HEARING BEFORE THE UNITED STATES SENATE SELECT COMMITTEE ON INTELLIGENCE

November 1, 2017

Testimony of Colin Stretch
General Counsel, Facebook

I. INTRODUCTION

Chairman Burr, Vice Chairman Warner, and distinguished members of the Committee, thank you for this opportunity to appear before you today. My name is Colin Stretch, and since July 2013, I've served as the General Counsel of Facebook. We appreciate this Committee's hard work to investigate Russian interference in the 2016 election.

At Facebook, our mission is to create technology that gives people the power to build community and bring the world closer together. We don't take for granted that each one of you uses Facebook to connect with your constituents, and that the people you represent expect authentic experiences when they come to our platform to share.

We also believe we have an important role to play in the democratic process—and a responsibility to protect it on our platform. That's why we take what's happened on Facebook so seriously. The foreign interference we saw is reprehensible and outrageous and opened a new battleground for our company, our industry, and our society. That foreign actors, hiding behind fake accounts, abused our platform and other internet services to try to sow division and discord—and to try to undermine our election process—is an assault on democracy, and it violates all of our values.

In our investigation, which continues to this day, we've found that these actors used fake accounts to place ads on Facebook and Instagram that reached millions of Americans over a two-year period, and that those ads were used to promote Pages, which in turn posted more content. People shared these posts, spreading them further. Many of these ads and posts are inflammatory. Some are downright offensive.

In aggregate, these ads and posts were a very small fraction of the overall content on Facebook—but any amount is too much. All of these accounts and Pages violated our policies, and we removed them.

Going forward, we're making some very significant investments—we're hiring more ad reviewers, doubling or more our security engineering efforts, putting in place tighter ad content restrictions, launching new tools to improve ad transparency, and requiring documentation from political ad buyers. We're building artificial intelligence to help locate more banned content, and bad actors. We're working more closely with industry to share information on how to identify and prevent threats so that we can all respond faster and more effectively. And we are expanding our efforts to work more closely with law enforcement.

I'm here today to share with you what we know so far about what happened—and what we're doing about it. At the outset, let me explain how our service works and why people choose to use it.

II. FIGHTING ELECTION INTERFERENCE ON FACEBOOK

A. Understanding what you see on Facebook

1. The News Feed Experience: A Personalized Collection of Stories. When people come to Facebook to share with their friends and discover new things, they see a personalized homepage we call News Feed. News Feed is a constantly updating, highly personalized list of stories, including status updates, photos, videos, links, and activity from the people and things you're connected to on Facebook. The goal of News Feed is to show people the stories that are most relevant to them. The average person has thousands of things on any given day that they could read in their News Feed, so we use personalized ranking to determine the order of stories we show them. Each person's News Feed is unique. It's shaped by the friends they add; the people, topics, and news sources they follow; the groups they join; and other signals like their past interactions. On average, a person in the US is served roughly 220 stories in News Feed each day. Over the time period in question, from 2015 to 2017, Americans using Facebook were exposed to, or "served," a total of over 33 trillion stories in their News Feeds.

2. Advertising and Pages as Sources of Stories in News Feed. News Feed is also a place where people see ads on Facebook. To advertise in News Feed, a person must first set up a Facebook account—using their real identity—and then create a Facebook Page. Facebook Pages represent a wide range of people, places, and things, including causes, that people are interested in. Any user may create a Page to express support for or interest in a topic, but only official representatives can create a Page on behalf of an organization, business, brand, or public figure. It is against our terms for Pages to contain false, misleading, fraudulent, or deceptive claims or content. Facebook marks some official Pages—such as for a public figure, media company, or brand—with a "verified" badge to let people know they're authentic. All Pages must comply with our Community Standards and ensure that all the stories they post or share respect our policies prohibiting hate speech, violence, and sexual content, among other restrictions. People can like or follow a Page to get updates, such as posts, photos, or videos, in their News Feed. The average person in the US likes 178 Pages. People do not necessarily see every update from each of the Pages they are connected to. Our News Feed ranking determines how relevant we think a story from a Page will be to each person. We make it easy for people to override our recommendations by giving them additional controls over whether they see a Page's updates higher in their News Feed or not at all. For context, from 2015 to 2017, people in the United States saw 11.1 trillion posts from Pages on Facebook.

