ROA+ROV Deployment & Industry Development

Che-Hoo Cheng

APNIC

@TWNIC IP OPM 32

2019-06-20

Security matters as your network is connecting to Internet

- You do NOT want your own routes to be hijacked by anyone, maliciously or accidentally
- You also do NOT want to receive bad routing information from any of your BGP neighbors or propagate bad routing information to any of them
- Basic measures include:
  - Bogons and martians filtering
  - Max prefix count
  - IRR (Internet Routing Registry) database checking
  - *So on and so forth*
- Additional measure should include:
  - **ROA (Route Origin Authorization) / ROV (Route Origin Validation)**

Routing security is becoming more important than ever

- Route-hijacking cases (malicious and accidental) are more and more common
  - Big incentive for hackers
    - Hijack DNS, hijack websites, steal passwords and so on
  - Misconfiguration does happen from time to time
- And, it is extremely easy to do route-hijacking, if protection measure is not implemented
- A lot of route objects on IRR-DB are not authenticated properly and so cannot be fully trusted
- Need better authenticity for routing info, i.e. need to make sure that the route originators are the true “owners” of the relevant IP resources

Fat-finger / hijacks

- **Quad101 related route (101.101.101/24) was hijacked by AS268869 (FIBRA PLUS) for 3.5 mins on 8 May 2019**
  - https://blog.apnic.net/2019/05/30/public-dns-in-taiwan-the-latest-victim-of-bgp-hijack/
  - Origin AS should be AS131621 (TWNIC)
  - Implication can be huge if anycast is not done well

Fat-finger / hijacks

- **Amazon (AS16509) Route53 hijack – Apr 2018**
  - AS10279 (eNET) announced/originated more specifics (/24s) of Amazon Route53's prefix (205.251.192.0/21)
    - **205.251.192.0/24 ……. 205.251.199.0/24**
    - https://ip-ranges.amazonaws.com/ip-ranges.json
  - The motive?
    - During the period, DNS servers in the hijacked range only responded to queries for myetherwallet.com
    - Responded with addresses associated with AS41995/AS48693

Fat-finger / hijacks

- **Bharti (AS9498) originates 103.0.0.0/10**
  - Dec 2017 (~ 2 days)
  - No noticeable damage done – more than 8K specific routes!
- **YouTube (AS36561) Incident**
  - Feb 2008 (down for ~ 2 hours)
  - PT (AS17557) announced 208.65.153.0/24 (208.65.152.0/22)
    - Propagated by AS3491 (PCCW)

RPKI

- **RPKI** is a public key infrastructure (PKI) framework, designed to secure BGP routing
  - Based on X.509 PKI standards
- **RPKI** adds Internet number resources (INR) information to X.509 certificates issued to resource holders
  - Representing “ownership” and other status
  - Certification hierarchy follows INR delegation hierarchy

