Characterization of Blacklists and Tainted Network Traffic

Jing Zhang[1], Ari Chivukula[1], Michael Bailey[1], Manish Karir[2], and Mingyan Liu[1]

[1] University of Michigan
[2] Cyber Security Division, Department of Homeland Security

Motivation

- Network reputation blacklists
  - Scale: Hundreds of providers
  - Widely adopted: DNS, Mail Server, Browser, Anti-Virus …

- .com and .net TLD queries

Rank	Domain Name
4	manitu.net
5	sorbs.net
35	secureserver.net
53	zenon.net
73	immunet.com

Query Example:
[IP address].sorbs.net

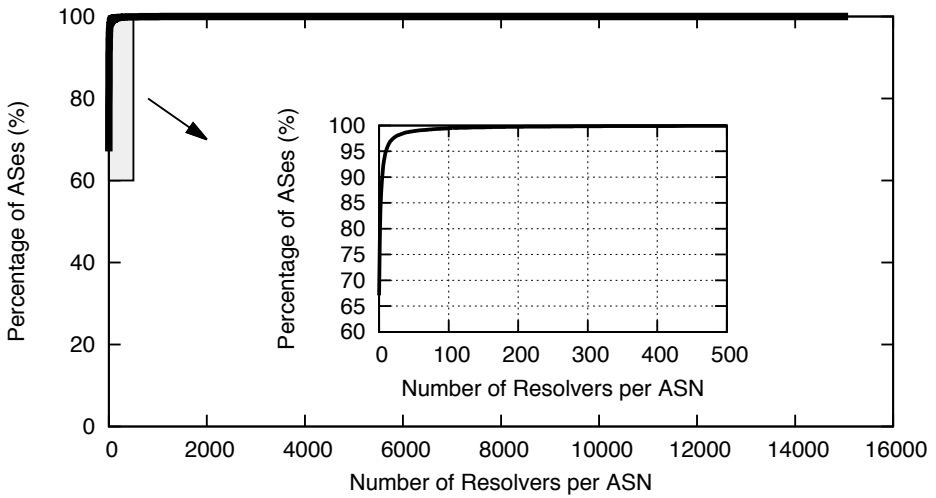

33% of ASes had at least one resolver queried the blacklisting domains

What is Missing?

- Researches on Reputation Blacklists
  - How to create them? [Antonakakis 2010, Craig 2012, Zhang 2008]
  - How effective are they? [Jung 2004, Sinha 2008]

- *What is missing?*
  - *Properties of the blacklists*
    - How dynamic are they?
    - How consistent are the bad networks?
    - What is the overlap between different lists?
  - *Impact of reputation*
    - What will happen if we apply filtering policies?

To answer these questions, we need:
- Multiple reputation blacklists
- Real-world network traffic

Agenda

1. Data Collection
   - Reputation Blacklists
   - Network Traffic
2. Properties of Reputation Blacklists
   - Timing
   - Region
   - Overlap
3. Impact of Reputation
   - Tainted Network Traffic
   - Heavy Hitting IP Addresses
4. Conclusions & Discussions

Agenda

1. Data Collection
   - Reputation Blacklists
   - Network Traffic
2. Properties of Reputation Blacklists
   - Timing
   - Region
   - Overlap
3. Impact of Reputation
   - Tainted Network Traffic
   - Heavy Hitting IP Addresses
4. Conclusions & Discussions

Data Collection

- The data in our study is collected at Merit Networks
  - A large regional ISP located in Michigan, USA
  - Over 100 customers, including educational, government, healthcare and non-profitable organizations
  - Load: 4 Gbps – 8 Gbps

- A period of one week starting from June 20, 2012

Data Collection

- Reputation Blacklists
  - Fetching directly from the publisher on a daily basis
  - Three broad classes of malicious network activities

