RPKI Deployment
– It is time to start doing it

Che-Hoo Cheng

APNIC

@HKNOG 7.0

2019-03-01

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:
  - **RPKI (Resource Public Key Infrastructure) / 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

- **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

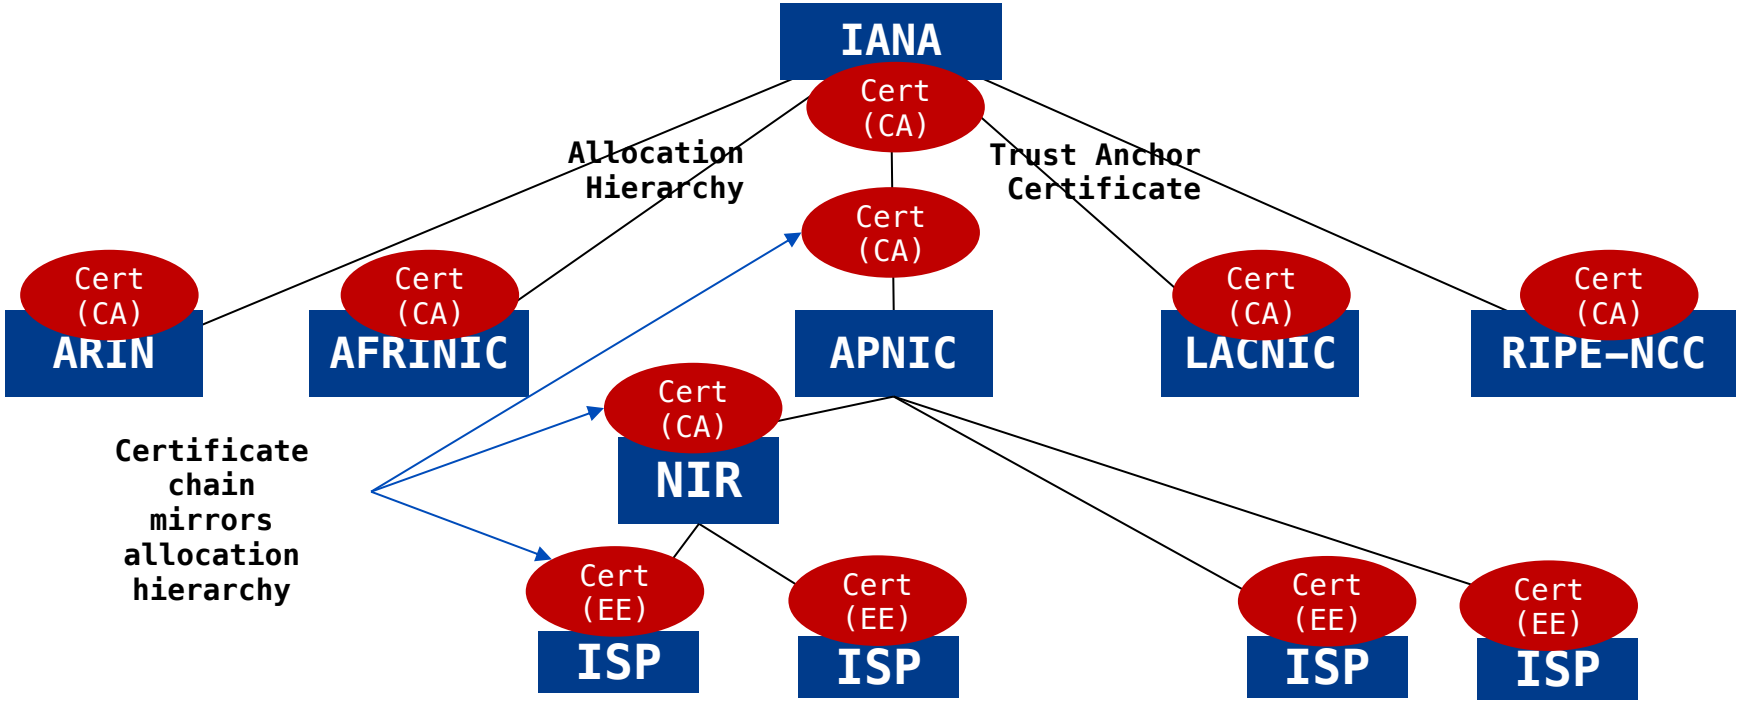

Source : http://isoc.org/wp/ietfjournal/?p=2438

RPKI Service Models

- **Hosted model**
  - **APNIC performs CA functions on behalf of members**
  - **Manage keys, repository etc**
  - **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 Authorisation (ROA)
  - Giving an ASN authority to route specific IP blocks
  - Signed by IP resource holder

