Prepared Statement

for the Record of

Ken Durbin, CISSP

Senior Strategist: Global Government Affairs & Cybersecurity

Symantec Corporation

Hearing on

"Preparing for the Future: An Assessment of Emerging Cyber Threats"

Before the

United States House of Representatives

Committee on Homeland Security

Subcommittee Cybersecurity, Infrastructure Protection, & Innovation

October 22, 2019

Chairman Richmond, Ranking Member Katko, my name is Ken Durbin, CISSP, and I am a Senior Strategist for Symantec Global Government Affairs and Cybersecurity. I have been providing Solutions to the Public Sector for over 30 years. My focus on Compliance and Risk Management (CRM) and its application in both the public and private sector has allowed me to gain insights into the challenge of balancing Compliance with the implementation of Cybersecurity Solutions. Additionally, I focus on the Standards, Mandates and Best Practices from NIST, OMB, DHS, etc. and their application to CRM. I spend a significant amount of my time on the NIST Cybersecurity Framework (CSF)[1] and the emerging Privacy Framework, the DHS Continuous Diagnostics and Mitigation (CDM) Program and the EU Global Data Protection Regulation (GDPR.)

Symantec Corporation is the world's leading cyber security company, allowing organizations, governments, and people to secure their most important data wherever it lives. Organizations across the world look to Symantec for strategic, integrated solutions to defend against sophisticated attacks across endpoints, cloud and infrastructure. Likewise, a global community of more than 50 million people and families rely on Symantec's Norton and LifeLock product suites to help protect their digital lives at home and across their devices. Symantec operates one of the world's largest civilian cyber intelligence networks, allowing it to see and protect against the most advanced threats. In my testimony I will discuss the current Threat Landscape, to include:

- Key findings from the 2019 Symantec Internet Security Threat Report (ISTR);
- Mobile Security privacy;
- Deepfakes risk to the Enterprise;
- Twitterbots in the 2016 election;
- Targeted Ransomware; and
- Stalkerware

The Threat Landscape

A review of the current threat landscape shows there are challenging new attacks and threats that need to be addressed. However, it also shows that it would not be wise to ignore the traditional threats we have been dealing with for years. Bad actors are finding new ways to attack using well established attack vectors. At the same time new technologies and campaigns are emerging to exert influence and drive behavior. I'll address both traditional and emerging threats in the following sections.

[1] NIST Cybersecurity Framework (CSF): Provides guidance to private companies on how best to prevent, detect, and respond to Cyber attacks

The Internet Security Threat Report

The Internet Security Threat Report (ISTR)[2] analyzes data from Symantec's Global Intelligence Network, the largest civilian threat intelligence network in the world, which records events from 123 million attack sensors worldwide, blocks 142 million threats daily, and monitors threat activities in more than 157 countries. The analysis provides insight into a wide variety of threats and identifies trends that help inform the public with the goal of helping them avoid risk. Highlights from the ISTR include:

- One out of ten URLS are malicious. That is up from one in sixteen in 2017. Clicking on a malicious URL continues to be a widely used attack vector by attackers.
- There was a 56% increase in web attacks over 2017. By the end of 2018, we blocked more than 1.3 million unique web attacks on endpoint machines every day.
- On average, 4,800 websites are compromised with Formjacking software each month. Formjacking is the use of malicious JavaScript code to steal payment card details and other information from payment forms on the checkout web pages of eCommerce sites. We blocked 3.7 million formjacking attempts on endpoint devices in 2018.
- Supply Chain attacks increased 78%. Supply chain attacks, which exploit third-party services and software to compromise a final target, take many forms, including hijacking software updates and injecting malicious code into legitimate software.
- 48% of malicious email attachments were MS Office Documents, up from just 5 percent in 2017. Cyber crime groups continued to use macros in Office files as their preferred method to propagate malicious payloads in 2018, but also experimented with malicious XML files and Office files with Dynamic Data Exchange (DDE) payloads.
- The number of attack groups using destructive malware rose 25%. Destructive malware is designed to inflict physical damage to an organizations network or facility. While still a niche area, the use of destructive malware continued to grow. Eight percent of groups were known to use destructive tools, up from 6 percent at the end of 2017.

