APNIC 44

IRR and RPKI: a problem statement

George Michaelson, Ideas Group APNIC

TAICHUNG, TAIWAN
7–14 September 2017

#apnic44

Don’t Panic!

- This is not a ‘policy discussion’
  - There is no ‘decision’ to be made here right now
  - This is intended as informational only
- You might want to think about the issues being flagged
  - It may inform ideas about future service in the RIR system
  - It is likely to come up in APRICOT APOPS or other operator meetings
- If you care about routing integrity, reputation of INR
  - This problem deserves some engagement

Problem Statement Summary

- Routing management practice is globally ‘fragmented’
  - Fragmented both ‘within’ & ‘between’ different communities of interest
  - Now fragmenting across RPSL and RPKI
- There is a role for the RIR system
  - We have a critical role in ‘custodianship’ and ‘identity’ related to the assignments and allocations we administer
  - We need to improve coordination a little as a result of INR transfer and final IPv4 policy effects
- Where we go from here is up to the community
  - Can we bootstrap a more ‘coherent’ conversation?

Coverage (by amount of addresses)

Technology	% covered in BGP (apnic)	% Overlap RPSL/RPKI (apnic)
RPSL v4	75% (81%)	93% (97%) rpki in rpsl
RPSL v6	91% (67%)	94% (82%) rpki in rpsl
RPKI v4	12% (2%)	12% (15%) rpsl in rpki
RPKI v6	9% (1%)	10% (1%) rpsl in rpki

- Not all BGP is in RPSL, but much more is in RPSL than in RPKI
- Most RPSL is outside of RPKI
- Most RPKI is covered in RPSL, but not all (10-12% lies outside)
  - A lot of the overlap is (semi)automatic route: object creation

Overlap in RPSL

	APNIC	IRINN	JPIRR	RADB	RIPE
APNIC		<1%	<1%	66%	1%
IRINN	42%		-	15%	-
JPIRR	<1%	-		85%	<1%
RADB	10%	<1%	8%		6%
RIPE	<1%	-	<1%	12%	

- Low overlaps (<1%, -) are good.
- High overlaps risk of diverging IRR statements
  - When are the different IRR updated? what keeps them in sync?

Fragmentation of routing practice

- Routing management practice is globally fragmented
  - Different Internet Routing Registry (IRR), each with their own policies
  - Emerging RPKI methods which specify similar data to IRR
  - Lack of coherence between and within these mechanisms
- It may be getting “worse”
  - Because we have more technology now: RPSL **and** RPKI
  - Because we have more participants in global routing
    - More ASN holders, participating in BGP
    - More address custodians who are not ASN holders who need assistance
    - More resources moving between regions, BGP speakers leading to cross-region data

Fragmentation within RPSL

- Two primary sources: RIPE WHOIS and RADB
  - A kind of ‘europe’ / ‘rest of the world’ dichotomy
  - Content can conflict. Which one is right?
- Other sources. APNIC, AfriNIC, JPNIC (…) less ‘globally’ applied
  - Content now visible in several sources.
  - ISP specific (eg NTT) with automated customer-AS routing
  - National-scope (JPIRR) with strong (annual) checks
- Lack of visible cohesion. What determines ground-truth?
  - If IRR conflict?
  - If IRR are incomplete?
  - If IRR include data with no visible linkage to origin assigning registry?

Trust Model in RPSL

- Trust in the public assertion of route: stems from belief in the integrity of the data being managed by the IRR publisher
- RADB (merit) has no innate linkage back to RIR registry data
  - From where does trust in the assertions vest?
    - They are ’fiat’ declarations by each individual resource holder in RADB
- RIPE RPSL includes ’foreign’ objects with open maintainer
  - under a public maintainer and does not have innate linkage back to the origin RIR for foreign AS or Inetnum/inet6num
  - All non-foreign objects in RIPE WHOIS relate to custodian acts in RIPE
    - But all foreign objects are ‘fiat’ declarations as for RADB.
- JPIRR annual check worth reflecting on.

