

On the classification and false alarm of invalid prefixes in RPKI based BGP route origin validation

Wenjie Xu

Department of Electronic Engineering
Tsinghua University
Beijing, China
jackxu199603@gmail.com

Deliang Chang

Department of Electronic Engineering
Tsinghua University
Beijing, China
chdlgs@gmail.com

Xing Li

Department of Electronic Engineering
Tsinghua University
Beijing, China
xing@cernet.edu.cn

Abstract—BGP is the default inter-domain routing protocol in today's Internet, but has serious security vulnerabilities [1]. One of them is (sub)prefix hijacking. IETF standardizes RPKI to validate the AS origin but RPKI has a lot of problems [2] [3] [4] [5], among which is potential false alarm. Although some previous work [4] [2] points it out explicitly or implicitly, further measurement and analysis remain to be done. Our work measures and analyzes the invalid prefixes systematically. We first classify the invalid prefixes into six different types and then analyze their stability. We show that a large proportion of the invalid prefixes very likely result from traffic engineering, IP address transfer and failing to aggregate rather than real hijackings.

Index Terms—BGP, RPKI, ROV

I. INTRODUCTION

Internet is an inter-connected network without a center. It is made up of more than 50000 autonomous systems(AS for short). To transmit packets across the ASes, Border Gateway Protocol(BGP for short) is designed and implemented [6]. In BGP, every autonomous system will announce the prefixes owned by itself and propagate routing information it learned from its neighbors according to a policy. When propagating the prefix, the ASes will maintain a path to the origin of the prefix and can choose among different paths.

However, as a fundamental part of Internet infrastructure, BGP has serious security vulnerabilities [1]. One of them is BGP prefix hijacking. In prefix hijacking, an AS may illegitimately announce a prefix not owned by itself and then those ASes who accept the announcement will transmit the packets to a wrong destination. Prefix hijacking may result from misconfiguration and malicious attack. Using prefix hijacking, the attackers can block web service, steal secret information and do man in the middle attack, etc [1]. Actually, BGP prefix hijacking is frequently appearing in recent news [7] [8]. In one of the most famous prefix hijacking events, Pakistan Telecom blocked Youtube for more than 2 hours [9], causing inconvenience to the Youtube visitors all over the world. To tackle the problem of prefix hijacking, IETF(Internet Engineering Task Force) standardizes a framework called RPKI(Resource Public Key Infrastructure) to validate the origination AS of a prefix [10]. In RPKI, trust anchors(the five regional Internet registries) sign prefixes and allocate the

signed prefixes to NIRs, LIRs or ISPs. And then prefixes will be assigned hierarchically to customers. Prefix owners can sign an object called Route Origin Authorization to authorize an AS to announce a prefix. A ROA consists of a prefix, prefix length, maximum length, the AS authorized to announce the prefix [10] and the trust anchor. ROAs are stored in distributed repositories and can be fetched to validate the BGP items. We now illustrate how ROAs can be used to validate the BGP items. For simplicity, we define BGP item as a two-element tuple (prefix, AS path), where AS path means a sequence of ASes that the announcement of the prefix traversed. And the last AS in the AS path is the origination AS of the prefix. Given a set of ROAs, there are three possible validation results of a BGP item as shown below.

Unknown The prefix in BGP item is not covered(Prefix A covering B means B is not longer than A and the first length of prefix A bits of the two prefixes coincide.) by any prefix in ROA.

Valid There exists a ROA item such that the prefix in BGP item is covered by the prefix in ROA, the length of the prefix in BGP item is no longer than the maximum length and the origination AS of the BGP item is the same as the authorized AS in ROA.

Invalid The prefix in the BGP item is covered by one prefix in ROA, but is not valid.

With ROAs validating all the BGP items, the prefix hijacking problem seems to be solved perfectly. However, up to now, RPKI has not been fully deployed and partial deployment may result in unexpected trouble like false alarm. False alarm prefix can reduce the trustability of RPKI and even make those ASes that discard false invalid prefixes lose Internet connection to the discarded prefixes. So a systematic analysis and evaluation of current invalid BGP prefixes are in urgent needs and of great significance. Our work, to the best of our knowledge, systematically classifies and evaluates the invalid BGP prefixes for the first time. We find that most of the invalid prefixes result from traffic engineering purposes like multi-homing and load-balancing. We also find a large part of the invalid prefixes are very likely transfer prefixes. And finally we build a website to publish our analysis and classification result to help the network operators design better routing policy.

