

IDR and SIDR
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2016

K. Sriram
D. Montgomery
US NIST
B. Dickson

K. Patel
Cisco
A. Robachevsky
Internet Society
March 14, 2016

Methods for Detection and Mitigation of BGP Route Leaks
draft-ietf-idr-route-leak-detection-mitigation-02

Abstract

In [[I-D.ietf-grow-route-leak-problem-definition](#)], the authors have provided a definition of the route leak problem, and also enumerated several types of route leaks. In this document, we first examine which of those route-leak types are detected and mitigated by the existing origin validation (OV) [[RFC 6811](#)]. It is recognized that OV offers a limited detection and mitigation capability against route leaks. This document proposes an enhancement that significantly extends the route-leak detection and mitigation capabilities of BGP. The solution involves carrying a per-hop route-leak protection (RLP) field in BGP updates. The RLP field is proposed be carried in an optional transitive path attribute. The solution is meant to be initially implemented as an enhancement of BGP without requiring BGPsec [[I-D.ietf-sidr-bgpsec-protocol](#)]. However, when BGPsec is deployed in the future, the solution can be incorporated in BGPsec, enabling cryptographic protection for the RLP field. That would be one way of implementing the proposed solution in a secure way. It is not claimed that the solution detects all possible types of route leaks but it detects several types, especially considering some significant route-leak occurrences that have been observed in recent years. The document also includes a stopgap method for detection and mitigation of route leaks for an intermediate phase when OV is deployed but BGP protocol on the wire is unchanged.

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 <http://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 September 15, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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](#) [3](#)
- [2. Related Prior Work](#) [3](#)
- [3. Mechanisms for Detection and Mitigation of Route Leaks](#) . . . [4](#)
 - [3.1. Route-Leak Protection \(RLP\) Field Encoding by Sending Router](#) [6](#)
 - [3.2. Recommended Actions at a Receiving Router for Detection of Route Leaks](#) [8](#)
 - [3.3. Possible Actions at a Receiving Router for Mitigation](#) . . [9](#)
- [4. Stopgap Solution when Only Origin Validation is Deployed](#) . . [9](#)
- [5. Design Rationale and Discussion](#) [10](#)
 - [5.1. Is route-leak solution without cryptographic protection a serious attack vector?](#) [10](#)
 - [5.2. Combining results of route-leak detection, OV and BGPsec validation for path selection decision](#) [12](#)
 - [5.3. Are there cases when valley-free violations can be considered legitimate?](#) [12](#)
 - [5.4. Comparison with other methods, routing security BCP](#) . . . [13](#)
- [6. Summary](#) [13](#)
- [7. Security Considerations](#) [14](#)
- [8. IANA Considerations](#) [14](#)
- [9. Acknowledgements](#) [14](#)

10.	References	14
10.1.	Normative References	14
10.2.	Informative References	14
Authors' Addresses	18

[1.](#) Introduction

In [[I-D.ietf-grow-route-leak-problem-definition](#)], the authors have provided a definition of the route leak problem, and also enumerated several types of route leaks. In this document, we first examine which of those route-leak types are detected and mitigated by the existing Origin Validation (OV) [[RFC6811](#)] method. OV and BGPsec path validation [[I-D.ietf-sidr-bgpsec-protocol](#)] together offer mechanisms to protect against re-originations and hijacks of IP prefixes as well as man-in-the-middle (MITM) AS path modifications. Route leaks (see [[I-D.ietf-grow-route-leak-problem-definition](#)] and references cited at the back) are another type of vulnerability in the global BGP routing system against which OV offers only partial protection. BGPsec (i.e. path validation) provides cryptographic protection for some aspects of BGP update messages, but in its current form BGPsec doesn't offer any protection against route leaks.

