CHAMPLAIN COLLEGE | LCDI Leahy Center for Digital Investigation

Amazon Echo Forensics

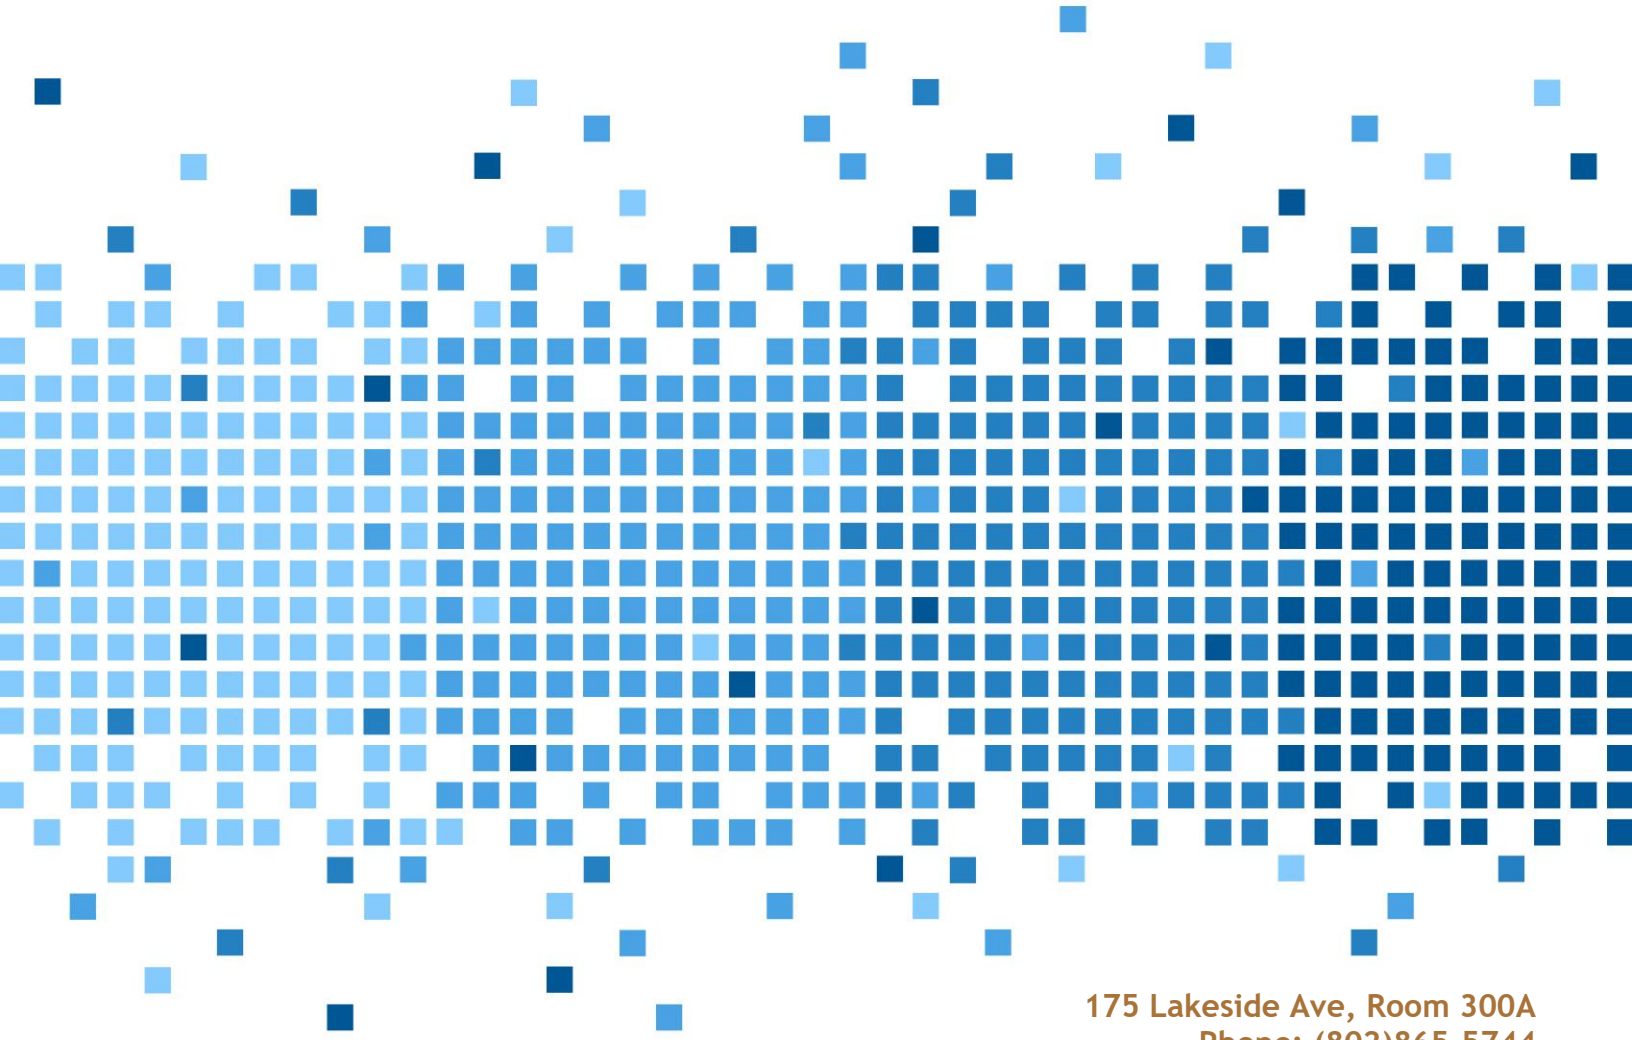

175 Lakeside Ave, Room 300A
Phone: (802)865-5744
Fax: (802)865-6446
http://www.lcdi.champlain.edu

03/31/2016

Disclaimer:

This document contains information based on research that has been gathered by employee(s) of The Senator Patrick Leahy Center for Digital Investigation (LCDI). The data contained in this project is submitted voluntarily and is unaudited. Every effort has been made by LCDI to assure the accuracy and reliability of the data contained in this report. However, LCDI nor any of our employees make no representation, warranty or guarantee in connection with this report and hereby expressly disclaims any liability or responsibility for loss or damage resulting from use of this data. Information in this report can be downloaded and redistributed by any person or persons. Any redistribution must maintain the LCDI logo and any references from this report must be properly annotated.

Contents

1 Introduction

The rapid advancement of technology has given rise to a myriad of devices that have the potential to pervade our daily lives. Human beings no longer require a keyboard and mouse to access the Internet. The Amazon Echo is one of these devices. Over the course of this project, the LCDI team sought to explore the Amazon Echo beyond its initial function -- to serve as a conduit between a user and the Internet -- by examining it as a potential source of forensically relevant digital information.

1.1 Background:

At the time of this report, the Amazon Echo has only been released to the public for a short time. We were unable to find any prior digital forensic research that has been conducted on the Echo. Our initial research into the Echo shows that its primary function is to pair with other services and devices in order to be used as a universal hub that allows instant access to them via voice commands from the user.

1.2 Purpose and Scope:

Our primary goal was to learn everything we could about the Amazon Echo. Since this is a relatively new device, any findings of forensic significance will be important moving forward. Part of our research included investigating devices that can be connected with the Echo for any residual artifacts. Any device with a web browser, the Alexa companion app, or IoT capability was within the scope of this project.

1.3 Research Questions:

- *What forensic artifacts can be recovered from devices and interfaces associated with the Amazon Echo?*

1.4 Terminology:

Acquisition – The process of copying data from a piece of evidence, to another location in a forensically sound manner so that the data may be analyzed at a later time. This is usually done by attaching some form of write blocking device to the storage media, and creating a copy of the data. The goal is to leave the original media intact while working on a copy of it. This allows for evidence to be verified at a later date. There are two different types of data acquisition methods: Physical and Logical.

Android (OS) – An open source operating system developed by Google and based on the Linux kernel for mobile devices with support for an expanding number of hardware devices.

