Email Campaign Explorer for Detecting Malicious Email Campaigns

Awalin Sopan Parnian Najafi
{awalin.sopan, parnian.najafi}@fireeye.com
FireEye, Inc

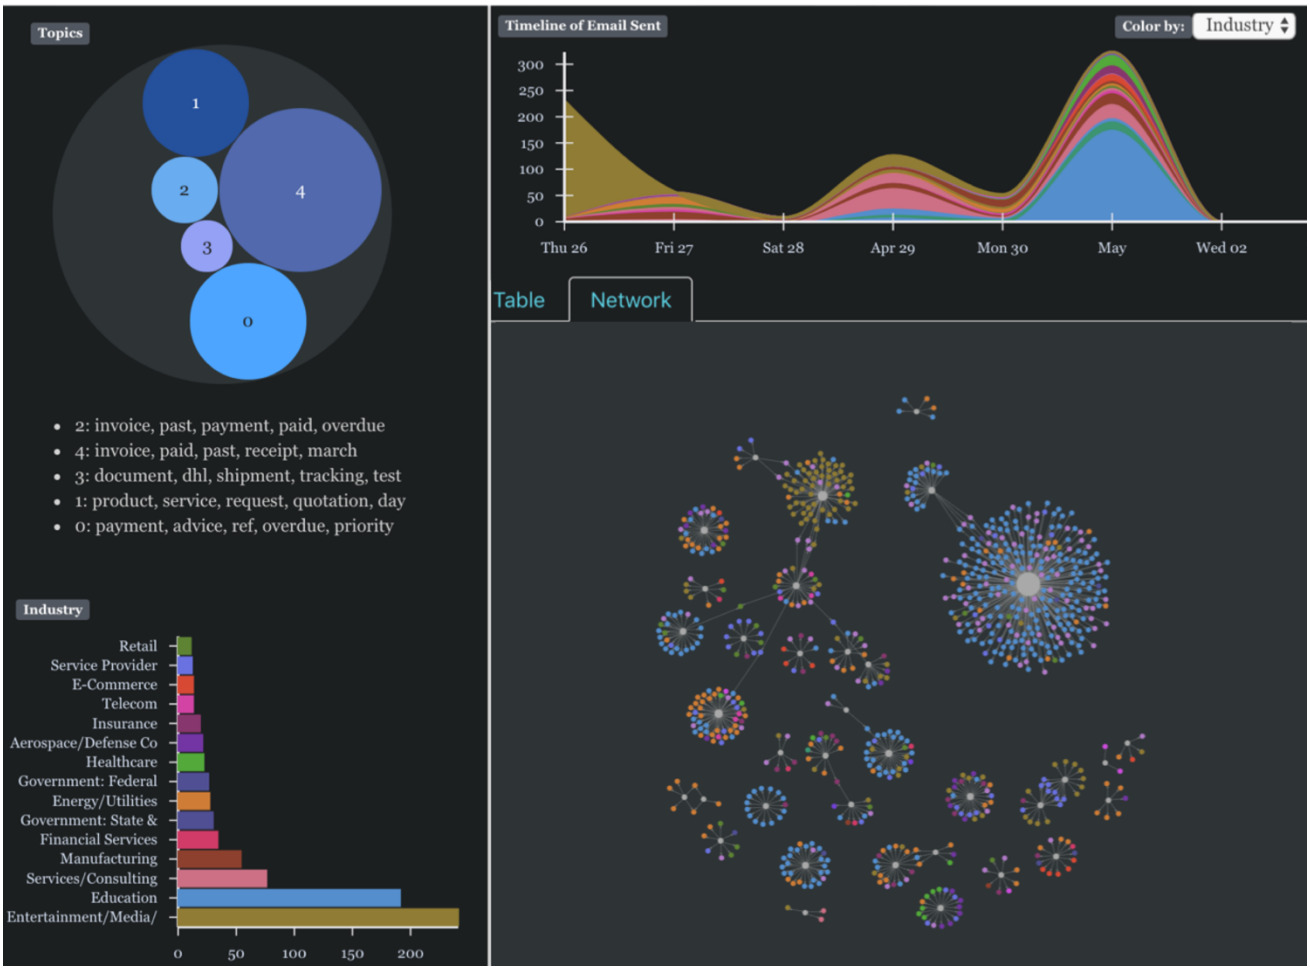

Figure 1: User interface of Campaign Explorer: Top-left: topics-cluster view in circle packed bubble chart showing topic as the most frequent topic with words like 'invoice, paid, receipt'. Top-right: time series visualization of the malicious emails color-coded by Industry of the email receiver, showing which industry received more malicious emails at a given time period. Bottom-left: Malicious email distribution grouped by and color-coded by Industry. Bottom-right: Email network of sender of malicious email and victim receiver, color coded by the industry of the email receiver. The active tab is the Network tab. Users can switch between Network and Table view tabs to see more details for each email.

Abstract—This paper demonstrates a work in progress for using visual analytic approach to detect email campaigns. We have developed a system that uses Natural Language Processing (NLP) to clusters emails with similar topic, network metrics to detect the important players, and statistical aggregation to find pattern and anomaly. Finally, we display these finding to the analysts in our exploratory visual analytics web-application.

1 Introduction

Email remains one of the most challenging entry points for organizations to protect as attackers develop new techniques to evade defenses. Email Campaign Explorer has been developed to aid in email traffic analysis by identifying potential trends or patterns of techniques used by malicious actors. There are a variety of attacks that can be done using email that can be categorized into malware attacks and malware-less attacks [1]. Malware attacks use email with malware-infected attachments to gain access to a user's computer. These attacks can be hidden within a variety of file types such as Microsoft documents and PDFs. Malware-less attacks rely on user action to gain access to information or company assets. This can be done by impersonating a trusted sender, tricking the user to personally send or enter in their credentials or linking to a malware-infected page [2]. When analysts investigate cyber-attacks, it is important to understand the temporal context [3] and textual content in case of email-related ones [4]. Hence, we developed a system-prototype to aid analysts in investigations by calculating important network-metrics and clusters emails based on the email subjects. Email Campaign Explorer presents the raw data in a table format and the details and aggregation in several coordinated visualizations to help the analysts' exploratory analysis process.

