Global Information Assurance Certification Paper

Privacy and the Internet of Things

GIAC (GSEC) Gold Certification

Author: Peter Milley, pbmilley@gmail.com
Advisor: Barbara Filkins
Accepted: October 21, 2017

Template Version September 2014

Abstract

The Internet of Things has gotten a lot of attention over the past year or so, and for good reason. From a security perspective, Internet-connected devices are easy targets, especially when they are not designed with security in mind. But, in addition to the concerns of botnets and DoS attacks, some newer devices also raise information privacy concerns.

Devices like Amazon's Echo and Google Home push the IoT envelope to new limits. Much like more traditional Web browsers, they potentially house a huge amount of information about us: our search habits, our networks, and our personal digital lives. All that is keeping this information private is the parent company's adherence to its usage policy. As consumers, we must trust that the company places a higher value on our privacy than on the potential profit to be made from sharing our information with others. New freedoms granted to Internet Service Providers by Congress, create the opportunity for our private information to become a much sought-after commodity.

In this paper, we will explore this issue in detail, and provide some much-needed answers to how consumers can protect themselves in this ever more unwelcoming landscape of technology versus privacy.

1. Introduction

This paper addresses Digital Assistants and the devices that use them. Digital assistants have existed in one form or another since the days of *Clippy*, the annoying little paperclip in Microsoft Word that would constantly interrupt your train of thought with what was supposed to be helpful advice but was more often just a distraction. They have come a long way since those early days, with much broader access and increased functionality, such as voice recognition.

Stand-alone devices like the Amazon Echo and Google Home have created a new playing field in the technology marketplace, bridging the gap between our computing devices and the Internet of Things (IoT), opening our connected homes to a world of information in a new and exciting way. But in our excitement over this new paradigm shift of information at our fingertips, have we overlooked privacy? Has our desire for the latest and greatest in technology and constant access to information caused us to become more and more complacent about how and where our personal information is gathered and used?

The goal of this paper is to delve into this subject and answer these questions, but first, let's begin by defining some terms.

- **Internet of Things (IoT):** Many modern devices, which for years were only expected to perform a simple task, are now getting "connected." The term encompasses all kinds of modern convenience electronics, from home A/V equipment to thermostats to kitchen appliances that can now be connected to home computer networks, and potentially the Internet ("The Internet of Things," n.d.).
- **Digital Assistants:** Most computerized electronics—desktop/laptop computers, tablets, and mobile phones—have some sort of application that helps to simplify our daily tasks. They find directions, look up recipes, perform web searches, take notes, and schedule meetings. They are tightly integrated into the operating system and connect to the web browser, email, and calendar apps for convenience (Kelly, 2015).

- **Home Assistant devices (sometimes known as smart speakers):** A standalone, Internet-connected device that brings the functionality of a digital assistant to a home or office setting. Amazon was the first in this market with the *Echo*—a device paired with the digital assistant Alexa. They have recently added the *Echo Show*, a device with video capabilities, to their lineup. Google is the other major player, with its *Google Home* device (Hardy, 2017). Apple is expected to release a Siri-powered device in December 2017 (Fox Rubin, 2017), and although Microsoft doesn't manufacture it, at least one smart speaker exists with Cortana integration (Pierce, 2017).

Convenience is the business driver leading appliance makers to connect their devices to the Internet. Technology has evolved to a point where we, as consumers, expect convenience from all of our products. We need the ability to turn lights on or off remotely, or save ourselves a trip by checking the connected refrigerator to see if we need milk while we are out, or unlock the front door to let in the repairman while we are at work, or peek in on Fido to see if he gets on the couch while we are away—all conveniences that, until very recently, were the stuff of science fiction.

