

IDR and SIDR
Internet-Draft
Intended status: Standards Track
Expires: January 6, 2016

K. Sriram
D. Montgomery
US NIST
B. Dickson
Twitter, Inc.
July 5, 2015

Methods for Detection and Mitigation of BGP Route Leaks
draft-sriram-idr-route-leak-detection-mitigation-01

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](#)] and BGPSEC path validation [[I-D.ietf-sidr-bgpsec-protocol](#)]. Where the current OV and BGPSEC protocols don't offer a solution, this document suggests an enhancement that would extend the route-leak detection and mitigation capability of BGPSEC. The solution can be implemented in BGP without necessarily tying it to BGPSEC. Incorporating the solution in BGPSEC is one way of implementing it in a secure way. We do not claim to have provided a solution for all possible types of route leaks, but the solution covers several, especially considering some significant route-leak attacks or occurrences that have been observed in recent years. The document also includes a stopgap method for detection and mitigation of route leaks for the phase when BGPSEC (path validation) is not yet deployed but only origin validation is deployed.

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 January 6, 2016.

Copyright Notice

Copyright (c) 2015 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	2
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
3.2.1.	Recommended Actions at a Receiving Router when the Sender is a Customer	8
3.2.2.	Recommended Actions at a Receiving Router when the Sender is a Peer	9
3.3.	Possible Actions at a Receiving Router for Mitigation . .	10
4.	Stopgap Solution when Only Origin Validation is Deployed . .	10
5.	Design Rationale and Discussion	11
5.1.	Is route-leak solution without BGPSEC a serious attack vector?	11
5.2.	Comparison with other methods, routing security BCP . . .	12
6.	Summary	12
7.	Security Considerations	13
8.	IANA Considerations	13
9.	Acknowledgements	13
10.	References	13
10.1.	Normative References	13
10.2.	Informative References	13
	Authors' Addresses	17

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](#)] and BGPSEC path validation [[I-D.ietf-sidr-bgpsec-protocol](#)]. For the rest of this document, we use the term BGPSEC as synonymous with path validation. The BGPSEC protocol provides cryptographic protection for some aspects of BGP update messages. OV and BGPSEC together offer mechanisms to protect against mis-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 and BGPSEC so far offer only partial protection.

For the types of route leaks enumerated in [[I-D.ietf-grow-route-leak-problem-definition](#)], where the current OV and BGPSEC protocols don't offer a solution, this document suggests an enhancement that would extend the detection and mitigation capability of BGPSEC. The solution can be implemented in BGP without necessarily tying it to BGPSEC. Incorporating the solution in BGPSEC is one way of implementing it in a secure way. We do not claim to provide a solution for all possible types of route leaks, but the solution covers several relevant types, especially considering some significant route-leak occurrences that have been observed frequently in recent years. The document also includes (in [Section 4](#)) a stopgap method for detection and mitigation of route leaks for the phase when BGPSEC (path validation) is not yet deployed but only origin validation is deployed.

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, 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, c2p, p2p, 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 in

encoded per prefix, as 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, 5, 6, and 7 are not detected by OV or even by BGPSEC. Type 2 route leak involves changing a prefix to a subprefix (i.e. more specific); such a modified update will fail BGPSEC checks. Clearly, Type 3 route leak involves mis-origination or hijacking, and hence can be detected by OV. In the case of Type 3 route leak, there would be no existing ROAs to validate a re-originated prefix or subprefix, 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	Detection Coverage and Comments
Type 1: U-Turn with Full Prefix	Neither OV nor BGPSEC (in its current form) detects Type 1.
Type 2: U-Turn with More Specific Prefix	In OV, the ROA maxLength may offer detection of Type 2 in some cases; BGPSEC (in its current form) always detects Type 2.
Type 3: Prefix Mis-Origination with Data Path to Legitimate Origin	OV by itself detects Type 3; BGPSEC does not detect Type 3.
Type 4: Leak of Internal Prefixes and Accidental Deaggregation	For internal prefixes never meant to be seen (i.e. 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 deaggregation, OV may offer some detection due to ROA maxLength. BGPSEC does not catch Type 4.
Type 5: Lateral ISP-ISP-ISP Leak	Neither OV nor BGPSEC (in its current form) detects Type 5.
Type 6: Leak of Provider Prefixes to Peer	Neither OV nor BGPSEC (in its current form) detects Type 6.
Type 7: Leak of Peer Prefixes to Provider	Neither OV nor BGPSEC (in its current form) detects Type 7.

