UBER

Testimony of

John Flynn
Chief Information Security Officer
Uber Technologies, Inc.

Subcommittee on Consumer Protection, Product Safety, Insurance, and Data Security
Committee on Commerce, Science and Transportation
United States Senate

February 6, 2018

Mr. Chairman, Ranking Member Blumenthal, and members of the Subcommittee, my name is John Flynn. Since July 2015, I have served as the Chief Information Security Officer for Uber Technologies, Inc. I am grateful for the opportunity to testify today regarding bug bounty programs, the 2016 data security incident at Uber, and lessons that we – and the broader technology community – have learned from that incident. I am honored to be on such an esteemed panel with people who have brought such an important security practice to companies worldwide.

Before addressing today's topics, I would like to tell you a little about myself. My parents were USAID diplomats and Peace Corps volunteers. After studying computer engineering at the University of Minnesota, I too joined the Peace Corps. As a Peace Corps volunteer, I served for more than two years in Belize, where I helped lead a program that ensured teachers had access to computers and I taught classes on information security. After the Peace Corps, I attended night classes to obtain a master's degree in computer science while working full time as a Security Engineer at the George Washington University here in Washington. Before joining Uber, I held positions as an Information Security Manager at Google, and as an Information Security Director at Facebook. I have spent over a decade working on highly technical data security issues, during a period in which data security has expanded dramatically as a field and as a paramount priority for the technology industry and the country.

I would like to focus on three topics in my testimony today. *First*, I have significant experience with bug bounty programs from working for multiple companies, and will explain the important role that such programs play in the never-ending battle against cyber threats. *Second*, I will provide my perspective on the 2016 data security incident at Uber. My primary involvement in that matter was on the technical side, working under our chief security officer, and leading the effort to determine how the intrusion occurred and then to close the gaps that intruders exploited. While I am in a strong position to address the technical aspects of that incident, I was not actively involved in the process of identifying the intruders or interacting with the intruders once they were identified by others. *Third*, we learned valuable lessons from the 2016 incident, and I will describe the additional layers of protection and other enhancements that we have

implemented to secure our users' data and minimize the risk of future intrusions.

***Importance of Bug Bounty Programs***

Bug bounty programs are a critically important tool and widely used as part of comprehensive data security programs. Of course, bug bounty programs do not take the place of dedicated internal security teams who work throughout the entire software development lifecycle to detect and repair vulnerabilities. At Uber, there are multiple teams of specialized experts constantly working to ensure that our systems are secure. My team consists of more than 100 people with experience in technical areas of security. Our security efforts generally involve the following: (1) controlling access to our systems and services; (2) using security by design principles during the planning process; (3) auditing and testing code during development and throughout its lifecycle; (4) monitoring for threats; and (5) managing ongoing reinforcement and patching processes to protect our systems and software from reported vulnerabilities.

Bug bounty programs are a useful addition to these steps. Let me briefly explain bug bounty programs. All complex systems have "bugs"—imperfections unintentionally written within the software's code. Sometimes these bugs create vulnerabilities, which could be exploited by an intruder to gain access to confidential data. Security teams across the industry, including those at Uber, invest heavily in preventing and identifying as many of these bugs as we can before code is updated in our products. However, due to the evolving nature of software, programmers continuously update code by augmenting, rewriting, and overwriting their prior work. That process inevitably results in unexpected errors and vulnerabilities. To help mitigate this reality, bug bounty programs allow companies to access additional skilled individuals to augment our in-house engineers. This outside perspective is also valuable in providing a fresh set of eyes and new ways of thinking to help our security teams address various challenges with innovative solutions.

Typically, a bug bounty program is an invitation for outside experts (commonly referred to as "researchers") to search voluntarily for vulnerabilities and report them to the company or government agency that is the sponsor of the particular bug bounty program. This is supposed to be done pursuant to specific guidelines, as well as defined parameters regarding the types of systems that should be searched. For example, Uber posts a "treasure map" online to tell our researchers where to look for bugs in our systems. It points our researchers to the systems we care the most about.

Companies typically offer rewards, or "bounties," in recognition of the work performed by the researchers. Monetary bounties vary in size, from hundreds of dollars to hundreds of thousands of dollars, depending on the severity of the bug. Companies may also offer physical items, such as branded apparel, commemorating bugs that are found, as a non-monetary reward for the researcher. "Street cred" and public recognition also go a long way to motivate researchers, so many companies publish information about the most impressive bugs found.

Not surprisingly, the security benefits of bug bounty programs have motivated many major technology companies, including Uber, Google, Facebook, Microsoft, and others, to implement bug bounty programs. Moreover, the U.S. Government also has recognized the value

of bug bounty programs to protect its sensitive information technology systems. For example, the U.S. Department of Defense has bug bounty programs such as "Hack the Pentagon" and "Hack the Air Force," which the Department has operated with great success. In addition, last July, the Computer Crime and Intellectual Property Section of the U.S. Department of Justice issued *A Framework for a Vulnerability Disclosure Program for Online Systems*, which provides helpful guidance on how to design and operate a bug bounty program.

