IoT Forensics- Finding reliable sources of Evidence from Amazon Echo

Adhirath Kapoor
adhirathlrkapoor@gmail.com
akap213@aucklanduni.ac.nz

IoT Forensics- Finding reliable sources of Evidence from Amazon Echo

Abstract

The IoT ecosystem is encompassing new devices at a constant rate. And now there are around twenty billion of these devices on the planet. One category of IoT devices, known as Home Assistants, find themselves being used in a lot of households. Amazon Echo, being the most popular device in this category also has a great potential for storing valuable evidence when a crime occurs. The biggest problem lies in retrieval of this evidence. This report deeply compares and contrasts the methodologies and findings from three major works entirely focused on forensics of Amazon Echo and tries to find out the evidentiary value of data retrieved and how this could prove relevant for a crime investigation.

1. Introduction

With a constant advancement in the field of technology, the phrase "Ease of Access" has been getting ameliorated every now and then. This amelioration has been mostly possible due to the advent of Internet of Things. Its incorporation in our daily lives has reduced manual efforts and thus can be assimilated into a lot of domains, ranging from the healthcare sector to the battlefield and warfare combat systems. One class of IoT devices where they act like a virtual butler by controlling other devices at home and acting like an intermediator between the users and devices is that of home assistants. Amazon Echo and Google Home are popular examples of such assistants. In addition to the operations they are designed for, they also find themselves being used in the field of Digital Forensics. The interactions between a user and their personal assistant might prove useful for an investigation. Home assistants' association with their users are established through a mobile companion application accessed through a smartphone or a web browser. This report deeply looks into the forensics of Amazon Echo through analysis of three articles which primarily deal with Amazon Echo and the Alexa Ecosystem, which is the virtual assistant for Echo and tries to answer one major question –

How can reliable sources of evidence be found from Amazon Echo during an investigation and prove useful for a criminal case?

2. Background of IoT Forensics

Digital Forensics, in order to combat crimes, follows certain procedures that are compliant with the law of that region. There is no certain guideline on how the procedures will be carried out but there are two very famous frameworks in the Forensics literature that are most commonly

used. These were provided by (Mckemmish, 1999) and (NIST, 2006) respectively. These frameworks can be applied to conventional digital devices because in order to carry out the investigations, the investigators construct major evidence only from those particular devices rather than their associations which might be other devices on the network or some peripherals they are connected to.

Since crime investigations these days are not just limited to conventional devices, it is important to understand if these conventional frameworks are reliable while investigating IoT devices and their associations. As described by (Zawoad & Hasan, 2015), IoT Forensics, in no way, is same as the conventional Digital Forensics.

The IoT Forensics comprises of three different classes of Digital Forensics. These classes are Device level Forensics, Network level Forensics and Cloud Forensics. Device level Forensics focuses on data extraction from the memory of IoT devices. Network level Forensics utilizes the logs which can give information about the devices connected to a network. Cloud level Forensics involves examination of the cloud-based data either through user credentials or cooperation from Cloud Providers.

Personal Assistants like Amazon Echo, when being investigated, follow similar footsteps of the investigation process described by (Zawoad & Hasan, 2015). There are three major works in the field of forensics literature which focus entirely on forensics of Amazon Echo, the device and its virtual assistant Alexa. The article written by (Chung, Park, & Lee, 2017) utilized a combination of Cloud and Network level forensics to identify APIs for the Alexa ecosystem that would find data with the most evidentiary value. The next work by (Orr & Sanchez, 2018) focuses on Echo's capabilities and how they can be utilized in the field of forensics. The authors then proceeded to extract evidentiary data from an Echo using the companion app and the mobile device cache. The latest article in this domain was by (Li, Choo, Sun, Buchanan, & Cao, 2019) who proposed a general framework for IoT devices based on the conventional digital forensics methodology and used Amazon Echo as a use case to test their framework. They tested Echo at all the three levels of forensics to determine reliable sources of evidence.

3. Examination of Echo

Echo's examination is carried out in three areas; at device level, in the network and at cloud level (Zawoad & Hasan, 2015). This is in coherence with the Zone 1-2-3 approach proposed by (Oriwoh, Jazani, Epiphaniou, & Sant, 2013) where the investigation is carried out in three zones, namely hardware, middleware and outside networks. The examination processes laid

out by researchers seem quite plausible in theory but vary greatly in practice from what is on the paper.

