

IDR Working Group
Internet-Draft
Intended status: Informational
Expires: January 3, 2019

K. Sriram, Ed.
USA NIST
July 2, 2018

Design Discussion of Route Leaks Solution Methods
draft-sriram-idr-route-leak-solution-discussion-00

Abstract

This document captures the design rationale of the route leaks solution document [draft-ietf-idr-route-leak-detection-mitigation]. The designers needed to balance many competing factors, and this document provides insights into the design questions and their resolution.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 3, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Related Prior Work	2
3. Design Rationale and Discussion	3
3.1. Explanation of Rules 1 and 2 in the solution document . .	3
3.2. Is route-leak solution without cryptographic protection an attack vector?	5
3.3. Combining results of route-leak detection, OV and BGPsec validation for path selection decision	6
3.4. Are there cases when valley-free violations can be considered legitimate?	6
3.5. Comparison with other methods (routing security BCPs) . .	7
3.6. Per-Hop RLP Field or Single RLP Flag per Update?	8
3.7. Prevention of Route Leaks at Local AS: Intra-AS Messaging	10
3.7.1. Non-Transitive BGP Community for Intra-AS Messaging .	10
3.8. Stopgap Solution when Only Origin Validation is Deployed	11
4. Security Considerations	12
5. IANA Considerations	12
6. References	12
6.1. Normative References	12
6.2. Informative References	12
Acknowledgements	15
Contributors	15
Author's Address	16

1. Introduction

This document captures the design rationale of the route leaks solution document [I-D.ietf-idr-route-leak-detection-mitigation]. The designers needed to balance many competing factors, and this document provides insights into the design questions and their resolution.

2. Related Prior Work

The solution described in [I-D.ietf-idr-route-leak-detection-mitigation] is based on setting an attribute in BGP route announcement to manage the transmission/receipt of the announcement based on the type of neighbor (e.g., customer, transit provider, etc.). Documented prior work related to this basic idea and mechanism dates back to at least the 1980's. Some examples of prior work are: (1) Information flow rules described in [proceedings-sixth-ietf] (see pp. 195-196); (2) Link Type described in [RFC1105-obsolete] (see pp. 4-5); (3) Hierarchical Recording described in [draft-kunzinger-idrp-IS010747-01] (see Section 6.3.1.12). The problem of route leaks and possible solution mechanisms based on encoding peering-link type information, e.g., P2C

(i.e., Transit-Provider to Customer), C2P (i.e., Customer to Transit-Provider), p2p (i.e., peer to peer) etc., in BGPsec updates and protecting the same under BGPsec path signatures have been discussed in IETF SIDR WG at least since 2011. [draft-dickson-sidr-route-leak-solns] attempted to describe these mechanisms in a BGPsec context. The draft expired in 2012. [draft-dickson-sidr-route-leak-solns] defined neighbor relationships on a per link basis, but in [I-D.ietf-idr-route-leak-detection-mitigation] the relationship is encoded per prefix, as routes for prefixes with different peering relationships may be sent over the same link. Also [draft-dickson-sidr-route-leak-solns] proposed a second signature block for the link type encoding, separate from the path signature block in BGPsec. By contrast, in [I-D.ietf-idr-route-leak-detection-mitigation], when BGPsec-based solution is considered, cryptographic protection is provided for Route-Leak Protection (RLP) encoding using the same signature block as that for path signatures (see Section 3.2.2 in [I-D.ietf-idr-route-leak-detection-mitigation]).

3. Design Rationale and Discussion

This section provides design justifications for the solution specified in [I-D.ietf-idr-route-leak-detection-mitigation], and also answers some questions that are anticipated or have been raised in the IETF IDR and SIDR working group meetings.

