

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: November 23, 2021

Z. Li  
Huawei  
L. Ou  
Y. Luo  
China Telcom Co., Ltd.  
S. Lu  
Tencent  
G. Mishra  
Verizon Inc.  
H. Chen  
Futurewei  
S. Zhuang  
H. Wang  
Huawei  
May 22, 2021

**BGP Extensions for Routing Policy Distribution (RPD)**  
**draft-ietf-idr-rpd-11**

Abstract

It is hard to adjust traffic and optimize traffic paths in a traditional IP network from time to time through manual configurations. It is desirable to have a mechanism for setting up routing policies, which adjusts traffic and optimizes traffic paths automatically. This document describes BGP Extensions for Routing Policy Distribution (BGP RPD) to support this.

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](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

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 November 23, 2021.

## Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Problem Statement</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">Inbound Traffic Control</a>	<a href="#">4</a>
<a href="#">3.2.</a>	<a href="#">Outbound Traffic Control</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Protocol Extensions</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">Using a New AFI and SAFI</a>	<a href="#">6</a>
<a href="#">4.2.</a>	<a href="#">BGP Wide Community and Atoms</a>	<a href="#">8</a>
<a href="#">4.2.1.</a>	<a href="#">RouteAttr TLV/sub-TLV</a>	<a href="#">8</a>
<a href="#">4.2.2.</a>	<a href="#">Sub-TLVs of the Parameters TLV</a>	<a href="#">12</a>
<a href="#">4.3.</a>	<a href="#">Capability Negotiation</a>	<a href="#">13</a>
<a href="#">5.</a>	<a href="#">Operations</a>	<a href="#">14</a>
<a href="#">5.1.</a>	<a href="#">Application Scenario</a>	<a href="#">14</a>
<a href="#">5.2.</a>	<a href="#">About Failure</a>	<a href="#">16</a>
<a href="#">6.</a>	<a href="#">Contributors</a>	<a href="#">16</a>
<a href="#">7.</a>	<a href="#">Security Considerations</a>	<a href="#">17</a>
<a href="#">8.</a>	<a href="#">Acknowledgements</a>	<a href="#">17</a>
<a href="#">9.</a>	<a href="#">IANA Considerations</a>	<a href="#">17</a>
<a href="#">9.1.</a>	<a href="#">Existing Assignments</a>	<a href="#">17</a>
<a href="#">9.2.</a>	<a href="#">Routing Policy Type Registry</a>	<a href="#">17</a>
<a href="#">9.3.</a>	<a href="#">RouteAttr Atom Type</a>	<a href="#">18</a>
<a href="#">9.4.</a>	<a href="#">Route Attributes Sub-TLV Registry</a>	<a href="#">18</a>
<a href="#">9.5.</a>	<a href="#">Attribute Change Sub-TLV Registry</a>	<a href="#">18</a>
<a href="#">10.</a>	<a href="#">References</a>	<a href="#">19</a>
<a href="#">10.1.</a>	<a href="#">Normative References</a>	<a href="#">19</a>
<a href="#">10.2.</a>	<a href="#">Informative References</a>	<a href="#">20</a>



Authors' Addresses . . . . . [20](#)

## **1. Introduction**

It is difficult to optimize traffic paths in a traditional IP network because of the following:

- o Complex and error prone configuration. Traffic can only be adjusted device by device. The configurations on all the routers that the traffic traverses need to be changed or added. There are already lots of policies configured on the routers in an operational network. There are different types of policies, which include security, management and control policies. These policies are relatively stable. However, the policies for adjusting traffic are dynamic. Whenever the traffic through a route is not expected, the policies to adjust the traffic for that route are configured on the related routers. It is complex and error prone to dynamically add or change the policies to the existing policies on the special routers to adjust the traffic.
- o Difficult maintenance. The routing policies used to control network routes are dynamic, posing difficulties to subsequent maintenance. High maintenance skills are required.