Classes	Blacklists
SPAM	CBL, BRBL, SpamCop, WPBL, UCEPROTECT
Phishing/Malware	SURBL, PhishTank, hpHosts
Active attacks	Dshield

- Network Traffic
  - Collected via NetFlow with a sampling ratio of 1:1
  - 118.4TB traffic with 5.7 billion flows and 175 billion packets

Agenda

1. Data Collection
   - Reputation Blacklists
   - Network Traffic
2. Properties of Reputation Blacklists
   - Timing
   - Region
   - Overlap
3. Impact of Reputation
   - Tainted Network Traffic
   - Heavy Hitting IP Addresses
4. Conclusions & Discussions

Timing Properties

- *Q1: How stable are the blacklists with respect to their size?*

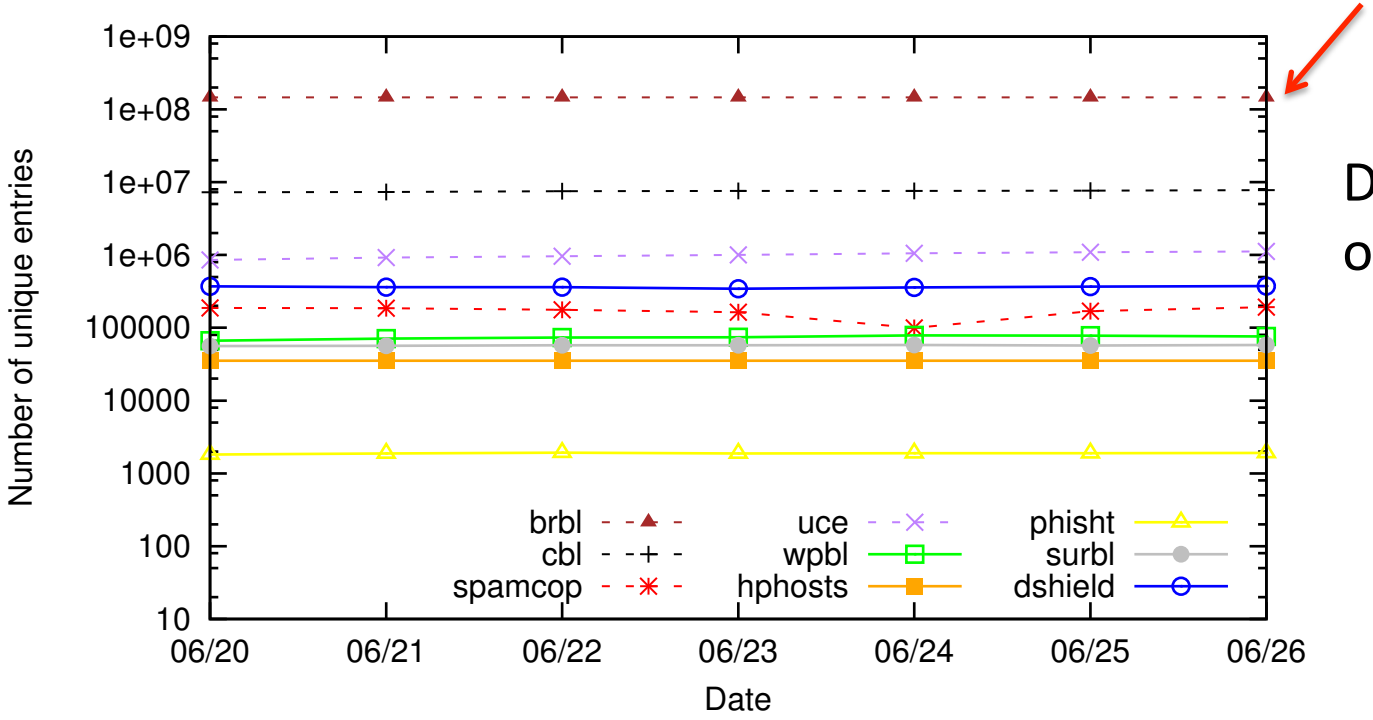

Daily number of unique IPs

- The size varied across different lists
- The size of each blacklist was consistent

