



# IRR 101

Job Snijders, DKNOG 8  
[job@ntt.net](mailto:job@ntt.net)



# What is this about

- Just a refresher on how things work today
- Using RPKI in context of provisioning



# Filtering recap

- 1) Reject RFC 1918 (private) IP space
- 2) Reject majority of Bogon/Private ASNs
- 3) Reject IXP Nets
- 4) Reject based on Peerlock rules
- 5) Allow what is registered in IRR, WHOIS, ... and RPKI?**
- 6) Reject the rest



# What is the IRR

- “Internet Routing Registry”
- What 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?

- Sources offered by Regional Internet Registries (RIPE, APNIC, ARIN, etc)
- Sources operated by “third parties” (like RADB, NTT, 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 ~$
```

# Same example, with bgpq3

```
job@vurt ~$ bgpq3 -h rr.ntt.net -l AS15562-in AS15562  
no ip prefix-list AS15562-in  
ip prefix-list AS15562-in permit 67.221.245.0/24  
ip prefix-list AS15562-in permit 165.254.255.0/24  
ip prefix-list AS15562-in permit 165.254.255.0/26  
ip prefix-list AS15562-in permit 165.254.255.26/32  
ip prefix-list AS15562-in permit 165.254.255.64/26  
ip prefix-list AS15562-in permit 165.254.255.132/32  
ip prefix-list AS15562-in permit 165.254.255.133/32  
ip prefix-list AS15562-in permit 165.254.255.144/28  
ip prefix-list AS15562-in permit 165.254.255.149/32  
ip prefix-list AS15562-in permit 165.254.255.160/28  
ip prefix-list AS15562-in permit 192.147.168.0/24  
ip prefix-list AS15562-in permit 204.2.30.0/23  
ip prefix-list AS15562-in permit 204.42.254.192/26  
ip prefix-list AS15562-in permit 209.24.0.0/16
```



# 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:                DUMMY-RIPE
admin-c:               DUMMY-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 '!AS15562:AS-SNIJDERS,1'  
A130  
AS15562 AS202539 AS205591 AS205593 AS206479  
AS206499 AS206551 AS234 AS267 AS31451 AS41731  
AS49697 AS51861 AS57436 AS60003 AS61438  
C
```



```
$ irrtree AS15562:AS-SNIJDERS
```

```
Processed: 0 objects (Elapsed Time: 0:00:00)
```

```
IRRTree (1.1.3) report for 'AS15562:AS-SNIJDERS'  
(IPv4), using rr.ntt.net at 2018-01-24 16:23
```

```
AS15562:AS-SNIJDERS (16 ASNs, 25 pfxs)
```

```
+-- AS-KING (12 ASNs, 8 pfxs)
```

```
|   +-- AS205591 (2 pfxs)
```

```
|   +-- AS206499 (2 pfxs)
```

```
|   +-- AS49697 (2 pfxs)
```

```
|   +-- AS51861 (1 pfxs)
```

```
|   +-- AS60003 (1 pfxs)
```

```
+-- AS-NETHER (2 ASNs, 4 pfxs)
```

```
|   +-- AS267 (2 pfxs)
```

```
|   +-- AS234 (2 pfxs)
```

```
+-- AS15562 (12 pfxs)
```

```
+-- AS57436 (1 pfxs)
```

# 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:"`
- The total list of prefixes from the above 2 steps is the input into ``bgpq3`` and ends up on routers
- The order matters, more on that later



# What when duplicate AS-SET?

- What if AS-STEALTH exists in multiple IRRs?
- AS-STEALTH exists in **both** RIPE and RADB
- The two are not managed by the same company

# So.. that's what that 'sources' is about

← → ↻ Secure [REDACTED]gtools/netgeek/autnum/1

Home Configtools EMACS ISIS Map Vendor Settings

ASN \* 15562

Name \* snijders-as

IPv4 as-set AS15562:AS-SNIJDERS

IPv6 as-set AS15562:AS-SNIJDERS

Peertype \* customer ▼

Contacts

Internet Routing Registry (IRR) Sources RIPE,RADB,NTTCOM

ACL Mail job@instituut.net

IPv4 Prefix Count 25

# Order is important

```
$ whois -h rr.ntt.net '!j-*'
[Querying rr.ntt.net]
[rr.ntt.net]
A340
NTTCOM:Y:939977-949960:949958
INTERNAL:Y:629-28859:28859
RADB:Y:3645499-3646474
RIPE:Y:40749570-40751305
ALTDB:Y:42283-44199
BELL:Y:248417-248891
LEVEL3:Y:777495-777945
RGNET:Y:139-139
GT:Y:1004-4246
APNIC:Y:6317163-6324052
JPIRR:Y:108306-129968
ARIN:Y:76179-79625
BBOI:Y:6715-6843
TC:Y:8802-17353
AFRINIC:Y:635751-651232
ARIN-WHOIS:Y:0-84
C
```

