

sidr  
Internet-Draft  
Expires: September 7, 2012

B. Dickson  
Brian Dickson  
March 6, 2012

Route Leaks -- Requirements for Detection and Prevention thereof  
draft-dickson-sidr-route-leak-reqts-02

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. Sometimes routes are announced to peers for which the local peering policy does not permit. And sometimes routes are propagated indiscriminantly, once they have been accepted.

This document is a requirements document for detection of (and prevention of) route leaks.

Together with the definitions document, it is intended to suggest solutions which meet these criteria, and to facilitate evaluation of proposed solutions.

The fundamental objective is to "solve the route leaks problem".

Author's Note

Intended Status: Informational.

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 7, 2012.

Copyright Notice

Copyright (c) 2012 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1.</a>	Rationale . . . . .	<a href="#">3</a>
<a href="#">1.2.</a>	Requirements . . . . .	<a href="#">3</a>
<a href="#">1.3.</a>	Terminology . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Peering Terms and Symbols . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Local Non-Leak Prefix Advertisement Matrix & Rules . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Route Leak Detection Requirements . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	Coloring Rules . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	Route Modification Rules . . . . .	<a href="#">6</a>
<a href="#">4.3.</a>	Single Party Rules . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">7</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">7</a>
<a href="#">7.</a>	Acknowledgements . . . . .	<a href="#">7</a>
<a href="#">8.</a>	References . . . . .	<a href="#">8</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">8</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">8</a>
	Author's Address . . . . .	<a href="#">8</a>

## 1. Introduction

### 1.1. Rationale

This document analyzes the particulars of situations which introduce route leaks, or propagates those leaks.

Using the definitions previously established, those conditions are reduced to a minimum set of requirements for the identification of route leaks.

Those conditions are validated at length, and all of the assumptions stated, and consequential conditions enumerated.

The result is a set of criteria for solving the route leak problem, preventing any single source of leakage regardless of intent or nature (operator, implementor, bad actor).

### 1.2. 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](#)].

### 1.3. Terminology

The reader is assumed to be familiar with the IETF.

## 2. Peering Terms and Symbols

We can represent the per-link peering categorizations with the following symbols:

Neighbor is:

- a. Transit Provider - T
- b. (Transit) Customer - C
- c. Peer - P
- d. Mutual Transit

In any neighbor relationship, the roles of the parties on either end of the link would be:

T-C

C-T

P-P

Mc-Mtp

Mtp-Mc

(where the last two, Mc/Mtp are a semantic and/or coloring distinction on routes, rather than two separate links.)

### 3. Local Non-Leak Prefix Advertisement Matrix & Rules

The following matrix shows what prefixes from a given source peering relationship, may be advertised to a given neighbor peering relationship without causing a route leak.

Src \ Dest	P	T	Mtp	Mc	C
P	-	-	-	Y	Y
T	-	-	-	Y	Y
Mtp	-	-	-	Y	Y
Mc	Y	Y	Y	-	Y
C	Y	Y	Y	-	Y

+-----+-----+-----+-----+-----+

Grouping the like items (by row and column) we get:

Src \ Dest	T/Mtp	P	Mc	C
T/Mtp	-	-	Y	Y
P	-	-	Y	Y
C/Mc	Y	Y	-	Y

When a prefix is sent to any T neighbor, the receiving neighbor sees it as C. Similarly, Mc is seen at Mtp.

The inverse of these is also true: C->T, Mtp->Mc.

And lastly, a prefix sent to a (P) will be received by the neighbor

as a (P).

This means that once a prefix has been sent to any of the two type sets "P" or "C/Mc", it must only subsequently be sent to "C" or "Mc" types.

This results in the regular expression for a valid (non-leaked) path:

$$\text{Origin} - (T - |Mtp - )*(P - )?(C - |Mc - )* \text{Destination}$$

Thus we have the basis for a simple set of rules, which would enable detecting and preventing route leaks.

#### 4. Route Leak Detection Requirements

Based on the advertisement rules, we now have enough information to specify the main rules that a Route Leak Detector would need to observe.

#### [4.1.](#) Coloring Rules

In no particular order, here are the requirements for coloring the path of a route.

- o Every BGP peering session (Link) MUST have a type associated with it.
- o Neighbors Agree - both sides of a BGP peering link must negotiate and agree on the link type.
- o Last Color Agrees with Link - the last color applied to the route must be the consistent with the link type.
- o If the Color used towards "Transit" is "Green", and the Color used towards "Peer" or "Customer" is "Yellow", then:
  - \* The entire Path must have a corresponding set of Colors, one for each AS-Hop.
  - \* The Path must be of the form (Green)\*(Green|Yellow)(Yellow)\*.
  - \* Once a Path has switched to Yellow, it cannot switch back to Green.

- \* Routes sent to T neighbors must mark the path Green.
- \* Only Green Routes may be sent to T or P neighbors.
- \* Routes sent to C or P neighbors must mark the path Yellow.
- \* A route learned via a P neighbor must be all Green followed by a single Yellow.
- \* A route learned via a T neighbor must be zero or more Greens followed by one or more Yellows.
- \* A route learned via a C neighbor must be one or more Greens (and no Yellows).

- \* Mutual Transit links must preserve the current color.
- \* Colors may be explicitly marked, or may be inferred as long as there is no room for ambiguity.

#### 4.2. Route Modification Rules

In addressing accidental route leaks, the secondary goal is to also prevent malicious route leaks.

The only additional rule for this is, that any additional BGP attributes implementing this would need to be included in the set of things cryptographically signed. This provides tamper evidence and prevention of substitution of values (on received routes).

This means that the assigning of colors must be handed by implementation based only on Link Type (and current Route color), with no over-ride by the operator possible, with a single exception: It should always be possible to "demote" a route from Green to Yellow, locally before or while sending.

Similarly, route-leak filtering of routes on both the send and receive direction, MUST be done based only on color vs link type. There cannot be an operator-exposed over-ride.

For an operator who has a need to make a routing announcement that violates the Link Type, the correct course of action would be to change the Link type. This would need to be done cooperatively with the party at the other end of the link.

#### 4.3. Single Party Rules

One objective in preventing Route Leaks from being initiated or propagated, is to examine the control points of the routing path itself.

By treating this as a path where the goal is to avoid any single point of failure, we can derive additional rules.

Here, the term "failure" is synonymous with "route leak". In other words, are there any points where a single error or omission can cause a route leak?

If there are any, the goal should be to replace those with equivalent elements which would require two errors or actions, by independent parties, to cause a route leak.

Here are some of the places where this is accomplished or needs to be done by solutions:

- o Sender/Receiver - both ends of a link need to agree on the type. Unilateral error here must fail "safe" -> BGP does not establish, with errors.
- o Always Validate Color Rules - while the blocking of leaked routes should occur automatically at the point of leak, failure to block a leak SHOULD be detected and the route SHOULD be blocked by the next recipient.

## [5. Security Considerations](#)

None per se.

## [6. IANA Considerations](#)

This document contains no IANA-specific material.

## [7. Acknowledgements](#)

To be added later.

## [8. References](#)

### [8.1. Normative References](#)



- [RFC1773] Traina, P., "Experience with the BGP-4 protocol", [RFC 1773](#), March 1995.
- [RFC1997] Chandrasekeran, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", [RFC 4456](#), April 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

## 8.2. 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](mailto:brian.peter.dickson@gmail.com)