RPKI is a different trust model

- Cryptographically secured assertions about Internet Number Resources
  - Strongly coupled to assignment registry through X.509 PKI
  - Testable assertions backed by cryptography.
- Route Origin Attestation/Authorization (ROA)
  - Can include ‘MaxLength’ of prefixes associated with origin-AS
  - Does not include cryptographic checks of authority for origin-AS
    - Only the routed prefix holder authorizes the object
- All X.509 objects have lifetimes and need to be renewed

RPSL and RPKI are not the same

- RPSL has no MaxLength parameter
  - ROA with maxlength equivalent to a ‘set’ of route: objects
  - How many? /32 maxlength /48 in Ipv6 == 65,536 subsidiary route: objects!
- RPKI does not demand consent of the origin-AS maintainer
  - Contrast with RPSL Route: object which demands both prefix and AS holder to authorize
    - Or some kind of local override mechanism to force the function.
  - RPKI Cryptographic checks don’t include authority of AS to route, only prefix holder’s assertion to be routed via that origin-AS

The Landscape is shifting

- Conversation in RIPE NCC region over ‘foreign’ objects
  - Between routing-wg and database-wg participants
  - Movement of data from RIPE WHOIS into AfriNIC and APNIC (and LacNIC)
- Lack of conversation with RADB
  - Some prior art in RPSS, RPSL-Auth models to distributed authentication behind objects but nothing solid deployed
- Increasing uptake of RPKI (ready to ROA, other RIR initiatives)
  - Automatic route: and route6: object creation for ROA, magnifies data in WHOIS (result of translating MaxLength to set of route: objects)

The Landscape is shifting

- What does it mean when IRR & RPKI disagree?
  - Where is ground-truth, when conflicting statements exist?
- Increasing requests to APNIC HM & Helpdesk to mediate
  - *’can you help me remove the objects in RIPE/RADB which I didn’t authorize?’*
  - *’how can I authorize an RPSL object for an ASN I don’t own?’*
- Increasing ‘frustration’ from anti-spam activists in NOG mailing lists
  - *“why can’t you guys keep on top of this? Who is the gate-keeper for origin-as?”*

Lets have a more coherent conversation

- There is no single venue for this conversation
  - The subject is under discussion in several places
    - on NOG mailing lists, at NOG meetings
    - WG in various RIR, IETF
    - At helpdesk as RIR try to assist custodians negotiate the systems
- The RIR/NIR have a key role
  - What we have: identity and authority data for INR custodians
    - Strong proofs in RPKI of evidence of control over a resource
  - What we do: Processes for discussion, lists, venues to talk and delivery of community consensus messages

Some previously canvassed technology

- Add some kind of external authorization check to RIPE, RADB IRR whois model
  - Breaks the ‘maintainer’ object nexus
  - External DB dependency cannot be coherent all the time (CAP theorem)
- ‘Proof of possession’ checks analogous to nonce-hash in web (google) models
  - Add remarks: field to whois object, google sees you as ‘in control’
  - RPKI signed assertion from nonce
- Automatic route: object creation from ROA?
  - Without clear understanding of MaxLength cannot be complete for all possible ROA states
  - Add MaxLength to RPSL?
- IRRToolset modifications?
- Signed RPSL specification from RIPE NCC
  - Similar to DKIM, signature over ‘normalised’ fields in RPSL objects
  - Doesn’t actually stipulate RPKI certificates to make the signatures

Lets re-engage as a community

- APRICOT 2018 APOPS
  - Opportunity to encourage a conversation
- Would you like to be a part of the conversation?
- What role would you like APNIC to have in global routing?
  - We will continue to offer RPSL/IRR and RPKI as part of our core role in INR custodianship
  - Do you want us to explore any of the technology fixes?
- What do you think? Lets take it to a list!

APNIC 44

#apnic44

TAICHUNG, TAIWAN

7-14 September 2017