Virtual Assistant Privacy and Security

Timothy Kang
Illinois Institute of Technology
10 West 35th Street
Chicago, IL 60616
tkang6@hawk.iit.edu

Christopher Steinberg
Illinois Institute of Technology
10 West 35th Street
Chicago, IL 60616
csteinberg@hawk.iit.edu

Ian Hernandez
Illinois Institute of Technology
10 West 35th Street
Chicago, IL 60616
ihernan2@hawk.iit.edu

ABSTRACT
With the appearance of many virtual assistants such as Google Home and Amazon Echo, many begin to question the privacy and security aspects of these types of technologies. Virtual assistants are always listening for specific keywords or commands before they begin executing their tasks. This can range from answering questions, to creating a list, or even queuing a video or song by communicating with other devices around the house. The questions that arise are: Are all our conversations, regardless of the keyword, being recorded and sent back to Google or other companies? Is the data that is being sent back encrypted and safe from unauthorized people? Is this data being stored temporarily or forever and what is it being used for? In addition, how secure is the Google Home itself? Is it vulnerable to known attack vectors? Are there other means of cracking into a google home? In order to answer these questions, specific tools must be used such as wireshark, a network traffic analyzer. To keep our results consistent, we are only using Google Homes. The goal of this is to not only monitor the network traffic but to determine the extent that Google uses these devices to collect data. In addition, observe and attempt to deploy known vulnerabilities to the Google Home.

Categories and Subject Descriptors
K.3.2 Computer and Information Science Education

General Terms
Security; Experimentation; Design; Performance; Documentation; Measurement; Theory; Verification

Keywords
Virtual Assistant; Google Home; Privacy; Security; Amazon Echo; Alexa; Amazon Dot; Siri; Cortana; IoT;

1. INTRODUCTION
With the advent of virtual assistants such as Amazon's Alex and Apple's Siri, Google has answered with Google Assistant. In an effort to make their AI's more accessible Amazon released the Amazon Echo and again Google's response was the Google Home. These devices have evolved from helping consumers manage daily tasks on mobile devices, to becoming embodied into their very own device, complete with microphones, speakers, and all the tools needed to make life easier. They occupy prime real estate in consumers lives, usually located in a central location within a household as to provide easy access to the unit in the event of a question, scheduling, or other spontaneous commands. Though the initial intention of virtual assistants was to make life easier, there are inherent concerns that are attached to them. Much like any piece of technology, the inherent concerns of these devices skirt the boundary of invasion of privacy, and security of personal information. Microphones that are always listening, data that is being stored by third party entities, are some of the issues that come to mind when thinking about these services. How much of this is true remains to be seen and much of the information on these topics are largely hyperbole.

1.1 Background
The background for digital assistants started with Amazon's Alexa and grew into other assistants like Google Assistant and Microsoft's Cortana. These Assistants are built to listen to voice activated commands, using keywords, or "hotwords" and then capture audio with the commands. They then transport this data offsite to a datacenter equipped with the proper hardware to convert the commands into code and then execute them. The servers then respond with answers or actions to the commands. This has turned into a method for tying aspects of life together by using voice commands. Alexa was a hit, and in turn this brought on the onset of other IOT virtual assistants that make life easier, as well as implicate a whole new world of security and privacy concerns.

2. HYPOTHESIS
The hypothesize for this paper is that Google encrypts all the data so that a person with ill intentions is unable to break into the Google Home and see the data that is being sent from the Home to Google. However, since "Google Home is a natural at syncing up with other Google products, such as a Gmail account," the Home can be compromised if one's Google account falls into the hands of another person [2]. Regardless of whether or not it is encrypted, having access to this account would serve as a big problem. We also hypothesize that these virtual assistants listen all the time, but no data is being transmitted until the keywords are spoken. Google has made security a priority for their company and they have vast resources at their disposal to stay at the forefront of cyber security. We predict that there is little chance of any currently known vulnerabilities to be present in the Google Home. Any potential break of the Home's security will most likely come from a day zero vulnerability. Lastly, we predict that data is stored indefinitely on Google's servers as this acts as a repository for

user history and command history, as well as allows the user to recall what is happening with their device usage.

3. PRIVACY POLICY

This section will discuss Google's privacy policy and how Google will handle any data retrieved from the device. Google has an extensive privacy and data retention policy concerning information from Gmail to the Google home. They have specific policies concerning the Home itself.

