Hochschule für Angewandte
Wissenschaften Hamburg
Hamburg University of Applied Sciences

A Reproducibility Study of "IP Spoofing Detection in Inter-Domain Traffic"

Jasper Eumann

October 9, 2019

iNET RG, Hamburg University of Applied Sciences

Overview

IP Spoofing

Mitigation in General

Detection in Inter-Domain Traffic

Results

False Positive Indicators

Conclusion

IP Spoofing

IP spoofing

- IP spoofing injects packets that include a forged IP source address which is not its own
- Replys are directed to the address in the packet and not to the origin

Abuse potential

In combination with a distributed amplification, in which small requests trigger much larger replies, this leads to serious denial of service attacks in the current Internet [5, 10].

Amplification and reflection attack using a DNS server

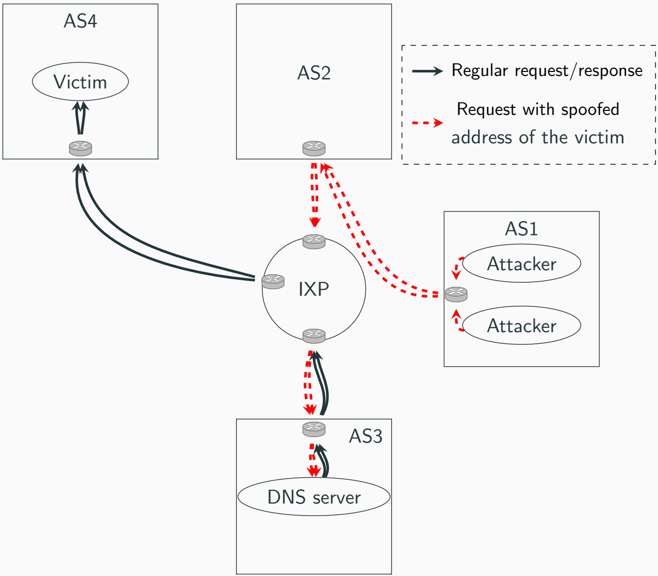

Mitigation in General

IP spoofing mitigation

- The most effective mitigation of reflection attacks is ingress filtering at the network of the attacker [3, 1]
- This solution is not sufficiently deployed [4]
- Can only be used in the area near the attacker

A border router blocks incoming traffic using ingress filtering

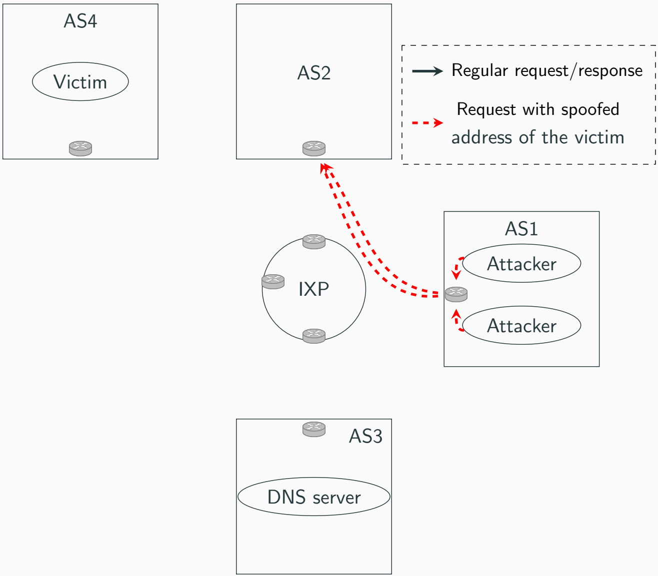

Detection in Inter-Domain Traffic

Spoofing detection in inter-domain traffic

- Packets passing through an IXP are forwarded by a peering AS
- Use expectation of "covered" prefixes to filter packets
- Complicated by transit providers

Customer cone

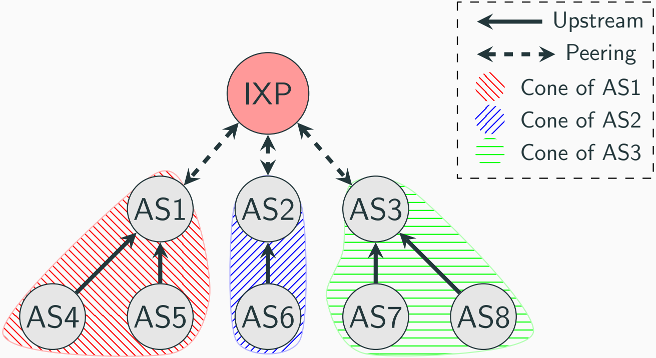

A customer cone includes all ASes that receive (indirect) upstream via the IXP member (AS1, AS2, AS3)