3.1. Explanation of Rules 1 and 2 in the solution document

In Section 3.3 in [I-D.ietf-idr-route-leak-detection-mitigation], Rules 1 and 2 were stated and the route leak mitigation policy was based on these rules to preserve the property of stable route convergence (i.e., avoid possibility of persistent route oscillations). Rule 1 is stated as follows:

- o Rule 1: If ISP A receives a route r1 from customer AS C and another route r2 from provider (or peer) AS B (for the same prefix), and both routes r1 and r2 contain AS C and AS X (any X not equal to C) in the path and contain [X] in their RLP Attributes, then prioritize the customer (AS C) route over the provider (or peer) route.

The rationale for Rule 1 can be developed as follows.

Preference condition for route stability: Prefer customer routes over peer or provider routes (see pp. 25-27 in [Gao-Rexford]).

Topology condition for route stability: No cycle of customer-provider relationships (see pp. 25-27 in [Gao-Rexford]).

Route-Leak Detection Theorem: Let it be given that ISP A receives a route r1 from customer AS C and another route r2 from provider AS B (for the same prefix), and each of the routes r1 and r2 contains AS C and AS X in its AS path and contains [X] in its RLP Attribute. Then, clearly r1 is in violation of [X]. It follows that r2 is also necessarily in violation of [X].

Proof: Let us suppose that r2 is not in violation of [X]. That implies that r2's path from C to B to A included only P2C links. That would mean that there is a cycle of customer-provider relationships involving the ASes in the AS path in r2. However, any such cycle is ruled out in practice by the topology condition for route stability as stated above. QED.

Corollary 1: The route leak detection theorem holds also when "provider AS B" in the theorem is replaced by "peer AS B". (Here peer means a lateral peer.)

Proof: Since r2 contains [X] in the RLP Attribute set by an AS prior to peer AS B, it follows that r2 is in violation of [X]. QED.

It can be observed that Rule 1 follows from the combination of the Theorem, Corollary 1 and the preference condition for route stability (stated above).

In Section 3.3 in [I-D.ietf-idr-route-leak-detection-mitigation], Rule 2 is stated as follows:

- o Rule 2: If ISP A receives a route r1 from peer AS C and another route r2 from provider AS B (for the same prefix), and both routes r1 and r2 contain AS C and AS X (any X not equal to C) in the path and contain [X] in their RLP Attributes, then prioritize the peer (AS C) route over the provider (AS B) route.

The rationale for Rule 2 can be developed as follows.

Corollary 2: The route leak detection theorem holds also when "customer AS C" in the theorem is replaced by "peer AS C".

Proof for Corollary 2: Let us suppose that r2 is not in violation of [X]. That implies that r2's path from C to B to A included only P2C links. This results in a topology in which A's lateral peer B is also A's transit provider's transit provider. This gives rise to possibility of looping of routes since A can send routes to its transit B, B can forward the routes to its transit C, and C can

forward the routes to its peer A. But such looping is forbidden by the topology condition stated above.

It can be observed that Rule 2 follows from Corollary 2. In essence, if the provider route (r2) is a detoured (longer) version of the lateral peer route (r1), and violates the same RLP [X] as does the peer route, then prefer the shorter route (r2) via the peer.

Rules 1 and 2 are

3.2. Is route-leak solution without cryptographic protection an attack vector?

It has been asked if a route-leak solution without BGPsec, i.e., when RLP Fields are not protected, can turn into a new attack vector. The answer seems to be: not really! Even the NLRI and AS_PATH in BGP updates are attack vectors, and RPKI/OV/BGPsec seek to fix that. Consider the following. Say, if 99% of route leaks are accidental and 1% are malicious, and if route-leak solution without BGPsec eliminates the 99%, then perhaps it is worth it (step in the right direction). When BGPsec comes into deployment, the route-leak protection (RLP) bits can be mapped into BGPsec (using the Flags field) and then necessary security will be in place as well (within each BGPsec island as and when they emerge).

