

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: May 29, 2015

J. Dong  
M. Chen  
Huawei Technologies  
R. Raszuk  
Mirantis Inc.  
November 25, 2014

Extensions to RT-Constrain in Hierarchical Route Reflection Scenarios  
draft-dong-idr-rtc-hierarchical-rr-02

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 describes the problem scenario and proposes a solution to address the RT-Constrain issue in 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 May 29, 2015.

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. Problem Statement . . . . .	2
3. Proposed Solution . . . . .	3
3.1. Add-path Based Solution . . . . .	4
4. IANA Considerations . . . . .	4
5. Security Considerations . . . . .	4
6. Acknowledgements . . . . .	4
7. Normative References . . . . .	5
Appendix A. Another Possible Solution . . . . .	5
Authors' Addresses . . . . .	6

1. Introduction

The Route Target (RT) Constrain mechanism specified in [RFC4684] 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 advertisement rules of RT membership information as defined in section 3.2 of [RFC4684] cannot guarantee a correct route distribution graph.

This document describes the problem scenario and proposes a solution to address the RT-Constrain issue in hierarchical RR scenarios.

2. Problem Statement




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, here 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.

### 3. Proposed Solution

### 3.1. Add-path Based Solution

During the discussion in the IDR working group, the add-path based solution is proposed. It makes use of the add-path mechanism as defined in [I-D.ietf-idr-add-paths] for RTC route advertisement. The solution is summarized as follows:

- o The route-reflector clients which themselves are also route-reflectors SHOULD be identified, then BGP add-paths [I-D.ietf-idr-add-paths] SHOULD be enabled for RT membership NLRI on the BGP sessions between the higher layer RR and the lower layer RRs to ensure that sufficient RT-Constrain routes can be advertised by the higher layer RR to the lower layer RRs to pass BGP loop detection. In this case normal BGP path advertisement rules as defined in [RFC4271] SHOULD be applied. The number of RT-Constrain routes to be advertised is a local decision of operators.
- o When advertising an RT membership NLRI to a route-reflector client which is not a lower layer RR, the advertisement rule as defined in section 3.2 of [RFC4684] SHOULD be applied.

With the above advertisement rule, RR-1 in figure 1 SHOULD advertise to RR-2 the RT-Constrain routes received from both RR-2 and RR-3, then the RTC route from RR-3 will pass the BGP loop detection on RR-2, thus the route distribution graph can be set up correctly.

### 4. IANA Considerations

This document makes no request of IANA.

### 5. Security Considerations

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

### 6. Acknowledgements

The authors would like to thank Yaqun Xiao for the discussion about RT-Constrain in hierarchical RR scenario. Many people have made valuable comments and suggestions, including Susan Hares, Jeffrey Haas, Stephane Litkowski, Vitkovsky Adam, Xiaohu Xu, Uttaro James, Shyam Sethuram and Saikat Ray.

## 7. Normative References

- [I-D.ietf-idr-add-paths]  
Walton, D., Retana, A., Chen, E., and J. Scudder,  
"Advertisement of Multiple Paths in BGP", draft-ietf-idr-  
add-paths-10 (work in progress), October 2014.
- [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.

## Appendix A. Another Possible Solution

This section provides another possible solution which was discussed among authors and IDR participants.

Since the advertisement of RT-Constrain route is to set up a route distribution graph and not to guide the data packet forwarding, actually all the available RT-Constrain routes should be considered in setting up the route distribution graph, not just the best one. Thus the following advertisement rule for RT membership information is proposed to replace the rule i and ii in section 3.2 [RFC4684]:

- o When advertising an RT membership NLRI to a route-reflector peer (either client or non-client), if the best path as selected by the path selection procedure described in Section 9.1 of [RFC4271] is the path received from this peer, and there are alternative paths received from other peers, then the most disjoint alternative route SHOULD be advertised to this peer. The most disjoint alternative path is the path whose CLUSTER\_LIST and ORIGINATOR\_ID attributes are diverse from the attributes of the best path.

With the above advertisement 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 correctly.

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

Robert Raszuk  
Mirantis Inc.  
615 National Ave. #100  
Mt View, CA 94043  
USA

Email: robert@raszuk.net