Amplification and reflection attack using a DNS server

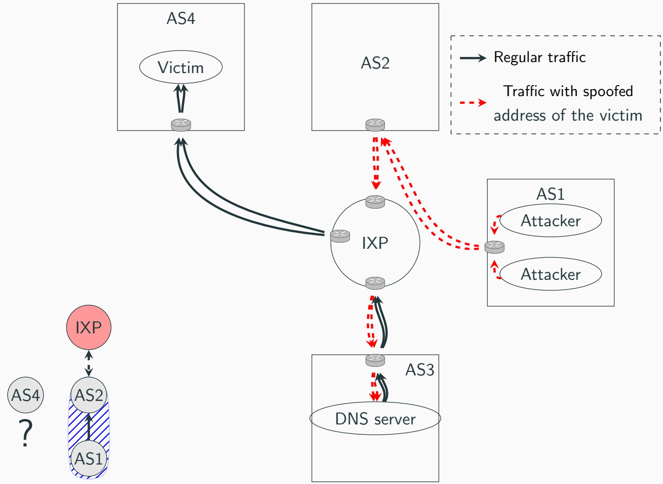

IMC'17 methodology

- *Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses* published at ACM IMC'17
  - passive detection of packets with spoofed IP address
  - minimize false positive inferences [6, § 1]
- Each packet that enters an IXP via an IXP member is checked via a customer cone that covers the prefix of the origin AS
- Paper presents three cone approaches

Customer cone approaches

1. **Naive Approach**: Uses public BGP information and considers a packet is valid if it originates from an AS that is part of an announced path for its source prefix

```
BGP4MP|1522454399|A|206.197.187.10|14061| 185.160.179.0/24 |
 14061 1299 12880 49148 |IGP|206.197.187.10|0|0||||
```

Customer cone approaches

1. **Naive Approach**: Uses public BGP information and considers that a packet is valid if it originates from an AS that is part of an announced path for its source prefix
2. **CAIDA Customer Cone**: Represents the business relationships rather than the topology. Build from AS relationships data provided by CAIDA [8]

Customer cone approaches

1. **Naive Approach**: Uses public BGP information and considers that a packet is valid if it originates from an AS that is part of an announced path for its source prefix
2. **CAIDA Customer Cone**: Represents the business relationships rather than the topology. Build from AS relationships data provided by CAIDA [8]
3. **Full Cone**: Built from public BGP announcements. This approach adds transitive relationships between peers. (Main method examined in the IMC'17 paper)

Manual intervention

- The authors of IMC'17 added "missing" links to the full cone by hand (based on whois information)
- In our opinion only a full scriptable method is usable in practice
- We show the properties of the cone approaches without manual intervention

.

Classification classes

The full pipeline sorts packets into four classes:

- **Bogon**: Address from a private network or other ineligible routable prefixes [9, 2, 11]
- **Unrouted**: Source is not included in any announcement
- **Invalid**: Packet with a spoofed source address
- **Regular**: Regular traffic without anomalies

Classification pipeline

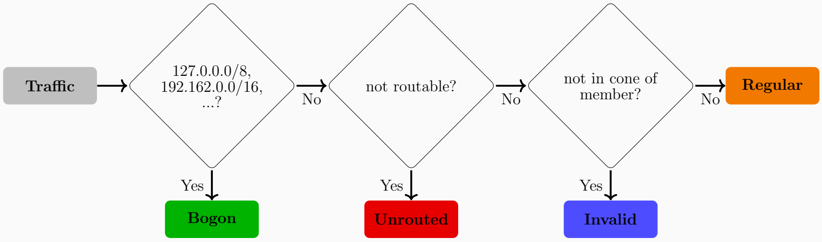

Reproduction procedure

1. Collect sampled flows data at an IXP
2. Apply scripts [7] kindly provided by the IMC'17 authors
   - We extended the implementation with missing functionality
3. Enhance cone construction with features for classifying payloads of spoofed traffic using libpcap[1]

[1] `https://www.tcpdump.org/`

Results

Comparison of classification results for invalid traffic

		IMC 2017		Reproduced Results	
		Bytes	Packets	Bytes	Packets
	Bogon	0.003%	0.02%	0.0009%	0.0022%
	Unrouted	0.004%	0.02%	0.00001%	0.0001%
Invalid	Naive	1.1%	1.29%	0.579%	1.537%
	CAIDA	0.19%	0.3%	0.955%	1.563%
	Full	0.0099%	0.03%	0.2%	0.488%

Time series of classified traffic distributions (Full)