Timing Properties

- *Q2: How persistent are the blacklisted IP addresses?*

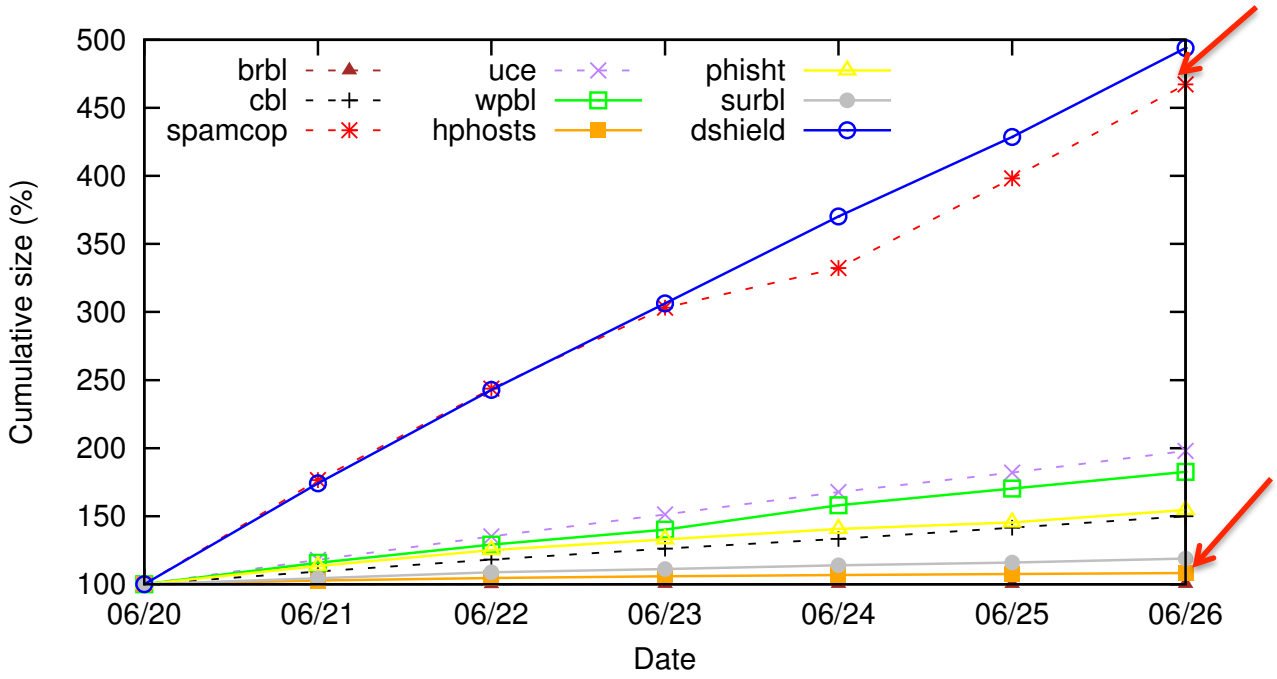

Cumulative size over one week

- Spamcop and Dshield updated aggressively (500% turnover)
- Some lists are relatively static (< 110% turnover)

Regional Characteristics

- *Q3: What is the distribution of malicious IPs over registries?*

	Spam					Phishing/Malware			Active
	BRBL	CBL	Spamcop	UCE	WPBL	hpHosts	Phisht	SURBL	Dshield
AFRINIC	3.02	7.70	5.89	6.37	4.19	0.20	0.58	0.04	2.19
APNIC	25.20	47.14	51.94	48.45	51.27	8.45	11.56	5.58	36.19
ARIN	6.23	1.05	2.53	1.84	6.17	53.32	43.93	54.70	13.54
LACNIC	17.11	16.19	12.15	15.89	10.59	1.66	5.32	1.44	8.54
RIPENCC	48.44	27.93	27.50	27.44	27.77	36.37	38.6	38.24	39.53

