Workgroup: idr
Internet-Draft:
draft-zzhang-idr-rt-derived-community-02
Published: 4 March 2022
Intended Status: Standards Track
Expires: 5 September 2022
Authors: Z. Zhang J. Haas K. Patel
Juniper Networks Juniper Networks Arrcus

Extended Communities Derived from Route Targets

Abstract

This document specifies a way to derive an Extended Community from a Route Target and describes some example use cases.

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 5 September 2022.

Copyright Notice

Table of Contents

1. Introduction

Consider a VPN with 10 PEs. A Route Target (say RT1) [RFC4360] is configured for the VPN and all PEs will import VPN routes with RT1 attached. The RT is an Extended Community (say EC1), with its sub-type being 0x02. While RT1 and EC1 have the same encoding, typically when we mention a Route Target, its property of being able to control the route propagation and importation is implied. When we just mention an Extended Community, that property is not implied.

Now consider that another BGP route needs to be imported by some but not all those PEs. The route could be of any SAFI/type (does not need to be a VPN prefix), but it needs to be associated with the VPN on those importing PEs. The exact meaning of "association" here does not matter, but the key is that those PEs need to know that the route is related to that VPN.

To control the propagation to and importation by those PEs, a different Route Target (say RT3) is attached to the route. For those PE to associate the route with the VPN, an Extended Community (say EC2) is attached. Even though RT1/EC1 is already associated with the VPN, EC2 needs to be different from RT1/EC1, because if EC1 was used, the route would be propagated to and imported by all the 10 PEs. EC2 cannot be the same as RT3 either, because there could be other routes to be propagated to and imported by those same set of PEs yet those other routes are not related to the VPN.

While EC2 can be any Extended Community (that is not a RT) configured on the originating and receiving PEs to map it to the VPN, it is convenient if EC2 is derived from the RT1/EC1, e.g. the sub-type of RT1/EC1 is changed to a new known value while everything else remains the same. We call this a Route Target derived Extended Community, or RT-derived EC. A new sub-type is assigned specifically for this purpose (see IANA considerations).

This document only specifies a way to derive an Extended Community from a Route Target Extended Community using IANA-assigned Extended Community sub-types (or Extended Community Type in case of IPv6-Address-Specific Extended Community). Any protocol/feature that can take advantages of the convenience of generic derivation may use them, or not use them at its own discretion, and how they are used is outside the scope of this document.

2. Use Cases

The following are a few examples of use cases. To reiterate, these are example scenarios where generic RT-derived ECs could be used (when the routes to which they are attached provide enough context). It is not the intention of this document to mandate that it must be used.

2.1. EVPN EVI-RT Extended Community

Section 9.5 "EVI-RT Extended Community" of [I-D.ietf-bess-evpn-igmp-mld-proxy] describes a situation similar to the above. As a solution, four EVPN specific EVI-RT ECs are defined, each mapping to a type of Route Target for the corresponding EVPN instance.

As a theoretical alternative, a RT-derived EC described in this document could be used instead - just derive a generic EC from the EVI RT. Note that this document does not attempt to change the existing procedures in [I-D.ietf-bess-evpn-igmp-mld-proxy], but merely use it for illustration purposes.

2.2. Leaf Discovery with Controller Signaled BGP-MVPN

In Section 2 "Alternative to BGP-MVPN" of [I-D.ietf-bess-bgp-multicast-controller], BGP MCAST-TREE SAFI signaling can be used for a controller to program multicast forwarding state in VRFs of ingress/egress PEs, instead of relying on distributed BGP-MVPN signaling. For the controller to learn egress PEs of a VPN customer multicast tree (so that it can build/find a corresponding provider tunnel), egress PEs signal leaf information to the controller via Leaf Auto-Discovery routes. The routes carry a Route Target for the controller (so that only the controller receives them), and an EC derived from the VPN's Route Target (so that the controller knows which VPN they are for).

2.3. Translated Route-target Extended Communities in [I-D.ietf-idr-legacy-rtc]

In Section 3.1 of [I-D.ietf-idr-legacy-rtc], a similar mechanism is described, as quoted below:

"The translation of the IRTs is necessary in order to refrain from importing "route-filter" VRF routes into VPN VRFs that would import the same route-targets. The translation of the IRTS is done as follows. For a given IRT, the equivalent translated RT (TRT) is constructed by means of swapping the value of the high-order octet of the Type field for the IRT (as defined in [RFC4360])."

3. Security Considerations

This document specifies a way to derive an Extended Community from a Route Target Extended Community and does not specify how derived Extended Communities are used. As a result, this document does not need security considerations. Any potential security concerns need be addressed by documents that specify the actual usage.

4. IANA Assignments

IANA has assign a new sub-type "RT-derived-EC" with value 0x15 in the following registries:

- *Transitive Two-Octet AS-Specific Extended Community Sub-Types
- *Transitive Four-Octet AS-Specific Extended Community Sub-Types
- *Transitive IPv4-Address-Specific Extended Community Sub-Types
- *Non-Transitive Opaque Extended Community Sub-Types
- *EVPN Extended Community Sub-Types

IANA has also assigned a new type "RT-derived-EC" with value 0x0015 in the following registry:

- *Transitive IPv6-Address-Specific Extended Community Types

If and when additional Extended Community types are defined with a Route Target sub-type, the "RT-derived-EC" sub-type may also be registered for those new types, preferably with the same value.

5. References

5.1. Normative References

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <https://www.rfc-editor.org/info/rfc4360>.

5.2. Informative References

[I-D.ietf-bess-evpn-igmp-mld-proxy] Sajassi, A., Thoria, S., Mishra, M., Drake, J., and W. Lin, "IGMP and MLD Proxy for EVPN", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-igmp-mld-proxy-19, 4 March 2022, <https://www.ietf.org/archive/id/draft-ietf-bess-evpn-igmp-mld-proxy-19.txt>.

[I-D.ietf-bess-bgp-multicast-controller] Zhang, Z., Raszuk, R., Pacella, D., and A. Gulko, "Controller Based BGP Multicast Signaling", Work in Progress, Internet-Draft, draft-ietf-bess-bgp-multicast-controller-07, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-bess-bgp-multicast-controller-07.txt>.

[I-D.ietf-idr-legacy-rtc] Mohapatra, P., Sreekantiah, A., Patel, K., Pithawala, B., and A. Lo, "Automatic Route Target Filtering for legacy PEs", Work in Progress, Internet-Draft, draft-ietf-idr-legacy-rtc-08, 12 September 2017, <https://www.ietf.org/archive/id/draft-ietf-idr-legacy-rtc-08.txt>.

Authors' Addresses

Zhaohui Zhang
Juniper Networks

Email: zzhang@juniper.net

Jeff Haas
Juniper Networks

Email: jhaas@juniper.net

Keyur Patel
Arrcus

Email: keyur@arrcus.com