

Internet Engineering Task Force (IETF)  
Request for Comments: 7908  
Category: Informational  
ISSN: 2070-1721

K. Sriram  
D. Montgomery  
US NIST  
D. McPherson  
E. Osterweil  
Verisign, Inc.  
B. Dickson  
June 2016

## Problem Definition and Classification of BGP Route Leaks

### 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 a common definition of the term has been lacking. This document provides a working definition of route leaks while keeping in mind the real occurrences that have received significant attention. Further, this document attempts to enumerate (though not exhaustively) different types of route leaks based on observed events on the Internet. The aim is to provide a taxonomy that covers several forms of route leaks that have been observed and are of concern to the Internet user community as well as the network operator community.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7908>.

## 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. Working Definition of Route Leaks . . . . .	3
3. Classification of Route Leaks Based on Documented Events . .	4
3.1. Type 1: Hairpin Turn with Full Prefix . . . . .	4
3.2. Type 2: Lateral ISP-ISP-ISP Leak . . . . .	5
3.3. Type 3: Leak of Transit-Provider Prefixes to Peer . . . .	5
3.4. Type 4: Leak of Peer Prefixes to Transit Provider . . . .	5
3.5. Type 5: Prefix Re-origination with Data Path to Legitimate Origin . . . . .	6
3.6. Type 6: Accidental Leak of Internal Prefixes and More- Specific Prefixes . . . . .	6
4. Additional Comments about the Classification . . . . .	7
5. Security Considerations . . . . .	7
6. Informative References . . . . .	7
Acknowledgements . . . . .	11
Authors' Addresses . . . . .	11

## 1. Introduction

Frequent incidents [Huston2012] [Cowie2013] [Toonk2015-A] [Toonk2015-B] [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. In order to pursue solutions to "the route-leak problem" it is important to first provide a clear, technical definition of the problem and enumerate its most common forms. Section 2 provides a working definition of route leaks, keeping in view many recent incidents that have received significant attention. Section 3 attempts to enumerate (though not exhaustively) different types of route leaks based on observed events on the Internet. Further, Section 3 provides a taxonomy that covers several forms of route leaks that have been observed and are of concern to the Internet user community as well as the network operator community. This document builds on and extends earlier work in the IETF [ROUTE-LEAK-DEF] [ROUTE-LEAK-REQ].

## 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 announcement from an Autonomous System (AS) 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, transit provider, peer). For literature related to AS relationships and routing policies, see [Gao], [Luckie], and [Gill]. For measurements of valley-free violations in Internet routing, see [Anwar], [Giotsas], and [Wijchers].

The result of a route leak can be redirection of traffic through an unintended path that 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. Our aim here is to have a working definition that fits enough observed incidents so that the IETF community has a basis for developing solutions for route-leak detection and mitigation.

3. Classification of Route Leaks Based on Documented Events

As illustrated in Figure 1, a common form of route leak occurs when a multihomed customer AS (such as AS3 in Figure 1) learns a prefix update from one transit provider (ISP1) and leaks the update to another transit provider (ISP2) in violation of intended routing policies, and further, the second transit provider does not detect the leak and propagates the leaked update to its customers, peers, and transit ISPs.

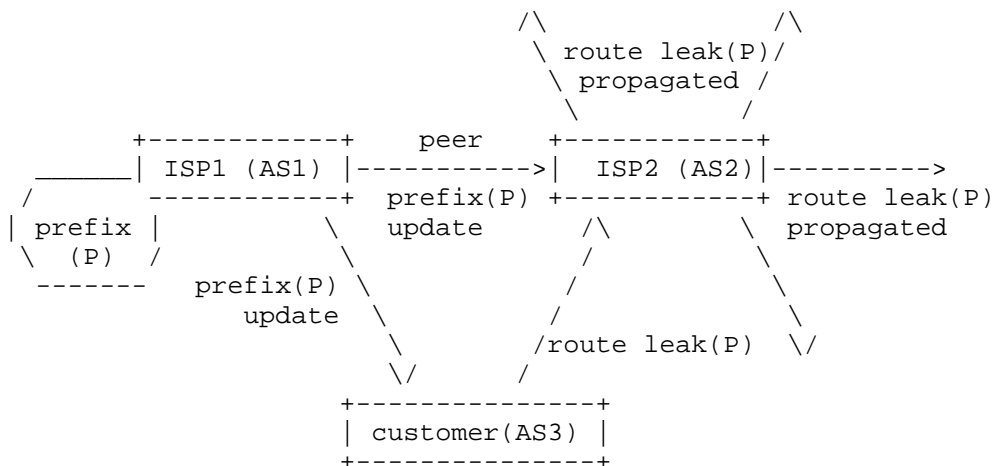


Figure 1: Basic Notion of a Route Leak

This document proposes the following taxonomy to cover several types of observed route leaks while acknowledging that the list is not meant to be exhaustive. In what follows, the AS that announces a route that is in violation of the intended policies is referred to as the "offending AS".