Mobile Security

The average smartphone user these days has between 60 and 90 apps on their device, and most of them request some sort of information about the user and the device. They may want to know your name, your email address, or your real-world address. But because smartphones are so powerful, they can also get quite a bit more than that, such as your exact location. Some apps will even request access to the device's camera or microphone despite having no legitimate need to use them.

[2] https://www.symantec.com/security-center/threat-report

In order to find out what kind of data your apps may be looking for, we analyzed the top 100 free apps as listed on the Google Play Store and Apple App Store on May 3, 2018[3]. For each we looked at two main things: how much personal information was the user sharing with the app and which smartphone permissions the app accessed.

Email addresses are the most common piece of personally identifiable information (PII) apps were accessing, as 48 percent of the iOS and 44 percent of the Android apps did so. Username was next, which was accessed by 33 percent of iOS and 30 percent of Android apps, followed by phone numbers, which were accessed by 12 percent of iOS and 9 percent of Android apps. Finally, 4 percent of iOS and 5 percent of Android apps accessed the user's physical address.

It is often reasonable and necessary to grant apps permission to access various features on a smartphone. For example, if you want to take a picture using an app, the app will need permission to use your device's camera. However, not all permissions are the same. We took a closer look at permissions that could provide access to data or resources that involve the user's private information or could potentially affect the user's stored data or the operation of other apps.

Camera access was the most requested permission, with 46 percent of Android and 25 percent of iOS apps seeking it. That was followed by location tracking, which was sought by 45 percent of Android and 25 percent of iOS apps. Twenty five percent of Android apps requested permission to record audio, while 9 percent of iOS apps did so. Lastly, 15 percent of Android apps sought permission to read SMS messages and 10 percent sought access to phone call logs. Neither of these permissions are available in iOS.

Apps have permissions because the user granted them by hitting an "*I Agree*" button – usually without considering if certain permissions make sense, and often without pausing to consider the request at all. For example: the Android flashlight app "*Brightest Flashlight LED - Super Bright Torch*", which has 10 million installs, asks for permissions including precise user location, access to user's contacts, and permission to directly call phone numbers. It is hard to imagine why a *flashlight* app has a legitimate need to copy all of your contacts, call all of your friends, or know exactly where you are located. Consumers should pause before the agree to permissions – and app developers should be very clear about what permissions their app needs and why it needs them.

Deepfakes

"Deepfakes" are audio or video tracks created or altered by artificial intelligence (AI) systems and used to make the public believe they are authentic. Most of the popular examples of

[3] https://www.symantec.com/blogs/threat-intelligence/mobile-privacy-apps

Deepfakes show politicians or actors saying or doing things designed to embarrass or harm reputations. As a result, Deepfakes are not typically viewed as a threat to Enterprise security.

This is short-sighted. Enterprises do need to pay attention to Deepfakes; fake content like videos, photos, audio recordings or emails represent a serious risk to individuals as well as the organization. The technology behind Deepfakes has advanced to the point decisions might be made based on a Deepfake, or decisions *not* made because an authentic video is thought to be a Deepfake. Deepfakes are particularly dangerous because there is such a low barrier of entry *and* because they are difficult to detect. Until recently, altering videos was expensive and required significant resources, specialized equipment, and money. Today, if someone has access to the internet, a gaming PC and the right software they can produce convincing Deepfakes. Specialized applications have reduced creating Deepfakes to a point and click exercise, reducing the need for advanced skills.

Deepfakes are created using a process based on Generative Adversarial Networks (GAN). Essentially, a GAN consists of two machine-learning networks that work in an ongoing feedback loop where one network creates the Deepfake and the second one tests the output. The networks pass the Deepfake back and forth making alterations to make it as realistic as possible. Since the GAN is "learning" throughout the process, the Deepfake becomes harder to spot with the naked eye.