3. Advertising to Promote Pages. Page administrators can create ads to promote their Page and show their posts to more people. The vast majority of our advertisers are small- and medium-sized businesses that use our self-service tools to create ads to reach their customers. Advertisers choose the audience they want to reach based on demographics, interests, behaviors or contact information. They can choose from different ad formats, upload images or video, and write the text they want people to see. Advertisers can serve ads on our platform for as little as $0.50 per day using a credit card or other payment method. By using these tools, advertisers agree to our

Self-Serve Ad Terms. Before ads appear on Facebook or Instagram, they go through our ad review process that includes automated checks of an ad's images, text, targeting and positioning, in addition to the content on the ad's landing page. People on Facebook can also report ads, find more information about why they are being shown a particular ad, and update their ad preferences to influence the type of ads they see.

B. Promoting Authentic Conversation

Our authenticity policy is the cornerstone of how we prevent abuse on our platform, and was the basis of our internal investigation and what we found.

From the beginning, we have always believed that Facebook is a place for authentic dialogue, and that the best way to ensure authenticity is to require people to use the names they are known by. Fake accounts undermine this objective, and are closely related to the creation and spread of inauthentic communication such as spam—as well as used to carry out disinformation campaigns like the one associated with the Internet Research Agency (IRA).

We build and update technical systems every day to better identify and remove inauthentic accounts, which also helps reduce the distribution of material that can be spread by accounts that violate our policies. Each day, we block millions of fake accounts at registration. Our systems examine thousands of account attributes and focus on detecting behaviors that are very difficult for bad actors to fake, including their connections to others on our platform. By constantly improving our techniques, we also aim to reduce the incentives for bad actors who rely on distribution to make their efforts worthwhile.

Protecting authenticity is an ongoing challenge. As our tools and security efforts evolve, so will the techniques of those who want to evade our authenticity requirements. As in other areas of cybersecurity, our security and operations teams need to continually adapt.

C. Protecting the Security of the 2016 Election and Learning Lessons Quickly

1. The Evolution of Facebook's Security Protections. From its earliest days, Facebook has always been focused on security. These efforts are continuous and involve regular contact with law enforcement authorities in the United States and around the world. Elections are particularly sensitive events for our security operations, and as the role our service plays in promoting political dialogue and debate has grown, so has the attention of our security team.

As your investigation has revealed, our country now faces a new type of national cyber-security threat—one that will require a new level of investment and cooperation across our society. At Facebook, we're prepared to do our part. At each step of this process, we have spoken out about threats to internet platforms, shared our findings, and provided information to investigators. As we learn more, we will continue to identify and implement improvements to our security systems, and work more closely with other technology companies to share information on how to identify and prevent threats and how to respond faster and more effectively.

2. Security Leading Up to the 2016 Election.

a. Fighting Hacking and Malware. For years, we had been aware of other types of activity that

appeared to come from Russian sources—largely traditional security threats such as attacking people's accounts or using social media platforms to spread stolen information. What we saw early in the 2016 campaign cycle followed this pattern. Our security team that focuses on threat intelligence—which investigates advanced security threats as part of our overall information security organization—was, from the outset, alert to the possibility of Russian activity. In several instances before November 8, 2016, this team detected and mitigated threats from actors with ties to Russia and reported them to US law enforcement officials. This included activity from a cluster of accounts we had assessed to belong to a group ("APT28") that the US government has publicly linked to Russian military intelligence services. This activity, which was aimed at employees of major US political parties, fell into the normal categories of offensive cyber activities we monitor for. We warned the targets who were at highest risk, and were later in contact with law enforcement authorities about this activity.

