University of Amsterdam

System & Network Engineering

BGP Origin Validation (RPKI)

July 5, 2013

Authors:

Remy de Boer
<Remy.deBoer@os3.nl>

Javy de Koning
<Javy.deKoning@os3.nl>

Supervisors:

Jac Kloots
<Jac.Kloots@SURFnet.nl>

Marijke Kaat
<Marijke.Kaat@SURFnet.nl>

Abstract

Border Gateway Protocol (BGP) is the routing protocol that is used in the core of the Internet. In the past years it has been shown over and over again that configuration errors made in BGP configurations can lead to large portions of the Internet being unavailable. The most well known example being the outage of Youtube due to a Pakistan Telecom configuration in 2008. BGP origin validation (Resource Public Key Infrastructure, RPKI) is a potential solution to this issue. However, in its current form RPKI cannot be used as a source for making routing decisions due to the large amount of configuration errors and a low adoption rate. In this research we will give a better insight into the current state of RPKI adoption. The goal of the research described in this paper is to find the origin of invalid routes, and analyzing common mistakes that lead to invalid routes. The decision was made to build a dashboard for BGP operators and Regional Internet Registries (RIRs) that shows daily statistics about the cause of invalids, origin of invalids and changes over time on a per Autonomous System (AS), per RIR and world wide basis. When analyzing the results from July 1st, 2013, it becomes clear that approximately 40% of the invalid routes originate from only three ASes. The LACNIC region currently has the highest adoption rate but also has the highest percentage of invalid routes. By gathering data on a day to day basis and visualizing this data using Google Charts, we provide insight into the origin of invalid prefixes and the mistakes that are being made while configuring RPKI. RPKI live data is presented on our dashboard located at http://rpki.surfnet.nl/.

Acknowledgements

We would like to thank our supervisors Jac Kloots and Marijke Kaat for their valuable input during the conduction of this project and the subsequent period of writing this report. We would also like to show our gratitude to the SURFnet team. They've contributed to a very pleasant working environment and shared good knowledge which has contributed to this report.

Contents

1. Introduction

The Internet is of great importance to our society and economy. We use it for knowledge gathering, e-mailing, chatting, social media and shopping. An Internet outage can have huge consequences because of the integration of the medium in our daily lives. The Internet is divided into logical organizational structures we call Autonomous Systems (AS). At the core of the Internet network reachability information is exchanged among ASes using the Border Gateway Protocol (BGP). To prevent the spread of false routing information, like prefix hijacking attempts, an ISP will typically perform filtering on routes learned from a peer. Because the BGP routing table currently holds approximately 490.000 routes[1] it is almost impossible to manually filter routes learned from other ISP's.

Not validating or filtering routes between ISP's can lead to large portions of the Internet to be unavailable when mistakes are made in the BGP configurations. On Sunday, 24 February 2008, Pakistan Telecom (AS17557) started an unauthorized announcement of the prefix 208.65.153.0/24. One of Pakistan Telecoms upstream providers, PCCW Global (AS3491) did not filter incoming routes and forwarded this announcement to the rest of the Internet, which resulted in the hijacking of YouTube traffic on a global scale [2].

A potential solution could be the use of route origin validation based upon the Resource Public Key Infrastructure (RPKI). RPKI was designed by the Secure Inter-Domain Routing (SIDR) working group of the Internet Engineering Task Force (IETF) in 2007 to secure the Internet routing infrastructure. Owners of resources (IP prefixes and AS numbers) can create Route Origin Authorizations (ROAs) to prove that they are the legitimate owners of these resources. These ROAs are stored in a central repository. From this repository, usually called a trust anchor, other ASes download the ROAs and store them in their local cache which cryptographically validates the ROAs. Finally the router downloads the cryptographically validated ROA data from the local cache and compares this to the BGP announcements for validation. Thereby allowing an operator to change the routing policy depending on the validity of the route. Depending on the validity, the operator could decide to assign a higher/lower priority to the route or ignore it all together.

Widespread RPKI adoption could simplify routing decision-making. However, measurements performed by SURFnet show that there are a high number of invalid routes. These invalid routes are likely to be caused by misconfigurations, thereby preventing ISP's from using RPKI to adjust its routing policy for those routes. If we could find the most common mistakes made by BGP operators, this research could lead to improved instructions that prevent these mistakes in the future. To get to the heart of the problem, we must find out where the invalid routes are originating from, so that the operators of the concerning ASes can be informed.

1.1. Research Question

This research will focus on finding the origin of invalid routes, and analyzing common mistakes that lead to invalid routes. This results in the following research question:

"How can we reliably determine which ASes are advertising invalid routes due to misconfigurations and how can we monitor this over the course of time?"

To improve the adoption rate of RPKI we would also like to answer the following sub-question:

"What are common causes that lead to misconfigured Route Origin Authorizations (ROAs)"

1.2. Related Work

Previous research has been done by Matthias Wahlisch, Olaf Maennel and Thomas C. Schmidt in 2012 [3]. This research focused on distinguishing between misconfigurations and route hijacking attempts. Matthias Wahlish has also published an article titled "One day in the life of RPKI" [4], this article discusses some RPKI statistics dating from 2011.

1.3. Scope

The scope of this research is aimed at discovering the origin of invalid routes over time, by analyzing the validity states of current prefixes that are covered by ROAs. Common causes that lead to misconfigurations will be analyzed and operator instructions or tools will be improved, should this be necessary.

