FACEBOOK, INC.

Moderator: Tom Reynolds
August 21, 2018 4:30 p.m. PT

Mark Zuckerberg: Thanks everyone for joining today. One the most important responsibilities we have as a company is to keep people safe and to stop anyone from abusing our services.

Our security program has always focused on preventing traditional hacking and cybersecurity threats. But in 2016, in addition to identifying these kinds of threats, we also faced coordinated information operations with networks of fake accounts spreading divisive content and misinformation.

Since then, we've been investing heavily to improve safety, security and privacy -- and to defend against these coordinated and inauthentic campaign. This has been a lot of hard work, and while still early, we're starting to see it pay off and we're identifying more of it before the elections.

We've strengthened our ad policies to make advertising on our services much more transparent. Thanks to advances in artificial intelligence, we're now much more proactive in finding and removing bad content, as well as fake accounts. In the past few months, we've been able to proactively identify and remove inauthentic accounts, Pages, and Groups coming from countries including Russia, Mexico and Brazil. And we're working more closely with outside experts, governments and other companies to prevent interference during elections. This last part of critical since no one company can win this fight on its own.

Now for today's news. This morning we removed more than 650 Pages, Groups and account for coordinated inauthentic behavior on Facebook and Instagram. These were networks of accounts there were misleading people about who they were and what they were doing. We ban this kind of behavior because authenticity matters. People need to be able to trust the connections they make on Facebook.

We're still investigating and there's a lot that we don't know. And as a company we don't have all the investigative tools and intelligence that governments have, which makes it hard to always attribute particular abuse to particular countries or groups. But based on what we do know, we believe these Pages, Groups and accounts were part two separate sets of campaigns:

- First, a set of activity from Iran, including from some with ties to state-owned media;
- And second, a set of activities of the US government and others have publicly linked to Russian military intelligence services.

We've been investigating some of these campaigns for months now – which highlights the tension we face in every investigation between removing bad actors quickly and improving our defenses over time, because if we remove them too early, it's harder to understand their playbook and the extent of their network. And it can also make it harder for law enforcement who are running their own investigations, as well.

As I have said before, security is not something that you ever fully solve. Our adversaries are sophisticated and well-funded, and we have to constantly keep improving to stay ahead. But the shift we've made from reactive to proactive detection is a big change -- and it's going to make Facebook safer for everyone over time.

Now I'm going to hand it over to Nathaniel and Guy to share some more details.

Nathaniel Gleicher: Thanks, Mark. As noted, we're talking about two separate sets of activities where we took action today. I want to offer some details on each, and we have additional facts and context in our Newsroom post.

First, regarding the inauthentic activity originating in Iran, we saw content targeting people across multiple Internet platforms in the Middle East, the U.S., the U.K., and Latin America. FireEye -- one of the outside cybersecurity firms we work with -- gave us a tip in July about Liberty Front Press, a network of Facebook Pages and accounts on Facebook and other Internet services. They've published an initial analysis and they will release a full report of their findings. We wanted to take this opportunity just to thank them for their work.

Our investigation, building off of this tip from FireEye, was able to expand this initial network of assets linked to Liberty Front Press and link it to Iranian state media through publicly available website registration information, as well as the use of related IP addresses and Facebook Pages sharing the same admins. For example, one part of the network, "Quest 4 Truth", claims to be an independent Iranian media organization, but is in fact, linked to Press TV -- an English-language news network affiliated with Iranian state media. The Liberty Front Press accounts that we found as part of this investigation were created beginning in 2013.

Beginning in 2017, they increased their focus on the U.K. and U.S. These accounts and Pages that have been linked to Liberty Front Press as part of this investigation typically posed as news and civil society organizations, sharing information in multiple countries without revealing their true identity. In total, we found 74 Pages, 70 accounts, and three Groups on Facebook, as well as 76 accounts on Instagram that were associated with Liberty Front Press.

We also identified about $6,000 in spending for ads on Facebook and Instagram, paid for in U.S. and Australian dollars. The first ad was created in January of 2015, and the last was created in August of 2018. That's one investigation in this first set.

Second, we also found links between Liberty Front Press and another set of accounts and Pages that we had seen and disrupted over the past several years. These accounts posed as news organizations and also engaged in traditional cybersecurity attacks, including attempts to hack people's accounts and spread malware. The assets associated with this set include 12 Pages and 66 accounts on Facebook, with approximately 15,000 followers.

