

RPKI deployment experience at IIJ

Matsuzaki 'maz' Yoshinbu
<maz@iij.ad.jp>

IIJ Overview

- Internet Initiative Japan Inc.
 - A pioneering and techie ISP established in 1992
 - Eating our own dog food
 - IPv6, DNSSEC, and RPKI ☺
- Focus on Enterprise Market
 - Huge enterprises, as well as government and academic customers
 - Some consumer services such as broadband and mobile
- Provides transit to other ISPs
 - <https://www.peeringdb.com/asn/2497>

Timeline

- 2005 Some RPKI workshops
- 2019 Our engineers decided to give it a try
- 2020/MAR ROV: Trials in the test environment
ROV: Trials in the production network
- 2020/JUL The engineers found deployment feasible
- 2020/OCT ROA: Trials with some prefixes
ROV: Configure RPKI cache on routers (RTR)
- 2020/NOV ROA: Starting registration for our PA
ROV: Starting rejection of invalid routes


Route Origin Authorization (ROA)

- An object states which AS is authorized to originate a particular IP prefix
- Major Components
 - IP prefix
 - Max length
 - AS number
- Creating ROAs at RPKI CA
 - Hosted RPKI (RIR/NIR system) or Delegated RPKI (Your Own)

Responsible for Issuing ROA NOC or LIR

- NIR/RIR account is required for ROA creation
- Our LIR team manages RIR/NIR accounts
 - Those accounts are authorized to manage IP addresses
 - Creating ROAs
 - And can also return IP addresses
- Our NOC should be able to create ROAs as they are related to routing
- → NIR/RIR accounts were given to our NOC team
 - And the NOC team carefully manages our ROAs.

ROA Life Cycle at IIJ

- 
- AS0 ROA
 - A prefix is not in use (just after transfer, and etc.)
 - No BGP origination from IIJ/AS2497
 - LIR team create a AS0 ROA for the prefix
 - AS2497 ROA
 - The prefix is starting to be used/announced
 - NOC team deletes the AS0 ROA, and creates AS2497 ROA



If necessary

Our Basic ROA Policy

- Create ROAs for all our PAs
 - ROA maxlength is the same as the announcing prefix length
 - RFC7115 suggests this
 - IIJ does not deaggregate
- Punching Holes
 - Announcing part of a PA block from another AS (e.g., customer)
 - Create a ROA corresponding to the customer AS
 - ROA maxlength to be discussed with the customer

Various IP prefix cases

- Historical resources
 - RIR/NIR coordination was not good on some occasions
 - It's improving though
- A prefix re-allocated entirely to a customer AS
 - Originating from their AS
 - Ask the customer about their intention to create a ROA
- Punching Holes
 - May be used by some customers for DoS protections, etc.
 - Creating an exact ROA is essential
 - Ask customers for information needed to create a ROA

IIJ/AS2497 ROA coverage is .. still under 40%

- Why?
 - → Many customer PI blocks
 - Customer holds PI block and requests IIJ/AS2497 to announce it
- Customers need to create a ROA by their own
 - Need to explain its necessity and risks
 - Need to get them access to the RIR/NIR RPKI system
- Continuous efforts are needed

Route Origin Validation (ROV)

- Route filters based on ROAs
 - Can apply policy according to match status to ROA
 - Reject, set BGP attribute or just monitoring
- Relying Party (RPKI Cache)
 - Collect and verify ROA
 - Send Validated ROA Payloads (VRP) to routers by RTR protocol
 - Router verifies incoming routes based on VRP

Our Basic ROV Policy

- Deploy on our eBGP routers
 - For peers and upstreams
 - For customers is TBD
 - Currently, a strict prefix and as-path route filter is applied
- Drop ROV invalids
 - Treat valid, unknown, and unverified as equivalent

Relying Party (RPKI Cache)

- Each router is connected to two servers
 - Servers deployed at different POPs
 - Servers with different software implementations
- One server serving about 20 routers
- Each server fetches ROA
 - It would be much friendlier to have only the representative servers perform the fetch, but we have not yet been able to implement that much

Convincing Sales and Support team

- They understand the purpose, but
- Concerned about impact on reachability
 - Destinations that cannot be reachable due to R0V
- About 3000 Invalid prefixes (as of 2020/SEP)
 - 0.3% of full routes

