SIDROPS Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: March 9, 2019

K. Sriram
O. Borchert
D. Montgomery
USA NIST
J. Snijders
NTT Communications

September 5, 2018

Origin Validation Policy Considerations for Dropping Invalid Routes
draft-sriram-sidrops-drop-invalid-policy-01

Abstract

During incremental deployment of RPKI and Route Origin Authorizations (and possibly under some transient conditions), network operators would wish to have a meaningful policy for dropping Invalid routes. Their goal is to balance (A) dropping Invalid routes so hijacked routes can be eliminated, versus (B) tolerance for missing or erroneously created ROAs for customer prefixes. This document considers a Drop Invalid if Still Routable (DISR) policy that is based on these considerations. The key principle of DISR policy is that an Invalid route can be dropped if a Valid or NotFound route exists for a subsuming less specific prefix.

Table of Contents

1. Introduction

During incremental deployment of RPKI [RFC6481] and Route Origin Authorizations [RFC6482] (and possibly under some transient conditions), network operators would wish to have a meaningful policy for dropping Invalid routes (see [RFC6811] for validation state definitions). Their goal is to balance (A) dropping Invalid routes so hijacked routes can be eliminated, versus (B) tolerance for missing or erroneously created ROAs for customer prefixes. This document considers a Drop Invalid if Still Routable (DISR) policy that is based on these considerations. The key principle of DISR policy is that an Invalid route can be dropped if a Valid or NotFound route exists for a subsuming less specific prefix.

The DISR policy applies in addition to (1) preferring Valid when more than one route exists for the same prefix, and (2) NotFound routes are always included in the best path selection process. Note that the existence of a NotFound route excludes the possibility of an alternate Valid or Invalid route for the same prefix or a subsuming less specific prefix.

This document also provides an algorithm for best path selection policy that considers Origin Validation (OV) outcome and includes the DISR policy.

2. Drop Invalid if Still Routable (DISR) Policy

When origin validation (OV) is performed on a BGP route, there are three possible outcomes: (1) Valid, (2) Invalid, or (3) NotFound (see definitions in [RFC6811]). During partial/incremental deployment of RPKI and Route Origin Authorizations, it is natural to always include Valid and NotFound routes in the path selection decision process. (Note: Valid and NotFound are mutually exclusive, i.e., there cannot be two routes for a prefix where one is Valid and the other is NotFound. The same is also true about Invalid and NotFound.) If Invalid routes are always dropped from consideration, then there would be no tolerance for missing or erroneously created ROAs for customer prefixes. Then, the question arises: Should an Invalid route be dropped only if another Valid or NotFound route exists for subsuming a less specific prefix? This policy is called Drop Invalid if Still Routable (DISR).

2.1. Motivation for the DISR Policy

Consider these scenarios:

Scenario 1: A transit ISP A (AS A) created a ROA for a /22 prefix they announce. They also announce a /24 prefix (subsumed in the /22) that is owned by directly-connected customer X (has no AS). But ISP A neglected to create a ROA for X's /24 prefix. Clearly, the announcement of X's /24 will be Invalid. ISP A happens to propagate to neighbors the /22 and the /24.

Scenario 2: Customer X (AS X) announces a /22 prefix to transit ISP A and a /24 prefix (subsumed in the /22) to transit ISP B. X is attempting to do traffic engineering (TE). X created a ROA for the /22, but neglected to have ROA coverage for the /24. Clearly, X's announcement of the /24 will be Invalid. X happens to propagate the /24 to ISP B; ISP B does not participate in OV and propagates the Invalid route to its neighbors.

In each of the above scenarios, DISR policy (applied at routers elsewhere in the Internet) ensures that traffic for the more specific (/24) still reaches the correct destination, i.e., customer X (albeit possibly via a suboptimal / non-TE path). Any actual hijacks of the /24 prefix would be dropped at all eBGP routers that employ the DISR policy.

