

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: January 9, 2020

Y. Gu  
Huawei  
H. Chen  
China Telecom Co., Ltd.  
D. Ma  
ZDNS  
S. Zhuang  
Huawei  
July 8, 2019

BMP for BGP Route Leak Detection  
draft-gu-grow-bmp-route-leak-detection-03

Abstract

According to RFC7908 [RFC7908], Route leaks refer to case that the delivery range of route advertisements is beyond the expected range. For many current security protection solutions, the ISPs (Internet Service Providers) are focusing on finding ways to prevent the happening of route leaks. However, the real-time route leak detection if any occurs is important as well, and serves as the basis for leak mitigation. This document extends the BGP Monitoring Protocol (BMP) [RFC7854] to provide a routing security scheme suitable for ISPs to detect BGP route leaks at the prefix level.

Requirements Language

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 RFC 2119 [RFC2119].

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 <https://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 January 9, 2020.

## Copyright Notice

Copyright (c) 2019 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 (<https://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. Terminology . . . . .	2
2. Introduction . . . . .	3
2.1. Actions Against Route Leaks . . . . .	3
2.2. Challenges of the Current Actions against Route Leaks . .	4
3. Route Leak Detection (RLD) Design Considerations . . . . .	5
4. BMP Support for RLD . . . . .	5
4.1. RLD TLV Format . . . . .	5
4.2. RLD TLV usage . . . . .	6
4.3. Coordination with iOTC and RLP . . . . .	7
5. Acknowledgements . . . . .	8
6. Contributors . . . . .	8
7. IANA Considerations . . . . .	8
8. Security Considerations . . . . .	8
9. References . . . . .	9
9.1. Normative References . . . . .	9
9.2. Informative References . . . . .	9
Authors' Addresses . . . . .	10

## 1. Terminology

BMP: BGP Monitoring Protocol

BMS: BGP Monitoring Station

C2P: Customer to Provider

ISP: Internet Service Provider

P2P: Peer to Peer

RIB: Routing Information Base

RLP: Route Leak Protection

RLD: Route Leak Detection

## 2. Introduction

RFC 7908 defines "Route Leak" as: A route leak is the propagation of routing announcement(s) beyond their intended scope, which can result in possible situations such as eavesdropping, device overload, route black hole and so on. More specifically, the intended scope of route announcements is usually defined by local route filtering/distribution policies within devices. These policies are designed to realise the pair-wise peering business relationships between ASes (autonomous systems), which include Customer to Provider (C2P), Peer to Peer (Peer to Peer), and Provider to Customer (P2C). In a C2P relationship, the customer pays the transit provider for traffic sent between the two ASes. In return, the customer gains access to the ASes that the transit provider can reach, including those which the transit provider reaches through its own transit providers. In a P2P relationship, the peering ASes gain access to each other's customers, typically without either AS paying the other AS Relationships, Customer Cones, and Validation [Luckie].

More precisely, the route leaks we discuss in this draft, referring to Type 1, 2, 3, and 4 Route Leaks defined in RFC 7908, can be summarized as: a route leak occurs when a route received from a transit provider or a lateral peer is propagated to another transit provider or a lateral peer.

### 2.1. Actions Against Route Leaks

There are several types of approaches against route leak from different perspectives. In this draft, we mainly discuss the following three types:

- o Route leak prevention: The approach to prevent route leak from happening. Commonly used methods include inbound/outbound prefix/peer/AS filtering policies configured at the ingress/egress nodes of ASes per the propagation of BGP routes.
- o Route leak detection: The approach to detect the existence of route leaks that happen at either the local AS, or upstream AS per the propagation of BGP routes. An intuitive way of detecting route leak is by checking the business relationship information at both the ingress and egress nodes of the local AS along the BGP route propagation path with the route leak violation rules defined in RFC7908 [RFC7908]. Thus, it requires the knowledge of the actual route propagation trace, as well as the resulting business

relationship information at the ingress and egress nodes. With the above information collected, the analysis can be done by the routing device or a centralized server. This draft specifies one such method.

- o Route leak mitigation: The approach to mitigate route leaks that already happened at either the local AS, or upstream AS per the propagation of BGP routes. Commonly used methods include reject, drop or stop propagating the invalid routes once detected the existence of leaks.