Regional Distribution of IPs for each blacklists (%)

- APNIC (Asia/Pacific) and RIPENCC (Europe) have more IPs that involved into SPAM and Active attacks
- ARIN (North America) and RIPENCC (Europe) are the most common regions for Phishing/Malware

Overlap

- *Q4: How many IPs is each blacklist are overlapped with others?*

	Spam					Phishing/Malware			Active
	BRBL	CBL	Spamcop	UCE	WPBL	hpHosts	Phisht	SURBL	Dshield
BRBL	100.0	75.2	94.6	89.8	93.8	5.3	10.0	30.7	33.2
CBL	3.9	100.0	98.1	91.7	70.2	0.5	0.7	6.2	9.3
Spamcop	0.1	2.3	100.0	12.6	21.5	0.1	0.1	0.8	1.2
UCE	0.6	12.1	69.4	100.0	50.6	0.3	1.5	1.2	4.8
WPBL	0.0	0.7	8.8	3.7	100.0	0.0	0.2	0.9	0.4
hpHosts	0.0	0.0	0.0	0.0	0.0	100.0	33.7	7.3	0.0
Phisht	0.0	0.0	0.0	0.0	0.0	1.8	100.0	1.7	0.0
SURBL	0.0	0.0	0.3	0.1	0.7	11.8	52.8	100.0	0.1
Dshield	0.1	0.4	2.4	1.8	2.2	0.4	0.7	0.3	100.0

The average % (of column) overlap between blacklists (row, column)

- The overlap within the same class of blacklists was significantly larger than the overlap among different types
- The two largest blacklists – BRBL and CBL, covered most of the entries in other Spam-related lists

Agenda

1. Data Collection
   - Reputation Blacklists
   - Network Traffic
2. Properties of Reputation Blacklists
   - Timing
   - Region
   - Overlap
3. Impact of Reputation
   - Tainted Network Traffic
   - Heavy Hitting IP Addresses
4. Conclusions & Discussions

