Tracking Down Sources of Spoofed IP Packets

Osvaldo Fonseca, Ítalo Cunha
Elverton Fazzion
Universidade Federal de Minas Gerais

Brivaldo Junior
Ronaldo A. Ferreira
Universidade Federal de Mato Grosso do Sul

Ethan Katz-Bassett
Columbia University

ABSTRACT

The lack of authentication in the Internet's data plane allows hosts to falsify (*spoof*) the source IP address in packets, which forms the basis for amplification denial-of-service (DoS) attacks. We propose techniques to identify networks that allow its hosts to send spoofed traffic. Our techniques systematically vary BGP announcements from multiple locations to induce changes to Internet routes and to the set of networks routed to each location. Preliminary evaluation in the real Internet indicates operators can correlate observations over multiple announcements to constrain the set of networks that may allow spoofed packets, possibly allowing targeted intervention.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network security**.

KEYWORDS

IP spoofing, traffic filtering, routing policies, topology discovery

ACM Reference Format:
Osvaldo Fonseca, Ítalo Cunha, Elverton Fazzion, Brivaldo Junior, Ronaldo A. Ferreira, and Ethan Katz-Bassett. 2019. Tracking Down Sources of Spoofed IP Packets. In *The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19 Companion), December 9–12, 2019, Orlando, FL, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3360468.3368175

1 INTRODUCTION

The lack of authentication in the Internet's data plane allows hosts to falsify the source IP address in packet headers, which forms the basis for amplification denial-of-service (DoS) attacks. The spoofed source addresses make the origins of such attacks seemingly untraceable, complicating mitigation efforts to squelch the attack or targeted efforts to convince networks to disallow spoofed traffic. Proposed techniques to identify routes taken by spoofed packets require changes to routers, cooperation of other networks, or wide deployment to provide accurate identification; none has been deployed and increased our ability to locate the origins of spoofed traffic. *We propose a novel approach to identifying networks that allow spoofed traffic that requires no changes to routers, no cooperation from other networks, and can be deployed by a single network.*

A network operator can estimate the volume of spoofed traffic received at each of its network's peering links [3] and the set of networks routed toward each peering link (a *catchment*). An operator can change the announcements for an IP prefix to induce changes to routes toward their prefixes and, more importantly, in the catchment of each peering link. The catchment changes, in turn, impact the volume of spoofed traffic observed at each peering link. Figure 1 provides intuition for how such measurements can be combined to identify networks that allow spoofed packets.

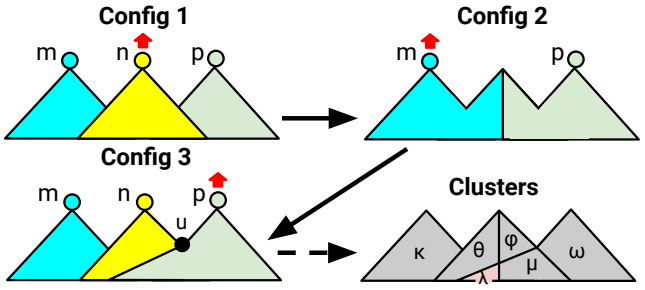

Figure 1: Example of catchments (colored triangles) for three announcement configurations performed by an origin network (not shown) peering with networks *m*, *n*, and *p*. The bottom right diagram shows the resulting clusters.

- In configuration 1 the operator announces a prefix through three peering links with networks *m*, *n*, and *p*; measures the catchment (colored polygons) and traffic arriving on each peering link; and identifies that the spoofed traffic is concentrated on the link with *n*, i.e., sent by networks in *n*'s catchment (red arrow).
- The operator later withdraws the announcement to *n* (configuration 2), remeasures catchments and traffic volumes, and identifies that the spoofed traffic is now concentrated on the peering link with *m*.
- Configuration 3 announces the prefix from *n* again, but poisoning AS *u* (which will force AS *u* to ignore the route from *n* and choose the route from *p* instead). The operator can measure catchments and traffic to identify that the spoofed traffic is concentrated on the peering link with *p*.

Finally, the operator can intersect the measured catchments to partition networks into *clusters* (bottom right), and correlate clusters with observed spoofed traffic (red arrows) to identify that the spoofed traffic is concentrated on networks comprising λ.

We present techniques to vary IP prefix announcement configurations that allow an operator to systematically induce changes to routes toward their prefixes and, more importantly, in the catchment of each of its network's peering links. Our techniques allow measurement of multiple mappings of *catchment* to *volume of spoofed traffic*. Operators can correlate mappings measured across multiple announcement configurations to identify networks that do not allow spoofed traffic or that may allow spoofed traffic.

We evaluate our techniques running experiments on the PEERING platform [4]. We deploy 705 different announcement configurations generated with our techniques from seven of PEERING's

CoNEXT '19 Companion, December 9–12,2019, Orlando, FL, USA

ACM ISBN 978-1-4503-7006-6/19/12.
https://doi.org/10.1145/3360468.3368175

peering links. Our experiments show our techniques effectively force route changes. We show that correlating information across multiple announcement configurations on PEERING allows us to partition the Internet into sets of three networks on average, indicating our techniques could constrain the set of networks that allow spoofed traffic to a small size.

We expect our techniques will ease identification of networks that do not employ BCP38 (ingress filtering) [2] and allow spoofed traffic, helping Internet bodies focus efforts and drive adoption.

2 INDUCING ROUTING CHANGES

To track down networks that allow spoofed traffic, an operator needs to generate a large number of mappings of catchment to volume of spoofed traffic. We present three techniques to systematically induce route changes.

