sidr
Internet-Draft
Expires: April 26, 2013

B. Dickson
Brian Dickson

October 23, 2012

Route Leaks -- Definitions
draft-dickson-sidr-route-leak-def-03

Abstract

The Border Gateway Protocol, version 4, (BGP4) provides the means to advertise reachability for IP prefixes. This reachability information is propagated in a peer-to-peer topology. Routes may be announced to neighbors, contrary to the receiver's local peering policy. If that occurs, those routes may then be propagated indiscriminantly, once they have been accepted.

This document considers the situations that can lead to routes being leaked, and tries to find acceptable definitions for describing these scenarios.

The purpose of these definitions is to facilitate analysis of what a route leak is, and what the scope of the problem space for route leaks is.

This, in turn, is intended to inform a requirements document for detection of (and prevention of) route leaks. And finally, the definitions and requirements are intended to allow proposed solutions which meet these criteria, and to facilitate evaluation of proposed solutions.

The ultimate goal is to "solve the route leaks problem".

Author's Note

Intended Status: Informational.

Status of this Memo

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 April 26, 2013.

Copyright Notice

Table of Contents

1. Introduction

1.1. Assumptions

Much of this document assumes the observer has total knowledge of the state of everything in the hypothetical examples presented.

It is understood that participants in the real world routing scenarios will not have that knowledge.

The purpose of presuming that total knowledge here, is to illustrate how little is needed to identify leaked routes.

In particular, it is hoped that this leads to a correspondingly simple set of definitions with useful real-world meaning.

1.2. Rationale

Generally speaking, a route-leak occurs when a route goes somewhere it should not. In other words, that somewhere along the path, a route was sent that somehow violated the implicit or explicit policy between two neighbors, without being blocked by the recipient. Route leaks cause harm, in a variety of ways. They expose traffic to Man-In-The-Middle (MITM) attacks. They may result in traffic congestion, latency, or even black-holing of traffic.

It is a leak if any receiver in the propogation path did not want the route, from a generic policy perspective. It does not matter which party caused the situation - a leaks are in the eye of the receivers. By their nature - unintentional, unwanted, and harmful - route leaks are bad.

By first establishing a more precise definition of route leak, the intent is to find requirements for mechanisms for stopping route leaks, and then finding solutions that meet those requirements.

1.3. Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.4. Terminology

The reader is assumed to be familiar with BGP version 4, both from a protocol perspective and from an operational perspective. BGP4 is defined in [RFC1771], and updated or enhanced by a variety of other RFCs.

The following additional terminology is used throughout this document:

Route (or synonomously, prefix): an NLRI in BGP, including all its attributes. (This term subject to change by GROW.)

Neighbor (or "peer", not capitalized): A toplogically adjacent Autonomous System, with whom routes are exchanged.

Link: A BGP connection to a Neighbor. A Neighbor may be reached via one or more links, where each link may have a different classification, and/or local policy.

Link Classification: The "intent" of a given BGP peering session, which addresses only the categories of route announced and accepted, and which is further modified by Local Policy.

Local Policy: The set of rules, as applied on a single Neighbor Link, specifying which routes are announced, which routes are accepted, and what attributes are changed to affect choice of BGP Best Path per prefix.

Path: Also known as AS_PATH (or optionally AS4_PATH), the sequence of ASNs through which a route has passed from Originator to recipient.

Hijacked Route: A route which has been originated by a party other than the owner of the prefix. This could be via a forged ASN, or from another ASN.

Validated Origination: a route whose origination has been validated, e.g. via cryptographic means, such as using an ROA.

2. Scope Limitations

The following issues are not in the scope of route leaks. Each item in the list includes the rationale for excluding it.

- Hijacked Routes - Origin Validation (proposed work in the SIDR WG) addresses the issue of Hijacked Routes. By limiting Route Leak efforts to Validated Routes, we are able to presume the origin is correct, and narrow the scope.
- Violations of Local Policy - issues between adjacent ASNs which do not propogate any further, or which do not violate the Link Classification.
- Other-ASN Relationship - The "correctness" of a given prefix received over a Link, is determined only by the Link Classifications of each Link in the Path. The existence of other Links, to Neighbors with ASNs on a given Path (which may have

differing Link Classifications), is a classic "apples to oranges" comparision. It is incorrect to compare ASNs outside the context of the AS path, so we exclude those comparisons from this work.

Essentially, the only elements being considered are the Path, and Link Classifications at each hop in the Path.

3. Route Leak Definitions

Route Leak Initiation: A Route announced over a Link by a Neighbor, which does not match the Link Classification, where one of the following is true:

- the Neighbor is the Originator
- the Neighbor received the Route, where the received Route was not a Route Leak

In lay terms, this means that the Neighbor is the party that caused the route leak, by announcing a route contrary to the Link Classification (and consequently also violated the Local Policy).