2. Introduction to RPKI

In order to understand the outcome of this research it is important to understand the RPKI infrastructure and the validation process. The infrastructure is discussed in detail in RFC6481[5] and the validation process is presented in RFC6811[6]. This section will cover the basics that are essential to this research.

2.1. PKI Infrastructure

Since IP addresses are handed out in a hierarchical structure it seems fair that certificates for these IP blocks are distributed in the same manner. The IANA is the root Certificate Authority (CA) and the RIR's are the authorized subordinate CA's as shown by Figure 2-1.

The root (IANA) has a self-issued certificate. The RIRs can assign AS numbers and IP prefixes to child sites or Local Internet Registries (LIRs). The LIRs can then choose to either use a hosted system maintained by their RIR or host their own CA for publication of their ROAs (authorization to announce a prefix). In either case the LIR can create ROAs for resources assigned to it. As such, each LIR can contact its RIR and request a resource certificate and use that to create ROAs. These ROAs and the LIRs resource certificate can then be stored at the LIR CA or at a hosted CA operated by the RIR.

2.2. Deployment Infrastructure

The certificate deployment structure follows a 3-tiered approach. It is presented as follows from top to bottom:

1. **Global RPKI:** The RPKI authoritative data (certificate objects) is published in public repositories. All Regional Internet Registries (RIR) host their own publication point. These public repositories are a trust anchor for local caches. The repositories hosted by the RIRs do not necessarily need to be used, it is also possible to add other trust anchors. Global RPKI refers to the infrastructure that was explained in section 2.1.

2. **Local cache:** Each autonomous system that participates in the global routing process should have a local set of one or more caches with verified ROAs. ROAs are downloaded from trust anchors and are cryptographically validated locally by the cache server. This prevents routers from doing the more CPU intensive cryptographic work.

3. **Routers:** For communication between the local cache and the routers the RPKI to Router Protocol[7] (rtr) is used. A router fetches data from a local cache using the rtr-protocol. The router functions as an rtr-client of the cache or rtr-server.

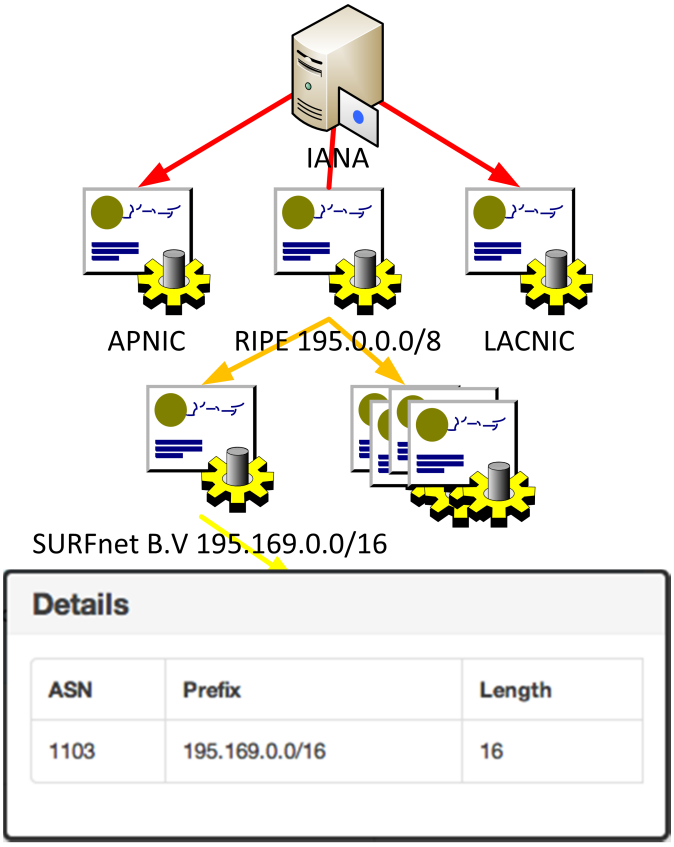

Figure 2-1: PKI infrastructure

2.3. Validation process

To check if an AS is advertising prefixes that are authorized by the legitimate prefix holder some validation and comparison steps need to take place. A logical presentation of this process is shown in Figure 2-2.

1. **Route Origin Authorization (ROA):** The digitally signed objects that associate exactly one AS with one or more blocks of IP addresses. These ROAs contain:
   - IP Address
   - Prefix length
   - Maximum length
   - AS number

   ROAs are stored in the central repositories operated by RIRs.

2. **Local cache:** Individual organizations and ASes can download their own local copy of the signed objects from the RPKI system using rsync. The local rtr-server

should then verify the signatures and process them. Validated ROAs are called "Validated ROA Payload" or "VRP".

3. **Routers:** Finally a router can import the VRPs from the rtr-server and use this information to determine if an incoming BGP announcement is valid. To check the validity the following rules apply[1]:
   - **Valid:** At least one VRP matches the prefix (Prefix length $\leq$ VRP max length, and AS number matches).
   - **Invalid:** At least one VRP covers the prefix, but no VRP matches it. (Prefix identical or more specific than VRP prefix)
   - **Unknown:** No VRP found that covers the prefix.