Tainted Network Traffic

- *Q5: What fraction of traffic carries a negative reputation?*

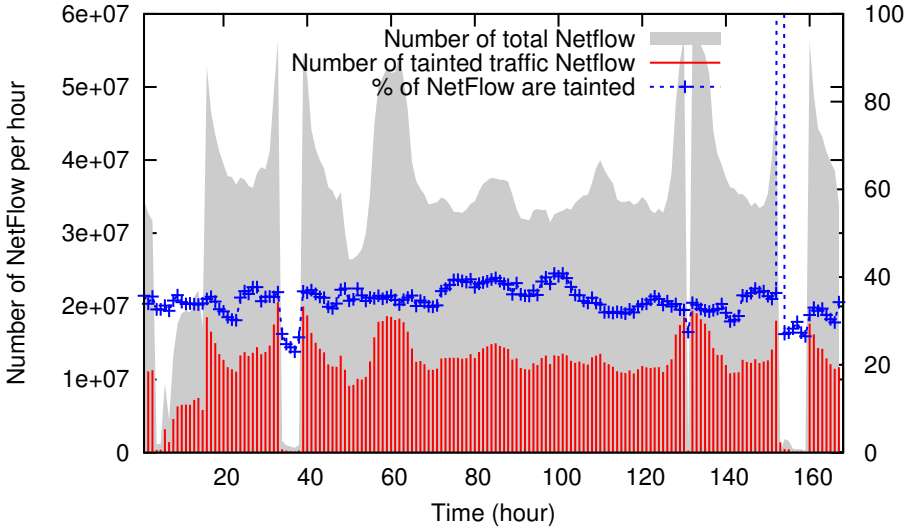

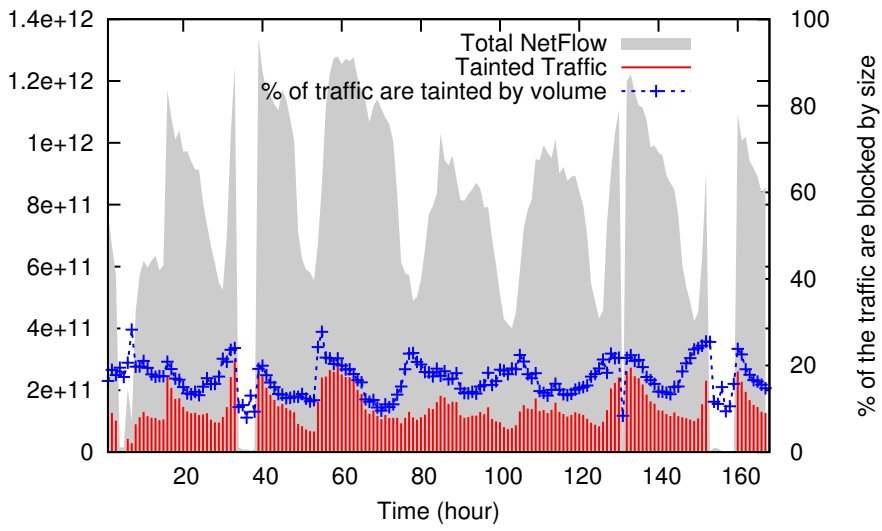

Tainted Traffic: The NetFlow who have a malicious source IP or malicious destination IP

- A surprisingly high proportion – 40% of flows (left) or 17% of traffic bytes (right), are tainted

Tainted Traffic by Blacklist

- *Q6: Whether a list, or a class of lists, have the greatest impact on our traffic?*