IANA ➔ RIR (➔ NIR) ➔ ISP ➔ …

RPKI hierarchy

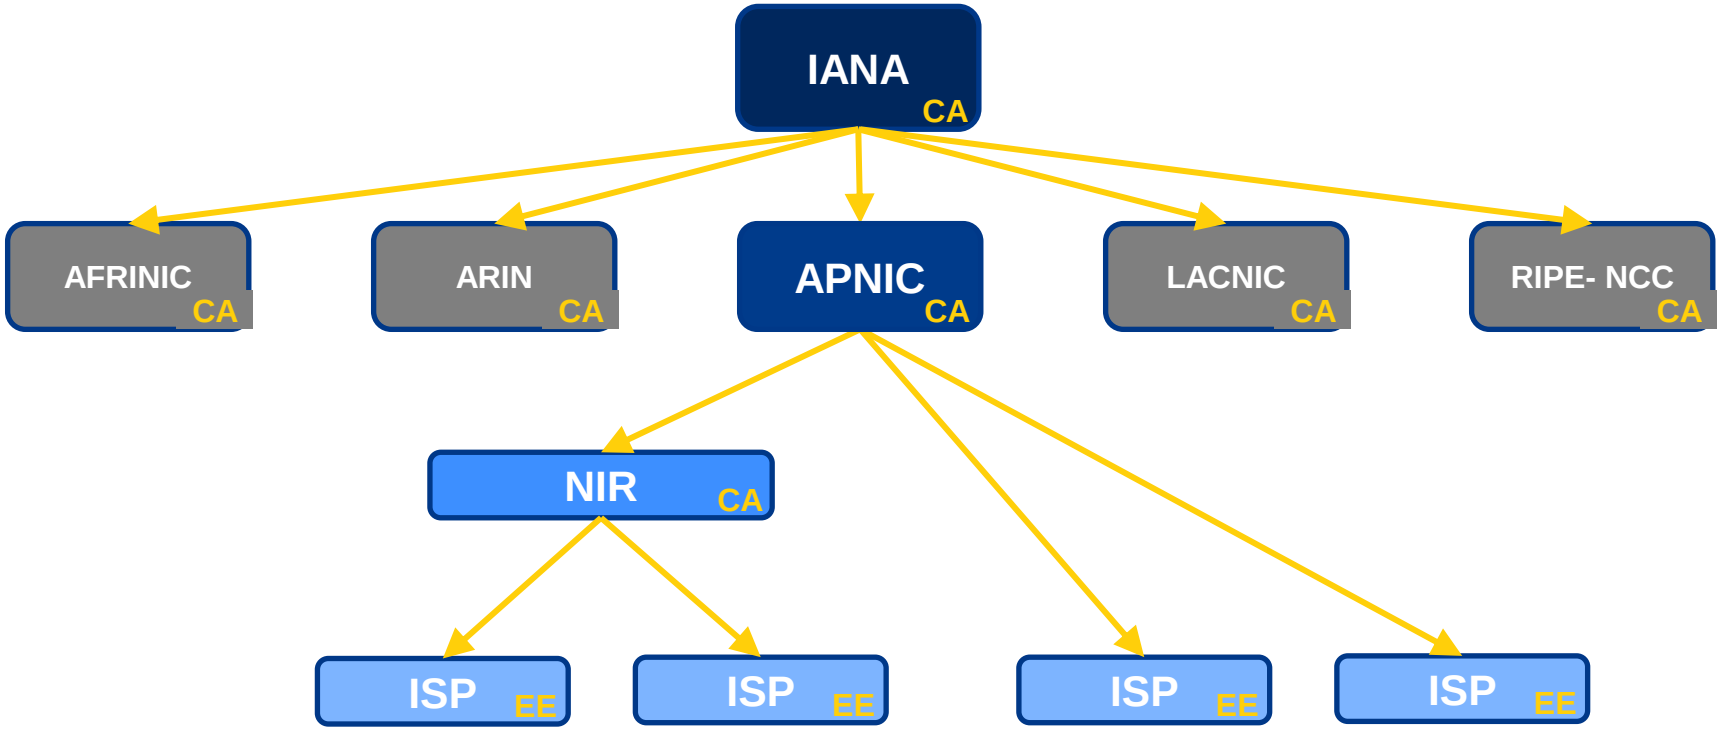

RPKI service models

- **Hosted model**
  - **APNIC performs CA functions on behalf of members**
  - **Manage keys, repository and so forth**
  - **Generate certificates for resource delegations**
  - **This “Member CA” is separate from the “APNIC CA”**
- Provisioning model
  - Member operates full RPKI system including CA
  - Communication with APNIC via “up-down” provisioning protocol
    - Either rsync (to be deprecated) or RRDP (preferred)
  - This is live at JPNIC, CNNIC and TWNIC (IDNIC in progress)

RPKI objects

- Resource certificates
  - Extension of standard X.509 certificates
  - Providing authority to use given IPv4/6 and ASN resources
  - Signed by issuing registry (serving as CA)
- Route Origin Authorization (ROA)
  - Giving an ASN authority to route specific IP blocks
  - Signed by IP resource holder
- “Anysign”, “ghostbuster” and more…
  - Other useful objects proposed and coming later