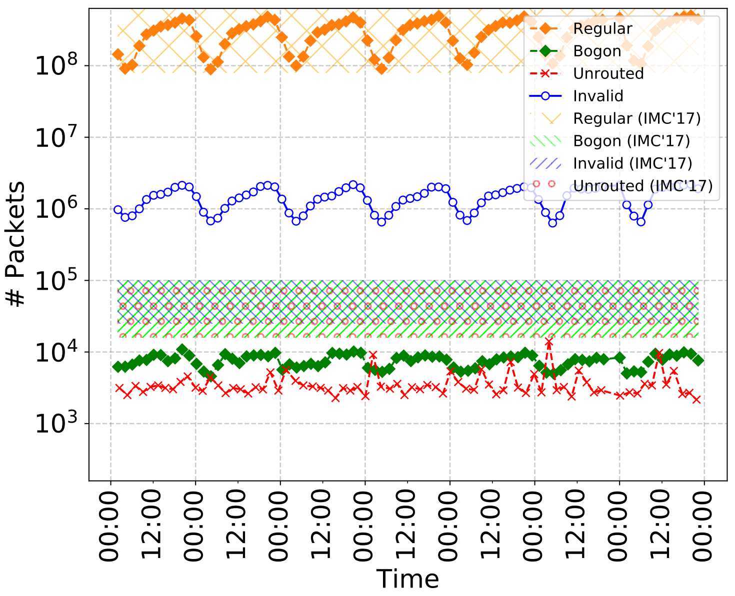

Time series of classified traffic distributions

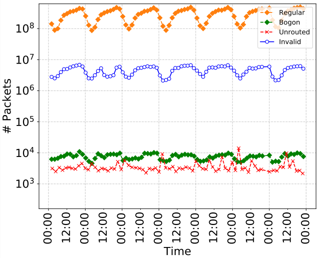

Naive

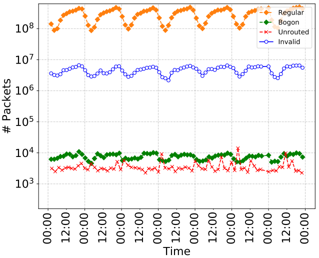

CAIDA

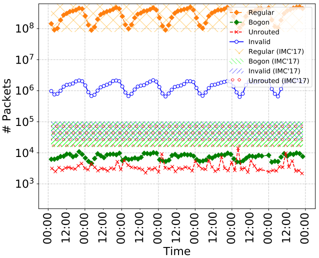

Full

CCDF: Fractions of invalid traffic per IXP member AS (Full)

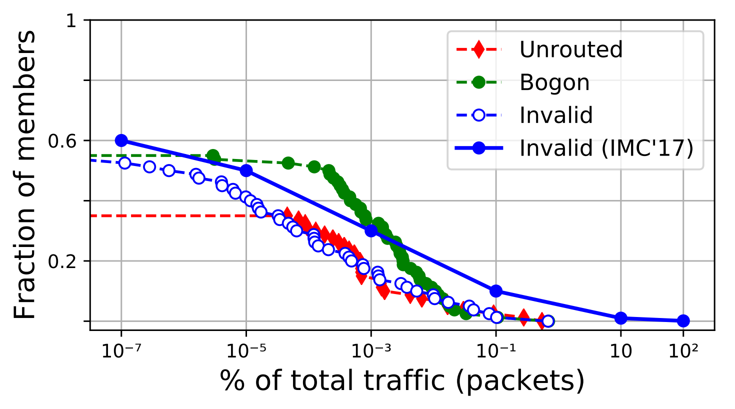

CCDF: Fractions of invalid traffic per IXP member AS

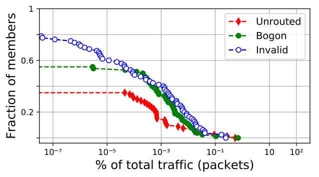

Naive

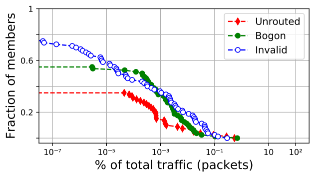

CAIDA

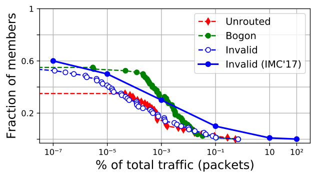

Full

CDF: Packets sizes by category (Full)