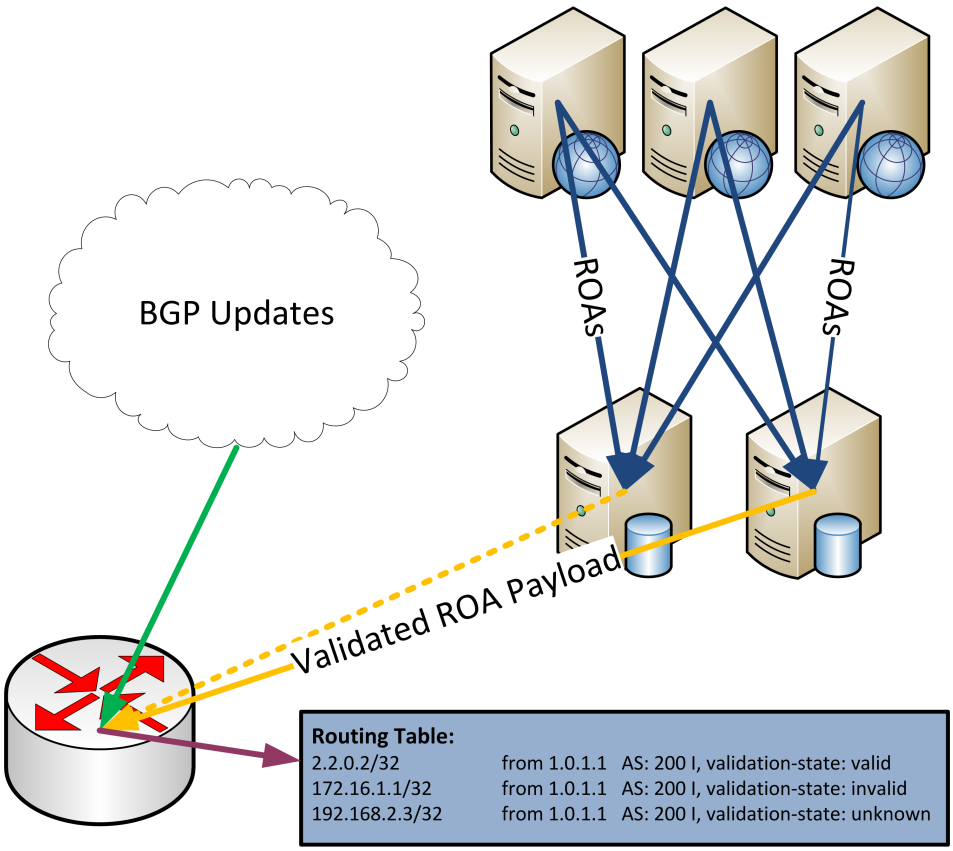

Figure 2-2: RPKI infrastructure and validation process.

[1]RFC6811 "BGP Prefix Origin Validation"

3. Research method

In order to completely answer the research question, it is relevant to see which routes are in the routing table and how incoming BGP routes are validated by a router. Therefore we will simulate a real-world BGP router. We will need to store information from the routing table in a database so we can generate statistics from the data at a later stage.

3.1. Information gathering

Before we can start generating detailed RPKI statistics, we require data from several sources. We need at least:

- AS number
- Prefix
- Validation state:
  - VRP AS number
  - VRP prefix
  - VRP max prefix length
- Associated Regional Internet Registry

The gathered data will be analyzed on a day-to-day basis. This BGP information will be gathered and validated using Bash and Python scripts that gather information from different sources:

- **Routing Table:** The RIS dump[8] provided by RIPE (http://www.ripe.net/) will be used as the source for our routing table. This dump provides the "IPv4/IPv6 Address to Origin Mapping". Only prefixes received from at least 5 RIS peers are added to the routing table, to create a routing table representative of the average BGP routing table. For importing the routing table to our database we had two options, one would be to use the BGP updates provided by the BGPMon XML feed. This feed presents a live stream of BGP updates in XML format and updates approximatly 200,000 prefixes on a daily basis. The other option is to use the RIS dump provided by RIPE. We tested both options and eventually chose to use the RIS dump, because this dump provides the most complete overview of entries in the BGP table on a daily basis. Using the RIS dump we are not able to provide real-time statistics, however for our goals this was not necessary because we generate statistics on a daily basis. The RIS dump provides a good source of all active BGP prefixes that are advertised on the Internet. Each dump can easily be imported to a different SQL table on a daily basis, allowing for easy access to historical data.
- **VRPs:** The cryptographically validated ROA Payloads are imported from the RIPE RPKI validator[9] using the CSV export function of the validator.

- **Region:** The region is established using the RIR associated with each prefix. The prefixes are mapped to a RIR using the IPv4/6 IP-address allocations provided by IANA [11]. Note that these allocations are not 100% accurate but should give a good indication.

This information is stored in a database, with two new tables created every day. One table to store the routes and one to store the ROAs present on that day. After the information has been imported into the database it is processed. The validity state of each route is determined by running a validator script we created. This script queries the routing table for prefixes that match a certain ROA. Once the validity state has been determined, a RIR is associated with the route by comparing the first bits of the prefix to those defined in the IANA IP-address allocations. A logical overview of this process is shown in Figure 3-1.