For the types of route leaks enumerated in [[I-D.ietf-grow-route-leak-problem-definition](#)], where the current OV method doesn't offer a solution, this document proposes an enhancement that would significantly extend the detection and mitigation capabilities of BGP. The solution involves carrying a per-hop route-leak protection (RLP) field in BGP updates. The RLP field is proposed be carried in an optional transitive path attribute. The solution is meant to be initially implemented as an enhancement of BGP without requiring BGPsec. However, when BGPsec is deployed in the future, the solution can be incorporated in BGPsec, enabling cryptographic protection for the RLP field. That would be one way of implementing the proposed solution in a secure way. It is not claimed that the solution detects all possible types of route leaks but it detects several types, especially considering some significant route-leak occurrences that have been observed in recent years. The document also includes a stopgap method (in [Section 4](#)) for detection and mitigation of route leaks for an intermediate phase when OV is deployed but BGP protocol on the wire is unchanged.

[2.](#) Related Prior Work

The basic idea and mechanism embodied in the proposed solution 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 said 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. Dickson developed the initial Internet draft of these mechanisms in a BGPsec context; see [[draft-dickson-sidr-route-leak-solns](#)]. The draft expired in 2012. [[draft-dickson-sidr-route-leak-solns](#)] defined neighbor relationships on a per link basis, but in the current draft the relationship is encoded per prefix, as routes for prefixes with different business models are often 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 the current draft 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 3.1](#)).

3. Mechanisms for Detection and Mitigation of Route Leaks

Referring to the enumeration of route leaks discussed in [[I-D.ietf-grow-route-leak-problem-definition](#)], Table 1 summarizes the route-leak detection capability offered by OV and BGPsec for different types of route leaks. (Note: Prefix filtering is not considered here in this table. Please see [Section 4](#).)

A detailed explanation of the contents of Table 1 is as follows. It is readily observed that route leaks of Types 1, 2, 3, and 4 are not detected by OV or BGPsec in its current form. Clearly, Type 5 route leak involves re-origination or hijacking, and hence can be detected by OV. In the case of Type 5 route leak, there would be no existing ROAs to validate a re-originated prefix or more specific, but instead a covering ROA would normally exist with the legitimate AS, and hence the update will be considered Invalid by OV.

Type of Route Leak	Current State of Detection Coverage
Type 1: Hairpin Turn with Full Prefix	Neither OV nor BGPsec (in its current form) detects Type 1.
Type 2: Lateral ISP-ISP-ISP Leak	Neither OV nor BGPsec (in its current form) detects Type 2.
Type 3: Leak of Transit-Provider Prefixes to Peer	Neither OV nor BGPsec (in its current form) detects Type 3.
Type 4: Leak of Peer Prefixes to Transit Provider	Neither OV nor BGPsec (in its current form) detects Type 4.
Type 5: Prefix Re-Origination with Data Path to Legitimate Origin	OV detects Type 5.
Type 6: Accidental Leak of Internal Prefixes and More Specifics	For internal prefixes never meant to be routed on the Internet, OV helps detect their leak; they might either have no covering ROA or have an AS0-ROA to always filter them. In the case of accidental leak of more specifics, OV may offer some detection due to ROA maxLength.

Table 1: Examination of Route-Leak Detection Capability of Origin Validation and Current BGPsec Path Validation

In the case of Type 6 leaks involving internal prefixes that are not meant to be routed in the Internet, they are likely to be detected by OV. That is because such prefixes might either have no covering ROA or have an AS0-ROA to always filter them. In the case of Type 6 leaks that are due to accidental leak of more specifics, they may be detected due to violation of ROA maxLength. BGPsec (i.e. path validation) in its current form does not detect Type 6. However, route leaks of Type 6 are least problematic due to the following reasons. In the case of leak of more specifics, the offending AS is itself the legitimate destination of the leaked more-specific prefixes. Hence, in most cases of this type, the data traffic is neither misrouted nor denied service. Also, leaked announcements of Type 6 are short-lived and typically withdrawn quickly following the announcements. Further, the MaxPrefix limit may kick-in in some

receiving routers and that helps limit the propagation of sometimes large number of leaked routes of Type 6.

