

Source/Destination Routing Using BGP-4
draft-xu-src-dst-bgp-00

Abstract

This document describes the changes necessary for BGP-4 to route traffic from a specified prefix to a specified prefix.

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 <http://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 September 22, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Theory of Routing	3
3.	Extended NLRI Encodings	3
4.	Dealing with Ambiguity	4
5.	Src-Dst Capability	4
6.	Compatibility Considerations	5
7.	Deployment Issues	5
8.	Security Considerations	5
9.	IANA Considerations	5
10.	References	5
10.1.	Normative References	5
10.2.	Informative References	6
	Authors' Addresses	6

[1.](#) Introduction

This specification builds on BGP-4 [[RFC4271](#)]. It defines the extended NLRI encodings for an appended source prefix, to define routes from a source prefix to a destination prefix.

Traditionally, routing protocols make routing decisions solely based on destination IP addresses, packets towards the same destination will be delivered to the same next hop no matter where they come from. However, considering policy-based routing, traffic engineering and security, source information is also important for making routing decisions.

In this document, we extend the NLRI field to support source prefix. This implies not simply routing "to a destination", but routing "to that destination AND from a specified source". Traffic within the network could be source/destination routed as well, or could be implicitly or explicitly routed from "any prefix", `::/0`.

1.1. 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](#)].

2. Theory of Routing

The primary function of BGP is to exchange network reachability information, compute the routes towards destination prefixes, and select the best routes according the pre-defined selection rules. BGP-4 can support only those policies which conform to the destination-based forwarding paradigm.

In this context, the route is qualified by a source prefix. Intrinsically, in traditional routing model, the object being routed to is a destination prefix; in the new routing model, the object being routed might be a destination prefix given that the packet sports a certain source prefix.

Routes that lack a source prefix match any source prefix (i.e., ::/0), by definition.

3. Extended NLRI Encodings

In order to carry the source prefix information in an UPDATE message, the existing NLRI encodings are extended by prepending the source prefix.

The NLRI encodings specified in [[RFC4271](#)] and [[RFC4760](#)] are extended as following:

```
+-----+
| Type (4 octets)           |
+-----+
| Length (1 octet)         |
+-----+
| Prefix (variable)        |
+-----+
```

Extended NLRI Encodings based on [RFC4271](#) and [RFC4760](#)

and the NLRI encoding specified in [[RFC3107](#)] is extended as the following:


```
+-----+
| Type (4 octets)          |
+-----+
| Length (1 octet)        |
+-----+
| Label (3 octets)        |
+-----+
| ...                     |
+-----+
| Prefix (variable)       |
+-----+
```

Extended NLRI encodings based on [RFC3107](#)

Type: Assigned by IANA.

Length: Indicates the length in bits of the IP address prefix.

Label: Carrying label information as defined in [[RFC3107](#)]

Prefix: The Prefix field contains an IP address prefix, followed by enough trailing bits to make the end of the field fall on an octet boundary.

4. Dealing with Ambiguity

Ambiguity could happen when there are two routes: A and B, where source prefix of A is more specific than source prefix of B, and destination prefix of B is more specific than destination prefix of A.

In this context, the matching rule follows that in [[I-D.baker-ipv6-ospf-dst-src-routing](#)], the FIB lookup MUST yield the route with the longest matching destination prefix that also matches the source prefix constraint. In the event of a tie on the destination prefix, it MUST also match the longest matching source prefix among those options.

5. Src-Dst Capability

The capability to carry both source and destination prefixes in BGP update messages (src-dst capability) is a new BGP capability [[RFC5492](#)]. The Capability Code for this capability is specified in the IANA. The Capability Length field of this capability is zero.

6. Compatibility Considerations

To be compatible with [[I-D.ietf-idr-add-paths](#)], the Type field (defined in Section [Section 3](#)) should be carefully defined by IANA.

7. Deployment Issues

Router without src-dst capability should discard the BGP messages with extended NRLI, and it falls back to traditional destination-based routing when this happens.

8. Security Considerations

While source/destination routing could be used as part of a security solution, it could be considered similar to an access list that is managed by and scales with routing.

9. IANA Considerations

The Type field in Section [Section 3](#), and the new capability code should be defined by IANA.

10. References

10.1. Normative References

- [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, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), DOI 10.17487/RFC5492, February 2009, <<http://www.rfc-editor.org/info/rfc5492>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", [RFC 3107](#), DOI 10.17487/RFC3107, May 2001, <<http://www.rfc-editor.org/info/rfc3107>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<http://www.rfc-editor.org/info/rfc4760>>.

10.2. Informative References

[I-D.ietf-idr-add-paths]

Walton, D., Retana, A., Chen, E., and J. Scudder,
"Advertisement of Multiple Paths in BGP", [draft-ietf-idr-add-paths-10](#) (work in progress), October 2014.

[I-D.baker-ipv6-ospf-dst-src-routing]

Baker, F., "IPv6 Source/Destination Routing using OSPFv3",
[draft-baker-ipv6-ospf-dst-src-routing-03](#) (work in
progress), August 2013.

Authors' Addresses

Mingwei Xu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-1572
Email: xumw@tsinghua.edu.cn

Shu Yang
Graduate School at Shenzhen, Tsinghua University
Division of Information Science and Technology
Shenzhen 518055
P.R. China

Phone: +86-755-2603-6059
Email: yang.shu@sz.tsinghua.edu.cn

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

