

Network Working Group	L. Blunk	
Internet-Draft	M. Karir	
Intended status: Standards Track	Merit Network	
Expires: May 19, 2008	C. Labovitz	
	Arbor Networks	
	November 16, 2007	

[TOC](#)

MRT routing information export format draft-ietf-grow-mrt-05.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 19, 2008.

Abstract

This document describes the MRT format for routing information export. This format was developed in concert with the Multi-threaded Routing Toolkit (MRT) from whence the format takes its name. The format can be used to export routing protocol messages, state changes, and routing information base contents.

Table of Contents

- [1.](#) Requirements notation
- [2.](#) Introduction

- [3.](#) Basic MRT Format
 - [4.](#) MRT Control Types
 - [4.1.](#) NULL Type
 - [4.2.](#) START Type
 - [4.3.](#) DIE Type
 - [4.4.](#) I_AM_DEAD Type
 - [4.5.](#) PEER_DOWN Type
 - [5.](#) MRT Routing Information Types
 - [5.1.](#) BGP Type
 - [5.1.1.](#) BGP_NULL Subtype
 - [5.1.2.](#) BGP_UPDATE Subtype
 - [5.1.3.](#) BGP_PREF_UPDATE Subtype
 - [5.1.4.](#) BGP_STATE_CHANGE Subtype
 - [5.1.5.](#) BGP_SYNC Subtype
 - [5.1.6.](#) BGP_OPEN Subtype
 - [5.1.7.](#) BGP_NOTIFY Subtype
 - [5.1.8.](#) BGP_KEEPALIVE Subtype
 - [5.2.](#) RIP Type
 - [5.3.](#) IDRP Type
 - [5.4.](#) RIPNG Type
 - [5.5.](#) BGP4PLUS and BGP4PLUS_01 Types
 - [5.6.](#) OSPF Type
 - [5.7.](#) TABLE_DUMP Type
 - [5.8.](#) TABLE_DUMP_V2 Type
 - [5.9.](#) BGP4MP Type
 - [5.9.1.](#) BGP4MP_STATE_CHANGE Subtype
 - [5.9.2.](#) BGP4MP_MESSAGE Subtype
 - [5.9.3.](#) BGP4MP_ENTRY Subtype
 - [5.9.4.](#) BGP4MP_SNAPSHOT Subtype
 - [5.9.5.](#) BGP4MP_STATE_CHANGE_AS4 Subtype
 - [5.9.6.](#) BGP4MP_MESSAGE_AS4 Subtype
 - [5.10.](#) BGP4MP_ET
 - [5.11.](#) ISIS Type
 - [5.12.](#) ISIS_ET Type
 - [5.13.](#) OSPFv3 Type
 - [5.14.](#) OSPFv3_ET Type
 - [6.](#) IANA Considerations
 - [6.1.](#) Type Codes
 - [6.2.](#) Subtype Codes
 - [7.](#) Security Considerations
 - [8.](#) References
 - [8.1.](#) Normative References
 - [8.2.](#) Informative References
 - [§](#) Authors' Addresses
 - [§](#) Intellectual Property and Copyright Statements
-

1. Requirements notation

[TOC](#)

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 [\[RFC2119\] \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

2. Introduction

[TOC](#)

Researchers and engineers often wish to analyze network behavior by studying routing protocol transactions and routing information base snapshots. To this end, the MRT format was developed to encapsulate, export, and archive this information in a standardized data representation. The BGP routing protocol, in particular, has been the subject of extensive study and analysis which has been significantly aided by the availability of the MRT format. The MRT format was initially defined in the [MRT Programmer's Guide \(Labovitz, C., "MRT Programmer's Guide," November 1999.\)](#) [MRT PROG GUIDE].

This memo serves to document the MRT format as currently implemented in publicly available software. The format has been extended since it's original introduction in the MRT toolset and these extensions are also included in this memo. Further extensions may be introduced at a later date through additional definitions of the MRT Type field and Subtype fields.

3. Basic MRT Format

[TOC](#)

All MRT format messages have a common header which includes a timestamp, Type, Subtype, and length field. The header is followed by a message field. The MRT common header is illustrated below.

4. MRT Control Types

The MRT format defines five Control Type messages. These messages are OPTIONAL and MAY be used to communicate the state of the MRT message source and it's peering sessions. The message field MAY contain an OPTIONAL message string for diagnostic purposes. The message string encoding MUST follow the UTF-8 transformation format. The Subtype field is unused for these Types and SHOULD be set to 0.

