The ISP Column
A column on things Internet

December 2023
Geoff Huston,
George Michaelson

Models of Trust for the RPKI

This is a report on a feasibility study looking at an alternative trust anchor structure for the Resource Public Key Infrastructure (RPKI).

Background

In the early days of the Internet in the 1980's when the Internet address plan used the Class A, B and C address structure, it appears that Jon Postel, in his role as as the Internet's record keeper, the Internet Assigned Numbers Authority, or *IANA*, kept his own record of the allocations of Class A address blocks, but delegated the administration (and record-keeping) for the allocation of myriad of Class B and C address blocks to the Internet Network Information Centre, or *InterNIC* (operated by Stanford Research International, as the SRI-NIC). When the contract to operate the InterNIC was reassigned to Government Solutions in 1990 or thereabouts, it appears that much the same arrangements applied: IANA looked after the Class A address blocks records, and the InterNIC looked after the records of allocations of Class B and C address blocks.

In around 1993 the Internet dropped the Class A, B and C address structure, and moved to a variably sized allocation policy, but the split of responsibilities was unchanged, with the IANA holding the records of /8 address block assignments and the InterNIC looking after all smaller-sized blocks.

With the formation of the first Regional Internet Registry (RIR) by the European networking community in 1993, IANA delegated three Class A-sized blocks to the RIPE Network Coordination Centre (RIPE NCC), thereby allowing them to perform allocations from these address blocks and maintain their own registry of these allocations. In the ensuring years a second RIR was formed for the Asia Pacific Region (APNIC). There were further changes when the InterNIC database and the address administration roles were moved from Network Solutions (the successor to Government Solutions) to ARIN, which was notionally a registry for the Americas, but also maintained all the legacy Internet address allocation records from the InterNIC, as well as performing the address registry function for Africa. In subsequent years LACNIC was formed as a RIR for Latin America and the Caribbean, and Afrinic as the RIR for Africa.

The original two address registries, the IANA record of Class A allocations and the records of Class B and C allocations held by the InterNIC had become six address registries:

- The IANA ran a registry, holding the original Class A address allocations, the record of which /8 address blocks were assigned to which RIR and the residual pool of unallocated address blocks.

- Each RIR ran a registry of address allocations that were performed from the /8 address blocks that were allocated to them by the IANA.

- And ARIN also held the registry of the historical allocations that were performed before the formation of the RIR system.

The subsequent RIR activity included a program to perform early registration transfers (ERX) where many of these historical address allocation records were shifted from ARIN to the RIR that operated the address registry in that region. This transfer of registration records included IPv4 address blocks, some IPv6 address blocks and Autonomous System Numbers (ASNs).

This presented the address registration system with a basic question: how should this inter-RIR movement of resources be described in the address registries?

One option would've been to use the IANA as one-level-up registry, where each entry in an IANA-managed registry listed the RIR that was the current holder of the registration record for each address block. This approach apparently presented a number of issues, including a number of historical Class A allocations where the holder has expressed a preference to leave their registry record with the IANA and not pass this over to an RIR (for example the address blocks 38/8, 48/8 and 53/8 are still described in IANA registry entries that refer to end user allocations, and not to an RIR). So, this arrangement was not going to uniformly applied in all cases. This approach also implied that any movement of the registration record between RIRs would have also to be recorded in the IANA registry. The double registration would've added a further opportunity for conflicting information and would've increase the registry administration workload without any apparent incremental benefit.

Another option was for the IANA registry to simply record what the IANA actually did at the time. For example, the IANA allocated 210/8 to APNIC in June 1996, and this action is recorded in the IANA registry. Any subsequent movement of registration records between the RIRs would be a matter for the RIRs to administer and would not directly alter the IANA role as the registry operator of the central pool of resources, nor the historical record of the IANA's allocation actions. This is the option that was adopted, more by default than a deliberate decision by the RIRs and the IANA at any point in time.

