Open-source Measurement of Fast-flux Networks While Considering Domain-name Parking

Leigh B. Metcalf
Dan Ruef
Jonathan M. Spring

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

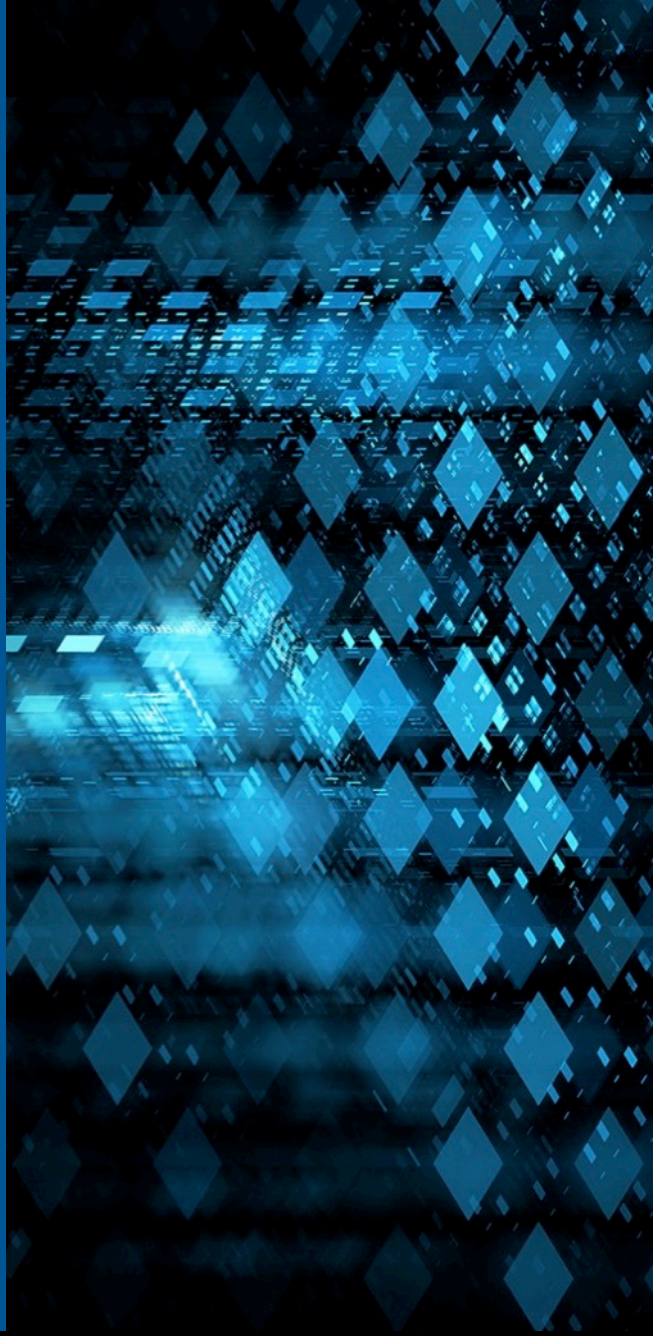

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM17-0800

Background

Fast Flux – “a network of compromised computer systems with public DNS records that are constantly changing”

Domain Parking – resolves to IP not controlled by domain owner

- on private IP address space, for suppressing errors
- NOT the same as the more common redirect to routeable IP addresses for generating ad revenue

Fast flux was widely studied, but never measured at scale

Parking would interfere with fast flux measurements

- (This is mentioned as an aside in the early Honeynet Project work on Fast Flux)

Data sources and tools

Security Information Exchange

- Passive DNS records stored from Jan 1, 2012 to June 30, 2017
- Monthly, unique FQDNs between 550 million and 1,000 million

Autonomous System Number, from Routeviews & RIPE RIS

- Produce data structure labeling every IP address with its ASN(s)

Blacklists

- Create 6-month chunks of unique blacklist domains / IPs
- See our WISCS paper for details of collection and content

Alexa list of popular domains

Analysis Pipeline

- Open-source tool in which much of the analysis is implemented

Method – Parking

1. Extract IPv4 A RRsets with answer in private IP space
2. Remove whitelisted domains
3. Find all RRsets associated with remaining domains within three weeks into future
4. Flag any domain that transistions from private to public IP space
   - Technically, from non-routeable to routeable