Android Studio – The official integrated development environment for Android platform development. It allows data to be pushed and pulled from an Android Device.

APK – The android application format.

Artifacts – Any data generated by user interaction that can be collected and examined. Any user data retrieved from the browser is considered an artifact, including cookies, caches, geolocation, search history, etc.

Bulk Extractor – A digital forensics tool used to scan a disk image, directory of files or nearly any digital media and extracts useful information. It creates histograms and lists of features found. It ignores system structure and automatically detects, decompresses and recursively reprocesses a variety of compressed data types.

Encase – A digital forensic investigation suite created by Guidance Software. We used this as our primary tool for parsing data for forensically relevant artifacts.

E01 – The extension of an image file for EnCase.

Image – Often refers to a copy of a hard drive, or disk image, which is compressed into a series of files. Physical images include all information (zeroes and ones) on the hard drive whether the space is being used or not, and ends up being close to the same size as the actual hard drive itself. As opposed to a physical image, a logical image only acquires the parts of the hard drive that have active data and dismisses the rest of the drive. Compared to a physical image, the size can be extremely small or the same size as the evidence drive.

Parse – The process of dividing a computer language statement into parts that can be made useful for the computer. A parser in a program compiler is a program that takes each program statement that a developer has written and divides it into parts (for example, the main command, options, target objects, their attributes, and so forth) that can then be used for developing further actions or for creating the instructions that form an executable program.

7-zip – An open source software used primarily to compress and extract files. It has a command line and graphic user interface.

2 Methodology and Methods

For this project we will be using utilizing a data generation procedure to explore information that can be extracted from the Amazon Echo and mobile devices that interact with it. Interfacing with the Echo will be the first avenue of contact. The Amazon Echo requires devices that are connected through the Alexa app to initially setup the device and interact with Amazon's servers. After contact, we will be utilizing a mobile platform- a Nexus 7 tablet - to collect data. The Nexus 7 is a typical handheld platform running Android 4.4 KitKat. The connectivity and configuration of the device will be key for finding additional data from the Amazon Alexa application.

The Amazon Echo mobile application and web interface are also supported on iOS, however that platform was outside the scope of our project.

Images of the mobile device will be taken to compare the data added during the data gen and the data stored on the device. Conclusions will be drawn from the final results of comparing potential application configurations.

2.1 Equipment Used

Table 1: Hardware

Device	OS Version
Nexus 7 Tablet	Android 4.4.4 KitKat
Amazon Echo	Proprietary Amazon Web Services
LCDI Workstations	Windows 7 Enterprise

2.2 Software Used:

Table 2: Software

Software	Version	Comments
Encase	7.10	
Microsoft Windows	7	

Application: Alexa	1.0.138-prod_713061	
Bulk Extractor	3.1	
Android Studio	Version 1.0.32	Android Debug Bridge
7-Zip	15.14 (x64)	

3 Data Collection:

This stage of the project focused on generating forensic artifacts that could be obtained from the Amazon Echo and its periphery software. In the case of the Echo, this involved beginning with a blank slate; this is because the minute the Echo is set up it starts responding to commands and generating data. The bulk of the generated data was on its capability to respond to voice commands.

Our largest data generation involved a formatted Nexus 7 tablet, a laptop and factory new Amazon Echo. These devices were isolated to their own network to minimize interference from outside connections. The Echo was turned on and set up using an associated mobile application that was installed on the tablet. The app continued to propagate new information as the Amazon Echo was asked a series of basic questions such as “where is Burlington, Vermont?”. During this process the Echo webpage, which requires Amazon credentials to access, was also open and being watched for activity. After the questioning period the Echo and all other devices devices were powered off.Our focus then shifted to the tablet as the most likely source of evidence. In order to answer specific questions, more targeted data generation tests were conducted using specific commands and smaller test groups. This was typically done to verify artifacts.

We were unable to obtain data directly from the Echo , although it was noted that there are jtag points on the device as pointed out by Bill Finlayson on Twitter (https://twitter.com/bill_billbill_/status/617547530092412928). Since its purpose is to serve as a cloud-based, voice activated assistant, it has no obvious input/output ports. The best course of action towards answering our research questions was to shift our investigation towards the periphery software.

