



Global Routing Security

Job Snijders
NTT Communications (AS 2914)

job@ntt.net

LACNIC 29



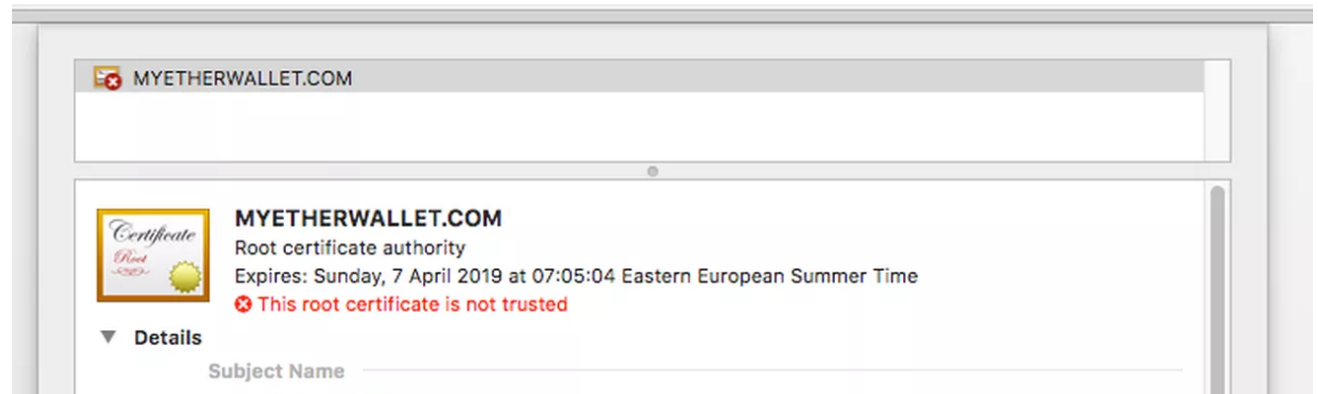
What is this about

- Relevance of routing security
- How IRRs work in other regions
- Information unique to LATAM
- Proposed future work
- Time for questions!

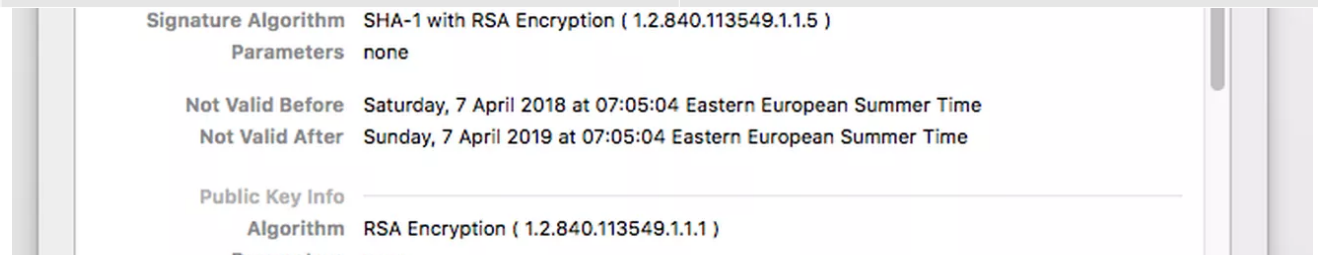
BGP Hijacking is lucrative



Amazon Route53 / MyEtherWallet.com hijack




Auth Nameserver	Original	Hijacked
205.251.192.73 ns-73.awsdns-09.com	205.251.192.0/23 AS 16509	205.251.192.0/24 AS 10297 205.251.193.0/24 AS 10297
205.251.195.239 ns-1007.awsdns-61.net	205.251.194.0/23 AS 16509	205.251.195.0/24 AS 10297
205.251.197.218 ns-1498.awsdns-59.org	205.251.196.0/23 AS 16509	205.251.197.0/24 AS 10297
205.251.199.201 ns-1993.awsdns-57.co.uk	205.251.198.0/23 AS 16509	205.251.199.0/24 AS 10297



It could've been worse!

- The AS 10297 upstreams (NTT, Cogent, Level3) & Equinix route server blocked the hijack attack
- Some peers of AS 10297 (Google, Hurricane Electric, BBOI) accepted the hijack
- Hijack impact was limited **thanks to filters**, but still an absolute disaster for all involved
- More info:
<https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/>

Mistakes happen...