RPKI – ROA

- Route Origin Authorization
  - Digitally signed object – list of prefixes and nominated ASN

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

  - Multiple ROAs can exist for the same prefix

RPKI application: ROV

- Route Origin Validation

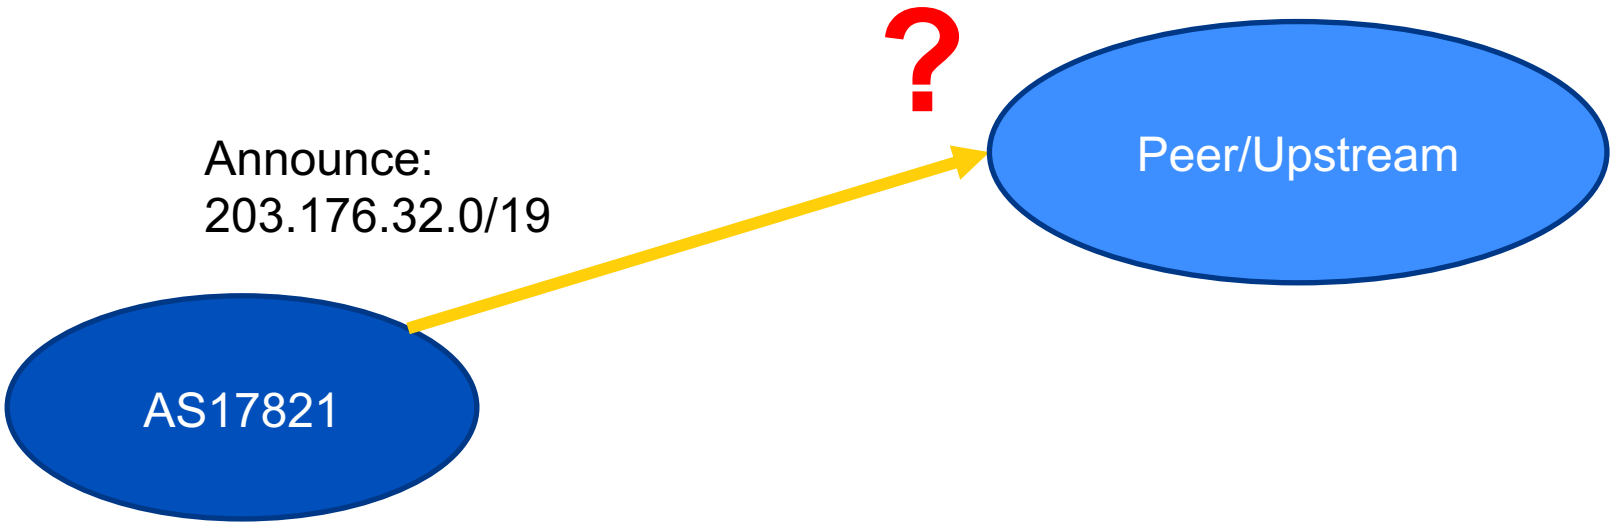

RPKI Validator

- Gathers ROAs from the distributed RPKI database
- Validates each entry's (ROA) signature
  - Creates a validated cache

