Amazon Echo Dot v2 Penetration Test Summary

Matthew Sutton
Tim Gekas
James Autry

University of Nebraska Omaha

CYBR4580 - Cybersecurity Capstone
Proctor: Dr. Matthew Hale
05/04/2017

Table of Contents

Executive Summary

Intelligent, voice-controlled systems are a new and emerging technology now being found in many homes. In 2014, Amazon launched their new voice controlled assistant called Alexa. It has been integrated into the Amazon Echo and into the Amazon Echo Dot series. These are devices which sit in a user's home, constantly listening for the user to make a request that Alexa could fulfill. Examples of requests that can be fulfilled are playing music, giving information on the weather, or controlling another Internet of Things (IoT) device in the user's home.

Our primary objective is to execute a full-scope pentest into the Amazon Echo Dot (2nd Generation). Specifically, this refers to compromising the security of the device itself - not the security of Alexa, whose implementation is stored on official Amazon servers. However, there are devices that are made to interface with the Echo Dot (namely, the Alexa Voice Remote) that we will also be inspecting to see if they can be used to compromise the Echo Dot.

Voice-controlled AI devices such as the Echo Dot are a newly emerging technology, and there is still a lot of research and testing that needs to be done before these IoT devices can be said to have been properly vetted for widespread use. In pentesting the Echo Dot, we hope to accomplish one of two ends: either improve the assurance that these devices are secure enough for widespread use, or improve the security of the device so that it can eventually be suitable for widespread use.

Goals and Objectives

- Gain valuable experience pentesting as a team
- Plan a thorough investigation into the security of the Echo Dot
- Test the following strategies:
  - Intercept the Echo Dot's Wi-Fi communications and execute Man-in-the-Middle based attacks
  - Attempt to be the first known group to root an Echo Dot v2
  - Find a vulnerability that could lead to compromise of the Echo Dot through the Skills API
  - Develop secure usage practices by finding the threshold decibel levels required to converse with Alexa
- Provide a technical writeup concerning the results of our investigation into the Echo Dot, both for the benefit of Amazon so that they may patch any vulnerabilities we find, and for the rest of the security community, so that they may build off our work

Assessment Activity Summary

To assess the Echo Dot, we first developed five of the most common user stories for entities that interface with the Echo Dot. User stories were used to determine what the Echo Dot does.

US1 - End users giving *verbal requests to Alexa*
US2 - End users wanting *security and confidentiality from their Echo Dot*
US3 - End users downloading and installing *Alexa Skills*
US4 - Developers creating new *Alexa skills*
US5 - Amazon employees ensuring there are *no firmware update vulnerabilities*

From these user stories, we assessed their standard procedures and developed acceptance criteria that need to be met to assure our user stories have no vulnerabilities. To frame the threat landscape, we created a use/misuse case diagram to assess for potential vulnerabilities. Additionally, an activity diagram was created for each user story to visualize the user's actions as they complete a user story's action. These visual aids were all used to identify potential vulnerabilities in the Echo Dot's functionality.

US1 Activity Diagram
Giving a verbal request to the Echo Dot

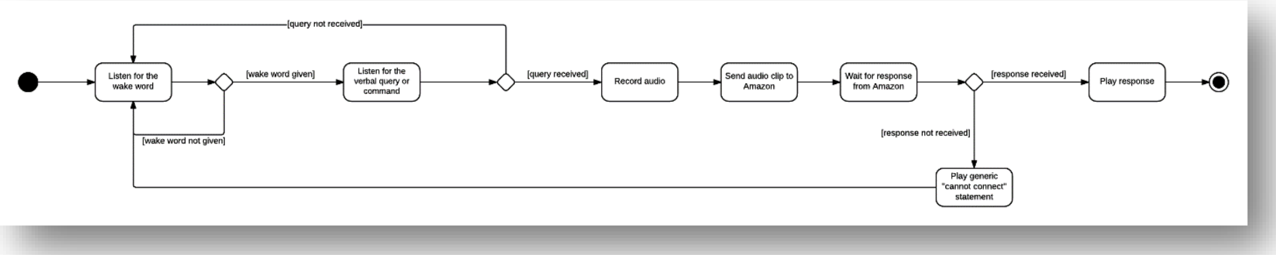

After identifying weak-spots, we set forward to gain more information about how the device operates. From our gained knowledge of the product we deduced there were three architectural layers to the Echo Dot: hardware, firmware, and the Amazon backend. Architectural diagrams were created for each layer to further footprint how each layer's components interface with each other. By knowing how the components are organized, we planned methods of exploitation. When one component is compromised other sibling components may follow. All this planning led us to the exploitation phase of our penetration test.