4 Data Extraction

4.1 Android Tablet Data Processing

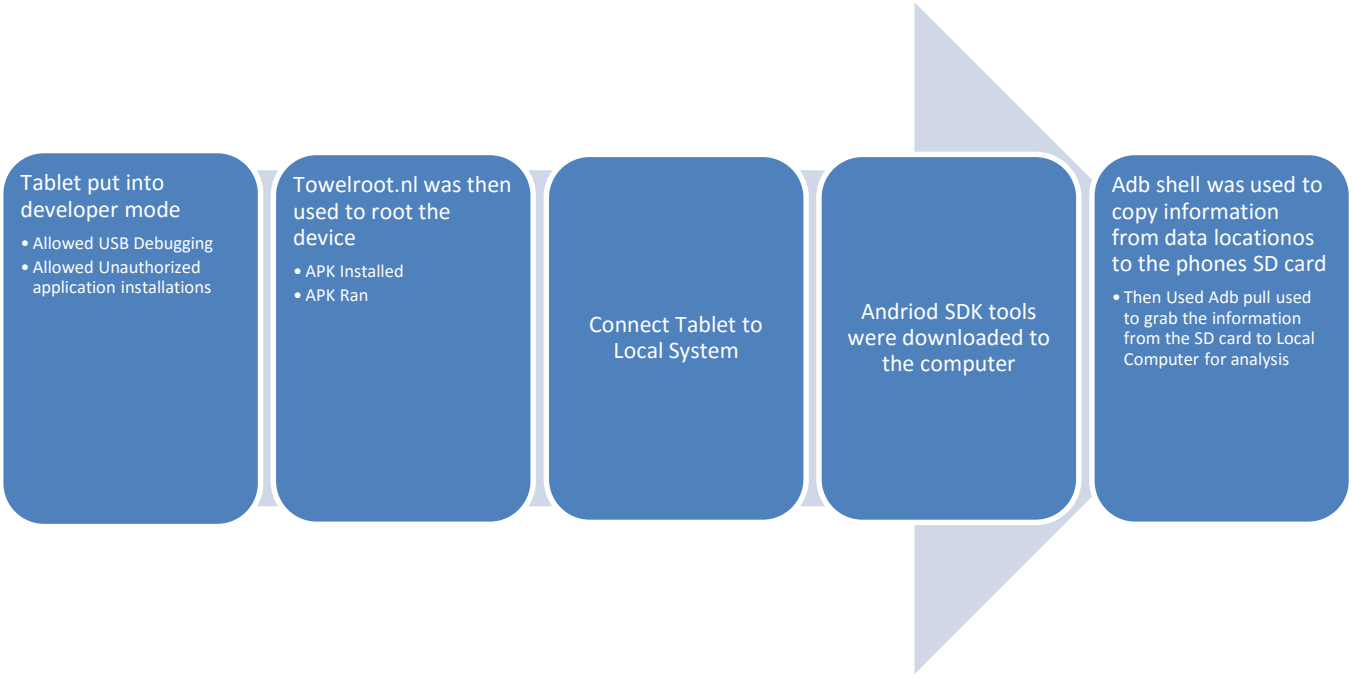

5 Analysis

Our initial investigation showed that the Amazon Echo is mostly a cloud based service, and thus we weren't sure how much user data would be retrievable from the associated applications. While most data processing is done over the Internet, a small amount of information is retained on the mobile application. When a voice prompt is given to the Amazon Echo, a corresponding "card" database entry is created on Amazon's servers. When prompted by the mobile application, this information can be downloaded to the phone's cache.

The term 'card' comes from the file storage type found in the application. Once files were pulled from the Nexus tablet the contents were added to Encase. These files included /*data/data/com.amazon.dee.app* and *com.amazon.dee*.app.

In EnCase, the phone's files can be viewed in their original forms. The file system structures also remain unchanged. We began by processing *com.amazon.dee.app*. The structure of the Amazon application was typical to most application data: this includes user permissions and the email address of the user's Google account. Once this was examined we looked to the app's main directory.