Route Leak Propagation: A Route announced over a Link by a Neighbor, where the Route that the Neighbor received was either a Route Leak Initiation, or a Route Leak Propagation.

Once a Route has become a Route Leak Initiation, any further announcement of that Route is a Route Leak Propagation.

NB: A Route Leak Propagation may appear to match the Link Classification, since the Path appears similar to non-leaked routes for the first two ASNs in the Path.

Link Classifications: a Link may be classified as:

- Customer
- Transit
- Peer
- Special (which includes Mutual Transit, Sibling, and other non-trivial arrangements)

Special (e.g. Mutual Transit): a Link where the two parties agree to provide full routes, and to advertise each others' customers routes the same as they would advertise their own customers' routes. Semantically, this behaves the same as having two parallel Links between the same two Neighbors, where one Link Policy is Transit and the other Link Policy is Customer. Recall, Link Classification is the superset of Local Policy - the term "full routes" here means simply that any route in addition to customers' routes, is permitted.

4. Peer Links and Routes

A Peer Classification is a Link over which the two parties send ONLY their respective Customer Routes (and their Customer's Routes, and so on).

A Link which is classified as a Peer, will see us as a Peer Classification as well. The relationship is symmetric in nature.

5. Customer Links and Routes

A Customer Link Classification: The Customer sends us only their own (locally originated) Routes, and the Customer's Customer's Routes (and Customer^Nth Routes). The Customer relationship is transitive.

A Transit Link Classification: The Transit provider sends all Routes. This include the Transit Provider's Customers, the Transit Provider's Peers, and if there are any, the Transit Provider's Transit Provider's Routes. The Transit Provider relationship is also transitive.

Transit and Customer are the opposite ends of the same Link, by definition.

The Transit Link Classification is a superset of the actual Local Policy of a specific Customer. This means that while a Transit Link Classification means "we send all routes", the actual Local Policy for a specific Customer might differ, and the Customer might only receive some Routes, or none at all. Similarly, the Classification means that we are prepared to accept the Customer's own Routes, as well as those of the Customer's Customers. However, the Local Policy might be to accept only a specific subset of the Customer's Routes.

5.1. Customer's Customer

It is important to define when a Route is a Customer's Customer Route.

A Customer's Customer Route: the Path to be from the Customer's Customer, to the Customer, to us. Similarly, Customer^Nth Paths must proceed directly from Customer^N to Customer^(N-1) to Customer to us. It is not sufficient for the Origin of the Route to be the ASN of a Customer's Customer. Each Link must be a Customer Classification (or Special, e.g. Mutual Transit, which is a superset of Customer).

In particular, if the Path were to include any Link which were not a

Customer Link, the Route would NOT be a Customer^N.

NB: It is sufficient that the Customer's Customer relationship is declared. The "Customer" relationship, in the context of route leaks, is restrictive. Erroneous or inadvertent classification as Customer cannot result in a route leak.

6. Non-Leak-Initiation Links

To help identify the exact conditions where a Route Leak Initiation can occur, it is helpful to exclude Link Classifications where it is axiomatically impossible to cause a Route Leak Initiation.

Since a Transit Classification, by definition, can receive all routes, a Transit Link cannot be the source of a Route Leak Initiation. By the same logic, a Special (e.g. Mutual Transit) Classification cannot be the source of a Route Leak Initiation.

This leads to a more precise definition of a Route Leak Initiation.

7. Route Leak Initiation

Route Leak Initiation: A Non-Customer Route which is received over a Peer or Customer Link.

8. Route Leak

Route Leak: any Route where, somewhere in the Path, a Non-Customer Route was received over a Peer or Customer Link. (This is synomous with "was sent over a Peer or Transit Link".)

It should be observed that a route which is not a route leak, has an as-path that matches the following pattern:

{C|S}*P?{T|S}*

Where C is Customer, T is Transit, P is Peer, and S is Special, and "{ | }" denotes either/or, "*" means zero or more occurrences of, and "?" means zero or one occurrences of.

9. Security Considerations

None per se.

10. IANA Considerations

This document contains no IANA-specific material.

11. Acknowledgements

To be added later.

12. References

12.1. Normative References

[RFC1033] Lottor, M., "Domain administrators operations guide", RFC 1033, November 1987.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, March 1998.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", RFC 5011, September 2007.

[RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, March 2008.

12.2. Informative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Author's Address

Brian Dickson
Brian Dickson
703 Palmer Drive,
Herndon, VA 20170
USA

Email: brian.peter.dickson@gmail.com