April 25, 2018

Twitter's Responses to Chairman Grassley's April 10, 2018, Data Privacy Questions

1. **What are your current policies and procedures with respect to the sharing of data with third party developers, including how you notify users of such sharing and/or request their consent?**

   a. **How have these policies and procedures evolved/changed since 2010?**
   b. **Do you intend to make any changes in light of recent events?**

Twitter is public by its nature. That is the key feature of our platform and what sets us apart from many other internet companies. Through this public platform, we are committed to providing a service that fosters and facilitates free and open democratic debate, and we do so by making it possible for people to react to, comment on, engage with, and criticize content that they or other accounts choose to share. Public Tweets are viewable and searchable by any person who accesses our platform. While we do offer people the ability to share non-public, protected Tweets with their authorized followers, the vast majority of people on Twitter choose to engage with public content and to post their content for public engagement. Users thus make decisions about what they are sharing with the world, as opposed to just their friends or followers.

Twitter is also unique in the level of transparency that we provide about public Tweet data. Through our free public and commercial application programming interfaces ("APIs") and paid data products, we make available to people using Twitter, developers, researchers, and other third parties access to public Tweets. We encourage developers to create products using this data for purposes that serve the public interest and the general Twitter community. This service is a hallmark of our commitment to transparency, collaboration, and innovation.

To protect the data of people using Twitter, we implemented and enforce a strict "know your customer" policy for our Enterprise data products. Each year, we receive and review approximately 1,000 applications for access to our Enterprise data products (*i.e.*, paid access to public Tweet data). While the majority of those applications are for uses consistent with our policies, we require all customers to go through a thorough review process that involves our Trust & Safety and Legal teams.

Through that process, we examine both the party seeking access and the stated purpose for seeking access to public Twitter content. On average, less than half of the applications require additional information or review by our policy team. However, when we identify an application that we determine to be inconsistent with our policies—or one that represents a potential risk to Twitter users or to the general Twitter community—we reject the application and decline to provide access to our Enterprise APIs. Approximately 5% of the applications we receive each year are rejected on that basis.

We provide access to public data to third parties engaged in research, marketing, and advertising, but we strictly prohibit—and will decline applications for—the use of Twitter data for a number of other purposes. As stated in our Developer Agreement and Policy, developers are prohibited from using Twitter content for targeting, segmenting, or profiling individuals based on health, negative financial status or condition, political affiliation or beliefs, race, ethnicity, religious or philosophical affiliation or beliefs, sex life or sexual orientation, trade union membership, data relating to any alleged or actual commission of a crime, or any other sensitive categories of personal information prohibited by law.

We further prohibit any entity accessing Twitter content from conducting any analyses or research that isolates a group of individuals or any single individual for any unlawful or discriminatory purpose or in a manner that would be inconsistent with our users' reasonable expectations of privacy. Our policies also prohibit the use of Twitter content (and information derived from such content) for surveillance purposes, or for investigating or tracking Twitter users or their content, or for tracking, alerting, or other monitoring of sensitive events (*e.g.*, protests, rallies, or community organizing meetings). *See* https://developer.twitter.com/en/developer-terms/agreement-and-policy.

Even for permitted users and uses, we restrict the types of data that we make accessible through our APIs. For example, we do not provide a user's phone number or email address alongside data provided via any of our APIs. In addition, we provide access to non-public communications such as Direct Messages ("DMs") only to the users who send and receive them, as well as to developers that users have expressly authorized to manage their DMs (*e.g.*, automated customer service).

We provide an Ads API specifically for customers who advertise with Twitter so that they can manage their ad campaigns. Via this API, we share programmatically the same type of data that would be displayed for an advertiser on ads.twitter.com when a user sees or engages with an advertiser's ad. We share this data with advertising customers about their ad campaigns, as well as with developers whom advertising customers have expressly authorized to manage their Twitter ad campaigns. *See* https://developer.twitter.com/en/docs/ads/general/overview/adsapi-application.