Figure 5.1 shows a rough outline for the file structure of the /data/data/com.amazon.dee.app directory:

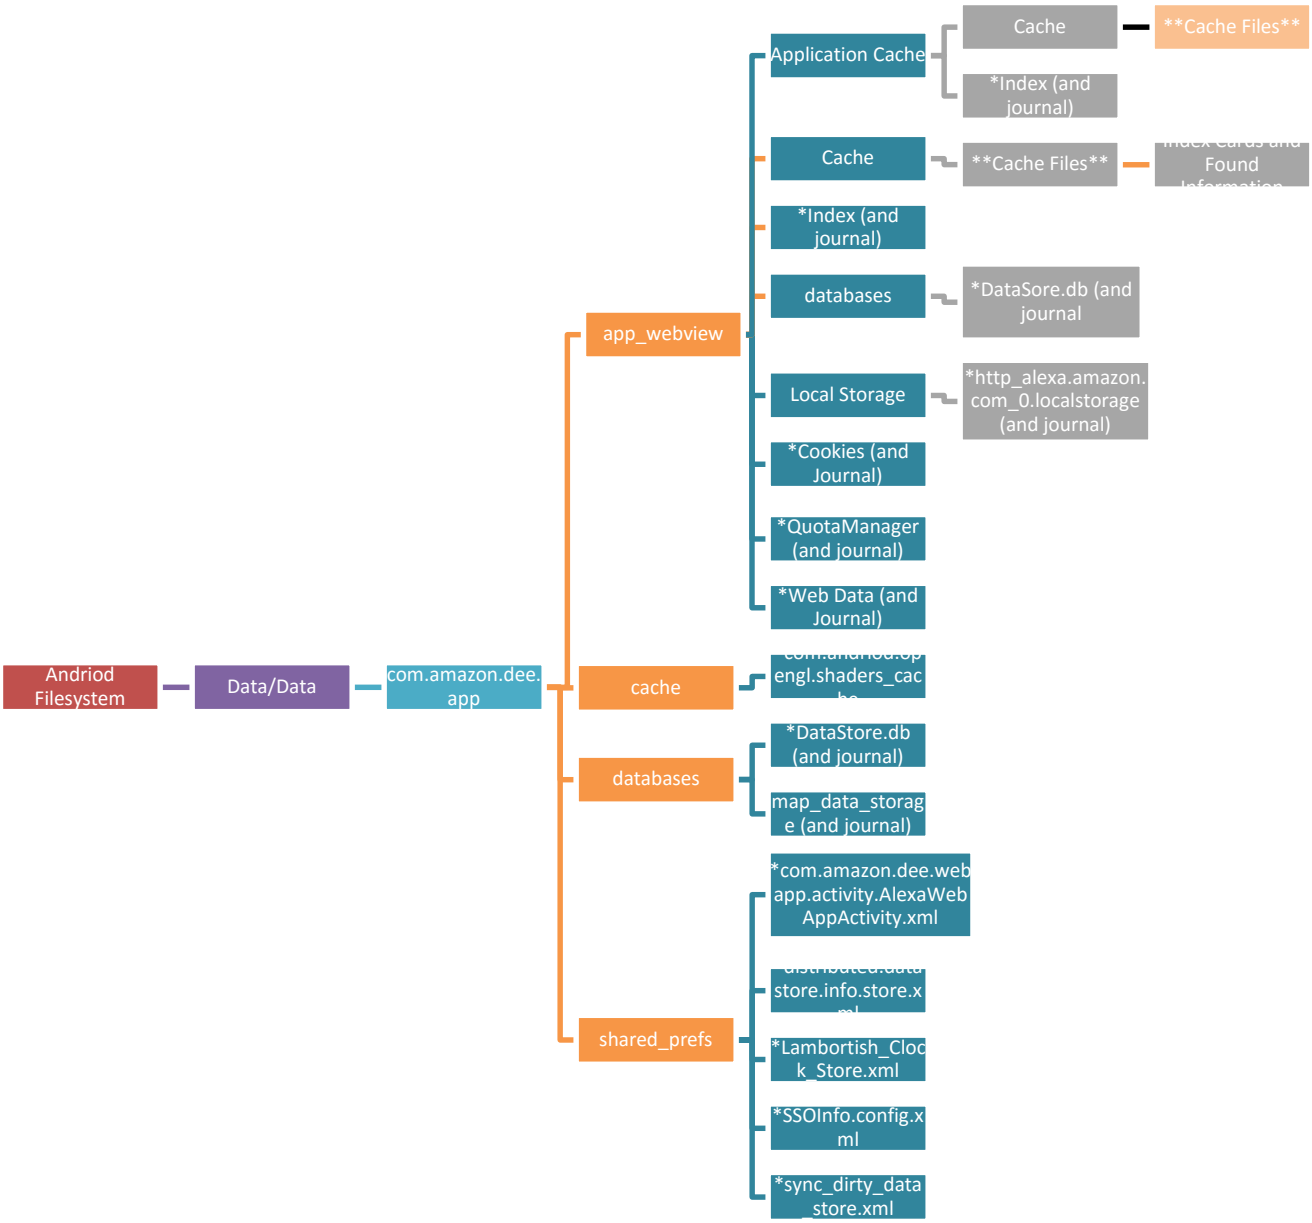

The first file that we noticed was called **sound**, a .wav file with no extension, stored in */data/data com.amazon.dee.app/cache*. This file will only be created if the user listened to one of their prior voice commands from the "History" tab of the mobile application.When this is done, the Echo downloads the file from Amazon's servers and plays it back. This .wav file, which can be opened with any applicable media player, is the only time that a sound file will be physically stored on the device.

The next location of interest is the *com.amazon.dee.app/web_view/cache*; this location has multiple files that are passively created while the app is running. These files are use gzip compression and, when decompressed with 7zip, compose of various images and text files. The amount of information contained in the cache varies with how often it is cleared,, and we were unable to determine the maximum amount of data it can hold.

Further, we were able to confirm the existence of several 'cards' in the application data. These cards contain a variety of forensic artifacts including the text-to-speech conversion of what the Echo heard, its vocal response to the command, and the URL of the recorded sound file's location on the Amazon server. All the information that can be viewed on the home screen can be found in these card files.

The cards are presented in .json format. There are a few common tags across the cards although some are unique to specific actions. An example of a "StandardCard" is shown below:

```
    {1}
root
   array{20}
      cardMetricAttributes:null
      cardType:StandardCard
      creationTimestamp:1456426374073
      deleteCardAction{5}
         actionType:DeleteCardAction
         cardId:AIGGLSORA58RK#1456426374073
         mainText:Remove card