Realistically, BGPsec may take a much longer time being deployed than OV. Hence solution proposals for route leaks should consider both scenarios: (A) OV only (without BGPsec) and (B) OV plus BGPsec. Assuming an initial scenario A, and based on the above discussion and Table 1, it is evident that in our proposed solution method, we need to focus primarily on route leaks of Types 1, 2, 3, and 4. In [Section 3.1](#) and [Section 3.2](#), we describe a simple addition to BGP that facilitates detection of route leaks of Types 1, 2, 3, and 4. The simple addition involves a Route-Leak Protection (RLP) field, which is carried in an optional transitive path attribute in BGP. When BGPsec is deployed, the RLP field will be accommodated in the existing Flags field (see [[I-D.ietf-sidr-bgpsec-protocol](#)]) which is cryptographically protected under path signatures.

[3.1](#). Route-Leak Protection (RLP) Field Encoding by Sending Router

The key principle is that, in the event of a route leak, a receiving router in a transit-provider AS (e.g. referring to Figure 1, ISP2 (AS2) router) should be able to detect from the update message that its customer AS (e.g. AS3 in Figure 1) SHOULD NOT have forwarded the update (towards the transit-provider AS). This means that at least one of the ASes in the AS path of the update has indicated that it sent the update to its customer or lateral (i.e. non-transit) peer, but forbade any subsequent 'Up' forwarding (i.e. from a customer AS to its transit-provider AS). For this purpose, a Route-Leak Protection (RLP) field to be set by a sending router is proposed to be used for each AS hop.

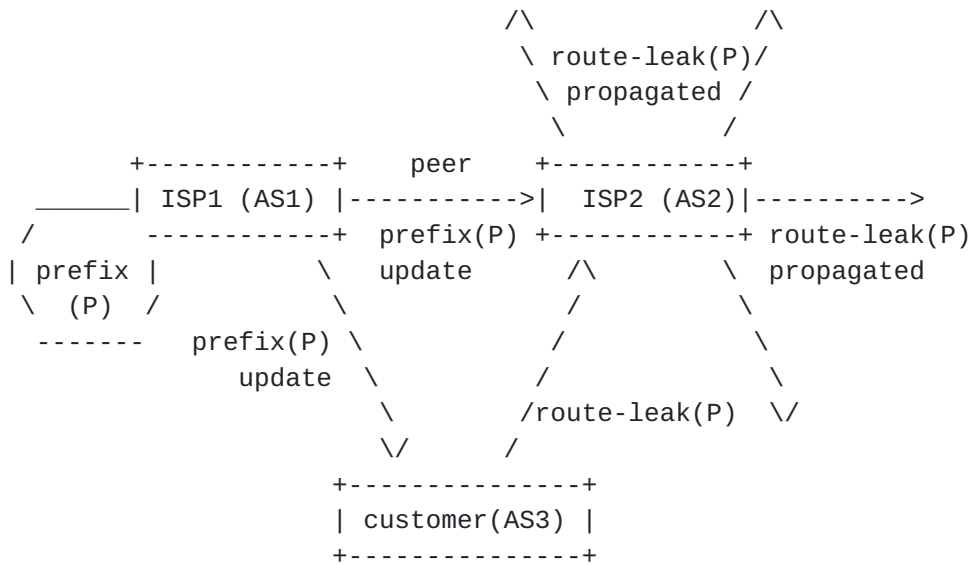


Figure 1: Illustration of the basic notion of a route leak.

For the purpose of route-leak detection and mitigation proposed in this document, the RLP field value SHOULD be set to one of two values as follows:

- o 00: This is the default value (i.e. "nothing specified"),
- o 01: This is the 'Do not Propagate Up or Lateral' indication; sender indicating that the route SHOULD NOT be forwarded 'Up' towards a transit-provider AS or to a lateral (i.e. non-transit) peer AS.

The RLP indications SHOULD be set on a per prefix and per neighbor AS basis. This is because updates for prefixes with different business models are often sent over the same link between ASes.