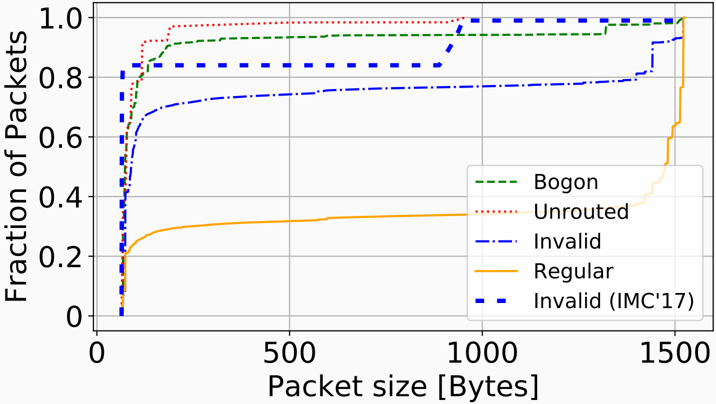

CDF: Packets sizes by category

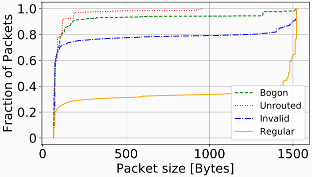

Naive

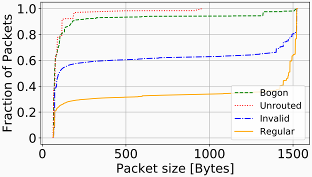

CAIDA

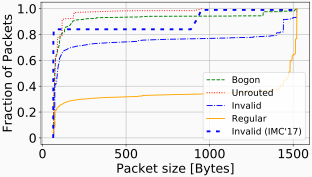

Full

Traffic mix per protocol and dst port of invalid packets (Full)

ICMP							total 0.37%
UDP	53 1.18%	123 $< 0.1\%$	161 0.35%	443 19.73%	ephe. 0.94%	other 0.81%	total 20.36%
TCP	80 3.50%	443 62.29%	27015 0.00%	10100 0.00%	ephe. 6.75%	other 13.67%	total 79.45%

False Positive Indicators

False positive indicators

Idea: Check if we actually identified invalid traffic

1. SSL over TCP
2. HTTP responses
3. ICMP echo replies
4. TCP packets carrying ACKs
5. Malformed packets (e.g., transport port 0)

False positive indicators by approach

	Naive	CAIDA	Full
SSL over TCP	3.985%	4.166%	6.395%
HTTP response	0.174%	0.134%	0.117%
ICMP echo reply	0.056%	0.070%	0.043%
TCP ACK	86.188%	69.197%	76.079%
malformed	0.000%	0.000%	0.001%

Conclusion

Conclusion

- The manual intervention has a significant effect on the results
- Without strong adjustments the methodology cannot be used in automatically fashion

Questions?

Thanks for your attention!

F. Baker and P. Savola.
Ingress Filtering for Multihomed Networks.
RFC 3704, IETF, March 2004.

M. Cotton and L. Vegoda.
Special Use IPv4 Addresses.
RFC 5735, IETF, January 2010.

P. Ferguson and D. Senie.
Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing.
RFC 2827, IETF, May 2000.

David Freedman, Brian Foust, Barry Greene, Ben Maddison, Andrei Robachevsky, Job Snijders, and Sander Steffann.
Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide.
RIPE Documents ripe-706, RIPE, June 2018.

Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti.
Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem.
In *Proc. of the 2017 Internet Measurement Conference*, IMC '17, pages 100–113, New York, NY, USA, 2017. ACM.

Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann.
Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses.
In *Proceedings of the 2017 Internet Measurement Conference*, IMC '17, pages 86–99, New York, NY, USA, 2017. ACM.

Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann.
transitive closure cone, 2018.
Accessed: 2019-08-28.

Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy.
AS Relationships, Customer Cones, and Validation.
In *Conference on Internet Measurement Conference*, IMC'13, pages 243–256, New York, NY, USA, 2013. ACM.

Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear.
Address Allocation for Private Internets.
RFC 1918, IETF, February 1996.

References v

Fabrice J. Ryba, Matthew Orlinski, Matthias Wählisch, Christian Rossow, and Thomas C. Schmidt.
Amplification and DRDoS Attack Defense – A Survey and New Perspectives.
Technical Report arXiv:1505.07892, Open Archive: arXiv.org, June 2015.

J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, and M. Azinger.
IANA-Reserved IPv4 Prefix for Shared Address Space.
RFC 6598, IETF, April 2012.

Top port UDP DST distribution of invalid packets

Naive	443 12.140%	53 4.040%	4500 1.800%	3074 1.218%	ephemeral 34.012%	other 44.664%
CAIDA	443 30.921%	53 3.637%	3074 1.296%	1193 0.951%	ephemeral 28.181%	other 33.507%
Full	443 77.174%	53 5.472%	16759 1.645%	161 1.406%	ephemeral 5.129%	other 8.157%