3.1 Data Collection

According to the Google Privacy Policy, Google will collect information from you personally as well as device specific information. They state that personal information can consist of name, email, credit card numbers, billing information, address, and other such identifiers that are provided to them by the user. Google also states that they collect device specific information for any device that is connected to google, this is not limited to, hardware model, operating system, mobile network information, phone number, and Google will associate this information with the user's Google account. Google also collects information from its voice service, which the Google Home uses extensively, this includes SMS data, email data, call data, callers phone number, time, duration of call, recorded messages, voicemail messages, etc. While Google claims that you can remove this data, they cannot guarantee that it will be completely removed as there may be residual data in their backup systems.[5]

3.2 Data Use

In Google's privacy policy they claim they do not give out private data to any other companies outside of Google except in certain circumstances. Any data that comes from the user that does not personally identify the user is shared with affiliates of Google and third party partners. This would include relevant data that would improve ads and services. [5]

3.3 Data Storage

Google does not give a specific amount of time that they store private information on their servers for. They do state that it can be removed at any time, but do not provide specific instructions on how to do so. Google also mentions that private data can be stored on any Google server, including those overseas or out of the country.[5]

4. PRIVACY AND SECURITY CONCERNS

This section will cover the privacy and security concerns that plague these technologies.

4.1 Privacy Concern

When it comes to privacy, the Fourth Amendment of the U.S. Constitution is typically mentioned. However, can this amendment and the privacy rules that appeared over the course of the past 40 years be useful to guide the privacy concerns of virtual assistants? For example, the "third party doctrine" resulted after two Supreme Court cases in the 1970s [1]. This doctrine states that if the user shares information "with anyone or anything that constitutes a 'third party'," the expectation of privacy is completely changed [1]. Simply put, this expectation is gone the moment you make a phone call, access a web page, or in this case, communicate with Google using the Home. This clearly is a visible concern that follows the use of virtual assistants. It is stated that the Home and other virtual assistants are always listening and only when a keyword such as "Hey Google" is spoken, the virtual assistant then sends data out. For example, "Amazon's devices store roughly 60 seconds of audio before a wake word but transmit a fraction of a second of audio before the wake word, plus your interaction with Alexa, to Amazon [2]." Similarly Google states in their privacy policy for the Google Home that the Home, "listens in short (a few seconds) snippets for the hotword.[5]" Also that "Google Home records what you say, and sends that recording (including the few-second hotword recording) to Google[5]". The question that arises is: is this true? Another thing to note is that if this is truly the case, what happens in the event that someone says a phrase that is similar to the keyword? If this device can "mishear" consistently, can be a privacy concern where data is being sent out without the user's acknowledgement? In addition to this, the private data is being sent out to a remote location. How long is this data being stored and what is being done with it? This may not apply to all virtual assistants but Google is known to "share anonymized information with developers [2]" and since the user has no control over this, these questions clearly show how this can be a concern.

4.2 Security Concern

There's no misunderstanding when it comes to the fact that data is being sent from the virtual assistant to a remote location. However, security concerns that arise are: is this data encrypted? Can any hacker simply get a hold of the data and see things they should not be seeing? This concern is amplified by the fact that "[getting] into one device that is connected to Google, [allows you to] have a person's entire life in your hands" since "Google has visibility into your day to day activities [3]." Even if everything is being encrypted, there is still the possibility that the Home can be broken into; whether it is done by finding a vulnerability with the device or finding a way to access the Google Account that is linked with it.

5. THE SYSTEM

In order to discern what the Google Home is actually doing, a system has to be built to monitor its activities. The system will rely on packet sniffing software, specifically Wireshark. Wireshark captures packets that are sent and received over a specific interface. This allows the user to filter the traffic that occurs over the network and determine what is actually going on. Since the software captures packets, if the data within is not encrypted it can be easily read and monitored. If it is encrypted it will have to be decoded in order to determine what has been sent over the network.

5.1 Design

The design of the system will be using a specific hardware setup. The system will consist of a Google Home, a Windows based router, wireshark, and a wireless access point. This system will be housed in an isolated environment for the purpose of not contaminating tests with outside disturbances. The hardware of the windows router is not of any importance to the integrity of the testing and due to this we will be using a generic Fujitsu laptop running Windows Server 2016. The server will be set up to use Windows Server DHCP features (Fig. 1).