The MRT Control Types are defined below:

0	NULL
1	START
2	DIE
3	I_AM_DEAD
4	PEER_DOWN

4.1. NULL Type

[TOC](#)

The NULL Type message causes no operation, A sender may wish to send these for synchronization or keep-alive purposes.

4.2. START Type

[TOC](#)

The START Type indicates a sender is about to begin sending MRT messages

4.3. DIE Type

[TOC](#)

A DIE Type signals that the receiver should shut down.

4.4. I_AM_DEAD Type

[TOC](#)

A I_AM_DEAD indicates that the sender is shutting down.

[TOC](#)

4.5. PEER_DOWN Type

A PEER_DOWN indicates when one of the sender's peers is down. In practice, a sender will likely have multiple peers. The sender SHOULD use the Message field to convey the IP address of the peer represented in UTF-8.

5. MRT Routing Information Types

[TOC](#)

The following Types are currently defined for the MRT format. Types 5-12 were defined in the initial MRT Toolkit package. The BGP4MP Type, number 16, was initially defined in the Zebra routing software package. The BGP4MP_ET, ISIS, and ISIS_ET Types were initially defined in the Sprint Labs Python Routing Toolkit (PyRT).

5	BGP	*DEPRECATED*
6	RIP	
7	IDRP	*DEPRECATED*
8	RIPNG	
9	BGP4PLUS	*DEPRECATED*
10	BGP4PLUS_01	*DEPRECATED*
11	OSPF	
12	TABLE_DUMP	
13	TABLE_DUMP_V2	
16	BGP4MP	
17	BGP4MP_ET	
32	ISIS	
33	ISIS_ET	
48	OSPFv3	
49	OSPFv3_ET	

5.1. BGP Type

[TOC](#)

The BGP Type indicates the Message field contains BGP routing information. The BGP routing protocol is defined in [RFC 4271 \(Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 \(BGP-4\)," January 2006.\)](#) [RFC4271]. The information in the message is dependent on the Subtype value. The BGP Type and all associated Subtypes are considered to be DEPRECATED by the BGP4MP Type. The following BGP Subtypes are defined for the MRT BGP Type.

5.1.4. BGP_STATE_CHANGE Subtype

[TOC](#)

The BGP_STATE_CHANGE Subtype is used to record changes in the BGP finite state machine. These FSM states and their numeric encodings are defined in [RFC 4271 \(Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 \(BGP-4\)," January 2006.\)](#) [RFC4271], Appendix 1. Both the old state value and the new state value are encoded as 2-octet numbers. The format of the MRT Message field is 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Source AS number          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Source IP address          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Old State          |          New State          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

5.1.5. BGP_SYNC Subtype

[TOC](#)

The BGP_SYNC Subtype is used to indicate a File Name where BGP Table Dump messages should be recorded. The View # corresponds to the View # provided in the TABLE_DUMP Type messages. The following format applies to this Subtype:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          View #          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          File Name... (variable)
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The File Name is terminated with a NULL (0) character.

5.1.6. BGP_OPEN Subtype

[TOC](#)

The BGP_OPEN Subtype is used to encode BGP OPEN messages. The format of the MRT Message field for this Subtype is the same as the BGP_UPDATE, however, the last field contains the contents of the BGP OPEN message.

5.1.7. BGP_NOTIFY Subtype

[TOC](#)

The BGP_NOTIFY Subtype is used to encode BGP NOTIFICATION messages. The format of the MRT Message field for this Subtype is the same as the BGP_UPDATE, however, the last field contains the contents of the BGP NOTIFICATION message.

5.1.8. BGP_KEEPAIVE Subtype

[TOC](#)

The BGP_KEEPAIVE Subtype is used to encode BGP KEEPAIVE messages. The format of the MRT Message field for this Subtype is the same as the BGP_UPDATE, however, the last field contains no information.

5.2. RIP Type

[TOC](#)

The RIP Type is used to export RIP protocol packets as defined in [RFC 1058 \(Hedrick, C., "Routing Information Protocol," June 1988.\)](#) [RFC1058]. The Subtype field is currently reserved for this Type and SHOULD be set to 0.