Combinatorial prefix withdrawals. A network with multiple routes toward an IP prefix will pick the most preferred . Withdrawal of this most preferred route *will* result in the network using a different route (e.g., all networks in n's catchment in the change between configurations 1 and 2 in Fig. 1). Our first technique iteratively tries all possible combinations of announcements and withdrawals.

Iterative AS-path prepending. When choosing between multiple available routes, BGP first compares the policy-defined local preference (LocalPref). If multiple routes have the same LocalPref, BGP attempts to break ties by AS-path length. This implies that if a network has multiple available routes with the same LocalPref, the preferred route is the shortest. If this shortest route is made longer, the network *will* choose another (now shorter) route. Our second technique uses BGP AS-path prepending to artificially increase AS-path lengths, prepending from each peering link in turn.

Targeted AS-path poisoning. To prevent routing loops, BGP checks if the network's AS number is present in an announcement's AS-path before accepting the route. BGP poisoning adds a target network's AS number to an announcement's AS-path to trigger this behavior. Poisoning a network's most preferred route *may* trigger that network to choose a different (unpoisoned) route; when the target network changes routes, any downstream network originally routing through the target may also change routes. Our third technique uses AS-path poisoning to trigger path changes in target networks.

Combinatorial prefix withdrawals and AS-path prepending provide a systematic way to explore alternate routes and discover the most preferred routes from all networks in the Internet, while BGP poisoning allows route exploration at specific networks.

3 PRELIMINARY EVALUATION

We evaluate our techniques deploying 705 announcement configurations using 7 peering links of the PEERING platform. We consider all $\sum_{x=0}^{3} \binom{7}{7-x} = 64$ combinations withdrawing up to 3 of the 7 links. This guarantees we discover at least four distinct routes from each network toward different PEERING locations. For each of these 64 configurations, we generate additional configurations prepending from each location in turn (294 additional configurations). This guarantees we discover one alternate route with the same LocalPref as the preferred route (when it exists) for each withdrawal combination. Finally, we generate an additional 283 configurations poisoning all networks that are 2 AS-hops away from PEERING.

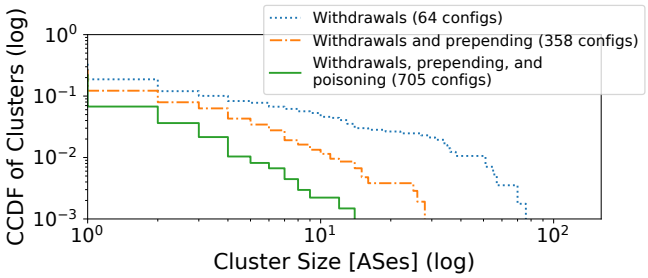

Figure 2: Distribution of cluster sizes after each phase.

We measure catchments at the granularity of ASes using BGP updates from RouteViews and RIPE RIS as well as traceroute measurements from 1600 RIPE Atlas probes (limited by our probing budget). We map traceroutes to ASes using Chen et al.'s heuristics [1]. Our catchment measurements cover a total of 1885 ASes.

Figure 2 shows the complementary distribution of cluster sizes, i.e., non-overlapping intersections across all catchment measurements (see bottom right part of Fig. 1). We show one distribution at the end of each phase, i.e., after withdrawals, after withdrawals and prependings, and after combining all techniques. We find all techniques are effective in inducing route changes, and that their combination significantly reduces cluster sizes. After deploying all 705 configurations, 99.9% of clusters have 12 or fewer ASes. Reducing cluster sizes at the tail is important, as it defines the worst-case accuracy of the location of sources of spoofed traffic.

4 DISCUSSION AND NEXT STEPS

Practical deployment. The number of possible configurations and the size of measured catchments depends on the number of links in an operator's network. Our techniques will perform better on large networks with rich connectivity; alternatively, multiple small network operators can join efforts and coordinate announcement configurations announcing a single prefix from all their networks concurrently. To avoid impacting production traffic due to the systematic route changes, we envision our techniques will be run using dedicated IP prefixes not carrying any production traffic. Dedicating multiple prefixes allows deployment of multiple announcement configurations in parallel and deployment of more configurations within the same time period.

Multiple attackers. Our exposition so far considers that a single or very few ASes allow spoofed traffic. In practice, we expect these networks will be spread across the Internet. An iterative approach is to constrain announcement propagation (e.g., by using poisoning or announcing to select peers) and apply our techniques to subsets of the Internet, then gradually increase announcement propagation as the number of networks allowing spoofed packets is reduced. More generally, however, we plan to evaluate an approach that solves a linear system of equations that equal the volume of spoofed traffic observed on a peering link to the volume of spoofed traffic sent by all networks in that link's catchment.

Localization speed. Our techniques generate an exponential number of configurations, which we currently choose based on our expectations of how many path changes they will induce. We plan to develop algorithms to intelligently choose which configurations to deploy, aiming at supporting active real-time mitigation of attacks.

REFERENCES

[1] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. 2009. Where the Sidewalk Ends: Extending the Internet AS Graph using Traceroutes from P2P Users. In *Proc. CoNEXT*.

[2] P. Ferguson and D. Senie. 2000. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (Best Current Practice). http://www.ietf.org/rfc/rfc2827.txt Updated by RFC 3704.

[3] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proc. Intl. Symp. on Research in Attacks, Intrusions and Defenses (RAID)*.

[4] B. Schlinker, K. Zarifis, Í. Cunha, N. Feamster, and E. Katz-Bassett. 2014. PEERING: An AS for Us. In *Proc. ACM HotNets*.