About 295 results (0.26 seconds)



Here's why you may have had internet problems today

CNNMoney - Nov 6, 2017

According to reports from **Down Detector**, a website that monitors internet **outages**, **Comcast** and **Level 3** connectivity was impacted nationwide beginning around 10 a. Pacific. Other internet service providers including Spectrum, Verizon, and AT&T showed a spike in connectivity issues, too, though they ...

Comcast's nationwide outage was caused by a configuration error

Engadget - Nov 7, 2017

Level3 Service Disruption Causes Nationwide Internet Outages

Patch.com - Nov 6, 2017

Comcast, Others Dealing With Internet Outage

Multichannel News - Nov 6, 2017

Comcast's Xfinity internet service is reportedly down across the US

Highly Cited - The Verge - Nov 6, 2017

There Were Widespread Internet Outages Reported Monday in the US

Blog - Slate Magazine (blog) - Nov 6, 2017



Engadget



Slate Magazi...




Patch.com



Multichannel ...



FierceCable



The Verge

[View all](#)

3 reasons to filter

- Creating filters based on public data, forces malicious actors to leave a trail in IRR, WHOIS or other data sources: **audit-ability**
- **Bugs happen:** your router may suddenly ignore parts of your configuration, you'll then rely on your BGP peer's filters
- **Everyone makes mistakes** – a typo is easily made



Filtering recap

- 1) Reject RFC 1918 (private) IP space
- 2) Reject Bogon/Private ASNs
- 3) Reject IXP Nets
- 4) Allow what is registered in IRR, WHOIS, RPKI**
- 5) Reject all other BGP announcements



What is the IRR

- “Internet Routing Registry”
- What companies like NTT uses as a source to generate per customer prefix filters
- Publicly available, to help debugging and provide transparency
- By making our source for filter generation publicly available, other parties can inspect what we take into consideration.

What sources are there?

- IRR Sources offered by Regional Internet Registries (RIPE, APNIC, ARIN, etc)
- IRR Sources operated by “third parties” (like RADB, NTT, etc)
- WHOIS sources (ARIN WHOIS, Registro.BR)
- RPKI sources (LACNIC, RIPE, etc)
- In total there are ~ 40 sources, but NTT only uses 14 of them
- The sources are **NOT** equal, some operate by different rules than others

A route object: the atom

```
$ whois -h rr.ntt.net 192.147.168.0/24
```

```
route:           192.147.168.0/24  
descr:          Job Snijders  
origin:         AS15562  
notify:        job@instituut.net  
mnt-by:        MAINT-JOB  
changed:       job@ntt.net 20161003  
source:        NTTCOM
```

(only the bold lines are relevant in the process)

Generating a prefix filter

```
job@vurt ~$ whois -h rr.ntt.net '!gAS15562'  
A212  
165.254.255.132/32 165.254.255.26/32  
165.254.255.0/25 165.254.255.144/28  
165.254.255.133/32 192.147.168.0/24  
165.254.255.160/28 165.254.255.149/32  
209.24.0.0/16 204.42.254.192/26  
165.254.255.0/24 67.221.245.0/24  
C  
job@vurt ~$
```

Grouping ASNs: AS-SETS

```
job@vurt ~$ whois -h rr.ntt.net AS15562:AS-SNIJDERS
```

```
as-set:                AS15562:AS-SNIJDERS
members:            AS15562          # Me
members:              AS57436          # Samer
members:              AS-KING          # Thomas King
members:              AS-NETHER       # Jared
tech-c:               DUMY-RIPE
admin-c:              DUMY-RIPE
notify:               job@instituut.net
org:                  ORG-SNIJ1-RIPE
mnt-by:               SNIJDERS-MNT
created:               2018-01-16T17:54:54Z
last-modified:        2018-01-16T17:58:36Z
source:               RIPE
```

Systematic access to AS-SETS

```
$ whois -h rr.ntt.net '!iAS15562:AS-SNIJDERS,1'  
A130  
AS15562 AS202539 AS205591 AS205593 AS206479  
AS206499 AS206551 AS234 AS267 AS31451 AS41731  
AS49697 AS51861 AS57436 AS60003 AS61438  
C
```



Wrapping it up:

- An AS-SET is resolved into all its member ASNs
- For each ASN we do a reverse lookup to find all route-objects where the ASN is the “origin:”



How one IRR source is unlike the other..

