

Securing Routing Information

Findings from an Internet Society Roundtable
September 2009

Executive Summary

A roundtable discussion of the current state and prospective solutions for securing routing information elicited a wide variety of observations, many shared views and some differences of opinion. Operators are aware of the risks and have mechanisms, at different levels of automation, to deal with route hijacking and errors that advertise false routes. RPKI is seen as an important step toward improving routing security, although it directly solves only the small part of the problems with respect to address origination, not AS paths. The suspect quality of the data on which validity of address prefix announcement is based is a serious problem that requires immediate attention, and will probably take some time to address. Efforts were identified as either short or long term.

Background

Routing security is a decades-old problem, with occasional occurrences of route hijacking on the Internet. The IETF believes it is making incremental progress in the Secure Inter-Domain Routing (SIDR) working group. Protocol development is far enough along that attention is shifting toward deployment by network operators.

ISOC hosted this invitation-only roundtable discussion on securing routing information to obtain and share a better understanding of the perspective of network operators. Capturing the operator perspective is important for any deployment that significantly improves Internet security. Equipment vendors, Regional Internet Registries and protocol developers were not brought to the discussion in order to ensure focus on operator voices.

The discussion was focused on how operators secure routing information now and their plans for improvement. Although participants were aware of, and variously participate in the development of, the Resource Public Key Infrastructure (RPKI) design in the Secure Inter-Domain Routing (SIDR) working group of the IETF, this discussion was intended to inform rather than compete with SIDR progress.

Participants agreed that the report would document observations and views, but without attribution. After participants described their individual perspectives, discussion sought to identify consensus and differing views.

Observations

When each operator explained his or her own view of the current and prospective situation, there was quite a bit of commonality in views. The recurring (or widely supported) views are these:

1. Network operators are sensitive to concerns about hijacking an address prefix, especially since the notorious case¹ of Pakistan Telecom advertizing Google's YouTube prefixes (which was intended to only block access in the country but black-holed traffic globally).
2. Most network operators routinely check the address prefix a customer asks them to advertize into the Internet's core BGP, through Shared Whols Project (SWIP) or Internet Routing Registries (IRRs, e.g. Routing Assets Databases RADBs). Filtering BGP updates is rarely done between settlement-free

Peers. In many cases, the ISP manages the records for their customers in these databases. The quality of database maintenance varies by database and by region. For example, additions are often made to the RADb, but not deletions.

3. Who can authorize an ISP to advertise a given address prefix is insufficiently clear. Due to mergers, historical practices, and distributed operations, the individual requesting routing for a customer can vary. Because sometimes a customer is known by different names, making the implication of records in address databases ambiguous, it is difficult to rely solely on the data in the registration databases in order to authenticate requests
4. Current data for IPv4 address allocations and assignments is not very accurate. Validating and correcting this data is a significant amount of work, and is generally considered not worth the delay and effort required. However, because the data set for IPv6 is still fairly small, using a clean IPv6 data set as the starting point is suggested.
5. There is wide variation in how much address validity checks are automated. An experiment demonstrated that IRR data can be queried directly from a process in the router, and used to manipulate routing. Other operators suggest that the query can be integrated into the provisioning system and used to generate access control lists (ACLs) or prefix-lists. All operators agreed that there is a natural progression from the manual provisioning process to full automation, and each operator has the flexibility to stop at the solution best suited for his or her needs, or progress towards increased automation as the comfort level increases.
6. The business case for funding more automation to improve address filtering is difficult for several reasons. One reason is that filtering address prefixes for your own customer primarily benefits competitor's customers by preventing their prefixes from being hijacked. Another reason is that Internet operations within a large telecommunications company have been arguing that Internet operations are safe enough to be in this business, so now saying that we need funding to make it safe is not attractive.
7. If the cryptographic operations to validate certificates binding address prefixes to autonomous systems (AS) numbers is in the routers, deployment must be coordinated with scheduled equipment replacement, cannot require its own roll-out.
8. Protecting just the origination of address prefix, which is the scope of SIDR work, is only a tiny part of securing routing information. While it is acknowledged that securing BGP is off in the future, there is interest in filtering routing information based on the relationships between AS paths and business interconnection. One barrier to partial solutions appears to be reluctance to reveal business relationships to those not involved in them.
9. What happens when a RPKI certificate expires?