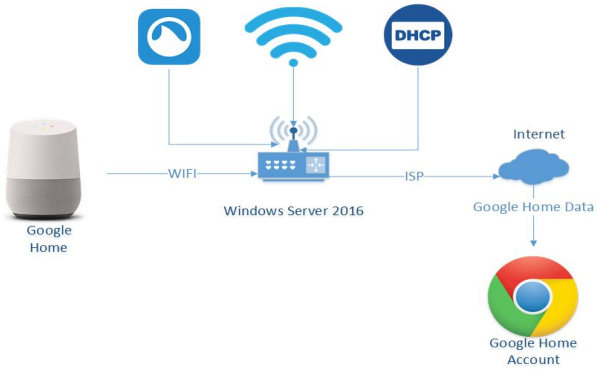

Fig. 1. This diagram shows how the system will be connected.

This allows us to use a graphical version of wireshark directly on the network gateway and we can monitor both, incoming traffic, as well as outgoing on both interfaces. The server will also be configured with a wireless LAN to incorporate an all in one monitoring system. This will isolate the Google Home and limit the amount of traffic captured to that of the Home. This will greatly simplify analyzation of network traffic and improve the accuracy of our data.

5.2 Implementation

This will cover the implementation of the system. It will describe the steps used to build the system. For this experiment, one system was used. One was setup in a test environment without access to outside noise. The idea was to isolate interactions with the Google home and provide data that is only test data.

Initial setup consisted of building the router first. The router hardware is a Fujitsu Lifebook T731. It contains an Intel i5 and 6GB of RAM. It also includes onboard wifi, and integrated ethernet LAN port. This provided an easy to use, low power DHCP router for the system, this also provided the option for a battery in the event the power goes out. Windows Server 2016 Standard 64-bit Desktop Experience was installed on the server to provide the necessary OS functionality. After OS installation the necessary roles were installed. These consisted of DHCP server role, and .NET framework 3.5. After the services were set up and installed, Firefox and Wireshark were installed to the server. Next the WLAN was configured. Using the command line the wireless adapter was changed to broadcast an SSID and DHCP was configured to provide an IP to the Google Home. After this the WAN ethernet port was configured to share internet with the Wireless LAN port to provide internet connectivity to the Google Home. The router is now ready to connect to the google home and the internet. Post router setup consisted of downloading the google home application from the google play store, then linking the account that was to be monitored to the app. After this was completed the next step was to connect to the google home via bluetooth and then follow the wizard to prep the home for use. After this was done the setup of the test environment was complete and the system was ready for monitoring.

6. TESTING

Now that the system has been implemented, it is time to test the Google Home and then begin capturing the packets. To test the Google Home, we developed a set of tests to see what the Google Home did with the data provided to it. Initial tests were designed to perform basic functions of the Google Home and then analyze the data into graphs to graphically show what the device is actually doing. The first test was to stream some music with the resulting graph (figure. 2).

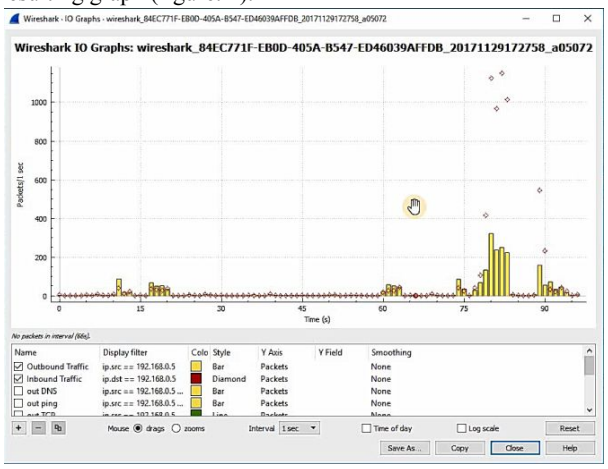

Fig. 2.1. A graph displaying Inbound traffic after asking the Google Home to stream music. This is represented via the diamond marks.

The second test was to test outgoing data, for instance asking the Google Home to define a word or a phrase. This forces the Google Home to send the data for processing and shows how much of it was sent (figure. 2.1).