The format of the MRT Message field for the RIP Type is 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Source IP address                       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Destination IP address                   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               RIP Message Contents (variable)         |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

5.3. IDRP Type

[TOC](#)

The IDRP Type is used to export Inter-Domain-Routing Protocol (IDRP) protocol information as defined in the ISO/IEC 10747 standard. The Subtype field is unused. This Type is deprecated due to lack of deployment of IDRP.

5.4. RIPNG Type

[TOC](#)

The RIPNG Type is used to export RIPNG protocol packets as defined in [RFC 2080 \(Malkin, G. and R. Minnear, "RIPng for IPv6," January 1997.\)](#) [RFC2080]. The Subtype field is currently reserved for this Type and SHOULD be set to 0.

The format of the MRT Message field for the RIPNG Type is 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Source IP address                       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Destination IP address                   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               RIPNG Message Contents (variable)       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

5.5. BGP4PLUS and BGP4PLUS_01 Types

[TOC](#)

The BGP4PLUS and BGP4PLUS_01 Types were defined to support IPv6 BGP routing information. The BGP4PLUS Type was specified based on the initial Internet Draft for Multiprotocol Extensions to BGP-4. The BGP4PLUS_01 Type was specified to correspond to the -01 revision of this Internet Draft. The two Types share the same definitions in terms of their MRT format specifications.

The Subtype field definitions are shared with the BGP Type, however, the address fields in the BGP_UPDATE, BGP_OPEN, BGP_NOTIFY, BGP_KEEPALIVE, and BGP_STATE_CHANGE Subtype messages are extended to 16 octets for IPv6 addresses. As with the BGP Type, the BGP4PLUS and BGP4PLUS_01 Types are deprecated as they superseded by the BGP4MP Type.

5.6. OSPF Type

[TOC](#)

This Type supports the OSPF Protocol as defined in [RFC 2328 \(Moy, J., "OSPF Version 2," April 1998.\)](#) [RFC2328]. The Subtype field may contain two possible values:

The format of the TABLE_DUMP Type is illustrated below.

[illegible]

The View field is normally 0 and is intended for cases where an implementation may have multiple RIB views (such as a route server). The Sequence field is a simple incremental counter for each RIB entry. A typical RIB dump will exceed the 16-bit bounds of this counter and implementation should simply wrap back to zero and continue incrementing the counter in such cases.

The Prefix field contains the IP address of a particular routing RIB entry. The size of this field is dependent on the value of the Subtype for this message. For AFI_IPv4, this field is 4 octets, for AFI_IPv6, it is 16 octets in length. The Prefix Length field indicates the length in bits of the prefix mask for the preceding Prefix field.

The Status octet is not used in the TABLE_DUMP Type and SHOULD be set to 1.

The Originated Time contains the 4-octet time at which this prefix was heard. The value represents the time in seconds since 1 January 1970 00:00:00 UTC.

The Peer IP field is the IP address of the peer which provided the update for this RIB entry. As with the Prefix field, the size of this field is dependent on the Subtype. AFI_IPv4 indicates a 4 octet field and an IPv4 address, while a Subtype of AFI_IPv6 requires a 16 octet field and an IPv6 address. The Peer AS field contains the AS number of the peer.

Attribute length is the length of Attribute field and is 2-octets. The Attribute field contains the attribute information for the RIB entry.

5.8. TABLE_DUMP_V2 Type

The TABLE_DUMP_V2 type updates the TABLE_DUMP type to include 32-bit ASN support and full support for BGP Multiprotocol extensions. It also improves upon the space efficiency of the TABLE_DUMP type by employing an index table for peers and permitting a single MRT record per NLRI entry. The following subtypes are used with the TABLE_DUMP_V2 type.

- 1 PEER_INDEX_TABLE
- 2 RIB_IPV4_UNICAST
- 3 RIB_IPV4_MULTICAST
- 4 RIB_IPV6_UNICAST
- 5 RIB_IPV6_MULTICAST
- 6 RIB_GENERIC

An initial PEER_INDEX_TABLE MRT record provides the BGP ID of the collector, an optional view name, and a list of indexed peers. Following the PEER_INDEX_TABLE MRT record, a series of MRT records are used to encode RIB table entries. The header of the PEER_INDEX_TABLE Subtype is shown below. The View Name is optional and, if not present, the View Name Length MUST be set to 0. The View Name encoding MUST follow the UTF-8 transformation format.