This arrangement worked tolerably well for a while, but IPv4 address exhaustion and the desire to fairly apportion the residual shards of address space across the five RIRs would cause some revision to his model. The administration of a number of the early IANA /8 address allocations were re-assigned to a different RIR in the mid 2010's in an effort to equal the residual balance of any unassigned address blocks within those /8's across the RIRs/ The IANA records this re-assignment of administrative responsibility to the 'new' RIR. For example, the address block 154/8 is listed by the IANA as "Administered by AFRINIC", while the /8 address block was originally assigned to ARIN in 1992, and a large collection of registry entries that reflect prior allocations from that address block are held in the ARIN registry.

Further movement of address registration records between RIRs has resulted from the policy decision in a number of RIR communities to recognise inter-RIR address transfers, where the registry record is shifted from one RIR to another to match the address transfer.

The result is that the IPv4 address space is far more fragmented between the RIRs than the IANA registry might lead you to believe. To illustrate this, we've mapped the IPv4 address space onto a Hilbert curve, using a unique colour for each RIR, and where each pixel represents a /24 of IPv4 address space, and each colour represents each RIR. The map derived from the IANA registry is shown in Figure 1, and the map genmerated from the RIRs' registries is shown in Figure 2. The major differences between the two maps are the fragmentation of the old Class B legacy space across the RIRs and also the low part of the Class C legacy space.

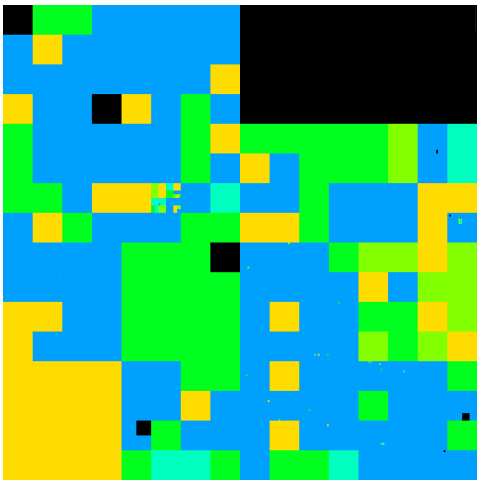

Figure 1 – Hilbert Curve of IANA IPv4 address registry

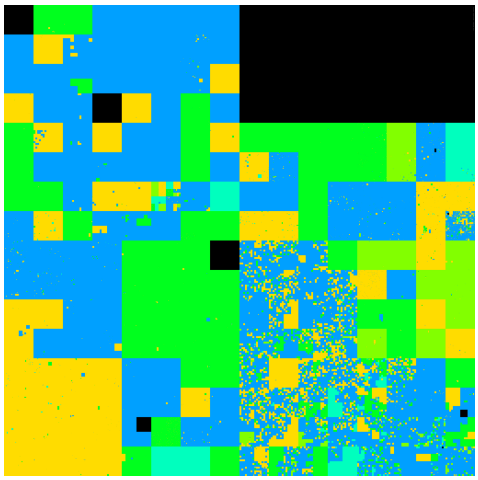

Figure 2 – Hilbert curve of RIR IPv4 address registries

There is a relevant question here that is pertinent to the design of the trust structure in the RPKI: How can an observer determine which RIR is authoritative for any given resource?

RPKI Trust Anchor(s)

A Trust Anchor in the PKI is a self-signed certificate that is trusted by an PKI Relying Party. The certificate validation process in the PKI attempts to establish a linked issuer/subject chain of PKI certificates that links a Trust Anchor to the certificate being validated.

In theory, if the PKI has multiple Trust Anchors, then any Trust Anchor is as good as any other if an appropriate validation chain of issuer/subject linked certificates can be established. However, this arrangement admits the possibility of rogue certificates being treated as valid in some circumstances, as we have experienced with multiple trust anchors used in the Web PKI. A stricter trust model admits just single trust anchor, and there is no inherent uncertainty as to which trust anchor to use for certificate validation.

