Protecting networks against targeted attacks

Detecting advanced persistent threats

Wim Mees – Royal Military Academy
wim.mees@rma.ac.be

I. INTRODUCTION

The threats, faced by corporate and operational networks, evolve rapidly. In the past attackers would scan public networks for accessible services and try to exploit them. Later came viruses that tricked the human computer user into executing some code that would then infect the user's host and spread itself further. Fortunately, most organizations managed over time to harden their networks against these attacks by putting a strong perimeter defense in place, by adding anti-virus software, applying strict updating policies to hosts and gateways, and by organizing security awareness trainings for their users.

Lately however, we see a rise in very specific targeted attacks, that are custom-tailored to the victim and exploit one or more zero day vulnerabilities. These attacks can easily bypass the security controls that are currently in place and therefore present a major concern to network security engineers. In this paper we will first show how these attacks are organized and how they are addressed in our risk management processes. Finally we will present present a novel way of detecting compromised hosts in our network, that were the victim of such a targeted attack.

II. ADVANCED PERSISTENT THREATS (APT)

A typical APT consist of a message that is specifically written for the intended victim, who can be any member of the target organization with network access, and contains a spoofed sender identity, so the recipient will likely open the message. Upon accessing the information communicated through the message, for instance in the form of a pdf file in annex or a reference to a page on a website, a vulnerability in the visualization software used by the victim, will be exploited. As a result, the victim's computer will download a malware and run it. The human user will typically not even notice the infection.

Since it is a custom-made malware, its signature will not be received by an antivirus software manufacturer and therefore the infection can go unnoticed for several months and sometimes even years.

The malware is controlled from the outside. It can be used to infect other hosts over the corporate network, scan the network for available services, collect data and finally exfiltrate the gathered information. In order to receive its orders from its master, the malicious software will periodically initiate a command & control (CnC) connection to a server on the Internet. It will tunnel this CnC channel through a protocol that is allowed by the organization's perimeter security controls, for instance the hypertext transfer protocol (HTTP).

III. RISK MANAGEMENT IN THE PRESENCE OF APT

Information system security is addressed as a risk management problem. Based on a system and scope characterization, the vulnerabilities for the identified assets are listed, threat agents and probabilities are associated with vulnerabilities, and finally impacts are estimated resulting in a set of risks that need to be addressed for the given information system. Depending on the computed risk levels, the risks will be eliminated, reduced, accepted or transferred. The

reduction of a risk can target its probability, its impact, or both. The elimination or the reduction of risks is often realized through the addition of security controls.

In the case of APTs, eliminating the risk is impossible with currently available security controls. Indeed, on the one hand our staff needs to be able to communicate with the outside world and on the other hand zero-day vulnerabilities will continue to be discovered in the years to come. Therefore the APT risk will be addressed by reducing its probability of occurring through increased user awareness, but most importantly by attempting to identify compromised hosts. Indeed, knowing that nowadays compromised hosts often remain under the control of the attacker for months, it is clear that the potential impact is huge. If we can locate a compromised host within hours or days, the impact will obviously be strongly reduced. It is precisely that identification of compromised hosts through the detection of their CnC channel which is the main topic of this paper.

IV. APT DETECTION

The challenge when trying to detect the CnC channel of a previously unseen APT, is to separate it from the background noise of user initiated connections while surfing the web, as well as of other software initiated connections, for instance originating from operating system and client software automated updates, VoIP clients or cloud services that tunnel through HTTP, automatically refreshing objects in an open browser window like stock tickers or publicity, etc.

Since we want to detect previously unseen malwares, we cannot simply look for hand-crafted signatures or precise behavioural descriptions that result from the reverse engineering of a specific captured malware sample. We will rather implement a number of detection algorithms that look for generic behavioural patterns that we have observed over a wide variety of malware instances. Since a single such pattern will inevitably produce a large number of false alerts, and since moreover some malwares will not exhibit a behaviour that matches the given pattern, we will be aggregating evidence over a number of detectors. As an aggregation operator we use the Ordered Weighted Averaging (OWA) operator.

A number of detectors have been implemented for evaluating the degree of suspiciousness of HTTP requests, for instance based on frequency analysis, time-domain impulse detection, inverse flow detection, geographic outlier detection, and high fan-in, fan-out detection.

The system will not automatically classify hosts as clean or infected. Because of the characteristics of an APT, an approach based on the automatic thresholding at a predefined level that ensures an acceptable probability of detection would inevitably result in a high probability of false alert. Therefore the system is conceived as a semi-automatic data exploration system, with a human expert in the loop. The system just draws the attention of the human expert to potentially suspicious events and provides him with the tools that explore these further.

V. CONCLUSION

We have presented in this paper the problem of advanced persistent threats. The best control measure that is currently available against this type of targeted attacks is to detect a compromised host as quickly as possible in order to reduce the impact when an infection occurs.

For that reason we have developed a novel approach for identifying compromised hosts by detecting their command & control channel. The system is designed to have a human network

security analyst in the loop, who is presented with a number of agent outputs and can then explore in more detail the information available about the identified hosts.

The system is still under development but has nevertheless already produced interesting results.