

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 21, 2014

J. Dong
M. Chen
Huawei Technologies
June 19, 2014

Extensions to RT-Constrain in Hierarchical Route Reflection Scenarios
draft-dong-idr-rtc-hierarchical-rr-00

Abstract

The Route Target (RT) Constrain mechanism specified in RFC 4684 is used to build a route distribution graph in order to restrict the propagation of Virtual Private Network (VPN) routes. In network scenarios where hierarchical route reflection (RR) is used, the existing RT-Constrain mechanism cannot build a correct route distribution graph. This document refines the route distribution rules of RT-Constrain to address the hierarchical RR scenarios.

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 <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 December 21, 2014.

Copyright Notice

Copyright (c) 2014 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	2
2. Proposed Solution	3
3. IANA Considerations	4
4. Security Considerations	4
5. Acknowledgements	4
6. Normative References	4
Authors' Addresses	4

1. Introduction

The Route Target (RT) Constrain mechanism specified in RFC 4684 is used to build a route distribution graph in order to restrict the propagation of Virtual Private Network (VPN) routes. In network scenarios where hierarchical route reflection (RR) is used, the existing RT-Constrain mechanism cannot build a correct route distribution graph.




Figure 1. RT-Constrain with Hierarchical RR

As shown in Figure 1, hierarchical RRs are deployed in the network, RR-2 and RR-3 are route-reflectors of their connecting PEs, and are also the clients of RR-1. If each PE advertises RT membership information of RT-1 to the upstream RR, after the best path selection, both RR-2 and RR-3 would create the CLUSTER_LIST attribute, prepend their local CLUSTER_ID and then advertise the best path to RR-1 and their clients respectively.

On receipt of the RT-Constrain routes from RR-2 and RR-3, RR-1 will select one of the received routes as the best route, assume the route received from RR-2 is selected by RR-1 as the best path. Then RR-1 needs to advertise the best path to both RR-2 and RR-3 to create the route distribution graph of VPN-1. RR-1 would prepend its CLUSTER_ID to the CLUSTER_LIST of the path, and according to the rules in Section 3.2 of [RFC4684], it sets the ORIGINATOR_ID to its own router-id, and the NEXT_HOP to the local address for the session. Then RR-1 would advertise this route to both RR-2 and RR-3. On receipt of the RT-Constrain route from RR-1, RR-2 checks the CLUSTER_LIST and find its own CLUSTER_ID in the list, so this route will be ignored by RR-2. As a result, RR-2 will not form the outbound filter of RT-1 towards RR-1, hence will not advertise VPN routes with RT-1 to RR-1.

2. Proposed Solution

The problem described in the above section is that the best path is sent back to the BGP speaker which advertised the path and get discarded due to the BGP loop detection mechanisms. Since the advertisement of RT-Constrain route is to set up a route distribution graph and not to guide the data packet forwarding, all the available paths can be considered in setting up the route distribution graph, not just the best path. Thus in addition to the rules specified in section 3.2 of [RFC4684], the following rule applies in the advertisement of RT-Constrain routes:

- o When advertising an RT membership NLRI to a route-reflector client, if the best route as selected by the path selection procedure described in Section 9.1 of [RFC4271] is the path received from this client, and there are alternative paths received from other peers, the most disjoint alternative route SHOULD be advertised to that client; The most disjoint alternative path is the path whose CLUSTER_LIST and ORIGINATOR_ID attributes are different from the attributes of the best path.

With this additional rule, RR-1 in Figure 1 would advertise to RR-2 the RT-Constrain route received from RR-3, although the best route is received from RR-2. Thus RR-2 will not discard the RT-constrain

route received from RR-1, and the route distribution graph can be set up completely.

3. IANA Considerations

This document makes no request of IANA.

4. Security Considerations

This document does not change the security properties of BGP based VPNs and [RFC4684].

5. Acknowledgements

The authors would like to thank Yaqun Xiao for the discussion about RT-Constrain in hierarchical RR scenario.

6. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, November 2006.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: mach.chen@huawei.com