RPKI & applications

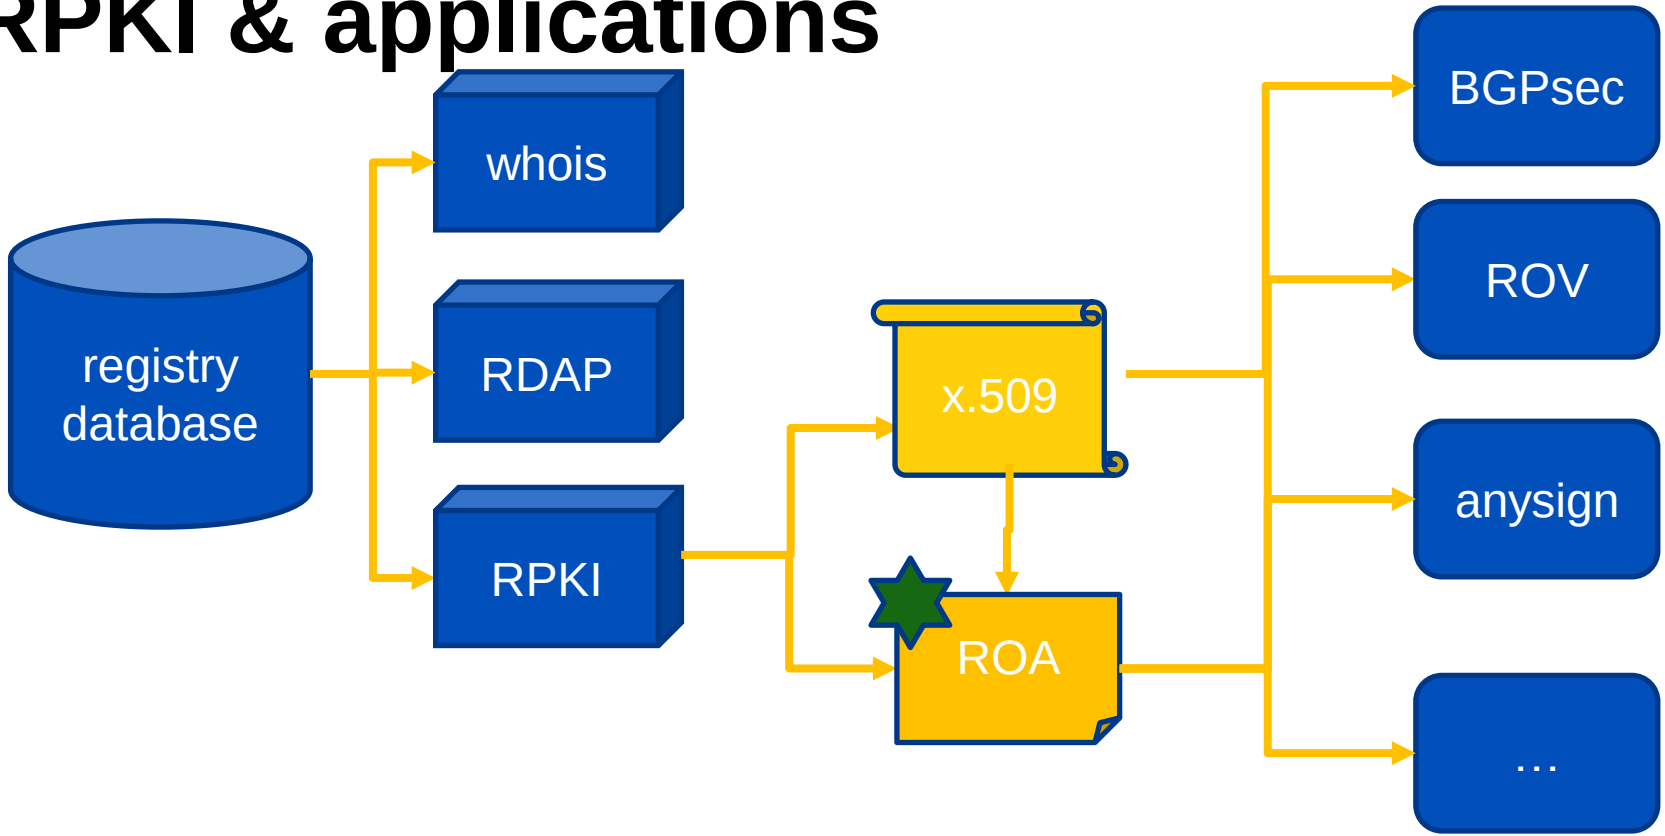

RPKI application: ROA

- Route Origin Authorization
  - List of prefixes with ASN authorized to announce
  - Signed by the prefix holder

Prefix	203.176.32.0/19
Max-length	/24
Origin ASN	**AS17821**

- RPKI validates the integrity of the ROA
  - It is provably created by the holder of the prefix
  - Can now be used to construct route filters for prefix-OriginAS pair in BGP
- Multiple ROAs can exist for the same prefix

Use of ROA: ROV

- Route Origin Validation

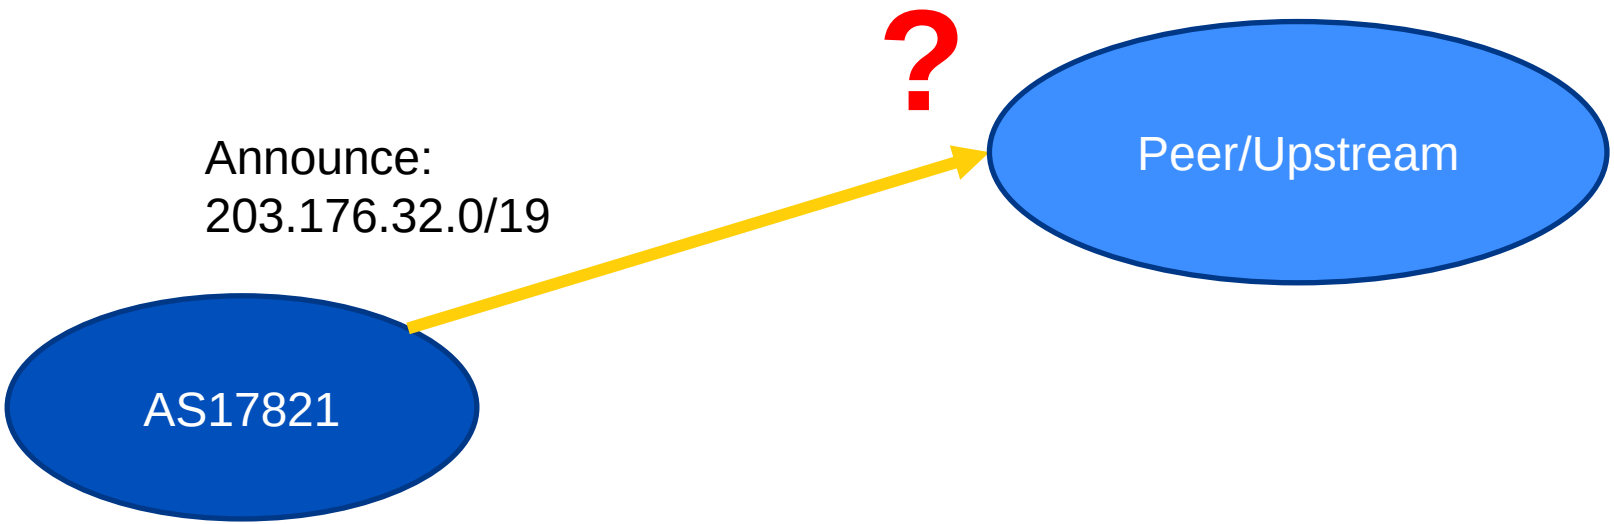

RPKI validator

- Gathers and validates ROAs from the distributed RPKI databases
  - Using rsync or RRDP (preferable)
  - Maintains a validated cache representing complete global state
- Can then perform ROV for routers using RPKI-Router (RTR) protocol

