Network Working Group
Internet Draft
Intended status: Standards Track
July 4, 2011
Expires: Jan 4, 2012

Wim Henderickx
Adam Simpson
Alcatel-Lucent
July 4, 2011

Generalized Redirect Action in BGP Flow Specification Routes
draft-simpson-idr-flowspec-redirect-00.txt

Status of this Memo

Copyright Notice

Abstract

Flowspec is an extension to BGP that allows for the dissemination of traffic flow specifications. This has several applications, but one of key interest to many network operators is network-wide distribution of traffic filtering rules as part of a threat mitigation strategy.

Every flowspec route is effectively a rule, consisting of a matching part (encoded in the NLRI field) and an action part. The current standards support common filter actions including discard, rate limit, sample, etc. and all of these actions are encoded in BGP extended communities. For policy-based forwarding the current standards also define a redirect-to-VRF action (again encoded in a BGP extended community), but for some flowspec applications this can be complex to implement, particularly in networks where L3 VPNs are not prevalent.

This draft proposes a generalized flowspec redirect action that allows a more complete set of policy-based forwarding actions to be signaled with a flowspec route. This generalized action is encoded in a BGP path attribute and uses a TLV-style encoding for future extensibility. Two redirect action TLVs are defined in this draft: one for redirecting matched packets towards a remote IPv4 destination and the other for redirecting matched packets towards a remote IPv6 destination. Many routers already support these filter actions in the datapath and so the proposed flowspec extensions are simply filling a control plane gap.

Table of Contents

1. Introduction.....3
2. Terminology.....3
3. The Generalized Flowspec Redirect Attribute.....3
4. Security Considerations.....5
5. IANA Considerations.....5
6. References.....6
   6.1. Normative References.....6
   6.2. Informative References.....6
7. Acknowledgments.....6

1. Introduction

Flowspec is an extension to BGP that allows for the dissemination of traffic flow specifications. This has several applications, but one of key interest to many network operators is network-wide distribution of traffic filtering rules as part of a threat mitigation strategy.

Every flowspec route is effectively a rule, consisting of a matching part (encoded in the NLRI field) and an action part. The current standards support common filter actions including discard, rate limit, sample, etc. and all of these actions are encoded in BGP extended communities. For policy-based forwarding the current standards also define a redirect-to-VRF action (again encoded in a BGP extended community), but for some flowspec applications this can be complex to implement, particularly in networks where L3 VPNs are not prevalent.

This draft proposes a generalized flowspec redirect action that allows a more complete set of policy-based forwarding actions to be signaled with a flowspec route. This generalized action is encoded in a BGP path attribute and uses a TLV-style encoding for future extensibility. Two redirect action TLVs are defined in this draft: one for redirecting matched packets towards a remote IPv4 destination and the other for redirecting matched packets towards a remote IPv6 destination. Many routers already support these filter actions in the datapath and so the proposed flowspec extensions are simply filling a control plane gap.

2. Terminology

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].

3. The Generalized Flowspec Redirect Attribute

All of the actions defined in the current flowspec standards, [RFC5575] and [IPv6-FLOW], are encoded using BGP extended communities. While it would be desirable to define new BGP extended communities for the new types of redirect action called for in this document the maximum length of extended communities (7 octets of data) makes this very limiting.

This document therefore proposes a new BGP path attribute called the Generalized Flowspec Redirect Attribute. It is a transitive, optional attribute of non-extended length. The value field of the path

attribute contains exactly one Redirect Action TLV, which has the structure shown in Figure 1.

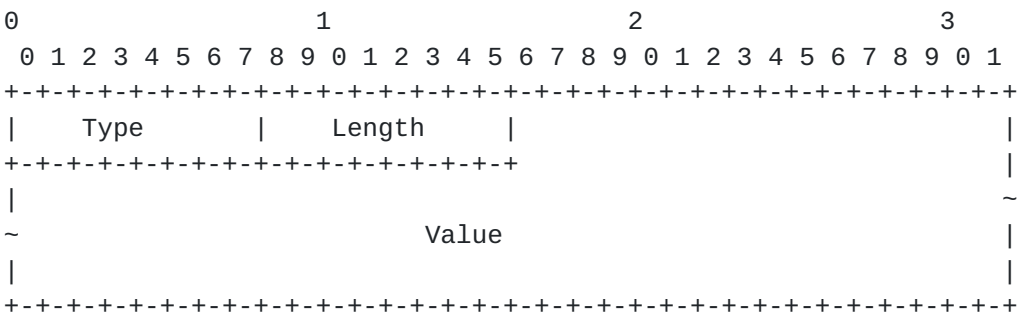

Redirect Action TLV

Figure 1

The Redirect Action TLV has the following fields:

- Type: A single octet encoding the TLV Type. Only type 1, "Remote IPv4 Address", and type 2 "Remote IPv6 Address" are defined in this document.

