

Global Routing Operations
Internet-Draft
Intended status: Informational
Expires: September 10, 2015

K. Sriram
D. Montgomery
US NIST
D. McPherson
E. Osterweil
Verisign, Inc.
March 9, 2015

Problem Definition and Classification of BGP Route Leaks
draft-ietf-grow-route-leak-problem-definition-01

Abstract

A systemic vulnerability of the Border Gateway Protocol routing system, known as 'route leaks', has received significant attention in recent years. Frequent incidents that result in significant disruptions to Internet routing are labeled "route leaks", but to date we have lacked a common definition of the term. In this document, we provide a working definition of route leaks, keeping in mind the real occurrences that have received significant attention. Further, we attempt to enumerate (though not exhaustively) different types of route leaks based on observed events on the Internet. We aim to provide a taxonomy that covers several forms of route leaks that have been observed and are of concern to Internet user community as well as the network operator community.

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 10, 2015.

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. Working Definition of Route Leaks	3
3. Classification of Route Leaks Based on Documented Events . .	3
4. Summary	6
5. Security Considerations	7
6. IANA Considerations	7
7. Acknowledgements	7
8. Informative References	7
Authors' Addresses	9

1. Introduction

Frequent incidents [Huston2012][Cowie2013][Cowie2010][Madory][Zmijewski][Paseka][LRL][Khare] that result in significant disruptions to Internet routing are commonly called "route leaks". Examination of the details of some of these incidents reveals that they vary in their form and technical details. Before we can discuss solutions to "the route leak problem" we need a clear, technical definition of the problem and its most common forms. In Section 2, we provide a working definition of route leaks, keeping in view many recent incidents that have received significant attention. Further, in Section 3, we attempt to enumerate (though not exhaustively) different types of route leaks based on observed events on the Internet. We aim to provide a taxonomy that covers several forms of route leaks that have been observed and are of concern to Internet user community as well as the network operator community.

2. Working Definition of Route Leaks

A proposed working definition of route leak is as follows:

A "route leak" is the propagation of routing announcement(s) beyond their intended scope. That is, an AS's announcement of a learned BGP route to another AS is in violation of the intended policies of the receiver, the sender and/or one of the ASes along the preceding AS path. The intended scope is usually defined by a set of local redistribution/filtering policies distributed among the ASes involved. Often, these intended policies are defined in terms of the pair-wise peering business relationship between ASes (e.g., customer, provider, peer). For literature related to AS relationships and routing policies, see [Gao][Gill][Luckie]. For measurements of valley-free violations in Internet routing, see [Giotsas][Wijchers].

The result of a route leak can be redirection of traffic through an unintended path which may enable eavesdropping or traffic analysis, and may or may not result in an overload or black-hole. Route leaks can be accidental or malicious, but most often arise from accidental misconfigurations.

The above definition is not intended to be all encompassing. Perceptions vary widely about what constitutes a route leak. Our aim here is to have a working definition that fits enough observed incidents so that the IETF community has a basis for starting to work on route leak mitigation methods.

3. Classification of Route Leaks Based on Documented Events

As illustrated in Figure 1, a common form of route leak occurs when a multi-homed customer AS (such as AS1 in Figure 1) learns a prefix update from one provider (ISP1) and leaks the update to another provider (ISP2) in violation of intended routing policies, and further the second provider does not detect the leak and propagates the leaked update to its customers, peers, and transit ISPs.

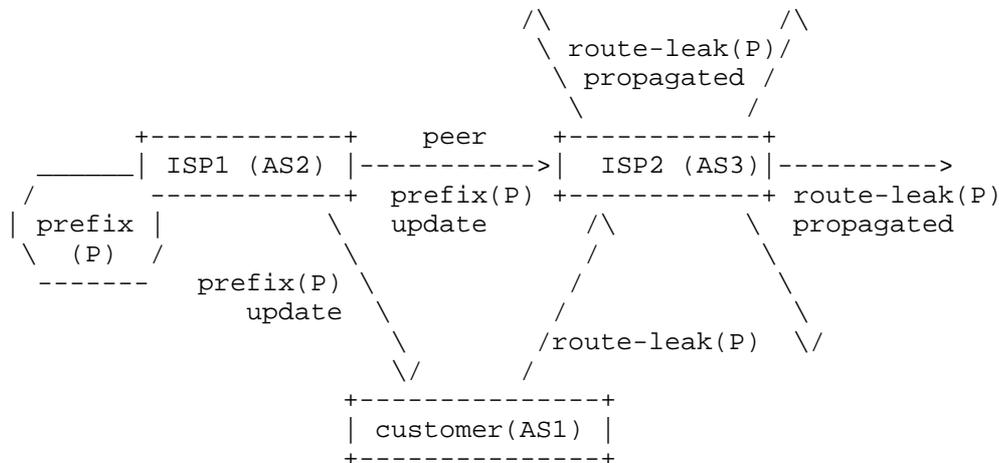


Figure 1: Illustration of the basic notion of a route leak.

We propose the following taxonomy for classification of route leaks aiming to cover several types of recently observed route leaks, while acknowledging that the list is not meant to be exhaustive. In what follows, we refer to the AS that announces a route that is in violation of the intended policies as the "offending AS".