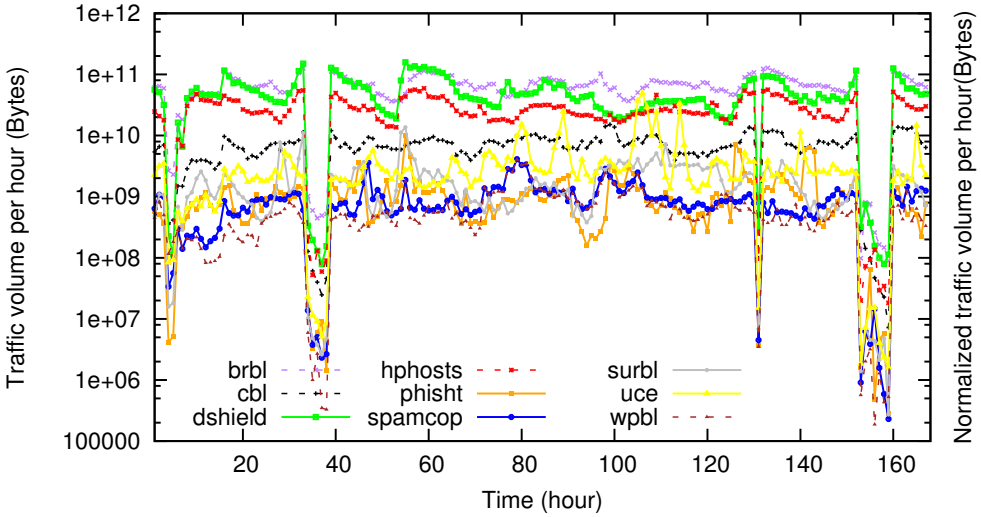

Total traffic bytes

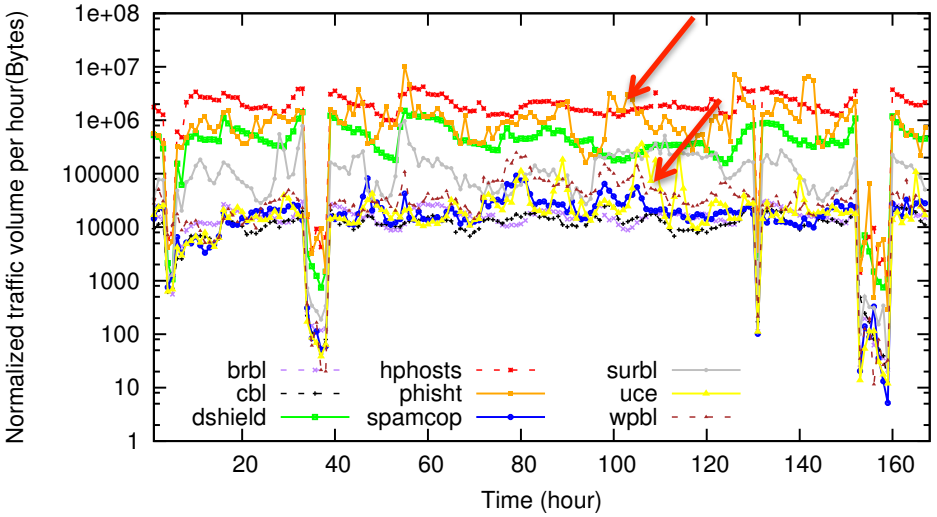

Normalized (traffic per IP)

Variance among the tainted traffic volumes, ranging from more than 10 GB per hour to tens of MB per hour

Each IP in Phishing/Malware and Active attack blacklists contributed two orders of magnitude higher tainted traffic than IPs in SPAM-related blacklist

Local v.s. Global

- *Q7: What fraction of global blacklists are touched by local traffic?*

	Spam					Phishing/Malware			Active
	BRBL	CBL	Spamcop	UCE	WPBL	hpHosts	Phisht	SURBL	Dshield
Touched entries	4,142,394	577,583	44,383	134,024	16,288	13,989	983	14,043	105,918
% of the list	2.8%	7.7%	29.3%	39.5%	51.2%	25.2%	24.4%	13.9%	22.1%

Blacklists entries touched by our network traffic

- Only a small fraction of malicious IP addresses were touched by a regional ISP's traffic
- Confirm the differences between local and global perspectives

Heavy Hitting IPs

- *Q8: Is there any IPs that are responsible for a disproportional large fraction of tainted traffic?*

Internal Heavy Hitters:

- Akamai
- Education Institutions
- Library
- Medical Centers

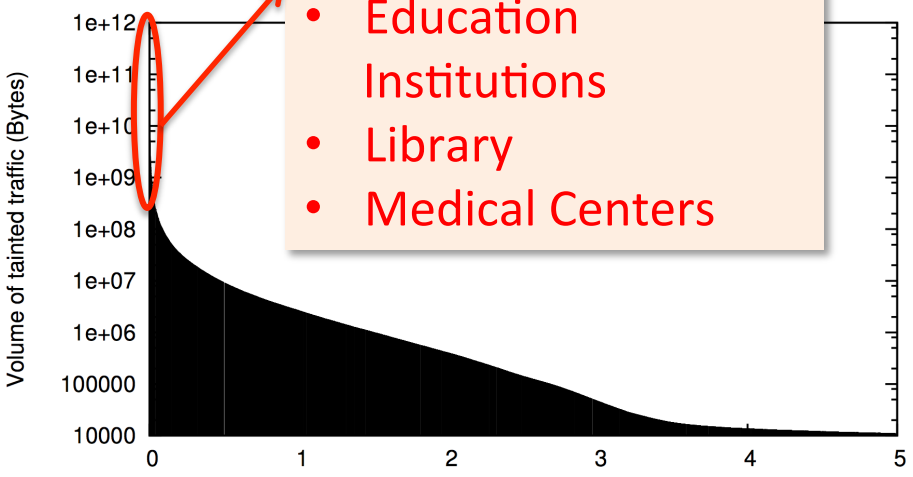

Internal IPs (%)

