

Workgroup: idr
Internet-Draft:
draft-zzhang-idr-rt-derived-community-04
Published: 7 February 2023
Intended Status: Standards Track
Expires: 11 August 2023
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 11 August 2023.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Route Target Type/sub-type Conventions](#)
- [3. Use Cases](#)
 - [3.1. EVPN EVI-RT Extended Community](#)
 - [3.2. Leaf Discovery with Controller Signaled BGP-MVPN](#)
 - [3.3. Translated Route-target Extended Communities in \[I-D.ietf-idr-legacy-rtc\]](#)
- [4. Security Considerations](#)
- [5. IANA Assignments](#)
- [6. Acknowledgements](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Authors' Addresses](#)

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 [[RFC5701](#)]). 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. Route Target Type/sub-type Conventions

It may be expected by some people that Route Targets are Extended Communities with sub-type 0x02 (or with Type 0x0002 in case of IPv6 Address Specific Extended Community). However, the only official specification are in [[RFC7153](#)] [[RFC7432](#)] and are only for the following types:

- *Type 0x00 (Transitive Two-Octet AS-Specific EC)
- *Type 0x01 (Transitive IPv4-Address-Specific EC)
- *Type 0x02 (Transitive Four-Octet AS-Specific EC)
- *Type 0x06 (EVPN AS-Specific EC)
- *Type 0x0002 (Transitive IPv6-Address-Specific Route Target)
- *Type 0x0011 (Transitive IPv6-Address-Specific EC, UUID-based Route Target))

While it may be desired to follow the unwritten convention and assign sub-type 0x02 for future Route Targets of future types of ECs, there is no guarantee of that. For example, Type 0x0011 is assigned for UUID-based Route Target that imposes as an IPv6 Address Specific EC (even though UUID is not an IPv6 address).

IANA has assigned sub-type 0x15 (or Type 0x0015 in case of IPv6 Address Specific EC) to indicate an EC is derived from a Route Target that has sub-type 0x02 (or type 0x0002 in case of IPv6 Address Specific EC), and this document will further request another Type TBD (say, 0x0016) to derive from the UUID-based Route Target.

All those can only be registered with the known types listed above. When a new type is defined and registered, the corresponding 0x02 sub-type may be registered for Route Target purpose or for something else, and there is no guarantee that the 0x15 sub-type will not be registered for something else as well (than for RT derivation). As a result, the mapping between sub-type 0x02 and 0x15, type 0x0002 and

0x0015, type 0x0011 and TBD are only defined for the known types listed above.

Of course, when a new type is defined and registered, it is desired to proactively register sub-type 0x02 and 0x15 at the very beginning for Route Target and Route Target Derivation purposes, should the review process catch it.

3. 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.

3.1. 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.

3.2. 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).

3.3. 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])."

4. 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.

5. IANA Assignments

This document requests IANA to assign a new Type value (0x0016 suggested) for "UUID-RT-derived-EC".

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.

6. Acknowledgements

The authors thank Robert Raszuk for his valuable comments and suggestions.

7. References

7.1. 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>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", RFC 7153, DOI 10.17487/RFC7153, March 2014, <<https://www.rfc-editor.org/info/rfc7153>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.

7.2. Informative References

- [I-D.ietf-bess-evpn-igmp-mld-proxy] Sajassi, A., Thoria, S., Mishra, M. P., Drake, J., and W. Lin, "Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-igmp-mld-proxy-21, 22 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-bess-evpn-igmp-mld-proxy-21.txt>>.
- [I-D.ietf-bess-bgp-multicast-controller] Zhang, Z. J., Raszuk, R., Pacella, D., and A. Gulko, "Controller Based BGP Multicast Signaling", Work in Progress, Internet-Draft, draft-ietf-bess-bgp-multicast-controller-09, 11 April 2022, <<https://www.ietf.org/archive/id/draft-ietf-bess-bgp-multicast-controller-09.txt>>.
- [I-D.ietf-idr-legacy-rtc] Mohapatra, P., Sreekantiah, A., Patel, K., Burjiz, 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