BGP runs on broken configurations
Why BGP is still insecure after 8 years of RPKI

by Niels Overkamp – s1783831

Introduction

If the internet is a network of networks, the Border Gateway Protocol (BGP) is the glue holding them together. The networks, called Autonomous Systems (AS), send BGP messages informing each other of routes to IP subnets. BGP itself is, however, devoid of security features. The protocol will not stop an AS from announcing that is the origin of an IP subnet it does not own. In 2011 a potential solution was deployed called the Resource Public Key Infrastructure (RPKI) which uses a PKI to specify which AS is allowed to announce which IPs. However in 2019, 8 years after the initial deployment, only about 10% of the BGP traffic is secured this way. This is further complicated by the fact that a large portion of the traffic that is secured, isn't actually secured in a valid way.

How did this system end up in this state? A paper by researchers from multiple different universities and big internet research institutes from 2019 [1] tries to give an insight on the history of RPKI.

RPKI

The goal of this protocol is to allow BGP routers to verify whether the *origin* of a route is allowed to announce a certain IP prefix. This is done using Route Origin Authorizations (ROAs), which are stored and made available by the Regional Internet Registries (RIRs). A ROA could look like this:

```
130.89.0.0/16; AS1133
```

We say that any BGP message announcing a route to an IP prefix in `130.89.0.0/16` is *covered* by this ROA. If we now get an announcement for a route to `130.89.0.0/16` we can validate whether the origin AS matches the AS in a ROA. [2] This groups all BGP announcements into three RPKI validity states:

- Valid: The IP prefix is covered by a ROA, and the AS and prefix match

[1] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. van Rijswijk-Deij, J. P Rula, N. Sullivan, "RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins", Proceedings of the Internet Measurement Conference, October 2019, pp. 406–419, `https://doi.org/10.1145/3355369.3355596`

[2] In the MaxLength intermezzo below we will discuss what happens if we get an announcement for a longer prefix, such as `130.89.192.0/18`

exactly;

- Invalid: The IP prefix is covered, but the AS or the prefix does not match;
- Unknown: The IP prefix is not covered by any ROA

The idea of RPKI is that routers can now use ROAs [3] to drop the invalid announcements and only letting legitimate or unknown routes into their routing tables. However, later on we will see why only a few operators are actually dropping invalid announcements.

A slow but steady adoption

Chung et al present a number of graphs showing the adoption of RPKI over time from different angles:

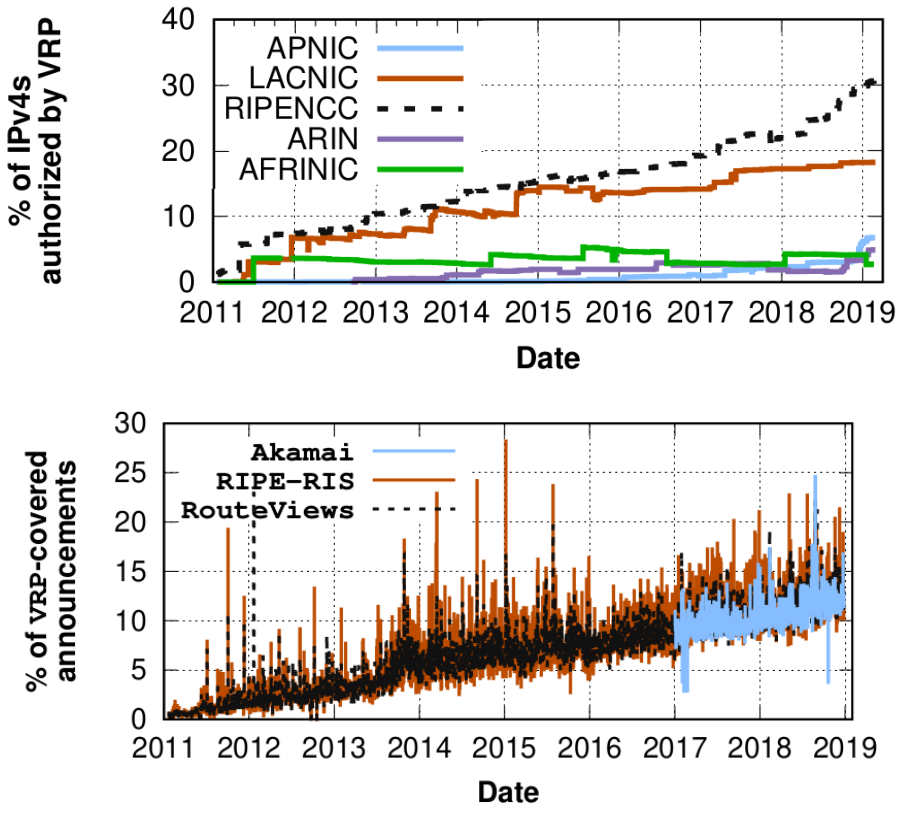

Adoption of RPKI in % of IPv4 space owned by RIRs covered by a ROA and in % of BGP announcements covered by a ROA