Our privacy policy, which is accessible on our website to both users and non-users, explains our data-sharing policies. Users who open a Twitter account are given an opportunity to review our privacy policy and are asked to provide their consent. In addition to notifying our users about how Twitter shares public data via our platform and APIs, our privacy policy also notifies users about how they can control the data we share:

> We share or disclose your personal data with your consent or at your direction, such as when you authorize a third-party web client or application to access your account or when you direct us to share your feedback with a business. If you've shared information like Direct Messages or protected Tweets with someone else who accesses Twitter through a third-party service, keep in mind that the information may be shared with the third-party service.
>
> Subject to your settings, we also provide certain third parties with personal data to help us offer or operate our services. For example, we share with advertisers the identifiers of devices that saw their ads, to enable them to measure the effectiveness of our advertising business. We also share device identifiers, along with the interests

> or other characteristics of a device or the person using it, to help partners decide whether to serve an ad to that device or to enable them to conduct marketing, brand analysis, interest-based advertising, or similar activities. You can learn more about these partnerships in our Help Center, and you can control whether Twitter shares your personal data in this way by using the "Share your data with Twitter's business partners" option in your Personalization and Data settings. (This setting does not control sharing described elsewhere in our Privacy Policy, such as when we share data with our service providers.) The information we share with these partners does not include your name, email address, phone number, or Twitter username, but some of these partnerships allow the information we share to be linked to other personal information if the partner gets your consent first.

See https://twitter.com/en/privacy.[1] We explain in our privacy policy and identify in our help center who our current select partners are and how they are permitted to use the data. In addition, our privacy policy links to the user's Personalization and Data settings, where users are able to opt out of sharing any such data with our select partners. Through our Personalization and Data settings, we also provide users with the ability to opt out of receiving personalized ads on Twitter based on off-Twitter data, personalization across devices, personalization based on the places a user has been, and tracking where a user sees Twitter content across the web.

While policies and procedures for providing Twitter user data to Enterprise partners in a way that protects our users have always been in place, they have evolved over the years to accommodate new and emerging uses, and to address common questions and concerns. In 2014, for example, Twitter began incorporating a partner's use case (*i.e.*, the specified purpose for accessing data) into the contractual language of the licensing agreement (a process that was memorialized in other documentation, but was not a part of the agreement up until that point). The revised approach resulted in greater scrutiny applied to our review and approval of developer use cases, reflected in the application review process that is in place today. We continue to adapt and improve this process.

In addition, over the years, we have clarified and detailed with greater specificity prohibited uses of Twitter user data. In the past, the standard Developer Agreement or commercial licensing terms were subject to Twitter's "Developer Rules of the Road," which provided high-level guidance requiring developers to respect our users' reasonable expectations of privacy. In June 2016, Twitter updated its standard licensing terms to include additional express restrictions designed to bolster user data protections (discussed below in response to Question 2). We have adopted and implemented additional restrictions in recent years, and we continue to reexamine and update relevant policies and procedures to ensure that data of those who use our services remain secure and free from abuse.

[1] On April 24, 2018, Twitter updated its privacy policy, effective May 25, 2018. The language above is drawn from Twitter's updated Privacy Policy, which expands on the language from Twitter's current Privacy Policy (also available at the link above). In both versions, users are linked to the Help Center and their Personalization and Data settings.

Our user protection policies are well-established and long-standing. We also continually refine these policies. For example, in June 2016 we introduced clarifications to further restrict developers from using the data they obtain from us for monitoring certain sensitive categories, including, but not limited to, political affiliations and beliefs, health status, race, and national origin. We have also expanded our use case review process to include a customer background check and ongoing due diligence and monitoring of existing customers. In addition, our use case review now applies to a greater number of developers, and includes not only Enterprise partners but also Premium API endpoints (*see* https://developer.twitter.com/en/pricing).

2. How do you ensure that user data shared with third party developers is not improperly transferred or used?

New Enterprise and Premium API customers and data partners go through a review process before they are granted access to Twitter user data. That review examines the company's history, the proposed use of the data, and privacy and security considerations designed to prevent misuse of the data.

Before developers receive access to any Twitter content, all Enterprise partner applications are manually reviewed to ensure that the proposed use of Twitter user data falls within the scope of permitted uses and is in compliance with our developer agreements and policies. We also conduct standard background checks and reviews of prospective customers, to ensure we have a full understanding of a prospective customer's business and end users.

If, in the course of our manual review, we determine that the application warrants further review (*e.g.*, the company, proposed use, developer geographic location, or ascertained end users fall into predetermined high-risk categories), our Trust & Safety Team conducts an additional review of the developer's application and proposed use of Twitter user data. This step often requires additional information or commitments from the developer about their intended uses and plans for ensuring compliance with Twitter's policies.

We routinely reject use cases that do not comply with our rules, and we often require customers to revise and resubmit applications in order to ensure that they are in compliance with our policies. Each year, we review approximately 1,000 use cases, 40% of which are then escalated for further review by our Trust & Safety Team. Once we approve the developer's proposed use case, the stated use is incorporated into the agreement and is binding on the developer. The developer's use of Twitter user data is limited to the scope set forth in that agreement. Deviation from the pre-approved use case in any form is grounds for enforcement actions, including immediate termination of the agreement and of further access to Twitter data.