Pipeline config file available in the paper

Method – Fast-flux networks

1. Network is a graph of IPs, ASNs, and FQDNs that are associated by DNS A records, and IPs mapped to ASNs
2. Whitelist any non-routeable IP addresses
3. Whitelist any FQDN that is on the Alexa top 24,000 for 330 of the past 365 days (or is a subdomain of such a domain)
4. Any network with over 500 IPs, 23 ASNs, and 667 FQDNs is marked as a fast-flux network
   - The time frame is roughly within 3-5 hours
   - Roughness is due to layers of deduplication in DNS data

Results – Flux

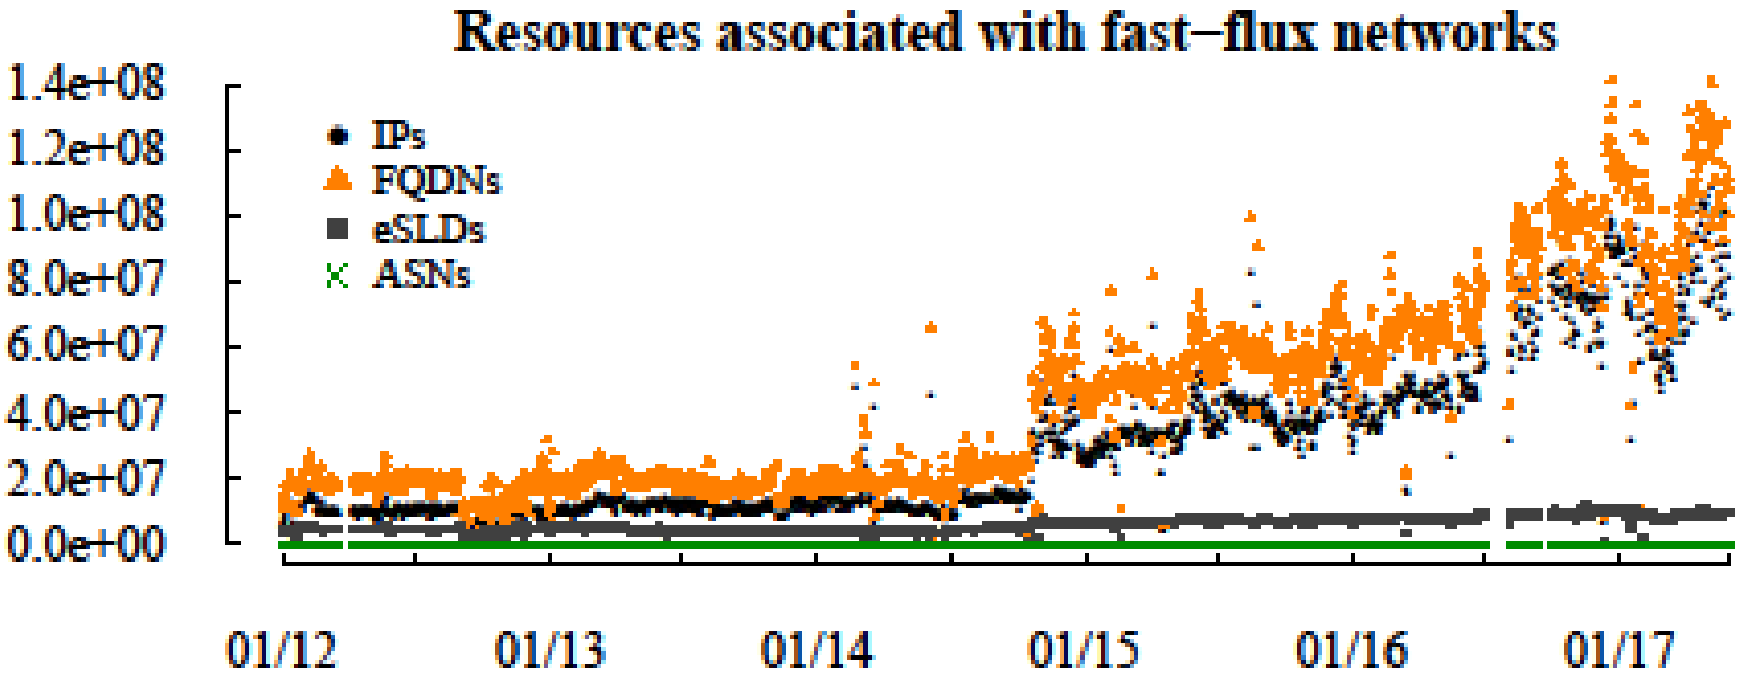

Figure 3: Total unique network resources of different types associated with fast-flux networks every day. Gap in July 2016 is a collection error.

