Towards Detecting BGP Route Hijacking using the RPKI

Matthias Wählisch
Freie Universität Berlin
waehlisch@ieee.org

Olaf Maennel
Loughborough University
olaf@maennel.net

Thomas C. Schmidt
HAW Hamburg
t.schmidt@ieee.org

ABSTRACT

Prefix hijacking has always been a big concern in the Internet. Some events made it into the international world-news, but most of them remain unreported or even unnoticed. The scale of the problem can only be estimated.

The Resource Publication Infrastructure (RPKI) is an effort by the IETF to secure the inter-domain routing system. It includes a formally verifiable way of identifying who owns legitimately which portion of the IP address space. The RPKI has been standardized and prototype implementations are tested by Internet Service Providers (ISPs). Currently the system holds already about 2% of the Internet routing table.

Therefore, in theory, it should be easy to detect hijacking of prefixes within that address space. We take an early look at BGP update data and check those updates against the RPKI—in the same way a router would do, once the system goes operational. We find many interesting dynamics, not all can be easily explained as hijacking, but a significant number are likely operational testing or misconfigurations.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing Protocols*

Keywords

BGP, RPKI, secure inter-domain routing, deployment

1. INTRODUCTION

The Border Gateway Protocol (BGP) was designed with next to no security in mind and is therefore quite vulnerable to attacks such as prefix hijacks, protocol attacks, and accidental misconfiguration [1, 2].

This lack of security made it easy to announce address space that does belong to someone else—without the owners permission. "Evil guys" use this to inject spam, to launch attacks, or perform other illegal activities. Given the critical importance of the Internet, and the widespread use of this hijacking, there has been a lot of concern around that topic [3].

The IETF Secure Inter-Domain Routing (SIDR) working group has taken up the task to secure the BGP routing system. It builds upon ideas to cryptographically verify BGP update messages [2]. At first, this must include a mechanism of certifying who owns what address space. RPKI certificate hierarchy [4] forms a hierarchical relationship that follows the address space allocation, e.g., IANA gives address space to ARIN, ARIN to ISPs, ISPs to customers. In this way Route Origin Authorizations (ROAs) can be cryptographically signed and published in repositories. On the other end, relying party tools download that information, verify and upload it to the router. The router, now, receives BGP update messages and is able to check them against those validated ROAs. This is called *BGP Origin Validation*. If a ROA is found following a longest common prefix match, the prefix update can either be *valid* or *invalid*. It is up to the ISP to decide about the consequences, but, the intended meaning of *invalid* is that the legitimate owner did *not* agree to the announcement.

In this poster, we aim at first steps in understanding to what extent BGP hijacking actually occurs in the Internet. To achieve this we are downloading ROAs in the same way an ISP is downloading from the public repositories. We then look at BGP tables and updates from RIPE's RRC00 [5] and RouteViews [6] collector, and match these real-world BGP data against the validated ROA information. We expect that the *invalid* messages fall into two categories: (1) hijacks or similar illegitimate activities or (2) system testings or misconfiguration, including operators not understanding the RPKI system. To our surprise it is challenging to distinguish between those two categories.

The RPKI implements the concept of positive attestation. Consequently, all legitimate origin ASes of an RPKI-enabled address block must be included in signed ROAs. An operator may easily miss some more-specific prefixes from customers or additional (e.g., sibling) ASes. An invalid BGP update, thus, is by no means obviously a hijack and the properties of system testing or misconfigurations are hardly distinguishable from prefix hijacks. Currently, up to 20% of the verifiable routing table are *invalid*, which we will analyze in detail in the next section.

2. INVALID $\neq$ INVALID?

We look at one month, April 2012, of BGP update data and try to develop a methodology that estimates the occurence of real BGP hijacking. On April 1st, 2012, the RRC00 table dump contained 432,864 prefixes of which 6,843 (1.6%) prefixes were *valid*, but also 1,665 (0.4%) prefixes were *invalid*. We now try to understand better what causes invalid prefixes. Prefixes are *invalid*, if (i) the BGP origin AS does not comply with the ROA origin AS (20% of all *invalid* cases), (ii) the announced prefix-length is longer than the specified max-length in the ROA (70%), or (iii) both (10%). RouteViews exhibit similar results.

When looking at those prefixes, which turned invalid because the origin AS differed, we noticed that in over 40% of the cases we found the valid origin AS as the first upstream AS (second last

SIGCOMM'12, August 13–17, 2012, Helsinki, Finland.
ACM 978-1-4503-1419-0/12/08.