Later in the summer we also started to see a new kind of behavior from APT28-related accounts—namely, the creation of fake personas that were then used to seed stolen information to journalists. These fake personas were organized under the banner of an organization that called itself DC Leaks. This activity violated our policies, and we removed the DC Leaks accounts.

b. Understanding Fake Accounts and Fake News. After the election, when the public discussion of "fake news" rapidly accelerated, we continued to investigate and learn more about the new threat of using fake accounts to amplify divisive material and deceptively influence civic discourse. We shared what we learned with government officials and others in the tech industry. And in April 2017, we shared our findings with the public by publishing a white paper that described the activity we detected and the initial techniques we used to combat it.

As with all security threats, we have also been applying what we learned in order to do better in the future. We use a variety of technologies and techniques to detect and shut down fake accounts, and in October 2016, for example, we disabled about 5.8 million fake accounts in the United States. At the time, our automated tooling did not yet reflect our knowledge of fake accounts focused on social or political issues. But we incorporated what we learned from the 2016 elections into our detection systems, and as a result of these improvements, we disabled more than 30,000 accounts in advance of the French election. This same technology helped us disable tens of thousands more accounts before the German elections in September. In other words, we believe that we're already doing better at detecting these forms of abuse, although we know that people who want to abuse our platform will get better too and so we must stay vigilant.

3. Investigating the Role of Ads and Foreign Interference. After the 2016 election, we learned from press accounts and statements by congressional leaders that Russian actors might have tried to interfere in the election by exploiting Facebook's ad tools. This is not something we had seen before, and so we started an investigation that continues to this day. We found that fake accounts associated with the IRA spent approximately $100,000 on more than 3,000 Facebook and Instagram ads between June 2015 and August 2017. Our analysis also showed that these accounts used these ads to promote the roughly 120 Facebook Pages they had set up, which in turn posted more than 80,000 pieces of content between January 2015 and August 2017. The Facebook accounts that appeared tied to the IRA violated our policies because they came from a

set of coordinated, inauthentic accounts. We shut these accounts down and began trying to understand how they misused our platform.

a. Advertising by Accounts Associated with the IRA. Below is an overview of what we've learned so far about the IRA's ads:

- **Impressions (an "impression" is how we count the number of times something is on screen, for example this can be the number of times something was on screen in a person's News Feed):**
  - 44% of total ad impressions were before the US election on November 8, 2016.
  - 56% of total ad impressions were after the election.
- **Reach (the number of people who saw a story at least once):**
  - We estimate 11.4 million people in the US saw at least one of these ads between 2015 and 2017.
- **Ads with zero impressions:**
  - Roughly 25% of the ads were never shown to anyone. That's because advertising auctions are designed so that ads reach people based on relevance, and certain ads may not reach anyone as a result.
- **Amount spent on ads:**
  - For 50% of the ads, less than $3 was spent.
  - For 99% of the ads, less than $1,000 was spent.
  - Many of the ads were paid for in Russian currency, though currency alone is a weak signal for suspicious activity.
- **Content of ads:**
  - Most of the ads appear to focus on divisive social and political messages across the ideological spectrum, touching on topics from LGBT matters to race issues to immigration to gun rights.
  - A number of the ads encourage people to follow Pages on these issues, which in turn produced posts on similarly charged subjects.

b. Content Posted by Pages Associated with the IRA. We estimate that roughly 29 million people were served content in their News Feeds directly from the IRA's 80,000 posts over the two years. Posts from these Pages were also shared, liked, and followed by people on Facebook, and, as a result, three times more people may have been exposed to a story that originated from the Russian operation. Our best estimate is that approximately 126 million people may have been served content from a Page associated with the IRA at some point during the two-year period. This equals about four-thousandths of one percent (0.004%) of content in News Feed, or approximately 1 out of 23,000 pieces of content.

Though the volume of these posts was a tiny fraction of the overall content on Facebook, **any amount is too much**. Those accounts and Pages violated Facebook's policies—which is why we

removed them, as we do with all fake or malicious activity we find. We also deleted roughly 170 Instagram accounts that posted about 120,000 pieces of content.

Our review of this activity is ongoing. Many of the ads and posts we've seen so far are deeply disturbing—seemingly intended to amplify societal divisions and pit groups of people against each other. They would be controversial even if they came from authentic accounts in the United States. But coming from foreign actors using fake accounts they are simply unacceptable.

