

Bgpdump2: A Tool for Full BGP Route Comparison

Yasuhiro Ohara

yasu@nttv6.jp

NTT Communications Corporation

Motivation

- Need to handle BGP full route dump data
 - Obtain an instance of a full-route routing table (needed for a routing research)
 - Health check, past analysis, future model in ISP
- Comparison on full routes between ISPs
- Lack of a flexible tool (i.e., flexible to modify)

Existing tools

- libbgpdump
 - written in C. <<https://bitbucket.org/ripenc/bgpdump/wiki/Home>>
- zebra-dump-parser
 - written in Perl. <<https://github.com/rfc1036/zebra-dump-parser>>
- java-mrt library
 - written in Java. <<https://github.com/paaguti/java-mrt>>
- UCLA bgpparser
 - written in C++. <<http://irl.cs.ucla.edu/software/bgpparser.html>>
- mrtparse
 - written in Python. <<https://github.com/YoshiyukiYamauchi/mrtparse>>

bgpdump2 summary

- Full scratch in C (4000- lines)
- Open-source software
 - <https://github.com/yasuhiro-ohara-ntt/bgpdump2>
- Capability:
 - supported: bz2, MRT TABLE_DUMP_V2 format, ipv4 / ipv6
 - show statistics per peers (e.g., #routes, #nexthops)
 - routing table construction, longest-match table lookup
- Good for detailed investigation of a RIB file
- Good for full-route comparison

simple display

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 |  
head -10  
0.0.0.0/0 203.189.128.233 origin_as: 9902 as-path[2]: 23673 9902  
1.0.0.0/24 213.144.128.203 origin_as: 15169 as-path[2]: 13030  
15169  
1.0.0.0/24 198.129.33.85 origin_as: 15169 as-path[2]: 293 15169  
1.0.0.0/24 5.101.110.2 origin_as: 15169 as-path[3]: 202018 1299  
15169  
1.0.0.0/24 162.243.188.2 origin_as: 15169 as-path[2]: 393406 15169  
1.0.0.0/24 95.85.0.2 origin_as: 15169 as-path[2]: 200130 15169  
1.0.0.0/24 192.241.164.4 origin_as: 15169 as-path[2]: 62567 15169  
1.0.0.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169  
1.0.0.0/24 66.185.128.1 origin_as: 15169 as-path[2]: 1668 15169  
1.0.0.0/24 173.205.57.234 origin_as: 15169 as-path[3]: 53364 3257  
15169
```

libbgpdump compatible mode

- not finished: e.g., support for community, localpref, etc., are not yet completed.

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.
```

```
20150319.0000.bz2 -m | head -3
```

```
TABLE_DUMP2|1426723200|B|203.189.128.233|23673|  
0.0.0.0/0|23673 9902|INCOMPLETE|203.189.128.233|0|0|0|  
NAG||
```

```
TABLE_DUMP2|1426723200|B|213.144.128.203|13030|  
1.0.0.0/24|13030 15169|INCOMPLETE|213.144.128.203|0|1|0|  
NAG||
```

```
TABLE_DUMP2|1426723200|B|198.129.33.85|293|1.0.0.0/24|  
293 15169|INCOMPLETE|198.129.33.85|0|0|0|NAG||
```

Speed

74M routeviews/oregon-ix2/rib.20150319.0000.bz2

```
% /usr/bin/time -p bash -c 'bzipcat routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '
```

```
real    28.98 user    28.80 sys     0.10
```

```
% /usr/bin/time -p bash -c 'bzipcat routeviews/oregon-ix2/rib.20150319.0000.bz2 | ./zebra-dump-parser/zebra-dump-parser.pl > /dev/null 2>&1 '
```

```
real    428.29 user    467.67 sys     1.07
```

```
% /usr/bin/time -p bash -c './libbgpdump-1.4.99.11/bgpdump routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '
```

```
real    148.71 user    148.38 sys     0.24
```

```
% /usr/bin/time -p bash -c './bgpdump2/src/bgpdump2 routeviews/oregon-ix2/rib.20150319.0000.bz2 > /dev/null 2>&1 '
```

```
real    89.09 user    88.74 sys     0.19
```