3.1 Device Examination

This level of forensics is similar to carrying out conventional forensics on a device. The main difference here is that the device being examined, Amazon Echo is very constrained in terms of memory and security. Hardware analysis of Echo is difficult as any physical access to its compact configuration might lead to destruction of data and the fact that court of law requires evidence in its original untampered state makes it a complex task. (Chung et al., 2017) mentioned this approach being adopted in the papers they cited and stated that such analysis might lead to an internal port access but this approach may not yield evidentiary data. (Orr & Sanchez, 2018) stated that this approach has more cons than pros and heavily disregarded it. (Li et al., 2019) examined the Echo and were able to find some configuration data and settings of Echo through UART port access. They also performed analysis of firmware of Echo through EnCase and this gave them a disc image of Echo. From the experiments conducted by all the three works at device level, it can be clearly inferred that this approach needs to be used when access to the companion networks and cloud is missing. The data of evidentiary value resulting from this level of forensics is very limited and might not be preservable if the process goes wrong.

3.2 Network Examination

Through this approach, investigation is carried out on the network, the device is connected to. For an Echo, unless the user credentials are known, traffic analysis will only provide encrypted traffic flows. (Chung et al., 2017) used this approach to identify cloud native artefacts in the form of APIs. They analysed the traffic through Charles Debugging Proxy. However, at this stage, the APIs were inaccessible because the traffic was encrypted.

(Orr & Sanchez, 2018) didn't perform network level forensics in their study. (Li et al., 2019) utilized Zenmap to locate IP and MAC addresses of Echo and its associations. They stated that for more complex cases, port pings are more efficient than Zenmap.

3.3 Examination at Cloud Level

Examination at cloud level is the most important phase during IoT forensics. This is due to the fact that cloud stores most of the data for an IoT device. But cloud level forensics is not as easy as it seems. Cloud providers tend to keep customers' data secure and cooperation from them is

difficult as they would argue that this would lead to leakage of private data. However, for an Echo, all the three works attempted to use this level of forensics to get the most evidentiary data from it. (Chung et al., 2017) split this approach into two subtypes – examination of the cloud through first accessing the user credentials and then examination of the client-side companion application to understand client behaviour. (Orr & Sanchez, 2018) also followed a similar approach by examining the companion application and the web browser but they didn't use any forensic tools to perform this analysis as they had knowledge of the user credentials. They however worked with Blacklight, a forensic tool in order to access the mobile device cache as this would help the investigators find evidence in the case, user credentials were not known. (Li et al., 2019) gained access into the cloud through the Companion application and the Alexa Voice Server. Latter's access by the authors was a slightly different approach than the previous works. This access was established by first investigating the firmware, which is a part of device level forensics and then getting credentials like the Product ID and Client ID.

In this section, it was observed that all the three works used different approaches while performing Forensics on Amazon Echo. After reviewing all three, it can be inferred that the most crucial and heavily adopted approach in Echo Forensics is the cloud level examination. The authors were able to examine the cloud only when they had the user credentials. In absence of user credentials, this would be tedious, but (Orr & Sanchez, 2018) examined the cache of the smartphone on which the smartphone was installed. Authors stated jailbreaking the device would lead to finding of more data as not all the cache files were accessible through the forensic tools. Network Level examination yielded different results for (Chung et al., 2017) & (Li et al., 2019). One article focused on identifying the relevant APIs which open up a portal to the Cloud data and the other tried to find network associations of an Echo to find other devices that may be a part of the investigation. Thus, the main observation in this section is that although cloud is a repository for huge amount of evidentiary data for an IoT Forensics case, but in cases, where User Login is missing and the cloud providers refuse to cooperate with the investigators, other approaches need to be adopted and more research needs to be undertaken to find ways out to get data out of the device itself. This can be improved with the suggestion made by (Li et al., 2019) for making the devices forensic friendly.

The next section discusses the findings of each of the works, the differences and how they can be collaborated to give the investigator a clear vision on what data would be most useful to solve a case involving Echo.

4. Findings of Analysis

All the three works came up with interesting findings about the data that could be very useful for an Echo. Although, a lot of data was extracted, but the focus was to consider only the data that would be considered useful for investigation purposes.