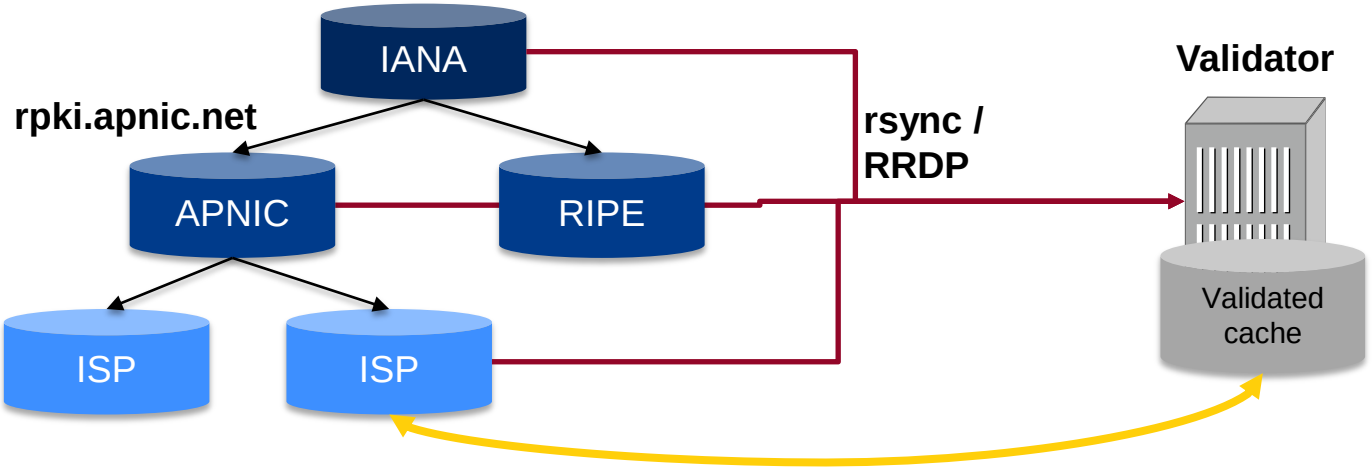

RPKI validator options

- Available validators
  - Dragon Research toolkit
    - https://github.com/dragonresearch/rpki.net
  - RIPE validator :
    - https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources
  - Routinator
    - https://github.com/NLnetLabs/routinator
  - RTRlib (bird, FRR, Quagga…)
    - https://rtrlib.realmv6.org/

Route validation states

- **Not Found (Unknown)**
  - No ROA found, probably not created yet
  - This will be “default” for some time.
- **Valid**
  - ROA exists
  - Prefix, Origin ASN and prefix-length match those found in validated cache
- **Invalid**
  - ROA exists
  - Prefix found, but Origin ASN is wrong, Prefix-length longer than Max-length, or certificates are expired or otherwise invalid.
  - Some action needed