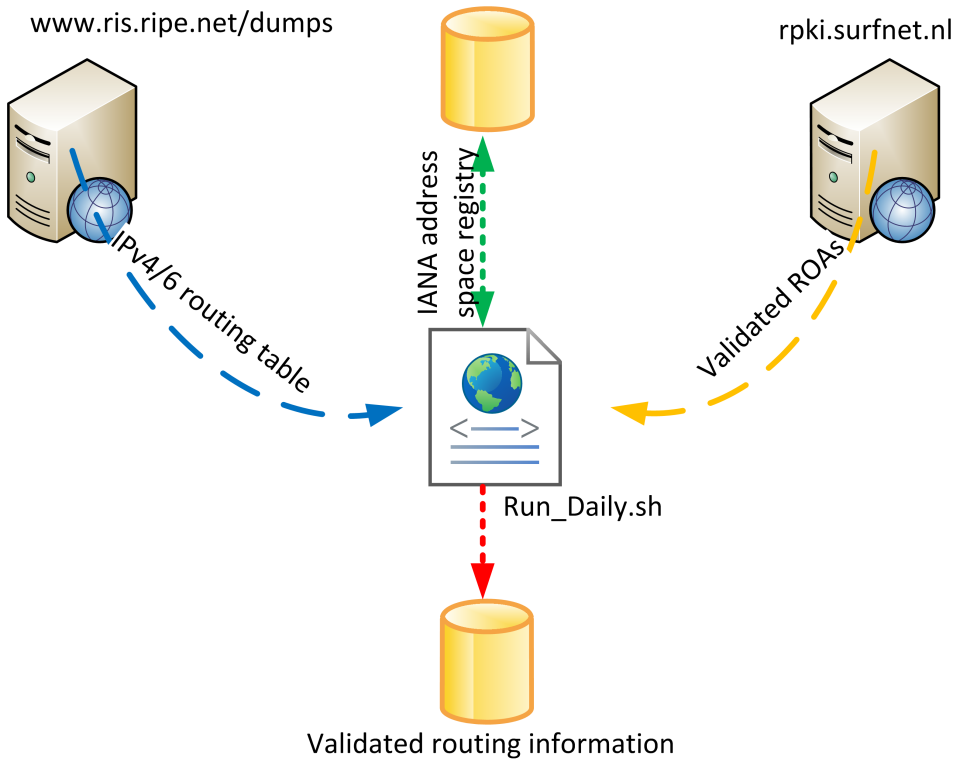

Figure 3-1: Logical overview of the information gathering procedure.

3.2. Generating statistics

The statistics are generated and shown on a web dashboard using PHP and the Google Charts JavaScript API. The values for all the charts are queried from the database using PHP. The JavaScript charts provide a convenient way to present the graphs in an interactive way.

On the dashboard different types of statistics are shown:

- Global RPKI statistics regarding RPKI usage, such as number of valid and invalid routes as well as the distribution of causes of routes being invalid.
- A top-10 list of ASes advertising the most valid and invalid prefixes.
- Comparisons between RPKI use in IPv4 and IPv6 and the distribution of the cause for invalid routes.
- Statistics and trends per specific AS number and RIR.
- A list of valid and invalid prefixes, where the invalid prefixes can be displayed depending on what made the prefix invalid, eg. range length exceeded or AS number mismatch.
- Map charts visualizing RPKI statistics per RIR, showing which regions of the world have the highest RPKI adoption rate and the highest number of valid and invalid routes.
- Trends for RPKI usage, detailing how much the usage of RPKI has increased or decreased. The cause for invalid prefixes is also stored, to show what kind of configuration errors are more prominent during a certain time span.

As some charts require a large amount of database queries, all the dynamic PHP pages are cached into a static HTML page, which is then served to the user. The only pages that are not cached are the statistics pages for each AS, as it would take a considerable amount of time to generate them all and most of the AS pages will likely not be viewed daily.

The code for our data collection and statistics generation scripts can be found here: https://github.com/remydb/RPKI-Dashboard More information about using this code can be found in appendix B.

4. Results

In this section the results of the data analysis will be presented. All the results mentioned in the coming subsections are based on the data collected on July 1st, 2013, unless stated otherwise. Up-to-date data from this project is also available for viewing at http://rpki.surfnet.nl.

4.1. RPKI adoption and invalid routes

The routing table contains 493,712 prefixes, including both IPv4 and IPv6. The validation state has been determined for 14,455 prefixes meaning that these prefixes match at least one ROA. This means that 2.93% of the prefixes in the routing table are covered by at least one ROA. Figure 4-1 shows the distribution between Valid, Invalid and Unknown routes.

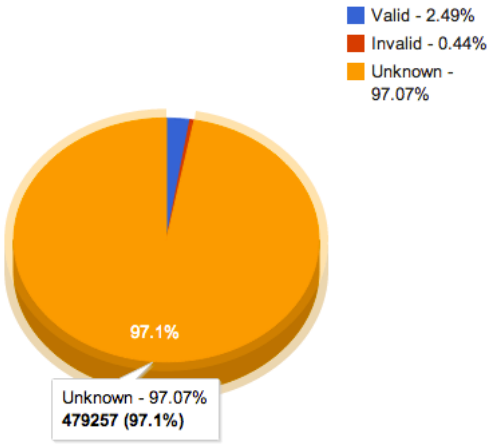

Figure 4-1: State of RPKI adoption as of July 1st, 2013.