An early proposed model of trust in the RPKI used a single trust anchor operated by the IANA. The issue with this model is that if the IANA uses only the information contained in the IANA-operated assignment registries to issue subordinate certificates, then the description of the fragmentation of address blocks across the RIRs cannot be represented in these subordinate certificates, as this information is not contained in IANA-managed registry entries.

The viable approach if we want to avoid having the IANA issuing certificates based on knowledge for which the IANA is not authoritative (which is not a robust situation if we are talking about trust at the apex of a PKI) is to fold this inter-RIR movement of registry entries into the certificate structure at a lower level of the RPKI certificate hierarchy.

For example, if an address holder with a registry entry held by ARIN has an address block drawn from 154/8, then, as Afrinic is the certified holder of 154/8, Afrinic would need to issue a subordinate certificate with the subject of ARIN with a resource section that includes this holder's addresses, thereby allowing ARIN to then issue a certificate to the address holder.

What happens if this address holder then concludes a transfer to a holder where the new registry entry is held by the RIPE NCC? Does the certificate chain replicate the history of each of these individual transfers? Or is this chain "wound back" to Afrinic, and even through Afrinic was not a party to this transfer it would need to revoke the resources in the certificate issued to ARIN and install the resources in the certificate it issues to the RIPE NCC. These alternatives are shown in Figure 3.

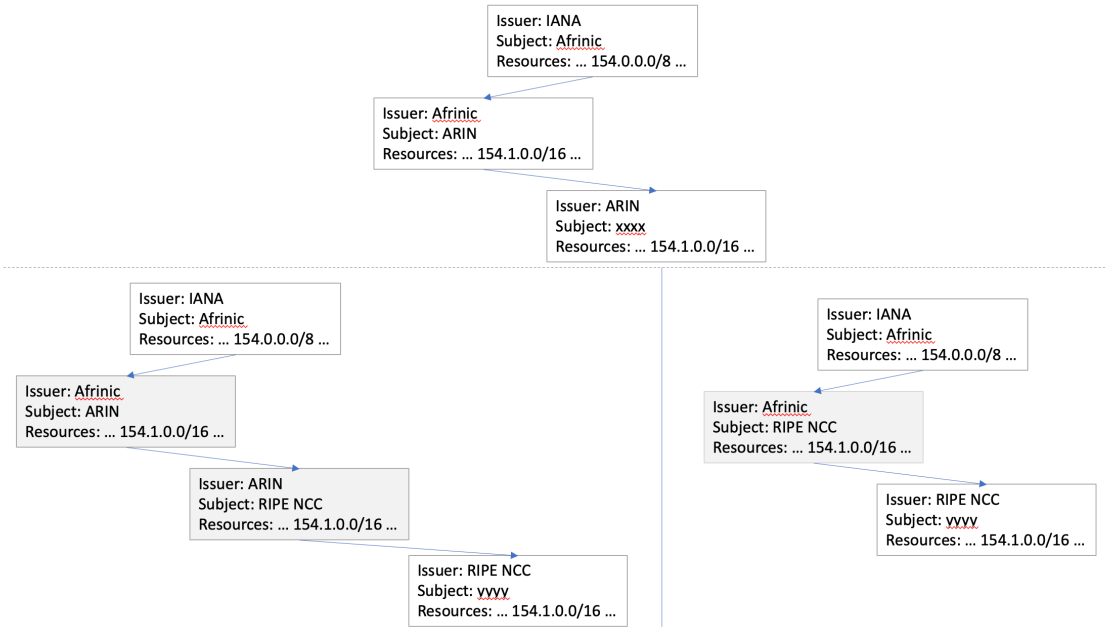

Figure 3 – Single Trust Anchor Model

Rather than head down this admittedly somewhat complex path, the adopted approach to Trust Anchors in the RPKI has been to use a framework with five trust anchors, one used by each RIR. Each of these RIR's self-signed certificates includes a resource list of all IP number resources. The transfer of resources between RIRs can be simplified in the RPKI system in that the disposing RIR need only revoke the certificate in question, and the RIR that receives the registration entry issues a new certificate (Figure 4).