The third part of this investigation uncovered another distinct set of accounts and Pages, the first of which was created in 2011 that shared content about Middle East politics, primarily in Arabic and Farsi. They also shared content about politics in the U.K. and U.S. in English. While we

believe this set was operated from Iran, our investigation to date has not revealed any connections to either Liberty Front Press or Iranian state media. This group had 168 Pages and 140 accounts on Facebook. There are about 813 accounts -- 813,000 accounts that followed at least one of these Pages.

We're still investigating, and we've shared what we know with the U.S. and U.K. governments. Since there are U.S. sanctions involving Iran, we've also briefed the U.S. Treasury and State Departments. These sanctions allow companies to provide people in Iran with Internet services for personal communications, including the government and its affiliates, but Facebook takes steps to prevent people in Iran and other sanctioned countries from using our ad tools.

For example, our systems screen every advertiser to identify their current location and whether they're named on the U.S. government's list of sanctioned individuals. Based on what we learn in this investigation and from government officials, we'll make changes to better detect people who try to evade our sanctions compliance tools and to prevent them from advertising.

Finally, we've also removed Pages, Groups, and accounts that can be linked to sources that the U.S. government has previously identified as Russian military intelligence services. This is unrelated to the activities we found in Iran; these are distinct sets. While these are some of the same bad actors we moved for cybersecurity attacks before the 2016 U.S. election, the more recent activity that we've taken action on that we're -- that we're discussing today focused on politics in Syria and Ukraine.

For example, they were associated with the Inside Syria Media Center, which the Atlantic Council and other organizations have identified as covertly spreading pro-Russian and pro-Assad content. To date, we have not found activity by these activities targeting the U.S.
We're working closely with U.S. law enforcement on this investigation, and we appreciate their help. These investigations are ongoing, and we anticipate that as they continue, we will learn more. But as such, we can only discuss what we know at this time. These are some high-level key points. We've also shared more information and detail in our Newsroom post.

Now let me turn to Guy Rosen.

Guy Rosen: Thank you, Nathaniel.

To close out, as Mark mentioned, our efforts to stop bad actors on the platform, they're an arms race, and our security work continues to improve. Since 2016, we've made significant progress on stopping fake accounts. We've worked to develop systems to discover coordinated operations like these ones, as well as financially motivated operations. We're also putting new transparency products in place for both ads and for Pages, along with new verification processes. In fact, our ads verification systems flagged some political ads these actors attempted to run.

These efforts work in tandem with our progress to reduce the spread of misinformation and fake news. Put another way, our new products and the improved systems operate at scale, effectively shrinking the haystack, which allows our investigative teams to more effectively sift through and look for the specific needles.

We know we can always do better, so we continue to make investments in technology and in people as we work to improve people's experience on Facebook. And as we've done in the past, we'll share with the public and the relevant authorities as we find more.

Q&A

Julia Boorstin, CNBC: Hi. Thanks so much, guys. Mark, I don't know if you can give us a sense of how many more potential situations there are like this. You identified these two different groups. Do you think -- are you aware of others in the works? You obviously said you don't want to give away too much, but do you think these -- this is the first of many of such things you're going to be disclosing and shutting down?

Mark Zuckerberg: Thanks, Julia. A couple of weeks ago, we updated on a separate coordinated inauthentic behavior campaign that we had found that was helping to promote events which we were able to cancel before some of those events had taken place. I think it's safe to say that we have a number of investigations that are going on and we'll update you when we know more.

Nathaniel Gleicher: We're always looking for more activity like this. This is something that we have a core investigative team that's focused on. And we've said before that we expect we're going to find more and that when we do as we learn about it we're going to update law enforcement, we're going to update government, and we're going to -- when we can -- we're going to update the public.

Deepa Seetharaman, Wall Street Journal: Hi. Thank you, all, for doing the call. This -- or the tip that had sparked this entire investigation seems to have come from an outside firm. Have you guys learned anything from this situation that would allow you to find this kind of content by yourself?

Nathaniel Gleicher: That's a great question. Just to be precise, we're actually talking about four distinct investigations here -- three investigations involving activity linked to Iran, one investigation involving activity linked to Russia.

One of the core investigations, which is Liberty Front Press, was kicked off by a tip from, as you mentioned, FireEye, a cybersecurity company that we work with. Some of these investigations were the result of our own team's internal investigations; some of this work was a result of working with companies like FireEye; some of this was the result of open source investigation; and some of it was working with law enforcement. For several of these most critically, we actually found them only recently in July, and so we're moving very quickly to action them.