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

In the case of Type 4 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 4 leaks that are due to accidental deaggregation, they may be detected due to violation of ROA maxLength. BGPSEC does not catch Type 4. However, route leaks of Type 4 are least problematic due to the

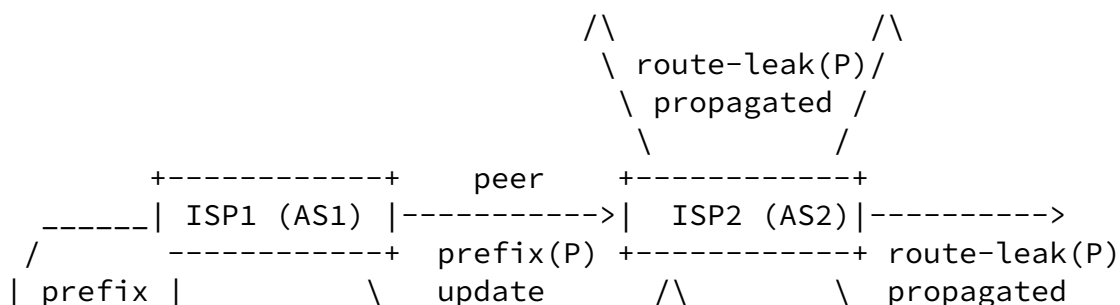
following reasons. In the case of accidental deaggregation, the offending AS is itself the legitimate destination of the leaked more-specific prefixes. Hence, in most cases of this type, the data traffic is neither misrouted nor denied service. Also, leaked announcements of Type 4 are short-lived and 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 4.

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, 5, 6, and 7. 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, 5, 6, and 7. 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 provider AS (e.g. referring to Figure 1, ISP2 (AS2) router) should be able to detect from the prefix-update that its customer AS (e.g. AS3 in Figure 1) SHOULD NOT have forwarded the update (towards the 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 peer AS, but forbade any subsequent 'Up' forwarding (i.e. from a customer AS to its 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.



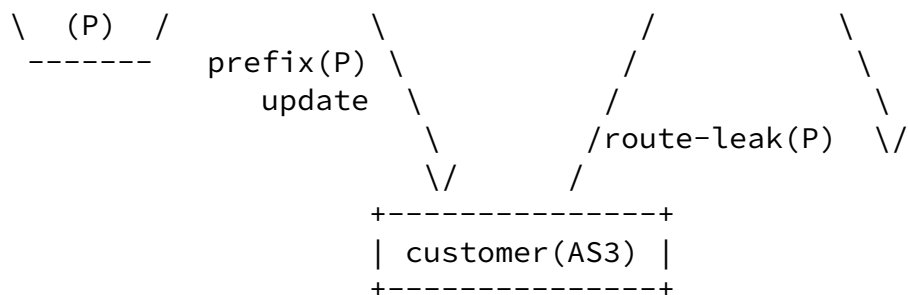


Figure 1: Illustration of the basic notion of a route leak.

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' indication; sender indicating that the prefix-update SHOULD NOT be forwarded 'Up' towards a provider AS.

There are two different scenarios when a sending AS SHOULD set the '01' indication in a prefix-update: (1) when sending the prefix-update to a customer AS, and (2) to let a peer AS know not to forward the prefix-update 'Up' towards a provide AS. In essence, in both scenarios, the intent of '01' indication is that any receiving AS along the subsequent AS path SHOULD NOT forward the prefix-update 'Up' towards its (receiving AS's) provider AS.

One may argue for additional RLP indications: for example, '10' to specify 'Propagate to Customers Only', and possibly '11' to signal 'Do Not Propagate' (i.e. NO_EXPORT). But in the interest of keeping the methodology simple, the choice of two RLP field values as defined above (00 - default, and 01 - 'Do not Propagate Up') is all that is needed. This two-state specification in the RLP field can be shown to work for detection and mitigation of route leaks of Types 1, 2, 5, 6, and 7, which are the focus here (see [Section 3.2](#) and [Section 3.3](#)).

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

The recommended receiver actions differ slightly depending on whether the update is received from a customer or a peer. When detecting route leaks of Type 1, 2, and 7, the receiving router is dealing with a customer as the sender. When detecting route leaks of Type 5 and 6, the receiving router is dealing with a peer as the sender.

[3.2.1.](#) Recommended Actions at a Receiving Router when the Sender is a Customer

We provide here an example set of receiver actions that work to detect and mitigate route leaks of Types 1, 2, and 7. 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 AS.
2. It is Valid in accordance with the Origin Validation (OV) and BGPSEC protocols. (Note: BGPSEC validation is not applicable if update is not signed).