Further, let us consider the worst-case damage that can be caused by maliciously manipulating the RLP Field values in an implementation without cryptographic protection (i.e., sans BGPsec). Manipulation of the RLP bits can result in one of two types of attacks: (a) Upgrade attack and (b) Downgrade attack. Descriptions and discussions about these attacks follow. In what follows, P2C stands for transit provider to customer (Down); C2P stands for customer to transit provider (Up), and p2p stands for peer to peer (lateral or non-transit relationship).

(a) Upgrade attack: An AS that wants to intentionally leak a route would alter the RLP encodings for the preceding hops from 1 (i.e., 'Do not Propagate Up or Lateral') to 0 (default) wherever applicable. This poses no problem for a route that keeps propagating in the 'Down' (P2C) direction. However, for a route that propagates 'Up' (C2P) or 'Lateral' (p2p), the worst that can happen is that a route leak goes undetected. That is, a receiving router would not be able to detect the leak for the route in question by the RLP mechanism described here. However, the receiving router may still detect and mitigate it in some cases by applying other means such as prefix filters [RFC7454] [NIST-800-54]. If some malicious leaks go undetected (when RLP is deployed without BGPsec) that is possibly a

small price to pay for the ability to detect the bulk of route leaks that are accidental.

(b) Downgrade attack: RLP encoding is set to 1 (i.e., 'Do not Propagate Up or Lateral') when it should be set to 0 (default). This would result in a route being mis-detected and marked as a route leak. By default, RLP encoding is set to 0, and that helps reduce errors of this kind (i.e., accidental downgrade incidents). Every AS or ISP wants reachability for prefixes it originates and for its customer prefixes. So, an AS or ISP is not likely to change an RLP value 0 to 1 intentionally. If a route leak is detected (due to intentional or accidental downgrade) by a receiving router, it would prefer an alternate 'clean' route from a transit provider or peer over a 'marked' route from a customer. It may end up with a suboptimal path. In order to have reachability, the receiving router would accept a 'marked' route if there is no alternative that is 'clean'. So, RLP downgrade attacks (intentional or accidental) would be quite rare, and the consequences do not appear to be grave.

3.3. Combining results of route-leak detection, OV and BGPsec validation for path selection decision

Combining the results of route-leak detection, OV, and BGPsec validation for path selection decision is up to local policy in a receiving router. As an example, a router may always give precedence to outcomes of OV and BGPsec validation over that of route-leak detection. That is, if an update fails OV or BGPsec validation, then the update is not considered a candidate for path selection. Instead, an alternate update is chosen that passed OV and BGPsec validation and additionally was not marked as route leak.

If only OV is deployed (and not BGPsec), then there are six possible combinations between OV and route-leak detection outcomes. Because there are three possible outcomes for OV (NotFound, Valid, and Invalid) and two possible outcomes for route-leak detection (marked as leak and not marked). If OV and BGPsec are both deployed, then there are twelve possible combinations between OV, BGPsec validation, and route-leak detection outcomes. As stated earlier, since BGPsec protects the RLP encoding, there would be added certainty in route-leak detection outcome if an update is BGPsec valid (see Section 3.2).

3.4. Are there cases when valley-free violations can be considered legitimate?

There are studies in the literature [Anwar] [Giotsas] [Wijchers] observing and analyzing the behavior of routes announced in BGP updates using data gathered from the Internet. The studies have

focused on how often there appear to be valley-free (e.g., Gao-Rexford [Gao] model) violations, and if they can be explained [Anwar]. One important consideration for explanation of the violations is per-prefix routing policies, i.e., routes for prefixes with different peering relationships may be sent over the same link. One encouraging result reported in [Anwar] is that when per-prefix routing policies are taken into consideration in the data analysis, more than 80% of the observed routing decisions fit the valley-free model (see Section 4.3 and SPA-1 data in Figure 2). [Anwar] also observes, "it is well known that this model [the basic Gao-Rexford model and some variations of it] fails to capture many aspects of the interdomain routing system. These aspects include AS relationships that vary based on the geographic region or destination prefix, and traffic engineering via hot-potato routing or load balancing." So, there may be potential for explaining the remaining (20% or less) violations of valley-free as well.