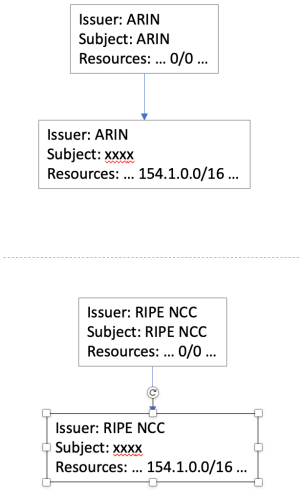

Figure 4 – Multi-Trust Anchor Model

This approach is simple, it can be implemented without strict process ordering of certification steps between the RIRs when processing resource transfers and provides consistent trust anchors for Relying Parties.

But it is a multi-trust anchor scenario, and a Relying Party has no a-priori knowledge as to which RIR's trust anchor is the "correct" trust anchor to use for certificate validation for any given number resource. Like the Web PKI, if any of the trusted operators can be coerced into falsely issuing certificate then there is no way that a relying party can avoid being misled.

Can we Improve this Situation?

The question is: Can we improve this situation and make it harder, if not impossible, for a RIR to be coerced into issuing a certificate for which it does not hold the authoritative registration record? And can we achieve this without a large-scale change of the current administrative processes used in address management in the RIRs? In other words, can we slide in a new single trust anchor at the apex of the RPKI without making any changes in the processes used by the RIRs, nor in the function of Relying Party tools?

What we want here is a single trust anchor which issues five subordinate certificates (one for each RIRs) that uniquely positions each number resource into just one RIR's certificate.

To drive this process, we need to know which RIR currently holds the authoritative registration record for each number resource. The RIRs already produce this information for each RIR in the form of a "daily stats report" and an NRO process performs a daily sweep of these reports to produce a single union report for the entirety of the Internet's number space.

This report is the missing "map" that resolves the fragmentation of the number space, and for every number resource there is a single RIR in this map that holds the registry record. It's this map that forms the foundation of this feasibility work in trying use a single RPKI trust anchor for the entire number framework.

A Prototype Single Trust Anchor

The trust anchor in this feasibility prototype itself is a conventional self-signed RPKI certificate, which has the value "`CN=example-nro-ta`" in both Issuer and Subject fields of the certificate.

The resources attribute of this certificate has the union of all listed internet number resources for each of the RIRs. There is no "all number resources" claim in this certificate: it only lists those resources which have been assigned into the RIR system by the IANA. Some resources, such as the documentation prefixes which are marked as reserved by the IANA, the multicast IP address ranges, and unassigned (by IANA) address ranges such as 127.0.0.0/8 and 240.0.0.0/4 are not included since they do not show in the IANA registry as being administered by any RIR and so should not appear in the products of any RIR's RPKI signed products.

This Trust Anchor certificate signs over instances of each of the RIR operational certificates. The operational certificates are the more-specific states which lie under the existing RIR 0/0 and 0::/0 Trust Anchor certificate for each RIR. Using an openssl facility to generate an X509 Certificate Signing Request (CSR) from a valid certificate, the single Trust Anchor keypair is used to sign over these operational certificates, changing the Issuer to the new "`CN=example-nro-ta`" but leaving the Subject and Subject Information Access attributes for each CA certificate unchanged. This means that a Relying Party will find all the same RIR-generated RPKI products unchanged, and can therefore validate all products found there, signed over by the RIR keypair.

This means that a top-down walk of the RPKI, starting with this single Trust Anchor certificate will find the subordinate RIR objects, re-issued by the single Trust Anchor but with the same public key, such that their respective subordinate products can be validated in this top-down pass across the entire RPKI publication tree.

This proposed single Trust Anchor overlay is shown in Figure 5.

The intention here is to leave the existing RIR RPKI products unchanged, and to allow existing RPKI client tools to also continue to operate without any changes to their operation, as long as they perform a traversal of the RPKI tree in a top-down manner. If the client tools support RFC 8360 (RPKI Validation Reconsidered), then any instances of duplicated resources in the RIRs' subordinate products would be resolved according to the per-RIR resource lists used by the single Trust Anchor.

