

GROW  
Internet-Draft  
Updates: [1997](#), [7947](#) (if approved)  
Intended status: Standards Track  
Expires: April 11, 2018

N. Hilliard  
INEX  
W. Hargrave  
LONAP  
October 8, 2017

The BGP NO\_EXPORT\_VIA\_RS Community for Route Servers  
draft-hilliard-grow-no-export-via-rs-00

## Abstract

This document describes a BGP Well-known Community called NO\_EXPORT\_VIA\_RS. This community allows BGP route server clients to instruct an Internet Exchange BGP Route Server to tag prefixes destined for other route server clients with the NO\_EXPORT Well-Known Community. This mechanism allows route server clients to transitively control distribution of their prefixes in other Autonomous Systems connected to the Internet Exchange Route Server.

## 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 [[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 <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 April 11, 2018.

## Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

BGP Large Communities

October 2017

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 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

|                      |   |                   |
|----------------------|---|-------------------|
| <a href="#">1.</a>   | Introduction . . . . .                            | <a href="#">2</a> |
| <a href="#">2.</a>   | The BGP NO_EXPORT_VIA_RS Community . . . . .      | <a href="#">3</a> |
| <a href="#">3.</a>   | Operator Implementation Recommendations . . . . . | <a href="#">4</a> |
| <a href="#">4.</a>   | Vendor Implementation Recommendations . . . . .   | <a href="#">4</a> |
| <a href="#">5.</a>   | Security Considerations . . . . .                 | <a href="#">4</a> |
| <a href="#">6.</a>   | IANA Considerations . . . . .                     | <a href="#">4</a> |
| <a href="#">7.</a>   | References . . . . .                              | <a href="#">4</a> |
| <a href="#">7.1.</a> | Normative References . . . . .                    | <a href="#">4</a> |
| <a href="#">7.2.</a> | Informative References . . . . .                  | <a href="#">5</a> |
|                      | Authors' Addresses . . . . .                      | <a href="#">5</a> |

## [1.](#) Introduction

Internet Exchange Route Servers [[RFC7947](#)] use BGP to broker network reachability information among their clients. BGP operators often wish to control distribution of their prefixes in adjacent autonomous systems. When using bilateral BGP sessions, prefix distribution control can be achieved using the BGP NO\_EXPORT community defined in [[RFC1997](#)].

[[RFC1997](#)] states that the NO\_EXPORT community must not be announced outside a BGP confederation boundary, so all BGP speakers that strictly interpret this document will refuse to forward prefixes tagged with this community to a EBGp peer.

Some Internet Exchange route server operators ignored the strict interpretation of [[RFC1997](#)] on this point because although it is technically more correct to interpret the NO\_EXPORT community on the route server, the point of a route server is to act as a broker rather than a router. Consequently it has been argued that it is more useful for route server clients if prefixes tagged with the

NO\_EXPORT community are passed unaltered through the route server rather than being blocked. This approach gives route server clients the flexibility of being able to use the NO\_EXPORT community to signal adjacent ASNs, even if this behaviour is not technically correct according to [\[RFC1997\]](#).

As there is no consensus among IXP route server operators about which is the more appropriate behaviour, the Internet Exchange Route Server [\[RFC7947\]](#) specification document stayed quiet on the issue, noting only in [Section 2.2.4](#) that a route server may interpret, modify or remove specific BGP communities, as defined by local policy. This formally allowed the practice of treating the NO\_EXPORT community as a transparent transitive attribute, but did not define whether NO\_EXPORT should be interpreted or passed through the route server.

This allowed an operational ambiguity to continue, namely that there was no consistent way for a route server client to limit propagation of any prefixes announced to a route server.

This document resolves this ambiguity by creating a new BGP Well-Known Community called NO\_EXPORT\_VIA\_RS, which allows BGP route server clients to instruct an Internet Exchange BGP Route Server to tag prefixes destined for other route server clients with the NO\_EXPORT Well-Known Community. This mechanism provides an unambiguous and consistent way for route server clients to transitively control distribution of their prefixes in other Autonomous Systems connected to the Internet Exchange Route Server.

## [2.](#) The BGP NO\_EXPORT\_VIA\_RS Community

The BGP NO\_EXPORT\_VIA\_RS Community is a community interpreted only by [\[RFC7947\]](#) Internet Exchange Route Servers. If a route server client announces a prefix tagged with NO\_EXPORT\_VIA\_RS to a route server, the route server MUST attach the [\[RFC1997\]](#) NO\_EXPORT community to all outbound announcements of that prefix to other route server clients. The route server SHOULD strip the NO\_EXPORT\_VIA\_RS community from the outbound prefix announcement.

If a route server client announces a prefix tagged with both the NO\_EXPORT and NO\_EXPORT\_VIA\_RS communities to a route server, the route server MUST ignore the NO\_EXPORT community, and otherwise MUST handle the prefix and its attributes as specified in the previous

paragraph.

If a BGP speaker that is not a route server receives a prefix tagged with NO\_EXPORT\_VIA\_RS from a BGP peer, the BGP speaker MUST ignore the NO\_EXPORT\_VIA\_RS community, MUST NOT treat the UPDATE as an error according to [[RFC7606](#)] and SHOULD strip the NO\_EXPORT\_VIA\_RS community from the prefix.

This document formally updates the NO\_EXPORT definition in [[RFC1997](#)] and overrides the ambiguity in [Section 2.2.4 of \[RFC7947\]](#) in respect of this specific community.

### [3.](#) Operator Implementation Recommendations

It is possible to implement the semantics described above in many BGP implementations by using standard BGP policy language statements.

### [4.](#) Vendor Implementation Recommendations

Route server implementations MUST provide a configuration switch to implement the behaviour specified in [Section 2](#). This configuration switch SHOULD be enabled by default.

### [5.](#) Security Considerations

The purpose of the NO\_EXPORT\_VIA\_RS BGP Community is to control the NO\_EXPORT Community, so there are no additional security considerations outside those associated with the NO\_EXPORT community.

Network administrators should note the recommendations in [Section 11](#) of BGP Operations and Security [[RFC7454](#)].

### [6.](#) IANA Considerations

IANA is requested to register NO\_EXPORT\_VIA\_RS in the "BGP Well-known Communities" registry.

NO\_EXPORT\_VIA\_RS (= TBD) (IANA suggested: 0xFFFFF05)

### [7.](#) References

## 7.1. Normative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", [RFC 1997](#), DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", [RFC 7606](#), DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.

Hilliard & Hargrave

Expires April 11, 2018

[Page 4]

---

Internet-Draft

BGP Large Communities

October 2017

- [RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", [RFC 7947](#), DOI 10.17487/RFC7947, September 2016, <<https://www.rfc-editor.org/info/rfc7947>>.

## 7.2. Informative References

- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", [BCP 194](#), [RFC 7454](#), DOI 10.17487/RFC7454, February 2015, <<https://www.rfc-editor.org/info/rfc7454>>.

### Authors' Addresses

Nick Hilliard  
Internet Neutral Exchange Association CLG  
4027 Kingswood Road  
Dublin 24  
Ireland

Email: [nick@inex.ie](mailto:nick@inex.ie)

Will Hargrave  
LONAP Ltd  
5 Fleet Place  
London EC4M 7RD  
United Kingdom

Email: [will@lonap.net](mailto:will@lonap.net)