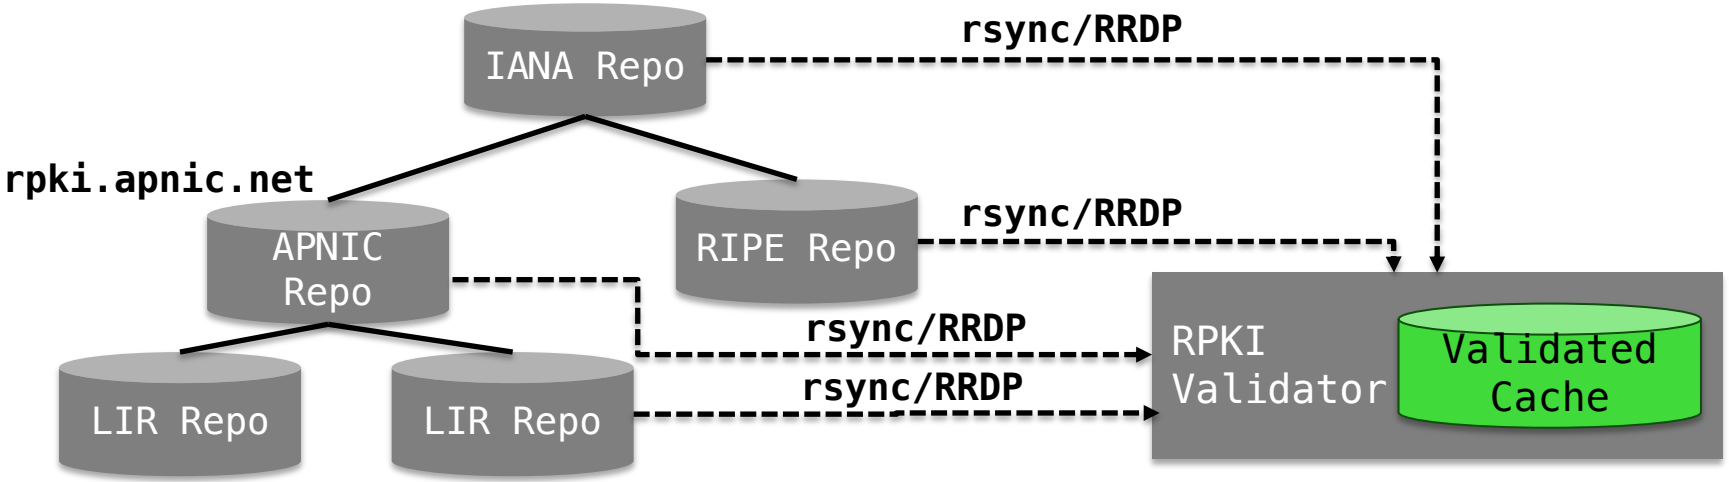

RPKI Validator Options

- 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/

RPKI Validation States

- **Valid**
  - Prefix, Origin ASN and prefix-length match those found on database
- **Not Found (Unknown)**
  - No valid ROA found
    - Neither valid nor invalid (perhaps ROA not created)
- **Invalid**
  - Prefix is found on database, but Origin ASN is wrong, OR
  - Prefix-length is longer than the Max-length