In 2015, when I joined the company, one of the first things we did to improve security was launch a bug bounty program. This was a private "beta" program and included about two hundred researchers who helped us identify and remediate nearly 100 bugs. Following the success of our beta program, we launched a public bug bounty program in March 2016. Our current program, hosted by HackerOne, offers a combination of public recognition and monetary bounties as incentives for researchers to search our products and websites for potential bugs.

Since its initial launch, this bug bounty program has assisted Uber in resolving more than 800 system vulnerabilities. The program's monetary payout stands at approximately $1.3 million in total. For us, this bug bounty program has been incredibly valuable, achieving very significant improvements in our data security posture for a relatively modest expenditure. I believe many other companies and agencies have had a similar experience with bug bounty programs.

Our bounties typically range from a few hundred dollars to several thousand dollars—depending on the impact and severity of the bug. Given the large number of companies with bug bounty programs, monetary payments can help incentivize bug hunters to focus on Uber's bugs. That is, companies compete for the time and attention of these outside researchers, and relatively modest monetary incentives help ensure that researchers focus their attention on our software. Again, I think many companies and agencies have reached this same view.

The vulnerabilities found by our researchers demonstrate the concrete value of bug bounty programs. As we have publicly shared, one researcher discovered a bug in the SSH authentication system used between different internal services. If exploited, the bug could have allowed escalation of internal privileges. This would have allowed people to access systems they did not have privileges to access. Another researcher who participated in our public bug bounty program found a "remote code execution" bug on one of our websites. This was an important issue because remote code execution gives attackers the ability to run commands on a target computer. In this case, the researcher demonstrated the ability to execute commands on a system within our data center. Potentially, a malicious attacker could have used this vulnerability to access sensitive user data.

Uber's bug bounty program unquestionably has increased the scale and speed at which we are able to identify and eliminate cybersecurity threats. We are constantly refining our tools to prevent the bugs that are found from being written into our code in the first place.

Over the nearly three years we have been running this program, more than 500 researchers have participated. Through our bug bounty program, we can benefit from a vast, diverse, worldwide pool of talent, often beyond our ability to hire.

Of course, operating a bug bounty program is not without its challenges. Security researchers can be an eccentric group, and within this community there are individuals with varying degrees of technical experience and professionalism who engage through bug bounty programs. Researchers sometimes express concern with the amount of the bounty that is paid, believing that their discovery may be worth more than we determine was appropriate, based on our program guidelines. Other times, a researcher may identify a bug that we already know and are working to fix. The researcher sometimes takes issue with not receiving a monetary reward for those already identified bugs. Occasionally, a person may contact the company to report a vulnerability (without exploiting it), completely unaware of our bug bounty program, and make a demand for compensation. We try to work with such persons to submit their report through the bug bounty program in exchange for a fair reward under the program guidelines.

***2016 Uber Data Security Incident***

The 2016 data security incident unfolded in a way that is entirely different from the typical bug bounty program scenario. On November 14, 2016, Uber's security team received emails from an anonymous individual who claimed to have accessed Uber data and demanded a six-figure payment. Uber investigated and determined that the individual and another person working with him had obtained access to certain archived copies of Uber databases and files located on Uber's private cloud data storage environment on Amazon Web Services ("AWS"). In line with standard protocol, Uber assembled an incident response team. This team included technical experts whom I directed, and we worked quickly to determine the means of access, shut down the compromised credential, and take various steps to secure our systems against a further attack. To the best of Uber's knowledge, the intruders' access began on October 13, 2016, and there was no further access by the intruders after November 15, 2016.

For the Subcommittee's information, I would like to explain in greater detail how Uber responded to this security incident. As with any security incident, the first step was to validate the claims that the intruder had made. Very often these situations are hoaxes. The Uber security team requested data from the intruder, which he provided, and then confirmed that the data were Uber's. With that validation, we initiated an incident response procedure. Incident response to any data incident is an orchestrated affair. The first steps involve fast, intense work with limited information and a very short time to eliminate the threat. We set up a command center where members of the team could work in parallel and discuss issues in real time.

The overall effort was led by our former Chief Security Officer, Joe Sullivan, to whom I reported. I led the technical work to identify how the intrusion occurred and remove the vulnerability. Joe Sullivan and others led what we call "attribution"—the process of identifying the intruders.

During the technical effort, we immediately began the process of determining where the data at issue resided and how the intruder gained access. Within 24 hours, we determined that the data came from back-up files stored in an AWS S3 bucket. S3 stands for "simple storage service."