display peer table

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -  
P | head -12
```

```
Collector BGP ID: 128.223.51.102
```

```
View Name Length: 0
```

```
View Name:
```

```
Peer Count: 58
```

```
# peer_table[0] changed: 0.0.0.0 asn:0 [129.250.0.11|::]
```

```
# peer_table[1] changed: 10.10.10.252 asn:53364  
[173.205.57.234|::]
```

```
# peer_table[2] changed: 0.0.0.0 asn:0 [192.241.164.4|::]
```

```
# peer_table[3] changed: 4.69.184.193 asn:3356 [4.69.184.193|::]
```

```
# peer_table[4] changed: 5.101.110.2 asn:202018 [5.101.110.2|::]
```

```
# peer_table[5] changed: 12.0.1.63 asn:7018 [12.0.1.63|::]
```

```
# peer_table[6] changed: 64.57.28.241 asn:11537 [64.57.28.241|::]
```

```
# peer_table[7] changed: 66.185.128.1 asn:1668 [66.185.128.1|::]
```


display per peer

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -  
p 19 | head -8
```

```
1.0.0.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169
```

```
1.0.4.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257  
4826 56203
```

```
1.0.5.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257  
4826 56203
```

```
1.0.6.0/24 129.250.0.11 origin_as: 56203 as-path[4]: 2914 3257  
4826 56203
```

```
1.0.7.0/24 129.250.0.11 origin_as: 56203 as-path[6]: 2914 3257  
4826 56203 56203 56203
```

```
1.0.38.0/24 129.250.0.11 origin_as: 24155 as-path[2]: 2914 24155
```

```
1.0.43.0/24 129.250.0.11 origin_as: 24155 as-path[4]: 2914 1299  
10026 24155
```

```
1.0.44.0/24 129.250.0.11 origin_as: 24155 as-path[4]: 2914 1299  
10026 24155
```

routing table lookup (longest-match)

- Routing Table (PATRICIA) can be created.
- one address lookup.

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -  
p 19 -l 8.8.8.8
```

```
looking up an address: 8.8.8.8
```

```
8.8.8.0/24 129.250.0.11 origin_as: 15169 as-path[2]: 2914 15169
```

- a list of address lookup contained in a file.

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -  
p 19 -L ../test-lookup.txt
```

```
50.0.0.0/16 129.250.0.11 origin_as: 7065 as-path[3]: 2914 8121  
7065
```

```
100.0.0.0/16 129.250.0.11 origin_as: 701 as-path[2]: 2914 701
```

```
150.0.0.0: no route found.
```

```
200.0.0.0: no route found.
```

per peer statistics

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.20150319.0000.bz2 -k | head -20
```

```
peer[0]:
```

```
Number of routes: 0
```

```
Number of routes per plen:
```

```
/0 : 0 /1 : 0 /2 : 0 /3 : 0 /4 : 0  
/5 : 0 /6 : 0 /7 : 0 /8 : 0 /9 : 0  
/10: 0 /11: 0 /12: 0 /13: 0 /14: 0  
/15: 0 /16: 0 /17: 0 /18: 0 /19: 0  
/20: 0 /21: 0 /22: 0 /23: 0 /24: 0  
/25: 0 /26: 0 /27: 0 /28: 0 /29: 0  
/30: 0 /31: 0 /32: 0
```

```
Number of nexthops: 0
```

```
Number of origin_as: 0
```

```
Number of unique as paths: 0
```

```
Number of as path len: 0
```

```
peer[1]:
```

```
Number of routes: 525361
```

```
Number of routes per plen:
```

```
/0 : 0 /1 : 0 /2 : 0 /3 : 0 /4 : 0  
/5 : 0 /6 : 0 /7 : 0 /8 : 16 /9 : 12  
/10: 33 /11: 93 /12: 264 /13: 504 /14: 1000
```

diff between peers

```
% ./src/bgpdump2 ../routeviews/oregon-ix2/rib.  
20150319.0000.bz2 -u -p 1 -p 3
```