One major design factor is that the Route-Leak Protection (RLP) encoding is per prefix. Hence, the solution is consistent with ISPs' per-prefix routing policies. Large global and other major ISPs will be the likely early adopters, and they are expected to have expertise in setting policies (including per prefix policies, if applicable), and make proper use of the RLP indications on a per prefix basis. When the large ISPs participate in this solution deployment, it is envisioned that they would form a ring of protection against route leaks, and co-operatively avoid many of the common types of route leaks that are observed. Route leaks may still happen occasionally within the customer cones (if some customer ASes are not participating or not diligently implementing RLP), but such leaks are unlikely to propagate from one large participating ISP to another.

3.5. Comparison with other methods (routing security BCPs)

It is reasonable to ask if techniques considered in BCPs such as [RFC7454] (BGP Operations and Security) and [NIST-800-54] may be adequate to address route leaks. The prefix filtering recommendations in the BCPs may be complementary but not adequate. The difficulty is in ISPs' ability to construct prefix filters that represent their customer cones (CC) accurately, especially when there are many levels in the hierarchy within the CC. In the RLP-encoding based solution described here, each AS sets RLP for each route propagated and thus signals if it must not be subsequently propagated to a transit provider or peer.

AS path based Outbound Route Filter (ORF) described in [I-D.ietf-idr-aspath-orf] is also an interesting complementary technique. It can be used as an automated collaborative messaging system (implemented in BGP) for ISPs to try to develop a complete

view of the ASes and AS paths in their CCs. Once an ISP has that view, then AS path filters can be possibly used to detect route leaks. One limitation of this technique is that it cannot duly take into account the fact that routes for prefixes with different peering relationships may be sent over the same link between ASes. Also, the success of AS path based ORF depends on whether ASes at all levels of the hierarchy in a CC participate and provide accurate information (in the ORF messages) about the AS paths they expect to have in their BGP updates.

3.6. Per-Hop RLP Field or Single RLP Flag per Update?

The route-leak detection and mitigation mechanism described in [I-D.ietf-idr-route-leak-detection-mitigation] is based on setting RLP Fields on a per-hop basis. There is another possible mechanism based on a single RLP flag per update.

Method A - Per-Hop RLP Field: The sender (eBGP router) on each hop in the AS path sets its RLP Field = 1 if sending the update to a customer or lateral peer (see Section 3.2 in [I-D.ietf-idr-route-leak-detection-mitigation]). No AS (if operating correctly) would rewrite the RLP Field set by any preceding AS.

Method Z - Single RLP Flag per Update: As it propagates, the update would have at most one RLP flag. Once an eBGP router (in the update path) determines that it is sending an update towards a customer or lateral peer AS, it sets the RLP flag. The flag value equals the AS number of the eBGP router that is setting it. Once the flag is set, subsequent ASes in the path must propagate the flag as is.

To compare Methods A and Z, consider the example illustrated in Figure 1. Consider a partial deployment scenario in which AS1, AS2, AS3 and AS5 participate in RLP, and AS4 does not. AS1 (2 levels deep in AS3's customer cone) has imperfect RLP operation. Each complying AS's route leak mitigation policy is to prefer an update not marked as route leak (see Section 3.3 in [I-D.ietf-idr-route-leak-detection-mitigation]). If there is no alternative, then a transit-provider may accept and propagate a marked update from a customer to avoid unreachability. In this example, multi-homed AS4 leaks a route received for prefix Q from transit-provider AS3 to transit-provider AS5.




