Internet of Things (IoT) Mobility Forensics

Roger Baker and Chase Miller
{rjbaker, cnmiller}@cmu.edu
Carnegie Mellon University

***Abstract*— In this paper, we examine the challenges associated with conducting forensic investigations with Internet of Things (IoT) devices. The purpose of this paper is to be a catalyst in preparing investigators and researchers into thinking about solving IoT Forensics problems they may come in contact with before they need to. To help mitigate these problems, we introduce a framework to aid those interested in the forensic value of IoT devices. Our process, a precursor to live investigations, ensures that investigators get the best preparation before faced with such problems or roadblocks. Establishing a standards of procedures for each alike IoT device also helps investigations. Developing a sort of community driven database of IoT device forensic interest datasheets and writeups is the optimal solution at this time.**

I. INTRODUCTION

The Internet of Things (IoT) introduces a new level of connectivity and convenience to our everyday lives. Vehicles, refrigerators, televisions, lights, and more all make up some of the billions of devices that are now connected to the Internet. As these devices become more pervasive, they play a larger role in our lives. Each object potentially contains data of forensic interest. A programmable thermostat, for example, could log information and help firefighters investigate details about a house fire. Security sensors can help shed light on a break in for investigators. These are the obvious cases in which this extra knowledge means a lot to different stakeholders, however, many more scenarios can be imagined when considering the breadth of devices that exist.

The problem these devices pose, though, is evident when considering their characteristics. Engineered to be small, affordable and lightweight, design considerations are made to meet consumer needs. Recently, this market has seen a significant shift from devices having removable media (like microSD cards) to hosting the storage internally on an integrated chip. This makes doing forensics a bit more complex than traditional forensics. These integrated chips typically use some low-level communication protocol like I2C, SPI, or CAN. While there are proven methods in reading these integrated chips, the problem is not solved. New storage devices, custom file formats, and non-standard file systems make this problem even more difficult. Combine the possibility that the data can be encrypted or obfuscated while at rest (in storage), and the problem becomes even more challenging.

Another problem within this space is the ability to have discovered content admissible in court. Without a methodical plan of acquiring forensically interesting data, it can be difficult to convince a judge and/or jury that a.) data has been collected in a non-compromising way and b.) that evidence supports and is warranted for the investigation. Without prior research into the exact IoT device, or something very similar, its hard to get to that level within the time frame of an investigation.

Developments within this field happen, however, few would be deemed groundbreaking. There exists few periodicals dedicated to the digital forensics field. Most notably are IEEEs Transactions on Information Forensics and Security and Elseviers Digital Investigation. These journals offer a variety of advancements in the field, with varying specifications.

As it relates to the subject at hand, Hegarty's "Digital Evidence Challenges in the Internet of Things" identifies the problems the pervasiveness of IoT has manifested.[1] This paper does a great job in exposing the gap previously discussed, as well as reiterate general procedures in handling unfamiliar devices. What this paper inherently lacks, however, is a detailed explanation in handling specific devices.

II. METHODOLOGY

Forensic investigations of Internet of Things (IoT) devices are a little-documented branch of digital forensics. In 2013, Oriwoh [2] published a landmark paper detailing the challenges and approaches associated with IoT forensics. Although this paper does propose two approaches to IoT forensics, it does not break these

approaches down in a way that would be easily reproducible by any forensic investigator. As a result of this, combined with the challenges we faced when conducting our own IoT forensic investigation of the iRobot Roomba 980, we propose the following, step-by-step process for conducting a forensic investigation on an IoT device. For simplicity, we have separated the host-based and network forensic processes.

A. Host-Based Forensic Process

We begin the process of identifying necessary steps in working with the IoT device with simple intelligence gathering, and with enough knowledge maturity, we begin to experiment with the device and components independently. Keep in mind that this framework should be adapted in combination with any investigator's standards of procedures. A lot of the work and research from traditional forensics can be applied here, high-level at a minimum.

1) Profile the Device: The first step in conducting a forensic investigation on an Internet of Things device is to positively identify the device that you are working with. Unlike traditional digital forensics, which is often well defined and typically conducted on well-documented devices with standard operating procedures (SOPs), investigations dealing with IoT devices are not always obvious. For example, a hidden "nanny cam" can take form of a desk clock, stuffed animal or any other common object. Positively identifying the device before beginning the investigation can significantly aid the investigative process. Identification should include the make, model, and version number of the device that is being investigated. A brief understanding of the purpose and intention of the device is also necessary at this stage.

2) Collect Device Documentation: Before beginning the investigation, research should be conducted to identify any available documentation on the IoT device. Documents such as schematics or manuals can provide investigators with valuable information such as what type of storage the device contains and where on the device you would find it. In some cases, manufacturers require you to be an authorized repair specialist to receive such documentation. This is typically a quick registration with the company in exchange for potentially valuable documentation.

Also not to be overlooked is the plethora of unauthorized or unofficial information available on the internet. This includes hobby sites, blogs, and community forums who share the same passion and interest. While researchers should always exercise caution when utilizing unofficial advice, this content is often helpful in that you can see what others have done, where they failed and what they could have done better.