2 User Interface

The Email Campaign Explorer has three major user interface components (Figure 1):

1. A time series visualization
2. Aggregated statistics of various metrics
3. A tabbed view of the email data in:

a) Tabular view and
b) Network view

Time Series Visualization: The timeseries view shows the total number of emails per day by stacking them by Industry or Country of the email receivers. This stacked area chart can be color coded based on the user Country or Industry. This feature works helps analysts find a specific campaign activity affecting a country or industry. By default, it shows the last 24 hours of data.

Aggregated Statistics: The leftmost vertical panel shows various aggregated statistics of the data (see Figure 1 (left)) including:

- The topic cluster distribution based on email subject
- Distribution of emails by:
  - Receivers' countries
  - Receivers' industries
  - Senders' domains

Email topics have been derived using Scikit-learn library and are shown in a circle-pack diagram. Email topics have irregular distribution. For instance, for a sample used lots of emails contained shipment related information. The top ten categories are shown in other distribution charts such as top countries and industries. The top ten categories are shown in other distribution charts such as top countries and industries (Figure 2). This way we can see which industries or countries are attacked the most in the selected timeframe.

Tabbed Panel-> Table view: This is the detail view of all the email data. By default, the tab with the table view is visible to the user. Our users are analysts who are used to explore data in spreadsheet-like user interface, to keep this familiarity, we decided to use the table view as the default tab. Certain columns are shown by default in the web-app, The columns that are shown by default are sender, receiver, md5, industry, country, customer, subject, data source, URL/Attachment in the email, date sent, signature (that detected the email as malicious), domain (of the sender's email), and node-degree (node-degrees are calculated using NetworkX Python package, in this case it is the number of receivers a specific sender has targeted) can be shown from the columns section on the right.