Results – Parking

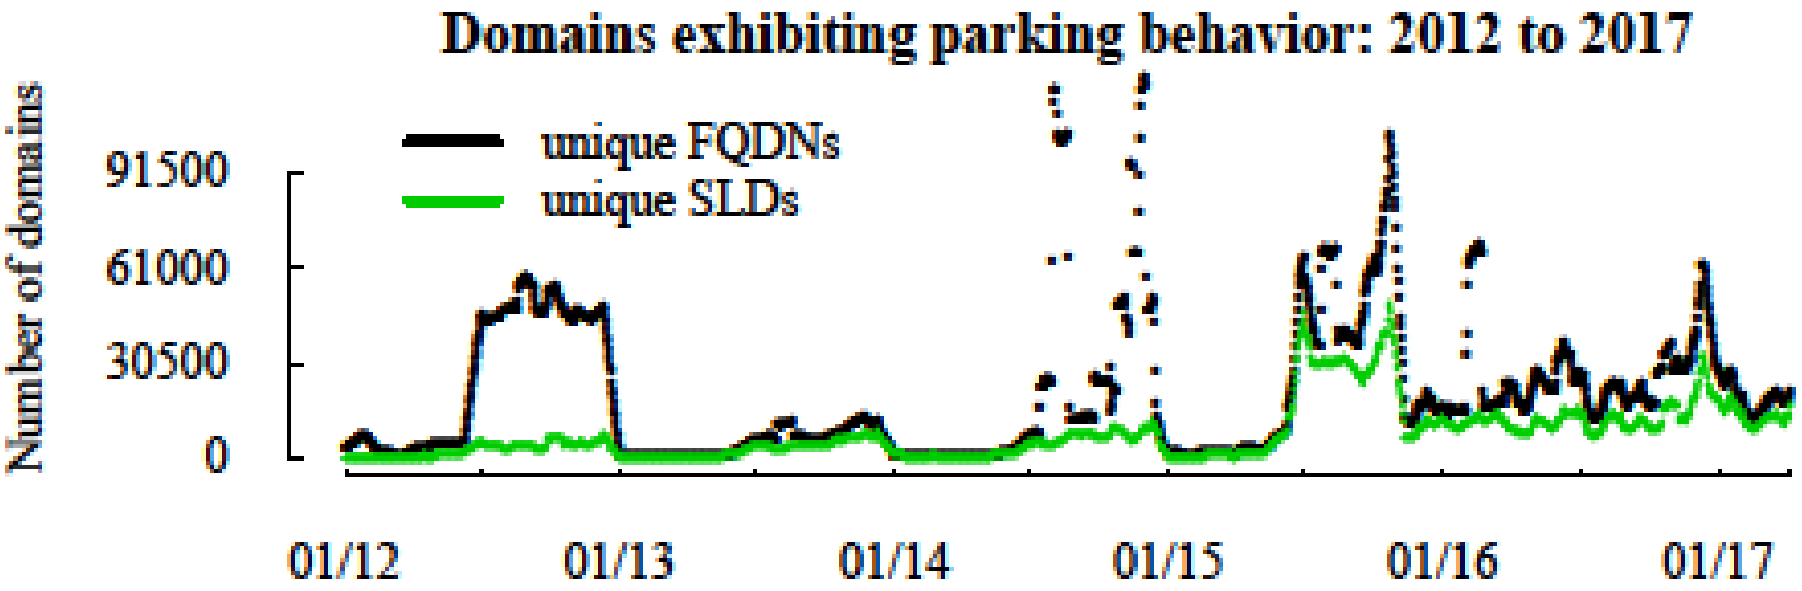

Figure 2: Three-week rolling window of unique domains and eSLDs exhibiting parking behavior.