3) Identify Storage Media: Depending on the fruitfulness of the documentation collection from Step 2, the forensic investigator should now have a general idea of what, if any, storage media exists on the device and how to access it. If device schematics were not available, the device should be carefully disassembled and inspected for any storage media. The key here is identifying storage media that is non-volatile, that is, data integrity is maintained between power cycles. Tradtional storage media is removable, such as microSD cards or similar memory sticks. As design considerations change, storage media is more often seen as integrated chips (ICs), embedded on the surface of a printed circuit board (PCB). This adds another level of complexity that we'll address later.

If no non-volatile storage is identified through the disassembly process, it is likely that the device either does not store data, or it stores data off-site - either with a companion mobile device application or via cloud storage. If that is suspected, proceed to the IoT Network Forensic Process.

4) Media Acquisition: Getting the data from the storage device can be tricky. For traditional media devices, this step is fairly trivial and mimics traditional forensics, using a write-blocking device reader. In the case of an integrated chip storing the data, some success has been made in retrieving this data.

Most ICs use some form of low-level communication protocol like I2C, SPI or CAN to transfer data to and from the main CPU on the board. These protocols operate in a master-slave fashion, with the storage device being the slave. In order to retrieve the data, performing a process called "chip-off" allows the investigator to desolder the IC from the PCB and place it in their own circuit. This allows the investigator to control the flow of data, and act as the master. There exists a good amount of hardware tools out there to help facilitate this process, as well as robust software to control the hardware [4] [3].

Once the storage media is identified and extracted if necessary, a forensically-sound image of all storage devices should be taken. Taking this exact copy of the devices storage allows investigators to analyze it

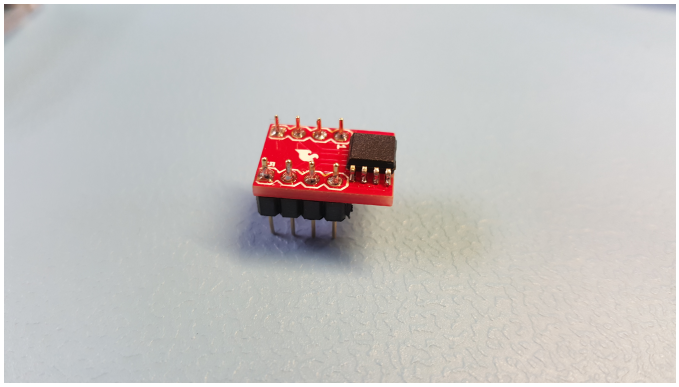

Fig. 1: IC is removed and place on breakout board

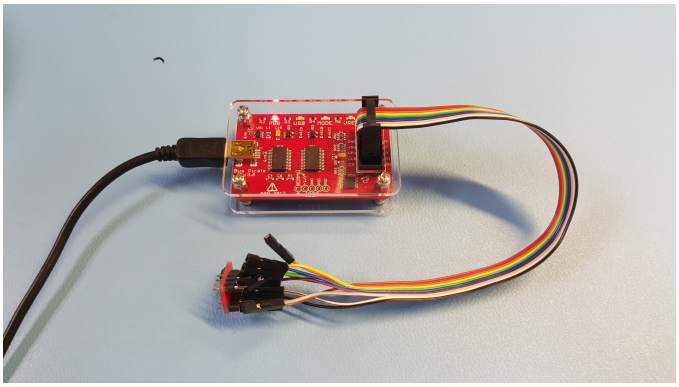

Fig. 3: IC on breakout board connected to Bus Pirate

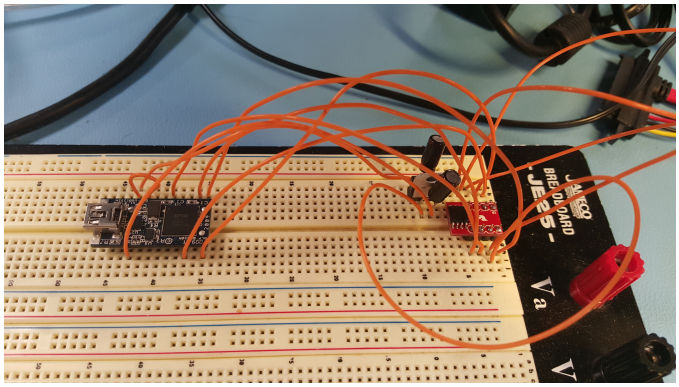

Fig. 2: Creating custom circuit on breadboard

without the risk of damaging the original data or even the device itself. A cryptographic hash (typically SHA256 or MD5) should also be calculated to identify possible changes to the data. For added insurance, using multiple acquisition tools can help guarantee you've captured correct data. This is illustrated in Figure 2 and Figure 3, using two different hardware tools.

