An AI and data-driven approach to unwanted network traffic inspection

Francesca Soro - S255922

The growing number of connected devices on the Internet makes the end users more and more vulnerable to cyberattacks. Malicious entities in the network are constantly fostering the rise of new botnets and crafting new threats that, if successful, may directly impact critical infrastructures and people's everyday life. Detecting these attacks in real-time is paramount to properly counteract them, but it is also complex, due to the volume, variety and velocity of the data travelling on the Internet.

In this scenario, collecting and analyzing unwanted darknet traffic - often referred to as Internet Background Radiation - may be a path to take to detect new and potentially malicious phenomena. Observing the traffic hitting these fully passive probes, a network analyst can identify for instance heavy-hitter sources, service requests to vulnerable target or new coordinated events. Moreover, coupling the fully passive probes together with active honeypots further enriches the visibility on malicious events. In contrast to the darknets, honeypots are indeed able to reply to unwanted requests, providing a broader knowledge on the threat scenario. A manual inspection of these traffic traces is however impossible. For this reason, in this thesis I present a framework to automatically extract knowledge from unwanted traffic data captured on darknets and honeypots. Given the characteristics of the dataset at my disposal, I make extensive use of big-data and machine learning techniques to reach my goal.

I first provide a characterization of the traffic hitting fully passive probes located in three different parts of the world, highlighting the type of traffic they receive, the most targeted services and their differences and similarities. As a second step, I enrich my scenario with an active honeypot infrastructure, capable of replying to service requests with different levels of complexity. I demonstrate that actively engaging with the senders increases the volume of traffic, and more complex responses push the attackers to reach further attack stages.

After a thorough characterization of the most relevant network events, I proceed to the data analytics phase: I first depict the activity on the network as a graph, on top of which I test different community detection algorithms to group together similar activity patterns. By means of such techniques I am able to distinguish between communities devoted to vertical and horizontal scans, for instance, or recognize some more fine-grained patterns. As a final step, I benchmark a well-known set of anomaly detection algorithms against a novel AI technique, providing a set of custom metrics to quantify their detection capabilities, even when no ground truth is available.

My results demonstrate how the adoption of big-data and machine learning techniques ease the network traffic monitoring and analysis task, highlighting potentially critical events that would otherwise go unnoticed.