That's why we've given the ads and posts to Congress—because we want to do our part to help investigators gain a deeper understanding of foreign efforts to interfere in the US political system and explain those activities to the public. These actions run counter to Facebook's mission of building community and everything we stand for. And we are determined to do everything we can to address this new threat.

D. Mobilizing to Address the New Threat

We are taking steps to enhance trust in the authenticity of activity on our platform, including increasing ads transparency, implementing a more robust ads review process, imposing tighter content restrictions, and exploring how to add additional authenticity safeguards.

1. Promoting Authenticity and Preventing Fake Accounts. We maintain a calendar of upcoming elections and use internal and external resources to best predict the threat level to each. We take preventative measures based on our information, including working with election officials where appropriate. Within this framework, we set up direct communication channels to escalate issues quickly. These efforts complement our civic engagement work, which includes voter education. In October 2017, for example, we launched a Canadian Election Integrity Initiative to help candidates guard against hackers and help educate voters on how to spot false news.

Going forward, we're also requiring political advertisers to provide more documentation to verify their identities and disclose when they're running election ads. Potential advertisers will have to confirm the business or organization they represent before they can buy ads. Their accounts and their ads will be marked as political, and they will have to show details, including who paid for the ads. We'll start doing this with federal elections in the US and then move onto other elections in the US and other countries**.** For political advertisers that don't proactively identify themselves, we're building machine learning tools that will help us find them and require them to verify their identity.

Authenticity is important for Pages as well as ads. We'll soon test ways for people to verify that the people and organizations behind political and issue-based Pages are who they say they are.

2. Partnering with Industry on Standards. We have been working with many others in the technology industry, including with Google and Twitter, on a range of elements related to this investigation. Our companies have a long history of working together on other issues such as child safety and counter-terrorism.

We are also reaching out to leaders in our industry and governments around the world to share information on bad actors and threats so that we can make sure they stay off all platforms. We

are trying to make this an industry standard practice.

3. Strengthening Our Advertising Policies. We know that some of you and other members of Congress are exploring new legislative approaches to political advertising—and that's a conversation we welcome. We are already working with some of you on how best to put new requirements into law. But we aren't waiting for legislation. Instead we're taking steps where we can on our own, to improve our own approach to transparency, ad review, and authenticity requirements.

a. Providing Transparency. We believe that when you see an ad, you should know who ran it to be able to understand what other ads they're running—which is why we show you the Page name for any ads that run in your News Feed.

To provide even greater transparency for people and accountability for advertisers, we're now building new tools that will allow you to see the other ads a Page is running as well—including ads that aren't targeted to you directly. We hope that this will establish a new standard for our industry in ad transparency. We try to catch material that shouldn't be on Facebook before it's even posted—but because this is not always possible, we also take action when people report ads that violate our policies. We're grateful to our community for this support, and hope that more transparency will mean more people can report violating ads.

b. Enforcing Our Policies. We rely on both automated and manual ad review, and we're now taking steps to strengthen both. Reviewing ads means assessing not just what's in an ad but also the context in which it was bought and the intended audience—so we're changing our ads review system to pay more attention to these signals. We're also adding more than 1,000 people to our global ads review teams over the next year and investing more in machine learning to better understand when to flag and take down ads. Enforcement is never perfect, but we will get better at finding and removing improper ads.

c. Restricting Ad Content. We hold people on Facebook to our Community Standards, and we hold advertisers to even stricter guidelines. Our ads policies already prohibit shocking content, direct threats and the promotion of the sale or use of weapons. Going forward, we are expanding these policies to prevent ads that use even more subtle expressions of violence.

III. CONCLUSION

Any attempt at deceptive interference using our platform is unacceptable, and runs counter to everything we are working toward. What happened in the 2016 election cycle was an affront to us, and, more importantly, to the people who come to Facebook every day to have authentic conversations and to share. We are committed to learning from these events, and to improving. We know we have a responsibility to do our part—and to do better. We look forward to working with everyone on this Committee, in the government, and across the tech industry and civil society, to address this important national security matter so that we can prevent similar abuse from happening again.