The above mentioned actions can be used separately or combinely, depending on the entities (routing devices, network manager) that execute the actions, and the relative positions of the executing entities from the leaking point (local or downstream).

## 2.2. Challenges of the Current Actions against Route Leaks

draft-ietf-idr-bgp-open-policy [I-D.ietf-idr-bgp-open-policy] updates the BGP Open negotiation process with a new BGP capability to exchange the BGP Roles between two BGP speakers, and also proposes to use a new BGP attribute, called the iOTC (Internal Only To Customer) Path attribute to mark routes according to the BGP Roles established in Open Message. The iOTC attribute of the incoming route is set at the ingress node of the local AS, and is conveyed with the BGP Update to the egress node of the local AS for outbound filtering to prevent route leaks in the local AS. This attribute is removed at the egress node before the BGP Update is sent to eBGP neighbors. For representation simplification, we use iOTC to refer to the method specified in draft-ietf-idr-bgp-open-policy [I-D.ietf-idr-bgp-open-policy].

draft-ietf-grow-route-leak-detection-mitigation [I-D.ietf-grow-route-leak-detection-mitigation] describes a route leak detection and mitigation solution based on conveying route-leak protection (RLP) information in a well-know transitive BGP community, called the RLP community. The RLP community helps with detection and mitigation of route leaks that happen at the upstream AS (per the BGP routes propagation), as an inter-AS solution. For representation simplification, we use RLP to refer to the method specified in draft-ietf-grow-route-leak-detection-mitigation [I-D.ietf-grow-route-leak-detection-mitigation].

The above two drafts provide solutions for route leak prevention, detection and mitigation. To summarize:

- o iOTC is used for route leak prevention of the local AS. It does not provide the detection or mitigation of route leaks of either local As or upstream AS per the BGP routes propagation.
- o iOTC is peer/AS-level route leak prevention, due to the fact the BGP Role negotiation is peer-level. It's does not provide prefix-level route leak prevention.
- o RLP is used for route leak detection and mitigation of route leak that happens in the upstream AS (per the BGP Update distribution). It is prefix-level detection and mitigation.

Thus, there lacks method for local AS route leak detection.

### 3. Route Leak Detection (RLD) Design Considerations

Considering the challenges facing the existing approaches, this draft proposes a method called Route Leak Detection (RLD). It utilizes the BGP Monitoring Protocol (BMP) to convey the RLD information from to the BMP server to realize centralized leak detection. BMP is currently deployed by OTT and carriers to monitor the BGP routes, such as monitoring BGP Adj-RIB-In using the process defined in RFC7854 [RFC7854], and monitoring BGP Adj-RIB-Out using the process defined in draft-ietf-grow-bmp-adj-rib-out [I-D.ietf-grow-bmp-adj-rib-out]. On the other hand, the RLD information is in fact a representation of the business relationships between the local AS and its neighboring AS. It does not involve any information disclosure issue regarding third parties. Thus, a single ISP can deploy RLD without relying on any information from either other ISPs or other third parties.

### 4. BMP Support for RLD

#### 4.1. RLD TLV Format

A RLD TLV is defined for the BMP Route Monitoring Message. Considering that the AS relationships are sometimes per route based instead of per peer/AS based, this TLV is appended to each route, following the BGP Update Message. The order of placing the RLD TLV among other BMP supported TLVs is out of the scope of this draft. The TLV format is defined as follows:

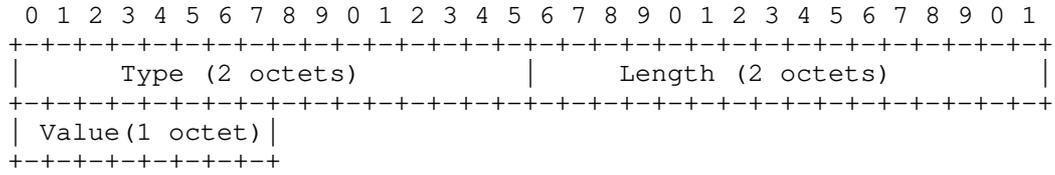


Figure 1: RLD TLV

- o Type (2 octets) = TBD1, the RLD TLV represents the prefix-level business relationship between the transmitter AS and the receiver AS. The local AS is a transmitter or a receiver, depending on if the route is an incoming route from a neighbor AS or an outgoing route to a neighbor AS.
- o Length (2 octets): Defines the length of the Value field. It SHOULD be set to 0x01, considering the Value field is of 1 octet fixed length.
- o Value (1 octet): Currently 4 values are defined to represent the business relationships, which are specified in Table 1.