- Not all IRRs are equal
- They differ in terms of ownership, purpose, policy, validation
- All of IRR is “garbage in, garbage out”
- Some RIRs offer good training materials on how to use the IRR
- Some IRRs have fancy web interfaces, some require interaction via email

Differences #1

- In NTTCOM, any customer can create any route object for any prefix (if it hasn't been covered by another route object in NTTCOM)
- In RADB anyone that pays \$500 per year can create any route object for any prefix (if it hasn't been covered by another route object in RADB)

Differences #2

- In ARIN, any ARIN member can create any route object for any prefix (if it hasn't been covered by another route object in ARIN)
 - ARIN staff is working to fix this!
- In ARIN WHOIS, only the owner of the IP block can specify an Origin AS
 - More information:
<https://medium.com/@jobsnijders/a-new-source-for-authoritative-routing-data-arin-whois-5ea6e1f774ed>

Differences #3

- In RIPE, only the owner of the IP block can create/designate route objects. Except when it isn't RIPE managed space... then anyone can create any route object for any prefix (if it hasn't been covered by another route object in RIPE)
- In the future RIPE will show the difference between route-objects for which it is authoritative and and which ones it isn't by showing: "source: RIPE" and "source: RIPE-NONAUTH"

Differences #4

- In the APNIC and AfriNIC database you can only create route-objects for APNIC/AfriNIC managed space **AND** with approval from the IP block owner, but not approval from the ASN owner.
 - This is the most sane approach, cleanest data

LATAM challenges


1. LACNIC does not offer an IRR...
 - But there is excellent RPKI data
2. Not all countries in LACNIC region have RPKI..
 - But there are excellent WHOIS databases such as `registro.br`
3. In absence of common, trustworthy, IRR, the creation of AS-SETS is cumbersome
 - RPKI does not yet fill this gap

What about RPKI?

- A RPKI ROA kind of looks like a route object
- It has a “prefix” and an “origin”
- RPKI is trustworthy data, we know for sure that the owner of the IP space created the ROA
- RPKI ROAs are “higher” (more important) than IRR route-objects
- 23% of LACNIC prefixes are **RPKI VALID!**

<https://rpki-monitor.antd.nist.gov/?p=4&s=0>

Provisioning use case for RPKI data?

← → ↻ ⓘ Not secure | localcert.ripe.net:8088/roas ☆ 

RPKI Validator Home Trust Anchors **ROAs** Ignore Filters Whitelist BGP Preview Export and

Validated ROAs

Validated ROAs from **APNIC from AFRINIC RPKI Root, APNIC from ARIN RPKI Root, APNIC from IANA RPKI Root, LACNIC RPKI Root, APNIC from RIPE RPKI Root, ARIN RPKI Root, AfrinIC RPKI Root, LACNIC RPKI Root, RIPE NCC Pilot (RRDP prefetch), RIPE NCC RPKI Root, RIPE NCC RPKI Root (RRDP prefetch), RIPE NCC prepdev (RRDP prefetch), altca, apnic-testbed.**

Show entries

Search:

ASN	Prefix	Maximum Length	Trust Anchor
15562	2001:67c:208c::/48	48	RIPE NCC RPKI Root

Simple RPKI ROA example

```
job@vurt ~$ ftp -VM -o - \  
    http://localcert.ripe.net:8088/export.json \  
    | jq '.roas[] | select(.asn | contains("AS15562"))? | .prefix' \  
    | uniq  
"2001:67c:208c::/48"  
job@vurt ~$
```


Simple registro.br example

```
job@vurt ~$ whois -h registro.br 200.160.4.6 | grep -A 1 inetnum
inetnum:      200.160.0/20
aut-num:      AS22548
```

```
job@vurt ~$ grep AS22548 nicbr-asn-blk-latest.txt
```

```
AS22548|Núcleo de Inf. e Coord. do Ponto BR - NIC.BR|
05.506.560/0001-36|200.160.0.0/20|2001:12ff::/32
```

- Computer parseable Registro.br data dump:
 - <ftp://ftp.registro.br/pub/numeracao/origin/nicbr-asn-blk-latest.txt>(Thank you Frederico Neves!)
- 56% of NICBR WHOIS entries exact match with BGP DFZ, quite accurate!