Out of the 14,455 prefixes that are covered by a ROA, 12,274 (84.91%) result in a valid route. More important is the number of invalids and their cause. Figure 4-2 shows that a non-matching AS number is the most common cause for an invalid route. This means that there is at least one ROA that covers the prefix but none of the covering ROAs has an AS number that matches the AS that is originating the route. Therefore, the prefix is originated from an unauthorized AS. Although this looks like a hijacking attempt at first, taking a closer look will show an indication of a configuration error. This will be covered in more detail using an example in section 4.4.

The second most common cause for an invalid route is that the maximum prefix length field in the ROA is exceeded. This means that an AS announces a more specific prefix than was allowed in the maximum length field of the ROA. Third, it is often observed that no maximum length is specified in the ROA. When no maximum length is specified in the ROA, every more specific announcement is invalidated. Finally, 368 entries in our

routing table show that both the AS number and the prefix length do not match any of the configured ROAs.

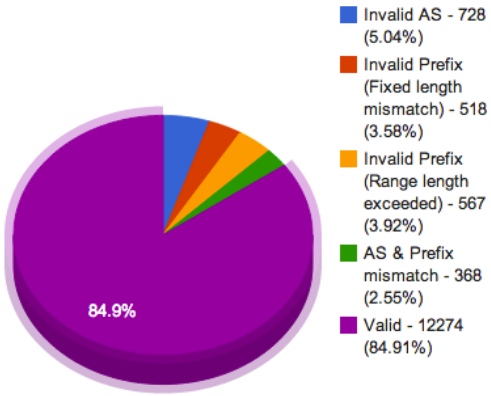

Figure 4-2: Distribution of prefixes covered by ROAs.

Figure 4-3 shows that the number of invalid routes has increased by $((2181-2176)/2176*100) \approx \mathbf{0.23\%}$ between the June 19th, 2013 and June 27th, 2013 while the valid routes have increased with $((12274-12115)/12115*100) \approx \mathbf{1{,}31\%}$ over the same period of time. Although these figures show a trend going in the right direction, we hope to see the number of valid prefixes grow faster in the near future because otherwise RPKI will not be fully implemented in the next few years and therefore not suitable for applying routing policies.

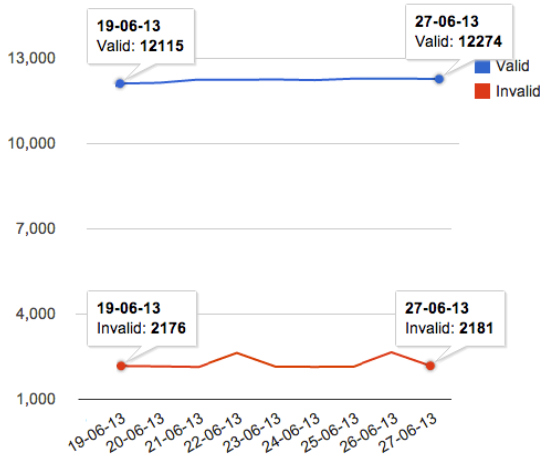

Figure 4-3: Distribution of prefixes covered by ROAs.

4.2. IPv4 vs. IPv6

This section shows a comparison between IPv4 and IPv6 regarding RPKI adoption. Figure 4-4 shows that currently 2.85% of the IPv4 prefixes match at least one ROA while 5.68% of the IPv6 prefixes are covered by at least one ROA. The adoption rate of RPKI is higher for IPv6 prefixes.

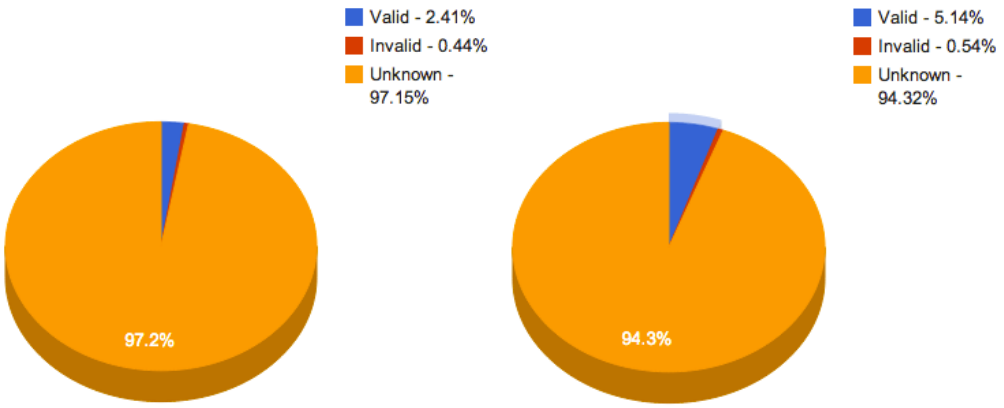

(a) IPv4 (b) IPv6

Figure 4-4: Comparison between IPv4 and IPv6 RPKI adoption

Another interesting difference is the amount of invalid prefixes. Figure 4-5 shows that 9.6% of the IPv6 routes in the routing table are invalid while 15,6% of the IPv4 routes are invalid. We should take into account that the dataset for IPv6 is relatively small compared to IPv4, as there are a lot less IPv6 prefixes being announced. A small amount of changes in the IPv6 dataset could cause significant changes for the IPv6 statistics.

