United States House of Representatives
Committee on Oversight and Reform
Subcommittee on National Security
"Securing U.S. Election Infrastructure and Protecting Political Discourse"
May 22, 2019

Testimony of Kevin Kane
Twitter, Inc.

Chairman Lynch, Ranking Member Hice, and Members of the Subcommittee:

Thank you for the opportunity to appear before you today.

The purpose of Twitter is to serve the public conversation. We serve our global audience by focusing on the needs of the people who use our service, and we put them first in every step we take. People from around the world come to Twitter to engage in an open and free exchange of ideas. We must be a trusted and healthy place that supports free and open democratic debate.

The public conversation occurring on Twitter is never more important than during elections, the cornerstone of democracies across the globe. Our service shows the world what is happening, democratizes access to information and—at its best—provides people insights into a diversity of perspectives on critical issues; all in real time.

Twitter continues to demonstrate a strong commitment to transparency regarding our election integrity efforts. We conducted a comprehensive retrospective review of the activity on the service related to the 2018 U.S. midterm elections and released the findings in full. This report includes information about the lessons we learned from the 2016 elections, the improvements to the platform that we continue to implement, our engagement with key stakeholders, and the activity we saw on Election Day, including examples of the nearly 6,000 Tweets we removed as attempted voter suppression information, primarily of domestic origin, and information regarding potential service manipulation activity by foreign actors. I have attached the full retrospective review to my testimony and it can be found electronically at: https://blog.twitter.com/content/dam/blog-twitter/official/en_us/company/2019/2018-retrospective-review.pdf

IMPROVEMENTS TO TWITTER RELATED TO GLOBAL ELECTIONS

Twitter is committed to improving the collective health, openness, and civility of public conversation on our service. Twitter's health is built and measured by how we help encourage more healthy debate, conversations, and critical thinking. Conversely, abuse, malicious automation, and platform manipulation detract from it. Since January 2017, we have launched numerous product and policy changes, expanded our enforcement and operations, and strengthened our team structure, all designed to foster the health of the service and protect the people who use Twitter from abuse and malicious automation.

A. Combating Malicious Automation and Protecting Conversation Health

Using the insights from our retrospective review of recent U.S. and global elections, Twitter continues to develop the detection tools and systems needed to combat malicious automation on our service. For example, Twitter prioritizes identifying suspicious account activity, such as exceptionally high-volume Tweeting with the same hashtag or mentioning the same @handle without a reply from the account being addressed; when we identify such activity, we may require an individual using the service to confirm human control of the account or their identity. Twitter has also increased its use of challenges intended to catch automated accounts, such as reCAPTCHAs, that require individuals to identify portions of an image or type in words displayed on screen, and password reset requests that protect potentially compromised accounts. In 2018, we challenged 426,040,548 account registrations in total. Twitter has also implemented mandatory email or phone verification for all new accounts.

Our motivation with these changes is to reduce the burden on people who use Twitter to report spam and malicious automation. These technological improvements have brought about a corresponding reduction in the number of spam reports from people on Twitter. Reports of spam are reports from people who use Twitter after receiving an interaction, such as a follow, mention, or Direct Message, from a suspected spam account.

We also removed locked accounts from follower counts to ensure these figures more accurately reflect the actual reach and popularity of an account. Accounts are locked when our systems detect unusual activity and force a password change or other challenge. If the challenge has not been met or the password has not been changed, the account is immediately locked, barring it from sending Tweets, Retweets or liking posts from others. Removing locked accounts from follower counts helps ensure that people can trust the reliability of the information on profiles on Twitter.

B. Cross-Functional Team

Our improvements in preparation for the 2018 U.S. midterm elections and elections across the globe included important structural changes. In particular, Twitter created an internal cross-functional analytical team whose mission is to monitor site and service integrity. Drawing on expertise across the company, this team can respond rapidly to escalations of inauthentic, malicious automated or human-coordinated activity on the service. The team's work enables us to better understand the nature of malicious activity and mitigate it more quickly.

The primary focus of the cross-functional analytical team is election readiness. The team examines, responds to, and escalates instances of suspected inauthentic, election-related coordinated activity in political conversation and conduct in-depth analyses of relevant Twitter data.

To supplement its own analyses, Twitter's analytical team also receives and responds to reports from across the company and from external third parties. The results from all of the team's analyses are shared with key stakeholders at Twitter and provide the basis for policy changes, product initiatives, and the removal of accounts.

Through real-time review and detection of anomalous and potentially malicious automated and human-coordinated activity, the team works to identify and address any attempts by bad faith actors to interfere with the electoral process and is better informed about where and how to deploy resources to proactively review potential malicious activity. Accounts are escalated for review in real time if exhibiting anomalous patterns of behavior. These efforts significantly improve our ability to detect malicious automated and human-coordinated activity surrounding political content as well as the speed with which we address those issues.