Given the low barrier of entry and that they are difficult to detect, Enterprises need to understand the risks Deepfakes pose to their organization. For example: A Deepfake of a CEO announcing a massive layoff could cause their stock price to sink. A Deepfake could be used to order an employee to wire funds, or transfer intellectual property out of the company. Until a proven method to identify or block Deepfakes is developed organizations will be best served educating employees about the danger of Deepfakes and implementing rapid response plans that can be executed as soon as a Deepfake is identified.

Twitterbots

In October 2018, Twitter released a massive dataset of content posted on its service by the Internet Research Agency (IRA) beginning in May of 2014. The IRA is the Russian company behind the social media propaganda campaign directed against the 2016 U.S. elections. Symantec conducted an in-depth analysis of the dataset to learn more about how the campaign operated.

The dataset consisted of 3,836 Twitter accounts and nearly 10 million tweets. These accounts amassed almost 6.4 million followers and followed 3.2 million accounts. The sheer volume of data was enormous, more than 275 GB.

Our research[4] led to a number of interesting findings:

1. The operation was carefully planned, with accounts often registered months before they were used. The average time between account creation and first tweet was 177 days. The average length of time an account remained active was 429 days.
2. A core group of main accounts was used to push out new content. These were often "fake news" outlets masquerading as regional news outlets or pretending to be political organizations.
3. A much larger pool of auxiliary accounts was used to amplify messages pushed out by the main accounts. These accounts usually pretended to be individuals.
4. Some operatives may have been making money on the side by using monetized URL shorteners to create links. If they did monetize the URLs one account in particular could have generated almost $1 million.

We divided the accounts into two main categories; main accounts and auxiliary accounts. Each category had different characteristics and played a different role. We identified 123 main accounts, each having at least 10,000 followers. Main accounts tended to not be followers of other accounts. They were primarily used to publish new tweets.

We identified 3,713 auxiliary accounts, each having less than 10,000 followers. Auxiliary accounts tended to be followers of thousands of other accounts. Their main purpose was to retweet messages from other accounts. Since auxiliary accounts were used to amplify targeted messages it makes sense they were the larger category.

A particularly effective account in the dataset was called TEN_GOP. Created in November 2015, the account masqueraded as a group of Republicans in Tennessee. It appears to have been manually operated. In less than two years TEN_GOP managed to rack up nearly 150,000 followers. Despite only tweeting 10,794 times, the account garnered over 6 million retweets. Only a small fraction (1,850) of those retweets came from other accounts within the dataset. In other words, almost all of its retweets came from accounts outside the dataset, meaning many could have been real Twitter users.

The Twitterbot campaign is often referred to as the work of trolls, but the release of the dataset makes it obvious that it was far more than that - it was highly professional. It was planned months in advance and the operators had the resources to create and manage a vast disinformation network. And aside from the sheer volume of tweets generated over a period of years, its orchestrators developed a streamlined operation that automated the publication of new content and leveraged a network of auxiliary accounts to amplify its impact.

[4] https://www.symantec.com/blogs/threat-intelligence/twitterbots-propaganda-disinformation

Targeted Ransomware

Ransomware continues to be one of the most dangerous cyber threats facing any organization. The threat has changed significantly over the past two years, as criminals are increasingly targeting enterprises. During 2018, while the overall number of ransomware infections was down 20 percent, attacks against organizations (as opposed to against individuals) rose by 12 percent. Alarmingly, Enterprises accounted for 81 percent of all ransomware infections in 2018. Targeted Attacks have been particularly hard on State and local government organizations. In March of 2018 the city of Atlanta was attacked and ransomware encrypted servers that made over a third of the 424 citywide services inaccessible. The cleanup costs for the attack are expected to run to over $10 million. The Colorado Department of Transportation spent $1.5 million to clean up after they were attacked. Two Florida cities that were attacked took another route – they paid the ransom, which totaled $1 million between them.