ROV at border routers

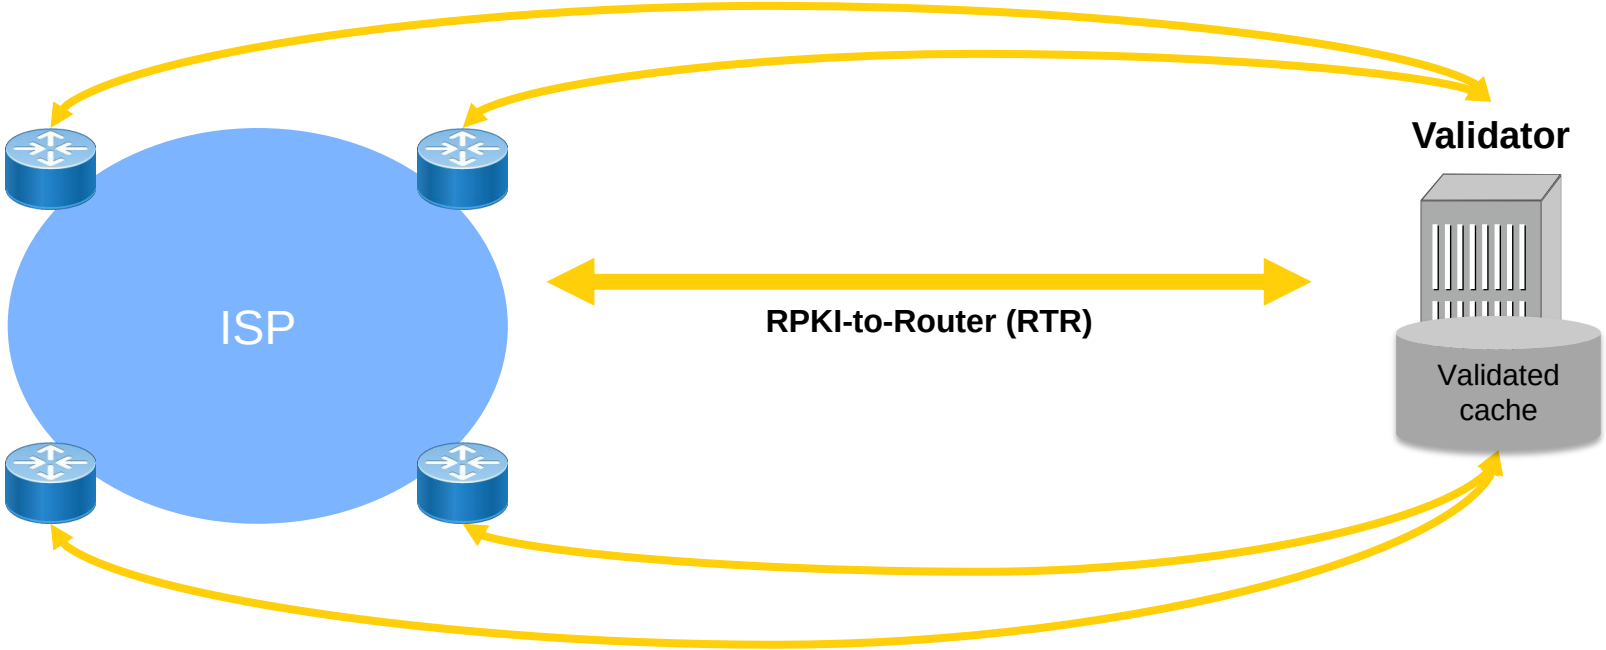

Options when receiving invalid routes

- For End/Stub Networks:
  - Drop them, OR
  - Give them lower LOCAL_PREF, OR
  - Do nothing (not recommended)
- For Transit Networks:
  - For inbound routes from upstreams / peers:
    - Give them lower LOCAL_PREF, OR
    - Drop them, OR
    - Do nothing (not recommended)
  - For outbound routes to customers:
    - Tag them before re-distributing them to customers and allow customers to make their own choices

ROV at IXP – RS and/or shared validator

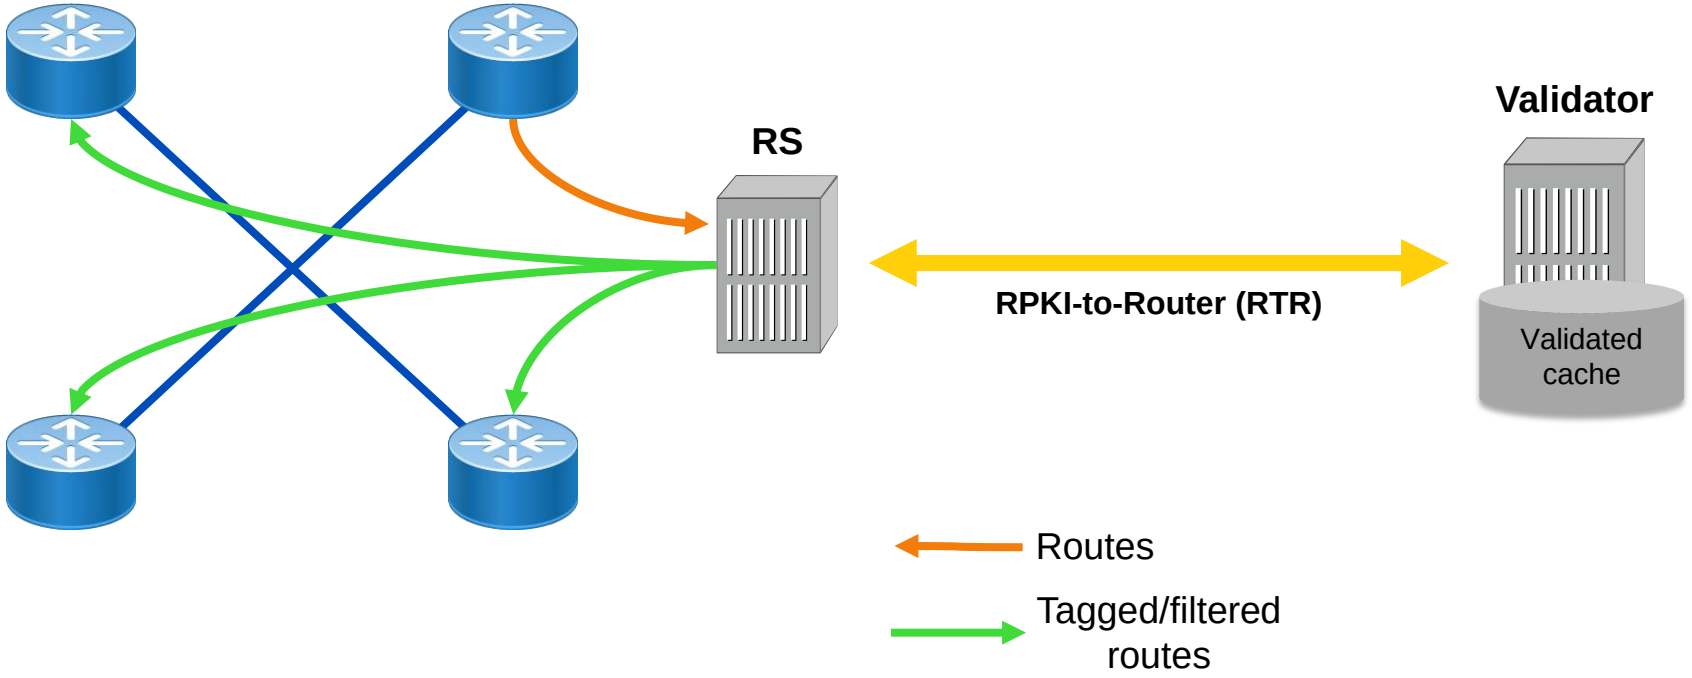

ROV at IXP – RS usage options

- Similar to the case of Transit Networks
- Lower LOCAL_PREF, OR
- Filtering
  - Do not advertise **Invalid** routes
  - Need to publish on RS policy
- Tagging
  - Apply community tags based on the validation state
    - let individual member ASNs act on the validation states
  - Example:
    - **Valid** (*ASN:65XX1*)
    - **Not Found** (*ASN:65XX2*)
    - **Invalid** (*ASN:65XX3*)

ROV at IXP – Examples in Asia Pacific

- Shared validator provided by:
  - JPNAP & BKNIX
- Other IXPs?
  - IXPs are good locations to place shared validator as they are just one hop away from their participants and they are mostly trustable
  - You may push your IXPs to support it to ease your burden of setting up your own Validator/Cache
  - IXP Manager SW (https://www.ixpmanager.org) now supports easy ROV deployment on RS at IXPs

ROA+ROV – Why do we do it?