II. RELATED WORK

Currently there are mainly two lines of efforts to tackle the threat of prefix hijacking. The first is detection approach. For example, Zheng Zhang, et al., designed a system called iSPY to detect IP prefix hijacking on its own [11]. Xiaoliang Zhao, et al., analyzed BGP multiple origin conflicts and gave the potential reasons [12]. Xingang Shi, et al., utilized the correlation of control plane information and data plane information to detect the hijacking prefix [13]. However, detection approach may suffer from false alarm.

The second is validation approach. To tackle the threat of prefix hijacking thoroughly, IETF(Internet Engineering Task Force) standardizes a framework called RPKI [10]. ASes can use signed ROAs to validate the origin of prefixes. However, RPKI is not perfect and faces some unexpected problems in the deployment process. Danny Cooper, et al., flipped the threat model and analyzed the risk of misbehaving RPKI authorities [3]. Ethan Heilman, et al., design tools to detect potential harm to BGP prefix and propose some modification of RPKI to improve its transparency [2]. Yossi Gilad, et al., show that MaxLength can be harmful to RPKI [5]. Yossi Gilad, et al., point out that partial deployment of RPKI can result in false invalid BGP prefix [4]. However, to the best of our knowledge, there is no systematic measurement and analysis of invalid prefix. To fill the gap, we collect BGP prefixes for 3 months and analyze them by classifying them according to the AS path structure, aggregation structure and AS commercial relationship.

III. CLASSIFICATIONS OF INVALID PREFIXES

Now we describe our classifications of invalid prefixes based on AS-path structure, prefix aggregation structure and AS-relationship structure. The following six types of invalid prefixes are in fact false alarms, which mean invalid BGP items providing legitimate connections. Note that in all the invalid prefix illustration figures the AS following the prefix is the BGP origin AS and TAxX represents a trust anchor.

A. Invalid load-balancing prefix

In this scenario, an AS may first got a ROA with a relatively short maximum length. However, the AS may then announce to be the origin of a more specific prefix for load balancing reasons. For illustration, as shown in figure 1, AS1 may be assigned a prefix 123.121.0.0/23 and to secure the prefix, relevant authority may sign a ROA with maximum length 23 for the (AS1, 123.121.0.0/23) pair. However, after some time, AS1 may announce two more specific prefixes 123.121.0.0/24 and 123.121.1.0/24 through its two providers AS2 and AS3 respectively for load balancing reason. Then, AS4 using RPKI for BGP prefix validation will determine those two more specific prefixes are invalid.

B. Invalid failing to aggregate prefix

This type of invalid prefixes differ from invalid load-balancing prefixes in that invalid failing to aggregate prefixes are announced according to exactly the same export policy.




Fig. 1. An illustration of invalid load-balancing prefix

For illustration, as shown in figure 2, AS1 may be first assigned a prefix 123.121.0.0/23 and a corresponding ROA with maximum length 23. Then, for some reason such as ignorance or customer requirement, AS1 may then announce a more specific prefix 123.121.0.0/24. However, because 24 is larger than the maximum length 23 in ROA, 123.121.0.0/24 is considered invalid by AS4, which uses RPKI for prefix validation.




Fig. 2. An illustration of invalid failing to aggregate prefix

C. Invalid multihoming prefix

In this scenario, a customer doing multihoming may announce invalid prefix. For illustration, as shown in figure 3, the provider AS2 may assign a subprefix 123.11.0.0/24 of its own prefix 123.11.0.0/23 to AS1 and at the same time, AS2 got a ROA item: AS2, 123.11.0.0/23,24, TAxX. However, AS1 may have other provider AS3 besides AS2 and propagate prefix through AS3. Then AS4 will treat 123.11.0.0/24 as invalid prefix.




Fig. 3. An illustration of invalid multihoming prefix

D. Invalid singlehoming prefix

In this scenario, a customer doing singlehoming may also announce invalid prefix. For illustration, as shown in figure




Fig. 4. An illustration of invalid singlehoming prefix

4, AS2 has a prefix 123.11.0.0/23 and to protect its prefix from being hijacked, AS2 also gets a ROA item: AS2, 123.11.0.0/23, 24, TAx. However, AS2 further assigns a subprefix of 123.11.0.0/23 to its customer AS1. And AS1 announces the subprefix through its provider AS2. For some reasons like ignorance or to attract more traffic, AS2 does not aggregate the subprefix announced by the customer AS1. Then AS3 considers 123.11.0.0/24 as invalid.