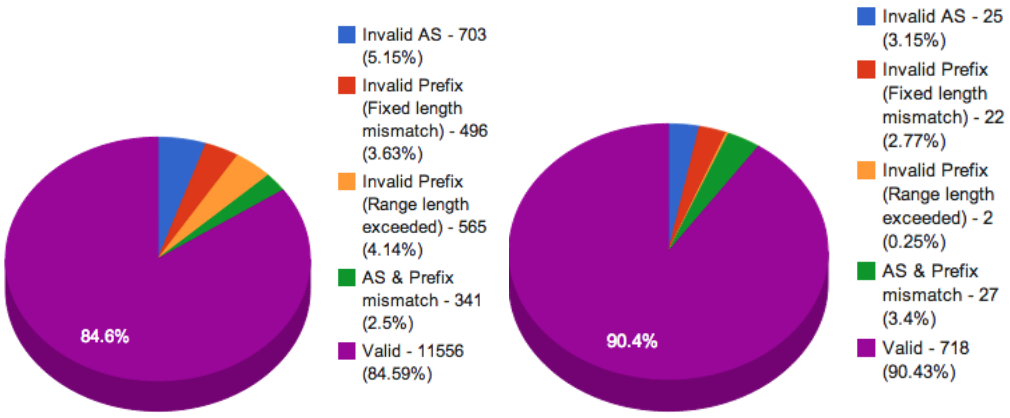

(a) IPv4 (b) IPv6

Figure 4-5: Comparison between IPv4 and IPv6 RPKI validation states

4.3. Regional Internet Registry comparison

The IANA Address Space Registry[11] was used to determine to which RIR a specific prefix has been assigned.

RIR	Total	Valid	Invalid	Unknown	RPKI adoption
AFRINIC	10700	0 (0%)	0 (0%)	10700 (100%)	0%
APNIC	116458	84 (0.07%)	214 (0.18%)	116160 (99.74%)	0.26%
ARIN	181974	200 (0.11%)	30 (0.02%)	181744 (99.87%)	0.13%
LACNIC	56398	5565 (9.87%)	1184 (2.1%)	49649 (88.03%)	11.97%
RIPE	128785	6421 (4.99%)	1147 (0.89%)	121217 (94.12%)	5.88%

Table 4-1: Comparision between Regional Internet Registry RPKI statistics using data gathered on 02-07-2013

Table 4-1 shows that LACNIC currently has the highest adoption rate of RPKI. This can mostly be credited to AS10620 which belongs to Telmex Colombia S.A. and is announcing 2184 valid IPv4 prefixes. What also should be noted is that APNIC is the only RIR that is responsible for more invalids than valids. Nepal Telecommunications Corporation is mostly responsible for these numbers because they are originating 119 invalid routes from AS23752.

4.4. Top 10

A partial goal of this research is to determine which ASes are advertising invalid routes due to misconfigurations. Therefore we have created a list of ASes sorted by the number of invalid routes that are announced from an AS. Currently that results in the top 10 shown in Figure 4-6.

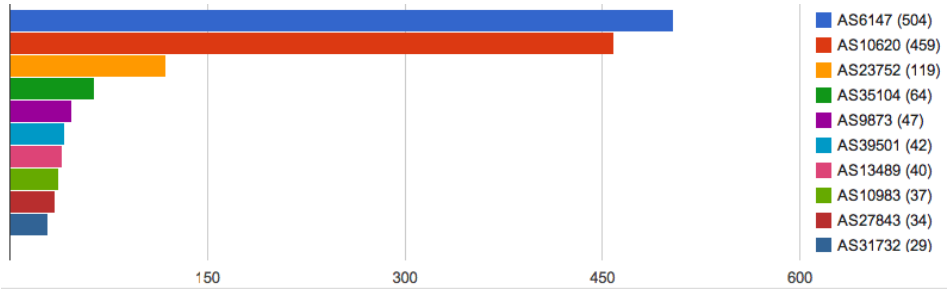

Figure 4-6: Top 10 of invalid routes announced per AS.

To find out the root cause of these potential misconfigurations we attempted to contact the top 10 ASes advertising invalid routes. Unfortunately, none of the operators of any of the ASes responded to our enquiries, so the root cause of their invalid routes remains uncertain, though the statistics often give a good idea of what may have gone wrong. For example, an AS which advertises a /16 prefix and has a ROA with a fixed length set

for that prefix may also advertise more specific /24s. These /24 prefixes will be marked invalid and the cause of this is the unnecessary advertisements for the /24 prefixes which are already advertised as an aggregated /16. If, for some reason, it is necessary to advertise the more specific /24s, the operator should configure a higher maximum prefix length or should configure additional ROAs for the /24s.

Another possibility could be a situation observed for AS10620 on 01-07-13, namely the advertisements of prefixes which have a ROA specifying a different AS number than the one they are being advertised from. AS10620 belongs to the Colombian company Telmex Colombia, but it is not the only AS belonging to Telmex. Telmex also owns AS14080. The ROA for 181.48.0.0/13-24 states AS14080 to be the originating AS, however these prefixes are actually originating from their other AS, namely AS10620. As such, a lot of these prefixes are marked invalid. If Telmex would add ROAs for each of these prefixes with the originating AS set to AS10620, hundreds of currently invalid routes would turn valid. With this information we can safely assume that these invalids are caused by configuration errors. This holds true for all the ASes listed in our top-10, as when we inspected each of these more closely we found that most of their invalids were caused by one of the common configuration errors mentioned in this report.

