

Global Routing Operations
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2017

C. Petrie
RIPE NCC
T. King
DE-CIX Management GmbH
July 8, 2016

Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format
with BGP Additional Paths Extensions
draft-ietf-grow-mrt-add-paths-01

Abstract

This document updates the Multi-threaded Routing Toolkit (MRT) export format for Border Gateway Protocol (BGP) routing information by extending it to support the Advertisement of Multiple Paths in BGP extensions.

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 January 9, 2017.

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.

Table of Contents

1.	Introduction	2
2.	Requirements Language	2
3.	Rationale	3
4.	MRT Subtypes for Types BGP4MP/BGP4MP_ET	3
5.	MRT Subtypes for Type TABLE_DUMP_V2	3
5.1.	AFI/SAFI specific RIB Subtypes	4
5.2.	RIB_GENERIC_ADDPATH Subtype	4
6.	IANA Considerations	5
6.1.	BGP4MP/BGP4MP_ET Subtype codes:	5
6.2.	TABLE_DUMP_V2 Subtype codes:	5
7.	Security Considerations	5
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	6
8.3.	URIs	6
	Authors' Addresses	7

[1.](#) Introduction

The MRT record format [[RFC6396](#)] was developed to provide researchers and engineers a means to encapsulate, export, and archive routing protocol transactions and routing information base snapshots.

The Advertisement of Multiple Paths in BGP [[I-D.ietf-idr-add-paths](#)] defines a BGP extension to allow the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previous ones.

This document contains an optional extension to the MRT format [[RFC6396](#)] and introduces additional definitions of MRT subtype fields to permit representation of multiple path advertisements [[I-D.ietf-idr-add-paths](#)].

[2.](#) 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 [RFC 2119](#) [[RFC2119](#)].

3. Rationale

MRT parsers are usually stateless. In order to parse BGP messages which contain data structures that depend on the capabilities negotiated during the BGP session setup, the so-called MRT subtypes are utilized. The Advertisement of Multiple Path [[I-D.ietf-idr-add-paths](#)] extension for BGP alters the encoding of the BGP NLRI format for withdraws and announcements. Therefore new BGP4MP/BGP4MP_ET subtypes as defined in [[RFC6396](#)] are required to signal to a MRT parser how to parse the NLRI.

In [section 4.3 \[RFC6396\]](#) of the MRT specification RIB subtypes are specified. Prefix length and prefix fields are encoded in the same manner as the BGP NLRI encoding. In order to support path identifier information as defined in [[I-D.ietf-idr-add-paths](#)] new subtypes need to be added.

The following two sections define the required subtypes.

4. MRT Subtypes for Types BGP4MP/BGP4MP_ET

This document defines the following new Subtypes:

- o BGP4MP_MESSAGE_ADDPATH
- o BGP4MP_MESSAGE_AS4_ADDPATH
- o BGP4MP_MESSAGE_LOCAL_ADDPATH
- o BGP4MP_MESSAGE_AS4_LOCAL_ADDPATH

The fields of these message types are identical to the equivalent non-additional-path versions specified in [section 4.4 \[RFC6396\]](#). These enhancements continues to encapsulate the entire BGP message in the BGP message field.

5. MRT Subtypes for Type TABLE_DUMP_V2

This document defines the following new Subtypes:

- o RIB_IPV4_UNICAST_ADDPATH
- o RIB_IPV4_MULTICAST_ADDPATH
- o RIB_IPV6_UNICAST_ADDPATH
- o RIB_IPV6_MULTICAST_ADDPATH

- o RIB_GENERIC_ADDPATH

The fields of these message types are identical to the equivalent non-additional-path versions specified in [section 4.3 \[RFC6396\]](#). However, for the specific case of the 4 AFI/SAFI specific RIB subtypes, the existing RIB Entries field is re-defined as detailed in the sections below.

5.1. AFI/SAFI specific RIB Subtypes

In order to preserve the record compaction achieved by using the most common subtypes, and allowing multiple RIB Entries to be stored in a single TABLE_DUMP_V2 record, the existing RIB Entries field is redefined for use within the new AFI/SAFI specific RIB Subtypes defined by this document as follows:

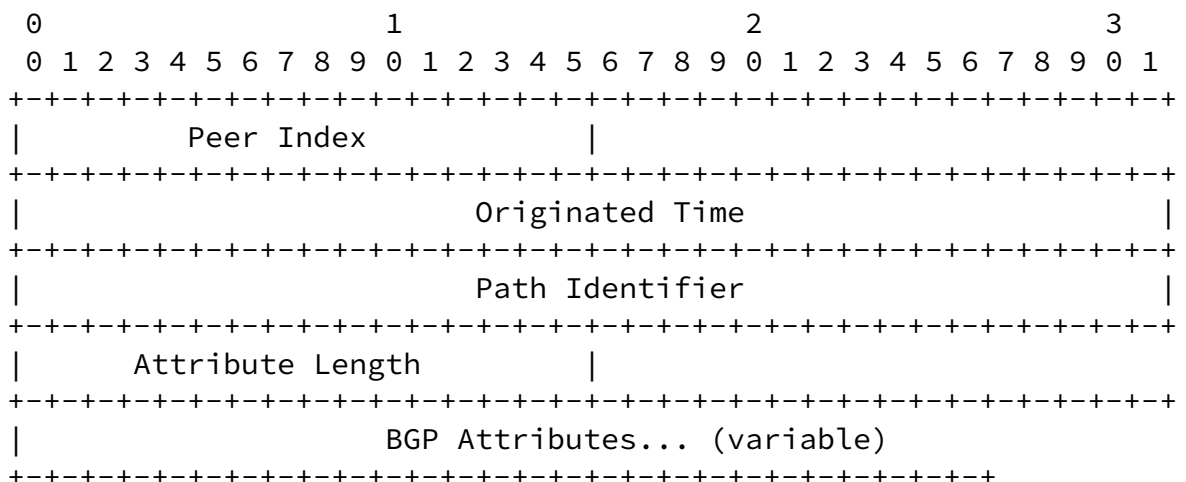


Figure 1: RIB Entries for AFI/SAFI-specific RIB Subtypes with additional-paths support

This adds a field to the RIB Entries record, to store the path identifier, when used with the RIB_IPV4_UNICAST_ADDPATH, RIB_IPV4_MULTICAST_ADDPATH, RIB_IPV6_UNICAST_ADDPATH and RIB_IPV6_MULTICAST_ADDPATH subtypes.

[5.2.](#) RIB_GENERIC_ADDPATH Subtype

The fields of this subtype are identical to the equivalent non-additional-path versions specified in [section 4.3.3 \[RFC6396\]](#). These fields continue to encapsulate the raw and additional-path enabled AFI/SAFI/NLRI in the record, and the raw attributes in the RIB Entries.

For clarity, the RIB Entries in this subtype are not redefined.

[6.](#) IANA Considerations

This document requests that IANA assign the following subtype codes to the MRT name space [\[1\]](#):

[6.1.](#) BGP4MP/BGP4MP_ET Subtype codes:

BGP4MP_MESSAGE_ADDPATH = 8 ([Section 4](#))

BGP4MP_MESSAGE_AS4_ADDPATH = 9 ([Section 4](#))

BGP4MP_MESSAGE_LOCAL_ADDPATH = 10 ([Section 4](#))

BGP4MP_MESSAGE_AS4_LOCAL_ADDPATH = 11 ([Section 4](#))

The values provided above are suggested as they are used in implementations.

[6.2.](#) TABLE_DUMP_V2 Subtype codes:

RIB_IPV4_UNICAST_ADDPATH = 8 ([Section 5.1](#))

RIB_IPV4_MULTICAST_ADDPATH = 9 ([Section 5.1](#))

RIB_IPV6_UNICAST_ADDPATH = 10 ([Section 5.1](#))

RIB_IPV6_MULTICAST_ADDPATH = 11 ([Section 5.1](#))

RIB_GENERIC_ADDPATH = 12 ([Section 5.2](#))

The values provided above are suggested as they are used in implementations.

[7.](#) Security Considerations

It is not believed that this document adds any additional security considerations.

However, the security considerations of [[RFC6396](#)] are equally applicable to this document, and this document permits the export of more detailed routing data.

An organization that uses the MRT format to store their BGP routing information should be aware that supporting these extensions permits more detailed network path information to be stored, and should consider the implications of this within their environment.

An organization that peers with public BGP collectors, and enables the additional-paths capability on a peering session, should be aware that it is exporting not only its best paths, but potentially other paths within its networks. The BGP peer should consider any and all implications of exposing this additional data.

[8.](#) References

[8.1.](#) Normative 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-15](#) (work in progress), May 2016.

- [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>>.
- [RFC6396] Blunk, L., Karir, M., and C. Labovitz, "Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format", [RFC 6396](#), DOI 10.17487/RFC6396, October 2011, <<http://www.rfc-editor.org/info/rfc6396>>.

8.2. Informative References

- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), DOI 10.17487/RFC2629, June 1999, <<http://www.rfc-editor.org/info/rfc2629>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<http://www.rfc-editor.org/info/rfc3552>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

8.3. URIs

- [1] <https://www.iana.org/assignments/mrt/mrt.xhtml>

Authors' Addresses

Colin Petrie
RIPE NCC
Singel 258
Amsterdam 1016 AB
NL

Email: cpetrie@ripe.net

Thomas King
DE-CIX Management GmbH
Lichtstrasse 43i
Cologne 50825
Germany

Email: thomas.king@de-cix.net