Firmware Architectural Diagram

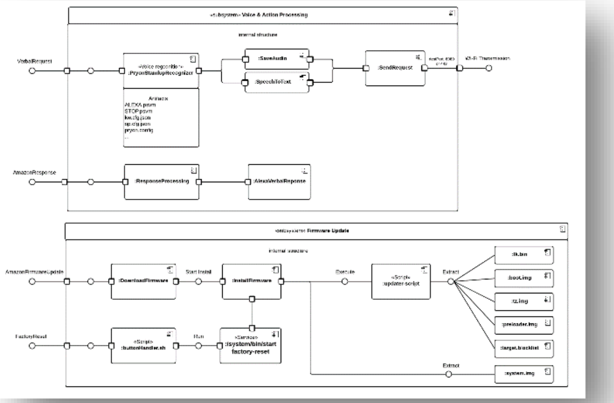

Threat Landscape

From all our Echo Dot foot printing, we believed the device would be susceptible to network-based attacks. Specifically, we were not confident in the Echo Dot's protections from man-in-the-middle attacks. All our user stories require network communications in some manner:

US1 - End users give requests to Alexa that are then transmitted to Amazon for processing.
US2 - End users use the Echo Dot for wireless IoT administration of their home.
US3 - End users download new skills to their Echo Dot to add features.
US4 - Developers create new skills and upload them to Amazon's servers.
US5 - Amazon pushes a firmware update to the Echo Dot.

It was concluded that network-based exploitation would be one of our most likely avenues of exploitation. As such, we obtained a Wi-Fi pineapple for network-based pentesting. On a similar train of thought, we also decided as part of our wireless exploitation to pentest the Echo Dot's interactions with the Alexa Voice Remote, a device which connects to the Echo Dot via Bluetooth.

Another avenue of exploitation was via Amazon's Skill API. We theorized that it could be possible to code malicious skills using the Alexa Skills Kit. Vulnerable code or functions in the skills kit could be deployed in a malicious fashion to exploit an end user.

Our final avenue of exploitation was to obtain root access to the Echo Dot to identify is full capabilities. Preliminary research led us to believe this was not possible at the current time. The bootloader is locked down and Amazon is not releasing its key. However, the Echo Dot is can be flashed via MediaTek Android Smartphone Flashing Toolkits given the proper data. We obtained the MediaTek tools for firmware based exploitation.

Assessment Findings Summary

Network Findings

Almost all traffic to and from the Echo Dot v2 is encrypted using TLS v1.2. A man-in-the-middle attack using SSLsplit failed to net any results. We do not possess Amazon's private-key and were unable to strip encryption. We have evaluated that the Echo Dot v2 properly secures all network traffic from eavesdroppers.