There are two different scenarios when a sending AS SHOULD set the '01' indication in an update: (1) when sending the update to a customer AS, and (2) when sending the update to a lateral peer (i.e. non-transit) AS. In essence, in both scenarios, the intent of '01' indication is that the neighbor AS and any receiving AS along the subsequent AS path SHOULD NOT forward the update 'Up' towards its (receiving AS's) transit-provider AS or laterally towards its peer (i.e. non-transit) AS. When sending an update 'Up' to a transit-provider AS, the RLP encoding should be set to the default value of '00'. When a sending AS sets the RLP encoding to '00', it is indicating to the receiving AS that the update can be propagated in any direction (i.e. towards transit-provider, customer, or lateral peer). This two-state specification in the RLP field can be shown to

work for detection and mitigation of route leaks of Types 1, 2, 3, and 4 which are the focus here (see [Section 3.2](#) and [Section 3.3](#)). The '10' and '11' values in the RLP field (assuming that two bits are used to encode it) are currently unassigned, and reserved for possible future use.

The proposed RLP encoding SHOULD be carried in BGP-4 [[RFC4271](#)] updates in an optional transitive path attribute. In BGPsec enabled routers, the RLP encoding SHOULD be accommodated in the existing Flags field in BGPsec updates. The Flags field is part of the Secure_Path Segment in BGPsec updates [[I-D.ietf-sidr-bgpsec-protocol](#)]. It is one octet long, and one Flags field is available for each AS hop, and currently only the first bit is used in BGPsec. So there are 7 bits that are currently unused in the Flags field. Two (or more if needed) of these bits can be designated for the RLP field. Since the BGPsec protocol specification requires a sending AS to include the Flags field in the data that are signed over, the RLP field for each hop (assuming it would be part of the Flags field) will be protected under the sending AS's signature.

[3.2.](#) Recommended Actions at a Receiving Router for Detection of Route Leaks

We provide here an example set of receiver actions that work to detect and mitigate route leaks of Types 1, 2, 3, and 4. This example algorithm serves as a proof of concept. However, other receiver algorithms or procedures can be designed (based on the same sender specification as in [Section 3.1](#)) and may perform with greater efficacy, and are by no means excluded.

A recommended receiver algorithm for detecting a route leak is as follows:

A receiving router SHOULD mark an update as a 'Route Leak' if ALL of the following conditions hold true:

1. The update is received from a customer or lateral peer AS.
2. The update has the RLP field set to '01' (i.e. 'Do not Propagate Up or Lateral') indication for one or more hops (excluding the most recent) in the AS path.

The reason for stating "excluding the most recent" in the above algorithm is as follows. An ISP should look at RLP values set by ASes preceding the immediate sending AS in order to ascertain a leak. The receiving router already knows that the most recent hop in the update is from its customer or lateral-peer AS to itself, and it does

not need to rely on the RLP field value set by said AS for detection of route leaks.

If the RLP encoding is secured by BGPsec (see [Section 3.1](#)) and hence protected against tampering by intermediate ASes, then there would be added certainty in the route-leak detection algorithm described above (see discussions in [Section 5.1](#) and [Section 5.2](#)).

A receiving router expects the RLP field value for any hop in the AS path to be either 00 or 01. However, if a different value (say, 10 or 11) is found in the RLP field, then an error condition will get flagged, and any further action is TBD.

[3.3.](#) Possible Actions at a Receiving Router for Mitigation

After applying the above detection algorithm, a receiving router may use any policy-based algorithm of its own choosing to mitigate any detected route leaks. An example receiver algorithm for mitigating a route leak is as follows:

- o If an update from a customer or lateral peer AS is marked as a 'Route Leak', then the receiving router SHOULD prefer an alternate unmarked route if available.
- o If no alternate unmarked route is available, then the route marked as a 'Route Leak' MAY be accepted.

A basic principle here is that the presence of '01' value in the RLP field corresponding to one or more AS hops in the AS path of an update coming from a customer AS informs a receiving router in a transit-provider AS that a route leak is likely occurring. The transit-provider AS then overrides the "prefer customer route" policy, and instead prefers an alternate 'clean' route learned from another customer, a lateral peer, or a transit provider over the 'marked' route from the customer in question.

[4.](#) Stopgap Solution when Only Origin Validation is Deployed

Here we describe a stopgap method 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 with regard to 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 [[draft-ietf-idr-aspath-orf](#)] can be helpful (see discussion in [Section 5.4](#)).

5. Design Rationale and Discussion

In this section, we will try to provide design justifications for the methodology specified in [Section 3](#), and also answer some questions that are anticipated or have been raised in IETF IDR/SIDR meetings.