ROV@Border Routers

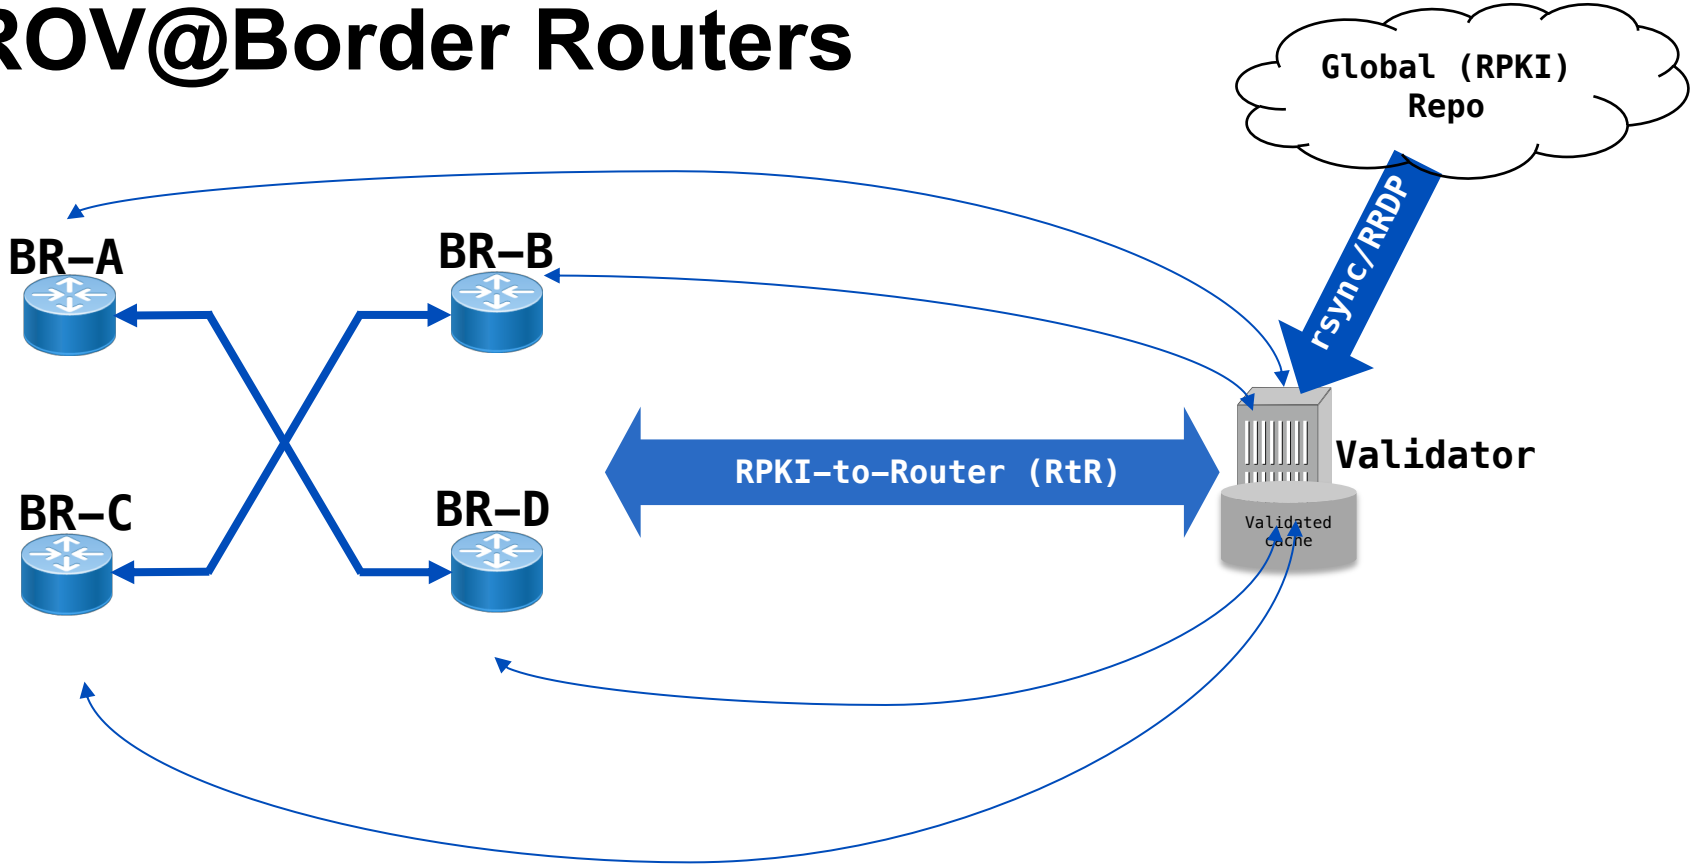

Options when seeing invalid routes

- 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@IXPs
– Route Server Scenario

