



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

RIPE NCC Updates on Routing Information Service (RIS)



BGPdump

Upcoming changes

BGPdump



- New version will be pushed soon
- Support for `_LOCAL` subtypes (Colin Petrie)
- Support for `_ET` types (Alexis Fasquel)
- Support for `_ADDPATH` subtypes (Colin Petrie)
- Sequence number output (Max Lapan)
- Output format changes - surprise!!!

BGPdump - Normal Output?



- First column
 - TABLE_DUMP
 - TABLE_DUMP2
 - BGP
 - BGP4MP
- Have you ever expected anything else in the first column?

```
BGP4MP|1279829997|W|193.203.0.139|3303|190.186.128.0/24|
BGP4MP|1279829997|W|193.203.0.139|3303|187.120.32.0/24|
BGP4MP|1279829997|W|193.203.0.139|3303|91.199.82.0/24|
BGP4MP|1279829997|A|193.203.0.139|3303|91.213.6.0/24|
BGP4MP|1279829997|A|193.203.0.139|3303|207.254.176.0/24|
BGP4MP|1279829999|A|193.203.0.97|286|110.172.163.0/24|
BGP4MP|1279829999|A|193.203.0.91|13237|110.172.163.0/24|
BGP4MP|1279829999|A|193.203.0.97|286|110.172.163.0/24|
BGP4MP|1279830000|A|193.203.0.21|8447|186.216.64.0/24|
BGP4MP|1279830000|A|193.203.0.21|8447|187.120.32.0/24|
BGP4MP|1279830000|A|193.203.0.21|8447|207.254.176.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.0.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.206.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|200.83.0.0/19|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.213.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.47.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.192.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.90.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.205.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.215.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.219.0/24|
BGP4MP|1279830000|A|193.203.0.91|13237|190.160.218.0/24|
```

Current output of BGPdump

BGPdump - _LOCAL



- First column
 - BGP4MP_LOCAL
- Change in semantics of 4th and 5th fields (source IP, source AS)
- Now **destination IP, destination AS**

```
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.1 | 1853 | 130.36.35.0/24 | 1853  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.1 | 1853 | 130.36.34.0/24 | 1853  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.1 | 1853 | 182.48.252.0/24 | 1853  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.1 | 1853 | 91.213.6.0/24 | 1853  
BGP4MP_LOCAL | 1279829996 | W | 193.203.0.97 | 286 | 187.120.32.0/20  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.97 | 286 | 187.120.32.0/20 | 286  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.97 | 286 | 110.172.163.0/24 | 286  
BGP4MP_LOCAL | 1279829996 | A | 193.203.0.97 | 286 | 187.120.32.0/20 | 286  
BGP4MP_LOCAL | 1279829997 | A | 193.203.0.139 | 3303 | 186.216.64.0/20 | 3303  
BGP4MP_LOCAL | 1279829997 | W | 193.203.0.139 | 3303 | 190.186.128.0/19  
BGP4MP_LOCAL | 1279829997 | W | 193.203.0.139 | 3303 | 187.120.32.0/20  
BGP4MP_LOCAL | 1279829997 | W | 193.203.0.139 | 3303 | 91.199.82.0/24  
BGP4MP_LOCAL | 1279829997 | A | 193.203.0.139 | 3303 | 91.213.6.0/24 | 3303  
BGP4MP_LOCAL | 1279829997 | A | 193.203.0.139 | 3303 | 207.254.176.0/20 | 3303  
BGP4MP_LOCAL | 1279829999 | A | 193.203.0.97 | 286 | 110.172.163.0/24 | 286  
BGP4MP_LOCAL | 1279829999 | A | 193.203.0.91 | 13237 | 110.172.163.0/24 | 13237
```

Locally originated output of BGPdump

see: RFC6396, section 4.4.5 and section 4.4.6

BGPdump - _ET



- First column
 - BGP4MP_ET
- Change in semantics of 2nd field (timestamp)
- Now high precision
- No longer an integer

```
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.127.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.126.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.80.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.75.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.74.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.72.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.70.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.69.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.68.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.45.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.44.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.42.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.41.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.30.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.29.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.28.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.27.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.26.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.25.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.23.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.22.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.21.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.20.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|74.125.6.0/24|
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|66.249.72.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|66.249.70.0/24
BGP4MP_ET|1445565695.585328|A|206.220.231.55|3856|66.249.60.0/24
```