5.1. Is route-leak solution without cryptographic protection a serious attack vector?

It has been asked if a route-leak solution without BGPsec, i.e. when RLP bits are not protected, can turn into a serious 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 bits 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 '01' (i.e. 'Do not Propagate Up or Lateral') to '00' (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](#)]. 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 '01' (i.e. 'Do not Propagate Up or Lateral') when it should be set to '00' (default). This would result in a route being mis-detected and marked as a route leak. By default RLP encoding is set to '00', 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 '00' to '01' 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.

5.2. 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 5.1](#)).

5.3. 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. In particular, 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 violations is per-prefix routing policies, i.e. routes for prefixes with different business models are often 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). The authors in [[Anwar](#)] also observe, "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 in the methodology described in this document is that the Route-Leak Protection (RLP) encoding is per prefix. So

the proposed 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 configuring policies (including per prefix policies, if applicable), and make proper use of the RLP indications on a per prefix basis. When said 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 said leaks would be much less likely to propagate from one large participating ISP to another.

5.4. Comparison with other methods, routing security BCP

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, AS operators signal for each route propagated, if it SHOULD NOT be subsequently propagated to a transit provider or peer.

AS path based Outbound Route Filter (ORF) described in [draft-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 business models are often 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.

6. Summary

It should be emphasized once again that the proposed route-leak detection method using the RLP encoding is not intended to cover all forms of route leaks. However, we feel that the solution covers several important types of route leaks, especially considering some significant route-leak attacks or occurrences that have been frequently observed in recent years. The solution can be implemented

in BGP without necessarily tying it to BGPsec. The proposed solution without BGPsec can detect and mitigate accidental route leaks, and the same with BGPsec can detect and mitigate both accidental and malicious route leaks. Carrying the proposed RLP encoding in an optional transitive path attribute in BGP is proposed, but in order to prevent abuse, the RLP encoding would require cryptographic protection. Incorporating the RLP encoding in the BGPsec Flags field is one way of implementing it securely using an already existing protection mechanism provided in BGPsec path signatures.

7. Security Considerations

The proposed Route-Leak Protection (RLP) field requires cryptographic protection in order to prevent malicious route leaks. Since it is proposed that the RLP field be included in the Flags field in the Secure_Path Segment in BGPsec updates, the cryptographic security mechanisms in BGPsec are expected to also apply to the RLP field. The reader is therefore directed to the security considerations provided in [[I-D.ietf-sidr-bgpsec-protocol](#)].

8. IANA Considerations

No updates to the registries are suggested by this document.

9. Acknowledgements

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

10. References

10.1. 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, <<http://www.rfc-editor.org/info/rfc4271>>.