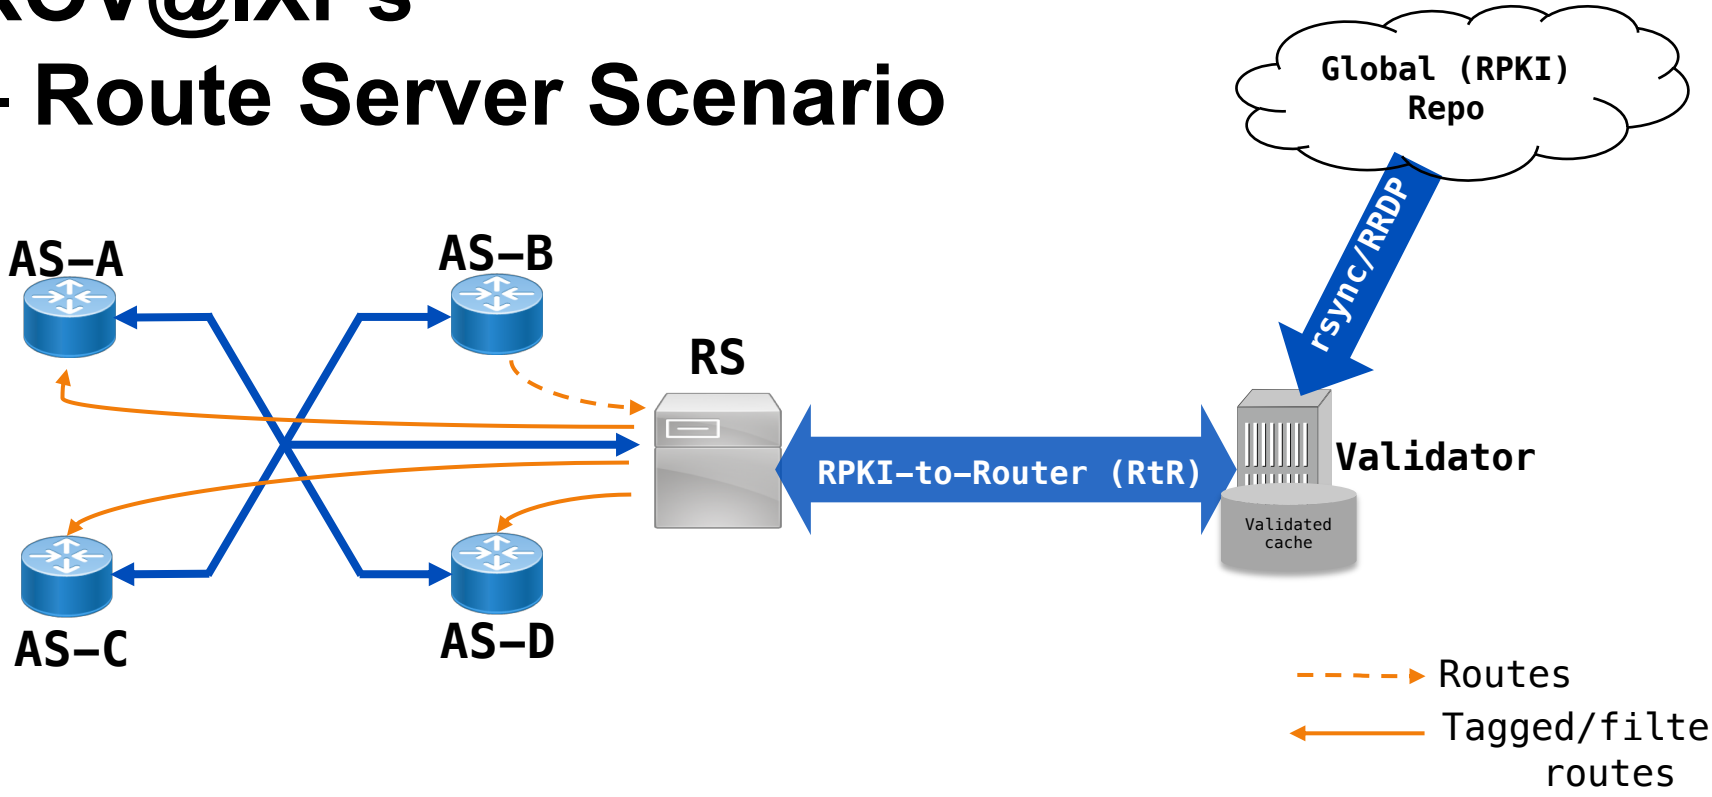

ROV@IXPs – RS Usage Options

- 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@IXPs – Shared Validator/Cache Scenario