Results – Is flux a known-bad?

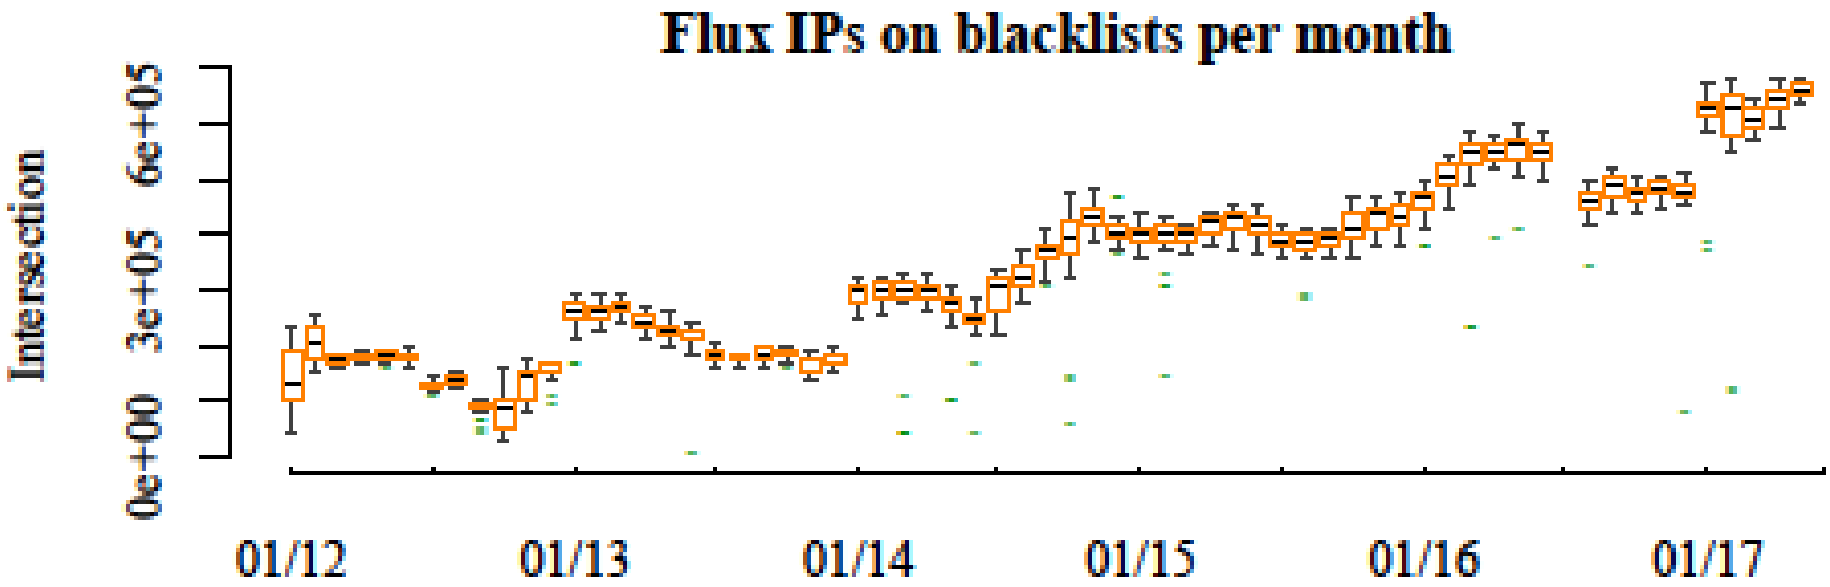

Figure 4: Summary statistics of fast-flux–blacklist intersection. Each day's flux from January through June 2013 is intersected with all 1H2013 blacklist data, for example, and the daily intersections are summarized in monthly box plots. The whisker length is 1.5 times the inter-quartile range (IQR).

Discussion

Fast flux networks are still widely used

- Not clear how to interpret lack of blacklist overlap

Domain parking on private IP address space is uncommon

- Unlikely collision in fast flux measurement

Ideally, defenders run our Pipeline analysis to identify fast-flux networks, and block resolutions to new domains when they would be added to a known network (see paper for example config)

Thanks for your time.

Questions?