One of the things that we've learned that is very clear is that no one company can solve this problem on its own. We have a team of investigators that are looking for this material. But there are always more eyeballs and more experts in the cybersecurity world that in any one company. And so one of the things we're certainly learning as you saw from our work with the Atlantic Council previously and with FireEye here today is that it's really important that we all work together to be able to find and disrupt this activity.

Mark Zuckerberg: One thing that I'd add to this, just on the strategy, is that it's as important that we focus on building relationships with folks in law enforcement and governments and other companies so that way we can exchange information and find threats as it is to build up our own capacity to do this in isolation. Like Nathaniel said, this is something where you need to be able to pool signals together to be able to do this work effectively. And we think that there's a lot of good work happening on both of those fronts.

Donie O'Sullivan, CNN: I guess my main question is, are you guys planning on sharing the archive of these Pages so that researchers and the reporters can learn from what they look like and figure out ways to identify them in the future?

Nathaniel Gleicher: Today we're going to be sharing samples of the relevant content, particularly to sort of help inform the public about what we've found. Part of what we're doing and part of why we're building out our ability to work with researchers like FireEye and the Atlantic Council is that they're able to more deeply analyze and independently analyze the nature of the content in these spaces and provide exactly the kind of context and analysis that you're talking about doing.

FireEye has already released a sort of initial analysis, and they're going to do a more comprehensive one in the days and weeks to come, so that will give -- that will answer some of that going forward. And given the sort of ongoing status of these investigations, we're working with law enforcement, we're working with lawmakers and the government bodies to be able to share as much as we can.

Jo Ling Kent, NBC News: Thank you. Thanks for doing this call, guys. This question is for Mark. Mark, what's your message to worried Facebook users who like the product, who use it every day, but who are also concerned about the continuing attempted influence campaigns that are coming from Russia and Iran?

Mark Zuckerberg: What I'd say is that we're extremely focused on helping to keep people safe. And we know that that's a big responsibility that we have, and it's one of the core principles that we focus on here at the company. When you operate services at the scale of the ones that we do, you're going to see all of the good things that people can do around the world, and you're going to see people try to abuse the services in every way possible as well.

And certainly that's been a theme that we've been focused on for the last couple of years -- is making sure that it is hard for anyone to abuse the services, including, now, nation-states and folks who have very sophisticated and well-funded efforts that aren't going to stop. We need to make sure that we continue strengthening the security operations that we have to be able to defend against those better.

We're committed to doing that. And I think that while the efforts here are still early, I think that this is starting to show some real progress, too.

Sarah Frier, Bloomberg: Hi. I'm wondering if you can give me any detail to what extent these influence campaigns were trying to target elections in particular and whether you've been able to link the activity you announced on July 31st to any particular country?

Nathaniel Gleicher: That was sort of two questions. But the first one, I mean, we're not really in a position to assess the motivation of these bad actors and what they were or were not attempting to accomplish. We've already alerted law enforcement to our findings and we plan to continue working with them and with federal officials who are better positioned to conduct those kinds of assessments.

We also -- you were asking about our earlier announcement -- we don't have anything additional to announce today in terms of further connection or identification based on the actors we identified in late July.

Ryan Nakashima, Associated Press: Hi. Thanks for doing this. Sorry, it sounds like I'm repeating a question. But in the start of the introduction, you mentioned the midterm elections, but when you talked about the accounts, you just said they're being inauthentic. What are the links that you can make right now that ties this activity to the elections, if you could?

Nathaniel Gleicher: When we were talking about the midterm elections, I think what we were talking about was building up our capacity to detect malicious activity. And that's something that we've heavily focused on. But here, as I said, we're not seeing and we're not really in a position to assess the direct targeted goals of these actors. What we're doing, though, is we're rapidly expanding our ability to identify these malicious actors and make it harder for them to operate under concealment on the platform.

Guy Rosen: We're continuing to invest to proactively go out and find these kinds of activities to both build a system that can take these down to scale and detect them, but also to have the investigative skills to be able to track these down, and ahead of all elections around the world. Whether it's the accounts that we announced we found ahead of the Mexican election, whether it's our efforts in Brazil, and whether it is our efforts ahead of the midterms, we are continuing to invest to make sure that we can get ahead of any threats that might be coming.

Olivia Gazis, CBS News: Hi, thanks very much for doing the call. I'd appreciate if you can speak a little bit to the issue of attribution and what distinguishes these examples from the July 31 example, why you were reluctant to issue attributions to a state actor in that case and why in this case you're able to clearly and candidly say they are Iran- and Russia-sponsored influence campaigns? Thanks.

