Written Testimony of Richard Salgado

Senior Counsel, Law Enforcement and Information Security, Google

Senate Judiciary Subcommittee on Crime and Terrorism

Hearing on "Extremist Content and Russian Information Online: Working with Tech to Find Solutions"

October 31, 2017

Chairman Graham, Ranking Member Whitehouse, distinguished Members of the Subcommittee: Thank you for inviting us to participate in today's hearing and for your leadership on these challenging and important issues.

My name is Richard Salgado. As the Director of Law Enforcement and Information Security at Google, I work with the thousands of people across teams at Google tasked with protecting the security of our networks and user data. Previously, I served as the Senior Counsel in the Computer Crimes and Intellectual Property Division at the Department of Justice, focusing on computer network cases involving hacking, illegal computer wiretaps, denial of service attacks, malicious code, and other computer crimes.

The spread of misinformation, extremist content, and hate speech online is a serious issue. Addressing these threats is a goal that Google shares with this Committee, and we are fully committed to working with government, civil society, and the private sector to do our part to meet this challenge.

Google's services can be used to provide real benefits to our users. We also recognize that detecting and preventing misuse of those services is critically important. We are deeply committed to working with Congress, law enforcement, others in our industry, and the NGO community to protect our users and products and to strengthen protections around elections and help combat disinformation.

I would like to take this opportunity today to walk through our systems, policies, and proactive efforts to combat the behavior of bad actors, including government-backed attackers, on our systems and against our users.

Protecting our platforms from government-backed interference is a challenge we began tackling as a company long before the 2016 presidential election. We've dedicated significant resources to help protect our platforms from such attacks by maintaining cutting edge defensive systems and by building advanced security tools directly into our consumer products. In 2007, we introduced our Safe Browsing tool, which helps protect our users from phishing, malware, or other attacks. Today, Safe Browsing is used on more than three billion devices worldwide. When we detect that a user's account has been targeted by a government-backed attacker, we show a warning that includes proactive steps the user can take to increase the security of the account. Earlier this month, we announced a new feature called the Advanced Protection Program. This program is designed to fit the needs of those who are at particularly high risk of targeted online attacks, such as employees of a political campaign, journalists, or prominent public figures. We update these tools constantly to try to stay ahead of evolving threats.

In addition to protecting our network security, we also have a range of tools to detect and prevent bad-actors from engaging in artificial amplification of content on our platforms. YouTube, for example, employs one of the most sophisticated spam and security breach detection systems available today, using an array of signals to catch bad actors as they try to artificially inflate the view counts of their videos or the number of subscribers to their channels. Bad actors attempting to spread misinformation online is a deeply concerning issue. We are constantly aiming to improve our products to surface relevant results to our users, and to empower our users to understand the information our products present.

On Google News, we mark up links with labels that help users understand what they are about to read, whether it is local content, an op-ed, or an in-depth piece, and encourage them to be thoughtful about the content they are looking at. Publishers who review third party claims or rumors can showcase their work on Google News through fact-check labels and in Google Search through fact-check cards. To help ensure Google does not monetize content designed to mislead users, we have implemented a new policy for our AdSense publishers that explicitly bans ads on any site that misrepresents who the publisher is or contains deceptive content. For Google Search, we updated our Search Quality Rater Guidelines and our evaluation test sets to help identify misleading information and unexpected offensive results, and have used this data to improve our search algorithms. This results in higher quality and more authoritative Search results.

With respect to the 2016 election, we have been looking across our products to understand whether individuals who appear to be connected to government-backed entities were disseminating information in the US for the purpose of interfering with the election. This was based on research conducted by Alphabet's Jigsaw group, the investigatory work of our information security team, and on leads provided by other companies.

We also included a broad review of all Ads from June 2015 until the election last November that were categorized as potentially "political" by our systems and had even the loosest connection to Russia, such as Russian IP address or billing address, or use of Russian currency. We suspected this broader look would ultimately include potentially legitimate activity and in fact, we have seen no evidence to show that any of these ads are connected to this effort.

On our advertising platforms, we ultimately found two accounts that appeared to be engaged in activity associated with known or suspected government-backed entities. The two accounts spent roughly $4,700 in connection with the 2016 US Presidential election.

We also focused on our other platforms. On YouTube, we found 18 channels with approximately 1,100 videos that were uploaded by individuals who we suspect are associated with this effort and that contained political content. These videos mostly had low view counts — just 3% had more than 5,000 views, and constituted only forty-three hours of YouTube content. While this is a relatively small amount of content — people watch over a billion hours of YouTube content a day, and 400 hours of content are uploaded every minute — we understand that any misuse of our platforms for this purpose is a serious challenge to the integrity of our democracy. The videos were not targeted to any particular segment of the US population as that is not a feature available on YouTube, but we did observe that links to those videos were frequently posted to other social media platforms. Similarly, we found a limited number of Gmail accounts that appear to have been used primarily to set up accounts on social media platforms.