- Contribute to Global Routing Security
  - Help reduce the effect of route hijacking or misconfiguration
  - Protect your own networks and your customers better
- Collaborative effort among network operators is key

ROA+ROV is NOT a bullet-proof solution

- It is just one small step for improving routing security
- But it helps improve the situation for routing security, especially if everybody does it
- Coupled with more and more direct peering, the protection for routing security should be more effective
- Highly recommend doing full MANRS as well

ROA stats – global snapshot

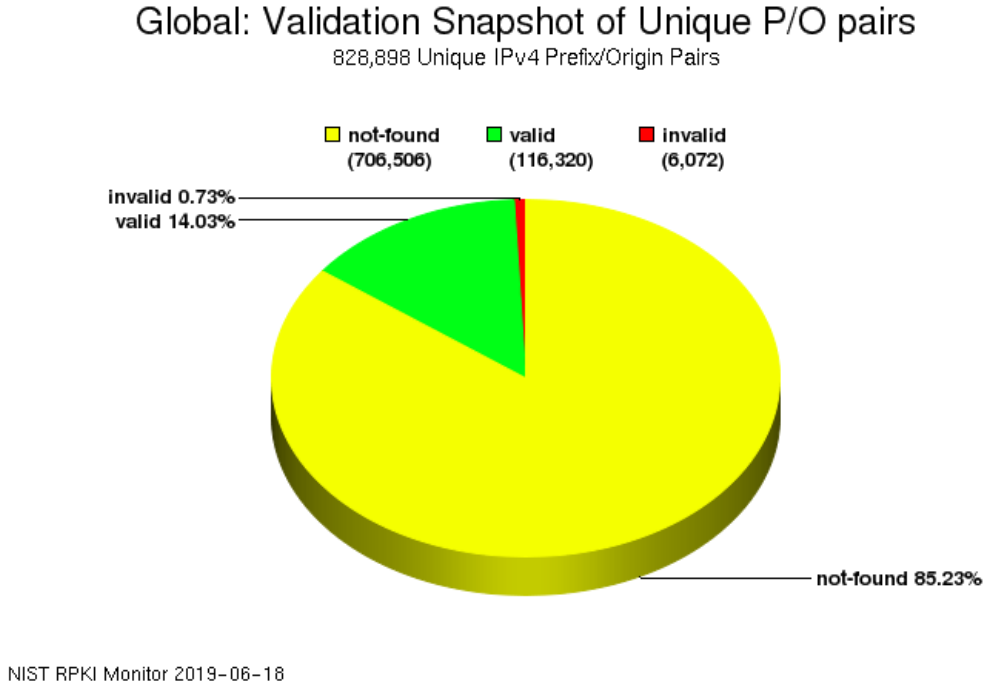

- Source: https://rpki-monitor.antd.nist.gov/?p=0&s=0

ROA stats – global trend

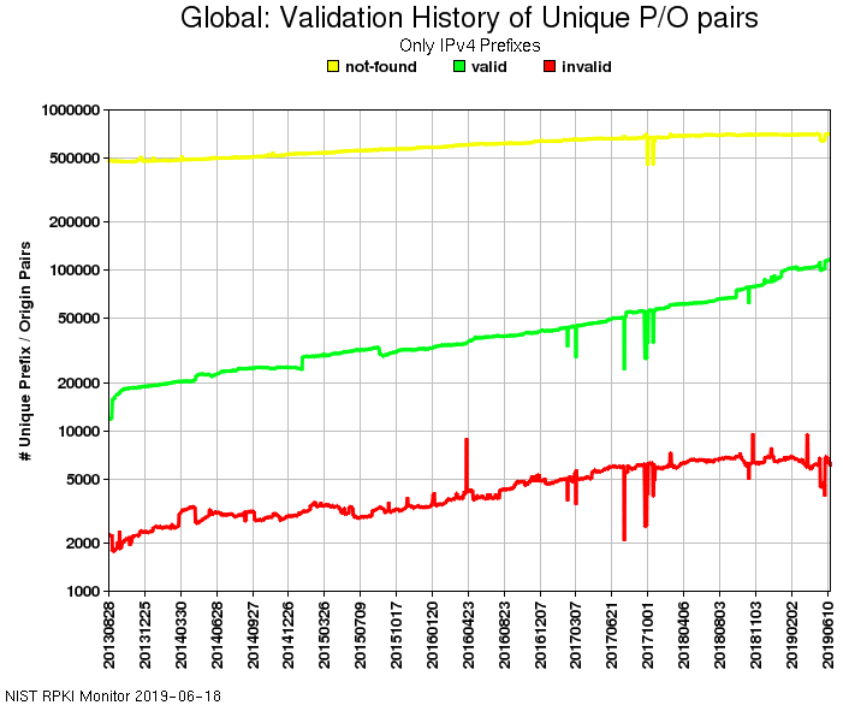

- Source: https://rpki-monitor.antd.nist.gov/?p=0&s=0

ROA stats – APNIC region snapshot

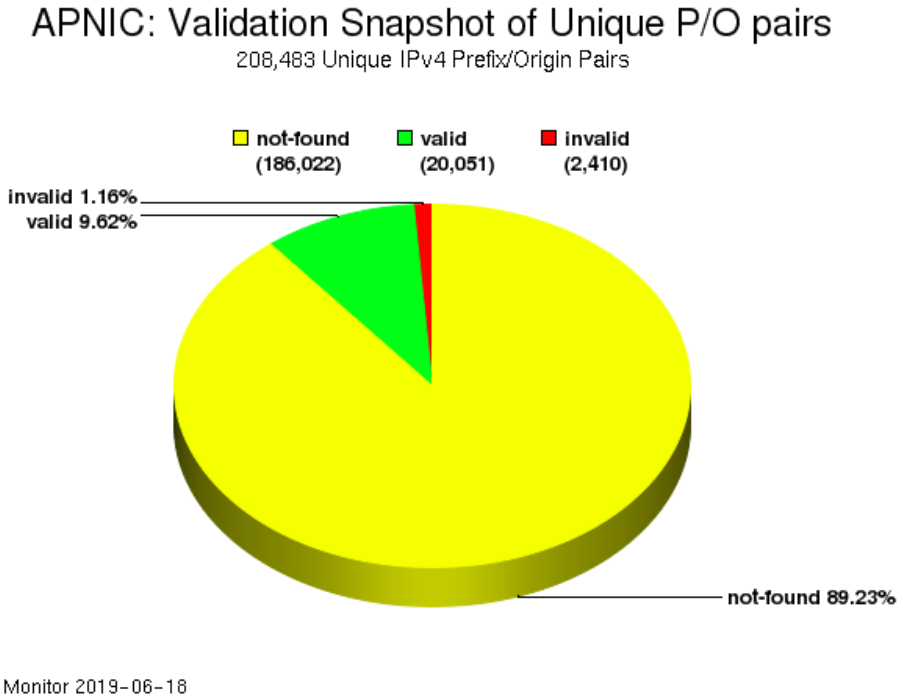

- Source: https://rpki-monitor.antd.nist.gov/?p=3&s=0

ROA stats – APNIC region trend

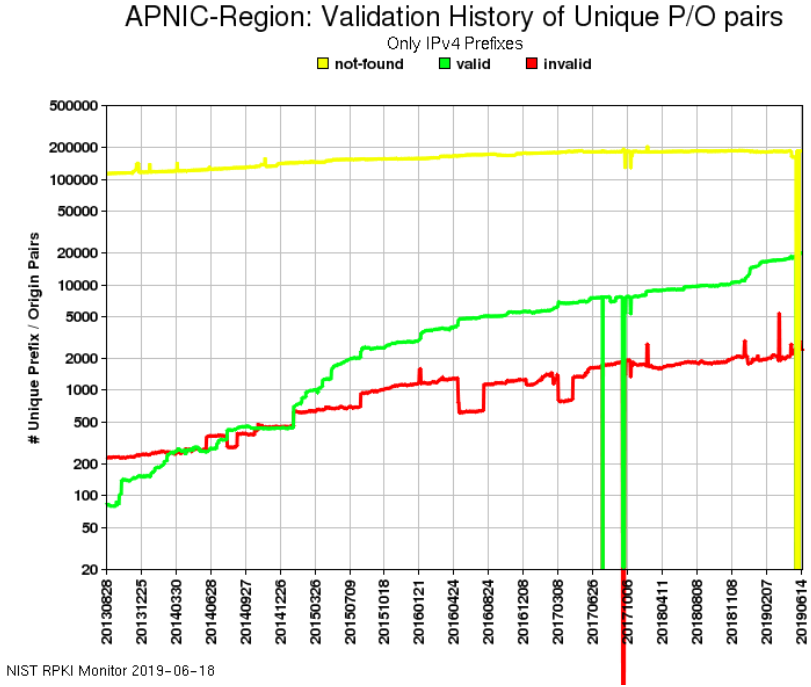

- Source: https://rpki-monitor.antd.nist.gov/?p=3&s=0

ROAs in APNIC region

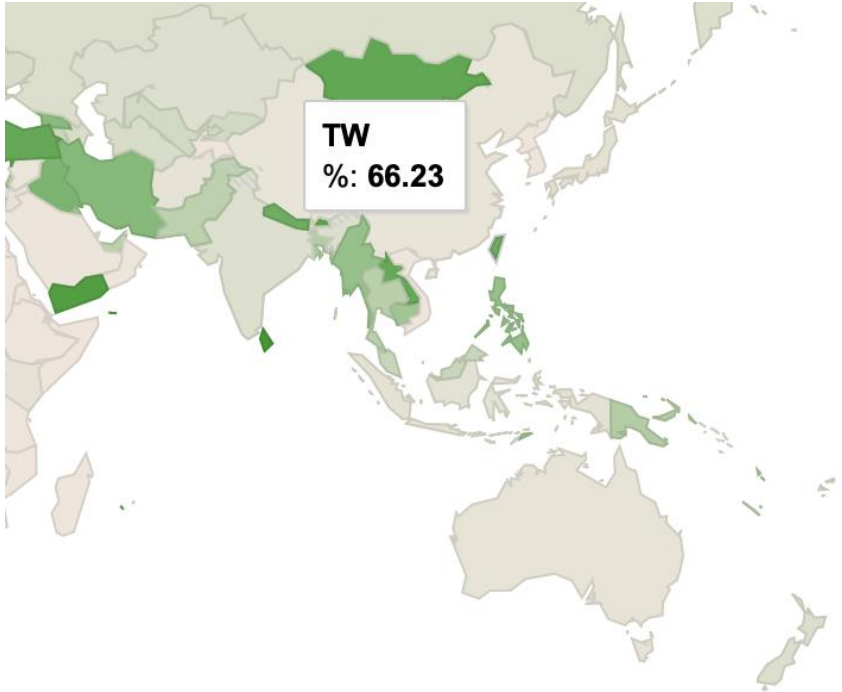

- Source (IPv4 prefixes covered): https://lirportal.ripe.net/certification/content/static/statistics/world-roas.html

ROA+ROV deployment steps

- **Create your own ROAs at relevant registries to better protect your own routes**
  - **And encourage your peers/customers to do the same**
  - **For APNIC members, it is easy to do it on MyAPNIC**
    - **We can help!**
    - **Please contact APNIC Helpdesk**
- Next step is to set up route origin validation (ROV) at your border routers
  - Using public or IXP validator, or your own
    - Do care about backup and resilience
  - And ask your IXP/upstream providers to implement ROV
    - Simple click to test your provider: https://ripe.net/s/rpki-test

Industry development on ROA+ROV