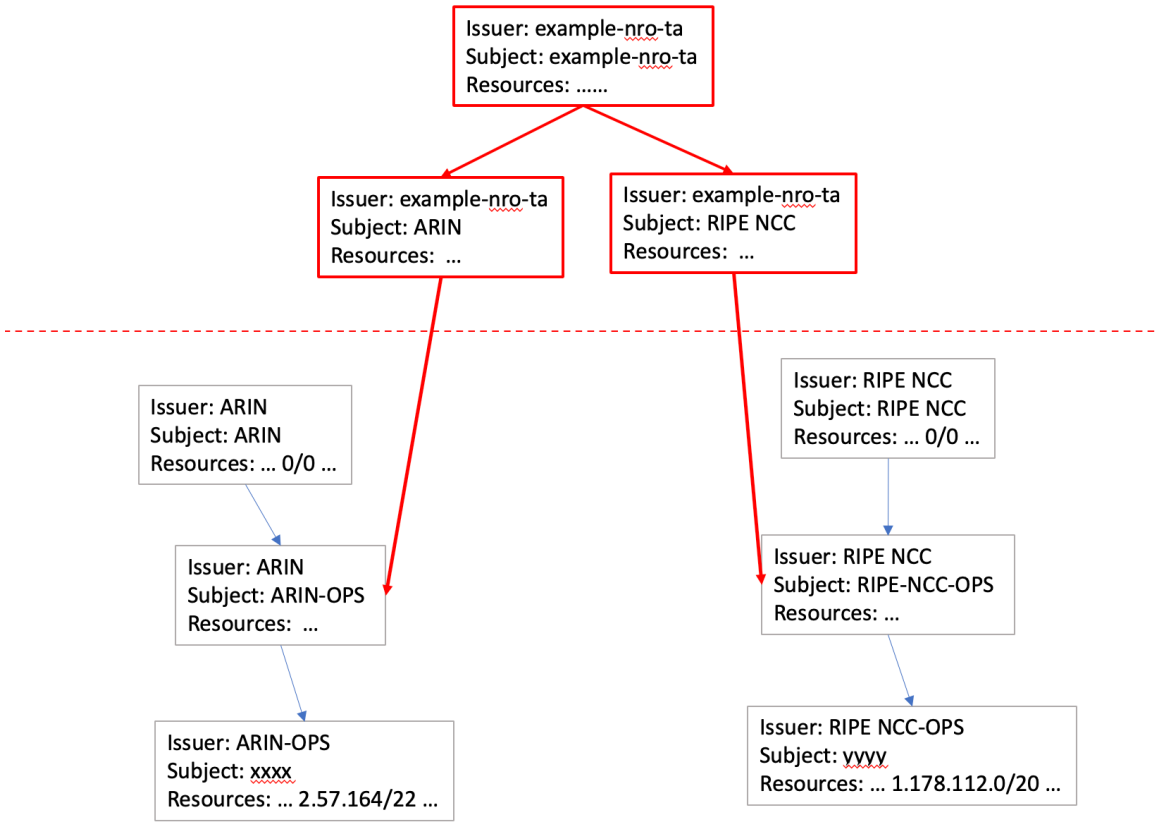

Figure 5 – Single Trust Anchor Overlay

The work in setting up this feasibility study is documented at https://labs.apnic.net/nro-ta/. The code modifications to OpenSSL are included there, along with the OpenSSL configuration details needed to issue the per-RIR subordinate certificate products under this single trust anchor.

Of course, this is just a proof-of-concept exercise, not a production service, and definitely should not be used to validate certificates in the public RPKI, but it does illustrate an approach to replacing the multiple Trust Anchors in the RPKI with a single Trust Anchor, removing the potential for ambiguities and clashes in validated subordinate certificate products in the RPKI. It also aligns all RPKI certificates at the upper level of the RPKI with the published registry states of the RIRs.

Disclaimer

The above views do not necessarily represent the views of the Asia Pacific Network Information Centre.