We believe that the relatively limited amount of activity we found is a result of the safeguards that we had in place in advance of the 2016 election. Google's products also don't lend themselves to the kind of targeting or viral dissemination tools that these actors seem to prefer. But we are committed to continuing to improve upon our existing security measures to help prevent abuse.

As part of our commitment, we are making our political advertising more transparent, easier for users to understand, and even more secure. In 2018, we will release a transparency report for election ads and pair that with a library of election ad content that will be accessible to researchers. We're also going to make it easier for users to understand who bought the election ads they see on our networks. Going forward, users will be able to find the name of any advertiser running an election-related ad on Search, YouTube, or the Google Display Network with one click on an icon above the ad. And we will be enhancing the safeguards we already have in place to ensure users are in compliance with our ads policies and US laws covering election ads by verifying the identity of anyone who wants to run an election ad or use our political-interest-based tools and confirming that person is permitted to run that ad. We remain open to working with governments on legislation that promotes electoral transparency.

In addition to our own work, we will continue engaging in our robust partnerships with NGOs and our peers to collaborate on solutions in this space. Google has supported significant outreach to increase security for candidates and campaigns across the United States, France, Germany, and other countries. We have funded and advised on the National Cybersecurity Alliance's "Lock Down Your Login" education and training programs that focus specifically on the online security of campaigns and elected officials. We have provided significant advisory and financial support to the bi-partisan Defending Digital Democracy Project, a program of the Belfer Center for Science and International Affairs at Harvard Kennedy School. These types of initiatives and programs are critical to fostering a culture of security among internet users in the election space. We look forward to growing the current partnerships and pursuing new efforts.

The Committee has also expressed interest in hearing about our efforts to counter the propagation of extremist content online, and I welcome the opportunity to describe them. We have learned a lot over these past months and years on the challenges of addressing extremist content. We have been developing rigorous policies and programs to defend the use of our platforms against the desire by bad actors to spread hate or incite violence. YouTube, in particular, has long had policies that prohibit terrorist recruitment, violent extremism, incitement to violence, and instructional content that could be used to facilitate substantial bodily injury or death. When we become aware of content that violates these policies, we immediately remove it.

We use a mix of technology and humans to enforce our guidelines. Users can also alert us to content that they think may violate our policies through our flagging mechanism found beneath every YouTube video. We have teams charged with reviewing flagged content in multiple languages and countries around the world. But we also rely on our Trusted Flagger program, which provides a bulk-flagging tool and expedited review when content is flagged by NGOs and government agencies that have expertise on issues like hate speech and terrorism. And of course we rely upon our technology, which has always been such a critical part of our solution. Our matching technology, for example, is able to prevent the dissemination of bad content, by catching re-uploads of known bad content before that content is available to the public.

But we understand our goal, and we keep going further to enhance our systems. In June of this year, we announced four steps we are taking to further combat terrorism and hate speech on our platforms.

- **The first is an investment in technologies for detection and removal of extremist content**. Google has been working on machine learning for years, and we recently deployed new classifiers that detect terrorist material and flag it for removal. Over 83 percent of the videos we have removed for violent extremism over the past month were taken down using machine learning tools (and before receiving a single human flag), which is up 8 percentage points since August.

- **We are also focused on improving and expanding our expertise on these issues**. We committed to expanding our aforementioned Trusted Flagger Program to an additional 50 NGOs. We have onboarded over 35 and are actively consulting counter-terrorism experts to help inform our approach to this content. This work helps us to better understand how these issues play out on our services.

- **Thirdly, we are taking a tougher stance on videos that may be offensive but do not violate our policies.** On August 24, we launched a new approach. These videos remain on YouTube but they will not be recommended or monetized, and will not have key features like comments, suggested videos, or likes.

- **Finally, we are also creating programs to promote counterspeech on our platforms.** We are expanding our counter-extremism work to use advertising to present counternarratives. One such approach is through a tactic called the Redirect Method, which uses curated video content to redirect people away from violent extremist propaganda and steer them toward content that counters those narratives. We also launched a program called Creators for Change, which is a program for Youtube creators to use their voices and creativity to speak out against hate speech, xenophobia, and extremism.

These efforts are not just happening on YouTube. We have similar policies against terrorist and hate speech content on Drive, G+, and Blogger.

We also collaborate across the industry. Last year we created a coalition of tech companies that share hashes (or "digital fingerprints") of terrorist content to stop its spread. This summer, we announced the Global Industry Forum to Counter Terrorism to formalize industry collaboration on research, knowledge sharing, and technology.

Ultimately, violent extremism and misinformation that has the intent to manipulate our election system is an attack on open societies, and on the values we share. We acknowledge the threat and affirm that addressing it is a critical challenge for us all.

Google and YouTube are committed to doing our part, but as we all recognize across government, civil society, and the private sector, we will only make progress by working together to address these complex issues at their root. That is why forums like this are so important to underscoring our shared goals and commitments. We look forward to continuing to work with the Committee as it takes on this important issue.

Thank you for your time. I look forward to your questions.