- o Type 1 "U-Turn with Full Prefix": A multi-homed AS learns a prefix route from one upstream ISP and simply propagates the prefix to another upstream ISP. Neither the prefix nor the AS path in the update is altered. This is similar to a straight forward path-poisoning attack [Kapela-Pilosov], but with full prefix. It should be noted that attacks or leaks of this type are often accidental (i.e. not malicious). The update basically makes a U-turn at the attacker's multi-homed AS. The attack (accidental or deliberate) often succeeds because the second ISP prefers customer announcement over peer announcement of the same prefix. Data packets would reach the legitimate destination albeit via the offending AS, unless they are dropped at the offending AS due to its inability to handle resulting large volumes of traffic.
 - * Example incidents: Examples of Type 1 route-leak incidents are (1) the Dodo-Telstra incident in March 2012 [Huston2012], (2) the Moratel-PCCW leak of Google prefixes in November 2012 [Paseka], and (3) the VolumeDrive-Atrato incident in September 2014 [Madory].
- o Type 2 "U-Turn with More Specific Prefix": A multi-homed AS learns a prefix route from one upstream ISP and announces a sub-prefix

(subsumed in the prefix) to another upstream ISP. The AS path in the update is not altered. Update is crafted by the attacker to have a subprefix to maximize the success of the attack while reverse path is kept open by the path poisoning techniques as in [Kapela-Pilosov]. Data packets reach the legitimate destination albeit via the offending AS.

- * Example incidents: An example of Type 2 route-leak incident is the demo performed at DEFCON-16 in August 2008 [Kapela-Pilosov]. An attacker who deliberately performs a Type 1 route leak (with full prefix) can just as easily perform a Type 2 route leak (with subprefix) to achieve a greater impact.
- o Type 3 "Prefix Hijack with Data Path to Legitimate Origin": A multi-homed AS learns a prefix route from one upstream ISP and announces the prefix to another upstream ISP as if it is being originated by it (i.e. strips the received AS path, and re-originate the prefix). This amounts to straightforward hijacking. However, somehow (not attributable to the use of path poisoning trick by the attacker) a reverse path is present, and data packets reach the legitimate destination albeit via the offending AS. But sometimes the reverse path may not be there, and data packets get dropped following receipt by the offending AS.
 - * Example incidents: Examples of Type 3 route leak include (1) the China Telecom incident in April 2010 [Hiran][Cowie2010][Labovitz], (2) the Belarusian GlobalOneBel route leak incidents in February-March 2013 and May 2013 [Cowie2013], (3) the Icelandic Opin Kerfi-Simmin route leak incidents in July-August 2013 [Cowie2013], and (4) the Indosat route leak incident in April 2014 [Zmijewski].
- o Type 4 "Leak of Internal Prefixes and Accidental Deaggregation": An offending AS simply leaks its internal prefixes to one or more of its transit ASes and/or ISP peers. The leaked internal prefixes are often deaggregated subprefixes (i.e. more specifics) of already announced aggregate prefixes. Further, the AS receiving those leaks fails to filter them. Typically these leaked announcements are due to some transient failures within the AS; they are short-lived, and typically withdrawn quickly following the announcements.
 - * Example incidents: Leaks of internal prefix-routes occur frequently (e.g. multiple times in a week), and the number of prefixes leaked range from hundreds to thousands per incident. One highly conspicuous and widely disruptive leak of internal prefixes happened recently in August 2014 when AS701 and AS705

leaked about 22,000 more specifics of already announced aggregates [Huston2014][Toonk].

- o Type 5 "Lateral ISP-ISP-ISP Leak": This type of route leak typically occurs when, for example, three sequential ISP peers (e.g. ISP-A, ISP-B and ISP-C) are involved, and ISP-B receives a prefix-route from ISP-A and in turn leaks it to ISP-C. The typical routing policy between laterally (i.e. non-hierarchically) peering ISPs is that they should only propagate to each other their respective customer prefixes.
 - * Example incidents: In [Mauch-nanog][Mauch], route leaks of this type are reported by monitoring updates in the global BGP system and finding three or more very large ISP ASNs in a sequence in a BGP update's AS path. Mauch [Mauch] observes that these are anomalies and potentially route leaks because very large ISPs such as ATT, Sprint, Verizon, and Globalcrossing do not in general buy transit services from each other. However, he also notes that there are exceptions when one very large ISP does indeed buy transit from another very large ISP, and accordingly exceptions are made in his detection algorithm for known cases.
- o Type 6 "Leak of Provider Prefixes to Peer": This type of route leak occurs when an offending AS leaks prefix-routes learned from its provider to a lateral peer.
 - * Example incidents: The incidents reported in [Mauch] include the Type 6 leaks.
- o Type 7 "Leak of Peer Prefixes to Provider": This type of route leak occurs when an offending AS leaks prefix-routes learned from a lateral peer to its (the AS's) own provider. These leaked prefix-routes typically originate from the customer cone of the lateral peer.
 - * Example incidents: Some of the example incidents cited for Type 1 route leaks above are also inclusive of Type 7 route leaks. For instance, in the Dodo-Telstra incident [Huston2012], the leaked routes from Dodo to Telstra included routes that Dodo learned from its providers as well as lateral peers.

4. Summary

We attempted to provide a working definition of route leak. We also presented a taxonomy for categorizing route leaks. It covers not all but at least several forms of route leaks that have been observed and are of concern to Internet user and network operator communities. We

hope that this work provides the IETF community a basis for pursuing possible BGP enhancements for route leak detection and mitigation.

5. Security Considerations

No security considerations apply since this is a problem definition document.

6. IANA Considerations

No updates to the registries are suggested by this document.

7. 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, Brian Dickson, Amogh Dhamdhere, Jakob Heitz, Geoff Huston, Randy Bush, Ruediger Volk, Andrei Robachevsky, 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.

8. 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/>>.

[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/>>.
- [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/lrsl-events-from-2003-to-2009/>>.
- [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/>>.
- [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/>>.
- [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/>>.
- [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: doug@nist.gov

Danny McPherson
Verisign, Inc.

Email: dmcpherson@verisign.com

Eric Osterweil
Verisign, Inc.

Email: eosterweil@verisign.com