Methods for Detection and Mitigation of BGP Route Leaks
draft-sriram-route-leak-detection-mitigation-00

Abstract

   In [I-D.ietf-sriram-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 BGPSEC protocol [I-D.ietf-sidr-bgpsec-protocol-09]. Where
   the current BGPSEC protocol doesn'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.

Copyright Notice

Table of Contents

1. Introduction

   In [I-D.ietf-sriram-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 BGPSEC protocol [I-D.ietf-sidr-bgpsec-protocol]. The BGPSEC
   protocol provides cryptographic protection for some aspects of BGP
   update messages. It offers 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-sriram-
route-leak-problem-definition] and references cited at the back) are
another type of vulnerability in the global BGP routing system
against which BGPSEC so far offers only partial protection.

For the types of route leaks enumerated in [I-D.ietf-sriram-route-
leak-problem-definition], where the current BGPSEC protocol doesn'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 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. Mechanisms for Detection and Mitigation of Route Leaks

Referring to the five types of route leaks discussed in [I-D.ietf-
sriram-route-leak-problem-definition], Table 1 summarizes the route-
leak detection capability offered by ROA-based origin validation (OV)
and BGPSEC path validation (PV) for different types of route leaks.
(Note: Route filtering is not considered here in this table. Please
see Section 3.)

A detailed explanation of the contents of Table 1 is as follows. It
is readily observed that Type 1 and Type 5 route leaks are not
detected by origin validation or even by BGPSEC path validation. It
is also easy to observe that route leaks of Types 2 and 3 can be
readily detected by the BGPSEC protocol (which checks for origin-AS
validation and AS path signature verification). Type 2 route leak
involves changing a prefix to a subprefix (i.e. more specific); such
a modified update will fail BGPSEC validation checks. In the case of
Type 3 route leak, there would be no existing ROAs to validate a re-
originated prefix or subprefix, and hence the update will be
considered Invalid. If BGPSEC is deployed, then in the case of a
Type 3 route leak, the update will be Invalid due to Invalid path
signatures as well as Invalid origin AS.

Type of Route Leak	Detection Coverage and Comments
Type 1: U-Turn with Full Prefix	BGPSEC in its current form [I-D.ietf-sidr-bgpsec-protocol] does not detect Type 1.
Type 2: U-Turn with More Specific Prefix	ROA maxLength may offer detection in some cases when only origin validation is used; BGPSEC detects Type 2.
Type 3: Prefix Hijack with Data Path to Legitimate Origin	Origin validation by itself detects and also BGPSEC detects 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, origin validation helps detect their leak; they might either have no covering ROA or have a ROA-AS0 to always filter them. In the case of accidental deaggregation, ROA maxLength may offer detection. BGPSEC path validation does not catch Type 4.
Type 5: Lateral ISP to ISP Leak	BGPSEC in its current form [I-D.ietf-sidr-bgpsec-protocol] does not detect Type 5.

Table 1: Examination of Route-Leak Detection Capability of Origin
        Validation and Current BGPSEC AS-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
origin validation. That is because such prefixes might either have
no covering ROA or have a ROA-AS0 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 path validation
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 not 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.

From the above, it is evident that in our proposed solution method,
we need to focus primarily on Type 1 and Type 5 route leaks. In
Section 2.1 and Section 2.2, we describe a simple addition to BGPSEC
that facilitates cryptographically-enabled detection of route leaks
of Type 1. Then in Section 2.3, we will explain how the same method
as described in Section 2.1 can be utilized between ISPs to detect
and mitigate lateral (ISP to ISP) route leaks (i.e. Type 5).

