakao primarily operates in South Korea under South Korea laws. As a result, all reporting is focused on South Korea rather than being broken down by country.

⁵ Because Daum Kakao is a South Korean company that reports on copyright requests received under South Korean law, DMCA based processes such as counter-notices are not relevant.

Chart 4a: Reporting on Trademark Requests - U.S.-Based Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on the format or medium of content being targeted by requests	Reporting on products targeted by requests	Reporting on specific parties that submitted requests	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in requests ¹	Reporting on the number of accounts and items impacted by requests ²	Reporting on how the company responded to requests
Automattic	✓	✓			✓					✓
Cheezburger ³	✓	✓	N/R	N/R		N/R	N/R	N/R	N/R	N/R
DreamHost	✓		✓ ⁴						✓ ⁵	✓
Etsy	✓	✓			N/A				✓ ⁶	
Facebook	✓	✓			✓ ⁷				✓	✓
Kickstarter	✓	✓			N/A			✓	✓	✓
Medium	✓	✓			N/A				✓	
Tumblr ⁸	✓	✓			N/A			✓	✓	
Twitter	✓	✓							✓	

¹ Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

² Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

³ In its most recent and available transparency report (2014), Cheezburger indicated that it received 0 "trademark-related requests". As a result, it is not possible to determine their performance in reporting on many of these best practices and the designation N/R is used here.

⁴ Although DreamHost reports the specific total numbers of different types of demands it receives, it then combines all of those different types of requests when reporting numbers broken down by country. The best practice would be to provide specific numbers in regard to each type of request both in total and when broken down by country.

⁵ Although DreamHost reports the specific total numbers of different types of demands it receives, it then combines all of those different types of requests when reporting the number of accounts impacted by demands. The best practice would be to provide specific numbers in regard to each type of request both in total and when broken down by number of accounts impacted by demands.

⁶ Although Etsy reports the specific total numbers of different types of demands it receives, it then combines all of those different types of requests when reporting the number of accounts impacted by demands. The best practice would be to provide specific numbers in regard to each type of request both in total and when broken down by number of accounts impacted by demands.

⁷ Facebook's report on trademark requests includes data on multiple products, namely Facebook and Instagram. As a result Facebook received credit for reporting on the products targeted by trademark requests. However, the best demonstration of this best practice would be for Facebook to report on trademark requests in relation all of its products.

⁸ Tumblr is a subsidiary of Oath. Oath does not report on trademark separately and directs readers instead to Tumblr's transparency report, but that report does not cover other Oath products like Yahoo! and AOL. As a result, for the purposes of this toolkit, Oath does not receive credit for reporting on trademark requests, while Tumblr is listed here as a separate entity from Oath in order to recognize the exceptional quality of its reporting in this area.

Chart 4b: Reporting on Trademark Requests - International Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on the format or medium of content being targeted by requests	Reporting on products targeted by requests	Reporting on specific parties that submitted requests	Specifying which laws pertain to specific demands	Reporting on the number of accounts and items specified in requests ¹	Reporting on the number of accounts and items impacted by requests ²	Reporting on how the company responded to requests
Daum Kakao	✓	✓	N/A ³		✓		✓			✓

¹ Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

² Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.

³ Daum Kakao primarily operates in South Korea under South Korea laws. As a result, all reporting is focused on South Korea rather than being broken down by country.

Chart 5a: Reporting on Network Shutdowns and Service Interruptions - U.S.-Based Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down network shutdowns/service interruptions by country	Reporting on the reasons behind network shutdowns/service interruptions	Reporting on products targeted by network shutdowns/service interruptions	Reporting on specific government agencies that submitted demands	Specifying which laws pertain to specific demands	Reporting on the number of accounts impacted by network shutdowns/service interruptions	Reporting on how the company responded to demands	Specifying date or date range of network shutdowns/service interruptions	Specifying duration of network shutdowns/service interruptions
Facebook	✓	✓	✓							✓	✓
Google	✓	✓	✓	✓	✓					✓	✓

Chart 5b: Reporting on Network Shutdowns and Service Interruptions - International Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down network shutdowns/service interruptions by country	Reporting on the reasons behind network shutdowns/service interruptions	Reporting on products targeted by network shutdowns/service interruptions	Reporting on specific government agencies that submitted demands	Specifying which laws pertain to specific demands	Reporting on the number of accounts impacted by network shutdowns/service interruptions	Reporting on how the company responded to demands	Specifying date or date range of network shutdowns/service interruptions	Specifying duration of network shutdowns/service interruptions
Liquid Telecom Kenya ¹	✓	✓	N/A	N/R	N/A	N/R	N/R	N/R	N/R	N/R	N/R
Millicom	✓	✓	✓	✓	N/A		✓			✓	✓
Telefonica	✓	✓	✓				✓	✓	✓		
Telenor	✓	✓	✓		N/A						

¹ In its most recent and available sustainability report for the period of 2016-2017, Liquid Telecom Kenya notes that “there were no requests for Internet censorship or blocking of Internet to any customer received from any agency in Kenya.” As a result, it is not possible to determine their performance in reporting on these best practices and the designation N/R is used here.

Chart 6: Reporting on Right to Be Forgotten Delisting Requests - U.S.-Based Companies

Company	Issuing reports specific to the type of demand	Reporting on types of demands using specific numbers	Breaking down demands by country	Reporting on categories of content targeted by requests	Reporting on the categories of websites targeted by requests	Reporting on specific parties that submitted requests	Specifying which laws pertain to specific demands	Reports on the number of accounts or items specified in requests ¹	Reporting on the number of accounts or items impacted by requests ²	Reporting on how the company responded to requests
Google	✓	✓	✓	✓	✓	N/A	✓	✓	✓	
Microsoft	✓	✓	✓			N/A	✓	✓	✓	

¹ Although this chart gives credit to companies that either report on the number of accounts specified in demands or the number of items specified in demands, the best practice for companies is to report on both.

² Although this chart gives credit to companies that either report on the number of accounts impacted by demands or the number of items impacted by demands, the best practice for companies is to report on both.



This report carries a Creative Commons Attribution 4.0 International license, which permits re-use of New America content when proper attribution is provided. This means you are free to share and adapt New America's work, or include our content in derivative works, under the following conditions:

- **Attribution.** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

For the full legal code of this Creative Commons license, please visit **creativecommons.org**.

If you have any questions about citing or reusing New America content, please visit **www.newamerica.org**.

All photos in this report are supplied by, and licensed to, **shutterstock.com** unless otherwise stated. Photos from federal government sources are used under section 105 of the Copyright Act.