3.1. Type 1: Hairpin Turn with Full Prefix

Description: A multihomed AS learns a route from one upstream ISP and simply propagates it to another upstream ISP (the turn essentially resembling a hairpin). Neither the prefix nor the AS path in the update is altered. This is similar to a straightforward path-poisoning attack [Kapela-Pilosov], but with full prefix. It should be noted that leaks of this type are often accidental (i.e., not malicious). The update basically makes a hairpin turn at the offending AS's multihomed AS. The leak often succeeds (i.e., the leaked update is accepted and propagated) 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.

- o Example incidents: Examples of Type 1 route-leak incidents are (1) the Dodo-Telstra incident in March 2012 [Huston2012], (2) the VolumeDrive-Atrato incident in September 2014 [Madory], and (3) the massive Telekom Malaysia route leak of about 179,000 prefixes, which in turn Level3 accepted and propagated [Toonk2015-B].

### 3.2. Type 2: Lateral ISP-ISP-ISP Leak

Description: The term "lateral" here is synonymous with "non-transit" or "peer-to-peer". 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 route from ISP-A and in turn leaks it to ISP-C. The typical routing policy between laterally (i.e., non-transit) peering ISPs is that they should only propagate to each other their respective customer prefixes.

- o Example incidents: In [Mauch-nanog] and [Mauch], route leaks of this type are reported by monitoring updates in the global BGP system and finding three or more very large ISPs' Autonomous System Numbers (ASNs) in a sequence in a BGP update's AS path. [Mauch] observes that its detection algorithm detects for these anomalies and potentially route leaks because very large ISPs do not, in general, buy transit services from each other. However, it 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 its detection algorithm for known cases.

### 3.3. Type 3: Leak of Transit-Provider Prefixes to Peer

Description: This type of route leak occurs when an offending AS leaks routes learned from its transit provider to a lateral (i.e., non-transit) peer.

- o Example incidents: The incidents reported in [Mauch] include Type 3 leaks.

### 3.4. Type 4: Leak of Peer Prefixes to Transit Provider

Description: This type of route leak occurs when an offending AS leaks routes learned from a lateral (i.e., non-transit) peer to its (the AS's) own transit provider. These leaked routes typically originate from the customer cone of the lateral peer.

- o Example incidents: Examples of Type 4 route-leak incidents are (1) the Axcelx-Hibernia route leak of Amazon Web Services (AWS) prefixes causing disruption of AWS and a variety of services that run on AWS [Kephart], (2) the Hathway-Airtel route leak of 336 Google prefixes causing widespread interruption of Google services in Europe and Asia [Toonk2015-A], (3) the Moratel-PCCW route leak of Google prefixes causing Google's services to go offline [Paseka], and (4) some of the example incidents cited for Type 1 route leaks above are also inclusive of Type 4 route leaks. For instance, in the Dodo-Telstra incident [Huston2012], the leaked routes from Dodo to Telstra included routes that Dodo learned from its transit providers as well as lateral peers.

### 3.5. Type 5: Prefix Re-origination with Data Path to Legitimate Origin

Description: A multihomed AS learns a 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 can be called re-origination or mis-origination. However, somehow a reverse path to the legitimate origination AS may be present and data packets reach the legitimate destination albeit via the offending AS. (Note: The presence of a reverse path here is not attributable to the use of a path-poisoning trick by the offending AS.) But sometimes the reverse path may not be present, and data packets destined for the leaked prefix may be simply discarded at the offending AS.

- o Example incidents: Examples of Type 5 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]. The reverse paths (i.e., data paths from the offending AS to the legitimate destinations) were present in incidents #1, #2, and #3 cited above, but not in incident #4. In incident #4, the misrouted data packets were dropped at Indosat's AS.

### 3.6. Type 6: Accidental Leak of Internal Prefixes and More-Specific Prefixes

Description: An offending AS simply leaks its internal prefixes to one or more of its transit-provider ASes and/or ISP peers. The leaked internal prefixes are often more-specific prefixes subsumed by an already announced, less-specific prefix. The more-specific prefixes were not intended to be routed in External BGP (eBGP). 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. However, these more-specific prefixes may momentarily cause the routes to be preferred over other aggregate (i.e., less specific) route announcements, thus redirecting traffic from its normal best path.

- o Example incidents: Leaks of internal 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 routes happened in August 2014 when AS701 and AS705 leaked about 22,000 more-specific prefixes of already-announced aggregates [Huston2014] [Toonk2014].

#### 4. Additional Comments about the Classification

It is worth noting that Types 1 through 4 are similar in that a route is leaked in violation of policy in each case, but what varies is the context of the leaked-route source AS and destination AS roles.

