Golden Prefixes
IRR Lockdown

Job Snijders <job@ntt.net>

Agenda

- What's the problem?
- IRR not ideal
- A possible solution: "Golden prefixes"
- Making the best of IRR: "IRR Lockdown"

Actual Frustrations

- The Youtube Hijack (oops! classic!)

  in 2008, AS17577 announces 208.65.153.0/24 -> end game is entire pakistan offline

- Route leaking through the OSPF/ISIS rabbit hole

  - Originating a full table with your own ASN: AS: HOPPA GANGNAM STYLE

- “BGP optimisers” route leaking

  - NO_EXPORT doesn’t always work (CSCum76994)

Crash-course IRR

```
route:          37.77.56.0/21
descr:          S.J.M. Steffann
origin:         AS57771
mnt-by:         STEFFANN-MNT
source:         RIPE
```

- Upload snippets of text to a database

- Clients query that database

```
hanna:d job$ bgpq3 -A AS-SNIJDERS
no ip prefix-list NN
ip prefix-list NN permit 165.254.255.
ip prefix-list NN permit 194.33.96.0/
```

What's wrong with IRR?

- Every breathing idiot can create any route object in RADB, ARIN, etc
- No guarantees that the "owner" of the space authorised that route object
  - Exception: RIPE, APNIC, AFRINIC…
- Lots and lots of stale data, even my study room is cleaner

RPKI issues

- Legal issues with obtaining root anchors
- Tooling is immature
- Local policy knobs limited
- Adds a new protocol in your network (RTR)
- Still risk of stale data

Possible solution?

Golden Prefixes

Golden prefixes

- SSL-pinning for BGP Prefixes
- Central repository
- Simple format:

```
Vurt:goldenprefixes job$ cat AS8283/list
2a02:898::/32
94.142.240.0/21
185.52.224.0/22
194.1.163.0/24
195.114.12.0/24

Vurt:goldenprefixes job$ grep 8283 auth
8283 C57E21E27E5BEC10
Vurt:goldenprefixes job$
```

Some useful configuration: youtube

```
prefix-set AS43515
 64.15.112.0/20,
 208.65.152.0/22,
 208.117.224.0/19,
 208.117.236.0/24,
<snip>
 208.117.251.0/24,
 208.117.254.0/24,
 208.117.255.0/24,
 216.239.60.0/24
end-set
!
```

```
route-policy golden-prefix-list
if destination in AS43515 and as-path originates-from '43515' then pass exit
if destination in AS43515 then drop exit
if destination in AS8283 and as-path originates-from '8283' then pass exit
if destination in AS8283 then drop exit
```

Applicable to all BGP sessions!

Advantages

- Legal could be more friendly (MIT or Apache license?)
- Proven technology:
    - route-maps & prefix-lists have been in use for more then a decade
- Transparency
    - All communication surrounding GP is publicly accessibly
    - Full logs for accounting are in git
- Local decision which ASNs are of interest
- No stale data

Participation process

1. Two introducers required
2. Exchange of PGP material with the “Auditor”
3. Auditor verifies the following:
    1. No duplicates? No overlap with existing prefixes?
    2. Has the route been stable for the last two months?
    3. Were procedures followed properly?
4. ??

Data consumption

1. Obtain a copy of “goldenprefixes” repository
2. Run the validator tools to verify integrity
3. Generate network config with the tools (run from crontab)
4. Network config is based on templates and settings:
    - Ignore AS 65503
    - Use these suffixes/prefixes on prefix-lists
5. Push to network device

(uiteraard in crontab of jenkins, elke 12 of 24 uur)

Now what?

There has been interest from various ISPs (large and small), so The Todo

- Gather community interest
- Develop strong policies / procedures
- Write some software
- Get it rolling with a few data producers & consumers

The NLNOG Foundation could take a leading role

What is an IRR Lockdown?

- Only honor route objects which come from an IRR source which properly authenticates

- Discard route objects for parts of the DFZ which come from the “locked down” IRR

The plan

Knowing that: RIPE administrates roughly 35 /8 blocks

NTT will **only** allow route objects covering RIPE administrated space to influence NTT prefix filters if they have passed RIPE authentication chain, resulting in:

- Ignore untrusted updates on NRTM streams
- Reject route object creation in NTTCOM registry for RIPE space

Examples: Untrusted NRTM updates

- Anything that RADB sends to NTTCOM over NRTM which cover part of the 35 /8s RIPE administrates

- Anything with “source: RIPE” from non-RIPE NRTM server

- Any route objects customers create which covers RIPE administrated space inside NTTCOM registry

Statistics (28 nov 2014)

- Total number of RIPE prefixes for which a route object ONLY exists in a foreign IRR AND which were observed in the DFZ: **1004 prefixes** (aggregated 522), spread over 280 ASNs.

- Total number of prefixes for which a route object exists in both RIPE IRR and a foreign IRR (with mismatching origins), AND where the foreign version is observered in the DFZ: **269 prefixes** spread over 119 ASNs.

Details: https://www.ripe.net/ripe/mail/archives/routing-wg/2014-November/002887.html

```
inetnum:        193.0.0.0 - 195.255.255.255
netname:        EU-ZZ-193-194-195
descr:          European Regional Registry
```

Good:

```
route:   193.0.0.0/21
descr:   RIPE-NCC
origin:  AS3333
mnt-by:  RIPE-NCC-MNT
source:  RIPE
```

BAD!

```
route:   193.0.0.0/21
descr:   RIPE-NCC
origin:  AS666
mnt-by:  RIPE-NCC-MNT
source:  RADB
```

Why would we ever honor the bad route object?!

Q & A for routing police?