This convenience comes at a price. The unfortunate reality is, the companies making these devices, although well steeped in the challenges of manufacturing physical products, are not as well versed in software development. They have had to quickly learn how to create software that enables the convenience functions we have come to expect in order to stay competitive in the marketplace. However, they do not have the years of experience that most large software firms have in creating secure code. Insecure shortcuts that major software vendors abandoned years ago are back, exposing regular problems with IoT devices. For example, toy maker Mattel recently released an Internet-connected doll with several security flaws. Among other issues, its associated app contained reusable authentication credentials and would automatically connect the user's mobile device to any WiFi network with "Barbie" in the name (Goodin, 2015).

Appliance makers create back-door access for support personnel or hard-coded passwords and encryption keys to simplify manufacturing and support with little regard for security. Furthermore, they rarely take into account the need for regular patch

maintenance and rely too heavily on the end user to make security changes to their products. Or worse, as evidenced recently by a manufacturer of smart door locks, break their own products during an update (Goodin, 2017; Spring, 2017).

The end result of these insecure practices is that IoT devices have become prime targets for the bad guys (Cheng, 2016; Greenberg, 2017; Ng, 2017). Black-hat hackers are extremely adept at finding and exploiting the weaknesses of these soft targets, using zero-day exploits as well as known vulnerabilities that are notoriously neglected in the world of IoT. In spite of the fact that most IoT devices contain relatively small amounts of computing power, an attacker can amass large numbers of them into powerful botnets with the strength to create distributed denial of service (DDoS) attacks against even large companies (Krebs, 2016).

So, why should this concern the average consumer? Well, in a word—privacy. Having one's thermostat used in a massive botnet attack against a major website probably won't cause more than a minor inconvenience for most home users, but allowing these attackers into their home network is much more serious. Think about that webcam, or the nifty new front door lock, or that top-of-the-line smart TV. Now imagine all of those things being controlled by some shady person on the Internet. These products can become a backdoor into your home computer network, where people keep their bank records, family photos and any amount of other private information. That all seems like a big risk for a little extra convenience.

To make matters worse, what if that random Internet stalker had access to your complete digital life—every Internet search, every shopping order placed, every email sent or received, and every Facebook post. This is what is at stake when digital assistants are exposed to the world-at-large. Digital Assistants have access to all of this information and more. They need this access to perform their intended role, but who is making sure the private information is kept private?

Somewhere in their lengthy End User License Agreements and Privacy Statements (yes, those things that no one reads before clicking the "I Agree" checkbox), companies describe how they plan to handle your personal information. Some of their terms may be surprising.

Peter Milley, pbmilley@gmail.com

As if those stakes weren't high enough, many of the digital assistants in modern day operating systems have the capability to be invoked with a special word or phrase (sometimes referred to as a "wake word" or "hotword"). This functionality is on by default for standalone devices like Echo and Google Home and is an option for handheld devices and personal computers. With this option enabled, these devices are listening at all times in order to react when called upon. But how much do they hear? Is what they hear recorded? If so, how long do these recordings hang around?

Again, we turn to the companies' customer agreements for answers, and again, you may find some of their stances surprising.

2. Digital Assistants and Privacy

Before we delve into the security of these devices, we should first get more familiar with how they work and how the companies treat the information they handle.

2.1. How do digital assistants work?

In order to effectively perform their primary function, digital assistants must gather and store relevant information. At the very minimum, this would include the question to be answered, but it may also include other details used to increase accuracy, or enhance the results in some way. This could include the asker's location, for instance, so that results can be narrowed to those in the immediate vicinity. One way to get this information is to prompt the user, but depending on the configuration, it could also be gleaned from the device's position information.

As mentioned earlier, most digital assistants are closely bound to the operating system for the device on which they run. Siri, for example, is a core function in iOS, with access to nearly every facet of an iPhone or iPad. Likewise, Cortana, which made its debut in Windows Phone 8.1, is tightly integrated into Microsoft's latest operating system, Windows 10. In fact, Cortana is *so* tightly integrated, that it can be difficult to disable without some serious gymnastics (Paul, 2016).

