Team develops a detector that stops lateral phishing attacks

29 October 2019, by Robert Florida

Credit: CC0 Public Domain

Lateral phishing attacks—scams targeting users from compromised email accounts within an organization—are becoming an increasing concern in the U.S.

Whereas in the past attackers would send phishing scams from email accounts external to an organization, recently there's been an explosion of email-borne scams in which an attackers compromise email accounts within organizations, and then uses those accounts to launch internal phishing emails to fellow employees—the kind of attacks known as lateral phishing.

And when a phishing email comes from an internal account, the vast majority of email security systems can't stop it. Existing security systems largely detect cyber attacks that come from the outside, relying on signals like IP and domain reputation, which are ineffective when the email comes from an internal source. Lateral phishing attacks are also costly. FBI data show, for instance, that these cyberattacks caused more than $12 billion in losses between 2013-2018. And in the last two years, the attacks have resulted in an increase of 136 percent in losses.

To alleviate this growing problem, Data Science Institute member Asaf Cidon helped develop a prototype of a machine-learning based detector that automatically detects and stops lateral phishing attacks.

The detector uses several features to stop attacks, including detecting whether the recipient deviates from someone an employee would usually communicate with; whether the email's text is similar to other known phishing attacks; and whether the link is anomalous. The detector can detect the vast majority of these attacks with a high precision rate and a low false positive rate—under four false positives for every one-million employee-sent emails.

Cidon was part of a research team that analyzed a dataset of 113 million employee-sent emails from nearly 100 businesses. They also characterized 147 lateral phishing incidents, each of which involved at least one phishing email. The study was conducted jointly with Barracuda Networks, a network security company that provided data on its customers to the researchers with the goal of developing a detector for lateral phishing.

The researchers also wrote a paper about the study, Detecting and Characterizing Lateral Phishing at Scale, which recently won a Distinguished Paper Award at Usenix Security 2019, a leading cybersecurity conference.

"The attacks analyzed in this study represent one of the most difficult types of cyber attacks to detect automatically, since they emanate from within an internal employee's account," said Cidon, an Assistant Professor of Electrical Engineering and Computer Science (jointly affiliated) at Columbia Engineering as well as a member of the Data Science Institute. "The key to stopping such targeted socially-engineered attacks is to use

machine-learning based methods that can rely on the unique context of the sender, recipient and organization."

When attackers launch a phishing attack, their objective is to convince the user that the email is legitimate and to cajole them into performing a certain action. What better way to convince a user that an email is legitimate, therefore, than by using a hacked email account from a colleague they know and trust. And in lateral phishing, attackers leverage a compromised email account to send phishing emails to other users in the organization, benefiting from the implicit trust of colleagues and the information in the hijacked user's account. The classifiers that Cidon helped to develop look for anomalies in communication patterns. For instance, the classifiers would flag an employee suddenly sending a burst of emails with obscure links or an employee systematically deleting emails from his or her sent items folders—trying to mask their scams.

Drawing on these kind of phishing attacks, as well as from a collection of user-reported incidents, the researchers used machine learning to quantify the scale of lateral phishing, identifying thematic content and recipient targeting strategies that attackers used. They then were able to characterize two strategies that attackers used to tailor their attacks: content and name tailoring. Content tailoring is how the attacker tailors the content of the email to compel the recipient to click on the link and fall for the phishing email. The most common content tailoring they discovered was a generic phishing content (for example, "You received a new document, click here to open"). But they also found that some attackers tailored the email to the specific context of the organization (e.g., "Please see the attached announcement about Acme's 25th year anniversary"). Name tailoring is how the attackers personalize the email to a recipient by using his or her name and role in the organization (e.g., "Bob, please review the attached purchase order," and in this instance Bob works in accounting).

Some key findings from their analysis of more than 100 million emails that compromised nearly 100 organizations include:

- More than 10 percent of incidents result in a successful additional internal compromise (this is orders of magnitude higher percentage than attacks originating externally).
- The majority of attacks are relatively simple phishing emails. But a significant percentage of attackers do heavily tailor their emails in accord with the recipient's role and the context of the organization.
- More than 30 percent of attackers engage in some kind of sophisticated behavior: either by hiding their presence in the attack (e.g., deleting outgoing emails) or by engaging with the recipient of the attack to ensure it is successful.

Cidon says these kinds of attacks represent the new frontier of cyber crime: highly personalized attacks where attackers are willing to spend days and weeks "doing reconnaissance."

"In this study we focused on link-based lateral phishing," adds Cidon. "There's still a large amount of work to do, however, in exploring attacks without links or attacks that combine other social mediums such as text messages and voice. But we hope our detector helps combat the growing scourge of lateral phishing attacks."

Provided by Data Science Institute at Columbia

APA citation: Team develops a detector that stops lateral phishing attacks (2019, October 29) retrieved 22 June 2021 from https://techxplore.com/news/2019-10-team-detector-lateral-phishing.html