```
:
```

```
-204.245.63.0/24 173.205.57.234 origin_as: 3549 as-path[3]:  
53364 3257 3549
```

```
-205.106.84.0/22 173.205.57.234 origin_as: 27066 as-path[6]:  
53364 40191 174 209
```

```
721 27066
```

```
+205.232.128.0/21 4.69.184.193 origin_as: 20208 as-path[4]:  
3356 7843 12271 2020
```

```
8
```

```
+205.251.194.0/23 4.69.184.193 origin_as: 16509 as-path[2]:  
3356 16509
```

```
:
```

Diffs a peer against others

- `script/compare-peer-with-others.sh`
- `marks`

udiff:

<: the prefix is in left and unreachable in right (maybe partially).

+: the prefix is only in right (but it is reachable in left).

): the prefix is in right, and is covered by a shorter prefix in left that is unreachable in right. (i.e., the shorter is '<')

>: the prefix is in right and unreachable in left (maybe partially).

-: the prefix is only in left (but it is reachable in right).

(: the prefix is in left, and is covered by a shorter prefix in right that is unreachable in left. (i.e., the shorter is '>')

Comparison among peers

- diff routes between 2914 and 3356 are shown like:
>70.61.2.0/24 4.69.184.193 origin_as: 30628 as-path[4]: 3356 7843 10796 30628
(70.61.4.0/24 129.250.0.11 origin_as: 10796 as-path[4]: 2914 2828 7843 10796
(70.61.5.0/24 129.250.0.11 origin_as: 10796 as-path[4]: 2914 2828 7843 10796
(70.61.6.0/24 129.250.0.11 origin_as: 33363 as-path[5]: 2914 2828 7843 33363 333
63
+71.29.112.0/21 4.69.184.193 origin_as: 7029 as-path[3]: 3356 2828 7029
>71.44.17.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
>71.44.53.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
>71.44.62.0/24 4.69.184.193 origin_as: 33363 as-path[2]: 3356 33363
-71.252.67.0/24 129.250.0.11 origin_as: 64512 as-path[3]: 2914 701 64512
+72.0.224.0/23 4.69.184.193 origin_as: 19940 as-path[2]: 3356 19940
+72.0.227.0/24 4.69.184.193 origin_as: 19940 as-path[2]: 3356 19940

missing prefix ranking in AS 2914 in Oregon IX

- `cat data/oregon-ix2-rib.20150319.0000.bz2-p19/route-diff-rib.20150319.0000.bz2-p19-*-diff.txt | grep '^>' | awk '{print $1;}' | sort -n | uniq -c | sort -n -r`

```
38 >94.176.2.0/24 94.176.2.0 ASTIMP-ASAstimpConsultingSRL,RO
38 >94.176.131.0/24 94.176.131.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >94.156.77.0/24 94.156.77.0 NETERRA-ASNeterraLtd.,BG
38 >94.156.185.0/24 94.156.185.0 NETERRA-ASNeterraLtd.,BG
38 >94.156.184.0/24 94.156.184.0 NETERRA-ASNeterraLtd.,BG
38 >93.123.18.0/24 93.123.18.0 NETERRA-ASNeterraLtd.,BG
38 >93.120.36.0/22 93.120.36.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >93.120.35.0/24 93.120.35.0 ASTIMP-
ASAstimpConsultingSRL,RO
38 >93.115.92.0/22 93.115.92.0 VOXILITY-ASVoxilityS.R.L.,RO
38 >93.115.88.0/22 93.115.88.0 VOXILITY-ASVoxilityS.R.L.,RO
```

ranking in whois descr appearance count

180

88 TFN-

TWTaiwanFixedNetwork,TelcoandNetworkServiceProvider.,TW

60 VOXILITY-ASVoxilityS.R.L.,RO

60 FPT-AS-

APTheCorporationforFinancing&PromotingTechnology,VN

42 SCRR-10796-TimeWarnerCableInternetLLC,US

37 TAIWANMOBILE-ASTaiwanMobileCo.,Ltd.,TW

33 XTGLOBALXTGLOBALNETWORKSLTD.,RO

32 DATAFRAMELO-DataframeLogistics,Inc.,US

28 ONE-NET-HKINTERNET-SOLUTION-HK,CN

24 VNPT-AS-VNVNPTCorp,VN

Apparently 2914:4429 community tag: do not advertise in Asia.

<http://www.us.ntt.net/support/policy/routing.cfm>

route count on multiple files