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Collector BGP ID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      View Name Length      |      View Name (variable)      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Peer Count      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The format of the peer entries is shown below. The PEER_INDEX_TABLE record contains Peer Count peer entries.

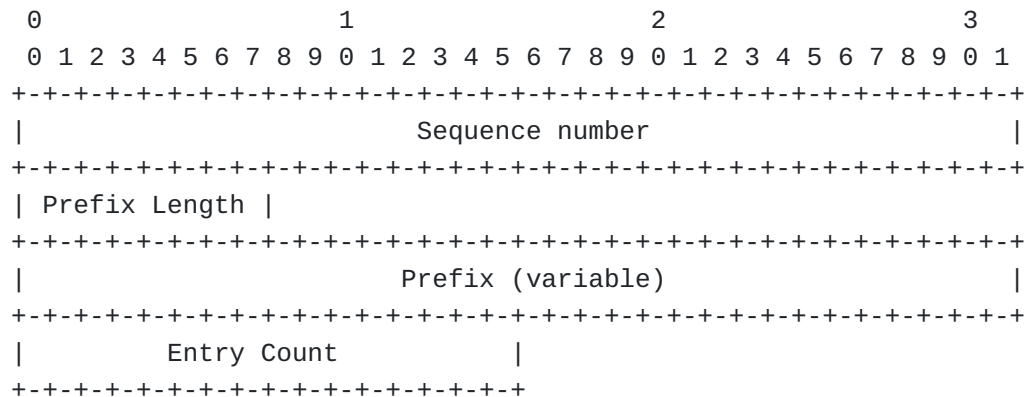
```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Peer Type  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Peer BGP ID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Peer IP address (variable)                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Peer AS (variable)                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The Peer Type field is a bit field which encodes the type of the AS and IP address as follows:

Bit 0 - unset for IPv4 Peer IP address, set for IPv6
 Bit 1 - unset when Peer AS field is 16 bits, set when it's 32 bits

The format for the RIB_IPV4_UNICAST, RIB_IPV4_MULTICAST, RIB_IPV6_UNICAST, and RIB_IPV6_MULTICAST headers are shown below. The Prefix Length and Prefix fields are encoded in the same manner as the BGP NLRI encoding for IPV4 and IPV6 prefixes. Namely, the Prefix field contains address prefixes followed by enough trailing bits to make the end of the field fall on an octet boundary. Note that the value of trailing bits is irrelevant.



The RIB_GENERIC header is shown below. It includes Address Family Identifier (AFI), Subsequent AFI and a single NLRI entry. The NLRI information is specific to the AFI and SAFI values. An implementation which does not recognize particular AFI and SAFI values SHOULD discard the remainder of the MRT record.

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Sequence number                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Address Family Identifier   |Subsequent AFI |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Network Layer Reachability Information (variable)   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Entry Count           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The RIB entry headers are followed by a series of RIB entries which are repeated Entry Count times. These entries share a common format as shown below. They include a Peer Index from the PEER_INDEX_TABLE MRT record, an originated time for the RIB entry, and the BGP path attribute length and attributes encoded as provided in a BGP Update message.