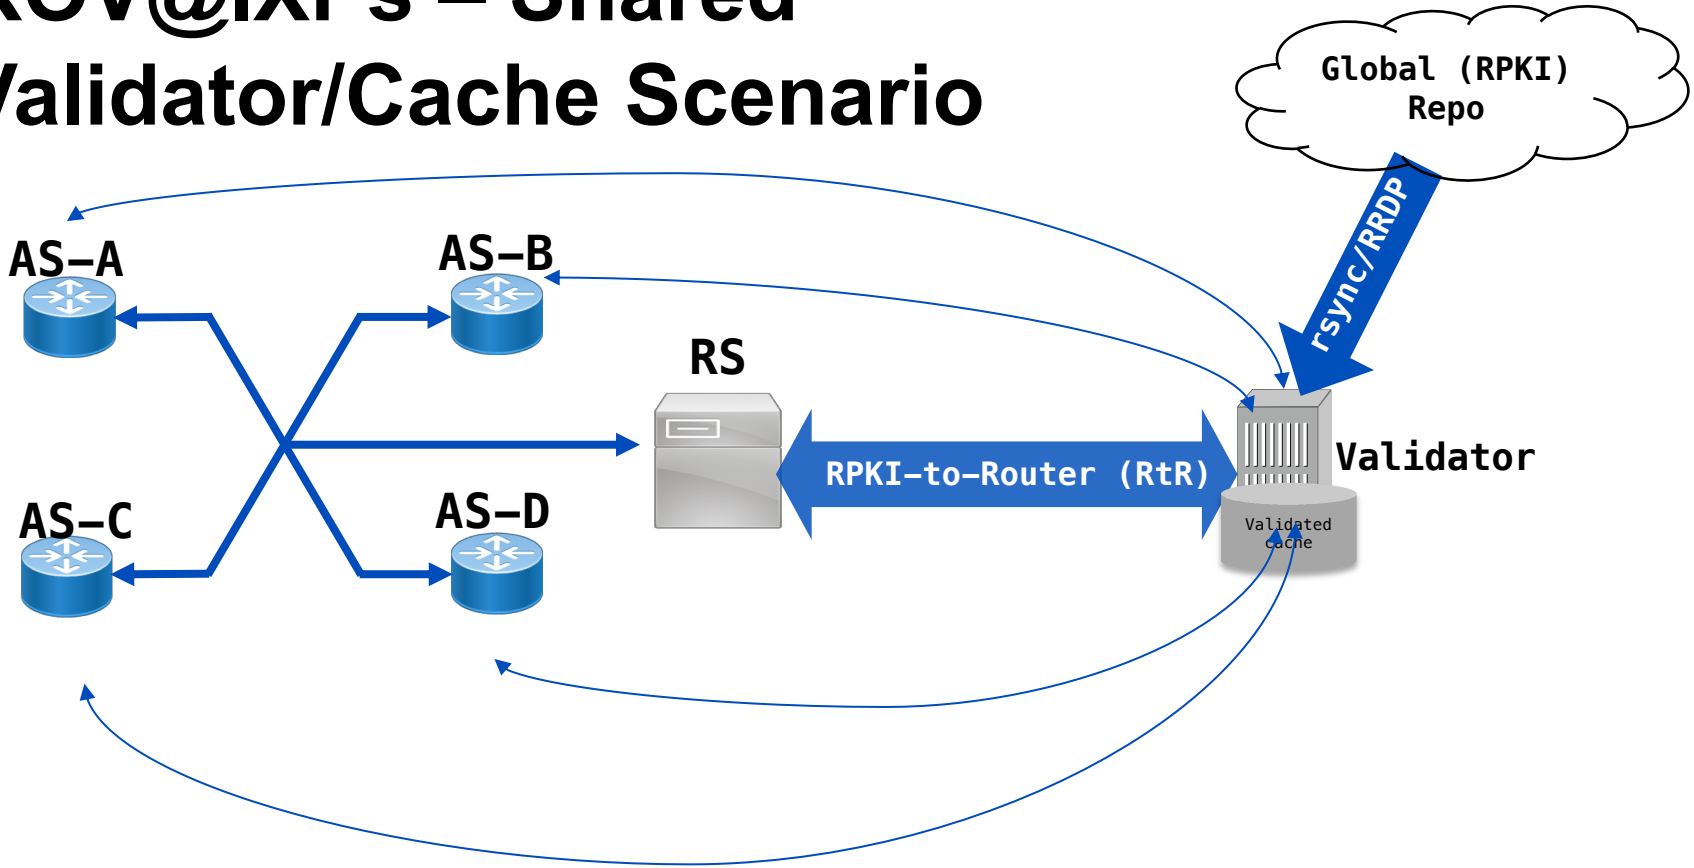

ROV@IXPs – Examples in Asia Pacific

- Shared Validator/Cache
  - JPNAP & BKNIX
- Other IXPs?
  - IXPs are good locations to place shared Validator/Cache 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 RPKI for easy deployment at IXPs

RPKI/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

RPKI is NOT a bullet-proof solution

- But it helps improve the situation, especially if everybody does it
- Coupled with more and more direct peering, the protection for routing security should be more effective