         subText:Learn more
         subTextRoute:help/node/201602230

      descriptiveText[1]
0:Digital forensics (sometimes known as digital forensic science) is a branch of
  forensic science encompassing the recovery and investigation of material found in
  digital devices, often in relation to computer crime. The term digital forensics was
  originally used as a synonym for computer forensics...

      giveFeedbackAction{9}
         actionType:GiveFeedbackAction
         mainText:Thank you! Your feedback helps Alexa understand you better.
         musicCustomerId:null
         route:beta-feedback
         serviceName:null
```

```
    subText:Send more detailed feedback.
    subTextRoute:null
    thirdPartyAppId:null
    thirdPartyAppName:null

id:AIGGLSORA58RK#1456426374073
imageAction:null
imageCaption:from Wikipedia
imageReference{3}
    fallbackIcon:null
    referenceType:UrlImageReference
    url:https://upload.wikimedia.org/wikipedia/commons/thumb/7/7a/Hard_disk.jpg/1440px-
        Hard_disk.jpg

nBestOptions:null
playbackAudioAction{5}
    actionType:PlayAudioAction
    mainText:Alexa heard: \"wikipedia digital forensics\"
    subText:null
    subTextRoute:null
    ur :/api/utterance/audio/data?id=AB72C64C86AW2:1.0/2016/02/25/18/B0F00715544703K4/52:50::T
    l   NIH_2V.8a8bb8b3-88f8-47d1-bf6e-a5e8055312ffZXV/1

primaryActions[1]
    0{5}
        actionType:OpenUrlAction
        mainText:Learn more on Wikipedia
        subText:null
        subTextRoute:null
        url:https://en.wikipedia.org/wiki/Digital_forensics


registeredCustomerId:AIGGLSORA58RK
secondaryActions:null
sourceDevice{2}
    serialNumber:B0F00715544703K4
    type:AB72C64C86AW2
```

```
subtitle:Wikipedia
thumbsUpDownActivityAction{5}
    actionType:ThumbsUpDownActivityAction
    activityId:AIGGLSORA58RK#1456426374073
    mainText:Did Alexa hear you correctly?
    subText:null
    subTextRoute:null

title:Digital forensics
wrapTitle:true
```