# Changing the order

```
$ bgpq3 -S RIPE,RADB -h rr.ntt.net -A AS-STEALTH | wc -l  
175
```

```
$ bgpq3 -S RADB,RIPE -h rr.ntt.net -A AS-STEALTH | wc -l  
1455
```

(bgpq3 is a small open source tool to generate prefix filters from the IRR, we don't use it, but if we would we'd have smaller configs ;-)



# 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)
- In ARIN WHOIS, only the owner of the IP block can specify an Origin AS

# 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 RIPE database, when you create route objects, both the IP block owner and the owner of the Origin ASN have to approve it. So double authentication is required.
- This is going to change in 2018, only the IP owner has to approve: aligns with APNIC, others

# Differences #5

- 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, and no approval from the ASN owner.
- This is the most sane approach, cleanest data
- JPIRR requires you to reconfirm a route object every year

# Summary: what the F\*\*\*



# What is this ARIN WHOIS thing?

- Remember from the fifth slide that we **only** care about the CIDR + Origin AS tuple?

Dashboard

Tickets & Messages 7

Your Account

Settings  
Profile and security information

Point of Contact records  
View and manage POCs

Organization Identifiers  
View and manage Org IDs

Associations Report  
Records connected to your account

IP Addresses

Search  
View and manage your networks

## View & Manage Network

### Information

#### NETWORK INFO

Net Range: 198.51.100.0 - 198.51.100.255

CIDR: 198.51.100.0/24

Origin AS: **AS19384**

Net Name: TEST-NET-2

Net Handle: NET-198-51-100-0

Parent: [NET-198-51-100-0](#)

Public Comments: THIS NETWORK IS NON-PORTABLE

Registered Date: 09-26-2008 15:36:10

Last Modified Date: 09-26-2008 15:36:10





# Using ARIN WHOIS

- It is a trustworthy authoritative source of data
- We download a 3.5GB XML dump from ARIN once a day
- We convert the XML into “route:” objects
- We load those IRR objects into rr.ntt.net
- This is to offer customers easier choices
  - Setting up IRR can be cumbersome

# Example output

```
$ whois -h rr.ntt.net 199.43.0.44
route:          199.43.0.0/24
descr:           NET-199-43-0-0-1
origin:        AS10745
remarks:         This route object represents authoritative
                  data retrieved from ARIN's WHOIS service.
remarks:         The original data can be found here:
                  https://whois.arin.net/rest/net/NET-199-43-0-
0-1
remarks:         This route object is the result of an
                  automated WHOIS-to-IRR conversion process.
mnt-by:          MAINT-JOB
changed:         job@ntt.net 20150715
source:          ARIN-WHOIS
```



# Filter generation time schedule

01:00 UTC - NTT ACL generator and ACL pre-generator are run and ingest from IRRd (rr.ntt.net)

