Security and Trust for Surveillance Cameras

Konstantin Boyarinov
British Columbia Institute of Technology
Burnaby, BC
kboyarinov1@my.bcit.ca

Aaron Hunter
British Columbia Institute of Technology
Burnaby, BC
aaron_hunter@bcit.ca

***Abstract*—We address security and trust in the context of a commercial IP camera. We take a hands-on approach, as we not only define abstract vulnerabilities, but we actually implement the attacks on a real camera. We then discuss the nature of the attacks and the root cause; we propose a formal model of trust that can be used to address the vulnerabilities by explicitly constraining compositionality for trust relationships.**

I. INTRODUCTION

In this paper, we are concerned with the notion of *trust* as it relates to the data obtained from a typical IP camera that might be used for surveillance. We give a concrete demonstration of the vulnerabilities, which shows exactly how easy it is to exploit the trust that users have in their surveillance cameras. It has been well-documented in the literature that such devices are vulnerable, which is a problem since cameras are used extensively in military and medical applications. We then suggest an abstract model of trust that may be useful in the analysis of such devices, demonstrating a useful connection with the Artificial Intelligence (AI) community.

II. CAMERA SECURITY

A. Practical Setup

In this study, we use the Reoloink RLC-410WS IP camera [1]. This is an inexpensive, commercial IP camera. While the camera comes with custom software, it also supports the open standard for sending video as proposed by the Open Network Video Interface Forum (ONVIF) [2]. By using an ONVIF-compliant camera, we hope to identify vulnerabilities that may exist across a variety of cameras used in a wide range of applications.

All of the vulnerability tests were conducted using three main devices: the attacking machine, the victim machine, and the camera itself. All of the devices are physically located on the same network.

B. Man In The Middle

The first vulnerability addressed is the susceptibility to a Man-in-the-middle (MiTM) attack. We use Address Resolution Protocol (ARP) spoofing as the first step in this attack. Any computer with access to the camera has an ARP table that records all of the IP addresses and the MAC addresses they are bound to. There can be a vulnerability here if there is no validation check on whether the MAC address that gets sent back by the actual destination address is correct. Using Ettercap, we were easily able to position our attacking machine as a MiTM.

C. Denial of Service

The fact that the attacking computer can be so easily positioned as a MiTM is problematic in terms of potential data leakage. Of course, the images sent from the camera are encrypted, so one might argue that this is not a terrible risk. The bigger problem is that once we have a MiTM, we can launch a variety of attacks. The first successful attack performed in this study was a Denail of Service(DoS) attack. By simply redirecting the camera output, we were able to block the desired recipient from receiving the camera's image feed. This was a straightfoward attack after establishing the MiTM.

D. Integrity Violation

While simply blocking the output of the camera from arriving at the victim machine can be damaging, a more troubling attack would be to block the data while simultaneously fooling the victim machine into believing the data is still live. We describe such an attack in this section.

The idea is simple. We simply capture a small collection of images from the camera, and then send them repeatedly in a loop to the victim machine. This is possible, because the camera that we are using sends video data in User Datagram Protocol (UDP) packets, which are then assembled and sequenced on the viewing machine. As such, in our attack, we saved 15 recent, sequential UDP packets containing video data. The result is a static image, as shown in Figure 1.

While this attack was somewhat successful, there were several problems:

1) Camera packets have a limited duration.
2) The state image flickers periodically.

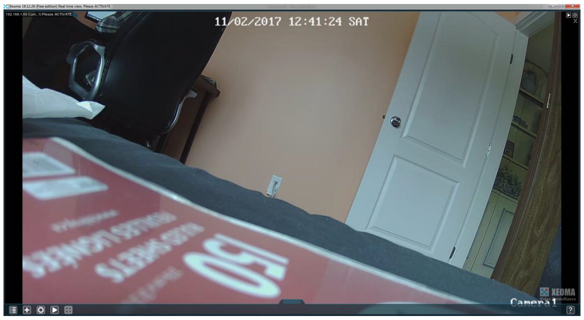

Fig. 1. Repeat Image

978-1-5386-0683-4/17/$31.00 ©2017 IEEE

3) If there is a person or moving object in the feedback loop, it will be detected authomatically.
4) The images have a time stamp. A user that looks at this time will immediately identify the attack.

We leave the improvement of this attack for future work.

E. Unauthorized Access

To this point, we have focused on manipulation of the images displayed to the user, by using a MiTM. These attacks are fundamentally related to trust; we return to this topic in the discussion section. However, we conclude with one final exploit that is not related to manipulating images at all.