Value	Business Relationship
0	P2C
1	C2P
2	P2P
3	I2I

Table 1: Business relationship value

4.2. RLD TLV usage

The RLD TLV, presenting the business relationship between the neighbor AS and the local AS of the incoming route, SHOULD be prepended to the Adj-rib-in at the ingress node of the local AS. The RLD TLV, representing the business relationship between the local AS and the neighbor AS of the outgoing route, SHOULD also be prepended to the Adj-rib-out at the egress node of the local AS. The BMP server, by analyzing the above two RLD TLVs of the same route, can use the rules defined in RFC7908 [RFC7908] to detect the existence of any route leak. As example is shown in Figure 2.



node), the BMP server can do validate the consistency of the iOTC attribute with the RLD. If contradiction detected, the BMP server may further check the bussiness contract for the actual business relationship.

- o For special prefixes that does not obey the peer/AS level business relationship negotiated through BGP Open Message, the BMP server can use the RLD TLV to detect such routes since the RLD TLV is set at prefix level.
- o For devices that do not support RLP, using RLD to collect the BGP routes, which conveys the RLD information from upstream ASes, allows the BMP server to detect and mitigate the route leaks that happen in the upstream AS. In other words, the detection and mitigation process can be also done in the BMP server, should the BMP server collects the BGP Update messages at the ingress or egress nodes.

#### 5. Acknowledgements

#### 6. Contributors

Haibo Wang

Huawei

Email: rainsword.wang@huawei.com

#### 7. IANA Considerations

This document defines the following new BMP Route Monitoring message TLV type (Section 4.1):

- o Type = TBD1, the RLD TLV represents the prefix-level business relationship between the transmitter AS and the receiver AS. The local AS is a transmitter or a receiver, depending on if the route is an incoming route from a neighbor AS or an outgoing route to a neighbor AS.

#### 8. Security Considerations

It is not believed that this document adds any additional security considerations.

## 9. References

### 9.1. Normative References

- [I-D.ietf-grow-bmp-adj-rib-out]  
Evens, T., Bayraktar, S., Lucente, P., Mi, K., and S. Zhuang, "Support for Adj-RIB-Out in BGP Monitoring Protocol (BMP)", draft-ietf-grow-bmp-adj-rib-out-06 (work in progress), June 2019.
- [I-D.ietf-grow-route-leak-detection-mitigation]  
Sriram, K. and A. Azimov, "Methods for Detection and Mitigation of BGP Route Leaks", draft-ietf-grow-route-leak-detection-mitigation-00 (work in progress), April 2019.
- [I-D.ietf-idr-bgp-open-policy]  
Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention using Roles in Update and Open messages", draft-ietf-idr-bgp-open-policy-05 (work in progress), February 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", RFC 7854, DOI 10.17487/RFC7854, June 2016, <<https://www.rfc-editor.org/info/rfc7854>>.
- [RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <<https://www.rfc-editor.org/info/rfc7908>>.

### 9.2. Informative References

- [Luckie] claffy, M. L. M. L. A. D. V. G. K., "AS Relationships, Customer Cones, and Validation", October 2013.

[Siddiqui]

Ramirez, M. S. S. D. M. M. Y. R. S. X. M. W., "Route Leak  
Detection Using Real-Time Analytics on local BGP  
Information", 2014.

Authors' Addresses

Yunan Gu  
Huawei  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: gyunan@huawei.com

Huanan Chen  
China Telecom Co., Ltd.  
109 Zhongshan W Ave  
Guangzhou 510630  
China

Email: chenhn8.gd@chinatelecom.cn

Di Ma  
ZDNS  
4 South 4th St. Zhongguancun  
Beijing, Haidian  
China

Email: madi@zdns.cn

Shunwan Zhuang  
Huawei  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: zhuangshunwan@huawei.com