2.1. 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 (AS3)
router) should be able to detect from the prefix-update that its
customer AS (e.g. AS1 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.

```
                                  /\              /\
                                   \ route-leak(P)/
                                    \ propagated /
                                     \          /
             +------------+    peer    +------------+
      _____| ISP1 (AS2) |----------->|  ISP2 (AS3)|---------->
     /       ------------+  prefix(P) +------------+ route-leak(P)
    | prefix |          \   update       /\          \  propagated
     \  (P)  /           \              /            \
      -------    prefix(P) \           /              \
               update  \             /                \
                        \          /route-leak(P)  \/
                         \/       /
                  +---------------+
                  | customer(AS1) |
                  +---------------+
```
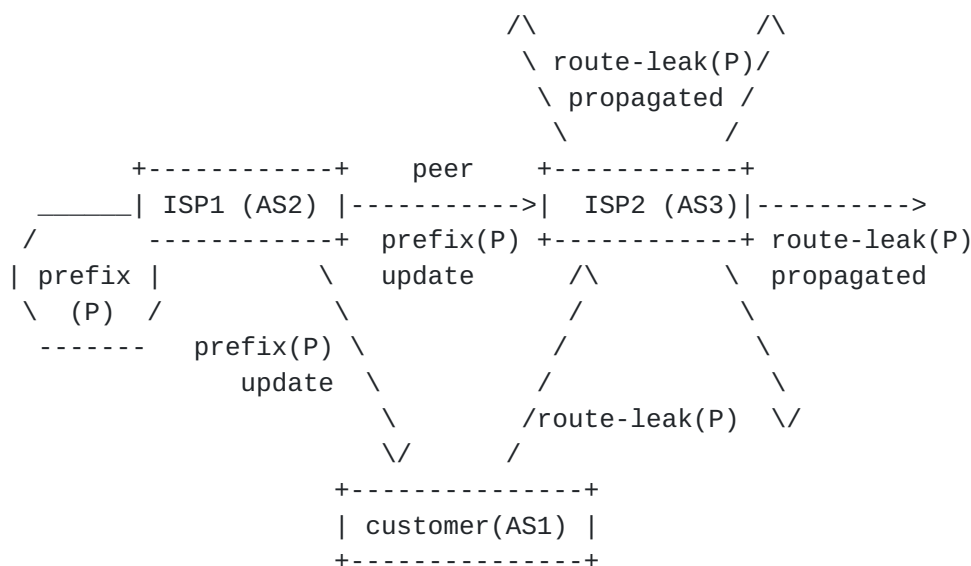
Figure 1: Illustration of the basic notion of a route leak.

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 an RLP field value (e.g. '10') to be used to
specify 'Up' (i.e. towards provider AS) directionality. 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 Type 1 (and also Type 5; see Section 2.3), which are
the focus here. (Please see Section 4 for further discussion about
the downside using 'Up' indication.)

In general, the proposed RLP encoding can be carried in BGP-4
[RFC4271] updates in any possible way, e.g., in a transitive
community attribute. We consider BGPSEC as an example, where the RLP
encoding can be accommodated in the existing Flags field and thereby
secured using the existing BGPSEC path signatures. The Flags field
is part of the Secure_Path Segment in BPGSEC 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.

2.2. Recommended Actions at a Receiving Router

We provide here an example set of receiver actions that work to
detect and mitigate route leaks of Type 1 (in particular). 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 2.1) 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 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 customer AS.

2. It is Valid in accordance with the BGPSEC protocol.

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 most
recent hop in the update is from its customer AS to itself, and hence
it does not need to rely on the RLP field value set by the customer
for detection of route leaks. (See further discussion in
Section 4.1.)

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.

The 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.

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.

2.3. Detection and Mitigation of Type 5 (Lateral ISP to ISP) Route Leaks

The sender and receiver actions described in Section 2.1 and
Section 2.2 clearly help detect and mitigate Type 1 route leaks.
With a slightly modified interpretation of the RLP encoding on the
receiver side, they can be extended to detect lateral ISP-to-ISP
route leaks (Type 5) also. A sending ISP router would set RLP field
value to '01' indication towards another ISP peer, following the same
sender principles as described in Section 2.1.

A recommended receiver algorithm for an ISP to detect a route leak of
Type 5 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 or a customer AS.

