The ISP Column
A column on various things Internet

October 2018
Geoff Huston

Securing the Routing System at NANOG 74

The level of interest in the general topic of routing security seems to come in waves in our community. At times it seems like the interest from network operators, researchers, security folk and vendors climbs to an intense level, while at other times the topic appears to be moribund. If the attention on this topic at NANOG 74 is anything to go by we seem to be experiencing a local peak.

A Legal Perspective

The first session in this topic was a presentation by Christopher Yoo of the University of Pennsylvania Law School on the topic of Legal Barriers to Securing the Routing Infrastructure. It seems odd to me that the use of the Resource Public Key Infrastructure (RPKI) should encounter any legal barriers, but the US National Science Foundation saw merit in funding a study in this area, and in retrospect I think I was more than a little naïve to think that this area does not have a considerable legal; dimension.

RPKI deployment is gathering pace in some parts of the world. Approximately one quarter of all announced prefix/origin pairs in BGP from prefixes administered by the RIPE NCC and LACNIC have an associated published ROA. The comparable number of prefixes with ROAs is far smaller for APNIC (around 5%) and with ARIN and AFRINIC its even smaller (2% for ARIN and 1% for AFRINIC) (https://rpki-monitor.antd.nist.gov/?p=0&s=0). The objectives of the UPenn study are to catalog the claimed barriers to RPKI adoption, independently evaluate the legal and institutional barriers, and suggest viable solutions that balance the interests of all stakeholders.

It's certainly the case that RPKI has a similar set of perception issues to those that we see with DNSSEC. Network infrastructure operators typically do not see an overwhelming case in favour of adoption of the technology, and in that light the case to deploy is balanced against the incremental cost of adding this to the production operational environment and the associated roisk factor of adding one more element to the network operation that can go wrong. But these are not legal barriers and sit more conventionally in the same cost and benefit structure that applies to many technology adoption decisions. Why is RPKI any different?

The RPKI is a conventional hierarchical public key structure, where a small number of putative trust anchors issue CA certificates to others under specified criteria, who in turn may themselves issue CA certificates to others under the same criteria. The assumption in a PKI is that trust is commutative, such that if I am prepared to trust a certificate authority, then I should be able to trust those authorities who have been certified by the original trusted certificate authority, and so on. The implication is that a relying party need only pre-load the public keys of the certificate authorities in whom it is prepared to trust in the role of a 'trust anchor', and all other certificates can be validated by establishing a certificate chain from a trust anchor to the certificate that follow these commutative pairwise relationships. The inference is that these trusted certificates play a pivotal role in the PKI, and if there are any issues with these trusted certificates, then the entire set of subordinate issued certificates and their derived products cannot be validated. This is not without consequence if the PKI is pivotal to the Internet's routing system. If a significant set of digitally signed products cannot be validated there is a distinct risk that security-aware routers may conclude that a set of routes are invalid and should be discarded.

Why should there be an issue with a public key certificate? The certificate is a digitally signed attestation that the entity whose public key is the subject of an issued certificate has met certain criteria ssociated with certificate issuance. In the case of the Resource PKI the criteria include the holding of IP addresses and autonomous system numbers, and the certificate enumerates these IP number resources. The particular condition that relying parties apply to certificates in the validation chain is that the resources in each certificate be either identical or a superset of its immediate subordinate certificate. The onus on all certificate authorities is to ensure that the listed number resources be accurate at all times.

If a certificate is published with an incorrect set of number resources then there is the possibility that its subordinate certificates cannot be validated, which, as we've noted has implications for the reachability of services in the Internet through the connection to the routing system. Thus, certificate issuance is not without its potential legal liabilities, and it is not unusual for certificate authorities to limit their liabilities.

ARIN has taken the decision to require all users of its trust anchor certificate to accept a Relying Party Agreement before obtaining a copy of this certificate and this agreement binds the recipient not to further distribute the certificate to any third party. As Christopher Yoo pointed out in his presentation "There are no direct legal precedents; no record of lawsuits against an RIR for RPKI, no record of lawsuits against providers of the roots for TLS, SSL, DNSSEC, or IRR. But past history does not guarantee future results (i.e. lack of past lawsuits in other contexts does not guarantee no future lawsuits over RPKI)". There are residual risks in providing trust anchors and US law requires an explicit act of agreement, as publication of terms and conditions is not sufficient in the US context.

It's true that the internet is no longer an experiment and there are very real liability issues for providers of services. We've often relied on various forms of disclaimers and other notices that intend to limit the liabilities of such service providers, but perhaps this form of "as-is" and "best effort" service provision is inadequate in today's context. If the intended role of the intended RPKI is so central to the internet's routing system, then perhaps it’s also necessary to provide appropriate levels of service support when providing these critical trust anchor certificate services.

Routing Incidents as seen by BGPmon

BGPmon's Andree Toonk described a number of recent routing incidents where routes have been hijacked. All of these incidents involved the advertisement of prefixes where the origin AS reflects some form of route hijack, and one observation is that all of these hijacks would've been detectable as invalid and not propagated through the routing system were ROA's being observed.

However there are a couple of observations that temper this rather optimistic view of the utility and value of ROAs. The first is that only 63 networks appear to reject routes where the ROA indicates an invalid origination of the route. Out of some 63,000 networks in today's routing system that’s a very small number. Hopefully, this situation will improve over time. The seconds observation is that the ROAs would only have been effective if these route leaks were inadvertent operational mistakes. If these route leaks were deliberate routing hijacks, then the attackers would've been able to create the hijacked route with the ROA-defined origin AS. While prudent use of the maxlength parameter in the ROA could've mitigated more specific attacks, the potential for routing disruption based on deliberate hijacks while preserving the origin AS still remains.