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Peer Index           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Originated Time                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Attribute Length           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               BGP Attributes... (variable)
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

There is one exception to the encoding of BGP attributes for the BGP MP_REACH_NLRI attribute (BGP Type Code 14) [RFC 4760]. Since the AFI, SAFI, and NLRI information is already encoded in the MULTIPROTOCOL header, only the Next Hop Address Length and Next Hop Address fields are included. The Reserved field is omitted. The attribute length is also adjusted to reflect only the length of the Next Hop Address Length and Next Hop Address fields.

5.9. BGP4MP Type

[TOC](#)

This Type was initially defined in the Zebra software package for the BGP protocol with multiprotocol extension support as defined by [RFC 4760 \(Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4," January 2007.\)](#) [RFC4760]. It supersedes the BGP,

BGP4PLUS, BGP4PLUS_01 Types. The BGP4MP Type has four Subtypes which are defined as follows:

0	BGP4MP_STATE_CHANGE	
1	BGP4MP_MESSAGE	
2	BGP4MP_ENTRY	*DEPRECATED*
3	BGP4MP_SNAPSHOT	
4	BGP4MP_STATE_CHANGE_AS4	
5	BGP4MP_MESSAGE_AS4	

5.9.1. BGP4MP_STATE_CHANGE Subtype

TOC

This record is used to encode state changes in the BGP finite state machine. As with the BGP_STATE_CHANGE Subtype, the BGP FSM states are encoded in the Old State and New State fields to indicate the previous and current state. The format is illustrated below:

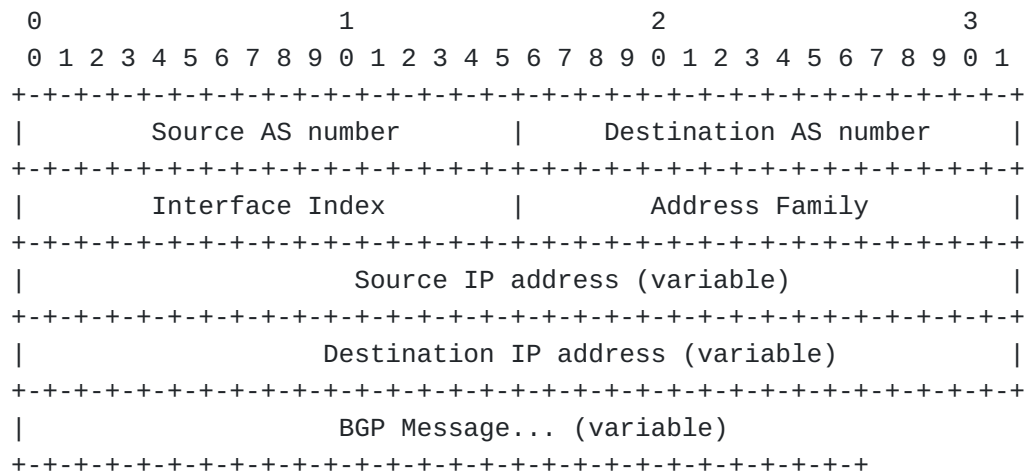
[illegible]

While BGP4MP_STATE_CHANGE message is similar to the BGP_STATE_CHANGE message, it also includes interface index and Address Family fields. As with the BGP_STATE_CHANGE message, the FSM states and their numeric encodings are defined in [RFC 4271 \(Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 \(BGP-4\)," January 2006.\)](#) [RFC4271], Appendix 1. The interface index provides the interface number of the peering session. The index value is OPTIONAL and MAY be zero if unknown or unsupported. The Address Family indicates what types of addresses are in the the address fields. At present, the following AFI Types are supported:

```
1  AFI_IPv4
2  AFI_IPv6
```

TOC

This Subtype is used to encode BGP Messages. It is similar to the BGP_UPDATE Subtype, except that it can be used to encode any Type of message (not just BGP UPDATES). In order to determine the BGP message Type, the entire BGP message, including the BGP header, is included in the BGP Message field. The BGP4MP_MESSAGE fields are shown below:



The interface index provides the interface number of the peering session. The index value is OPTIONAL and MAY be zero if unknown or unsupported. The Address Family indicates what types of addresses are in the the subsequent address fields. At present, the following AFI Types are supported:

- ```
1 AFI_IPv4
2 AFI_IPv6
```

Note that the Address Family value only applies to the IP addresses contained in the MRT header. The BGP4MP\_MESSAGE Subtype is otherwise transparent to the contents of the actual message which may contain any valid AFI/SAFI values. Only one BGP message may be encoded in the BGP4MP\_MESSAGE Subtype.

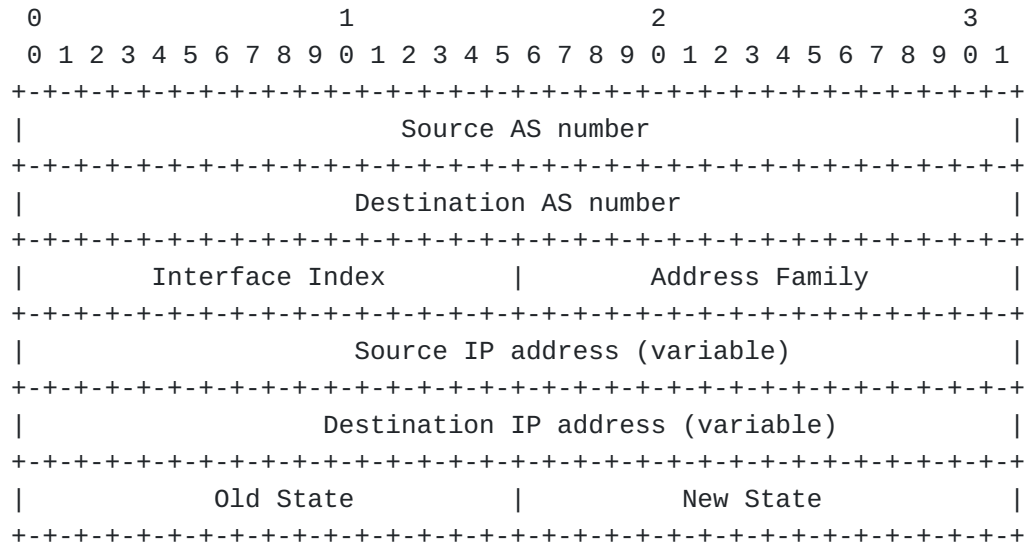
## TOC

This Subtype is similar to the TABLE\_DUMP Type and is used to record RIB table entries. It extends the TABLE\_DUMP Type to include true multiprotocol support. However, this type does not support 32-bit AS numbers and has not been widely implemented. This type is deprecated in favor of the TABLE\_DUMP\_V2 which includes 32-bit AS number support and a more compact format.



## TOC

This Subtype updates the BGP4MP\_STATE\_CHANGE Subtype to support 32BIT Autonomous System numbers. As with the BGP4MP\_STATE\_CHANGE Subtype, the BGP FSM states are encoded in the Old State and New State fields to indicate the previous and current state. Aside from the extension of the source and destination AS fields to 32 bits, this subtype is otherwise identical to the BGP4MP\_STATE\_CHANGE Subtype. The BGP4MP\_STATE\_CHANGE\_AS4 fields are shown below:



TOC

This Subtype updates the BGP4MP\_MESSAGE Subtype to support 32BIT Autonomous System numbers. The BGP4MP\_MESSAGE\_AS4 fields are shown below:



### 5.11. ISIS Type

This Type was initially defined in the Sprint Labs Python Routing and supports the IS-IS routing protocol as defined in [RFC 1195 \(Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments," December 1990.\)](#) [RFC1195]. There is no Type specific header for the ISIS Type. The Subtype code for this Type is undefined. The ISIS PDU directly follows the MRT common header fields.

---

### 5.12. ISIS\_ET Type

[TOC](#)

The ISIS\_ET Type extends the the ISIS Type to support microsecond timestamps. As with the BGP4MP\_ET Type, a 32-bit microsecond timestamp field is appended to the MRT common header after the length field. The ISIS\_ET Type is otherwise identical to the ISIS Type.

---

### 5.13. OSPFv3 Type

[TOC](#)

The OSPFv3 Type extends the original OSPF Type to support IPv6 addresses for the OSPFv3 protocol as defined in [RFC 2740 \(Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6," December 1999.\)](#) [RFC2740]. The format of the MRT Message field for the OSPFv3 Type is 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
+--+
| Address Family |
+--+
| Source IP address (variable) |
+--+
| Destination IP address (variable) |
+--+
| OSPF Message Contents (variable) |
+--+
```

### 5.14. OSPFv3\_ET Type

[TOC](#)

The OSPFv3\_ET Type extends the the OSPFv3 Type to support microsecond timestamps. As with the BGP4MP\_ET Type, a 32-bit microsecond timestamp field is appended to the MRT common header after the length field and

its length is included in the calculation of the length field value. The OSPFv3\_ET Type is otherwise identical to the OSPFv3 Type.

---

## 6. IANA Considerations

[TOC](#)

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to the MRT specification, in accordance with BCP 26, [RFC 2434 \(Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs," October 1998.\)](#) [RFC2434].

There are two name spaces in MRT that require registration: Type Codes and Subtype Codes.

MRT is not intended as a general-purpose specification for protocol information export, and allocations should not be made for purposes unrelated to routing protocol information export.

The following policies are used here with the meanings defined in BCP 26: "Specification Required", "IETF Consensus".

---

### 6.1. Type Codes

[TOC](#)

Type Codes have a range from 0 to 65535, of which 0-64 have been allocated. New Type Codes MUST be allocated starting at 65. Type Codes 65 - 32767 are to be assigned by IETF Consensus. Type Codes 32768 - 65535 are assigned based on Specification Required.