Discussion of what all participants brought to light led to the following list of topics, which was then roughly prioritized (ties indicated by letters) as context for the next day's discussion of consensus needs:

- 1 - which questions will be answered at which time? short-term or long-term
- 2a - protocol changes to BGP or not?
- 2b - will this use some new protocol or what we already have?
- 2c - do we need to change the BGP state machine?
- 2d - is this something for which we need new router software and/or hardware - Scudder said that adding signatures to routes is times-4 on memory.
- 3 - revocation
- 4 - P1: I would like to replace the widget that now checks against IRR with RPKI - P2: have the common parts done in open source so we can share.
- 5 - partial implementation - what to do with a prefix without a cert? prefer signed route
- 6 - business case
- 7 - IPv6 preference to IPv4 clean data
- 8 - behavior at bootstrapping or exception handling - what does a router do before it has a complete view of the world. e.g. converge first then validate
- 9 - dealing with inconsistencies - e.g. how to decide among redundant validators
- 10 - origin or path protection
- 10a - AS path relationship publishing question
- 11 - Whois data as the source? or IRR
- 12 - what needs to be added to RPKI?

Consensus needs and priorities

Consensus of the group was that the following list (developed interactively) were shared needs:

RPKI for origin validation should be pursued for both IPv4 and IPv6

The certificates for origin validation proposed in the RPKI design are considered useful and practical enough for deployment.

Uniqueness of IP address certification at the global level is required

Operators are adamant that they should not have to deal with conflicting origin certificates from multiple sources. One way to accomplish this would be a single RPKI hierarchy, but since that appears unlikely immediately, RIRs should not issue certificates that conflict with those from other RIRs, and should have a mechanism to resolve such conflicts.

IPv6 data cleanup now because it is easier - IPv4 harder and later

Operators recognize that cleaning the data on which origin validation depends will

take some time for the IPv4 data because of historical problems with tracking assignments. However, they expect that the recent beginning of IPv6 address allocations will enable that data to be cleaned before origination certificates are issued.

Authentication of resource holders (local solutions)

Solutions are needed for the problem that who is authorized to originate an address prefix is sometimes obscured by mergers; historical practices; and distributed operations, where employees of different branches of the customer organization request attachment. These solutions should include authentication and verification of individuals interacting with address allocators.

Need open source tools for certificate distribution and validation

Operators need open source tools to acquire, distribute and validate the origin certificates of RPKI.

Cost of (safe) business needs to be reduced, shared software tools would help

Because the benefit of each operator validating route origins accrues to customers who might connect through a different operator (by preventing a false origination), the costs of improving the safe business for all cannot fall on one or a few early adopters.

What can be done about path validation (short term: without changing BGP – long term: fix BGP) should be investigated

Operators agree that investigation is warranted for a system in which voluntary registration of AS path relationships might enable stronger routing security than just origin validation.

How invalidation of authority to route is done (including disputes) needs to be resolved

Operators are concerned about the process by which the authority to originate an address prefix is invalidated. They are worried that errors or different agendas in the certification chain must not introduce routing failures. They want mechanisms to resolve disputes when conflicting claims occur.

The consensus split of needs between those that should be accomplished in the short term from those in the long term was as follows. The interpretation of short term was between 2 and 4 years; of long term between 3 and 6 years, reflecting expected equipment replacement cycles.