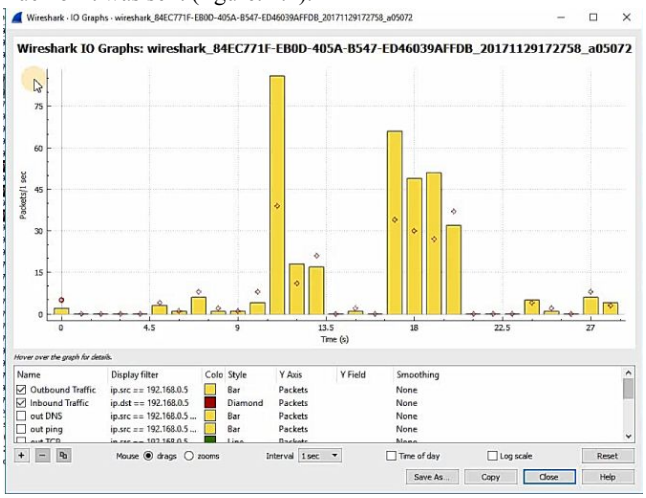

Fig. 2.1. A graph that shows outbound data after a question is asked of the Google Home. This is represented via the yellow bars.

Further testing consisted of testing what data is kept for ads. This was composed of asking the Google home about specific products or items (figure. 2.2).

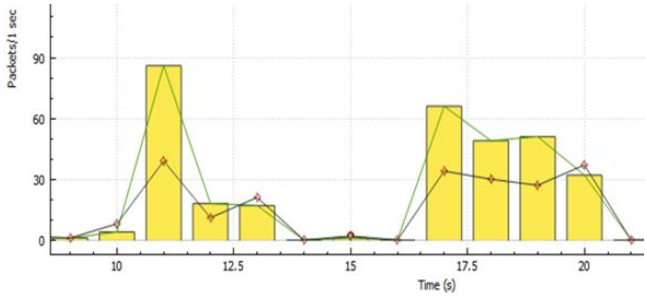

Fig. 2.2. Graph of the Google Home after asking about specific products. This is shown by the yellow bars.

7. SECURITY VECTORS

Saying that these virtual assistants have no vulnerabilities is far from true. Many people have attempted and at times, succeeded in exploiting a vulnerability in these machines. There are many vectors when it comes to the security aspects of virtual assistants. This includes bluetooth, physical, audio, and network vulnerabilities. For example, the BlueBorne attack exploited bluetooth vulnerabilities of both the Google Home and Amazon Echo. Multiple vectors will be discussed and tested.

7.1 BlueBorne

BlueBorne is one of the most recent bluetooth attacks that has found vulnerabilities in a large number of devices across multiple platforms. By exploiting a fault in the bluetooth protocol, BlueBorne is able to gain access to certain devices and perform tasks that include, man-in-the-middle, information leak, and even command line execution [6]. This exploit has been patched and you can check your own device on whether it is vulnerable to this attack by using the BlueBorne Scanner app from the Google Play Store. Since this has been patched, this exploit was attempted using an unopened Google Home that has not been configured and connected to a wireless network. In order for this attack to be done, a MAC address of the device is required. Our test setup was not successful in obtaining the MAC address through conventional scanners such as BTscanner, and hcitool. To overcome this, Blue Hydra was used, this tool detects both classic and low energy bluetooth devices over time and neatly displays a chart of nearby devices along with additional information.

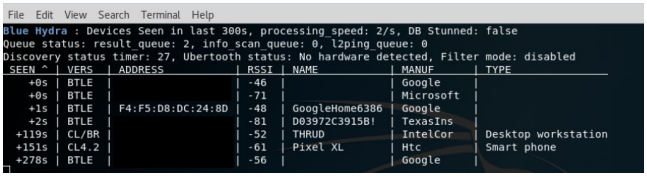

Fig 3.1 Displaying nearby bluetooth devices using Blue Hydra

As one can see in Fig. 3.1, the Google Home device has been detected by this service. In order to now attempt an attack on the device, a python script of the attack can be used and it is necessary to add the MAC address of the Google Home as the target. Unfortunately, two bluetooth errors occurred: one stated that the host was down and the other stated that the "software caused connection abort." These errors are indicative of a patched system or a system that is not affected by Blue Borne.

7.2 Dolphin

This attack exploited an audio vulnerability of virtual assistants such as Alexa, Siri, and Google Assistant. As stated earlier, virtual assistants use audio cues and commands. This attack leverages ultrasonic audio to create voice commands that are not audible to the human ear. However, these high frequency commands are still audible to the virtual assistant. This attack can do the following deeds: open a malicious website, engage in the act of spying, inject fake information, denial of service attacks, and conceal attacks [4].