Home assistant devices, on the other hand, are built around the digital assistants at their core. Without the digital assistant, or with that functionality restricted, there would hardly be a reason to own the device. Their prime directive is one of convenience. They provide their owner with information, manage their calendars, and place orders for goods

and services over the Internet. Therefore, limiting their reach in favor of privacy runs counter to their intended purpose. But what if they overstep their bounds? How would we know, and what's the risk?

2.2. EULAs and Privacy Statements

All of the companies that sell home assistant devices address privacy concerns through their End User License Agreements (EULAs) or their Privacy Statements. These customer contracts often have a section where they discuss the extent to which the company will gather and store Personally Identifiable Information (PII) or other private information about their customers. Deep in the legalese of these sections lie the answers to how one can expect to have their data handled by the company, as well as some description of how that data may be used and for what purpose.

For instance, Google Home's data policy states: "First and foremost, we use data to make our services faster, smarter, and more useful to you, such as by providing better search results and timely traffic updates. … Also, on surfaces where we show ads, we use data to show you ads that are relevant and useful, and to keep our services free for everyone. Google Home learns over time to provide better and more personalized suggestions and answers." ("Data security & privacy on Google Home - Google Home Help," 2017)

This is an interesting statement, with several implications. Let us break it down. First, this being the data use policy for the Google Home device, which has no screen on which to "show ads," the company is implying a connection to other Google products and services like the Chrome browser and their search engine—included in the "our services" part.

Second, it "learns over time," which of course means that the more information it collects and stores, the better it can serve you in all of your Google needs. This type of information is called "metadata" and is generally not considered private. However, when many data points are aggregated, it can reveal a lot about an individual user and their online habits. From the company's perspective, the more data they have about their customers, the better user experience they can provide, but over time this metadata becomes a digital fingerprint that can potentially be used to identify unique individuals.

Furthermore, in the general Google Privacy policy, it says: "We may share non-personally identifiable information publicly and with our partners – like publishers, advertisers or connected sites." ("Privacy Policy – Privacy & Terms – Google," 2017) Have you ever wondered why that item you were just researching shows up as an ad on Facebook a few minutes later? This statement may provide a clue. In a later section, we will discuss what options the user may have with regard to these terms.

2.3. Legal rights

Though the EULAs and privacy statements address some concerns by putting in writing how a company intends to deal with your personal information, what may be of a greater concern are the issues that *are not* covered in these contracts. Because these devices are so new, their legal boundaries have not been fully vetted in the courts. Consequently, the legal precedents surrounding their handling with regard to privacy have yet to be clearly defined.

There are court cases pending that could be instrumental in setting these precedents for all future litigation. For example, a case currently in the courts features an Amazon Echo that was present at a murder scene, collected for the evidence it may provide in convicting the guilty party (Sweetland Edwards, 2017). It is yet to be seen what kind of evidence can be gleaned from the device, so this case, in particular, could prove to be groundbreaking.

The laws as they are currently written do not favor privacy when it comes to these gadgets. Although Americans enjoy a constitutionally guaranteed expectation of privacy in their homes under the Fourth Amendment, it does not extend to cases where information is shared with a third party. This "third-party doctrine" is what allows law enforcement to gather call records from the phone company without a warrant, for example, because metadata is not considered private. In fact, this doctrine has been challenged many times since it was first used for a pair of Supreme Court cases in the 1970s (Supreme Court of the United States, 1976; Supreme Court of the United States, 1979).

However, as Villasenor (2013, para. 6) noted in his article on the subject: "Much has changed since the 1970s. Today, being engaged in the world involves using the Internet, mobile phones, apps, cloud-based services, GPS, and other technologies that

leave enormous amounts of information in the hands of third parties. Even before the NSA documents leaked by Edward Snowden started appearing on the home pages of the world's news sites, there was a robust discussion about the continued suitability of the third party doctrine."