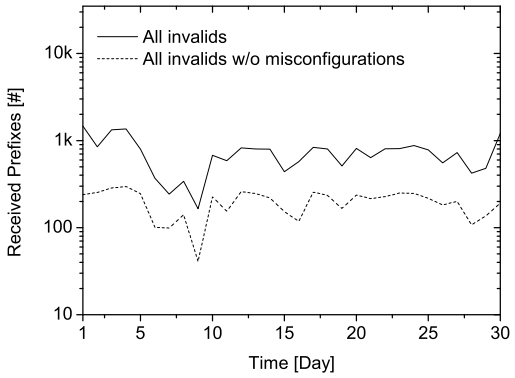

Figure 1: Number of invalid announcements for distinct prefixes per day

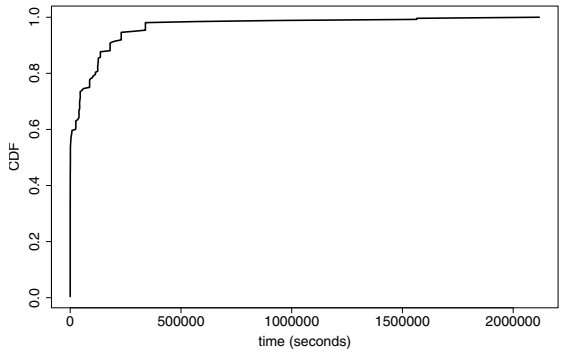

Figure 2: Duration of *invalid* in table.

AS on the path). This might lead to the speculation that operators of the upstream AS have "forgotten" to add at least some of their customers to the RPKI system. We also noticed this number has dropped significantly since January 2012, which might be due to increased training in how the RPKI works and how it secures the routing system.

Those considerations lead us to a set of assumptions, which we would like to understand better in order to distinguish misconfigurations from real illegitimate activity:

Misconfigurations are likely to cause (a) an *invalid* appearing on an AS path where the first upstream AS is one of the covering ROAs, (b) an *invalid* prefix that matches the ROA's origin AS, but is more specific compared to the maximum length defined in the ROA, or (c) an *invalid* announcement in which the ROA origin does not match the BGP origin AS, but both ASes are siblings.

Prefix Hijacking is indicated by (d) an *invalid* more-specific appearing temporarily and then disappearing again while a *valid* super-block is in the table, or (e) a *valid* prefix that is temporarily replaced by an *invalid* (at some observation points), while the corresponding ASes are not siblings.

Two points should be noted: (1) the proposed system for BGP origin validation is not intended to secure the routing system against attackers, who are faking their origin AS. For this reason alone, we do not consider BGP prefixes, where the origin AS is valid, but only the max-prefix length criteria is invalidating the announcement. We also believe that most of those cases are due to operators, who are currently in the process of adding their prefixes into the RPKI and have not done so for their whole network yet. (2) It should also be mentioned that hijacking often occurs in unallocated address space. However, unallocated address space is out of the scope of this work, as it is not secured by the RPKI.

We now look at BGP update dynamics for the whole month of April 2012. Figure 1 shows the number of invalid announcements for distinct prefixes per day, in which we exclude misconfigurations. Analyzing those in detail, only one *invalid* prefix appear as a more-specific of a *valid*, and 32 *invalid* prefixes occur, which temporarily replace an otherwise *valid* prefix (at some observation points). Figure 2 shows the duration of *invalids* in the table. While we would expect real hijacking to only last in the order of hours, we noticed that some "potential candidate prefixes" are still in the table for several days. We currently believe that many of those are still very likely operational tests and that our methodology needs further refinement. As a consequence this also implies, operators should be careful before turning on real filtering based on the RPKI.

The one prefix, which according to our methodology is most likely to be real hijcking, was in the table for 37.8 hours. While one prefix does does not sound too much, one should consider that the RPKI covers currently only 2% of the global routing table. If this would be a representative sample one might expect real hijacking in the order of 50 events per month.

3. CONCLUSION & FUTURE WORK

Securing the routing system is an important task for the Internet. Great progress has been made by the IETF SIDR working group, and prototype implementation of the system are being tested. In this poster we intended to look at BGP hijacking, but it turned out that we found many other interesting properties including how ISPs are starting to use the system. We argue that it is important to understand observable properties now, to support the efforts of the RPKI. For example, if we could detect common mistakes that operators are likely to make, this research could lead to improved training or to automated tools that watch-out for such misconfigurations.

In future work we intend to continue our investigation, and to improve our heuristics. We also intend to validate our findings by discussing with operators at forums such as NANOG and RIPE.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This work was partially financed by a gift from Cisco University Research Program Fund, grant 2011-89493(3696), and partially supported by the German BMBF within the project Peeroskop.

4. REFERENCES

[1] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, January 2010.

[2] G. Huston, M. Rossi, and G. Armitage, "Securing BGP - A Literature Survey," *Communications Surveys Tutorials, IEEE*, vol. 13, no. 2, pp. 199 –222, 2011.

[3] X. Hu and Z. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Security and Privacy, 2007. SP '07. IEEE Symposium on*, may 2007, pp. 3 –17.

[4] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," 2012, RFC 6480.

[5] "RIPE's Routing Information Service," http://data.ris.ripe.net/.

[6] "University of Oregon RouteViews project," http://www.routeviews.org/.