Figure 1: Example for comparison of Method A vs. Method Z

If Method A is implemented in the network, the two BGP updates for prefix Q received at AS5 are (note that AS4 is not participating in RLP):

U1A: Q [AS4 AS3 AS2 AS1] {RLP3(AS3)=1, RLP2(AS2)=0, RLP1(AS1)=1} from AS4

U2A: Q [AS3 AS2 AS1] {RLP3(AS3)=1, RLP2(AS2)=0, RLP1(AS1)=1} from AS3

Alternatively, if Method Z is implemented in the network, the two BGP updates for prefix Q received at AS5 are:

U1B: Q [AS4 AS3 AS2 AS1] {RLP(AS1)=1} from AS4

U2B: Q [AS3 AS2 AS1] {RLP(AS1)=1} from AS3

All received routes for prefix Q at AS5 are marked as route leak in either case (Method A or Z). In the case of Method A, AS5 can use additional information gleaned from the RLP fields in the updates to possibly make a better best path selection. For example, AS5 can determine that U1A update received from its customer AS4 exhibits

violation of two RLP fields (those set by AS1 and AS3) and one of them was set just two hops away. But U2A update exhibits that only one RLP field was violated and that was set three hops back. Based on this logic, AS5 may prefer U2A over U1A (even though U1A is a customer route). This would be a good decision. However, Method Z does not facilitate this kind of more rational decision process. With Method Z, both updates U1B and U2B exhibit that they violated only one RLP field (set by AS1 several hops away). AS5 may then prefer U1B over U2B since U1B is from a customer, and that would be bad decision. This illustrates that, due to more information in per-hop RLP Fields, Method A seems to be operationally more beneficial than Method Z.

Further, for detection and notification of neighbor AS's non-compliance, Method A (per-hop RLP) is better than Method Z (single RLP). With Method A, the bad behavior of AS4 would be explicitly evident to AS5 since it violated AS3's (only two hops away) RLP field as well. AS5 would alert AS4 and AS2 would alert AS1 about lack of compliance (when Method A is used). With Method Z, the alerting process may not be as expeditious.

3.7. Prevention of Route Leaks at Local AS: Intra-AS Messaging

Note: The intra-AS messaging for route leak prevention can be done using a non-transitive BGP Community or Attribute. The Community-based method is described below. For the BGP Attribute-based method, see [I-D.ietf-idr-bgp-open-policy].

3.7.1. Non-Transitive BGP Community for Intra-AS Messaging

The following procedure (or similar) for intra-AS messaging (i.e., between ingress and egress routers) for prevention of route leaks is a fairly common practice used by large ISPs. (Note: This information was gathered from discussions on the NANOG mailing list [Nanog-thread-June2016] as well as through private discussions with operators of large ISP networks.)

Routes are tagged on ingress to an AS with communities for origin, including the type of eBGP peer it was learned from (customer, provider or lateral peer), geographic location, etc. The community attributes are carried across the AS with the routes. These communities are used along with additional logic in route policies to determine which routes are to be announced to which eBGP peers and which are to be dropped. In this process, the ISP's AS also ensures that routes learned from a transit-provider or a lateral peer (i.e., non-transit) at an ingress router are not leaked at an egress router to another transit-provider or lateral peer.

Additionally, in many cases, ISP network operators' outbound policies require explicit matches for expected communities before passing routes. This helps ensure that that if an update has been entered into the RIB-in but has missed its ingress community tagging (due to a missing/misapplied ingress policy), it will not be inadvertently leaked.

The above procedure (or a simplified version of it) is also applicable when an AS consists of a single eBGP router. It is recommended that all AS operators SHOULD implement the procedure described above (or similar that is appropriate for their network) to prevent route leaks that they have direct control over.