3. The update has the RLP field set to '01' (i.e. 'Do not Propagate Up') 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. The provider AS already knows that the most recent hop in the update is from its customer AS to itself, and it does not need to rely on the RLP field value set by the customer for detection of route leaks.

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.

3.2.2. Recommended Actions at a Receiving Router when the Sender is a Peer

The sender and receiver actions described in [Section 3.1](#) and [Section 3.2.1](#) clearly help detect and mitigate route leaks of Types 1, 2, and 7. With a slightly modified interpretation of the RLP encoding on the receiver side, they can be extended to detect lateral ISP-ISP-ISP route leaks (Type 5) as well as leaks of provider prefixes to peer (Type 6). (Note: RLP encoding procedure by sending routers remains the same as described in [Section 3.1](#).)

A recommended receiver algorithm for an ISP to detect a route leak of either Type 5 or Type 6 is as follows:

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

1. The update is received from a lateral ISP peer.
2. It is Valid in accordance with the Origin Validation (OV) and BGPSEC protocols. (Note: BGPSEC validation is not applicable if update is not signed).
3. The update has the RLP field set to '01' indication for one or more hops (excluding the most recent) in the AS path.

In the above algorithm, the receiving AS interprets the '01' indication slightly strongly (i.e. stronger than in [Section 3.2.1](#)) to mean "the update SHOULD NOT have been propagated laterally to a peer ISP like me either". The rationale here is based on the fact that settlement-free ISP peers accept only customer prefix-routes from

each other. The receiving AS applies the logic that if a preceding AS (excluding the most recent) set '01' indication, it means that the

update was sent to a peer or a customer by the (preceding) AS, and the update should not be traversing a lateral peer-to-peer link subsequently.

[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 AS is marked as a Route-Leak, then the receiving router SHOULD prefer a Valid signed update from a peer or an upstream provider over the customer's update.

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 provider AS that a route leak is likely occurring. The provider AS then overrides the "prefer customer route" policy, and instead prefers a route learned from a peer or another upstream provider over the customer's route.

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

During a phase when BGPSEC path validation has not yet been deployed but only origin validation has been deployed, it would be good have a stopgap solution for route leaks. 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.

Sriram, et al.

Expires January 6, 2016

[Page 10]

Internet-Draft

Route Leak Detection and Mitigation

July 2015

- o Any prefix not in Cust_Prfx_List but announced by any of the ISP's customers is marked as a potential route leak. Then the ISP's router SHOULD prefer a Valid (i.e. valid according to origin validation) and 'not marked' update from a peer or an upstream provider over the customer's marked update for that prefix.

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 5.2](#)).

[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 anticipated questions.

[5.1.](#) Is route-leak solution without BGPSEC 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. That 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 BGPSEC. 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') to '00' (default) wherever applicable. It is true that in that case a receiving router would not be able to detect the leak for the specific prefix-route 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](#)] and AS path filters [[draft-ietf-idr-aspath-orf](#)]. If some

malicious leaks go undetected (for RLP without BGPSEC) that is possibly a small price to pay for the ability to detect the bulk of route leaks that are accidental.

[5.2.](#) 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 prefix-route propagated, if it SHOULD NOT be subsequently propagated to a provider/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 prefixes with different business models

are often sent over the same link between ASes. Also, the success of it depends on 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 to their provider ISP(s).

[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 malicious route leaks as well. 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, Andrei Robachevsky, Sue Hares, Keyur Patel, 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., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

10.2. Informative References

[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/>>.

Sriram, et al.

Expires January 6, 2016

[Page 13]

Internet-Draft

Route Leak Detection and Mitigation

July 2015

[[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, <<http://www0.cs.ucl.ac.uk/staff/V.Giotsas/files/giotsas.icc.2012.pdf>>.

[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., E. Osterweil, and B. Dickson "Problem Definition and Classification Route Leaks", [draft-ietf-grow-route-leak-problem-](#)

[definition-02](#) (work in progress), July 2015.

[I-D.ietf-sidr-bgpsec-protocol]

Lepinski, M., "BGPsec Protocol Specification", [draft-ietf-sidr-bgpsec-protocol-12](#) (work in progress), June 2015.

[Kapela-Pilosov]

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-pilosov-kapela.pdf/>.

[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](#), January 2013.

[RFC7454]

Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", [BCP 194](#), [RFC 7454](#), February 2015.

[Toonk]

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

[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>>.

Internet-Draft

Route Leak Detection and Mitigation

July 2015

[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
Twitter, Inc.

Email: bdickson@twitter.com

