Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
15 FEB 2015	**N/A**	

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Linking Social Media Reports to Network Indicators of DoS Attacks	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
Ritter /Evan Wright Rhiannon Weaver Alan	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited.

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**SAR**	**1**	

Linking Social Media Reports to Network Indicators of DoS Attacks

Evan Wright (CERT), **Rhiannon Weaver** (CERT), **Alan Ritter** (Ohio State University)

Introduction and Background

Social awareness of denial of service (DoS) as a cybersecurity threat has led to public reporting in fast-paced social media such as Twitter, but these reports are rarely linked to quantifiable network behavior. A data set of network-based vs. media-reported DoS attacks can help researchers determine the prevalence of DoS tactics such as IP spoofing, the intensity and duration that leads to media reporting, and the types of organizations are reported or under-reported. We used heuristics and machine learning methods to link DoS-related tweets to flow-based evidence, collected from a large private network, of DoS activity targeting entities extracted from those tweets. Preliminary results show promise for novel data visualization and future refinement for formal inference.

Methods and Data for Multi-Step Correlation

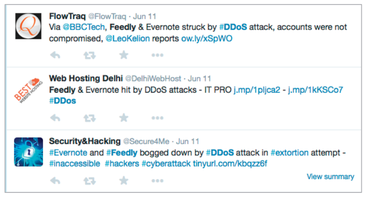

Tweets

We compiled tweets from a seed event table as well as those with the hashtag #DDoS (Ritter et al) for the ranges of April 7 - July 20 2014, and October 21 - November 9, 2014.

Date	Entity	Tweets
2014/06/11	feedly	26
2014/06/11	evernote	36

Entities

We employed clustering methods and natural language processing (NLP) to map the tweets into unique date-entity pairs.

Result: 533 unique date-entity pairs: *(D, E)*

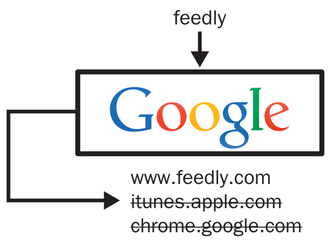

Domains

We used a Google API to return the top three domain names for each entity in the previous step (with a whitelist of Wikipedia, youtube, etc.).

Result: 355 *(D, E)* pairs mapped to at least one domain name: *(D, E, N)*

IP	Domain	Range
65.19.138.1	com.feedly	02/01/14-06/11/14
65.19.138.2	com.feedly	02/01/14-06/11/14
108.162.200.248	com.feedly	06/11/14-07/03/14

IP Addresses

We used the Security Information Exchange (SIE) passive DNS data (February–October 2014) to link domains from the previous step to IP addresses and dates.

Result: 345 *(D, E, N)* tuples mapped to at least one IP address: *(D, E, N, I)*

IP	Date	Inbound dIP	Outbound sIP
65.19.138.2	06/11/14:09	10450	117890
65.19.138.2	06/11/14:10	5361	60504

(Flag IP, hour with >500 unique sIP, dIP)

Network Activity

We used SiLK to find those IP addresses with evidence of backscatter from random spoofing: TCP SYN-ACK flags sent to >500 internal machines in a large private network, with corresponding RST flags sent from >500 internal machines (see Moore et al).

Result: 178 *(D, E, N, I)* tuples mapped to backscatter flow events: *(D, E, N, I, F)*

Analysis and Results

Preliminary Evaluation Data
Random sample of 30 *(D, E)* pairs yielding 21 unique entities *(E)*

	Targets	Non-Targets	Other
All *(E)*	**6**	**4**	**11**
Valid *(E, N)*	6	3	0
Spurious *(E, N)*	2	1	8
Valid *(E, N, I, F)*	4	0	0
Spurious *(E, N, I, F)*	1	0	4

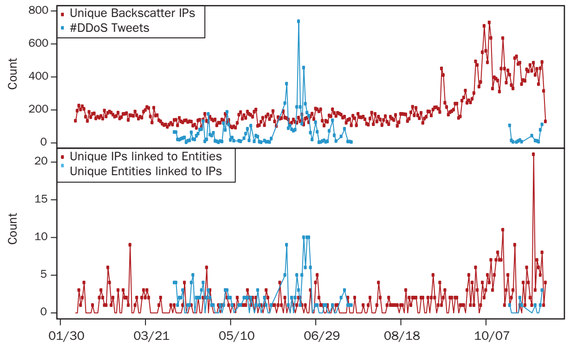

Individual Campaigns

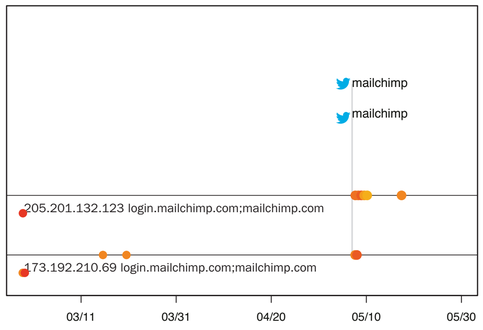

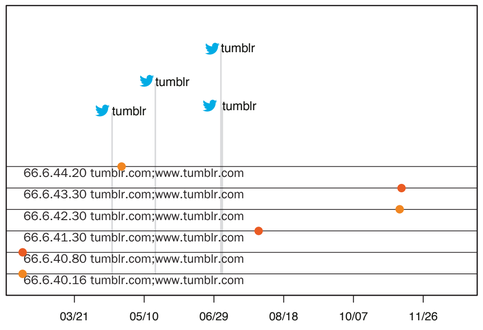

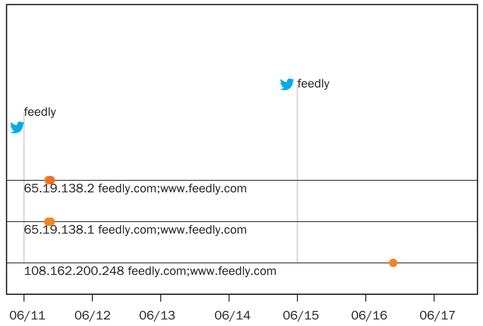

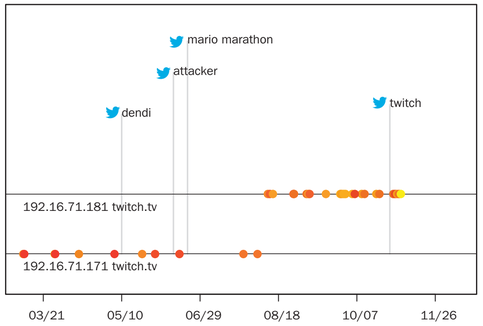

Future Work

In this analytic, there are many potential reasons for finding or failing to find connections in each of the correlation steps. We plan to conduct a detailed error analysis to learn these reasons and to understand their relative likelihoods for formal statistical inference. As suggested by the preliminary evaluation, our immediate next steps will focus on improving both the extraction of specific targeted entities from tweets, as well as the correlation of those entities to domains.

References

Moore, D., C. Shannon, D. Brown, G. Voelker, and S. Savage. 2006. “Inferring Internet Denial-of-Service Activity.” *ACM Transactions on Computer Systems* 24 (2): 115–39.
Ritter, Alan, Evan Wright, William Casey, and Tom Mitchell. 2014. “Weakly Supervised Event Extraction.” In review.

CERT | Software Engineering Institute | Carnegie Mellon University

Contact: Evan Wright | ewright@cert.org, Rhiannon Weaver | rweaver@cert.org