Following this interpretation, when an individual voluntarily installs a listening device, such as the one present in an Echo or Google Home, they are effectively waiving their Fourth Amendment right under this third-party rule (Baral, 2016; Sweetland Edwards, 2017). Some good news: the issue has reached the level of the U.S. Supreme Court, specifically the case *Carpenter v. United States* in which the legality of gathering mobile phone location data without a warrant is addressed (Kravets, 2017). This case is certain to bring the discussion of the third party doctrine again to the forefront.

While these limits are being explored in the U.S., there is likely to be even more confusion when it comes to these devices being used abroad. Many foreign countries, especially those that are members of the European Union (EU), have very strict privacy laws that make such situations more difficult to interpret. In fact, dealing with privacy issues in the EU will be even more restrictive when the General Data Protection Regulation (GDPR) goes into effect next year. Companies that do business in EU countries face substantial fines for non-compliance with the new rules once they are implemented ("Home Page of EU GDPR," 2017).

As a further example, see the following excerpt from the *Privacy Shield Framework*'s website, outlining the requirements for participation by U.S. businesses ("Privacy Shield | Privacy Shield," n.d.). The *Framework* exists explicitly to help U.S. businesses understand the requirements of handling personal data from the European Union and Switzerland.

> "The Privacy Shield Principles comprise a set of seven commonly recognized privacy principles combined with 16 equally binding supplemental principles, which explain and augment the first seven. Collectively, these 23 Privacy Shield Principles lay out a set of requirements governing participating organizations' use and treatment of personal data received from the EU under the Framework as well as the access and recourse mechanisms that participants must provide to individuals in the EU. Once an organization publicly commits to comply with the Privacy Shield Principles, that

commitment is enforceable under U.S. law." ("Requirements of Participation | Privacy Shield," n.d.)

2.4. What about security? Can they be hacked?

While the EULA and privacy agreements offer some protection for the end user, these contracts are only effective when the devices are being used lawfully and for their intended purpose. They won't offer much protection against evil-doers intent on exploiting their flaws and misusing the technology for nefarious purposes.

There is no evidence that any home assistant devices have been hacked yet—at least not directly. However, any piece of technology connected to the Internet is certainly at a higher risk than those that are air-gapped from the world at large. Furthermore, IoT devices are attractive targets for hackers, whether as a toehold into home networks, as attack devices (in botnets), or to be exploited in other ways.

But IoT devices may be vulnerable even without being attacked directly. Most Internet-enabled IoT devices are controlled, completely or in part, by applications that run in the public cloud. This makes them vulnerable without an attacker ever compromising the user's device directly. In reality, the device is only as secure as the online account that controls it. If an attacker compromises the device's cloud account, they can potentially control the device without the owner's permission or even their knowledge (see figure 1). Therefore, it is critical that this account is well secured.