4.5. Per AS statistics

To assist operators and RIRs with the correction of the aforementioned errors we have created a page that gives an overview of the RPKI statistics based on an AS number or RIR. This page shows operators the following information:

- Distribution of valid, invalid and unknown routes.
- Cause of invalids, shown in a pie chart.
- Monitoring of valid and invalid prefixes over the course of time in a line chart.
- A table showing all prefixes originating from the AS, including the validation state, potential cause of invalid state and AS number associated with the prefix in the whois database[10].

Using the data on this page, operators will be able to view what mistakes they have made and fix them accordingly. Each invalid or valid route will show its related ROAs and the AS number associated with it in the whois database. Combining these two pieces of information should give operators enough information to find the cause of the error and help them fix it. For example, if a prefix is invalid due to a mismatching AS, the ROA will show a different AS number and the whois query for the prefix could return a different AS number. Likewise, if the prefix is invalid due to a mismatching prefix length, the ROA will show that the prefix length does not match the length specified in the ROA and as such results in an invalid prefix.

Finally we try to give a better insight into the causes of invalids over the course of time by plotting these values in a line chart. Figure 4-7 gives an example. As is apparent from figure 4-7, there is definite fluctuation in the cause of invalids. On June 22nd, 26th

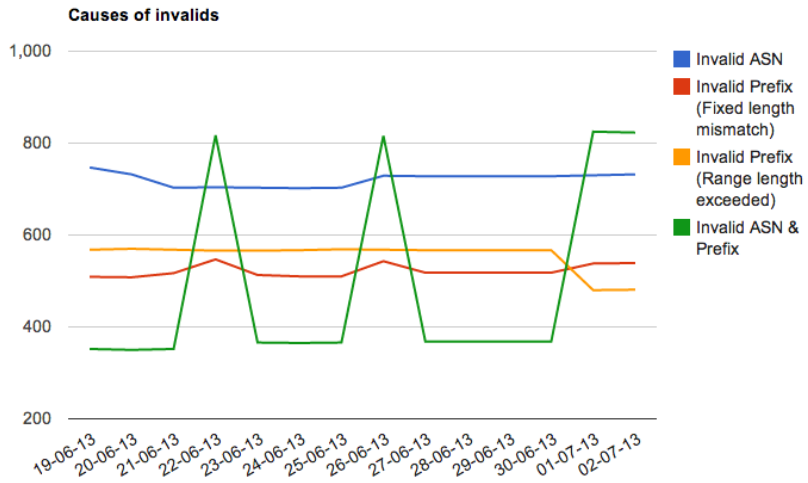

Figure 4-7: Monitoring the cause of invalids over the course of time.

and July 1st there are spikes in the graph for invalids caused by both a non-matching AS number and an exceeded prefix length. Though we do not know the exact reason for this, we expect it to be caused by advertisements being sent for certain ASes on these days, but not on the others. Prefixes derived from the RIPE RPKI validator show the same behaviour. Listing 4-1 shows the difference between a day without a peak of invalid prefixes (June 30th) and one with a peak (July 1st). The top 10 contains 7 different ASes on July 1st, compared to June 30th.

```
+--------+-------+ +--------+-------+
| June 30th, 2013| | July 1st, 2013 |
+--------+-------+ +--------+-------+
| ASN    | count | | ASN    | count |
+--------+-------+ +--------+-------+
|  35104 |    64 | |   2065 |    92 |
|   8649 |    16 | |   1942 |    82 |
| 197890 |    16 | |   1724 |    66 |
|   2199 |    13 | |  35104 |    64 |
|  27947 |    12 | |   2457 |    62 |
|   2471 |    12 | |   1937 |    39 |
|  24954 |     9 | |   1945 |    33 |
|  27817 |     9 | |   1723 |    32 |
| 197860 |     8 | | 197890 |    17 |
|   9180 |     7 | |   8649 |    16 |
+--------+-------+ +--------+-------+
```

Listing 4-1: Table showing the top 10 ASes announcing the most invalids caused by both a non-matching AS number and an exceeded prefix length.

5. Conclusion

To reliably monitor RPKI statistics over time, it is important to gather reliable data on a day to day basis. RIPE provides a good global BGP routing table for both IPv4 and IPv6. This routing table is publicly available through the RIPE RIS dumps and is used as a source for the data set used in this project. ROAs are downloaded from the public repositories of all five RIRs and then cryptographically verified by the local RPKI validator. All information is then combined into a database and RIR information is added based upon the IANA address space assignments.

The complete data set is used for generating statistics on a daily basis. These statistics are presented on a dashboard using the Google Chart API. The data presentation should give BGP operators and employees of regional internet registries a better overview of their RPKI configuration. Information about the cause for invalids, origin of invalids and changes over time are displayed on a per AS, per RIR and global basis.

Current statistics (July 2nd, 2013), show that RPKI is mainly being adopted in the RIPE (5.88%) and LACNIC (11.97%) region. All other regions have an adoption rate below 0.25%. There are 14,845 prefixes covered by ROAs present in the routing table of which 2575 are invalid. Three ASes are responsible for announcing approximately 40% of the invalid routes, AS10620 (460 invalids), AS6147 (418 invalids) and, AS23752 (120 invalids). The most common reason for invalid routes is a non-matching Autonomous System number (728 invalids, 28% of total invalid routes). This means that there is at least one ROA that covers the prefix but none of the covering ROAs has an AS number that matches the AS that is originating the route.