For this form of automated routing security to be effective at the BGP level its necessary to also validate the AS Path in some manner, as well as protecting the origination. The BGPSEC protocol was intended to perform this AS Path validation, but unfortunately the prospects for BGPSEC are not looking good. BGPSEC provides insufficient incentive when the protocol is partially deployed, and the concept of loading every router with certs and keys to sign BGP updates would make many network operators nervous, and justifiably so. It seems that either we need to figure out how ROAs alone can be of value to the routing world, or we need to look for other mechanisms that would provide some means for BGP speakers to assure themselves that the AS Path that they receive is at a very minimum a plausible one.

Internet Route Registries

For many years we thought that one way to control the potential anarchy in the routing realm was through the use of Internet Route Registries (IRRs). We've been using IRRs since the early 1990's and we have accumulated a lot of experience with them.

The fact that we still see a constant stream of routing anomalies over the same period point to an observation that IRRs have not managed to stop this problem. As Qrator's Alexander Azimov observes, the use of IRR filters can potentially filter some hijacks, and filter some route leaks, but many AS-Sets in IRRs are poorly maintained, and there are gaps in the use of filters in networks that permit widespread propagation of routes that supposedly should've been filtered by IRR entries. RPKI and ROAs can help, but it's hard to claim that any such assistance any more than palliative at present.

As NTT's Job Snijders points out many of the IRRs operate with a non-existent security model, in so far as anyone can create IRR objects and thereby corrupt the filters that are generated from such IRRs. Job asserts that the major IRRs did not perform validation of entries and had a potential issue here. The ITT operated by the RIPE NCC has shifted to prevent the entry of route objects that refer to address space that is not managed by the RIPE NCC. This measure improves the integrity of the data, but of course limits the applicability scope of the IRR as it can only describe a subset of the routing environment. A similar measure is being adopted at ARIN.

Job also observes that many IRR clients use one of two major aggregators, namely `whois.radb.net` and `rr.ntt.net`. One proposed measure here is to include the collection of valid ROAs within these aggregators and prefer the ROA-derived data over IRR entries whenever the two data sources conflict. This appears to be a reasonable step, given that the ROA is signed with a time-limited certificate and the current of the ROA is determined by the ROA creator.

Dropping ROA-Invalid Routes

This measure is the one that excites the issues of risk and liability. If all the use of RPKI and ROAs is limited to an exercise in route preference then while performance of forwarding to a destination may be compromised by problems in the ROAs, reachability is not impaired. When talk moves to the option of dropping routes where the ROA is invalid than things get far more serious.

We are now hearing calls to shift routing behavior from de-preferencing ROA-invalid routes to dropping those routes. Of course not everyone needs to do this for the measure to be effective. The Internet is not a densely interconnected mesh. It is more like a sparse hierarchy. If a small set of transit providers and some of the route servers at the major exchanges made this change, routes that are ROA-invalid would be prevented from propagating across the entire Internet.

One of the criticisms of a ROA-only filtering framework is that while it may detect inadvertent mis-origination of routes it is easily circumvented by a determined hijacker. The hijacker simply needs to attach the 'correct' origin AS to the synthetic route object. However, it’s perhaps not as hopeless as it may sound. The worst form of attack is to fake a more specific announcement of the larger route, and the way that this can be mitigated is to pay close attention to the max length parameter of the ROA. The pragmatic observation is that we need to use measures that make hijacking harder, and we don't necessarily need to go all the way and attempt to make all forms of route hijack impossible.

Where do we go from here?

What do the larger actors do about routing security? Chris Morrow of Google presented their plans and, interestingly, RPKI and ROAS are not part of their immediate plans. By early 2019 they are looking at using something like OpenConfig to comb the IRRs and generate route filters for their peers. RPKI and ROAS will come later.

Perhaps Google is behaving more conservatively than others here. There is a sense in routing security conversations that IRRs and IRR filtering are yesterday's solution. There is some visible enthusiasm for deploying some form of RPKI-based automated solution. But we need to recognize that a full route origination and AS Path protective approach, as specified in the BGPSEC protocol, is a long way off. Indeed, it may well be the case that it's so far off that the prospects of eventual useful deployment are non-existent.

Are ROAS useful? Yes, and it is possible to use the RPKI infrastructure and ROAs to comb IRR data feeds and drop inconsistent IRR entries, and even go further to drop routes where the ROA validation points to an invalid outcome.

But dropping routes generates the problem of risk and liability for RPKI operators and certificate publishers. Dropping routes associated with inadvertent mishaps is good. Dropping routes associated with some deliberate attempts to hijack address is also good. But dropping perfectly valid routes is not so good, and if valuable transactions were prevented as a result then we can expect the affected parties to seek some form of redress.

To what extent RPKI publishers and relying parties are liable for the consequences of their actions remains an area which has many legally untested aspects. It may be the case that the various disclaimers and clear statements of the roles and responsibilities of the various actors will contain this liability under law. But the Internet is now the world's public communications infrastructure, and the monetary valuation of transactions that happen across the network are likely to be incalculably large. There is no public service disclaimer in this privately operated internet, and while the risk of causing some form of epic failure in the routing system are low, the consequent costs of reparation could be considerable.

I look forward to further observations from Christopher Yoo's study on these legal issues. Securing the routing system is a very worthy goal, but if the scale of risks to the various actors involved in this exercise are so large that they are existential in nature, then even if the likelihood of a significant is low then we need to be extremely cautious as to how we proceed.

Disclaimer

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

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net