C. Advertising and Promoted Content

As we reported in our retrospective review of the 2018 U.S. midterm elections, we have devoted considerable resources to increasing transparency and promoting accountability in the ads served to Twitter customers. Twitter implemented an updated Political Campaigning Policy in June 2018 to provide clearer guidance about how we define political content and who can promote political content on our service. Under the revised policy, advertisers who wish to target the United States with federal political campaigning advertisements are required to self-identify as such and certify that they are located within the United States. Foreign nationals are not permitted to serve political ads to individuals who identify as being located in the United States.

Twitter account holders who wish to target the U.S. with federal political campaigning advertisements must also comply with a strict set of requirements. Among other things, the account's profile photo, header photo, and website must be identical to the individual's or organization's online presence. In addition, the advertiser must take steps to verify that the address used to serve advertisements with content related to a federal political campaign is genuine.

To further increase transparency and better educate those who access promoted content, accounts serving promoted Tweets with content related to a federal political campaign is now visually identified and contains a disclaimer. This feature allows people to more easily identify federal political campaign advertisements, quickly identify the identity of the account funding the advertisement, and immediately tell whether it was authorized by the candidate.

In June 2018, we launched the Ads Transparency Center, which is open to everyone on Twitter and the general public, and which includes comprehensive data on paid electioneering communications on the service within the United States. Twitter requires extensive information disclosures of any account involved in federal electioneering communications and provides specific information to the public via the Ads Transparency Center, including:

- Purchases made by a specific account;
- All past and current ads served on the service for a specific account;
- Targeting criteria and results for each advertisement;
- Number of views each advertisement received; and
- Certain billing information associated with the account.

These are meaningful steps that have enhanced the Twitter experience and protected the health of political conversations on the service.

We also implemented the next phase of our efforts to provide transparency with the launch of a U.S.-specific Issue Ads Policy and certification process in September 2018. In addition to the policy governing advertisers running campaigns for federal elections described above, Twitter implemented a new disclosure policy for advertisers promoting content about candidates running for federal, state, or local election, as well as those discussing issues of legislative national importance. To provide people with additional information about individuals

or organizations promoting issue ads, Twitter has established a process that verifies an advertiser's identity and location within the United States.

These advertisements will also be included in the Ads Transparency Center. We are also examining how to adopt political campaigning and issue ads policies around the world. We remain committed to continuing to improve and invest resources in this space.

In 2018, 96 political advertisers spent nearly $2.3 million to purchase 2,267 ads that resulted in nearly 170 million impressions. Additionally, 150 issue advertisers spent $2.2 million on 1,373 ads that resulted in approximately 198 million impressions.

D. Recent Updates to Twitter Rules

We have made a number of recent updates to the rules governing the use of our service to better protect the conversation around elections. On October 1, 2018, Twitter announced an update the Twitter Rules to provide clearer guidance around several key issues impacting the integrity of elections across the globe. As service manipulation tactics continue to evolve, we are updating and expanding our rules to better reflect how we identify fake accounts, and what types of inauthentic activity violate our guidelines. Some of the factors that we take into account when determining whether an account is fake include the use of stock or stolen avatar photos; the use of stolen or copied profile bios; and the use of intentionally misleading profile information, including profile location.

We also updated the Twitter Rules regarding attributed activity. Now, if we are able to reliably attribute an account on Twitter to an entity known to violate the Twitter Rules, we will remove additional accounts associated with that entity. We are expanding our enforcement approach to include accounts that deliberately mimic or are intended to replace accounts we have previously suspended for violating our rules. These steps allow us to take more aggressive action against known malicious actors, such as the Russian Internet Research Agency.

Additionally, we announced an update regarding hacked materials. Twitter rules prohibit the distribution of hacked material that contains private information or trade secrets or could put people in harm's way. According to the Twitter Rules, Twitter does not permit the use of our services to directly distribute content obtained through hacking that contains personally identifiable information, may put people in imminent harm or danger, or contains trade secrets. Direct distribution of hacked materials includes posting hacked content on Twitter (for instance, in the text of a Tweet or in an image), or directly linking to hacked content hosted on other websites.

We also expanded the criteria for when we will take action on accounts which claim responsibility for a hack, which includes threats and public incentives to hack specific people and accounts. We also may suspend accounts in which Twitter is able to reliably attribute a hack to the account distributing that content. Commentary about a hack or hacked materials, such as news articles discussing a hack, are generally not considered a violation of this policy. This includes, for example, journalistic and editorial discussion of a hacking and disclosures of legitimate public concern and which pose no physical harm.