---

### 6.2. Subtype Codes

[TOC](#)

Subtype Codes have a range from 0 to 65535. Subtype definitions are specific to a particular Type Code definition. New Subtype Code definition must reference an existing Type Code to which the Subtype belongs. As Subtype Codes are specific to Type Codes, new numbers must be unique for the particular Type Code to which the Subtype applies. Subtype Codes specific to the Type Codes 0 - 32767 are assigned by IETF Consensus. Subtype Codes specific to Type Codes 32768 - 65535 are assigned based on Specification Required.

---

[TOC](#)

## 7. Security Considerations

The MRT Format utilizes a structure which can store routing protocol information data. The fields defined in the MRT specification are of a descriptive nature and provide information that is useful to facilitate the analysis of routing data. As such, the fields currently defined in the MRT specification do not in themselves create additional security risks, since the fields are not used to induce any particular behavior by the recipient application.

---

## 8. References

[TOC](#)

### 8.1. Normative References

[TOC](#)

|           |                                                                                                                                                                                                                                                        |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [RFC1058] | Hedrick, C., " <a href="#">Routing Information Protocol</a> ," RFC 1058, June 1988 ( <a href="#">TXT</a> ).                                                                                                                                            |
| [RFC1195] | Callon, R., " <a href="#">Use of OSI IS-IS for routing in TCP/IP and dual environments</a> ," RFC 1195, December 1990 ( <a href="#">TXT</a> , <a href="#">PS</a> ).                                                                                    |
| [RFC2080] | <a href="#">Malkin, G.</a> and <a href="#">R. Minnear</a> , " <a href="#">RIPng for IPv6</a> ," RFC 2080, January 1997 ( <a href="#">TXT</a> ).                                                                                                        |
| [RFC2119] | <a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).                                         |
| [RFC2328] | <a href="#">Moy, J.</a> , " <a href="#">OSPF Version 2</a> ," STD 54, RFC 2328, April 1998 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).                                                                                       |
| [RFC2434] | <a href="#">Narten, T.</a> and <a href="#">H. Alvestrand</a> , " <a href="#">Guidelines for Writing an IANA Considerations Section in RFCs</a> ," BCP 26, RFC 2434, October 1998 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ). |
| [RFC2740] | <a href="#">Coltun, R.</a> , <a href="#">Ferguson, D.</a> , and <a href="#">J. Moy</a> , " <a href="#">OSPF for IPv6</a> ," RFC 2740, December 1999 ( <a href="#">TXT</a> ).                                                                           |
| [RFC4271] | <a href="#">Rekhter, Y.</a> , <a href="#">Li, T.</a> , and <a href="#">S. Hares</a> , " <a href="#">A Border Gateway Protocol 4 (BGP-4)</a> ," RFC 4271, January 2006 ( <a href="#">TXT</a> ).                                                         |
| [RFC4760] | <a href="#">Bates, T.</a> , <a href="#">Chandra, R.</a> , <a href="#">Katz, D.</a> , and <a href="#">Y. Rekhter</a> , " <a href="#">Multiprotocol Extensions for BGP-4</a> ," RFC 4760, January 2007 ( <a href="#">TXT</a> ).                          |

### 8.2. Informative References

[TOC](#)

|                  |                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------|
| [MRT PROG GUIDE] | <a href="#">Labovitz, C.</a> , " <a href="#">MRT Programmer's Guide</a> ," November 1999 ( <a href="#">HTML</a> ). |
|------------------|--------------------------------------------------------------------------------------------------------------------|

---

## Authors' Addresses

[TOC](#)

|        |                                                          |
|--------|----------------------------------------------------------|
|        | Larry Blunk                                              |
|        | Merit Network                                            |
| Email: | <a href="mailto:ljb@merit.edu">ljb@merit.edu</a>         |
|        |                                                          |
|        | Manish Karir                                             |
|        | Merit Network                                            |
| Email: | <a href="mailto:mkarir@merit.edu">mkarir@merit.edu</a>   |
|        |                                                          |
|        | Craig Labovitz                                           |
|        | Arbor Networks                                           |
| Email: | <a href="mailto:labovit@arbor.net">labovit@arbor.net</a> |

---

## Full Copyright Statement

[TOC](#)

Copyright © The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights

that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).