Nathaniel Gleicher: To be precise here, we're talking, as I said before, about four separate investigations. And during our -- aligned with our announcement in late July, we talked about how challenging attribution can be and the different pieces to attribution. And one of the things we said then is that we're best placed to be able to identify and comment on the sorts of technical indicators that we can identify in our platform. We're not in a position to conduct the sort of intelligence analysis that a law enforcement or intelligence organization might do about behaviors or motivations.

In this case, with these four assets, with these four investigations, for the first one, we were able to say that the -- that we saw technical indications that the accounts and the Pages were operated in Iran, and we also identified specific links between some of them and an Iranian state media entity.

For the second investigation, we were able to identify that the assets, the accounts and Pages were operated in Iran and had some links back to the first set.

For the third set, the third investigation, we were only able to identify that they were operated from Iran. We didn't see any links back to that -- either of the first two sets or to Iranian state media entities.

And for that fourth investigation, these assets have been previously identified -- not necessarily by us, but by intelligence services in the U.S. -- as linked to Russian intelligence agencies.

In each of these, what we've been able to do -- similar to what we did with the July announcement -- is outline the technical indicators that we've identified, the facts that we've been able to identify about either where these were operated from or the types of entities they were connected to. And that's what we've done here.

Hannah Kuchler, Financial Times: Hi, thanks for having the call. I wanted to ask, for the political content that was targeted at the U.K., was anything related to Brexit?

Nathaniel Gleicher: From what we found so far -- we're still working through all the content. It's a good question. But so far, we have not seen significant ad content in advance of Brexit. Of course we're going to work closely with the relevant authorities and governments here in the US and, as you ask, in the U.K., to provide them with the information they need to understand what we found, including any content that was shared by these accounts.

Ashley Gold, Politico: Hi. My question is for Mark. You're going to have Sheryl Sandberg up on the Hill on September 5th. What do you think is the most important message Facebook needs to be presenting to Congress when you come back? And you be testifying under oath again in light of these things that have popped up since the last time you were on the Hill.

Mark Zuckerberg: When we send people to testify we don't really think about it as a message we're trying to convey as much as trying to help answer questions to inform Congress and the people who the American public has entrusted to collect all the information and figure out what the country needs to do as a whole on this. That said, I do think that part of what we're trying to communicate overall and not just through testimony is all of the steps that we're taking to secure the integrity of elections on Facebook.

After 2016, when we found not just traditional cybersecurity threats and hacking and phishing, it was clear that we also faced these new kinds of coordinated information operations. We had a lot of work ahead of us to implement the strategies to be able to prevent that in elections around the world -- in order to put AI systems in place to find fake accounts, hire up -- we now have 20,000

people working on security and content review, implementing ads transparency to a higher standard than what's even on TV or print media today, verifying advertisers running political and issue oriented ads, working on deeper partnerships with law enforcement and government and other companies to be able to do signal sharing so that way we could run investigations like this better.

In all of these results around the world -- there have been a number of elections since the 2016 U.S. election. There was the French election, the German election, the Alabama special election, the Mexican election -- and in each of these elections, our systems have been able to find a lot of fake accounts that were attempting potentially to do bad things on the system. And we feel like each time we get better at identifying this kind of activity upfront and putting barriers in place to those who would try to abuse these systems.

We get that 2018 is a very important election year -- not only with the midterms here in the U.S. -- but also the Mexican election that just happened, the Brazilian election, leading into E.U. elections coming up, the Indian elections in early next year -- so this is really serious.

This is a top priority for our company. We're taking a lot of steps to be able to make sure that we do what we need to here. And I think that this investigation and what the four investigations here and what we found is just one more step in terms of the efforts that our team has taken to be able to find and remove this kind of bad content and add more hurdles in place to prevent this from happening.

Casey Newton, The Verge: Hi. I had a question about the third plank of your investigation. It's not clear to me reading your blog post what inauthentic behavior you found. Can you elaborate at all?

Nathaniel Gleicher: When you say the third plank, do you mean the third investigation involving potential activity emanating from Iran? I just want to make sure.

Casey Newton: That's right.

Nathaniel Gleicher: So what we see here is a set of accounts and Pages that were sharing content about Middle East politics in Arabic and Farsi and were also sharing content about politics in the U.K. and U.S. in English. These represented themselves as independent entities, but we were able to see through our investigation that they were actually all linked together.

What we're seeing here is a coordinated network, but the assets themselves were not presenting a coordinated front in terms of how they identified themselves.

END