7.3 ADB

This vector attempts to communicate with the Google Home via ADB (Android Debugging Bridge) and serves as a potential avenue of rooting or gaining escalated privileges at the operating system level. Typically this approach has had much success with Android OS for smartphones. The Google Home uses a modified Android OS, so two ADB vectors were tested. The first vector was over USB. Unfortunately the Google Home's micro USB port does not have proper drivers release by google and is most likely used for firmware updates or device recovery by Google techs. The device would be detected by windows only for a brief time before disappearing. This effectively thwarted any attempts to connect to the device over USB. Although there are theoretical approaches to using the micro USB port by using similar rooting methods to the Google Chromecast, since the hardware is almost identical. The other vector used, was ADB over TCP/IP. This vector uses the same debugging bridge but binds it to TCP port 5555 and attempts to open a connection to the device this way. After an Nmap scan, it was determined that there are only 4 ports open on the Google Home. 8008, which is used for unencrypted HTTP API calls, 8009, which is used for encrypted HTTPS API calls, 9000, which has an unknown purpose, and 10001, which also has an unknown purpose. ADB is successful in connecting to any of these ports, but does not establish an ADB connection. Since port 5555 is not open on the Google Home, this indicates that ADB is not running on the device itself and will render ADB connections useless.

7.4 HTTP and POST

Another vector that has a possibility to be exploited is using one of the only open ports on the home. Using port 8008 and utilizing HTTP POST, there are some possibilities to pass commands to the home. The extent as to what the Google OS will allow commands it not known and has not been tested.

7.5 Metasploit

The Google Home was tested using a basic scan using Rapid7's Metasploit. Performing this scan allowed us to test for various other exploits that we did not focus on. Initially an Nmap scan was performed on the Home and revealed four open ports. The initial discovery scan confirmed those findings as well as testing for which services were running on those ports. The services found were as follows, HTTP on port 8008, cslistener on port 9000, and scp-config on port 10001. An exploit scan was also performed without payloads and was returned with no exploits successful. The last attempt using Metasploit was a brute force attack. This yielded no results either.

8. ANALYSIS

Testing and data collection has reinforced our hypothesis that Google Home is only sending back information when triggered with its keywords. Based on the data, we can see that the number of packets being sent per second to Google differs based on the type or request it is. If it is a simple command or question that is being asked, a very slight jump can be seen. However, if it is a big request such as requesting the Home to play music, the amount of packets is noticeably different. There is no doubt that these kinds of requests require more data to be sent. Another thing that we observed was that while information is being sent back once the keyword has been said, for bigger requests, there are smaller spikes in packets even after the initial jump. For example, when asked to play music, since we only specified the genre instead of a specific song title, it began to play a playlist off a radio. On wireshark, we could see an enormous spike which was then followed by minor spikes as time went on. This shows that data is being sent to Google for requests that require Google to do continuous work.

9. CONCLUSION

The Google Home only transmits data to its servers when commands are made by the user. The entirety of the command is sent back to Google's server. Here the user's request is processed, and the appropriate response is sent back to the Google Home to relay back to the user. The Google server saves all requests made by the user on its servers. This even includes the voice recordings made by the user making the request. All requests made of Google Home are stored under 'my activity' history under the Google account tied to the Home. This history will remain indefinitely, unless the user specifically goes into their account and deletes their activity history. While the commands that were used for our testing purposes did not contain any private information, the same cannot be said about another user in their home setting. The Google account linked to the Home is very important and it is advised for that account to be kept safe and decrease the possibility of it being compromised.

REFERENCES

[1] Edwards, H. S. (2017). Alexa Takes the Stand: Listening Devices Raise Privacy Issues. Time, 189(18), 28-29.
[2] Pitsker, K. (2017). HOME SMART HOME. Kiplinger's Personal Finance, 71(10), 64-69.
[3] Rash, W. (2016). Security for Google's New Home Assistant May Get Lost on the IoT. Eweek, 12.
[4] Zhang G., Yan C., Ji X., Zhang T., Zhang T., Xu W. (2017) Dolphin Attack: Inaudible Voice Commands. ACM Conference on Computer and Communications Security (CCS), 1-15.
[5]Data security & privacy on Google Home. (n.d.). Retrieved November 29, 2017, from https://support.google.com/googlehome/answer/7072285?hl=en
[6]BlueBorne Cyber Threat Impacts Amazon Echo and Google Home, Retrieved February 14, 2018, from https://www.armis.com/blueborne-cyber-threat-impacts-amazon-echo-google-home/