It is desirable to have an automatic mechanism for setting up routing policies, which can simplify routing policy configuration. This document describes extensions to BGP for Routing Policy Distribution to resolve these issues.

## **2. Terminology**

The following terminology is used in this document.

- o ACL: Access Control List
- o BGP: Border Gateway Protocol [[RFC4271](#)]
- o FS: Flow Specification
- o NLRI: Network Layer Reachability Information [[RFC4271](#)]
- o PBR: Policy-Based Routing
- o RPD: Routing Policy Distribution
- o VPN: Virtual Private Network



### **3. Problem Statement**

Providers have the requirement to adjust their business traffic routing policies from time to time because of the following:

- o Business development or network failure introduces link congestion and overload.
- o Business changes or network additions produce unused resources such as idle links.
- o Network transmission quality is decreased as the result of delay, loss and they need to adjust traffic to other paths.
- o To control OPEX and CPEX, they may prefer the transit provider with lower price.

#### **3.1. Inbound Traffic Control**

In Figure 1, for the reasons above, the provider P of AS100 may wish the inbound traffic from AS200 to enter AS100 through link L3 instead of the others. Since P doesn't have any administrative control over AS200, there is no way for P to directly modify the route selection criteria inside AS200.



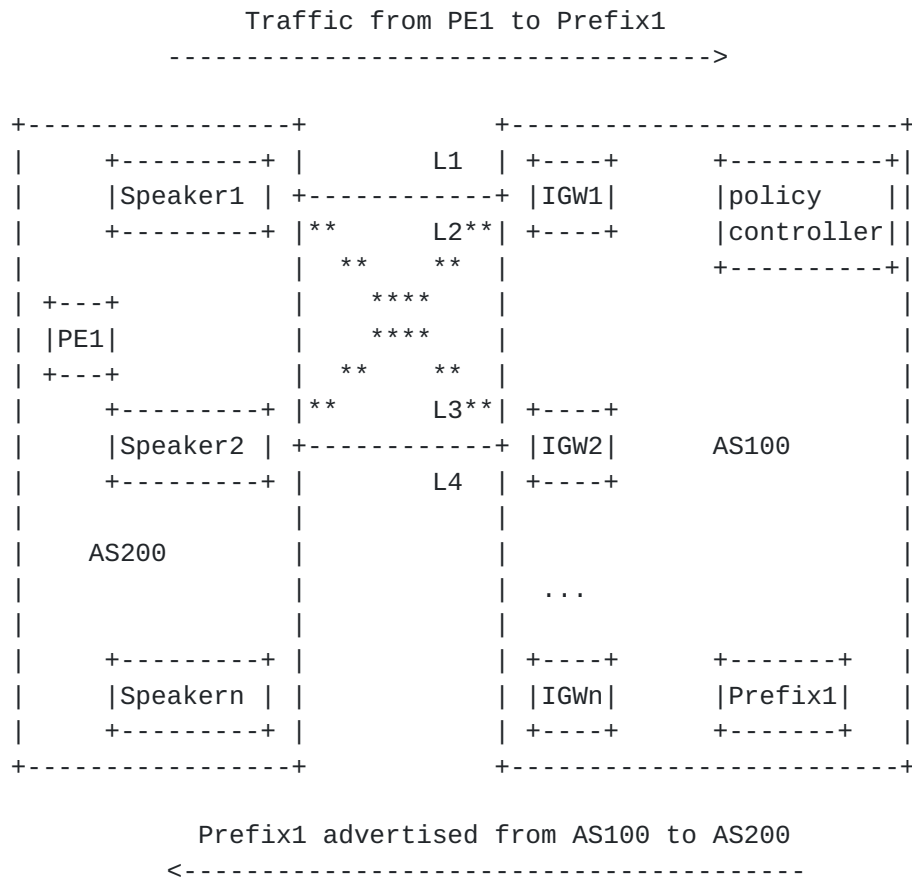


Figure 1: Inbound Traffic Control case

### 3.2. Outbound Traffic Control

In Figure 2, the provider P of AS100 prefers link L3 for the traffic to the destination Prefix2 among multiple exits and links to AS200. This preference can be dynamic and might change frequently because of the reasons above. So, provider P expects an efficient and convenient solution.