Comparing the contributions made by each work, all the three pieces of IoT Forensics literature presented some novel research methodologies that would be a very remarkable addition to the forensics domain. However among all the three, (Li et al., 2019) made a very significant contribution that would make it easier for an investigator to know about the different data types and under what level of forensics those types would fit. Figure 1(Li et al., 2019) contains all the data types.

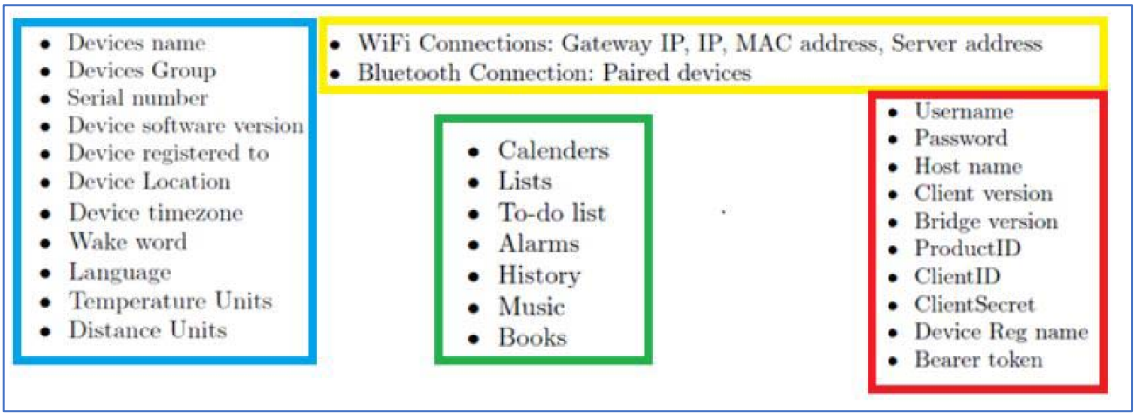

Fig 1 – Data types related to Amazon Echo
Source - (Li et al., 2019)

(Li et al., 2019), however in the article, didn't use the concept of parallelization to explain the significance of this list. The Blue List is an outcome of analysis of device level Forensics. The Yellow List is an outcome of Network Level Forensics. The Red List results from analysis of Firmware. The Green List is of utmost importance as it results from investigation of Client-Side Mobile companion application and the Cloud side Forensics. In their analysis, (Orr & Sanchez, 2018) found a list similar to the Green List through the examination of the client side companion application. This list, in addition to the contents of the Green List contained additional components which would prove useful for an investigation. (Chung et al., 2017) extracted APIs and categorized them based on the functions they served. The data resulting from the JSON scripts via these APIs was very similar to the components of Mobile companion application which was examined by (Orr & Sanchez, 2018). After having established a

background of what the authors extracted from their forensic examination, the significance of that extraction will be discussed next.

First, the components of the Green List will be discussed as it contains the data of most evidentiary value. Two components which are not explicitly a part of this list are added as well.

1) History

This is a component of **Settings** in the mobile companion app as well as the browser application and contains interactions between a user and their Echo. These interactions are either a command from the user or a response from the Echo. Associated with these interactions are the timestamps in a UNIX format. These timestamps refer to the exact time and date, these interactions occurred. (Li et al., 2019) stated that Echo was able to record private conversations, which are not initiated by the generic wake word, "Hey Alexa", but they were unable to identify them.

(Orr & Sanchez, 2018) examined **History** in a more elaborate manner and were able to find out these private conversations. These private conversations had no mention of the wake word. This finding was contrary to what the Amazon executives wrote in their letter to a Senator in the United States (Amazon, 2019). The authors found that these conversations were able to identify the presence of other people in the room. (Chung et al., 2017), in their analysis found two APIs with description similar to the functioning of **History** in the Green List. These two APIs were **utterance** which contained audio data of just the user and **activities** which contained interactions of Alexa with Echo. The authors didn't mention it but it is possible that contents of **utterance** were in coherence with the findings of (Orr & Sanchez, 2018), where they mentioned that Echo stores private recordings as well . Through this component of the green list, a timeline can be created which could bring up an interaction necessary for the investigation.