The first graph is constructed from the ROAs the RIRs published, and it shows

[3] Technically the routers would use a datastructure called Validated ROA Payloads (VRPs) derived from the ROAs, but for simplicity we will leave out this step in this article. The paper does make a distinction between the two.

a steady increase in IPv4s covered. At least, for IPs owned by RIPE NCC (Europe and Southwest Asia) and LACNIC (Latin America and some of the Caribbean). The remaining RIRs have not seen the same increase in adoption, with APNIC (Asia-Pacific), ARIN (North America and some of the Caribbean) and AFRINIC (Africa) all having less than 10% of IPv4 coverage.

The second graph shares the central observation: RPKI is adopted at a steady, but also slow pace. It is constructed from data sets containing BGP announcements from 3 different sources. After 8 years of RPKI deployment, it covers 10-15% of the BGP announcements. The steady increase is something we can credit in part to the nature of RPKI where it gives security benefits even if not many ASs are using it. Furthermore the authors note that the increasing trend in the graphs are encouring since earlier work showed that many network operators were not planning to deploy RPKI.

A lot of BGP traffic is (RPKI) incorrect

Earlier we stated that we should be able to use RPKI to drop invalid announcements, and with the increase in coverage this should make BGP more secure. Let's see how many advertisements that would mean are dropped.

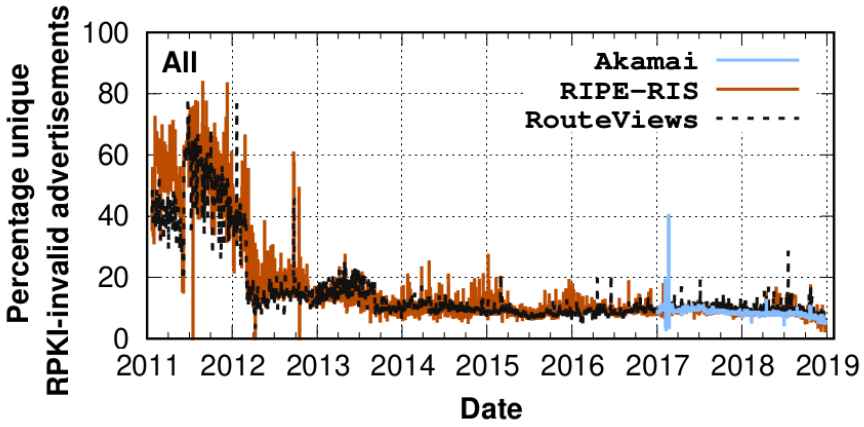

% of invalid announcements as seen from different datasets

The first thing you might notice is that a very large portion of the early BGP traffic was invalid. Chung et al. suspect that this was mostly misconfigurations on the network operators side, and that the sudden improvement was due to some RIRs improving their tooling and notifying operators of RPKI in-

valid announcements.

Then, you might notice that until late 2018 more than 5% of advertisements are still invalid with not much improvement, but that nearing the end of the period studied a slight improvement can be seen. The drop at the end can likely be attributed to efforts of Internet Exchange Points (IXPs) offering RPKI as a service. Furthermore, the practice of dropping invalid announcements is increasing. e.g. DE-CIX, one of the largest IXPs[4], started dropping invalid announcements in 2019.

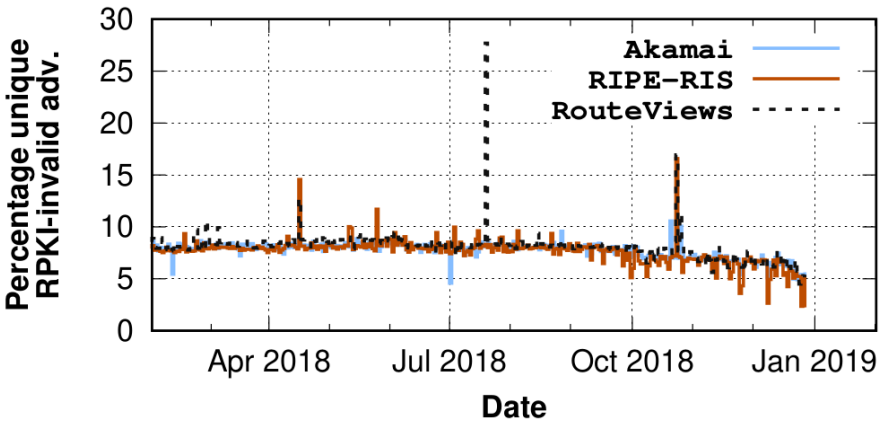

Figure 1: % of invalid announcements as seen from different datasets in 2018

Are these 2-5% of invalid announcements all malicious? If we have a closer look at the invalid BGP announcements we will see that many of them are actually likely to be legitimate advertisements.