In the 2018 U.S. midterm elections, the vast majority of violative content we removed from our service on Election Day was voter suppressive content. In general, Twitter looked for behavior that attempted to influence an election by deterring groups of eligible voters, particularly through voter intimidation or providing false information about voting or registering to vote. We removed nearly 6,000 Tweets we identified as attempted voter suppression, primarily of domestic origin.

Individuals using Twitter may not use Twitter for the purpose of manipulating or interfering in elections. This includes but is not limited to misleading information about how to vote or register to vote; misleading information about requirements for voting, including identification requirements; and misleading statements or information about the official, announced date or time of an election.

In April 2019, we issued a new policy regarding election integrity governing three categories of manipulative behavior and content related to elections. First, an individual cannot share false or misleading information about how to participate in an election or other civic event. This includes but is not limited to misleading information about how to vote or register to vote, requirements for voting, including identification requirements, and the official, announced date, or time of an election. Second, an individual cannot share false or misleading information intended to intimidate or dissuade voters from participating in an election. This includes but is not limited to misleading claims that polling places are closed, that polling has ended, or other misleading information relating to votes not being counted.

We also do not allow misleading claims about police or law enforcement activity related to polling places or elections, long lines, equipment problems, voting procedures or techniques which could dissuade voters from participating in an election, and threats regarding voting locations. Finally, we also do not allow the creation of fake accounts which misrepresent their affiliation, or share content that falsely represents its affiliation to a candidate, elected official, political party, electoral authority, or government entity.

We further expanded our enforcement capabilities for global elections by creating a dedicated reporting feature within the product to allow users to more easily report content that undermines the process of registering to vote or engaging in the electoral process. This reporting feature is being used first with 2019 Lok Sabha elections in India and the elections in the European Union.

E. Additional Safety Measures for Accessing Twitter's Application Programming Interfaces

To further address malicious automation and abuse on the service, we continue to update our developer policies and processes, which govern the access and use of public Tweet data made available to developers and other third parties through our application programming interfaces ("APIs").

We recognize that programmatic access to the Twitter service, including access to public Tweet data, could be manipulated, so we have taken steps to prevent the use of our APIs for products and services that are abusive or that disrupt the health of conversations. Applications to which we grant access to our APIs are prohibited from using the data to manipulate conversations or otherwise disrupt the integrity of the Twitter service or invade the privacy of people on Twitter. Between July and December 2018 alone we removed more than 162,000 applications that we determined to be in violation of our developer policies. Most violated our policies against producing spam via APIs. And we continue to invest in and improve our detection tools to stop misuse of public Twitter data.

In July 2018, we introduced a new measure designed to increase developers' accountability for applications that create and engage with Twitter content and accounts. Twitter now reviews and conducts compliance checks of all developers' stated use of the data that they wish to access. We have also added new protections aimed to prevent the registration of low quality and spam-generating applications. We believe that these additional steps will help protect the integrity of our service.

F. Engagement with Government Entities

Information sharing and collaboration are critical to Twitter's success in preventing hostile foreign actors from disrupting meaningful political conversations on the service. We have well-established relationships with law enforcement agencies active in this arena, including the Federal Bureau of Investigation Foreign Influence Task Force and the U.S. Department of Homeland Security's Election Security Task Force. We look forward to continued cooperation

with federal, state, and local government agencies on election integrity issues because in certain circumstances only they have access to information critical to our joint efforts to stop bad faith actors.

On Election Day in the 2018 U.S. midterms, Twitter virtually participated in an operations center convened by the U.S. Department of Homeland Security. The operations center also convened officials from the U.S. Department of Justice, the Federal Bureau of Investigation, and the Office of the Director of National Intelligence, in addition to federal, state, local, and private sector partners. In the lead up to Election Day, and throughout the course of the day itself, Twitter remained in constant contact with officials throughout all levels of government.

We also worked in close collaboration with the National Association of Secretaries of State (NASS) and the National Association of State Election Directors (NASED). Founded in 1904, NASS is the nation's oldest, nonpartisan professional organization for public officials, and is open to secretaries of states and lieutenant governors in the 50 states, D.C. and territories. In February, Twitter participated in a panel discussion convened by NASS on the Role of Social Media in Democracy and their New Voters Forum, broadcast on C-Span.

* * *

The people who use Twitter must have confidence in the integrity of the information found on the service, especially with respect to information relevant to elections and the democratic process. We continue to invest in our efforts to address those threats posed by hostile actors and foster an environment conducive to healthy, meaningful conversations on our service. We look forward to working with the Subcommittee on this important issue.