COLUMBIA UNIVERSITY
DATA SCIENCE INSTITUTE

Characterizing Network Behavior in Phishing Emails

Yifan Liu, Elisa Luo, Liane Young
Data Science Institute

Overview

- Phishing attacks are a targeted and pervasive threat to cybersecurity that aim to lure a targeted individual into revealing personal information.
- Past studies have primarily studied attacks from a behavioral perspective of the attackers or based on the content of emails [1]. Prior network-level analysis has focused on spam [2]
- Phishing attacks have evolved in complexity; both attackers and those defending from attacks have increasingly sophisticated methods.This project seeks to look at phishing attacks by focusing on network-level features, which may be more robust for detection and harder to manipulate by attackers

Data and Methods

- Email dataset: collected by Barracuda Networks in January 2020.
  - IP addresses identified in the “headers” field and mapped into an array, with the IP that was furthest from the receiver designated as the “sender” IP.
  - Private IP addresses were filtered from the dataset; they did not provide useful information for our analysis.
- Geolocation database: RIPE IPmap and Maxmind

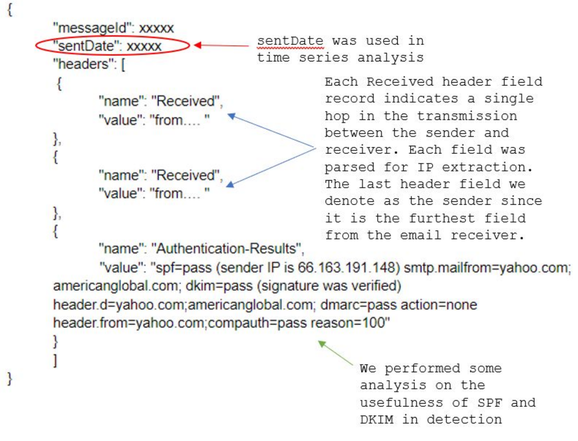

Figure 1. An example of an email header with annotations.

Where are attacks coming from?

The arrays of IP addresses parsed from the data were used to analyze several different factors:

1. The proportion of phishing and clean emails received over the course of the entire month, as a function of IP address space. Most IP address ranges that originate a significant amount of phishing emails also originate a lot of legitimate mail traffic, a few IP address ranges have significantly more malicious than legitimate mail (e.g., 20.*–30.*), and vice versa (e.g., 50.*–70.*). This characteristic suggests that it may be possible to use IP address ranges to distinguish phishing email from legitimate email.

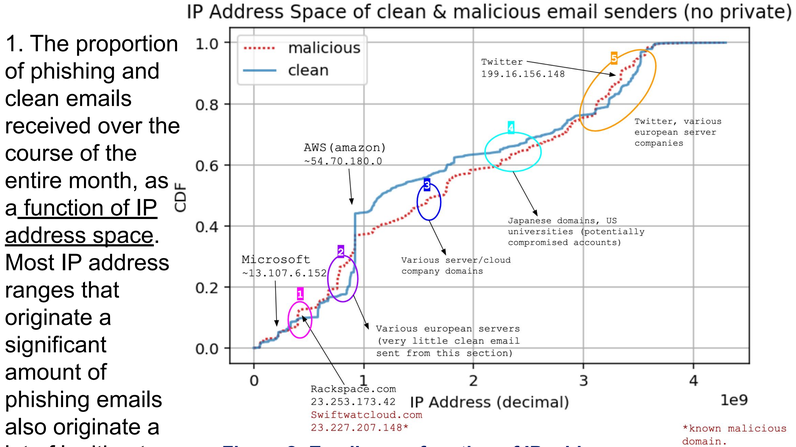

Figure 2: Emails as a function of IP address space

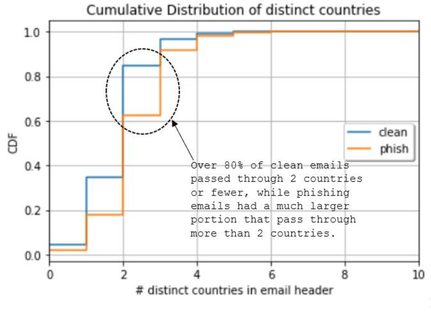

2. Clean emails skew towards having fewer distinct countries per header while phishing emails are more likely to pass through multiple countries from the recipient to the sender.

Figure 3. Analysis of the number of distinct countries an email passes through.

When do these attacks occur?

- Most clean emails sent follow a predictable diurnal pattern while malicious emails are usually sent in large, anomalous bursts or marked a significant deviation from an IP sender’s usual email sending pattern. In the show example from a Twitter IP, the sender started sending malicious emails at the end of the month.
- From this, various time-series season-trend decomposition and anomaly detection methods to identify time periods an IP may be sending phishing emails.

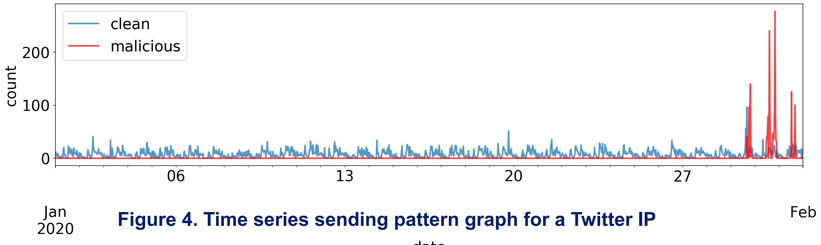

Figure 4. Time series sending pattern graph for a Twitter IP

Conclusions and Next Steps

- We have identified key network-level features that can be used in a classifier to detect phishing emails.
- These features will be used in a classification model to improve existing abilities to detect phishing attacks. We also have ongoing work in several other areas including expanding the geolocation analysis to look at patterns in countries that are used in phishing attacks.

Acknowledgments

We would like to acknowledge our faculty mentors Asaf Cidon and Ethan Katz-Bassett, and our project mentors Marco Schweighauser, Mohamed Ibrahim, and Grant Ho for their support and guidance.

References

[1] Grant Ho, Asaf Cidon, Lior Gavish, Marco Schweighauser, Vern Paxson, Stefan Savage, Geoff M. Voelker and David Wagner. Examining Lateral Phishing at Scale. Usenix Security 2019

[2] Ramachandran, Anirudh & Feamster, Nick. (2006). Understanding the Network-Level Behavior of Spammers. Proceedings of ACM SIGCOMM. 36. 10.1145/1151659.1159947.