E. Invalid provider prefix

A dual case of invalid singlehoming prefix is invalid provider prefix. Actually, some of the providers do not include its customer's AS number into the AS path when propagating the prefix announced by its customer. For illustration, as shown in 5, the provider AS2 assigns a prefix 123.11.0.0/24 to its customer AS1. To secure the prefix 123.11.0.0, a ROA (AS1, 123.11.0.0/24, 24, TAx) is signed and published in the ROA database. However, when propagating the routing information, AS2 announces 123.11.0.0/24 as the origin AS though it should have included AS1 as the origin AS. As a result, 123.11.0.0/24 will be considered as an invalid prefix.




Fig. 5. An illustration of invalid provider prefix

F. Invalid transfer prefix

Last but not least, due to active IP address transaction and the mobility of the organizations owning IP prefixes, IP address transfer is becoming more and more frequent and can result in invalid prefix. For illustration, as shown in figure 6, AS2 used to own the prefix 131.51.0.0/24 and the corresponding ROA. For some reason, the prefix transferred to AS5 but the ROA: (AS2, 131.51.0.0/23, 24, TAx) isn't revoked or modified. So when the routing information of 131.51.0.0/24 is propagated to AS3 from AS5, the prefix is considered invalid due to the obsolete ROA item.

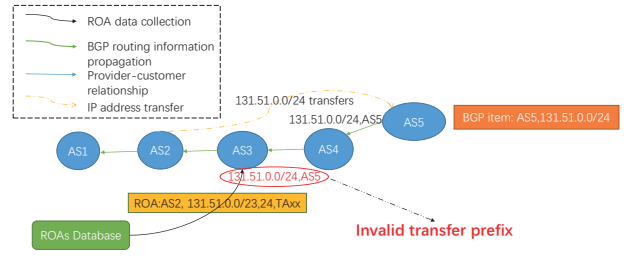


Fig. 6. An illustration of invalid transfer prefix

Validation Result	Number of Routing Items	Ratio
Unknown	635412	90.87%
Valid	58931	8.43%
Invalid	4949	0.71%

TABLE I
RPKI BASED ROUTE ORIGIN VALIDATION RESULT (DATA COLLECTED ON 16TH, MAY, 2018)

IV. DATASET AND ROUTE ORIGIN VALIDATION RESULT

The six types of invalid prefixes described in section III are all actually **legitimate**. That is to say, they are **false invalid** prefixes. To evaluate different types of invalid prefixes, we first collect BGP routing data for almost 3 months from February, 2018 and then do route origin validation. To collect BGP routing data, we set up a private AS. Then with our private AS, we do BGP peering with AS4538 (China Education and Research Network), which is adjacent to several core ASes and can collect most of the BGP data on today's Internet. Through AS4538, we collect BGP routing table, taking a snapshot of it every day, and BGP update data of the whole Internet. As for the ROA data, we use the software rpki-validator [14] provided by RIPE-NCC to collect and validate the ROA data. We also take everyday snapshot of the ROA data. After data collection, we build a prefix aggregation forest according to the aggregation relationship. We call the prefixes not covered by other prefixes maximal prefixes. Then with BGP routing table and ROA data, we do route origin validation. The validation result is shown in table I.

V. ARE INVALID PREFIXES REALLY INVALID?

Theoretical analysis in section III shows that there are six possible scenarios where traffic engineering, prefix deaggregation and address transfer can result in **false invalid** prefixes. And notably, there are 4949 invalid prefixes in the validation result. Although address hijacking is frequent these days, very unlikely we can detect thousands of hijackings at the same time. So the validation result is highly suspicious.

A. The classification result of real world BGP data

According to the classifications listed in section III, we design classification rules as shown in table II to classify the BGP prefixes.

We first sort the BGP prefixes, search the maximal prefix and build the prefix aggregation forest. For every node in the

Type of invalid prefix	Is the AS in ROA the same as BGP origin AS	Is the AS in ROA provider of BGP origin AS	Is BGP origin AS the provider of the AS in ROA	Multiple providers	Is there parent prefix or sibling prefix with different AS path	Is there parent prefix or sibling prefix with the same AS path
Invalid load-balancing prefix	Yes	No	No	—	Yes	—
Invalid failing to aggregate prefix	Yes	No	No	—	No	Yes
Invalid multihoming prefix	No	Yes	No	Yes	Yes	—
Invalid singlehoming prefix	No	Yes	No	No	Yes	—
Invalid provider prefix	No	No	Yes	—	—	—
Invalid transfer prefix	No	No	No	—	—	—

TABLE II
THE CLASSIFICATION RULES OF INVALID PREFIXES

prefix aggregation forest, we associate it with an AS path attribute for search in classification process. Then we apply the rules in table II to efficiently classify the prefixes. To find transfer prefix, we also use zmap [15] to scan the ip addresses under a prefix seeming to be transferred and we classify it as transfer prefix if we get active response. We show the classification result in table III. We observe that more than 60% of the invalid prefixes very likely result from traffic engineering, IP address transfer and failing to aggregate rather than real hijacking. And the rest of the invalid prefixes can be other types of false invalid prefixes or real hijacking.