The problem with IRR AS-SETs

- We don't really know what AS-SET belongs to what ASN
- There can be duplicate AS-SETs in different IRR databases
- We don't know if the owner of the ASN created the AS-SET

Conclusion: Commonly used, but far from ideal



RPKI “AS-Cones” as replacement for IRR AS-SET

- Ease of discovery
 - given ASN X – what list of your downstream customers I should use in my provisioning system?
- Guarantees that only the owner of the ASN could’ve created that list
- Unilateral declarations (just like AS-SETs)
- Per adjacent ASN granularity:
 - AS 15562 may announce a different set of downstreams to NTT than to GTT ^{27 /}



RPKI AS-Cones

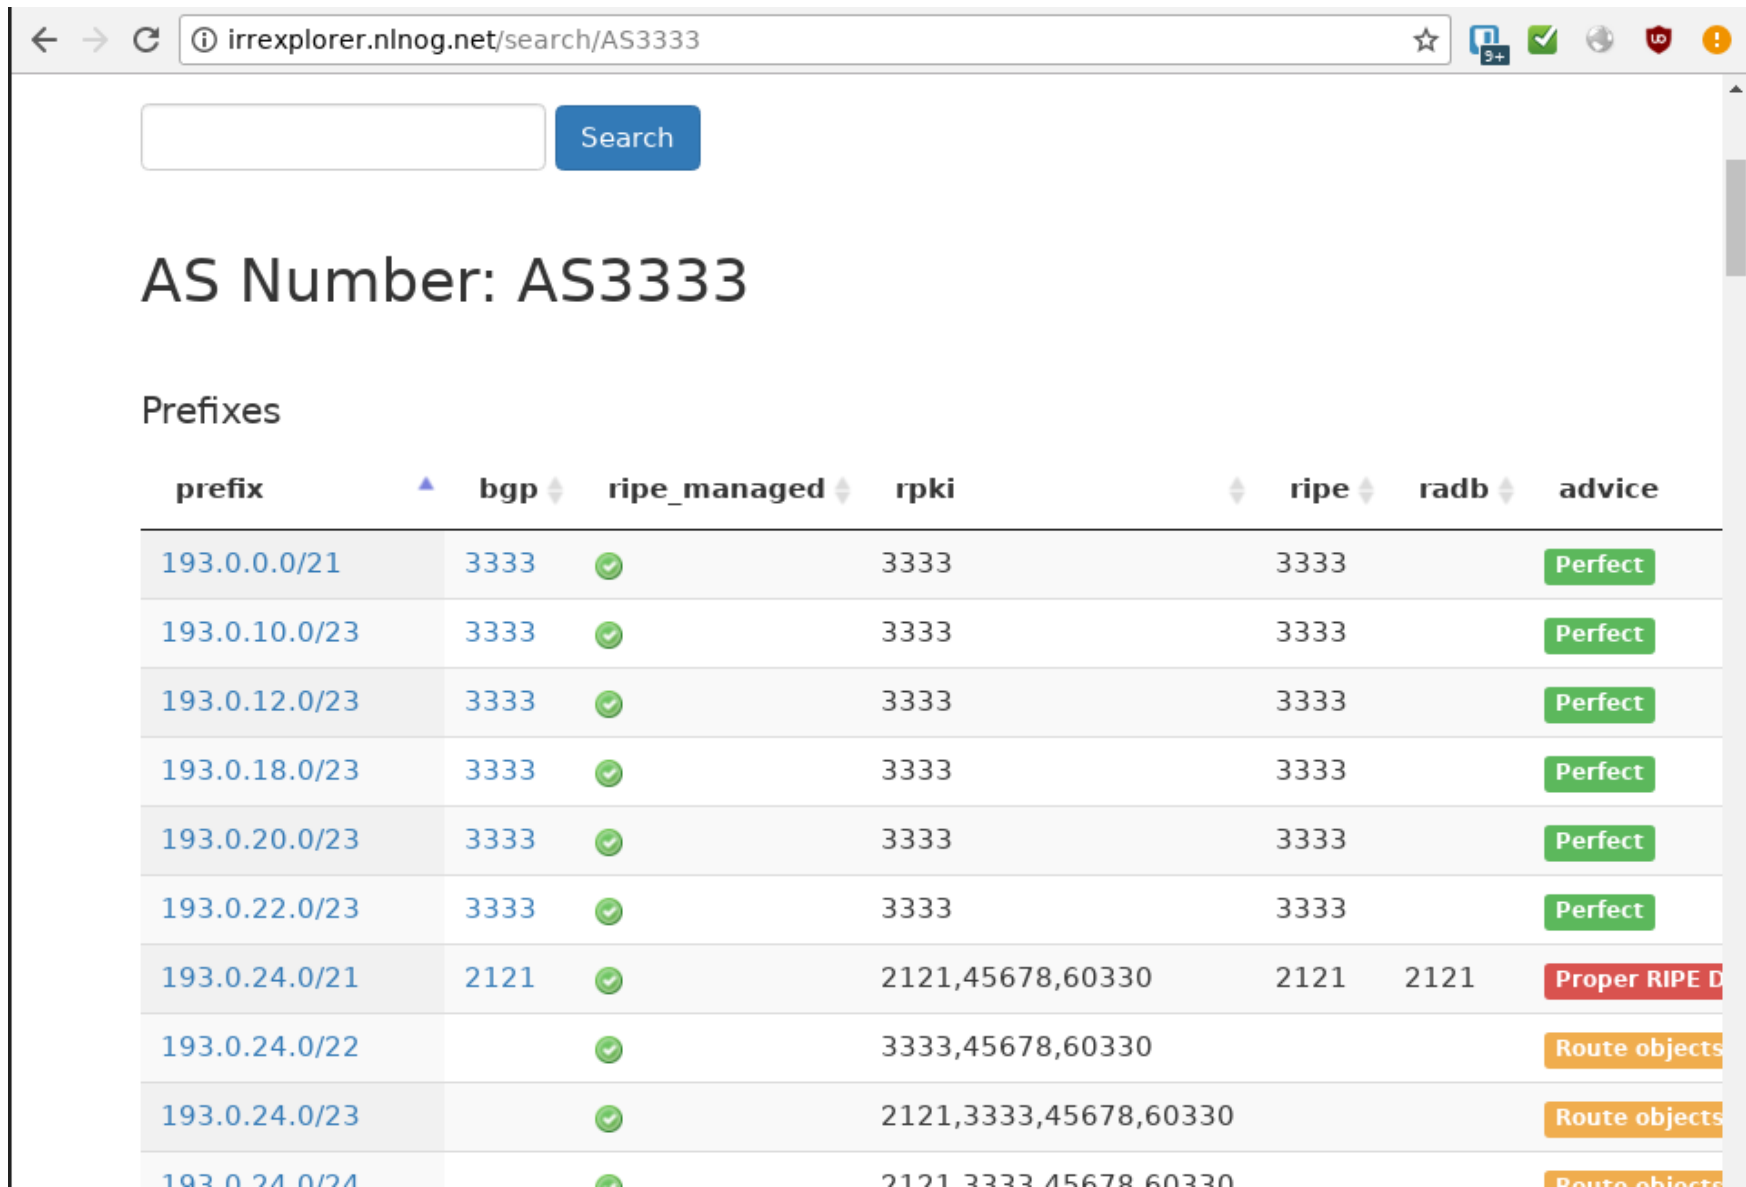
IETF Internet-Draft:

