Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

A. Azimov
E. Bogomazov
Qrator Labs
R. Bush
Internet Initiative Japan
K. Patel
Arrcus, Inc.

March 13, 2017

Route Leak Detection and Filtering using Roles in Update and Open messages
draft-ymbk-idr-bgp-eotr-policy-00

Abstract

[draft-ymbk-idr-bgp-open-policy] defines a BGP OPEN capability and consequent route marking which enforces a valley-free peering relationship. This document defines an eOTC (external Only To Customer) transitive BGP attribute which propagates the specific marking to automatically detect route leaks. The goal is to allow a distant AS to determine a violation of valley-free peering.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 14, 2017.

Table of Contents

1. Introduction

For the purpose of this document, BGP route leaks are when a BGP route was learned from transit provider or peer is announced to another provider or peer. See [I-D.ietf-grow-route-leak-problem-definition]. These are usually the result of misconfigured or absent BGP route filtering or lack of coordination between two BGP speakers.

[I-D.ietf-idr-route-leak-detection-mitigation] describes a method of marking and detecting leaks which relies on operator maintained markings. Unfortunately, in most cases, a leaking router will likely also be misconfigured to mark incorrectly.

It has been suggested to use white list filtering, relying on knowing the prefixes in the peer's customer cone as import filtering, in order to detect route leaks. Unfortunately, a large number of medium transit operators use a single prefix list as only the ACL for export filtering, without community tagging and without paying attention to the source of a learned route. So, if they learn a customer's route from their provider or peer - they will announce it in all

directions, including other providers or peers. This misconfiguration affects a limited number of prefixes; but such route leaks will obviously bypass customer cone import filtering made by upper level upstream providers.

This document specifies a way to to create automatic filters for detection of route leaks via new BGP Path Attribute which is set according to BGP Roles ([I-D.ymbk-idr-bgp-open-policy]) . While iOTC provides strong vendor-code-based enforcement of route leak prevention, route leaks could still exist as result of misconfigured old BGP implementations. Route leaks could also be result of malicious activity such as MITM attacks or DoS. The goal of this proposal is to allow a distant AS to determine a violation of valley-free peering that is made by mistake or by purpose.

2. BGP External Only To Customer attribute

The External Only To Customer (eOTC) attribute is a new optional, transitive BGP Path attribute with the Type Code <TBD1>. This attribute is four bytes and contains an AS number of the AS that added the attribute to the route.

There are four rules for setting the eOTC attribute:

1. If eOTC is not set and the sender's Role is Provider or Peer, the eOTC attribute MUST be added with value equal to the sender's AS number.

2. If eOTC is not set and the sender's Role is Complex and the prefix role is Provider or Peer, the eOTC attribute MUST be added with value equal to to the sender's AS number.

3. If eOTC is set, the receiver's Role is Provider or Peer, and its value is not the neighbor's AS number then the incoming route is route leak and MUST be given a lower local preference, or MAY be dropped.

4. If eOTC is set, the receiver's Role is Complex, the prefix role Role is Provider or Peer, and the eOTC value is not equal to the neighbor's AS number, then the incoming route is a route leak and MUST be given a lower local preference, or they MAY be dropped.

These four rules provide mechanism for route leak detection that is created by a distant party in the AS_Path.

3. Compatibility with BGPsec

For BGPsec [I-D.ietf-sidr-bgpsec-protocol] enabled routers, the Flags field will have a bit added to indicate that an eOTC attribute exists. The eOTC value will be automatically carried in AS field of the added Secure_Path Segment.

When a route is translated from a BGPsec enabled router to a non-BGPsec router, in addition to AS_PATH reconstruction, reconstruction MUST be performed for the eOTC attribute. If Flag bit was set in one of Secure_Path Segments, the eOTC attribute SHOULD be added with the AS number of the segment in which it appears for the first time.

4. IANA Considerations

This document defines a new optional, transitive BGP Path Attributes option, named "External Only To Customer", assigned value <TBD1> [To be removed upon publication: http://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml#bgp-parameters-2] [RFC4271]. The length of this attribute is 4.

5. Security Considerations

This document proposes a mechanism for detection of route leaks that are the result of BGP policy misconfiguration. If BGPSec is enabled it will also provide mechanism to detect leaks that are result of malicious activity.

Deliberate mis-marking of the eOTC flag could be used to affect the BGP decision process, but could not sabotage a route's propagation.

eOTC is a transitive BGP AS_PATH attribute which reveals a information about a BGP speaker's peering relationship. It will give a strong hint that some link isn't customer to provider, but will not help to distinguish if it is provider to customer or peer to peer. In addition it could reveal sequence of p2c to downstream ISPs. If eOTC is BGPsec signed, it can not be removed for peering confidentiality.

Still, any Tier-1 number in AS_PATH could be used in the same way to reveal possible p2c sequence.

6. References

6.1. Normative References

[I-D.ymbk-idr-bgp-open-policy]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Detection and Filtering using Roles in Update and Open messages", draft-ymbk-idr-bgp-open-policy-02 (work in progress), November 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://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, <http://www.rfc-editor.org/info/rfc4271>.

6.2. Informative References

[I-D.ietf-grow-route-leak-problem-definition]
Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", draft-ietf-grow-route-leak-problem-definition-06 (work in progress), May 2016.

[I-D.ietf-idr-route-leak-detection-mitigation]
Sriram, K., Montgomery, D., Dickson, B., Patel, K., and A. Robachevsky, "Methods for Detection and Mitigation of BGP Route Leaks", draft-ietf-idr-route-leak-detection-mitigation-03 (work in progress), May 2016.

[I-D.ietf-sidr-bgpsec-protocol]
Lepinski, M. and K. Sriram, "BGPsec Protocol Specification", draft-ietf-sidr-bgpsec-protocol-15 (work in progress), March 2016.

Authors' Addresses

Alexander Azimov
Qrator Labs

Email: aa@qrator.net

Eugene Bogomazov
Qrator Labs

Email: eb@qrator.net

Randy Bush
Internet Initiative Japan

Email: randy@psg.com

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com