- NTT – IRR improvement favoring Route Objects with valid ROAs
- Cloudflare – Public validator service & invalid routes filtering
- AT&T – Invalid routes filtering on peering connections
- Netnod – Invalid routes filtering and favouring of valid routes on IXP Route Servers
- AWS – BYOIP requires customers to set up ROAs
- Google – Will start to apply stricter filters to BGP announcements on all peering sessions by Sep 2019, using RPKI data (ROAs) where available to validate IRR data
- Big players are getting more and more serious with ROV…

Announcing invalid routes

- **Will get to fewer and fewer networks on Internet**
  - **Similar to being disconnected from bigger and bigger part of Internet**
- If it is just a mistake, updating the relevant ROA records (supposedly with proper authority) will solve the problem
  - Should always keep your ROA records updated
    - All can be managed at one place so should be easy
  - Can have ROA records for the same prefix under multiple Origin ASes at one time to help the cases of network migration and so on

Incentives for creating ROAs

- To have basic protection of your own routes from being hijacked at those networks which do ROV
- Industry push:
  - NTT – IRR improvement **favouring Route Objects with valid ROAs**
  - Netnod – Invalid routes filtering and **favouring of valid routes on IXP Route Servers**
  - AWS – **BYOIP requires customers to set up ROAs**
  - Google – Will start to apply stricter filters to BGP announcements on all peering sessions by Sep 2019, **using RPKI data (ROAs) where available to validate IRR data**
- More will be coming…
  - As a requirement for peering???

More about RPKI benefits

- Improved in-band verification of resource custodianship
  - Much safer than manually checking whois or IRR database
  - Ease of automation