Estimate Impact of Dropping Invalids

- In most cases, there are covering prefixes
 - Guessing that sub-prefixes for traffic control
- Excluding the above, 0.097% of full routes is likely to be completely unreachable by dropping Invalids
 - According to our NetFlow data, there is almost no traffic to any of these destinations
 - May be some Invalids for research purposes
- These factors convinced the team that dropping Invalid would not affect our customers

Requests from Our Support Team

- Tools to find out if the ROV is affected
 - Looking Glass
 - RPKI web UI
 - Dump of Invalids
- Reduction of invalid routes
 - Raise awareness of generating a proper ROA

Customer Announcements

- Decided not to make an announcement
 - Cannot prove absolutely no impact
 - Our AS operational policy
 - Invalid, so it deserves to be dropped
 - Instead, present our efforts through various community
 - JANOG and etc.

ROV Deployment

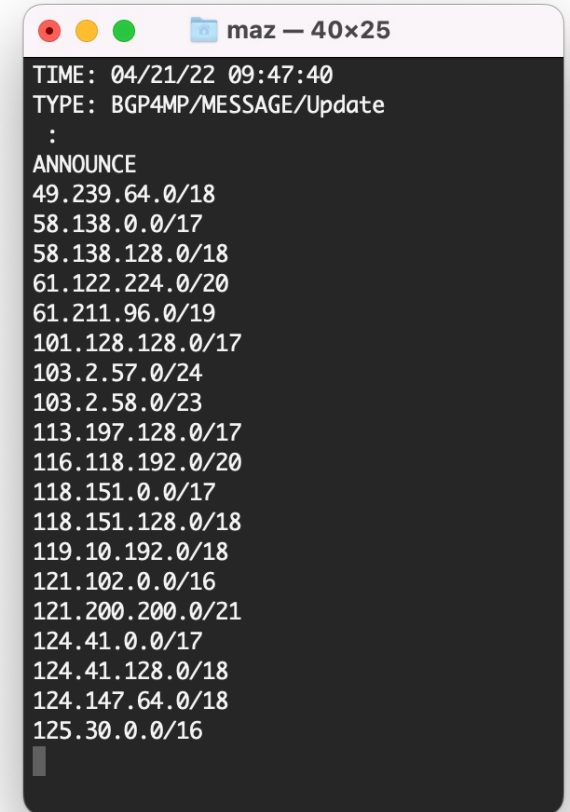
- Configuration applied to about 2,000 BGP peers
- Staged Deployment
 1. Apply LOCAL_PREF 0 to Invalids
 2. Gradually changed to DROP
 - APAC/Europe->US->Upstream
 - Paths for Invalids are directed to a specific upstream link
 3. Drop all Invalids
- No complaints 😊

Monitoring ROV

- Relying Party (RPKI Cache)
 - Various metrics
 - Easy integration with various monitoring tools
 - The challenge is what to alert on
 - Failures to communicate with CA are common
- Routers
 - No MIBs ☹️
 - Hard to monitor
 - Record of Dropped prefixes
 - Periodically run show command on router

An incident happened

- Happened on April 21, 2022
- One European AS started to originate IIJ PA blocks
 - Observed only at the Amsterdam node of RIPE RIS
 - <https://data.ris.ripe.net/rrc00/2022.04/updates.20220421.0945.gz>
- The announcements stopped when contacted
 - Cause unknown though
- No customer impact
 - Thanks to the ROAs, maybe 😊



```
maz — 40x25
TIME: 04/21/22 09:47:40
TYPE: BGP4MP/MESSAGE/Update
:
ANNOUNCE
49.239.64.0/18
58.138.0.0/17
58.138.128.0/18
61.122.224.0/20
61.211.96.0/19
101.128.128.0/17
103.2.57.0/24
103.2.58.0/23
113.197.128.0/17
116.118.192.0/20
118.151.0.0/17
118.151.128.0/18
119.10.192.0/18
121.102.0.0/16
121.200.200.0/21
124.41.0.0/17
124.41.128.0/18
124.147.64.0/18
125.30.0.0/16
```

Summary

- It took roughly one year to deployment
- Monitoring still needs to be improved
- Happy to have created ROAs
 - We have implemented the means currently available to us
- Not to generate Invalids when creating ROAs
 - Sub optimal prefixes
 - Punching Holes