Intermezzo: Prefix length

As stated before, if we have a ROA:

```
130.89.0.0/16; AS1133
```

any IP within this subnet is covered by this ROA. But we did not specify whether an announcement to a longer prefix within this subnet would be valid. Consider the announcement `130.89.0.0/18; AS1128, AS1133`. This announcement is covered by the ROA above, but even though the origin matches, it is not valid. RPKI enforces that the announced prefix needs to be exactly the same in the announcement and the ROA. This is to prevent forged-origin, sub-prefix hijacking.

[4] DE-CIX, Deutscher Commercial Internet Exchange, is a German IXP. It currently has an average throughput of about 7.3TB. `https://www.de-cix.net/en/locations/frankfurt/statistics`

If AS1128 wanted to intercept all traffic to `130.89.0.0/18`, they could announce the above BGP route. If AS1133 is only announcing `130.89.0.0/16` then all traffic to `130.89.0.0/18` would be routed via AS1128 due to longest prefix routing. Enforcing exact prefix matches prevents this from happening since no-one is authorized to announce a prefix longer than `/16`.

RPKI also allows operators to specify a range of prefix lengths that can be announced using the MaxLength prefix. AS1133 could for instance produce the following ROA:

```
130.89.0.0/16-20; AS1133
```

This would put them at risk of a sub-prefix hijack if they do not announce all prefixes themselves, but also give them more flexibility in what they can announce themselves.

Why are so many announcements wrong?

We know that a covered BGP announcement is valid if for one of the covering ROAs:

- the AS matches *and*
- the prefix length is no longer than MaxLength

This means that there are two reasons why an announcement could be *in*valid.

- The AS does not match
- The prefix is too specific (too long)

MaxLength is hard to understand

A survey from 2017 has shown that some network operators do not understand the MaxLength attribute or they might think that all announcements more specific than in the ROA would also be valid [5]. This is confirmed by the

[5] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman. Are We There Yet? On RPKI's Deployment and Security. NDSS, 2017

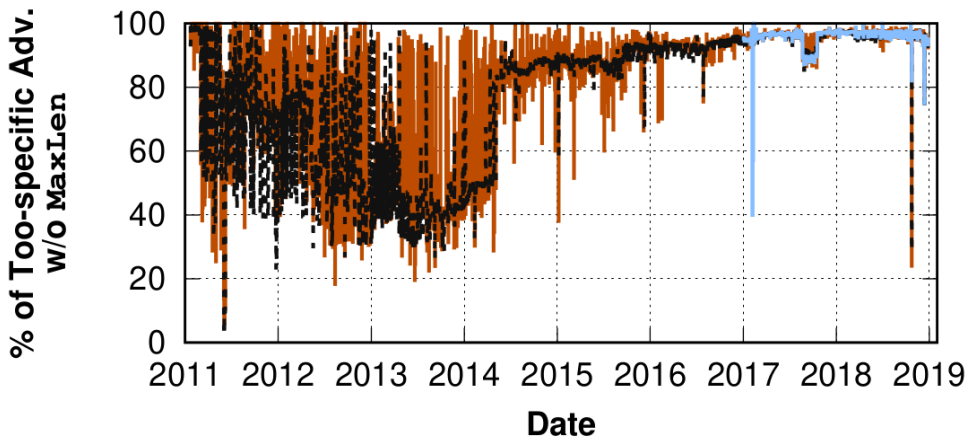

92% of too specific announcements do not have the MaxLength attribute set

graph from the paper by Chung et al. below.

It shows that almost all (92%) of too specific announcements do have left the MaxLength field empty, which means that it defaults to not allowing announcements of prefixes more specific than in the ROA. The steep increase in this percentage in 2014 was due to a single AS. They were making announcements that were more specific than the MaxLength field in their ROAs allowed. In 2014 they fixed this by increasing the MaxLength to 24.

This implies that the overwhelming majority of too specific announcements are not malicious but simply misconfigurations due to a misunderstanding of how RPKI and its MaxLength field works.

AS interrelations are complex

We now take a look at the announcements where the AS does not match any of the ROAs ASs. These could be caused by a malicious AS attempting to hijack traffic, but the graph below would suggest that this is mostly not the case.

The authors cross-referenced the BGP announcement data with a different datasets describing a few interrelations between ASs. It shows that about a third of 'Wrong AS' announcements are originated by the same ISP, and another third by a provider or customer of the one in the ROA. These are clearly legitimate traffic, but invalid due to a mistake in the ROAs.