- Secure Origin is the first step to preventing many attacks on BGP integrity
  - BGP Path remains a problem which is under development
  - Related information such as IRR Policy can now leverage strong proofs of validity (end the maintainer-authority problem in RADB/IRR)
- Instruction/information from the resource custodian can be cryptographically verified (e.g. LOA signing)

Some useful references

- https://blog.cloudflare.com/rpki-details/
- https://nlnetlabs.nl/projects/rpki/faq/
- https://2019.apricot.net/assets/files/APKS756/apricot2019_snijders_routing_security_roadmap_1551228895%20(2).pdf
- https://datatracker.ietf.org/meeting/100/materials/slides-100-sidrops-rpki-deployment-with-ixps-01
- https://datatracker.ietf.org/meeting/90/materials/slides-90-opsec-0
- https://www.ripe.net/support/training/ripe-ncc-educa/presentations/use-cases-stavros-konstantaras.pdf
- https://www.franceix.net/en/technical/france-ix-route-servers/

RPKI specifications

Some of over 42 RFCs on implementation of RPKI and BGPsec

- RFC3779 X. 509 Extensions for IP Addresses and AS Identifiers
- RFC6480 Infrastructure to support secure routing
- RFC6481 Profile for repository structure
- RFC6482 Profile for Route Origin Authorisation (ROA)
- RFC6483 Validation model
- RFC6484 Certificate Policy (CP) for RPKI
- RFC6485 Algorithms & Key sizes for RPKI
- RFC6486 Manifests for repositories in RPKI
- RFC6487 Profile for RPKI Certificates
- RFC6488 Signed object CMS template
- RFC6489 Key Rollover
- RFC6490 Trust Anchor Locator (TAL)
- RFC6492 RPKI Provisioning Protocol
- RFC7318 Policy Qualifiers in RPKI certificates
- RFC7382 Certificate Practice Statement (CPS)
- RFC8181 RPKI publication protocol
- RFC8182 RPKI Delta protocol (RRDP)
- RFC8183 Out-of-band RPKI setup protocol
- RFC8360 RPKI Validation Reconsidered

2019 should be a big year for ROA+ROV deployment…

Questions?