External Heavy Hitters:

- Amazon Web Services hosts
- Facebook CDN servers
- Pandora media servers
- EDGECAST IPs
- BOXNET servers

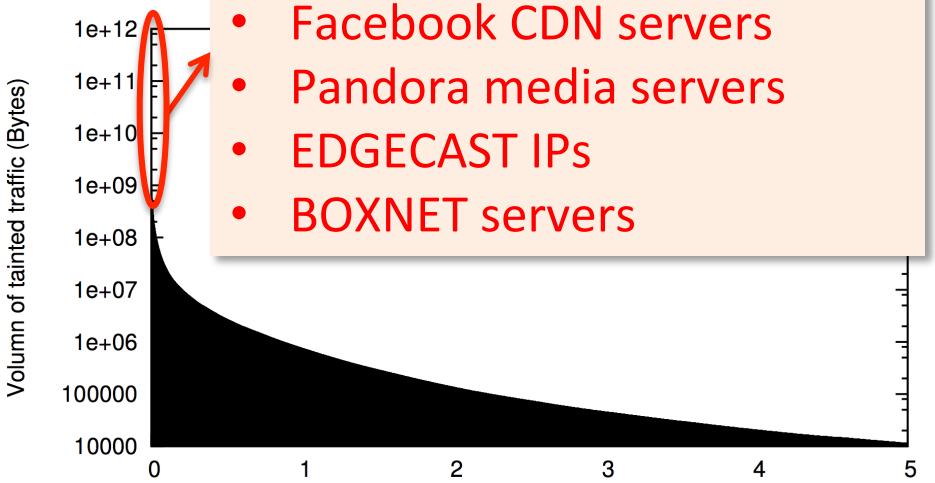

External IPs (%)

