

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 13, 2016

C. Petrie  
RIPE NCC  
October 11, 2015

**Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format  
with BGP Additional Paths Extensions  
draft-petrie-grow-mrt-add-paths-00**

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 April 13, 2016.

Copyright Notice

Copyright (c) 2015 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Requirements Language . . . . .	<a href="#">2</a>
<a href="#">3.</a>	Rationale . . . . .	<a href="#">2</a>
<a href="#">4.</a>	MRT Subtypes for Type BGP4MP . . . . .	<a href="#">3</a>
<a href="#">5.</a>	MRT Subtypes for Type TABLE_DUMP_V2 . . . . .	<a href="#">4</a>
<a href="#">5.1.</a>	AFI/SAFI specific RIB Subtypes . . . . .	<a href="#">4</a>
<a href="#">5.2.</a>	RIB_GENERIC_AP Subtype . . . . .	<a href="#">5</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">5</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">8.</a>	References . . . . .	<a href="#">6</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">6</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">6</a>
	Author's Address . . . . .	<a href="#">7</a>

## [1.](#) Introduction

Researchers and engineers often wish to analyze network behavior by studying routing protocol transactions and routing information base snapshots. To this end, the MRT record format [[RFC6396](#)] was developed to encapsulate, export, and archive this information in a standardized data representation.

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. The essence of the extension is that each path is identified by a path identifier in addition to the address prefix

This memo documents an optional extension to the MRT format [RFC6396](#) [[RFC6396](#)] and introduces additional definitions of MRT Subtype fields to permit representation of Multiple Path advertisements

## [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

When a BGP message requires information about the capabilities negotiated during the setup of the BGP session for a parser to interpret the message, this information is carried by the MRT subtypes.



The MRT specification defines the following BGP4MP subtypes:

- o BGP4MP\_MESSAGE
- o BGP4MP\_MESSAGE\_AS4
- o BGP4MP\_MESSAGE\_LOCAL
- o BGP4MP\_MESSAGE\_AS4\_LOCAL

These indicate to a parser whether the AS\_PATH and AGGREGATOR attributes should be interpreted according to the rules in [RFC6793](#) [[RFC6793](#)]

Additional Paths in BGP [[I-D.ietf-idr-add-paths](#)] alters the encoding of the BGP NLRI format for withdraws and announcements. Therefore new BGP4MP subtypes are required to signal to a parser how to parse the NLRI.

The MRT specification defines the following TABLE\_DUMP\_V2 subtypes:

- o RIB\_IPV4\_UNICAST
- o RIB\_IPV4\_MULTICAST
- o RIB\_IPV6\_UNICAST
- o RIB\_IPV6\_MULTICAST
- o RIB\_GENERIC

The existing TABLE\_DUMP\_V2 AFI/SAFI-Specific RIB Subtypes specify that the Prefix Length and Prefix fields are encoded in the same manner as the BGP NLRI encoding. These also require new subtypes to retain the path identifier information in Additional Paths.

The TABLE\_DUMP\_V2 RIB\_GENERIC subtype contains a single raw NLRI entry, the encoding of which is defined by the AFI and SAFI. Additional Paths alter the NLRI encoding. Therefore a new subtype is required to indicate the change in NLRI format.

#### **4. MRT Subtypes for Type BGP4MP**

This document defines the following new Subtypes:

- o BGP4MP\_MESSAGE\_AP
- o BGP4MP\_MESSAGE\_AS4\_AP



- o BGP4MP\_MESSAGE\_LOCAL\_AP
- o BGP4MP\_MESSAGE\_AS4\_LOCAL\_AP

The fields of these message types are identical to the equivalent non-additional-path versions specified in [RFC 6396](#) [[RFC6396](#)], and 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\_AP
- o RIB\_IPV4\_MULTICAST\_AP
- o RIB\_IPV6\_UNICAST\_AP
- o RIB\_IPV6\_MULTICAST\_AP
- o RIB\_GENERIC\_AP

The fields of these message types are identical to the equivalent non-additional-path versions specified in [RFC 6396](#) [[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:



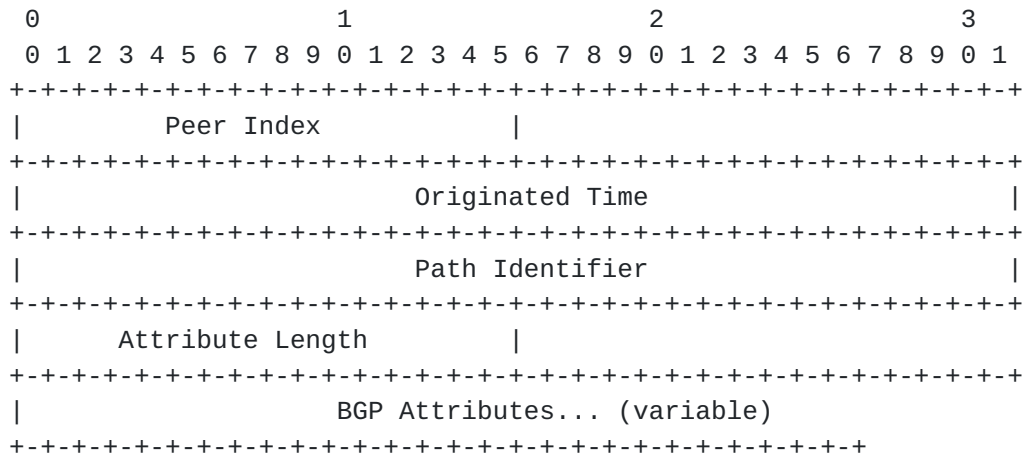


Figure 1: RIB Entries for AFI/SAFI-specific RIB Subtypes with additional-paths support

This adds a field to the RIB Entries record, to store the Path Identifier, when used with the RIB\_IPV4\_UNICAST\_AP, RIB\_IPV4\_MULTICAST\_AP, RIB\_IPV6\_UNICAST\_AP and RIB\_IPV6\_MULTICAST\_AP Subtypes

5.2. RIB\_GENERIC\_AP Subtype

The fields of this message types is identical to the equivalent non-additional-path versions specified in RFC 6396 [RFC6396], and continues to encapsulate the raw AFI/SAFI/NLRI in the record, and the raw attributes in the RIB Entries.

The RIB entries are unchanged, and should be interpreted according to RFC 6396 [RFC6396]

6. IANA Considerations

This document requests that IANA add the appropriate Type Codes and Subtype Codes (to be assigned). This is currently a placeholder.

7. Security Considerations

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

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

An organisation which 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 network that peers with public BGP collectors, and enable the additional-paths capability on a the peering session, should be aware that they are exporting not only their best paths, but potentially other paths within their network. The BGP peer should consider any 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-11](#) (work in progress), October 2015.
- [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>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", [RFC 6793](#), DOI 10.17487/RFC6793, December 2012,  
<<http://www.rfc-editor.org/info/rfc6793>>.

### 8.2. Informative References

- [I-D.narten-iana-considerations-rfc2434bis]  
Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [draft-narten-iana-considerations-rfc2434bis-09](#) (work in progress), March 2008.
- [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>>.

Author's Address

Colin Petrie  
RIPE NCC  
Amsterdam  
NL

Email: [cpetrie@ripe.net](mailto:cpetrie@ripe.net)