Short term	Long term
Cert-validation widget	Hardware changes
Open software tools	Path protection – with protocol changes
Origin protection	AS path relationships
Clean IPv6 data	Protocol changes
Partial implementation	Clean IPv4 data
Revocation	Bootstrapping-exception handling
Path protection - w/o protocol changes	

Differences

Some differences of opinion were found.

While there was consensus that data on who holds what address prefix needs cleaning, there were different opinions where the clean data should be. Should clean data be in an IRR, reached through RPSL to suit existing tools, or just in the new RPKI? Several participants explained that in regions such as ARIN, the data in whois, reached through SWIP, is more reliable than that in the IRR.

The disagreement as to the need or practicality of a single root of certificates was nearly as clear in this roundtable discussion as it has been in the SIDR WG. However, most agreed that the announcement by the NRO (Number Resource Organization) that the goal of a single root, signed by NRO, will be approached first through roots signed by each RIR. There was also disagreement as to how well Steve Kent's proposal for relying parties as trust anchors resolves the question of a single root for the RPKI.

There was disagreement whether SIDR's rejection of black-lists (and BOAs) would enable partial deployment, although there was consensus that partial deployment of RPKI would last a very long time, especially for portions of the IPv4 address space. There was no consensus on the use of AS-0 to prevent advertisement of (otherwise black-listed) prefixes. There was concern about the interpretation of the max-length (negation or no-op?) parameter in ROAs.

While there was not consensus, the point was raised that there is functionality missing from RPKI: the ability to signal from a neighbor to black-hole a particular prefix.

The biggest difference was whether a repository indicating which AS was upstream of each other would enable protecting path information in the short term, before security in BGP addressed path protection. Protecting information about the origin of a prefix was described as a tiny fraction of the problem of hijacked routes. The alternate view was that the complexity of path protection was why earlier efforts did not work out.

Path forward

Several participants said that an immediate benefit of the roundtable was to see the extent to which their problems were shared. It was also observed that the situation is different in different regions, and possibly for ISPs of different sizes. That each operator's control of origin information protects other operators' customers makes this a communal problem. But all operators benefit from incremental protection by each other.

The first step in the path forward is to publish a report of this roundtable, after ensuring that concerns of participants about what it might conclude are resolved, which will be done through discussion of draft(s) on the mailing list. The mailing list will be limited to roundtable participants until the report is finished.

The issues and changes which this panel identified cover a wide scope. Incorporating that feedback and implementing any of the ideas requires coordination throughout the industry, potentially across several IETF Working Groups as well as every RIR, not to mention the larger operator community and hardware/software

vendors. This panel is an attempt to start the discussion so that it is not bounded by the limited scope that each represents, as a coordinated effort among all of the different organizations involved will help to ensure the best possible solution. The panel discussed meeting again to perhaps identify where the different areas of work that they identified might belong within the IETF and RIR policy process, and possibly even collaborate on drafts.

Participants

The following people participated in the ISOC Securing Routing Information roundtable and contributed to this discussion of the relevant issues:

Jay Borkenhagen (AT&T), Wes George (Sprint), Monika Machado (Microsoft Networks), Jared Mauch (NTT America), Taka Mizuguchi (NTT Communications), Chris Morrow (Google), Heather Schiller (Verizon Business), Jason Schiller (Verizon Business), Tom Scholl (AT&T), Ryan Shea (Verizon Business), Ruediger Volk (Deutsche Telekom), Tomoya Yoshida (NTT Communications), Leslie Daigle (ISOC), Mat Ford (ISOC), John Schnizlein (ISOC)

About the Internet Society

The Internet Society is a non-profit organization founded in 1992 to provide leadership in Internet related standards, education, and policy. With offices in Washington, D.C., and Geneva, Switzerland, it is dedicated to ensuring the open development, evolution, and use of the Internet for the benefit of people throughout the world. More information is available at: <http://InternetSociety.org>

¹ http://www.circleid.com/posts/82258_pakistan_hijacks_youtube_closer_look,
<http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube/>