Tainted Traffic volume of top 5% of IPs

Top 50 IPs were responsible for ~40% of total tainted traffic

Heavy Hitters in the Blacklists

- *Q9: How are these heavy hitters distributed across blacklists?*

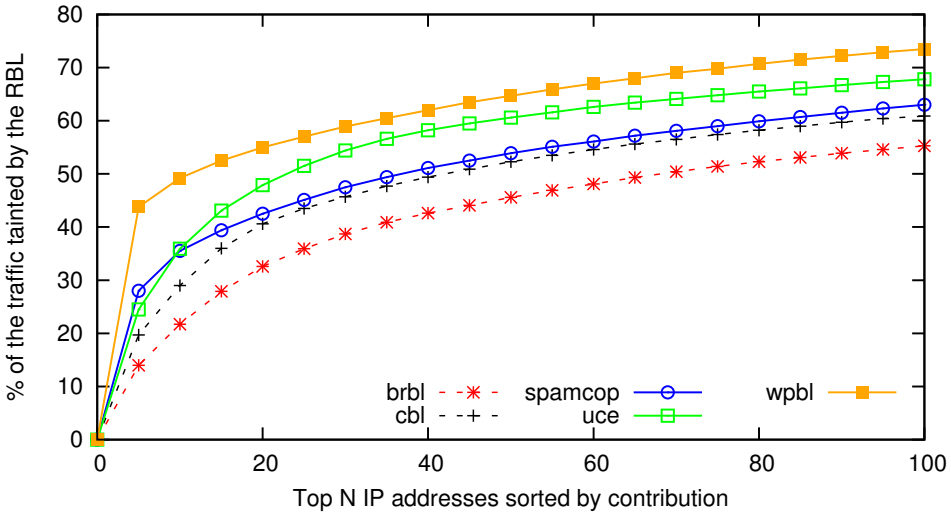

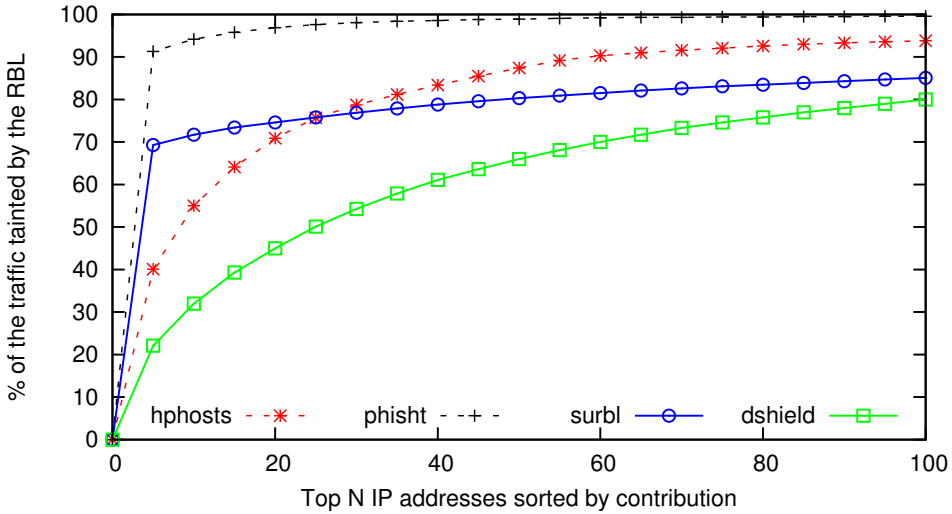

Cumulative contribution of the top N IPs per blacklists

- The top 50 IPs contributed more than half of the tainted traffic for each blacklists
- The contribution is even higher in Phishing/Malware lists (~80%)

Heavy Hitters in the Blacklists

- *Q9: How are these heavy hitters distributed across blacklists?*

	Spam					**Phishing/Malware**			**Active**
	BRBL	CBL	Spamcop	UCE	WPBL	hpHosts	Phisht	SURBL	Dshield
CDN	2	0	0	0	0	35	3	1	26
HOST	0	0	1	0	2	3	19	17	12
TOR	1	11	0	0	0	1	0	0	0
MAIL	0	0	0	3	5	0	1	0	1
VPN	3	0	0	1	0	0	0	0	0
Total	**10**	**13**	**1**	**4**	**7**	**39**	**23**	**18**	**39**

- 60 CDN servers and 51 hosting company IPs

Agenda

1. Data Collection
   - Reputation Blacklists
   - Network Traffic
2. Properties of Reputation Blacklists
   - Timing
   - Region
   - Overlap
3. Impact of Reputation
   - Tainted Network Traffic
   - Heavy Hitting IP Addresses
4. Conclusions & Discussions

Conclusions

- Characteristics of Reputation blacklists
  - While stable in size, the blacklisted IPs are highly dynamic, growing between 150% to 500% over a one week period
  - Classes of blacklists show significant internal entry overlap, but little similarity is seen between classes
  - Blacklists within the same classes share affinity for specific geographic distributions (e.g., RIPE and APNIC dominate SPAM; ARIN and RIPE dominate phishing and malware)

Conclusions

- Impact of Reputation
  - A surprisingly high proportion, up to 17%, of the collected network traffic is tainted by at least one of blacklists
  - Our network only saw traffic to a small portion, between 3% and 51%, of IP addresses within the blacklists
  - Heavy hitters account for a significant number of the tainted bytes to the network

Discussion

- False Positives
  - Some of the entries are likely false positives (e.g., Facebook CDNs)
  - Some of the entries are possibly decay entries (e.g. AWS hosts)
- Be more conservative
  - *Liberal approach:* tainted all the traffic with a union of the blacklists
    - 17% of total traffic bytes are tainted
  - Some blacklists are intended to taint one kind of application traffic
    - Reduce the taint traffic to 10.5% of total traffic bytes
  - Remove likely false positives
    - The volume of tainted traffic was reduced to 7.5% of total traffic

Thanks!