B. The stability of invalid prefixes

We also monitor the invalid prefixes from 28th, February, 2018 to 16th, May, 2018 and we find that as shown in table III, most of the invalid (prefix, origin AS)s in different types are actually long-lived(meaning the (prefix, origin AS) pair keeps existing during our data collection period), implying they are more likely just false alarms since the real hijackings tend to be short-lived. On the one hand, the potential false alarms diminish the reliability of RPKI, thus slowing down the deployment of RPKI. On the other hand, false alarm may affect the false invalid prefixes' reachability for RPKI adopters. So we build a website to publish the possible false alarm prefixes and help network operators make better routing policy to avoid losing reachability due to false alarm. The website can be accessed through 202.38.101.13:5000.

VI. CONCLUSION AND FUTURE WORK

We first describe our classifications of invalid prefixes based on the AS path structure, prefix aggregation structure, AS commercial relationship. Then we collect real world BGP data for about 3 months and design rules to classify them. We find that more than 60% of the prefixes belong to the six types we describe in section III, implying that they very likely result from traffic engineering, IP address transfer and failing to aggregate rather than real hijackings. We also find

Type of Invalid prefix	Number	Percentage in invalid prefix	Number of long-lived (invalid prefix, origin AS) pairs	Percentage of prefixes with long-lived (prefix, origin AS) pair in this type
Invalid load-balancing prefix	923	18.7%	770	83.4%
Invalid failing to aggregate prefix	703	14.2%	684	97.3%
Invalid multihoming prefix	378	7.6%	355	93.9%
Invalid singlehoming prefix	204	4.1%	177	86.8%
Invalid provider prefix	186	3.8%	147	79.0%
Invalid transfer prefix	737	14.9%	658	89.3%
Other invalid prefix	1818	36.7%	1695	93.2%

TABLE III
THE CLASSIFICATION RESULT AND STABILITY OF INVALID PREFIX(DATA COLLECTED ON MAY, 16TH, 2018)

that most of the invalid prefixes are long-lived, which justifies the implication. One possible direction of future work is to do survey over practitioners to verify whether the invalid prefixes are false alarms and find even more types of false alarms.

REFERENCES

- [1] S. Murphy, "Bgp security vulnerabilities analysis," Tech. Rep., 2005.

- [2] E. Heilman, D. Cooper, L. Reyzin, and S. Goldberg, "From the consent of the routed: Improving the transparency of the rpki," in *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4. ACM, 2014, pp. 51–62.
- [3] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg, "On the risk of misbehaving rpki authorities," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*. ACM, 2013, p. 16.
- [4] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, "Are we there yet? on rpki's deployment and security." in *NDSS*, 2017.
- [5] Y. Gilad, O. Sagga, and S. Goldberg, "Maxlength considered harmful to the rpki," in *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*. ACM, 2017, pp. 101–107.
- [6] Y. Rekhter, T. Li, and S. Hares, "A border gateway protocol 4 (bgp-4)," Tech. Rep., 2005.
- [7] Z. Julian, "In the news: A bgp hijacking technical post-mortem," <https://www.bishopfox.com/blog/2017/01/in-the-news-a-bgp-hijacking-technical-post-mortem/>, accessed 18, May, 2018.
- [8] A. Siddiqui, "What happened? the amazon route 53 bgp hijack to take over ethereum cryptocurrency wallets," <https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/>, accessed 18, May, 2018.
- [9] Y. Hijacking, "A ripe ncc ris case study," 2008.
- [10] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "Bgp prefix origin validation," Tech. Rep., 2013.
- [11] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "Ispy: detecting ip prefix hijacking on my own," in *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4. ACM, 2008, pp. 327–338.
- [12] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "An analysis of bgp multiple origin as (moas) conflicts," in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*. ACM, 2001, pp. 31–35.
- [13] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting prefix hijackings in the internet with argus," in *Proceedings of the 2012 Internet Measurement Conference*. ACM, 2012, pp. 15–28.
- [14] RIPE-NCC, "Ripe ncc rpki validator 2.24," 2018. [Online]. Available: <https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources>
- [15] Z. Durumeric, E. Wustrow, and J. A. Halderman, "Zmap: Fast internet-wide scanning and its security applications." in *USENIX Security Symposium*, vol. 8, 2013, pp. 47–53.