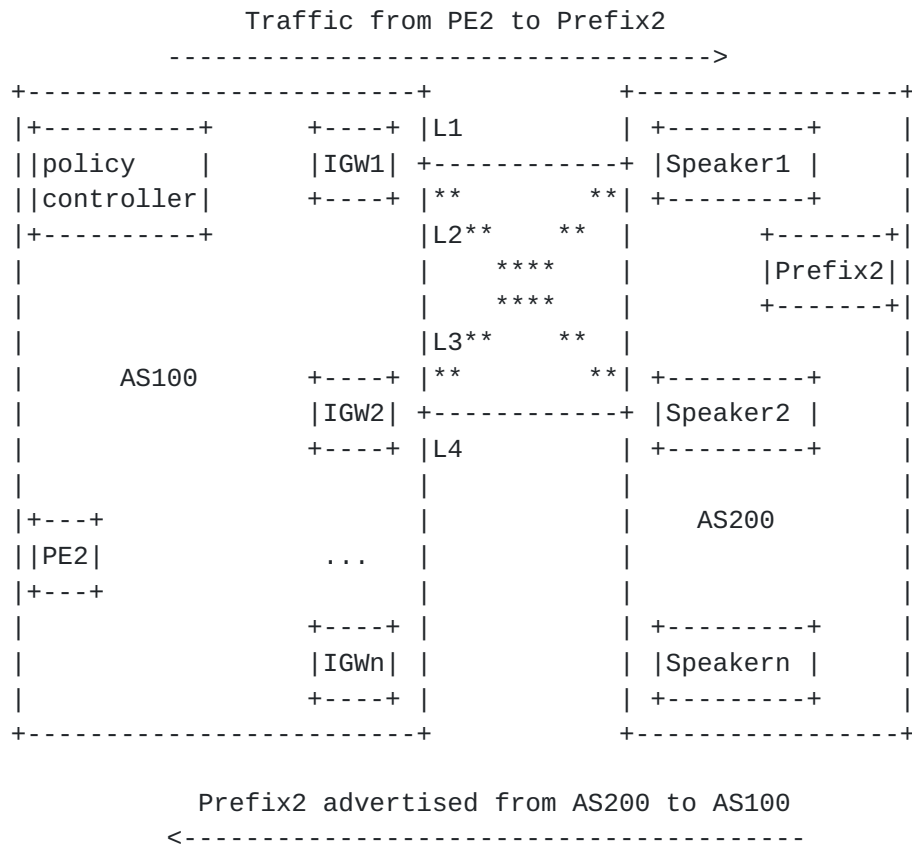


Figure 2: Outbound Traffic Control case

#### 4. Protocol Extensions

This document specifies a solution using a new AFI and SAFI with the BGP Wide Community for encoding a routing policy.

##### 4.1. Using a New AFI and SAFI

A new AFI and SAFI are defined: the Routing Policy AFI whose codepoint 16398 has been assigned by IANA, and SAFI whose codepoint 75 has been assigned by IANA.