A Type 5 route leak (i.e., prefix mis-origination with data path to legitimate origin) can also happen in conjunction with the AS relationship contexts in Types 2, 3, and 4. While these possibilities are acknowledged, simply enumerating more types to consider all such special cases does not add value as far as solution development for route leaks is concerned. Hence, the special cases mentioned here are not included in enumerating route-leak types.

#### 5. Security Considerations

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

#### 6. Informative References

[Anwar] Anwar, R., Niaz, H., Choffnes, D., Cunha, I., Gill, P., and N. Katz-Bassett, "Investigating Interdomain Routing Policies in the Wild", In Proceedings of the 2015 ACM Internet Measurement Conference (IMC), DOI 10.1145/2815675.2815712, October 2015, <<http://www.cs.usc.edu/assets/007/94928.pdf>>.

[Cowie2010] Cowie, J., "China's 18 Minute Mystery", Dyn Research: The New Home of 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: The New Home of 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 (TON), Volume 9, Issue 6, pp 689-692, DOI 10.1109/90.974523, 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, Volume 44, Issue 1, pp 28-34, DOI 10.1145/2567561.2567566, January 2014, <<http://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", 2012 IEEE International Conference on Communications (ICC), DOI 10.1109/ICC.2012.6363987, June 2012.
- [Hiran] Hiran, R., Carlsson, N., and P. Gill, "Characterizing Large-Scale Routing Anomalies: A Case Study of the China Telecom Incident", In Proceedings of the 14th International Conference on Passive and Active Measurement (PAM) 2013, DOI 10.1007/978-3-642-36516-4\_23, March 2013, <<http://www3.cs.stonybrook.edu/~phillipa/papers/CTelecom.html>>.
- [Huston2012] Huston, G., "Leaking Routes", Asia Pacific Network Information Centre (APNIC) Blog, March 2012, <<http://labs.apnic.net/blabs/?p=139/>>.
- [Huston2014] Huston, G., "What's so special about 512?", Asia Pacific Network Information Centre (APNIC) Blog, 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", 16th Defcon Conference, August 2008, <<https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-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", In Proceedings of the 2013 ACM Internet Measurement Conference (IMC), DOI 10.1145/2398776.2398780, 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", University of Arizona (UA) Network Research Lab: Projects Webpage, 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", In Proceedings of the 2013 ACM Internet Measurement Conference (IMC), DOI 10.1145/2504730.2504735, 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: The New Home of 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", 41st NANOG Conference, 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", CloudFlare Blog, November 2012, <<http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about/>>.
- [ROUTE-LEAK-DEF]  
Dickson, B., "Route Leaks -- Definitions", Work in Progress, draft-dickson-sidr-route-leak-def-03, October 2012.
- [ROUTE-LEAK-REQ]  
Dickson, B., "Route Leaks -- Requirements for Detection and Prevention thereof", Work in Progress, draft-dickson-sidr-route-leak-reqts-02, March 2012.
- [Toonk2014]  
Toonk, A., "What caused today's Internet hiccup", BGPMON Blog, August 2014, <<http://www.bgpmon.net/what-caused-todays-internet-hiccup/>>.
- [Toonk2015-A]  
Toonk, A., "What caused the Google service interruption", BGPMON Blog, March 2015, <<http://www.bgpmon.net/what-caused-the-google-service-interruption/>>.
- [Toonk2015-B]  
Toonk, A., "Massive route leak causes Internet slowdown", BGPMON Blog, June 2015, <<http://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/>>.
- [Wijchers] Wijchers, B. and B. Overeinder, "Quantitative Analysis of BGP Route Leaks", Reseaux IP Europeens (RIPE) 69th Meeting, November 2014, <<http://ripe69.ripe.net/presentations/157-RIPE-69-Routing-WG.pdf>>.
- [Zmijewski]  
Zmijewski, E., "Indonesia Hijacks the World", Dyn Research: The New Home of Renesys Blog, April 2014, <<http://research.dyn.com/2014/04/indonesia-hijacks-world/>>.

## Acknowledgements

The authors wish to thank Jared Mauch, Jeff Haas, Warren Kumari, Amogh Dhamdhere, Jakob Heitz, Geoff Huston, Randy Bush, Job Snijders, Ruediger Volk, Andrei Robachevsky, Charles van Niman, 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.

## Authors' Addresses

Kotikalapudi Sriram  
US NIST

Email: [ksriram@nist.gov](mailto:ksriram@nist.gov)

Doug Montgomery  
US NIST

Email: [doug@nist.gov](mailto:doug@nist.gov)

Danny McPherson  
Verisign, Inc.

Email: [dmcpherson@verisign.com](mailto:dmcpherson@verisign.com)

Eric Osterweil  
Verisign, Inc.

Email: [eosterweil@verisign.com](mailto:eosterweil@verisign.com)

Brian Dickson

Email: [brian.peter.dickson@gmail.com](mailto:brian.peter.dickson@gmail.com)