```
% ./bgpdump2 -c rib.20150101.bz2 rib.20150102.bz2 rib.  
20150103.bz2 rib.20150104.bz2 .....
```

```
#timestamp,peer1,peer2,...
```

```
1407758405,0,499472,503874,502451,502850,500615,499234  
,500185,0
```

```
#timestamp,peer1,peer2,...
```

```
1407762003,0,499591,503978,502564,502961,500720,499354  
,500311,0
```

```
#timestamp,peer1,peer2,...
```

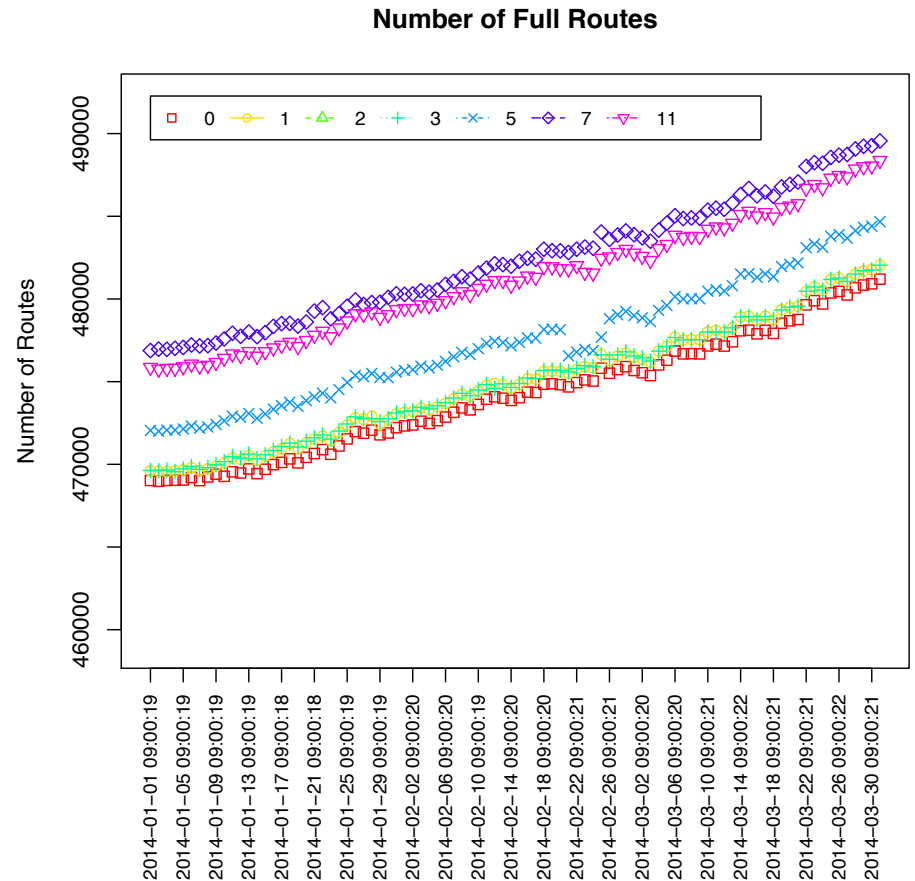