2. It is Valid in accordance with the BGPSEC protocol.

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 2.2) to
mean "the update SHOULD NOT have been propagated laterally to a peer
ISP like me either". This is strictly the case between settlement-
free large ISP peers.

The receiver algorithm for mitigation is up to the discretion of the
ISP. It may simply prefer another unmarked (i.e. not route-leak)
update from a different peer or an upstream ISP over a marked update.

3. Stopgap Solution when Only Origin Validation is Deployed

During a phase when BGPSEC 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.

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 an 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.

4. Design Rationale and Discussion

   In this section, we will try to provide design justifications for the
   methodology specified in Section 2, and also answer some anticipated
   questions.

4.1. Downside of 'Up (Towards Provider AS)' Indication in the RLP Field

   As we have shown in Section 2, route leak detection and mitigation
   can be performed without the use of 'Up' (i.e. from customer AS to
   provider AS) indication in the RLP field. The detection and
   mitigation action should primarily occur at a provider AS's router
   just as soon as a leaked update is received from a customer AS. At
   that point, a provider AS can be fooled if it merely looks to see if
   an offending customer AS has set an 'Up' indication in the RLP field.
   This is so since a customer AS intent on leaking a route can
   deliberately set "Not Specified (00)" indication in order to misguide
   its provider AS. So it seems better that a provider AS figures out
   that the update is moving in the 'Up' direction based only on its own

(configuration-based) knowledge that the update is coming from one of
its customer ASes. An 'Up' indication (if it were allowed) can be
also potentially misused. For example, an AS in the middle can
determine that a '01' (i.e. 'Do not Propagate Up') value already
exists on one of the preceding AS hops in a received update's AS
path. Then, said AS in the middle can deliberately set its own RLP
field to signal 'Up', in which case the update may be erroneously
marked as a route leak by a subsequent AS if it concludes that there
was a valley in the AS path of the update. So there appears to be
some possibility of misuse of 'Up' indication, and hence we proposed
not including it in the RLP specification in Section 2. However,
other proposals, if any, that aim to beneficially use an 'Up'
indication in the RLP field would be worth discussing.

4.2. **Possibility of Abuse of '01' (i.e.** 'Do not Propagate Up')
      Indication in the RLP Field

In reality, there appears to be no gain or incentive for an AS to
falsely set its own RLP field to '01' (i.e. 'Do not Propagate Up')
indication in an update that it originates or forwards. The purpose
of a deliberate route leak by an AS is to attract traffic towards
itself, but if the AS were to falsely set its own RLP field to '01'
value, it would be effectively repelling traffic away from itself for
the prefix in question (see receiver algorithms in Section 2.2 and
Section 2.3).

5. **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. Carrying the proposed
RLP encoding in a transitive community attribute in BGP is another
way, but in order to prevent abuse, the community attribute 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.

6. **Security Considerations**

The proposed Route Leak Protection (RLP) field requires cryptographic
protection. Since it is proposed that the RLP field be included in
the Flags field in the Secure_Path Segment in BPGSEC 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].

7. IANA Considerations

No updates to the registries are suggested by this document.

8. 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, Jakob Heitz, Geoff Huston, Randy
Bush, Ruediger Volk, Andrei Robachevsky, Chris Morrow, and Sandy
Murphy for comments, suggestions, critique at the IETF-90 in the
hall-ways and/or during the GROW WG meeting and/or on the GROW
mailing list. The authors are also thankful to Padma Krishnaswami,
Oliver Borchert, and Okhee Kim for their comments and review.

9. References

9.1. Normative References

[I-D.ietf-sidr-bgpsec-protocol]
          Lepinski, M., "BGPSEC Protocol Specification", draft-ietf-
          sidr-bgpsec-protocol-09 (work in progress), July 2014.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
          Protocol 4 (BGP-4)", RFC 4271, January 2006.

9.2. 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/>.

[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/>.

[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/>.

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

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

[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/>.

   [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/>.

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

   [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: dougm@nist.gov