As described in greater detail below (*see* response to Question 5), as part of our data compliance program, we also regularly investigate our data partners' use of Twitter data in order to assess whether their activities are within the scope of their agreement and consistent with our policies. We examine suspicious indicators that we detect through activity on our platform, and we investigate companies' use of Twitter content if and when concerns or suspicions arise. We also have a process for receiving, investigating, and acting on tips from external parties about potential misuses of Twitter user data.

In addition, our Developer Policy places a number of restrictions on how all developers—Enterprise partners as well as users of our Premium and public APIs—may use Twitter user data. We require that they obtain the user's express consent prior to taking a number of actions, including: (1) taking any actions through a user's account, including posting content, following/unfollowing other users, modifying profile information, starting a Periscope Broadcast, or adding hashtags or other data to the user's Tweets (authorization of a third-party application does not by itself constitute user consent); (2) republishing content accessed by means other than via Twitter APIs or other Twitter tools; (3) using a user's content to promote a commercial product or service; (4) storing non-public content such as DMs or other private or confidential information; and (5) sharing or publishing protected content and private or confidential information. In addition, our Developer Policy and Agreement strictly prohibits those who gain access to Twitter user data from sharing the data with other parties for any prohibited uses. *See* https://developer.twitter.com/en/developer-terms/agreement-and-policy. And, as noted above, we do not provide a user's phone number or email address alongside data provided via any of our APIs. We also prohibit developers who access data via our Ads API from sharing the data we provide to them with anyone other than the specific advertising customers they access the data on behalf of.

3. Do you limit the ability of third party developers to collect data beyond the scope of their application? If so, how do you ensure that third party developer agreements are limited in scope?

Other than the limited sampling of public data we offer through our free public APIs, we grant third-party developers access to Twitter user data for specific approved uses only. When we review an application, we evaluate the scope of the data requested against the developer's stated purpose for requesting access. That evaluation is designed to ensure that each party's access to the data is narrowly tailored to the stated objective for seeking Twitter user data. Additional information about how we enforce our policies is contained in response to Question 5, below.

4. What remedies do you have against a third party developer who exceeds the scope of their access?

If Twitter determines that developers are in violation of our policies—including by exceeding the scope of their permitted use of Twitter user data—we take appropriate action, which could include suspension or termination of the third-party developer's access to Twitter's APIs and data products. In addition, for customers who are in the process of renewing their developer agreement, Twitter will not process that renewal if the developer is found to have outstanding compliance issues.

We also vigilantly monitor applications that access our public APIs, and we do so primarily by leveraging automated detection tools. Those tools are able to detect prohibited activity and suspend those applications found to be engaging in such activity before they have the opportunity to impact the Twitter user's experience on the platform.

With respect to our Ads API, all third-party applications using the Ads API must meet our compliance standards before they are granted access to Twitter user data. If an Ads API developer is granted access to Twitter data, but is later found to have violated our rules and policies, Twitter will immediately suspend that developer's access to the Ads API. We will not lift the suspension unless and until the developer rectifies its prohibited use and demonstrates that it is in compliance.

We recognize that technology evolves and that bad actors' methods of misusing Twitter user data will likely evolve as well. We are committed to continuing to invest all necessary resources to ensure that our platform remains safe for our users, that our data partners remain compliant with our developer agreements and policies, and that our users' data and information are secure and protected.

5. Do you have protocols built into your third party developer platform to monitor data usage or access?

As we noted in Question 2, prior to providing access to our data, we carefully review and conduct risk assessments of all prospective Enterprise partners. Before they are granted access—and as part of the application process—each Enterprise partner is categorized based on a number of factors, including the customer's industry, country of origin, proposed end users and use case, data products to which the customer would have access, as well as the product's complexity. We monitor our Enterprise partners' use of Twitter user data according to their respective categorization.

Based on that categorization process, the majority of our Enterprise partners do not generally warrant escalated review or monitoring. Even for those customers who are not escalated, however, we nevertheless review the proposed uses and business purpose at the application stage and prior to renewal of their one-year contractual agreement with Twitter. All other Enterprise partners are escalated for further review and are more frequently and substantially investigated for compliance with our policies.