The AFI and SAFI pair uses a new NLRI, which is defined as follows:



```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+
|  NLRI Length  |
+---+---+---+---+---+---+
|  Policy Type  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Distinguisher (4 octets)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Peer IP (4/16 octets)                                     ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Where:

NLRI Length: 1 octet represents the length of NLRI. If the Length is anything other than 9 or 21, the NLRI is corrupt and the enclosing UPDATE message MUST be ignored.

Policy Type: 1 octet indicates the type of a policy. 1 is for Export policy. 2 is for Import policy. If the Policy Type is any other value, the NLRI is corrupt and the enclosing UPDATE message MUST be ignored.

Distinguisher: 4 octet value uniquely identifies the content/policy. It is used to sort/order the policies from the lower to higher distinguisher. They are applied in order. The policy with a lower/smaller distinguisher is applied before the policies with higher/larger distinguishers.

Peer IP: 4/16 octet value indicates IPv4/IPv6 peers. Its default value is 0, which indicates that when receiving a BGP UPDATE message with the NLRI, a BGP speaker will apply the policy in the message to all its IPv4/IPv6 peers.

Under RPD AFI/SAFI, the RPD routes are stored and ordered according to their keys. Under IPv4/IPv6 Unicast AFI/SAFI, there are IPv4/IPv6 unicast routes learned and various static policies configured. In addition, there are dynamic RPD policies from the RPD AFI/SAFI when RPD is enabled.

Before advertising an IPv4/IPv6 Unicast AFI/SAFI route, the configured policies are applied to it first, and then the RPD Export policies are applied.

The NLRI containing the Routing Policy is carried in MP\_Reach\_NLRI and MP\_UNREACH\_NLRI path attributes in a BGP UPDATE message, which MUST also contain the BGP mandatory attributes and MAY contain some BGP optional attributes.



When receiving a BGP UPDATE message with routing policy, a BGP speaker processes it as follows:

- o If the peer IP in the NLRI is 0, then apply the routing policy to all the remote peers of this BGP speaker.
- o If the peer IP in the NLRI is non-zero, then the IP address indicates a remote peer of this BGP speaker and the routing policy will be applied to it.

The content of the Routing Policy is encoded in a BGP Wide Community.

#### **4.2. BGP Wide Community and Atoms**

The BGP wide community is defined in [[I-D.ietf-idr-wide-bgp-communities](#)]. It can be used to facilitate the delivery of new network services and be extended easily for distributing different kinds of routing policies.

A wide community Atom is a TLV (or sub-TLV), which may be included in a BGP wide community container (or BGP wide community for short) containing some BGP Wide Community TLVs. Three BGP Wide Community TLVs are defined in [[I-D.ietf-idr-wide-bgp-communities](#)], which are BGP Wide Community Target(s) TLV, Exclude Target(s) TLV, and Parameter(s) TLV. The value of each of these TLVs comprises a series of Atoms, each of which is a TLV (or sub-TLV). A new wide community Atom is defined for BGP Wide Community Target(s) TLV and a few new Atoms are defined for BGP Wide Community Parameter(s) TLV. For your reference, the format of the TLV is illustrated below:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Type      |      Length      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Value (variable)                                     ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Format of Wide Community Atom TLV

##### **4.2.1. RouteAttr TLV/sub-TLV**

A RouteAttr Atom TLV (or RouteAttr TLV/sub-TLV for short) is defined and may be included in a Target TLV. It has the following format.









M-Type = 3: Match prefix within the range of the given masks.

Flags: 4 bits. No flags are currently defined.

IPv4 Address: 4 octets for an IPv4 address.

Mask: 1 octet for the mask length.