The next step of the investigation for my team was to determine how the intruder gained access to the AWS S3 bucket, which requires access credentials. We learned that the intruder found the credential contained within code on a private repository for Uber engineers on GitHub, which is a third party site that allows people to collaborate on code. We immediately took steps to implement multifactor authentication for GitHub and rotated the AWS credential used by the intruder. Despite the complexity of the issue and the limited information with which we started, we were able to lock down the point of entry within 24 hours.

Subsequently, we did a thorough review of our GitHub repositories. My technical team initiated the process of removing additional code from GitHub that could be considered sensitive, and confirming rotation of keys. We ceased using GitHub except for items like open source code. The incident response team also worked to identify the type of data downloaded to assess the risk.

In addition to the technical response, another team worked on attribution. Although I was not directly involved, I understand that the attribution team used various methods, including forensics, to gather further information on the intruders. This was a challenging endeavor because the intruders were extremely adept at covering their tracks.

Ultimately, the attribution team ascertained the real identity of both the original individual who contacted the company, and the second person working with him. I understand that the original individual was located in Canada, and that his partner, who actually obtained the data, was in Florida. I further understand that the attribution team made contact with both individuals and received assurances that the data had been destroyed.

As you know, Uber paid the intruders $100,000 through HackerOne and our bug bounty program. Our primary goal in paying the intruders was to protect our consumers' data. This was not done in a way that is consistent with the way our bounty program normally operates, however. In my view, the key distinction regarding this incident is that the intruders not only found a weakness, they also exploited the vulnerability in a malicious fashion to access and download data.

In 2017, after learning about the incident, new company leadership at Uber asked an independent cybersecurity firm, Mandiant, to conduct a thorough analysis of the data at issue. Mandiant's analysis showed that the data included information pertaining to approximately 57 million users worldwide, including approximately 25 million users in the United States. Of these, approximately 4.1 million users in the United States were drivers. For nearly all users, the downloaded files included names, email addresses and phone numbers. In some cases, the information also included information collected from or created about users by Uber, such as Uber user IDs, certain one-time locational information (e.g., the latitude and longitude corresponding to the location where the user first signed up for the Uber service), user tokens, and passwords encrypted using hashing and salting techniques. Of the driver accounts, approximately 600,000 thousand included driver's license numbers.

In their independent analysis, Mandiant found *no* indication that trip location history, credit card numbers, bank account numbers, Social Security numbers, or dates of birth were

compromised.

***Lessons Learned and Data Security Enhancements at Uber***

While the circumstances surrounding the 2016 security incident remain under investigation by the company and multiple regulators, and I am not privy to the details of those ongoing investigations, there are a number of lessons learned that I would like to highlight today.

First, I would like to echo statements made by new leadership, and state publicly that it was wrong not to disclose the breach earlier. The breach should have been disclosed in a timely manner. The company is taking steps to ensure that an incident like this does not happen again, with personnel changes and additional remedial actions. We are working to make transparency and honesty core values of our company. I would add that this is a change that I personally am gratified to see and wholeheartedly support.

Although we regret that we did not publicly report the incident in 2016, we did at that time take numerous steps internally to improve our security posture in response to the incident. As I noted previously, we immediately instituted multifactor authentication on Github. We then subsequently ceased using GitHub except for items like open source code. As to AWS, we were already using multifactor authentication for individual access accounts—which these intruders did not compromise. After the incident we expanded the use of multifactor authentication protocols for AWS service accounts using techniques such as IP restrictions, commonly referred to as "white listing." We have also taken other steps to enhance security for AWS data storage, such as refining Identity & Assessment Management permissions, improving our ability to authenticate someone before granting access to these systems and to confirm whether they are authorized to access them. We also added auto-expiring credentials to protect further against attacks using exposed, lost, or shared credentials. We continue to look to Amazon's evolving best practices and guidance to protect our AWS system.

We recognize that the bug bounty program is not an appropriate vehicle for dealing with intruders who seek to extort funds from the company. The approach that these intruders took was separate and distinct from those of the researchers in the security community for whom bug bounty programs are designed. While the use of the bug bounty program assisted in the effort to gain attribution and, ultimately, assurances that our users' data were secure, at the end of the day, these intruders were fundamentally different from legitimate bug bounty recipients.

Going forward, Uber is revisiting its incident response approach in circumstances such as these. We have hired Matt Olsen, a former general counsel of the National Security Agency and director of the National Counterterrorism Center, to help structure the security team and guide new processes going forward. I have already seen some of these changes take place, such as more stakeholders involved in the decision-making process for how to handle security incidents, and informing law enforcement of potential security incidents right away.

I would like to conclude by stating that we strongly support a unified, national approach to data security and breach standards. We are proactively engaged in the many conversations in both the technical and policy communities to help identify what the critical components of

federal data breach legislation should be, and are pleased to see this robust conversation taking place with various Members of Congress and your staff. We welcome the opportunity to be at the table to help all stakeholders understand the best practices.

* * *

Thank you again for the opportunity to appear and testify today. I would be happy to answer your questions.