3.8. Stopgap Solution when Only Origin Validation is Deployed

A stopgap method is described here for detection and mitigation of route leaks for the intermediate phase when OV is deployed but BGP protocol on the wire is unchanged. The stopgap solution can be in the form of construction of a prefix filter list from ROAs. A suggested procedure for constructing such a list comprises of the following steps:

- o ISP makes a list of all the ASes (Cust_AS_List) that are in its customer cone (ISP's own AS is also included in the list). (Some of the ASes in Cust_AS_List may be multi-homed to another ISP and that is OK.)
- o ISP downloads from the RPKI repositories a complete list (Cust_ROA_List) of valid ROAs that contain any of the ASes in Cust_AS_List.
- o ISP creates a list of all the prefixes (Cust_Prfx_List) that are contained in any of the ROAs in Cust_ROA_List.
- o Cust_Prfx_List is the allowed list of prefixes that is permitted by the ISP's AS, and will be forwarded by the ISP to upstream ISPs, customers, and peers.
- o A route for a prefix that is not in Cust_Prfx_List but announced by one of ISP's customers is 'marked' as a potential route leak. Further, the ISP's router SHOULD prefer an alternate route that is Valid (i.e., valid according to origin validation) and 'clean' (i.e., not marked) over the 'marked' route. The alternate route may be from a peer, transit provider, or different customer.

Special considerations regarding the above procedure may be needed for DDoS mitigation service providers. They typically originate or announce a DDoS victim's prefix to their own ISP on a short notice

during a DDoS emergency. Some provisions would need to be made for such cases, and they can be determined with the help of inputs from DDoS mitigation service providers.

For developing a list of all the ASes (Cust_AS_List) that are in the customer cone of an ISP, the AS path based Outbound Route Filter (ORF) technique [I-D.ietf-idr-aspath-orf] can be helpful (see discussion in Section 3.5).

Another technique based on AS_PATH filters is described in [Snijders]. This method is applicable to very large ISPs that have lateral peering. For a pair of such very large ISPs, say A and B, the method depends on ISP A communicating out-of-band (e.g., by email) with ISP B about whether or not it (ISP A) has any transit providers. This out-of-band knowledge enables ISP B to apply suitable AS_PATH filtering criteria for routes involving the presence of ISP A in the path and prevent certain kinds of route leaks (see [Snijders] for details).

4. Security Considerations

This document requires no security considerations. See [I-D.ietf-idr-route-leak-detection-mitigation] for security considerations for the solution for route leaks detection and mitigation.

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. 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>>.

6.2. Informative References

[Anwar] Anwar, R., Niaz, H., Choffnes, D., Cunha, I., Gill, P., and N. Katz-Bassett, "Investigating Interdomain Routing Policies in the Wild", ACM Internet Measurement Conference (IMC), October 2015, <<http://www.cs.usc.edu/assets/007/94928.pdf>>.

- [draft-dickson-sidr-route-leak-solns]
Dickson, B., "Route Leaks -- Proposed Solutions", IETF Internet Draft (expired), March 2012, <<https://tools.ietf.org/html/draft-dickson-sidr-route-leak-solns-01>>.
- [draft-kunzinger-idrp-ISO10747-01]
Kunzinger, C., "Inter-Domain Routing Protocol (IDRP)", IETF Internet Draft (expired), November 1994, <<https://tools.ietf.org/pdf/draft-kunzinger-idrp-ISO10747-01.pdf>>.
- [Gao] Gao, L. and J. Rexford, "Stable Internet routing without global coordination", IEEE/ACM Transactions on Networking, December 2001, <<http://www.cs.princeton.edu/~jrex/papers/sigmetrics00.long.pdf>>.
- [Gao-Rexford]
Freedman, M., "Interdomain Routing Policy", Princeton University COS 461 Lecture Notes; Slides 25-27, Spring 2011, <<http://www.cs.princeton.edu/courses/archive/spr11/cos461/docs/lec17-bgp-policy.ppt>>.
- [Gill] Gill, P., Schapira, M., and S. Goldberg, "A Survey of Interdomain Routing Policies", ACM SIGCOMM Computer Communication Review, January 2014, <<https://www.cs.bu.edu/~goldbe/papers/survey.pdf>>.
- [Giotsas] Giotsas, V. and S. Zhou, "Valley-free violation in Internet routing - Analysis based on BGP Community data", IEEE ICC 2012, June 2012.
- [I-D.ietf-idr-aspath-orf]
Hares, S. and K. Patel, "AS Path Based Outbound Route Filter for BGP-4", draft-ietf-idr-aspath-orf-13 (work in progress), December 2016.
- [I-D.ietf-idr-bgp-open-policy]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention using Roles in Update and Open messages", draft-ietf-idr-bgp-open-policy-03 (work in progress), June 2018.