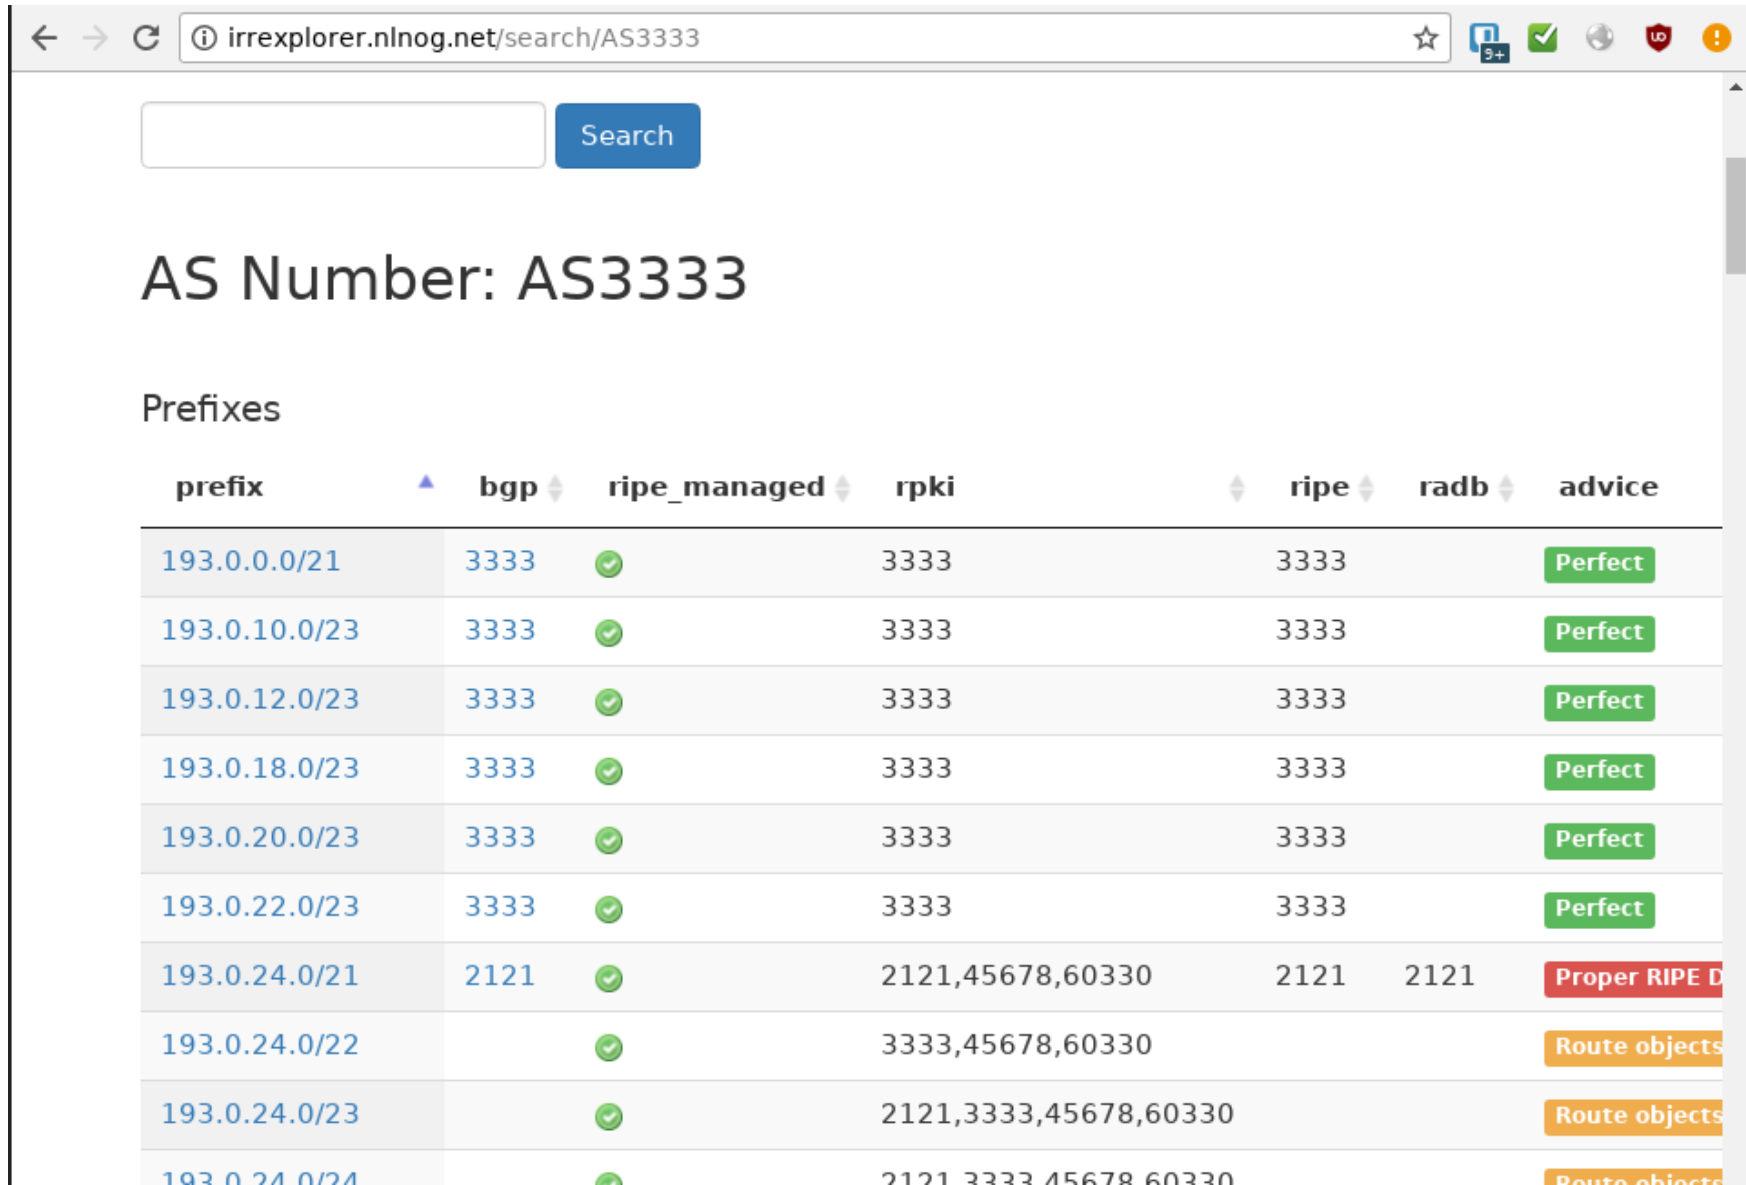
04:00 UTC - the newly generated ACLs are loaded into NTT AS 2914 border routers