Extended Timestamp output of BGPdump

see: RFC6396, section 3

BGPdump - _ADDPATH



- First column
 - BGP4MP_AP
- New extra column after prefix
- Contains **Path Identifier** for prefix
- 32-bit int

```
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2401:da80::/32|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:20f8::/48|11|51
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:20f8::/48|9|518
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:24fc::/48|11|51
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:24fc::/48|9|518
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a02:d740::/29|11|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a02:d740::/29|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a04:d140::/29|11|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a04:d140::/29|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2804:138b:c000::/36|11|5
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2804:138b:c000::/36|9|51
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:d8::/48|11|5186
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:d8::/48|9|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a03:2880:f00e::/48|11|5
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:4dc::/48|11|518
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2001:67c:4dc::/48|9|5186
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a01:5480:8000::/33|11|5
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a01:5480:8000::/33|9|51
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a03:d600::/32|11|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2a03:d600::/32|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2607:de00::/32|11|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2607:de00::/32|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2403:d800::/32|11|51861
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2403:d800::/32|9|51861 6
BGP4MP_AP|1446701617|A|fec0::a:cf:0:1|51861|2804:14d:8000::/40|11|51
```

Add-Path enabled output of BGPdump

see: <https://tools.ietf.org/html/draft-petrie-grow-mrt-add-paths>

BGPdump - Index (-p flag)



- New extra second column
- Sequence number of BGP message within MRT file
- Run bgpdump -p to get this output

```
BGP4MP|53|1279829713|A|193.203.0.97|286|209.169.32.0/24|
BGP4MP|54|1279829714|A|2001:7f8:30::2:1:1:3030|13030|20|
BGP4MP|57|1279829715|A|193.203.0.97|286|89.221.30.0/24|
BGP4MP|58|1279829715|W|193.203.0.1|1853|214.6.167.0/24|
BGP4MP|59|1279829715|A|193.203.0.1|1853|182.48.252.0/24|
BGP4MP|60|1279829715|A|193.203.0.1|1853|187.63.96.0/20|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.133.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.147.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.146.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.139.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.135.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.148.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.140.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.141.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.137.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.134.0/24|
BGP4MP|61|1279829715|A|193.203.0.1|1853|77.69.138.0/24|
```

Packet index output of BGPdump

BGPdump - Possible Combinations



```
/* Helper function that returns the format of a BGP4MP_MESSAGE as a string:
- BGP4MP (Standard)
- BGP4MP_ET (Extended Header)
- BGP4MP_LOCAL (Local)
- BGP4MP_ET_LOCAL (Local and Extended Header)
- BGP4MP_AP (Standard, Additional Paths)
- BGP4MP_ET_AP (Extended Header, Additional Paths)
- BGP4MP_LOCAL_AP (Local, Additional Paths)
- BGP4MP_ET_LOCAL_AP (Local, Extended Header, Additional Paths)
*/
```

BGP4MP format output permutations

- All the permutations for BGP4MP messages
- Also, TABLE_DUMP2 / TABLE_DUMP2_AP

BGPdump - Summary



- New MRT types and subtypes now supported
- Change in field types (ET)
- Change in field semantics (LOCAL)
- Change in number of fields (ADDPATH, Index)
- But of course, you wrote all your code to handle exceptional and unexpected code types cleanly, didn't you?



BGPdump

Should we stop this?
Is this getting messy?


BGPdump Output



- Inventing new output combinations whenever a new feature is added can break parsers
- Existing BGPdump output is based on `route_btoa.c`, from the original MRT suite, circa 2000
- It is **so** 15 years ago - things have moved on
- Should BGPdump stop trying to be compatible, and define a new, modern output format?

BGPdump Output

- Ideas:
 - Output structured serialised data
 - JSON, with version numbers and a published schema
 - Any modern application can load JSON into an appropriate data structure
- Example:
 - ExaBGP does this already



```
"message": {
  "update": {
    "attribute": {
      "origin": "igp",
      "as-path": [
        202194,
        198385,
        174,
        1299,
        7029,
        6316
      ],
      "confederation-path": [],
      "community": [
        [
          174,
          21100
        ],
        [
          174,
          22005
        ],
        [
          20219,
          19838
        ]
      ]
    },
    "announce": {
      "ipv4 unicast": {
        "91.206.52.60": {
          "66.19.194.0/24": {}
        }
      }
    }
  }
}
```

Example of ExaBGP JSON output

BGPdump Output



- There is no RFC for this ;-)
- Can we change it?
- Should we change it?
- Your feedback please!



RIS Sent Messages

Worthwhile?