Tabbed Panel-> Network View: In this force-directed network diagram of email sender-receiver, sender-nodes are depicted with grey circles and receiver email-accounts are color-coded based on their industry/country; node-size is proportional to their degree. Figure 3 shows tax-related phishing email activity by filtering the data by email-subjects containing the word 'tax'. Here the nodes and timeseries are color-coded by industry, and we see that 'Insurance' is the most attacked industry in our dataset and the most active sender is attacking various industries (the biggest node in the middle of the largest ego-network in the network diagram), where a few senders are only attacking financial services (green nodes).

Color Coding: Users can configure the color coding using the drop-down at top-right corner. Currently the available coding's are for Industry (default) and Country that uses a categorical color scheme. Based on this selection, the time series, network diagram, and histograms will be color coded accordingly.

Filtering: The column headers in the table view allow the users to filter the data with various complex conditions. Users can perform AND/OR operations using the table column headers. This advanced filtering feature is implemented after the users' feedback. Once the table is filtered, all the views are updated based on the data shown in the table.

3 Limitations and Future Work

This is a work in progress and uses a small email-dataset from our clients. The results do not reflect the entirety of the email campaign trends, but they are a good representation of the developments in the industry or countries. This is a prototype and further work is needed to make it scalable. The next version would use the feedbacks from analysts investigating Spam and phishing emails on a daily basis. We hope this combination of machine learning and exploratory visualization will allow them to investigate retroactively, correlate with other email campaigns to find the root cause and predict future attacks based on their observed email campaign patterns.

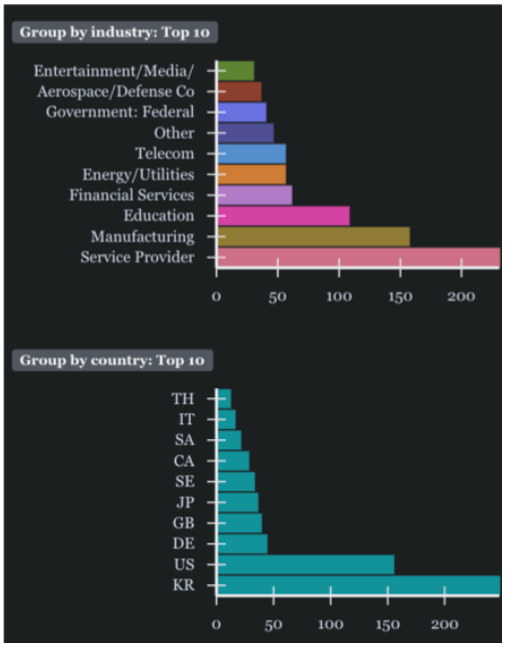

Figure 2: Aggregation of various features. Color code by receivers' industry.

Figure 3: Top: Filtered table to emails containing the text "tax" in the subject field. Bottom: a few senders posing as IRS, targeting various industries. The timeseries shows a spike on May 24, 2018 (some weeks after US tax return).

Acknowledgement

We cordially thank Raj Jammalamadaka & Prashant Shreedharan.

References

[1] Jason Milletary, Technical Trends in Phishing Attacks.
[2] Jason Hong. 2012. The state of phishing attacks. Commun. ACM 55, 1 (January 2012), 74-81. DOI: https://doi.org/10.1145/2063176.2063197
[3] Lu, J., Chen, K., Zhuo, Z. et al. A temporal correlation and traffic analysis approach for APT attacks detection. Journal of Cluster Computing (2017).
[4] Izzat Alsmadi, Ikdam Alhami. Clustering and classification of email contents, Journal of King Saud University - Computer and Information Sciences, Volume 27, Issue 1, 2015.