We undertake a review of our escalated customers at the application stage (*i.e.*, prior to onboarding), and regularly review these customers pursuant to their license agreements at the three-month mark of their contract with Twitter, and three months prior to their contract renewal date. That review includes, but is not limited to, revisiting and reviewing the customer's records, evaluating the scope of the customer's data use, comparing that data use to the customer's contractual use case, and examining external sources of information about the customer. We also conduct ongoing monitoring of our escalated customers through news alerts and notifications from our Trust & Safety Team.

If we identify deviations from the contractual scope of the customer's use case, or if other issues arise over the course of our review, Twitter launches an investigation into that customer's data use. For any specific suspected violations, we contact the customer and notify them of the investigation, and we request additional relevant information.

Where our investigation has identified a violation of our rules or a breach of the customer's contract, we have a number of options available. Those options include updating the customer's use case to ensure all of the customer's uses are comprehensively documented in the customer's agreement with Twitter (so long as the new use remains within the scope of our policies and procedures), educating the customer about permissible and impermissible uses of the data, temporarily suspending access to Twitter content, and permanently terminating the account.

Primarily through automated detection tools we have implemented, we also monitor data uses through our public APIs. As is the case with respect to spam and malicious activity on our platform, we are continuing our efforts to detect and prevent abuses of our public APIs. And we do so by leveraging our technological capabilities and improving our machine learning models to better detect and stop malicious actors. Our efforts have proved successful. Deploying those tools, since June 2017, we suspended more than 400,000 malicious applications for API abuse.

Our Trust & Safety and Platform Operations Teams also maintain a Proactive App Review program to regularly audit and review the activity of applications using our public APIs. Where warranted, we take enforcement action against applications found to be in violation of the Developer Agreement and Policy or the Twitter Rules. Since their introduction last year, those Teams' efforts have resulted in the suspension of more than 5,000 applications, many of which were large contributors and generators of spam and abuse on our platform. We continue to invest in both human and technical resources to increase the effectiveness and scale of this work.

With respect to our Ads API, we have established a tiered access-granting process in order to ensure that Ads API developers build their applications in compliance with the Ads API License Agreement. That process involves both automated and human review components. We initially grant third-party developers access to Twitter Ads API data on a very limited basis, which allows us to circumscribe the amount of data those developers can use while testing their application. Once the developers have completed their application integration process, a member of the Twitter API Partnerships Team will review a demonstration of the application with the developer. If we determine that the application complies with the Ads API License Agreement, we will increase the scope of the developer's access to Twitter Ads API data, which will allow the developer to pilot the product and generate case studies to establish the application's business value and performance. A Twitter employee will then review those case studies. If they are found to be within the scope of permitted uses and in compliance with our rules and policies, Twitter will lift the restrictions and allow the developer greater access to Twitter Ads API data.

6. **What audit procedures do you have to ensure compliance with third party developer agreements?**

   a. **How often have these audits been carried out?**
   b. **How compliant have third party developers been?**

The answer to Question 6 has been provided in response to Question 4 and Question 5.

7. **Have third party developers breached your terms/agreements in the past, and what remedies have you taken?**

The answer to Question 7 has been provided in response to Question 4.

8. What are your current policies and procedures with respect to notifying users of a data breach?

a. How have these policies and procedures evolved/changed since 2010?

b. Do you intend to make any changes in light of recent events?

In the event of a data-breach incident, we evaluate two key questions in determining whether notifying our users is necessary or warranted. First, we assess whether there is an applicable legal obligation to notify the users who were impacted by the breach. If such an obligation exists, we proceed to develop a user notice protocol based upon applicable law and best practices. Second, even if there is no legal requirement to notify, we evaluate whether issuing a notice to impacted accounts would advance transparency and greater security awareness or understanding on the part of the impacted users. We also consider whether those interests would be advanced by a user education effort beyond directly impacted users.

Our practices reflect longstanding Twitter positions on user transparency and choice. In the past, Twitter has notified users of security incidents even in circumstances where no data breach occurred and where no legal obligation obligated us to provide such notice. We have done so based on the analysis described above.

9. What are your current policies and procedures with respect to notifying users of an improper transfer of their data?

a. How have these policies and procedures evolved/changed since 2010?

b. Do you intend to make any changes in light of recent events?

In the event of an improper transfer of private data, Twitter follows the policies described in response to Question 8. For an improper transfer of data that a user has previously chosen to make public, such as a public Tweet, Twitter will take appropriate action based on a variety of factors. For example, as described in response to Question 2 above, Twitter regularly terminates API access for developers that violate Twitter's policies or otherwise engage in actions that raise privacy and/or security concerns.