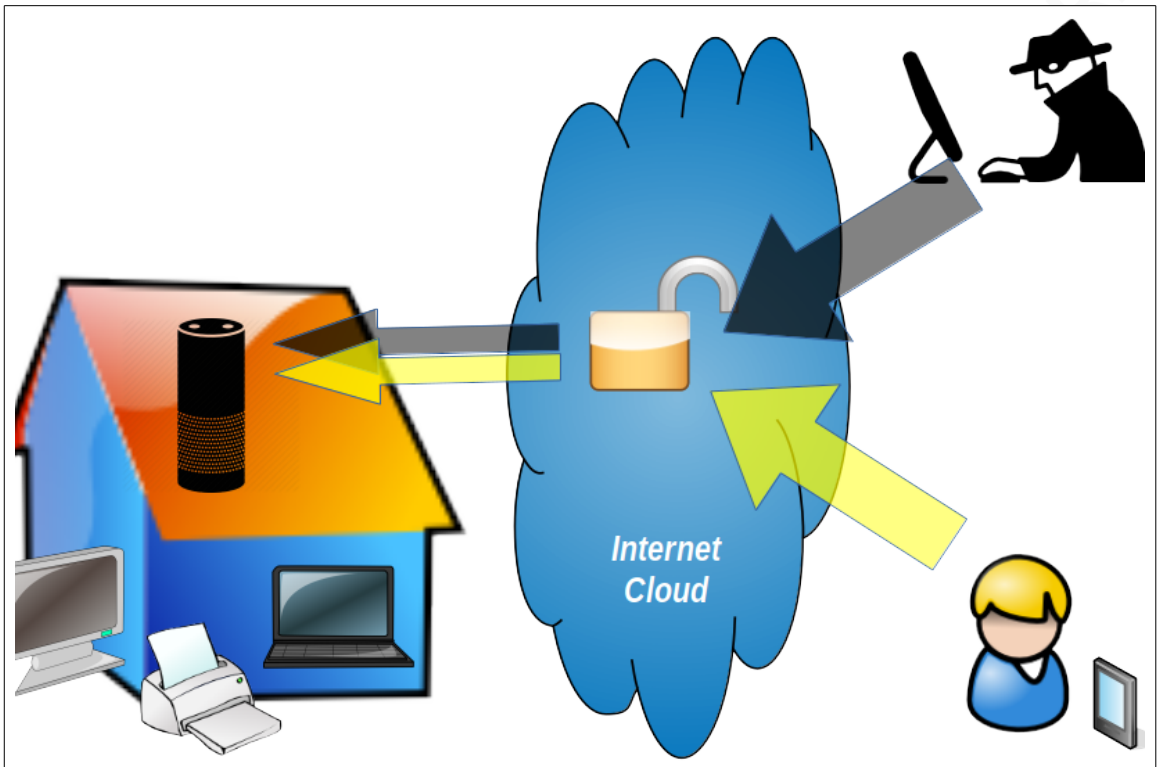

Figure 1: Cloud account

3. How to Take Action

The increasing popularity of these devices suggests that the average consumer is interested in functionality and novelty and may be unaware of any privacy concerns. In fact, most consumers seem resigned to a state of acceptance in the belief that someone is always watching. Perhaps they feel they needn't be concerned if they have nothing to hide. Still, others may think that properly securing their devices is simply too much work.

While this attitude may not be explicitly encouraged by IoT manufacturers, they have not gone out of their way to change it. They continue to add features to their products while letting flaws go unpatched. The Echo, for instance, now has a feature called *Drop In*, designed to "...let you easily connect with your closest friends and

family" ("Amazon.com Help: Alexa and Alexa Device FAQs," n.d.). This feature allows Echo users to eavesdrop on other Echo owners with little or no warning, depending on the configuration. Imagine this functionality in the wrong hands.

Manufacturers attitudes may start to change if more of them are held accountable for insecure practices, as the FTC is currently doing with D-Link (Ian Wong, 2017). But, until the laws begin to catch up with the technology, don't expect to look to the Federal Government for protection.

Consumers should be taking action to protect themselves, and not assume that the manufacturers have their best interests in mind. But what, as a consumer, can an individual do?

First, and foremost, all consumers should seriously consider whether the convenience of having one of these devices outweighs the risks of introducing it into their home networks. Yes, these devices are undoubtedly 'cool,' but until the implications of owning what is essentially a spy device are better understood and controlled, it may not be worth the risk.

If you already own a device or have determined it to be an indispensable addition to the home, here are some precautions that should be taken.

1. Disable the microphone or camera whenever possible. Each of the devices listed above has the ability to temporarily disable their listening modes. Familiarize yourself with the controls and how to operate them (McClelland, 2017).
2. Review and manage your history. This can be done either through the control app or via a special page in the online account associated with the device. This will show all of the searches performed and the questions asked of the device, and must be cleaned up and managed manually (McClelland, 2017).
3. Secure the online account! This is the Internet accessible gateway to the device. Make sure it's as secure as possible, including the use of two-factor or two-step authentication.
4. Change the "wake word." As of this writing, the wake words for Google Home and Echo are limited to just a few, but as fans of the television show *South Park* recently discovered, this change can be essential (Lynch, 2017). It could help to

thwart the casual guest, passer-by, or television program from controlling the device without your permission (McClelland, 2017).