<https://tools.ietf.org/html/draft-ss-grow-rpki-as-cones>

Discussion & feedback welcome in IETF GROW
Working Group!

Hoping for help and feedback from network
operators, LACNIC, NIC.br, and NICMx! :-)

http://irrexplorer.nlnog.net



irrexplorer.nlnog.net/search/AS3333

Search

AS Number: AS3333

Prefixes

prefix	bgp	ripe_managed	rpki	ripe	radb	advice
193.0.0.0/21	3333	✓	3333	3333		Perfect
193.0.10.0/23	3333	✓	3333	3333		Perfect
193.0.12.0/23	3333	✓	3333	3333		Perfect
193.0.18.0/23	3333	✓	3333	3333		Perfect
193.0.20.0/23	3333	✓	3333	3333		Perfect
193.0.22.0/23	3333	✓	3333	3333		Perfect
193.0.24.0/21	2121	✓	2121,45678,60330	2121	2121	Proper RIPE D
193.0.24.0/22		✓	3333,45678,60330			Route objects
193.0.24.0/23		✓	2121,3333,45678,60330			Route objects
193.0.24.0/24		✓	2121,3333,45678,60330			Route objects

Todo list for the community

- Going to IETF to define “AS-SETS in RPKI”
- Carriers like NTT should start using WHOIS & RPKI data in BGP-4 filter generation
- Make RPKI available in all LATAM countries
- Use RPKI to “drown out” proxy IRR objects
- Write a new IRRd (for `rr.ntt.net`) from scratch: **IRRdv4**
 - Allow for innovation, integration with the RIRs and NIRs