GeMask: 1 octet for match range's lower bound, must not be less than Mask or be 0.

LeMask: 1 octet for match range's upper bound, must be greater than Mask or be 0.

For example, tuple <M-Type=0, Flags=0, IPv4 Address = 1.1.0.0, Mask = 22, GeMask = 0, LeMask = 0> represents an exact IP prefix match for 1.1.0.0/22.

<M-Type=1, Flags=0, IPv4 Address = 16.1.0.0, Mask = 24, GeMask = 24, LeMask = 0> represents match IP prefix 1.1.0.0/24 greater-equal 24.

<M-Type=2, Flags=0, IPv4 Address = 17.1.0.0, Mask = 24, GeMask = 0, LeMask = 26> represents match IP prefix 17.1.0.0/24 less-equal 26.

<M-Type=3, Flags=0, IPv4 Address = 18.1.0.0, Mask = 24, GeMask = 24, LeMask = 32> represents match IP prefix 18.1.0.0/24 greater-equal 24 and less-equal 32.

Similarly, an IPv6 Prefix sub-TLV represents match criteria on IPv6 prefixes. Its format is illustrated below:

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type  4      |          Length (N x 20)          |M-Type| Flags |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          IPv6 Address (16 octets)          ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Mask      |   GeMask   |   LeMask   |M-Type| Flags |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~
. . .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          IPv6 Address (16 octets)          ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Mask      |   GeMask   |   LeMask   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Format of IPv6 Prefix sub-TLV







#### 4.2.2. Sub-TLVs of the Parameters TLV

This document introduces 2 community values:

**MATCH AND SET ATTR:** If the IPv4/IPv6 unicast routes to a remote peer match the specific conditions defined in the routing policy extracted from the RPD route, then the attributes of the IPv4/IPv6 unicast routes will be modified when sending to the remote peer per the actions defined in the RPD route.

**MATCH AND NOT ADVERTISE:** If the IPv4/IPv6 unicast routes to a remote peer match the specific conditions defined in the routing policy extracted from the RPD route, then the IPv4/IPv6 unicast routes will not be advertised to the remote peer.

For the Parameter(s) TLV, two action sub-TLVs are defined: MED change sub-TLV and AS-Path change sub-TLV. When the community in the container is MATCH AND SET ATTR, the Parameter(s) TLV can include these sub-TLVs. When the community is MATCH AND NOT ADVERTISE, the Parameter(s) TLV's value is empty.

A MED change sub-TLV indicates an action to change the MED. Its format is illustrated below:

[illegible]

### Format of MED Change sub-TLV

Type: 1 for MED Change.

Length: 5. If Length is any other value, the sub-TLV is corrupt and the enclosing UPDATE MUST be ignored.

OP: 1 octet. Three are defined:

OP = 0: assign the Value to the existing MED.

OP = 1: add the Value to the existing MED. If the sum is greater than the maximum value for MED, assign the maximum value to MED.

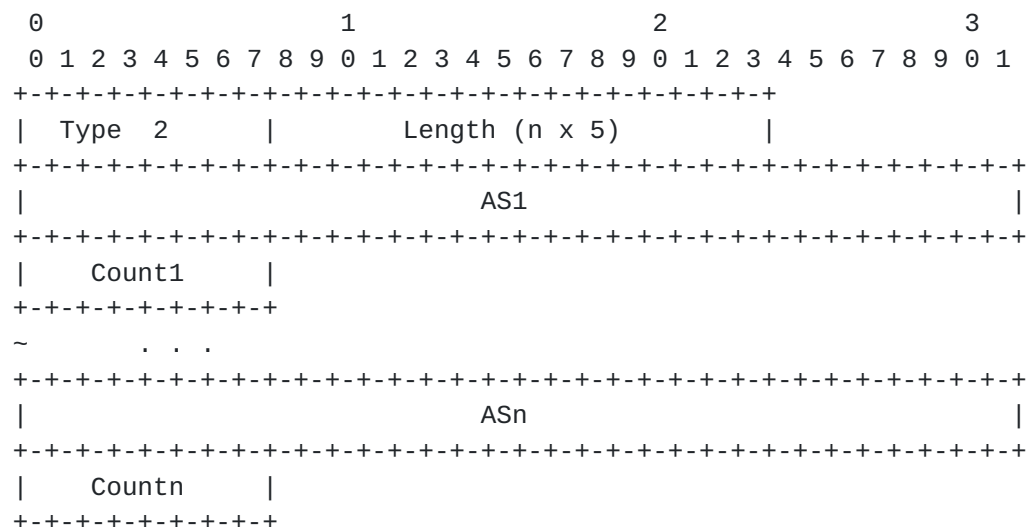


OP = 2: subtract the Value from the existing MED. If the existing MED minus the Value is less than 0, assign 0 to MED.

If OP is any other value, the sub-TLV is ignored.

Value: 4 octets.

An AS-Path change sub-TLV indicates an action to change the AS-Path. Its format is illustrated below:



### Format of AS-Path Change sub-TLV

Type: 2 for AS-Path Change.

Length:  $n \times 5$ . If Length is not a multiple of 5, the sub-TLV is corrupt and the enclosing UPDATE MUST be ignored.

ASi: 4 octet. An AS number.

Counti: 1 octet. ASi repeats Counti times.

The sequence of AS numbers are added to the existing AS Path.

### 4.3. Capability Negotiation

It is necessary to negotiate the capability to support BGP Extensions for Routing Policy Distribution (RPD). The BGP RPD Capability is a new BGP capability [[RFC5492](#)]. The Capability Code for this capability is 72 assigned by the IANA. The Capability Length field of this capability is variable. The Capability Value field consists of one or more of the following tuples:





```

+-----+
| Address Family Identifier (2 octets) |
+-----+
| Subsequent Address Family Identifier (1 octet) |
+-----+
| Send/Receive (1 octet) |
+-----+

```

### BGP RPD Capability

The meaning and use of the fields are as follows:

Address Family Identifier (AFI): This field is the same as the one used in [\[RFC4760\]](#).

Subsequent Address Family Identifier (SAFI): This field is the same as the one used in [\[RFC4760\]](#).

Send/Receive: This field indicates whether the sender is (a) willing to receive Routing Policies from its peer (value 1), (b) would like to send Routing Policies to its peer (value 2), or (c) both (value 3) for the <AFI, SAFI>. If Send/Receive is any other value, that tuple is ignored but any other tuples present are still used.

## 5. Operations

This section presents a typical application scenario and some details about handling a related failure.

### 5.1. Application Scenario

Figure 3 illustrates a typical scenario, where RPD is used by a controller with a Route Reflector (RR) to adjust traffic dynamically.



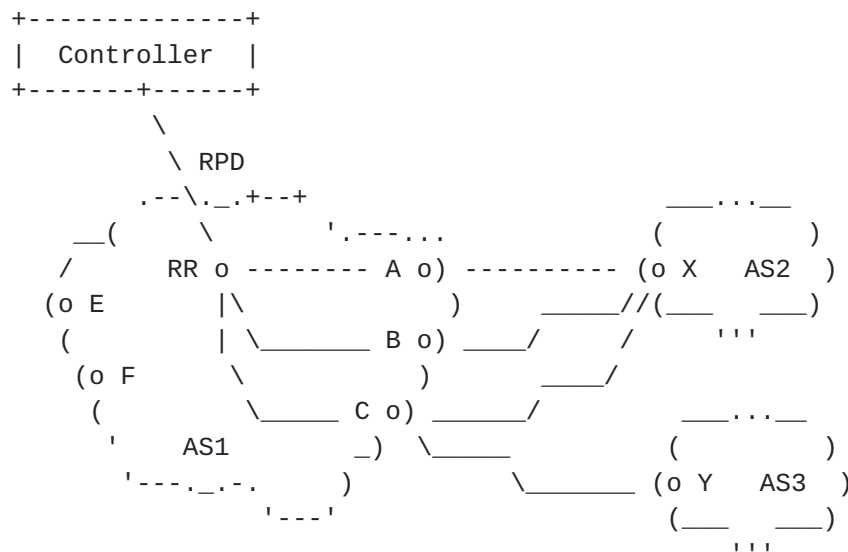


Figure 3: Controller with RR Adjusts Traffic

The controller connects the RR through a BGP session. There is a BGP session between the RR and each of routers A, B and C in AS1, which is shown in the figure. Other sessions in AS1 are not shown in the figure.

There is router X in AS2. There is a BGP session between X and each of routers A, B and C in AS1.

There is router Y in AS3. There is a BGP session between Y and router C in AS1.

The controller sends a RPD route to the RR. After receiving the RPD route from the controller, the RR reflects the RPD route to routers A, B and C. After receiving the RPD route from the RR, routers A, B and C extract the routing policy from the RPD route. If the peer IP in the NLRI of the RPD route is 0, then apply the routing policy to all the remote peers of routers A, B and C. If the peer IP in the NLRI of the RPD route is non-zero, then the IP address indicates a remote peer of routers A, B and C and such routing policy is applied to the specific remote peer. The IPv4/IPv6 unicast routes towards router X in AS2 and router Y in AS3 will be adjusted based on the routing policy sent by the controller via a RPD route.

The controller uses the RT extend community to notify a router whether to receive a RPD policy. For example, if there is not any adjustment on router B, the controller sends RPD routes with the RTs for A and C. B will not receive the routes.



The process of adjusting traffic in a network is a close loop. The loop starts from the controller with some traffic expectations on a set of routes. The controller obtains the information about traffic flows for the related routes. It analyzes the traffic and checks whether the current traffic flows meet the expectations. If the expectations are not met, the controller adjusts the traffic. And then the loop goes to the starter of the loop (The controller obtains the information about traffic ...).