5. Enable purchase locks (McClelland, 2017). This is especially important for Alexa and Siri because they are both tied directly to online stores where they can be used to make purchases.
6. Read the EULA and Privacy Statements. Yes, this can be a tedious task, but there are likely to be some features or services that are enabled by default and are only disabled by the user explicitly requesting it—known as "opting out." The best way to discover these opt-out features is to read the relevant agreements.

4. Conclusion

Technology is an ever-changing balance of benefit and burden, a double-edged sword that often leaves the consumer on the bleeding edge. We all know that technology moves quickly, continually proving Moore's law is alive and well (Moore, 1975). However, sometimes it moves so rapidly that it can be hard to discern which is the burden and which is the benefit, and the consumer is left to guess. The unfortunate result is one in which the consumer inevitably chooses convenience and novelty over privacy and safety.

In a perfect world, this responsibility would not be left to the consumer. Government regulations should protect end users with appropriate controls placed on the manufacturers to ensure product safety prior to their release. But government policy takes time, and laws must be vetted by the courts to be effective.

The onus to protect the consumer must be placed on the manufacturers, and they need to put as much, if not more, emphasis on designing and producing safe, secure, and trust-worthy products as they do feature enhancements. Until then, consumers must look out for themselves and learn how to stay safe in an ever more connected, and less private, world.

Peter Milley, pbmilley@gmail.com

References

Amazon.com Help: Alexa and Alexa Device FAQs. (n.d.). Retrieved September 22, 2017, from https://www.amazon.com/gp/help/customer/display.html?nodeId=201602230

Baral, S. (2016, December 20). Amazon Echo Privacy: Is Alexa listening to everything you say? *Mic*. Retrieved from https://mic.com/articles/162865/amazon-echo-privacy-is-alexa-listening-to-everything-you-say#.T0HKA44dX

Cheng, R. (2016, January 12). What lurks beneath the Internet of Things hype? Nagging security fears. Retrieved from https://www.cnet.com/news/what-lies-beneath-the-internet-of-things-hype-an-undercurrent-of-security-fears/

Data security & privacy on Google Home - Google Home Help. (2017). Retrieved from https://support.google.com/googlehome/answer/7072285?hl=en&ref_topic=7173611

Fox Rubin, B. (2017, June 5). Apple's new HomePod smart speaker brings Siri home. Retrieved from https://www.cnet.com/news/apples-new-homepod-smart-speaker-brings-siri-home/

Goodin, D. (2015, December 4). Internet-connected Hello Barbie doll gets bitten by nasty POODLE crypto bug. Retrieved from https://arstechnica.com/information-technology/2015/12/internet-connected-hello-barbie-doll-gets-bitten-by-nasty-poodle-crypto-bug/

Goodin, D. (2017, August 14). Update gone wrong leaves 500 smart locks inoperable. Retrieved from https://arstechnica.com/information-technology/2017/08/500-smart-locks-arent-so-smart-anymore-thanks-to-botched-update/

Greenberg, A. (2017, July 18). Hack brief: 'Devil's Ivy' vulnerability could affect millions of IOT devices. *Wired*. Retrieved from https://www.wired.com/story/devils-ivy-iot-vulnerability/

Hardy, J. (2017, June 8). Amazon Echo vs. Google Home: Which home assistant is best for you? [Web log post] Retrieved from https://www.affinitytechpartners.com/3n1blog/2017/6/8/amazon-echo-vs-google-home-which-home-assistant-is-best-for-you

Home Page of EU GDPR. (2017). Retrieved October 17, 2017, from http://www.eugdpr.org/eugdpr.org.html