10.2. 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>>.
- [Cowie2010] Cowie, J., "China's 18 Minute Mystery", Dyn Research/Renesys Blog, November 2010, <<http://research.dyn.com/2010/11/chinas-18-minute-mystery/>>.
- [Cowie2013] Cowie, J., "The New Threat: Targeted Internet Traffic Misdirection", Dyn Research/Renesys Blog, November 2013, <<http://research.dyn.com/2013/11/mitm-internet-hijacking/>>.
- [[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-ietf-idr-aspath-orf](#)] Patel, K. and S. Hares, "AS path Based Outbound Route Filter for BGP-4", IETF Internet Draft (expired), August 2007, <<https://tools.ietf.org/html/draft-ietf-idr-aspath-orf-09>>.
- [[draft-kunzinger-idrp-IS010747-01](#)] Kunzinger, C., "Inter-Domain Routing Protocol (IDRP)", IETF Internet Draft (expired), November 1994, <<https://tools.ietf.org/pdf/draft-kunzinger-idrp-IS010747-01.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>>.
- [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.
- [Hiran] Hiran, R., Carlsson, N., and P. Gill, "Characterizing Large-scale Routing Anomalies: A Case Study of the China Telecom Incident", PAM 2013, March 2013, <<http://www3.cs.stonybrook.edu/~phillipa/papers/CTelecom.html>>.
- [Huston2012] Huston, G., "Leaking Routes", March 2012, <<http://labs.apnic.net/blabs/?p=139/>>.
- [Huston2014] Huston, G., "What's so special about 512?", September 2014, <<http://labs.apnic.net/blabs/?p=520/>>.
- [I-D.ietf-grow-route-leak-problem-definition] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", [draft-ietf-grow-route-leak-problem-definition-04](#) (work in progress), February 2016.
- [I-D.ietf-sidr-bgpsec-protocol] Lepinski, M., "BGPsec Protocol Specification", [draft-ietf-sidr-bgpsec-protocol-14](#) (work in progress), December 2015.
- [Kapela-Piloso] Pilosov, A. and T. Kapela, "Stealing the Internet: An Internet-Scale Man in the Middle Attack", DEFCON-16 Las Vegas, NV, USA, August 2008, <<https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosoov-kapela.pdf>>.
- [Kephart] Kephart, N., "Route Leak Causes Amazon and AWS Outage", ThousandEyes Blog, June 2015, <<https://blog.thousandeyes.com/route-leak-causes-amazon-and-aws-outage>>.
- [Khare] Khare, V., Ju, Q., and B. Zhang, "Concurrent Prefix Hijacks: Occurrence and Impacts", IMC 2012, Boston, MA, November 2012, <<http://www.cs.arizona.edu/~bzhang/paper/12-imc-hijack.pdf>>.

[Labovitz]

Labovitz, C., "Additional Discussion of the April China BGP Hijack Incident", Arbor Networks IT Security Blog, November 2010,
<<http://www.arbornetworks.com/asert/2010/11/additional-discussion-of-the-april-china-bgp-hijack-incident/>>.

[LRL]

Khare, V., Ju, Q., and B. Zhang, "Large Route Leaks", Project web page, 2012,
<<http://nrl.cs.arizona.edu/projects/lrsl-events-from-2003-to-2009/>>.

[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>>.

[Madory]

Madory, D., "Why Far-Flung Parts of the Internet Broke Today", Dyn Research/Renesys Blog, September 2014,
<<http://research.dyn.com/2014/09/why-the-internet-broke-today/>>.

[Mauch]

Mauch, J., "BGP Routing Leak Detection System", Project web page, 2014,
<<http://puck.nether.net/bgp/leakinfo.cgi/>>.

[Mauch-nanog]

Mauch, J., "Detecting Routing Leaks by Counting", NANOG-41 Albuquerque, NM, USA, October 2007,
<<https://www.nanog.org/meetings/nanog41/presentations/mauch-lightning.pdf>>.

[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>>.

[Paseka]

Paseka, T., "Why Google Went Offline Today and a Bit about How the Internet Works", CloudFare Blog, November 2012,
<<http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about/>>.

[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>>.

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

[RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", [BCP 194](#), [RFC 7454](#), DOI 10.17487/RFC7454, February 2015, <<http://www.rfc-editor.org/info/rfc7454>>.

[Toonk] Toonk, A., "What Caused Today's Internet Hiccup", August 2014, <<http://www.bgpmn.net/what-caused-todays-internet-hiccup/>>.

[Toonk2015-A]

Toonk, A., "What caused the Google service interruption", March 2015, <<http://www.bgpmn.net/what-caused-the-google-service-interruption/>>.

[Toonk2015-B]

Toonk, A., "Massive route leak causes Internet slowdown", June 2015, <<http://www.bgpmn.net/massive-route-leak-cause-internet-slowdown/>>.

[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>>.

[Zmijewski]

Zmijewski, E., "Indonesia Hijacks the World", Dyn Research/Renesys Blog, April 2014, <<http://research.dyn.com/2014/04/indonesia-hijacks-world/>>.

Authors' Addresses

Kotikalapudi Sriram
US NIST

Email: ksriram@nist.gov

Doug Montgomery
US NIST

Email: dougmont@nist.gov

Brian Dickson

Email: brian.peter.dickson@gmail.com

Keyur Patel
Cisco

Email: keyupate@cisco.com

Andrei Robachevsky
Internet Society

Email: robachevsky@isoc.org