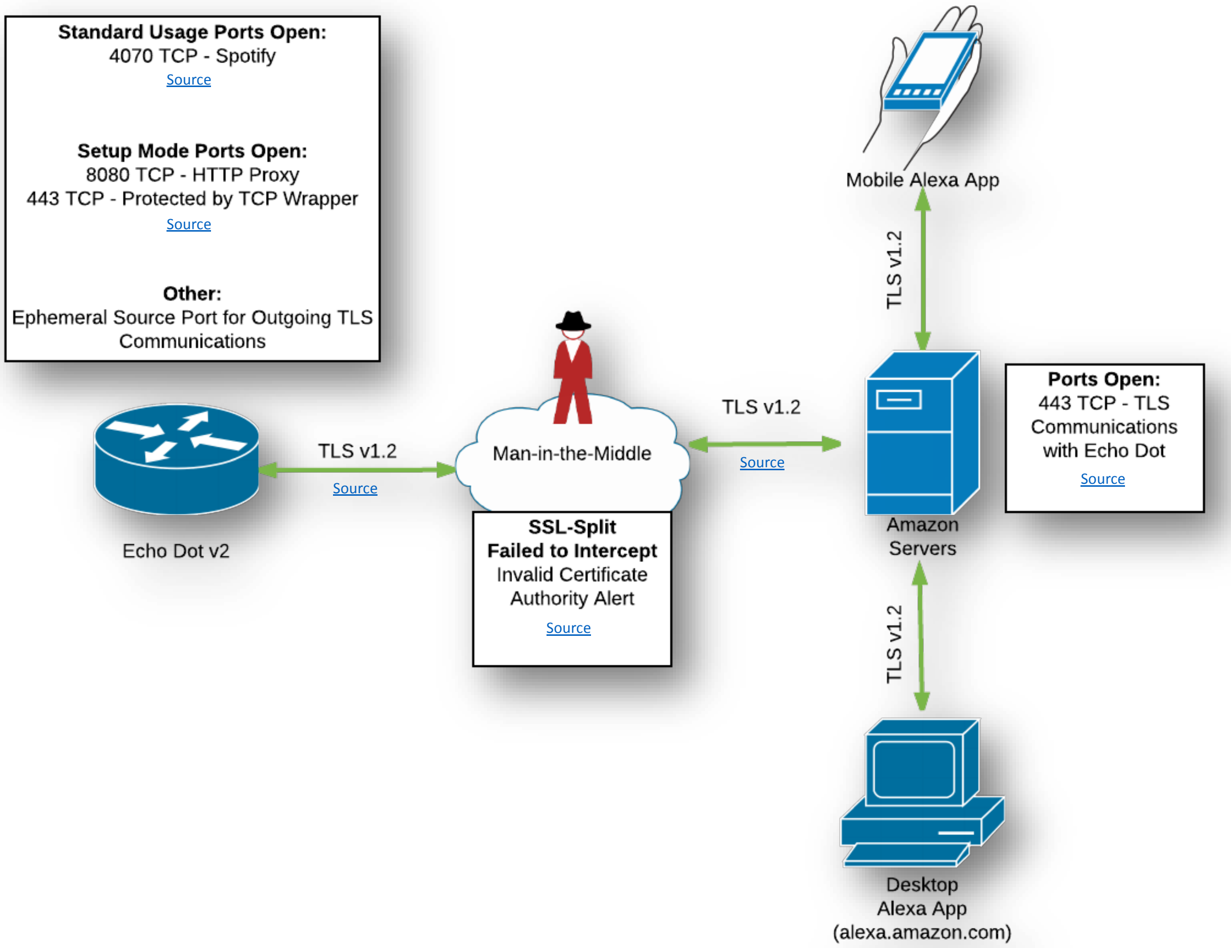

As another part of our exploration into wireless vulnerabilities, we also attempted to intercept and spoof the Bluetooth signal from an Alexa Voice Remote. However, we ran into technical issues in the interception phase and had very limited ability to intercept the remote's signal. As such, testing on this front has remained inconclusive for the present.

Skills API Findings

The Alexa Skill API provides a limited set of permissions that Skills can access. The end user always receives a permission request via the Alexa App when they enable a developer's Alexa Skill. The following table summarizes the permissions.

Permission	Description	Exploitable?	Reason
Device Address	Allows access to end user's full address, or country and postal code.	No	This permission can only expose the physical address if the end user provides that information. The end user can also restrict the address information to only the country and postal code.
List Read	Allows read-only access to the Alexa to-do and shopping lists.	No	This permission only allows read-only access to the shopping and to-do lists.
List Write	Allows read-write access to the Alexa to-do and shopping lists.	No	This permission does not allow Skills to actually order anything, only to add/remove items from the shopping and to-do lists.

All Alexa Skills must also pass the Amazon Skill Certification requirements to be published and made available for end users to enable. These requirements include answering questions about the behavior of the Skill and providing testing instructions for the certification team.

Since all Alexa Skills are tested by the Amazon certification team before being put into production, and given the limited permissions that Alexa Skills have access to, it is unlikely that a developer could create an Alexa Skill with malicious intent.

Root Access Findings

Rooting Method	Working?	Reasoning?	Workaround?
Bootloader	No	Amazon has locked down the bootloader. Fastboot getvar all shows unlock_status: false	None – wait until Amazon releases the unlock files.
MediaTek Smart Phone Flash Tools	No	Unable to acquire proper scatter file for flashing.	Buy a product with the same MT8163 V/B System-on-Chip and rip its scatter file to apply to the Echo Dot.

The Echo Dot v2 has a bootloader that is locked down. There is no access to common Android rooting tools like Android Debugging Bridge (adb). Until Amazon releases the unlock files for the bootloader another method must be used.

The Echo Dot v2 uses MediaTek hardware that includes a low-level USB preloader. This preloader can be used with MediaTek Flashing Tools to flash different firmware. We attempted to use the MediaTek Smart Phone Flash Tools to rip the firmware to see what is going on inside, but failed because we did not have the proper scatter file. Scatter files are partition specifications required for the flashing tools to execute. The Echo Dot requires a scatter file for the MT8163 V/B SoC. A scatter file from another MT8163 needs to be ripped and applied to the Echo Dot v2.