RPKI Status Globally – Snapshot

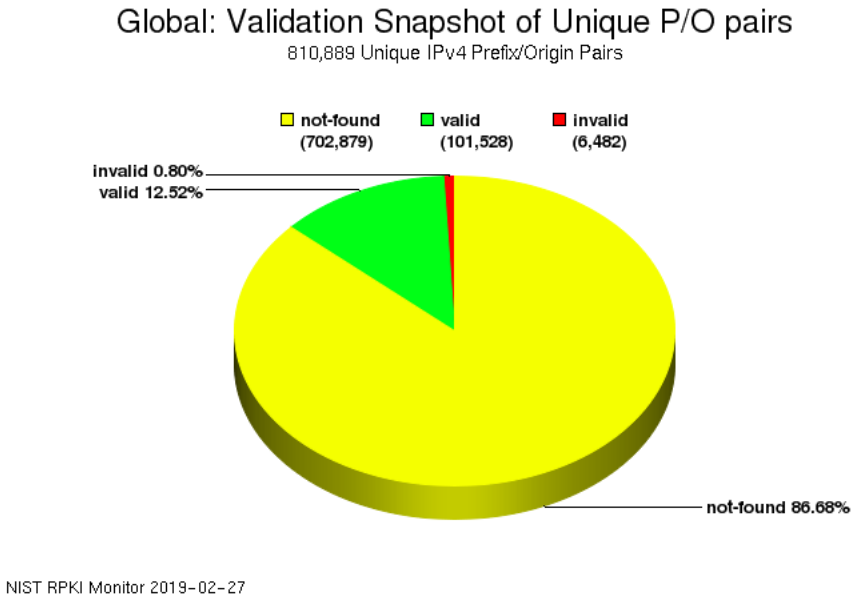

- Source: https://rpki-monitor.antd.nist.gov/?p=0&s=0

RPKI Status Globally – Trend

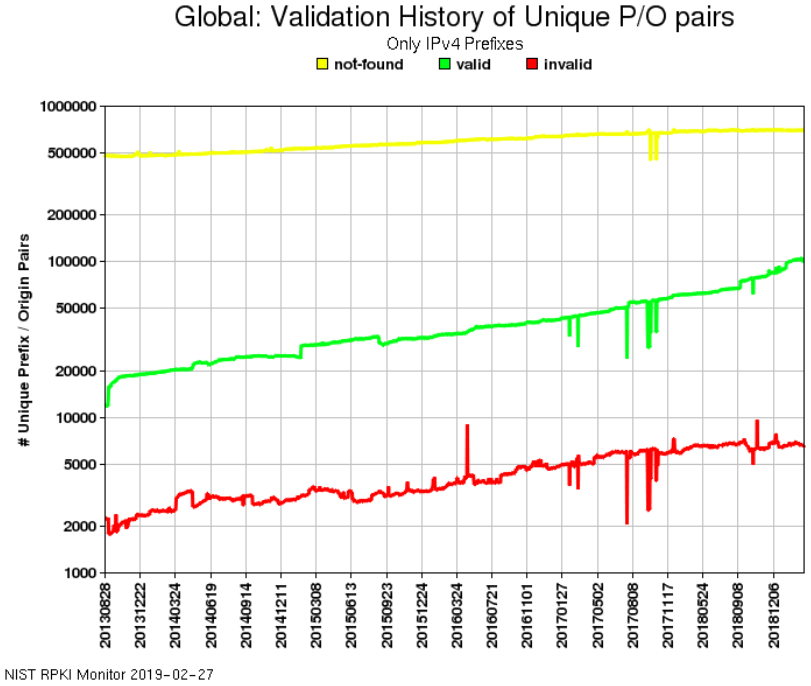

- Source: https://rpki-monitor.antd.nist.gov/?p=0&s=0

RPKI Status of APNIC Region – Snapshot

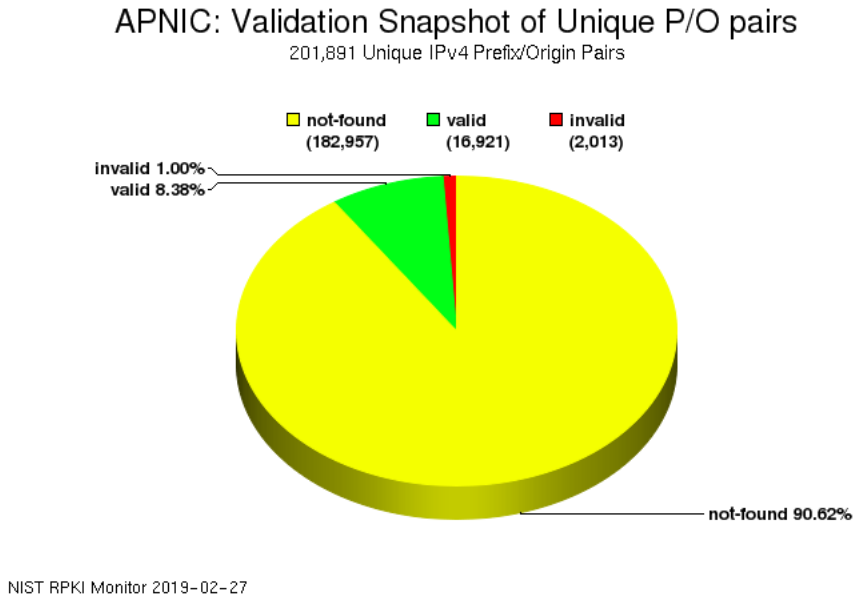

- Source: https://rpki-monitor.antd.nist.gov/?p=3&s=0

RPKI Status of APNIC Region – Trend

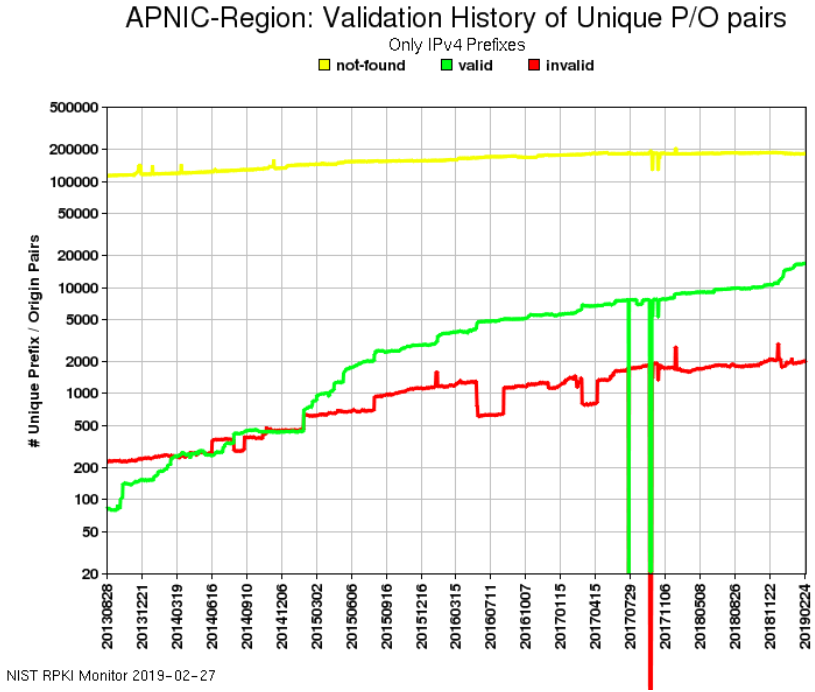

- Source: https://rpki-monitor.antd.nist.gov/?p=3&s=0

ROA Creation Statistics of APNIC Region

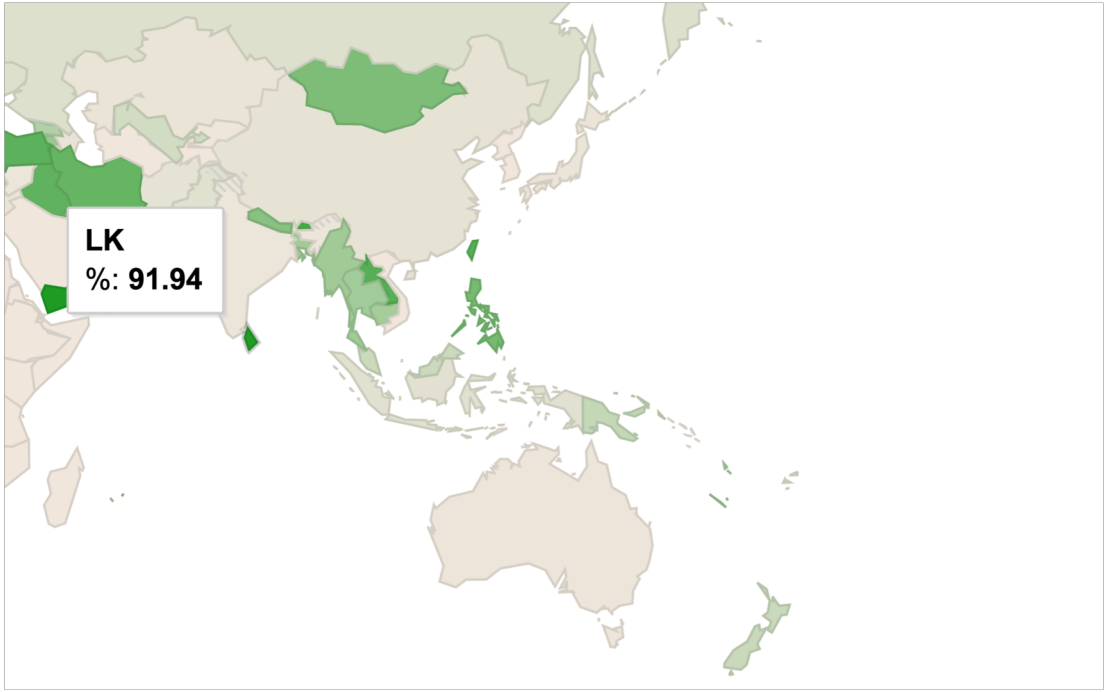

- Source: https://lirportal.ripe.net/certification/content/static/statistics/world-roas.html

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**
    - **You can contact APNIC Helpdesk at any time (https://www.apnic.net/get-ip/helpdesk/)**
- Next step is to do route origin validation (ROV) at your border routers
  - With or without your own validators

Industry Development on RPKI/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
- RPKI Operational Roundtable 2019 just held in San Francisco on Feb 17 (Sun) right before NANOG75
  - https://www.eventbrite.com/e/rpki-operational-roundtable-2019-tickets-8415556155?ref=enivtefor001&invite=MTU3ODY2NTgvcHdpbHNvbkBhcG5pYy5uZXQvMA%3D%3D%0A&utm_source=eb_email&utm_medium=email&utm_campaign=inviteformalv2&utm_term=eventpage
- Routing Security was one of the main themes at APRICOT 2019 held in Daejeon last week
  - https://2019.apricot.net/program/schedule/#/day/10
- Big players are getting more and more serious with RPKI/ROV…

Possible Implications to networks which are 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**
- More will be coming…
  - As a requirement for peering???

More About RPKI/ROV 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 RPKI deployment…

Questions?