Corrigendum: Cloud Provider Connectivity in the Flat Internet

Todd Arnold[†] Jia He[†] Weifan Jiang[†] Matt Calder[‡†]
Italo Cunha[♮†] Vasileios Giotsas[°] Ethan Katz-Bassett[†]

[†]Columbia University [‡]Microsoft [♮]Universidade Federal de Minas Gerais [°]Lancaster University

ABSTRACT

This corrigendum corrects and extends our results on the benefit of *peer locking* in mitigating the propagation of route leaks on the Internet, originally published in [2]. The updated results show even higher benefits of *peer locking* than originally reported, and an extended analysis covering additional *peer locking* deployment scenarios shows partial deployments also yield significant reduction in propagation of leaked routes.

CCS CONCEPTS

• **Networks → Logical / virtual topologies**; *Public Internet*; *Network architectures*; **Topology analysis and generation**.

KEYWORDS

Internet topology, AS relationships, Routing, Traceroute, BGP

1 INTRODUCTION

Our original results correctly filtered leaked routes announced directly to ASes deploying peer locking. However, the original results incorrectly allowed routes leaked to ASes that do not deploy peer locking to later propagate through ASes that do deploy peer locking. This led to an underestimation of the benefits of peer locking. As our original results already showed that peer locking significantly reduces the propagation of leaked routes, the conclusions remain unchanged. This corrigendum presents the corrected results.

Figures 7, 8, and 9 replace the corresponding figures in the original paper and show the increased resilience from peer locking. We make corresponding changes to Section 8.2 to reflect the new results in Figures 7 and 8. Although Figure 9 has been updated, the discussion in Section 8.3 is unchanged. We include the updated Section 8.2 and unchanged Section 8.3 below for completeness.

8.2 Resilience vs Peering Footprint

We run simulations where each cloud provider's routes are leaked by a misconfigured Autonomous System (AS). We also consider the cloud provider under different announcement configurations. We run 5000 simulations per configuration, choosing the misconfigured AS at random. Figure 8 shows the cumulative distribution function for the fraction of detoured ASes (*i.e.,* those ASes that route to the misconfigured AS) across all simulations for Google in the 2020 topology.

The misconfigured AS always leaks routes to all its neighbors. The *announce to all* line shows results when Google announces its routes to all neighbors. For comparison, the *average resilience* line shows the average fraction of ASes detoured for a random (legitimate) origin AS and a random misconfigured AS. For each of 200 randomly chosen misconfigured ASes, we randomly choose 200 victim origin ASes and calculate their average resilience. Our results show Google's footprint provides significantly stronger resilience compared to a random origin AS.

Manually inspecting cases where the leaker attracts traffic from more than 20% of ASes found leakers with multiple well-connected providers (*e.g.,* Tier-1 and Tier-2 Internet Service Providers (ISPs)). Google peers with many networks, and these networks will prefer leaked route from customers over peer routes from Google. To verify this, we also show results for a scenario where Google announces to all its neighbors, and different subsets of Google's neighbors deploy filters such that they discard routes for Google's prefixes that they receive from any network other than Google (a.k.a. *peer locking* [4]), limiting the propagation of leaked routes. We consider three scenarios in terms of which neighbors deploy peer locking: Tier-1 neighbors, Tier-1 and Tier-2 neighbors, and all neighbors. Figure 8 indicates that peer locking Tier-1 and Tier-2 neighbors would limit even the worst leaks to 20% of the ASes in the Internet, and global peer locking would make Google virtually immune to route leaks.

Figure 8 also shows results simulating Google only announcing its prefixes to Tier-1 and Tier-2 ISPs (including its provider in the September 2020 dataset [3], Tata). This scenario, which ignores Google's rich peering with lower tier and edge ASes, shows significantly reduced resilience against route leaks. In fact, since Google peers with most Tier-1 and Tier-2 ISPs (instead of buying transit), Google's resilience in this configuration is worse than that of a random origin AS. While adding peers improves resilience against route leaks as it makes routes shorter, changing a relationship such that an AS receives a route from a peer rather than from a customer *decreases* resilience as it makes announcements less preferred.

Figures 7a to 7d are similar to Figure 8 and show the fraction of ASes detoured when Microsoft, Amazon, IBM, and Facebook announce their routes under different configurations. The *average resilience* line is the same in all graphs. The results show that all cloud providers are resilient to route leaks. Peer locking is slightly more effective for Google because it has more peers and fewer transit providers; conversely, we note other cloud providers would be more resilient to leaks than Google if they announced their routes only to Tier-1, Tier-2, and providers.

IMC '20, October 27–29, 2020, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-8138-3/20/10…$15.00
https://doi.org/10.1145/3419394.3423613