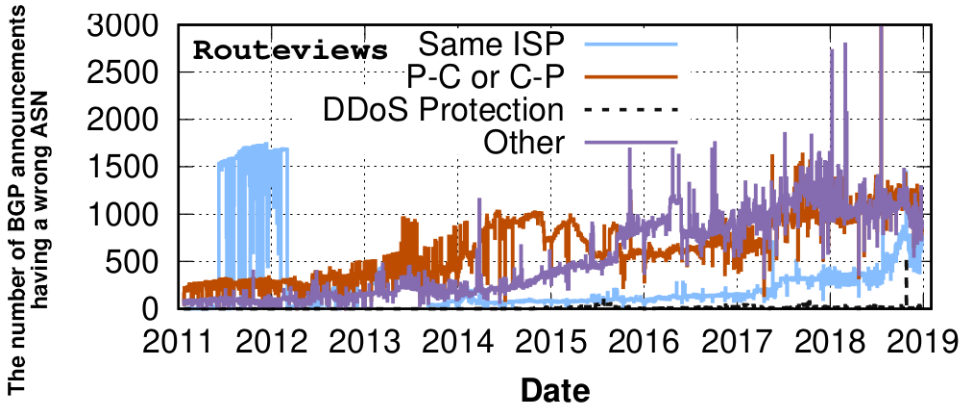

About a third of 'Wrong AS' announcements are originated by the same ISP, and another third by a provider or customer of the one in the ROA

Is it all misconfigurations?

It seems that a very large portion of the 2-5% invalid announcements we saw earlier in this blog are actually misconfigurations of the ROAs, and as such is hampering the growth of RPKI. If a large portion of actual legitimate advertisements are RPKI invalid, ASs will hesitate to drop the invalids (RPKI filtering). This will give less reason to ASs to fix their ROAs or to start using RPKI in the first place.

But not all advertisements could be linked to a misconfiguration reason. Below we show the plot of percentage of announcements falling in this category out of all advertisements. Between 1.07% and 1.39% are the values of the last datapoints. The authors state that these are likely to be hijack attacks. Further they present the finding that the percentage of HTTP/S traffic that is actually routed according to these invalid announcements turns out to be only 0.3%. While this is not an insignificant amount of traffic, it might be small enough for operators to decide to start dropping it.

Moving forward

To summarize, while RPKI could provide a significant increase in security of BGP, deployment is slow. A part of the cause is that a significant part of legit-

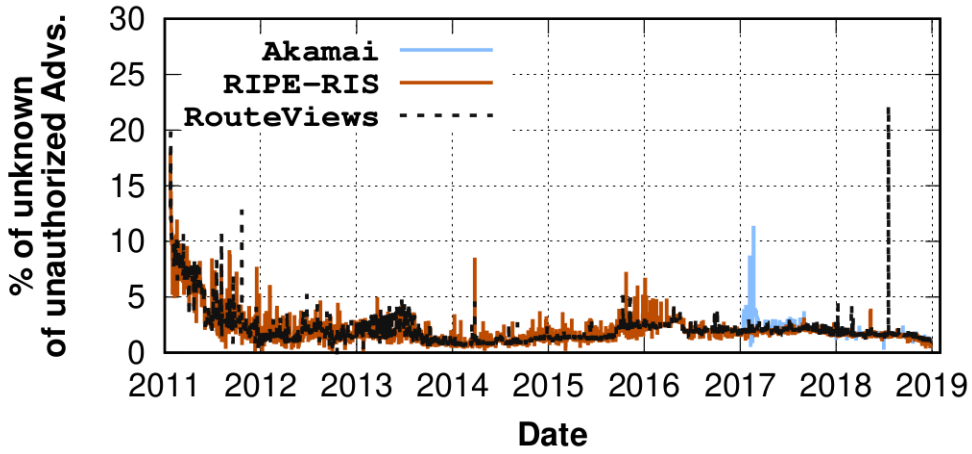

Percentage of announcements falling in the 'unknown' category out of all advertisements

imate BGP announcements are invalid. This is due to misconfigurations and misunderstanding how RPKI works among the network operators of the ASs.

A few solutions are possible. ASs could use manual rules, out of band information[6] or other workarounds to make exceptions for the legitimate invalid announcements. Alternatively ASs could decide to start dropping all invalids, as some of them have started doing. This will give operators a big incentive to improve their ROAs as they will notice things stop working. Furthermore with RPKI filtering being deployed more widely, this will give incentive to others to start using RPKI.

To readers interested in learning more about this subject, I recommend to read the full paper. A lot of the analysis of the data was left out here for brevity, but is explained well in the paper.

[6] e.g. using the datasets similar to the ones used by Chung et al. to determine AS relations