10. Do you restrict the ability of third party developers to access data for a political purpose?

As noted above, developers seeking access to Twitter user data are prohibited from using that data (or sharing the data with any entity) to target, segment, or profile individuals based on their political affiliation or beliefs. Applications for access to our Enterprise data products that state or suggest that they would use Twitter user data for those purposes are denied access to Twitter content via our APIs. While we may allow research institutions to access user data in order to assess the impact of a political movement or politically affiliated group, those use cases are carefully scrutinized and are narrowly tailored to the institution's stated purpose, and are still bound by our policies prohibiting tracking and monitoring users.

We also require our Ads API developers to comply with our Twitter Ads Policy regarding political campaigning, which requires political advertisers to comply with applicable laws regarding disclosure and ads content, eligibility restrictions, and blackout dates for the countries in which they advertise. *See* https://business.twitter.com/en/help/ads-policies/restricted-content-policies/political-campaigning.html.

11. What engagement do you have with political campaigns and do you have policies and procedures governing these engagements?

a. How have these policies and procedures evolved/changed since 2010?
b. Do you intend to make any changes in light of recent events?

During election periods, Twitter staff often meets with campaign staff in order to educate them about how to use the platform and obtain the best value for the campaign's advertising spend. This is a service we offer to all of our large advertisers; it is not unique to any campaign or to election campaigns. We have separate teams that exclusively service a particular political party, and we prohibit any employee from sharing confidential information of one advertiser with another. Political advertisers are all subject to our standard ads policies, and we ensure that there is parity in our advertising offerings and services to all candidates.

12. How do you monitor the ability of foreign entities to access user data? What restrictions are in place to limit such access?

Although the restrictions we place on third parties' access to our users' data apply to foreign and domestic entities alike, we apply a heightened degree of scrutiny to applications from countries that Twitter considers to be high-risk jurisdictions based on independent, external, and non-governmental indicators. Countries that place significant restrictions on freedom of speech and freedom of the press, or which exhibit high degrees of public corruption, are generally categorized as high risk. While Twitter does not categorically prohibit the use of data in certain markets (except as prohibited by law), we balance open, global access to our platform against potential risks in certain jurisdictions where third parties would seek access to our user data for surveillance purposes (*i.e.*, to monitor individuals' conduct and speech about controversial issues). Accordingly, pursuant to our standard review procedures, applications from those jurisdictions are escalated for further review by our Data Compliance Team and may be subject to additional restrictions (or may not be permitted to license data at all).

13. How do you monitor the ability of foreign entities to influence and interfere with U.S. elections?

As we noted in our January 19, 2018 update to Congress, our efforts to detect and stop malicious activity on our platform continue, particularly in the context of elections. Based on the understanding we have gained from our retrospective review of activity on our platform during the period leading up to the 2016 election, we have established an internal, cross-functional team dedicated to addressing election-related instances of abuse on Twitter.

The election team addresses this challenge in a number of ways. Among other things, to detect and promptly address impersonation attempts, the team will verify accounts of major party candidates for all statewide and federal offices, as well as all major national party accounts. In addition to monitoring and enforcing the Twitter Terms of Service and Twitter Rules, the election team will cooperate and communicate with federal and state election officials to swiftly escalate and address in real time attempts at election interference. And consistent with Twitter's commitment to curtailing malicious automation, spam, and false accounts on our platform, the election team will focus on deploying our proprietary tools specifically to detect and stop malicious election-related activity.

We were pleased to learn that the Federal Bureau of Investigation has launched a Task Force to assist companies and the public in identifying foreign manipulation efforts through social media platforms. We believe that federal law enforcement agencies are uniquely positioned to access, synthesize, and comprehend disparate sources of intelligence, and to alert the public, Congress, and social media companies of their findings in a way that provides a broader picture of the activity. The speed and nature of dissemination of information across the Internet—and across platforms—present substantial, novel, and unique challenges. As we continue improving our ability to recognize and tackle these critical challenges, we appreciate and welcome the input and guidance of outside entities such as the FBI. We will continue to work closely with other companies, civil society, experts, and law enforcement to consider these challenges for Twitter, the Internet, and society as a whole in light of the larger ecosystem of how information spreads.

14. Are you aware of any foreign entities seeking to influence or interfere with U.S. elections through your platforms?

As we noted in response in Question 13, our efforts to identify and halt attempts to interfere with the electoral process in any capacity (whether foreign or domestic) continue. We are taking all available steps to identify and stop any such activity on the platform. We also appreciate the efforts of the Foreign Interference Task Force work undertaken by the FBI, and, as we have done since the 2016 election, we will continue to work with law enforcement concerning any suspected election interference activity identified on our platform.