The communication between the victim machine and the camera requires authentication; the victim machine actually has a username and password that are used for this purpose. However, in testing, we discoveredthat passwords are sometimes sent in plain-text format. In particular, this happens when using Xeoma Surveillance App to search for cameras. We emphasize that the Reolink camera is ONVIF-compliant, which means that it meets the agreed-upon standards for video capture and transmission. The fact that such a camera can send passwords in plaintext throws some doubt about the completeness of the standard.

F. Result Summary

The results of these tests are summarized in the following table:

Attack	Result
MiTM	Feed obtained without disruption
DoS	Camera unable to send data
Integrity Violation	Fake feed displayed
Unauthorized Access	Camera reveals password

The first, second and fourth attacks can all be considered successful without qualification. The most interesting attack is the third, which was only partially successful.

III. Discussion

The notion of trust is key to many problems in security, including the development of *reputation systems* where agents are trusted based on past actions [3]. Such models are not only concerned with profiling the reliability of agents, but they are often concerned with defending against deception [4].

The vulnerabilities discovered in this paper are a result of chaining together trust relationships. The end user of the camera trusts the machine that they use to provide correct information. They are actually quite right to do so: the machine displays the images that it receives. The problem is that the machine is receiving the wrong images. The trust breaks down at the level of the network connection to the camera.

We suggest that this problem can be analyzed at a formal level, by using precise logics of knowledge and belief to capture the trust that one agent holds in another. This has been addressed in [5], where so-called *trust partitions* are used to formalize the fact that agents only trust others on specific domains. This model actually captures an important notion of *compositionality* for trust.

Consider an exchange of information between three agents: $D \to M \to U$. It is possible in this case that U trusts M over a particular report r, but U does not trust D over the same report. This chain of exchange allows U to incorporate information from D, even though they would not do so under normal circumstances. The vulnerabilities defined in this paper are linked to flawed assumptions around this issuel: you can not simply compose trust relationships between different agents. At a formal level, one way to fix this is to add formal postulates to constrain the function Ψ. This is the approach that has been taken in the belief revision community since the pioneering work in [6]. By specifying suitable postulates, we can dictate how agents must handle chains of trust in order to avoid the vulnerabilities we have outlined here. We leave this formal investigation for future work.

A. Future Work

There are multiple directions for future work. First, as noted previously, the integrity violation attack could be improved. This would involve looking in more detail at the format of the UDP packets to address issues with the time stamp, and also reducing flicker to obtain a more natural image.

It will also be important to test other cameras. Some of the vulnerabilities here are specific to the particular model under consideration; this is the case for the Xeoma software, for instance. It is also the case that other cameras might actually prevent the ARP spoofing by checking the MAC address of the machines on the network. The tests presented here would all be simple enough to test on a wide range of cameras.

At a theoretical level, we intend to develop a precise logical framework that captures the form of trust that is implicit in the use of surveillance cameras. This kind of logical model of trust is useful in a variety of applications, and it would be particularly useful to have a rigorous approach to analyzing and establishing trust relationships with cameras.

B. Conclusions

In this paper, we have presented a hands-on vulnerability assessment for a commercial IP camera, and we have demonstrated some clear vulnerabilities. We have suggested that these vulnerabilities are fundamentally related to trust, and we have proposed that a suitable formal model would be useful for understanding and mitigating these risks.

References

[1] “Rlc-410ws official product page,” https://reolink.com/product/rlc-410ws/.

[2] “ONVIF core specification,” https://www.onvif.org/specs/core/ONVIF-Core-Spec-v210.pdf, 2011.

[3] T. D. Huynh, N. R. Jennings, and N. R. Shadbolt, “An integrated trust and reputation model for open multi-agent systems,” *Autonomous Agents and Multi-Agent Systems*, vol. 13, no. 2, pp. 119–154, 2006.

[4] A. Salehi-Abari and T. White, “Towards con-resistant trust models for distributed agent systems.” in *IJCAI*, 2009, pp. 272–277.

[5] A. Hunter and R. Booth, “Trust-sensitive belief revision,” in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015, pp. 3062–3068.

[6] C. Alchourrón, P. Gärdenfors, and D. Makinson, “On the logic of theory change: Partial meet functions for contraction and revision,” *Journal of Symbolic Logic*, vol. 50, no. 2, pp. 510–530, 1985.