BGP Sent Messages



- Quagga/Zebra/Bird et al only store **received** BGP messages in MRT
- **But** the BGP session is bi-directional
- Our collectors do **send** BGP messages:
 - OPEN message at session startup
 - KEEPALIVE messages periodically
 - UPDATE messages for Routing Beacons
 - NOTIFY message at session tear-down
- These messages are not recorded

BGP Sent Messages



- MRT subtype to represent locally originated messages already exists in the spec (and now supported in BGPdump)
- Question - should we record this data?
- Simple Quagga patch
- Will record complete BGP conversation from both sides
- Is this useful?
- Or will it just add more noise?



New RIS Infrastructure

Data now being produced

New RIS Data




- Available via RIPEstat
- Three new RRCs
- MRT data is being collected (but not yet published)
- Looking glass not yet supported

Organisation name	
Contact name	
Contact e-mail*	
Contact phone	
NOC e-mail*	
NOC phone	
AS Number*	
Peering IPv4 Address*	
Peering IPv6 Address*	
RIS Route Collector*	<div><input checked="" type="checkbox"/> CIXP (RRC04) <input type="checkbox"/> DIX-IE (RRC06) <input type="checkbox"/> Netnod (RRC07) <input type="checkbox"/> MIX (RRC10) <input type="checkbox"/> NYIIX (RRC11) <input type="checkbox"/> MSK-IX (RRC13) <input type="checkbox"/> PAIX (RRC14) <input type="checkbox"/> PTTMetro (RRC15) <input type="checkbox"/> CATNIX (RRC18) <input type="checkbox"/> SwissIX (RRC20) <input type="checkbox"/> FranceIX Paris/Marseille (RRC21)</div>
AS Macros	
Router vendor	
OS version	

The RIS peering request form

https://labs.ripe.net/Members/colin_petrie/updates-to-the-ripe-ncc-routing-information-service

New RIS Data



New RIS Data



Visibility Location Details of AS3333						
RRC	IXP Location	Location	IPv4 peers seeing	IPv6 peers seeing	IPv4 Visibility	IPv6 Visibility
RRC00	RIPE-NCC Multihop	Amsterdam, Netherlands	14 of 14	10 of 10	100%	100%
RRC01	LINX	London, United Kingdom	11 of 11	10 of 10	100%	100%
RRC03	AMS-IX / NL-IX	Amsterdam, Netherlands	8 of 8	12 of 12	100%	100%
RRC04	CIXP	Geneva, Switzerland	7 of 7	4 of 4	100%	100%
RRC05	VIX	Vienna, Austria	6 of 6	7 of 7	100%	100%
RRC06	DIX-IE	Tokyo, Japan	2 of 2	2 of 2	100%	100%
RRC07	Netnod	Stockholm, Sweden	2 of 2	4 of 4	100%	100%
RRC10	MIX	Milan, Italy	9 of 9	7 of 7	100%	100%
RRC11	NYIIX	New York City, US	8 of 8	8 of 8	100%	100%
RRC12	DE-CIX	Frankfurt, Germany	15 of 15	20 of 20	100%	100%
RRC13	MSK-IX	Moscow, Russian Federation	11 of 11	5 of 5	100%	100%
RRC14	PAIX	Palo Alto, US	7 of 7	7 of 7	100%	100%
RRC15	PTTMetro	Sao Paulo, Brazil	10 of 10	10 of 10	100%	100%
RRC18	Catnix	Barcelona, Spain	2 of 2	1 of 1	100%	100%
RRC20	SwissIX	Zurich, Switzerland	12 of 12	10 of 10	100%	100%
RRC21	France-IX	Paris, France	12 of 12	13 of 13	100%	100%



New RIS Infrastructure

Streaming feedback

New RIS Streaming



- Demo at RIPE 70
- Architecture explained in RIPE Labs article
 - https://labs.ripe.net/Members/colin_petrie/updates-to-the-ripe-ncc-routing-information-service
- Feedback requested
- Very little received!

New RIS Streaming



- Questions:
 - Do you think the live BGP streaming is useful?
 - Are you interested in developing applications using this? What kind of applications can you think of?
 - Would you be interested in a full stream of all the available BGP data, or just a subset? What criteria would you want to filter it by? (e.g. RRC, origin ASN, prefix, etc.)
 - Would your application need the stream to be resumable, or is it not important if you miss messages when you are disconnected from the stream?
- Please give us your thoughts!



Questions



cpetrie@ripe.net
ris-users@ripe.net
routing-wg@ripe.net