## **5.2. About Failure**

A RPD route is not a configuration. When it is sent to a router, no ack is needed from the router. The existing BGP mechanisms are re-used for delivering a RPD route. After the route is delivered to a router, it will be successful. This is guaranteed by the BGP protocols.

If there is a failure for the router to install the route locally, this failure is a bug of the router. The bug needs to be fixed.

For the errors mentioned in [[RFC7606](#)], they are handled according to [[RFC7606](#)]. These errors are bugs, which need to be resolved.

Regarding to the failure of the controller, some existing mechanisms such BGP GR [[RFC4724](#)] and BGP Long-lived Graceful Restart (LLGR) can be used to let the router keep the routes from the controller for some time.

With support of "Long-lived Graceful Restart Capability" [[I-D.ietf-idr-long-lived-gr](#)], the routes can be retained for a longer time after the controller fails.

In the worst case, the controller fails and the RPD routes for adjusting the traffic are withdrawn. The traffic adjusted/redirected may take its old path. This should be acceptable.

## **6. Contributors**

The following people have substantially contributed to the definition of the BGP-FS RPD and to the editing of this document:

Peng Zhou  
Huawei  
Email: [Jewpon.zhou@huawei.com](mailto:Jewpon.zhou@huawei.com)



## **7. Security Considerations**

Protocol extensions defined in this document do not affect BGP security other than as discussed in the Security Considerations section of [\[RFC5575\]](#).

## **8. Acknowledgements**

The authors would like to thank Acee Lindem, Jeff Haas, Jie Dong, Lucy Yong, Qiandeng Liang, Zhenqiang Li, Robert Raszuk, Donald Eastlake, Ketan Talaulikar, and Jakob Heitz for their comments to this work.

## **9. IANA Considerations**

### **9.1. Existing Assignments**

IANA has assigned a new AFI of value 16398 from the registry "Address Family Numbers" for Routing Policy.

IANA has assigned a new SAFI of value 75 from the registry "Subsequent Address Family Identifiers (SAFI) Parameters" for Routing Policy.

IANA has assigned a new Code Point of value 72 from the registry "Capability Codes" for Routing Policy Distribution.

### **9.2. Routing Policy Type Registry**

IANA is requested to create a new registry called "Routing Policy Type". The allocation policy of this registry is "First Come First Served (FCFS)".

The initial code points are as follows:

Code Point	Description	Reference
0	Reserved	
1	Export Policy	This document
2	Import Policy	This document
3 - 255	Available	





### 9.3. RouteAttr Atom Type

IANA is requested to assign a code-point from the registry "BGP Community Container Atom Types" as follows:

TLV Code Point	Description	Reference
TBD1 (48 suggested)	RouteAttr Atom	This document

### 9.4. Route Attributes Sub-TLV Registry

IANA is requested to create a new registry called "Route Attributes Sub-TLV" under RouteAttr Atom TLV. The allocation policy of this registry is "First Come First Served (FCFS)".