- [I-D.ietf-idr-route-leak-detection-mitigation]
Sriram, K., Montgomery, D., Dickson, B., Patel, K., and A. Robachevsky, "Methods for Detection and Mitigation of BGP Route Leaks", draft-ietf-idr-route-leak-detection-mitigation-08 (work in progress), March 2018.
- [Luckie] Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., and kc. claffy, "AS Relationships, Customer Cones, and Validation", IMC 2013, October 2013, <<http://www.caida.org/~amogh/papers/asrank-IMC13.pdf>>.
- [Nanog-thread-June2016]
"Intra-AS messaging for route leak prevention", NANOG Email List - Discussion Thread , June 2016, <<http://mailman.nanog.org/pipermail/nanog/2016-June/thread.html#86348>>.
- [NIST-800-54]
Kuhn, D., Sriram, K., and D. Montgomery, "Border Gateway Protocol Security", NIST Special Publication 800-54, July 2007, <<http://csrc.nist.gov/publications/nistpubs/800-54/SP800-54.pdf>>.
- [proceedings-sixth-ietf]
Gross, P., "Proceedings of the April 22-24, 1987 Internet Engineering Task Force", April 1987, <<https://www.ietf.org/proceedings/06.pdf>>.
- [RFC1105-obsolete]
Lougheed, K. and Y. Rekhter, "A Border Gateway Protocol (BGP)", IETF RFC (obsolete), June 1989, <<https://tools.ietf.org/html/rfc1105>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<https://www.rfc-editor.org/info/rfc7454>>.
- [RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <<https://www.rfc-editor.org/info/rfc7908>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.

[Snijders]

Snijders, J., "Practical everyday BGP filtering with AS_PATH filters: Peer Locking", NANOG-47 Chicago, IL, USA, June 2016, <https://www.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf>.

[Wijchers]

Wijchers, B. and B. Overeinder, "Quantitative Analysis of BGP Route Leaks", RIPE-69, November 2014, <<https://ripe69.ripe.net/presentations/157-RIPE-69-Routing-WG.pdf>>.

Acknowledgements

The authors wish to thank Jared Mauch, Jeff Haas, Job Snijders, Warren Kumari, Amogh Dhamdhere, Jakob Heitz, Geoff Huston, Randy Bush, Alexander Azimov, Ruediger Volk, Sue Hares, Wes George, Job Snijders, Chris Morrow, Sandy Murphy, Danny McPherson, and Eric Osterweil for comments, suggestions, and critique. The authors are also thankful to Padma Krishnaswamy, Oliver Borchert, and Okhee Kim for their review and comments.

Contributors

The following people made significant contributions to this document and should be considered co-authors:

Alexander Azimov
Qrator Labs
Email: aa@qrator.net

Brian Dickson
Independent
Email: brian.peter.dickson@gmail.com

Doug Montgomery
USA National Institute of Standards and Technology
Email: dougm@nist.gov

Keyur Patel
Arccus
Email: keyur@arccus.com

Andrei Robachevsky
Internet Society
Email: robachevsky@isoc.org

Eugene Bogomazov
Qrator Labs
Email: eb@qrator.net

Randy Bush
Internet Initiative Japan
Email: randy@psg.com

Author's Address

Kotikalapudi Sriram (editor)
USA National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Email: ksriram@nist.gov