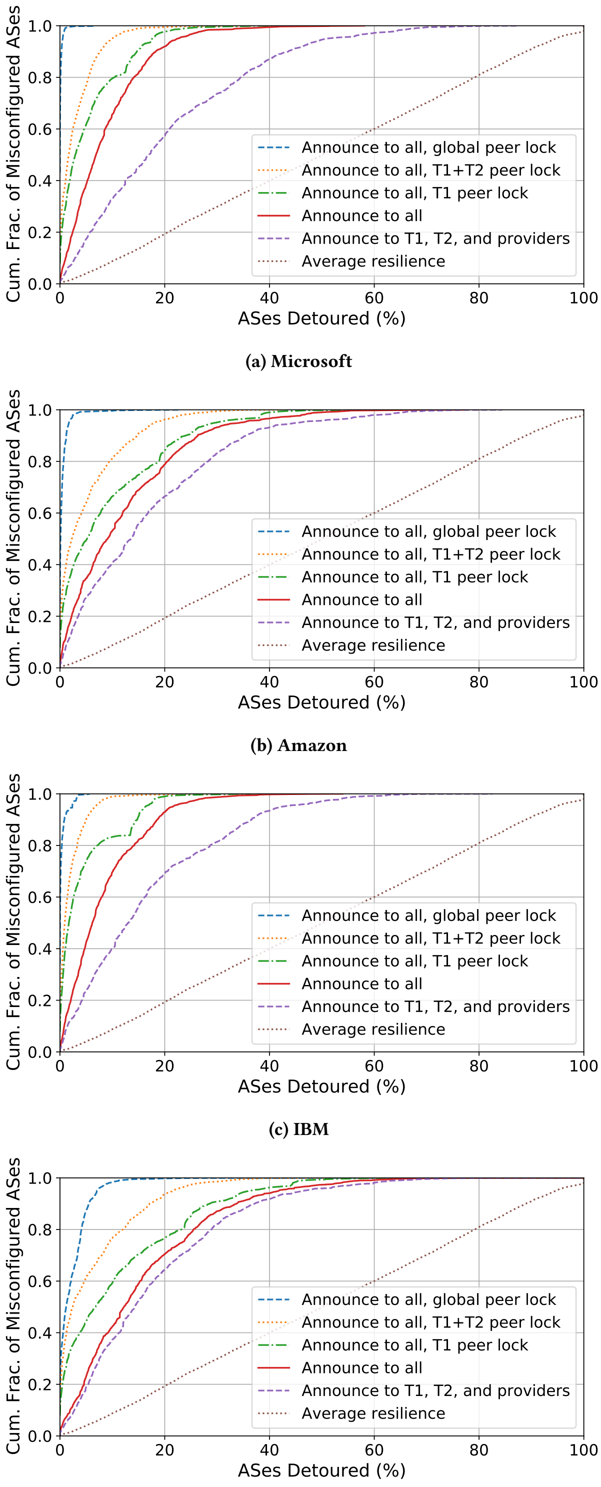

(a) Microsoft

(b) Amazon

(c) IBM

(d) Facebook

Figure 7: Percent of detoured ASes when cloud providers announce routes under different scenarios while a randomly selected misconfigured AS leaks the cloud provider's prefix.

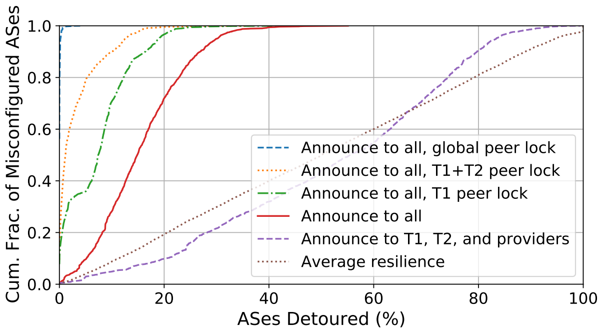

Figure 8: Percent of detoured ASes when Google announces routes under different scenarios while a randomly selected misconfigured AS leaks one of Google's prefixes. The results show that Google's peering footprint makes it resilient against route leaks.

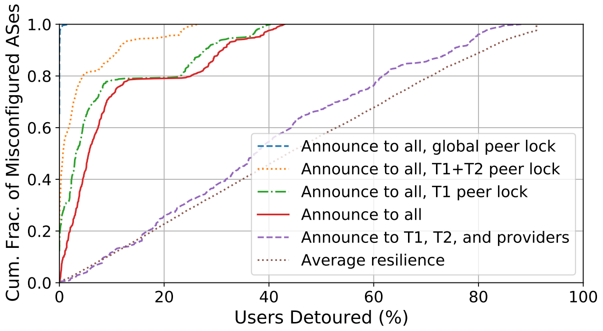

Figure 9: Percent of users in detoured ASes when Google announces routes under different scenarios. The results show that Google's peering footprint protects a large fraction of the user population from route leaks.

8.3 Fraction of Users Impacted

Figure 9 shows the fraction of users whose ASes have detoured routes for different route announcement configuration from Google. Figure 9 is similar to Fig. 8, but weights detoured ASes by their estimated population, as given by APNIC's population database [1]. Results are similar to the fraction of ASes detoured, with a slight skew to the left, indicating that some of the ASes that are detoured serve a relatively small fraction of users.

REFERENCES

[1] Asia-Pacific Network Information Centre (APNIC). [n.d.]. Visible ASNs: Customer Populations (Est.). https://stats.labs.apnic.net/aspop/.

[2] Todd Arnold, Jia He, Weifan Jiang, Matt Calder, Italo Cunha, Vasileios Giotsas, and Ethan Katz-Bassett. 2020. Cloud Provider Connectivity in the Flat Internet. In *Proc. ACM Internet Measurement Conference.*

[3] CAIDA. [n.d.]. CAIDA Serial-2 Dataset. http://data.caida.org/datasets/as-relationships/serial-2/.

[4] NTT. 2016. Deployment of NTT "Peer Locking" Route Leak Prevention Mechanism. http://instituut.net/~job/peerlock_manual.pdf.