The initial code points are as follows:

Code Point	Description	Reference
0	Reserved	
1	IPv4 Prefix Sub-TLV	This document
2	AS-Path Sub-TLV	This document
3	Community Sub-TLV	This document
4	IPv6 Prefix Sub-TLV	This document
5 - 255	Available	

### 9.5. Attribute Change Sub-TLV Registry

IANA is requested to create a new registry called "Attribute Change Sub-TLV" under Parameter(s) TLV. The allocation policy of this registry is "First Come First Served (FCFS)".

Initial code points are as follows:



Code Point	Description	Reference
0	Reserved	
1	MED Change Sub-TLV	This document
2	AS-Path Change Sub-TLV	This document
3 - 255	Available	

## 10. References

### 10.1. Normative References

- [I-D.ietf-idr-wide-bgp-communities]  
 Raszuk, R., Haas, J., Lange, A., Decraene, B., Amante, S., and P. Jakma, "BGP Community Container Attribute", [draft-ietf-idr-wide-bgp-communities-05](#) (work in progress), July 2018.
- [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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), DOI 10.17487/RFC5492, February 2009, <<https://www.rfc-editor.org/info/rfc5492>>.



- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

## **10.2. Informative References**

- [I-D.ietf-idr-long-lived-gr]  
Uttaro, J., Chen, E., Decraene, B., and J. G. Scudder,  
"Support for Long-lived BGP Graceful Restart", [draft-ietf-idr-long-lived-gr-00](#) (work in progress), September 2019.
- [I-D.ietf-idr-registered-wide-bgp-communities]  
Raszuk, R. and J. Haas, "Registered Wide BGP Community Values", [draft-ietf-idr-registered-wide-bgp-communities-02](#) (work in progress), May 2016.
- [RFC4724] Sangli, S., Chen, E., Fernando, R., Scudder, J., and Y. Rekhter, "Graceful Restart Mechanism for BGP", [RFC 4724](#), DOI 10.17487/RFC4724, January 2007, <<https://www.rfc-editor.org/info/rfc4724>>.
- [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>>.

## **Authors' Addresses**

Zhenbin Li  
Huawei  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: [lizhenbin@huawei.com](mailto:lizhenbin@huawei.com)



Liang Ou  
China Telcom Co., Ltd.  
109 West Zhongshan Ave, Tianhe District  
Guangzhou 510630  
China

Email: ouliang@chinatelecom.cn

Yujia Luo  
China Telcom Co., Ltd.  
109 West Zhongshan Ave, Tianhe District  
Guangzhou 510630  
China

Email: luoyuj@sdu.edu.cn

Sujian Lu  
Tencent  
Tengyun Building, Tower A , No. 397 Tianlin Road  
Shanghai, Xuhui District 200233  
China

Email: jasonlu@tencent.com

Gyan S. Mishra  
Verizon Inc.  
13101 Columbia Pike  
Silver Spring MD 20904  
USA

Phone: 301 502-1347  
Email: gyan.s.mishra@verizon.com

Huaimo Chen  
Futurewei  
Boston, MA  
USA

Email: Huaimo.chen@futurewei.com





Shunwan Zhuang  
Huawei  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: zhuangshunwan@huawei.com

Haibo Wang  
Huawei  
Huawei Bld., No.156 Beiqing Rd.  
Beijing 100095  
China

Email: rainsword.wang@huawei.com

