Unwedging Routing Security Activity: URSA Roundtable Meeting – Report

May 2015.

Proposals for improving routing security have been published and, to greater or lesser degree, implemented. However, there are still challenges in fostering uptake of new technologies and obtaining necessary operational information to address old security issues.

To elaborate possible steps forward in improving the state of Internet routing security, industry experts were invited to the Unwedging Routing Security Activity (URSA) roundtable in April 2015. Participants were invited from technical organizations that are impacted and/or need to act in order to make these advances.

This report reflects the overall nature of discussion and general conclusions from the meeting. Concretely, a straw proposal for near-term operational updates was outlined, and interest was expressed in the possibility of "stress-testing" any eventual improved infrastructure.

Routing security – the objective

Participants discussed what "good" (not perfect) routing looks like. One of the points that surfaced early in discussion was that it is difficult to implement more robust security measures in today's Internet because of both a lack of reliable information about other networks' routing policies, and the challenge of doing complex verification calculations in realtime routing hardware at full traffic volume. For example, BCP38[1] is a well-established "best practice" for ingress filtering, dealing with an old problem, and yet source verification isn't viable on inter-provider core network interfaces

There is a distinction to be made between:

- Ability to rely on routing announcements with confidence that packets will be forwarded to the appropriate end network (*e.g.*, no route hijacking).
- Ability to validate the source of inbound packets (*e.g.,* blocking at least some spoofed packets used for DDoS)

For today's networks, both are important but the second is necessary for improving defense against DDoS attacks, which are a regular operational challenge. In today's operational reality, not only is it impossible to authenticate the source IP address, it isn't even possible for any given network to confirm that a packet could legitimately have originated from a remote network, because it is impossible to know the remote network's routing policy choices.

[1] https://tools.ietf.org/html/bcp38

These issues identified key areas needing change, and provided the framework for proposing some near- to medium-term steps that could be taken towards improvement.

Current operational realities and requirements

Participants recognized that networks and networking practices are quite diverse. Common practices in an end-user customer network are quite different from those of a transit network that connects other networks. This diversity has been a hallmark of the Internet. Not only is it unlikely to change, the best approaches to improving the Internet should support diversity and not be based on expectations of uniform practices across the board.

The Internet Society's work on Mutually Agreed Norms for Routing Security (MANRS)[2] was presented and discussed. The group articulated challenges with practical implementation, of the norms given existing tools, available information, and pragmatic realities of every day network operation. Nonetheless, participants generally agreed with and supported the philosophy of the MANRS manifesto's and its call for collective stewardship of the Internet.

The group discussed the diverse uses of Internet Routing Registries (IRRs)[3] and the challenges with using them to improve security. Currently, IRRs operate in complete independence – some are associated with Regional Internet Registries (RIRs) and some are not. Currently, there is no way to determine which IRR a given operator uses, and many choices exist. For example, some operators in North America use Merit's "RADB"; others use the ARIN IRR; still others use both for different purposes. In general, IRR data is of uncertain reliability because

- There is motivation to put information into the IRR, but little penalty for failing to update it. Since there is no expiry mechanism, out-of-date information persists forever.
- The different, competing theories about which IRR to use and the common (though not ubiquitous) practice of mirroring information between several.
- No comprehensive, authoritative list of IRRs nor statement of their scope.
- There isn't a single unified purpose for IRRs
    - E.g., some organizations put backup information in one IRR and operational in another

What can be done, in the near term

Meeting participants discussed a number of possible incremental developments that operators could undertake:

- Preferring IRRs which used RPSL authorization (RFC2725). This might require, and encourage, more IRRs to do so.
- Whitelisting source IP packet filters through primary *and* backup paths (likely requires more automation.)

[2] See https://www.routingmanifesto.org/manrs/ .
[3] See http://www.irr.net/ for an overview

- Establishing and using a mechanism to publish and reliably locate the routing policy information necessary to implement ingress filtering, as well as other route/routing security improvements.

Route Origin Policy Location

IRRs can be used to store and serve reliable information for routing security. This would be useful for relying networks if the resource holder had a mechanism to indicate (securely) which IRR held information about its routing policies and feasible paths. If relying networks consistently used the information, resource holders would have an added incentive to keep the information in the IRR up to date.

With that information available in a reliable fashion, it would be appropriate to filter announcements for more specific routes than are provided in the IRR information (helps prevent route hijacking). The information could also be used to implement BCP38-style ingress filtering.

One approach to implementing this uses reverse DNS to store a URI to routing policy objects[4]. Relying networks could use the information to validate announcements and set up filters in their routers without disrupting current routing practices and models.

This would also reduce the need to mirror data between IRRs, because resource holders could clearly communicate where the one authoritative copy of information is (as opposed to trying to copy it everywhere reliant networks might be likely to look for it).

Stress-testing

Some interest was expressed in developing an activity to "stress test" known routing system resources, giving operators an opportunity to determine objectively whether or how their routing technology and practices withstand common types of routing failures/bad behavior. The goal would be to do this:

- publicly
- based on voluntary participation
- on the live Internet (not testbed)
- non-destructively

No single straw proposal was agreed on, although several were discussed.

Conclusions and next steps

The meeting provided an opportunity to share common perspectives from very different network operator viewpoints. In the interest of improving the state of route and routing

[4] See http://techreports.verisignlabs.com/docs/tr-1140006-1.pdf, "TASRS: Towards a Secure Routing System Through Internet Number Resource Certification", for one such approach.

security in the near terms, discussions will continue to concretely define a path for the route origin policy location approach, and a possible stress-test activity.

Appendix – Meeting Participation

Thanks to Verisign for hosting the meeting session.

Attendees of the meeting are employed by a number of industry organizations impacted by routing security considerations (Verisign, ARIN, NTT, Comcast, Verizon, Cogent, Internet Society), but they were invited, and participated, in their personal capacity as industry experts:

- Leslie Daigle – ThinkingCat, convenor
- Danny McPherson
- Andy Newton
- Jared Mauch
- Tony Tauber
- John Brzozowski
- Allison Mankin
- Eric Osterweil
- Anthony (Alby) Williams
- Dan Bruns
- Hank Kilmer
- Phil Roberts