- Length: A single octet encoding the length in octets of the TLV, including the type and length fields. The length is encoded as an unsigned binary integer.

- Value: A value field that contains type-dependent data. In a type 1 "Remote IPv4 Address" TLV the value field contains a 4-octet AS value followed by a 4-octet IPv4 address. In a type 2 "Remote IPv6 Address" TLV the value field contains a 4-octet AS value followed by a 16-octet IPv6 address.

When a system originates a flowspec route with the intent to redirect matched packets to a remote IPv4 address, it includes a Generalized Flowspec Redirect Attribute containing a "Remote IPv4 Address" Redirect Action TLV (type 1). The AS number in the TLV MUST be the AS number of the associated IPv4 address (this will typically be the same as the originator's AS number) or 0, if the AS number is unknown to the originator. A 2-octet ASN is encoded in the low-order 2 octets of the AS number field. The IPv4 address in the TLV is any routable /32 unicast IPv4 address.

When a system originates a flowspec route with the intent to redirect matched packets to a remote IPv6 address, it includes a Generalized Flowspec Redirect Attribute containing a "Remote IPv6 Address" Redirect Action TLV (type 2). The AS number in the TLV MUST be the AS number of the associated IPv6 address (this will typically be the same as the originator's AS number) or 0, if the AS number is unknown to the originator. A 2-octet ASN is encoded in the low-order 2 octets of the AS number field. The IPv6 address in the TLV is any routable /128 unicast IPv6 address.

A flowspec route MUST NOT have more than one Generalized Flowspec Redirect Attribute. Error handling must follow the procedures outlined in [OPT-TRANS].

A flowspec route that has a Generalized Flowspec Redirect Attribute in addition to one or more of the BGP extended community actions defined in [RFC5575] and [IPV6-FLOW] is valid but implementations MAY choose to ignore some or all of the BGP extended community actions when installing a filter entry for this type of route.

A router that receives a flowspec route with a Generalized Flowspec Redirect Attribute MAY check that the AS number in the Redirect Action TLV (if non-zero) is the origin AS associated with its route to the indicated remote IP address. In this case, if the AS numbers are found to be different the router SHOULD NOT install a filter entry for the flowspec route.

When a router receives and installs a flowspec route with a Generalized Flowspec Redirect Attribute the resultant filter entry should forward matched packets to the interface that is the IP next-hop towards the signaled "Remote IPv4 Address" or "Remote IPv6 Address". The remote address may be any number of IP forwarding next-hops away from the router installing the flowspec route. In certain deployments the IP next-hop towards the remote IP address may be an IP or MPLS tunnel.

4. Security Considerations

TBD

5. IANA Considerations

This document requests that IANA allocate a new BGP path attribute type number for the Generalized Flowspec Redirect Attribute. IANA should also establish and maintain a registry for Redirect Action TLVs and indicate the meaning of type 1 and type 2 in this context.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

6.2. Informative References

[RFC5575] P. Marques, N. Sheth, R. Raszuk, B. Greene, J. Mauch, D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, August 2009.

[IPV6-FLOW] R. Raszuk, B. Pithawala, D. McPherson, "Dissemination of Flow Specification Rules for IPv6", draft-ietf-idr-flow-spec-v6-00, June 2011.

[OPT-TRANS] J. Scudder, E. Chen, "Error Handling for Optional Transitive BGP Attributes", draft-ietf-idr-optional-transitive-03, Sept 2010.

7. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
2018 Antwerp, Belgium
Email: wim.henderickx@alcatel-lucent.be

Adam Simpson
Alcatel-Lucent
600 March Road
Ottawa, Ontario K2K 2E6
Canada
Email: adam.simpson@alcatel-lucent.com