The number of targeted ransomware attacks has multiplied as new groups move into this sector. Although targeted ransomware attacks account for a small percentage of overall ransomware attacks, they present a far greater risk as a successful targeted ransomware attack can cripple an ill-prepared organization. These attacks also typically involve much higher ransom demands, ranging from $50,00 to over $1 million.

Targeted attacks can result in hundreds of computers encrypted, backups destroyed, and business-critical data removed from the organization. Targeted attacks can shut down an organization, leading to loss of business, reputational damage, and multimillion-dollar clean-up bills. The number of organizations affected by targeted ransomware attacks has grown sharply over the past two and a half years. As recently as January 2017, Symantec observed just two organizations a month being attacked. However, recent months have seen that figure grow to above 50 organizations a month.

The SamSam ransomware group was the original targeted ransomware threat, but was joined in 2018 by another highly active targeted actor called Ryuk. In 2019 several additional groups were linked to a series of highly disruptive attacks in the U.S. and Europe. Current trends indicate that targeted ransomware is attracting a high degree of interest among cyber criminals, with new groups appearing at an accelerating pace, motivated no doubt by the success of some recent attacks. RobbinHood is another new family, first appearing in May 2019. It was reportedly used in the attack against the U.S. city of Baltimore that shut down several services, including municipal employees' emails, phone lines and online bill payments.

A group known as GoGalocker has used a new breed of targeted ransomware that appeared in early 2019. Traditional ransomware attackers cast a wide net using spam campaigns to improve their chances of finding a victim. GoGalocker selects targets and digs in deep. The attackers behind GoGalocker appear to be highly skilled, capable of breaking into the victim's network and deploying a wide array of tools in order to map the network, harvest credentials, elevate privileges, and turn off security software before deploying the ransomware. This process

permits the attackers to identify and access a large number of computers in order to later simultaneously infect them with the ransomware. By maximizing the number of assets, the attacker compromises the better the chances are the victim will pay the ransom.

Stalkerware

Stalkerware is a type of malware that is secretly loaded on an unsuspecting victim computing device giving almost total control of the device to a bad actor. The bad actor – who can be an ex-spouse, ex-boyfriend, or other stalker – would then know the victims exact location, be able to read their emails and texts, and even turn on their microphone or camera. Due to the control Stalkerware gives a bad actor, it is classified as a type of malware — malicious software designed to gain access to or damage your computer, often without your knowledge.

Stalkerware can affect PCs, Macs, and iOS or Android devices. Although Windows operating systems may be more susceptible to attacks, attackers are becoming better at infiltrating Apple's operating systems as well. Stalkerware typically infects a device when the victim accepts a prompt or pop-up without reading it first, downloads software from an unreliable source, opens email attachments from unknown senders, or pirate media such as movies, music, or games

So why is Stalkerware available in App Stores? Publishers of Stalkerware typically advertise their product as parental monitoring software to keep kids safe, and this can certainly be true when it is used appropriately by a responsible parent. However, any software surreptitiously loaded onto a device, no matter how well-meaning is malicious. Additionally, the features built into some of these Apps give more total control of a device than parents would need and make it ripe for abuse.

Conclusion

New threats are emerging every year - but that does not mean existing threats have gone away. We need to be vigilant in our defense against the traditional threats we have battled for years, while understanding emerging threats and planning defenses accordingly. Emails have been a persistent attack vector, yet attackers are finding new ways use the service against us. Ransomware is not new but the attacks are becoming more targeted and disruptive. Mobile security is a threat we allow by granting excessive permissions. Finally, Deepfakes and Twitterbots teach us that Cyber can be utilized to influence and force actions from a distance. The focus of the Cybersecurity, Infrastructure Protection, and Innovation committee is vital for our nation to understand the current threat landscape and ensure resources are allocated to determine how to defend against them. Thank you for the opportunity to testify before this committee, and I would be happy to take any questions you may have.