2) Calendars

This component contains e-email addresses associated with events added by the users as stated by (Orr & Sanchez, 2018), which is a different explanation than what the name suggests. (Orr & Sanchez, 2018) examined **Calendars** in the mobile application deeply and found email addresses which are unknown or not confirmed. These addresses could provide some evidences. (Li et al., 2019) found email addresses as a result of analysis of the firmware through a forensics tool. But in their study, they didn't mention the possibility of these addresses being associated with calendar events. Investigation of this component

yields only e-mail addresses which might provide only a limited evidentiary value. Further research needs to be conducted to find what else can be extracted from this component.

3) Lists

This component is also part of the Green List. **Lists** contain data pertaining to items added by the user in terms of a shopping list or a to-do list as simple as "Get prescriptions tomorrow". (Chung et al., 2017) found this component through accessing the database of the mobile companion apps of both the Android and iOS devices. Since the authors' work was more focused on finding the APIs to get a hold of the cloud data faster, this component of the "green list" was also accessible through the API **todos**. (Orr & Sanchez, 2018) found this component via a simple examination of Settings in the companion application. They explained the significance of this component for an investigation. In a case, it is possible that a user who might be a suspect or a victim would have entered a command that along with its timestamp could prove useful for an investigation.

4) Traffic

This component was not a part of the green list and wasn't mentioned explicitly by (Li et al., 2019), but the authors did mention the importance of geo-location in an investigation. They discussed the use of location for the address set by Echo as the source if any other geo-location is searched for by a user. (Chung et al., 2017) found the API, https://pitangui.amazon.com/api/traffic/settings during their analysis which was used to establish destination addresses during a commute. (Orr & Sanchez, 2018) explained the real significance of this component. **Traffic** keeps track only of the addresses added by the user. Sometimes users add extra stops between the source and destinations. All this information is stored in **Traffic**. This information can be used by investigators to confirm the location of the users. Although, (Orr & Sanchez, 2018) only talk about the commute behaviours of a user in this section, but from a technical standpoint , it is possible that the parameter they use, Geo-location, can identify the presence of the user through the mobile device connected to Alexa, even if no entries are created in **Traffic**.

5) Household Profile

This component is also not the part of Green List (Fig.1) (Li et al., 2019), but was discussed both from a technical standpoint and importance-wise by (Chung et al., 2017) and (Orr & Sanchez, 2018) respectively. The Household profile is accessible through cloud level

forensics. (Chung et al., 2017) found the API for household profile, https://pitangui.amazon.com/api/household. This API consists a list of household accounts associated with Amazon Echo. The authors didn't explain this component's significance. (Orr & Sanchez, 2018) described this list in a precise manner. Since this profile had different accounts associated with one Echo, it is possible that one of their activities might have an investigatory value. If during an investigation, the suspect is one of the members of this profile. Although determining that the suspect being a member of this profile will not add anything to the case investigation, but this will lead to further examination of other components of the newly formed green list.

The second most relevant finding of the examination of Echo by all the authors was the data present in the Cache of smartphones. The cache wasn't found easily through examination of just the application. (Orr & Sanchez, 2018) gained access into the cache via analysis of the device containing the application through a Forensics Tool. They found four files in the cache and among them, two were very significant from an investigator's point of view. These were com.amazon.echo.plist and LocalData.sqlite. The first file stored the user activities and settings and could be only opened when the device was jailbroken. The second file had the same information as **Lists**. The authors stated that this file should be examined under two conditions, absence of user credentials and when the investigators suspect that the Echo user has attempted to delete the information from **Lists**. While (Orr & Sanchez, 2018) used an iOS device for the examination, (Chung et al., 2017) examined caches through Android Web View, an iOS device and Chrome browser. For the Android device, they found the same data that (Orr & Sanchez, 2018) found for the iOS device. However, (Chung et al., 2017) didn't really explain the significance of the JSON scripts found in Chrome Cache.

In this section, two main findings were discussed; the Green List and the Application Cache. However, all the three works found small tidbits of information as well which could prove useful in an investigation, depending on the scenario. These tidbits were extracted as a result of three tier forensics. These tidbits are summarized in Table 1.