Measurements show that there are 10,417 Invalid prefix-origin pairs in the global Internet (based on NIST Routeviews/RPKI/ROA data analysis, February 2018). Of these, 6846 are Invalid due to maxlength violation. 6027 (of the 6846) are seen to be routable via Valid or NotFound routes for the same prefix (as in the Invalid

route) or a subsuming less specific prefix. Again, 5987 (of the 6027) are routes for which the corresponding Valid or NotFound routes (with the same or subsuming less specific prefix) have the exact same origin AS as in the Invalid route in question. These measurements show that Scenarios 1 and 2 described above do occur in significant numbers currently. So, the data lends support the efficacy of the DISR policy in terms of delivering the data traffic to the right destination (though not necessarily via the optimal/TE path).

The following is recommended in BCP 185 [RFC7115]: "Before issuing a ROA for a super-block, an operator MUST ensure that all sub-allocations from that block that are announced by other ASes, e.g., customers, have correct ROAs in the RPKI." However, as seen by the above measurement data, there are lapses in following this recommendation.

Network operators who do not wish to drop Invalid routes outright (during partial deployment or possibly in transient conditions), SHOULD consider employing the DISR policy. It helps eliminate actual prefix hijacks, while incentivizing creation of required ROAs and the adherence to the above recommendation from BCP 185. The stick used here is the possibility of data traveling via a suboptimal path, while the more aggressive stick of dropping all Invalid routes is held in abeyance.

3. Algorithm for Implementation of DISR Policy

An algorithm for implementation of the DISR policy is as follows.

Perform the following steps when a route is received:

1. Perform OV [RFC6811].

2. The second step consists of:

   * Modify LOCAL_PREF value: Add Kv if Valid; Add Knf if NotFound; Add Ki if Invalid (Kv > Knf >> Ki).

   * Store the route in RIB-in.

3. Apply route selection algorithm (this includes consideration of LOCAL_PREF and other parameters such as MED, etc. in the appropriate order [RFC4271]).

4. If selected route is Valid/NotFound, then add the route to Loc-RIB; Else, if Invalid, then add the route to Loc-RIB only if there is no existing route in the Loc-RIB for a subsuming Less Specific prefix.

Additional steps in the algorithm that are performed in reaction to addition/withdrawal of routes that influence DISR policy decisions and due to changes in RPKI:

a. When a Valid/NotFound route is added to Loc-RIB, check to see if there are any more-specific prefixes subsumed by the route prefix that are in Loc-RIB; If such more-specific prefix is Invalid, then remove it from Loc-RIB.

b. When a Valid/NotFound route is withdrawn from Loc-RIB, check to see if there are any more-specifics prefixes subsumed by the route prefix that are in RIB-in; If such more-specific prefix is Invalid, then rerun the route selection decision (Steps 3 and 4 above) for it.

c. When router learns of RPKI state change, then list all the prefixes effected by it. Rerun Steps 1 through 4 for those prefixes.

4. Security Considerations

This document addresses some aspects of best common practices for origin validation and related BGP policy. The security considerations provided in RFC 6811 [RFC6811] and BCP 185 [RFC7115] also apply here.

5. Normative References

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, DOI 10.17487/RFC6481, February 2012, <https://www.rfc-editor.org/info/rfc6481>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <https://www.rfc-editor.org/info/rfc6482>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/info/rfc6811>.

[RFC7115] Bush, R., "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)", BCP 185, RFC 7115, DOI 10.17487/RFC7115, January 2014, <https://www.rfc-editor.org/info/rfc7115>.

Acknowledgements

The authors wish to thank Sebastian Spies, Saku Ytti, Jeffrey Haas, Tim Bruijnzeels, and Jay Borkenhagen for comments and discussion related to this work. Also, thanks are due to Lilia Hannachi for her insightful analysis of global RPKI and BGP data that has been helpful in this work.

Authors' Addresses

Kotikalapudi Sriram
USA National Institute of Standards and Technology

Email: ksriram@nist.gov

Oliver Borchert
USA National Institute of Standards and Technology

Email: oliver.borchert@nist.gov

Doug Montgomery
USA National Institute of Standards and Technology

Email: dougm@nist.gov

Job Snijders
NTT Communications

Email: job@ntt.net