5.1 Fields of Note

The *creationTimeStamp* field shows the time zone of theWiFI network that Echo was connected to at the time of the command. This timestamp is in UNIX format. This is regardless of when the card itself propagated on the mobile device. If the device isn't connected or open at the time of creation the cards will then be pulled to the device when next updated.

The *maintext* field shows the results of the Echo's text-to-speech conversion.

The *sourceDevice* field represents the serial number which identifies the specific Echo that received the user query. Information in this field can connect the physical device to a user.

Finally, the url field gives a web address that begins with *api/utterance/*. In the current configuration, this address will lead to the location of the sound bit on Amazon's servers. It can be accessed by the user so long as the proper credentials for the corresponding account are supplied. The credentials needed to access the sound file are those used by the account that set up the device. The web domain is pitangui.amazon.com.

Pictures associated with the card can also be stored on and are retrievable from the mobile device, in the same location as the cards. Pictures are generated from a variety of queries like Wikipedia searches. Photos generated by a Wikipedia search can be linked with the Wikipedia article, as the icon for the article is the picture pulled. As shown above, there is also a URL to the referenced thumbnail.

5.2 Artifact Limitations

As we have been exploring and finding artifacts there can be large holes in the possible information. All of the artifacts found require user interaction to have then be cached on the device. If the app isn't opened the cache files aren't updated with the most recent information.

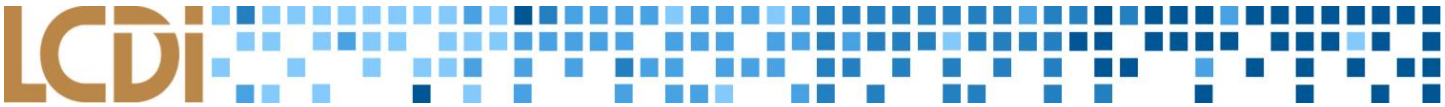

There also can be holes in the information stored if a third party app is used with the echo. The device may not record secondary responses given by the user. An example of this interaction can also be seen when using the amazon purchase feature. If the echo is used to buy a product from amazon the initial communication is recorded, then the echo asks for a confirmation code. This confirmation code being communicated is not saved in the cards above.

6 Conclusion

Upon the completion of our research, we were able to find a significant amount of user data on the Android application. This information includes timestamps for the commands given to Alexa as well as the Echo's response to user commands. However, this information can vary greatly depending on how the application is utilized.

Processing the application data is a fairly straightforward process once the format is understood. These artifacts could prove useful in future investigations and may also be relevant to research on Internet of Things devices, as they share the same potential to have a limited device interface.

The conclusions drawn from our research at time of publication were confirmed with multiple iterations of testing. However, with the software still being refined, updates may change the amount of information that can be collected.

7 Future Work

Future work on this project may include investigating third party devices and services that can interface with the Echo. These may serve as additional vectors for data collection. As time passes and the device receives updates, the research conducted in this project may be rendered obsolete. Revisiting the Echo at a later date may provide additional insight.

A project done by Chris Pazden, a Champlain College student, revolved around deciphering the index file structure. This research could provide additional information regarding the Echo's card files and determining their creation time on mobile devices.