Table 1

Potential Source of Evidence	**Relevance**	**Source**	**Level of Forensics**
Network Ports	Open Ports can analyse connection behaviours.	(Li et al., 2019)	Network Level

Stop.mp3, error.mp3	These identify the time when Echo was last used	(Li et al., 2019)	Device Level, Cloud Level
Login ID	It can be used to access the companion application or the service on a browser.	(Li et al., 2019)	Device Level (through analysis of firmware)
Third Party API	This API can identify associations of Echo and Alexa which are utilizing the device or the virtual assistant.	(Chung et al., 2017)	Network Level, Cloud Level
Phoenix API	This identifies compatible devices which work in collaboration with Alexa and Echo.	(Chung et al., 2017)	Network Level, Cloud Level
Home Cards	These are similar to **Lists**, however they don't have timestamps	(Orr & Sanchez, 2018)	Cloud Level

The potential sources of evidence summarized in Table 1 aren't as relevant as the Green List and the Cache, but it doesn't mean that they won't be useful for an investigation. All cases are not the same and it is quite possible that a suspect would have tampered with the Echo trying to render it useless, or simply would have destroyed the smartphone on which the companion app was installed. The device firmware was only examined by (Li et al., 2019), and this yielded Login ID. This finding would prove very useful when the smartphone was missing. Similarly, other sources presented in Table 1 can have scenario specific relevance.

5. Discussion & Results

In the above two sections, the methods of extracting evidentiary data and the relevance of that data were discussed. All of the three articles gave insights into the field of Echo Forensics from their own perspective. (Chung et al., 2017) worked with APIs and companion applications but their work was more API centric. This approach was unique and had the upper hand over the other two articles because in order to work with Alexa APIs, the investigator needs the user credentials only and doesn't even require the smartphone. The APIs identified by the authors were very useful for Echo Forensics. The major limitation of this approach was that the APIs won't yield any data in the absence of user credentials. The authors then studied the Caches in order to overcome this limitation but didn't explain their significance for an extraction. (Orr &

Sanchez, 2018) found the data files in the Cache and explained how they would be useful for an investigation. Although, their work didn't involve heavy usage of Forensics Tools, but through a deep analysis of the companion applications and the Alexa web application, they found some very interesting components which were residing in **Settings** section of the application. These components were able to give out important forensic details like the conversations between Users and Echo, other profiles associated with that particular Echo, commute behaviours of the users, etc. Their approach was comparatively easy to perform for an investigator, but again, in the absence of user credentials, it would fail. The authors took care of this limitation by finding evidentiary data from cache. In some cases, where cache files weren't accessible, they proposed that jailbreaking could be the answer. The analysis of Echo by (Li et al., 2019) was a little different but they didn't mention how they extracted the Green List (Fig.1). Their approach was a mix of all three forensic levels. They examined Echo's firmware to extract some basic information from it. They also analysed it with a forensic tool, and this gave out the login details for the user. Rest of their findings were similar to the above works.

All the three articles alone provided a good amount of information, but a collaboration of their techniques and findings would turn out to be a great investigators' handbook. In simple terms, one article's focus is how to find evidentiary data, the second article presents and lists it and the third majorly explains the significance of that data.

Let's go back to the infamous Arkansas murder case of 2015 (Taylor, 2017), where Amazon Echo was found as the only source of evidence. It was believed that suspect's Echo recorded some conversations which would prove him guilty. The investigators weren't aware of Echo Forensics and despite their best, couldn't find any data. They asked Amazon to provide them with cloud data. Amazon stated that such cooperation from them would lead to violation of consumer privacy. For a very long time, Amazon didn't provide any data to the Law Enforcement Agencies. In the end, investigators were provided with the history of transactions, but the audio files were not given by Amazon. In the future, if such cases arise, the investigators will have sufficient knowledge on how to deal with forensics of Echo and its associations. All the authors in their work mentioned the word "timeline". However, they didn't explain how important timeline construction would be for a case. In cases involving multiple devices, investigators, after collecting and examining relevant evidence can use the concept of timeline to reconstruct the complete crime scene. The phrases and the recordings obtained from Echo are just words and audio files until they have a timestamp associated with them. That is why during their experiment, (Orr & Sanchez, 2018) didn't rely on **Home Cards** which are more

easy to obtain but without the timestamps, they don't make any sense. And it is quite possible that if investigators just rely on Home Cards to find evidentiary data, they might make wrong judgement calls during the next steps of the investigation. Another important point that these articles raised was the need to make the devices forensic friendly. It should be understood by the device manufacturers that in today's world, security is not just a notion. As explained by (Li et al., 2019), most of the IoT devices have Memory Protect Technology embedded in them, which prevents investigators from accessing important files in the device. It is a well-known fact that IoT devices, due to their constrained nature can't afford to have state of the art security protocols protecting them from threats. This fact is being exploited by the malicious actors of technology so that they can wipe their footprints off. Thus, it is important that the IoT devices from now on should have a forensic friendly design.