The root cause of most configuration errors remains uncertain, we can not even be sure that all of the invalid routes we encounter are caused by configuration errors. We have attempted to contact the operators of the top 10 ASes advertising invalid routes. Unfortunately, none of the operators responded to our enquiries. However, the statistics give a good idea of common causes for configuration errors. A common mistake is leaving the optional "Max Length" field blank. When this field is left blank, only the exact same prefix length is accepted. Any more specific announcement is invalid, even if it is announced from the authorized AS. Another common mistake is configuring ROAs for AS X while announcing the prefixes from AS Y. Although this looks like a hijacking attempt, we've concluded that in most cases this is a configuration error because both AS numbers belong to the same organization.

The rate of RPKI adoption is still low and a lot of configuration errors occur without being corrected. Hopefully our research will help improve the rate of RPKI adoption and improve the correctness of configurations by providing insight into what mistakes are being made.

6. Future work

This project should give a good overview of the current RPKI adoption statistics. Nevertheless, further research can be done as to why certain RIRs and ASes adopt faster than others. The RPKI data which we gather every day might be interesting for further analysis in the near future. It would also be interesting to see how the performance of our application can be further improved for generating charts on the fly, which can take a substantial amount of time at the moment due to the large amount of MySQL queries being performed. As mentioned before this application was build for functionality but as the adoption of RPKI grows further, possible performance improvements will need to be researched, as the validating of prefixes is currently a very slow process which does not scale well at all.

A. Bibliography

[1] **'www.potaroo.net'**, "BGP Routing Table Analysis Report". http://bgp.potaroo.net/as1221/bgp-active.html

[2] **'RIPE NCC'**, "YouTube Hijacking: A RIPE NCC RIS case study", 2008. http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

[3] **'Matthias Wahlisch, Olaf Maennel, Thomas C. Schmidt'**, "Towards Detecting BGP Route Hijacking using the RPKI", 2012. http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p103.pdf

[4] **'Matthias Wahlisch'**, "One Day in the Life of RPKI", 2011. http://readonly.labs.ripe.net/Members/waehlisch/one-day-in-the-life-of-rpki

[5] **'G. Huston, R. Loomans, G. Michaelson'**, "A Profile for Resource Certificate Repository Structure", 2012. http://tools.ietf.org/html/rfc6481

[6] **'Multiple Authors'**, "BGP Prefix Origin Validation", 2013. http://tools.ietf.org/html/rfc6811

[7] **'R. Bush, R. Austein'**, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", 2013. http://tools.ietf.org/html/rfc6810

[8] **'RIS Whois Dump'**, "IPv4/IPv6 Address to Origin Mapping" http://www.ris.ripe.net/dumps/

[9] **'RIPE NCC'**, "RPKI Validator 2.9", 2013. http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources

[10] **'RIPE Whois database'**, "Database containing registration details of IP addresses and AS numbers" https://www.ripe.net/data-tools/support/documentation/ripe-database-fast-facts

[11] **'Address Space Registry'**, "The allocation of IPv4 and IPv6 address space.", 2013-05-20. http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.txt http://www.iana.org/assignments/ipv6-address-space/ipv6-address-space.txt

B. Scripts

All the scripts are available on Github (https://github.com/remydb/RPKI-Dashboard). The project consists of two parts, namely the scripts for collecting data and a web-dashboard (https://github.com/remydb/RPKI-Dashboard/tree/dashboard) for presenting the statistics of data.

To run this software there are some prerequisites:

- MySQL server
- Database called 'bgp'
- Python (v2.7), Python-ipaddr, Python-mysqldb
- PHP5, PHP5-mysql

To set up the dashboard, perform the following steps:

1. Clone the files from the 'master' branch
2. Set up a MySQL database
3. Edit the database credentials for the following files:
   a) run_daily.sh
   b) insertrirs.py
   c) validate_table.py
4. Clone the files from the 'dashboard' branch and place them in your website root (or where ever you want them to be).
5. Edit the 'include/functions.php' file and change the mysql username and password
6. Edit the 'createstatic.sh' file to place the static HTML files in the website root
7. Add rules to your crontab for running 'run_daily.sh' and 'createstatic.sh' every day, eg:

```
1 0 1 * * * /home/xxx/scripts/run_daily.sh > /dev/null
    2>&1
2 0 5 * * * /home/xxx/scripts/createstatic.sh > /dev/null
    2>&1
```

These lines were added to the crontab of the root user, as some file permissions have to be changed during the 'run_daily.sh' script. Note that there is a lot of time between the running of 'run_daily.sh' and the creating of the static pages, as the validation part of the information gathering can take quite a bit of time. The actual running time of the script depends on the amount of ROAs exported from the validator and the specs of the server running it. On our virtual machine, processing roughly 5400 ROAs took approximately 40 minutes.

C. Distribution of work

Remy has focussed on the following aspects of the project:

- Researching current solutions
- Developing the tools for the gathering of data
- Developing the dashboard and required functions
- Writing report

And Javy has focussed on:

- Researching current solutions
- Devising ways of visualizing statistics using Google Chart
- Analysis and presentation of results
- Presentation of work
- Writing report