Peter Milley, pbmilley@gmail.com

Ian Wong, J. (2017, January 6). FTC lawsuit against D-Link for shoddy security practices is good news for the Internet of Things — Quartz. Retrieved from https://qz.com/879852/ftc-lawsuit-against-d-link-is-good-news-for-the-internet-of-things/

The Internet of Things. (n.d.). Retrieved from ETH Zurich website: www.vs.inf.ethz.ch/res/show.html?what=iot

Kelly, H. (2015, July 28). Which is the best digital assistant: Siri, Cortana, Alexa or Google Now? Retrieved from http://money.cnn.com/2015/07/28/technology/digital-assistant-interview/index.html

Kravets, D. (2017, October 2). Supreme Court's new term: Surveillance, hacking, sports betting—and cake, too. Retrieved from https://arstechnica.com/tech-policy/2017/10/supreme-courts-new-term-surveillance-hacking-sports-betting-and-cake-too/

Krebs, B. (2016, September 16). KrebsOnSecurity Hit With Record DDoS [Web log post]. Retrieved from https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/

Lynch, J. (2017, September 14). The 'South Park' premiere set off a lot of fans' Alexa and Google Home devices with hilarious phrases. *Business Insider*. Retrieved from http://www.businessinsider.com/south-park-premiere-set-off-a-lot-of-alexa-and-google-home-devices-2017-9

McClelland, D. (2017, January 17). How to secure your Amazon Echo. Retrieved from http://www.techradar.com/how-to/how-to-secure-your-amazon-echo

Moore, G. (1975). *Progress in digital integrated electronics* [PDF document]. Retrieved from http://www.eng.auburn.edu/~agrawvd/COURSE/E7770_Spr07/READ/Gordon_Moore_1975_Speech.pdf

Ng, A. (2017, July 19). Net-connected devices open to hacks from widespread bug. Retrieved from https://www.cnet.com/news/iot-devices-hack-bug-vulnerability-devil-ivy-exploit/

Paul, I. (2016, July 26). You can't turn off Cortana in the Windows 10 Anniversary Update. *PC World*. Retrieved from

https://www.pcworld.com/article/3100358/windows/you-cant-turn-off-cortana-in-the-windows-10-anniversary-update.html

Pierce, D. (2017, May 8). The Invoke smart speaker brings Microsoft's AI to your living room. *Wired*. Retrieved from https://www.wired.com/2017/05/invoke-smart-speaker-brings-microsofts-cortana-ai-living-room/

Privacy Policy – Privacy & Terms – Google. (2017, April 17). Retrieved from https://www.google.com/policies/privacy/

Privacy Shield | Privacy Shield. (n.d.). Retrieved September 26, 2017, from https://www.privacyshield.gov/welcome

Requirements of Participation | Privacy Shield. (n.d.). Retrieved September 26, 2017, from https://www.privacyshield.gov/article?id=Requirements-of-Participation

Spring, T. (2017, August 14). Smart Locks Bricked by Bad Update. Retrieved from https://threatpost.com/smart-locks-bricked-by-bad-update/127427/

Supreme Court of the United States. (1976). *United States v. Miller, 425 US 435 - Supreme Court 1976.* Retrieved from https://scholar.google.com/scholar_case?case=15052729295643479698&hl=en&as_sdt=40000006

Supreme Court of the United States. (1979). *Smith v. Maryland, 442 US 735 - Supreme Court 1979.* Retrieved from https://scholar.google.com/scholar_case?case=3033726127475530815

Sweetland Edwards, H. (2017, May 4). Alexa Takes the Stand: Listening Devices Raise Privacy Issues. *Time*. Retrieved from http://time.com/4766611/alexa-takes-the-stand-listening-devices-raise-privacy-issues/

Villasenor, J. (2013, December 30). What You Need to Know about the Third-Party Doctrine - The Atlantic. Retrieved from https://www.theatlantic.com/technology/archive/2013/12/what-you-need-to-know-about-the-third-party-doctrine/282721/