6. Conclusion

With almost 20 billion IoT devices surrounding us from everywhere to make our lives easier, they are also recording our conversations, monitoring our commutes, analysing our behaviours, which from a user's perspective would be breach of privacy, but from a crime investigator's point of view, this could unlock the potential of solving the crimes in a much efficient manner. One important class of IoT devices is that of the home assistants which are not just accepting basic instructions from their owners, but also acting like an intermediator between the user and other smart devices in the home. Amazon Echo is one of them. Its conversations with the user, its interaction with other devices in the home like the smart thermostat, smart coffee makers and others is being stored in the cloud so that over time, its performance can be improved. However, this stored data proves very useful during the occurrence of a crime. The articles discussed in this report attempted to extract the meaningful excerpts of this data through different levels of forensics. All the three works used their own unique approach to get the evidentiary data. This data resides in different zones. And extraction from each zone has its own challenges. Throughout the report, these approaches and the findings were discussed. And the question that was raised in the section **Introduction** can now be safely answered. The most reliable sources of evidence in an Echo lies in the cloud and can be accessed either through companion application or through the APIs. Other means of finding evidentiary data could be through firmware analysis and scanning of ports. This would lead to discovery of other devices associated with the virtual Assistant, Alexa. The other part of the question was on how these sources can prove useful for a case. The contents of Green List (Fig.1) were able to record conversations between the user and Echo, analyse user's commute behaviours, find out what

the user was up to via Lists and other data that would prove useful for a case. All three articles had significant amount of information through which this question was answered. This report was focused on Amazon Echo, but the extraction methods and findings can be applied to other home assistants like Google Home as well. To sum up this report, as crimes are becoming complex due to the aid of technology, the investigators need to keep up with the non-standardised file systems of IoT devices that can't be accessed through the state-of-the-art forensic tools. Since the world is already filled with around twenty billion IoT devices, need of the hour from an Investigator's point of view is to utilize them in an efficient manner so as to solve crimes without hampering mass user privacy.

References

Amazon. (2019). *Amazon Letter to Senator Coons EEUU*. Retrieved from https://www.amazon.com/gp/help/customer/display.html?nodeId=202201630

Chung, H., Park, J., & Lee, S. (2017). Digital forensic approaches for Amazon Alexa ecosystem. *DFRWS 2017 USA - Proceedings of the 17th Annual DFRWS USA*. https://doi.org/10.1016/j.diin.2017.06.010

Li, S., Choo, K.-K. R., Sun, Q., Buchanan, W. J., & Cao, J. (2019). IoT Forensics: Amazon Echo as a Use Case. *IEEE Internet of Things Journal*, *6*(4), 6487–6497. https://doi.org/10.1109/jiot.2019.2906946

Mckemmish, R. (1999). What is Forensic Computing ? *Trends & Issues - Australian Institute of Criminology*.

NIST. (2006). Guide to integrating forensic techniques into incident response (NIST Special Publication 800-86). *NIST Special Publication*. https://doi.org/10.6028/NIST.SP.800-86

Oriwoh, E., Jazani, D., Epiphaniou, G., & Sant, P. (2013). Internet of Things Forensics: Challenges and approaches. *Proceedings of the 9th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing, COLLABORATECOM 2013*. https://doi.org/10.4108/icst.collaboratecom.2013.254159

Orr, D. A., & Sanchez, L. (2018). Alexa, did you get that? Determining the evidentiary value of data stored by the Amazon® Echo. *Digital Investigation*. https://doi.org/10.1016/j.diin.2017.12.002

Zawoad, S., & Hasan, R. (2015). FAIoT: Towards Building a Forensics Aware Eco System for the of Things. *Proceedings - 2015 IEEE International Conference on Services Computing, SCC 2015*. https://doi.org/10.1109/SCC.2015.46