5) Baseline Image Acquisition: If the device that is being investigated is brand-new, a baseline image should be taken before the device is used. This baseline image will help forensic investigators determine what data changes after use, indicating data with potential forensic value. However, forensic investigations are often conducted on devices that have already been used. In this case, investigators may need to acquire a new device, identical to the one being investigated, to obtain the baseline image.

6) Examine Artifacts on Connected Devices: Often, Internet of Things devices are connected to smartphones, computers, or other IoT devices. Perhaps the most common reason that an IoT device is connected to another device is for management purposes. Due to their form factor, most IoT devices do not contain onboard inputs. Instead, pairing the device with a smartphone and a mobile application allows users to easily interact with their device. However, this data that is exchanged between IoT devices and smartphones creates artifacts that can be of benefit to forensic investigators.

B. Network Forensic Process

Due to their proposed seamless integration into our everyday lives, the vast majority of Internet of Things devices communicate over wireless mediums. As a result, this section will focus primarily on wireless network forensic techniques.

1) Analyze Traffic During Setup Process: Many devices require a network connection to complete the setup process. Analyzing the traffic during this setup process can reveal much about the potential for forensically-interesting data. The setup process of our exemplar, the iRobot Roomba 980, is a great example of this. It first requires a series of buttons to be pressed on the device itself. Once these buttons are pressed, the device creates its own ad-hoc wireless network to which its corresponding mobile application, the "iRobot Home" application connects via WiFi. Once connected, the application instructs the user to enter their desired access point's SSID and password. Future communication is then handled through the specified access point. This setup process indicates that both the device and mobile application have the potential to contain forensically-valuable data, including the user's wireless access point credentials.

2) Analyze Traffic During Normal Operation: Analyzing traffic during the Internet of Things devices normal operation can reveal additional information that may benefit a forensic investigation. One example of this would be the device reaching out to a remote server.

If the device is continually transmitting data to a remote server during its normal operation, it is possible that data from the device is being stored remotely. In the case of a criminal forensic investigation, investigators may attempt to obtain and execute a search warrant for the identified server. It is also important to identify if and what type of encryption is being used on the transmitted data. If no encryption is being used, forensic artifacts may be able to be extracted from the data stream.

3) Analyze Traffic During Interaction with Other Devices: Analyzing traffic while the Internet of Things device is communicating with other devices, whether it be a smartphone, computer, or other IoT device, can reveal what type of forensic artifacts, if any, may be stored on the other devices. For example, some IoT devices may leverage the storage of other devices if their form factor does not allow for sufficient onboard storage.

III. Future Work

While conducting our research, we identified several areas where continued work could make substantial impacts on the Internet of Things forensics community. The first of these areas is to take the process detailed above in the methodology section and apply it to additional IoT devices. Much of our proposed process was constructed around our own experiences while working with the iRobot Roomba 980, one specific IoT device. It is highly likely that as our proposed forensic process is applied to additional devices, it will need to be tailored and expanded. When creating this process, we intentionally left it flexible enough to allow for such tailoring.

Additionally, we plan to open-source all of our data from this project including technical findings and our forensic methods and techniques. To facilitate this sharing, we are proposing an open-source, wiki-style database of IoT devices, with a focus on the device's security and forensic properties. Through this database, independent researchers can share their findings regarding individual IoT devices. We envision this database merging security and forensic data (e.g. encryption type, storage media, base images, etc.) with relevant device documentation (e.g. manuals, model/firmware numbers, schematics, etc.). Our goal is to have this database serve as the authoritative resource for forensic investigators during the reconnaissance phases of our proposed IoT forensics processes.

IV. Conclusions

As new, digital technologies come to market, the field of digital forensics continues to rapidly evolve. This is particularly true when it comes to the Internet of Things. In order to address this ever-changing device landscape, IoT forensic investigators need to introduce flexibility to their traditional digital forensic techniques.

Although there are significant differences between traditional digital forensics and IoT forensics, there are many similarities as well. We believe that the forensic processes outlined in this paper provides investigators with the flexibility they need to obtain data from new, unfamiliar devices, while still maintaining the general forensic structure which they are used to.

Acknowledgment

The authors would like to thank the INSuRE program, Purdue University's PURR, our Technical Director, Al Holt, our professor, Dr. Patrick Tague, and the support of the Information Networking Institute for the funding provided for the Roomba 980.

References

[1] Hegarty, R. C., D. J. Lamb, and A. Attwood. "Digital Evidence Challenges in the Internet of Things." Proceedings of the Tenth International Network Conference (INC 2014). Lulu.com, 2014.

[2] Oriwoh, E.; Jazani, D.; Epiphaniou, G.; Sant, P., "Internet of Things Forensics: Challenges and approaches," in Collaborative Computing: Networking, Applications and Worksharing (Collaboratecom), 2013 9th International Conference Conference on , vol., no., pp.608-615, 20-23 Oct. 2013

[3] flashrom. Flashrom Wiki. 1 March 2015. Web. http://www.flashrom.org/Flashrom.

[4] Bus Pirate - DP. Dangerous Prototypes. Web. http://dangerousprototypes.com/docs/Bus_Pirate.