06:00 UTC - ARIN starts creation of the full ARIN WHOIS database dump

12:00 UTC - irrexplorer.nlnog.net fetches & processes ARIN WHOIS database dump from ARIN

12:30 UTC - ARIN WHOIS data is loaded into IRRd at rr.ntt.net / rr1.ntt.net

# http://irrexplorer.nlnog.net



The screenshot shows a web browser window with the URL `http://irrexplorer.nlnog.net/search/AS3333`. The page features a search bar and a 'Search' button. Below the search bar, the text 'AS Number: AS3333' is displayed. Under the heading 'Prefixes', there is a table with columns: 'prefix', 'bgp', 'ripe\_managed', 'rpki', 'ripe', 'radb', and 'advice'. The table lists several IP prefixes and their associated AS numbers and RPKI status. The first six rows show prefixes under AS3333 with 'Perfect' RPKI status. The last three rows show prefixes under AS2121 with 'Proper RIPE D' status.

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



# New use case for RPKI data?

- 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

# New 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**, **NCC Pilot (RRDP prefetch)**, **RIPE NCC RPKI Root**, **RIPE NCC RPKI Root (RRDP prefetch)**, **RIPE NCC prepdev (RRDP prefetch)**, **altca**, **apnic-testbed**.

Show 10 entries

Search: 2

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

# Simple 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 ~$
```

# Using RPKI data: rtrsub

The screenshot shows the GitHub repository page for `job/rtrsub`. The browser address bar displays `https://github.com/job/rtrsub`. The repository name `job / rtrsub` is shown with a dropdown menu. Navigation tabs include `Code` (selected), `Issues 1`, `Pull requests 0`, `Boards`, `Reports`, and `Projects 0`. A button `Unwatch` is visible. The repository description is `RTR Substitution` with a link `Add topics`. Statistics show `63 commits`, `1 branch`, `22 releases`, and `1 contributor`. Action buttons include `Branch: master`, `New pull request`, `Create new file`, and `Upload files`. A list of files and their descriptions is shown:

File	Description
<code>rtrsub</code>	bump version
<code>template-examples</code>	make better use of the native bird rpki functions
<code>.gitignore</code>	improve bird template
<code>LICENSE</code>	Add LICENSE





# Future work #1

- Currently working with NIC.MX and NIC.BR for WHOIS access
- Continue to search for new quality data sources

# Future work #2

- Going to IETF, to define “AS-SETs in RPKI”
  - Design goal: avoid collisions & autodiscovery
- Write a new `IRRD` from scratch: **IRRDv4**
  - Have better admin capabilities
  - Allow for innovation, integration with the RIRs

# NLNOG Camp - June 1-3 Netherlands