```
1407765605,0,499535,503932,502515,502920,500673,499296  
,500264,0
```

so you can produce a figure like next slides.

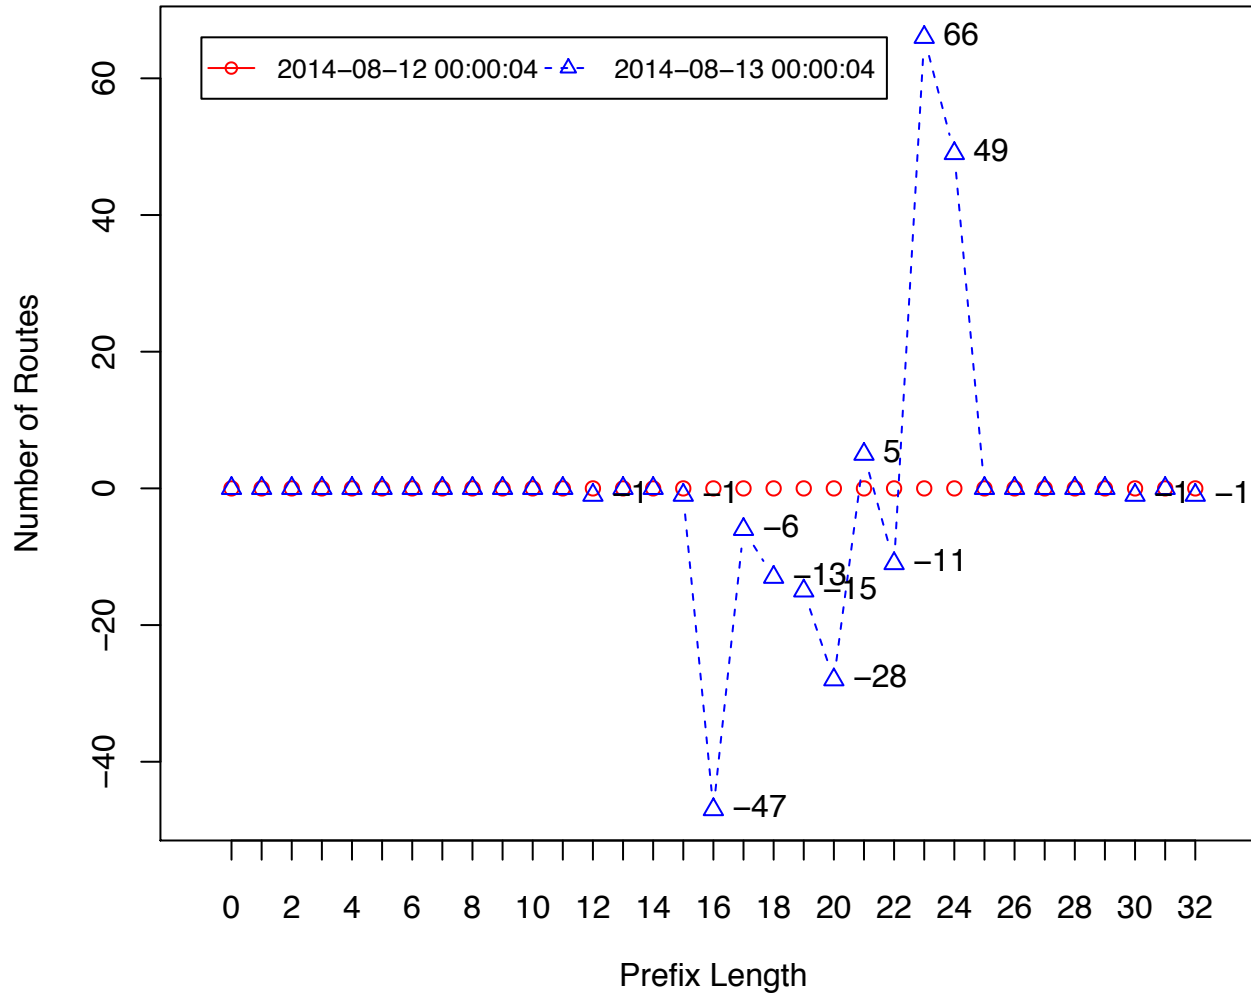
full route number

```
% sh full-route-number.sh -p 0 -p 1 -p 2 -p 3 -p 5 -p 7 -p 11 -p 13  
.././././routeviews/oregon2-summary/rib.20140{1,2,3}* .bz2
```

- peer[13] was empty.
- The R script is also provided.

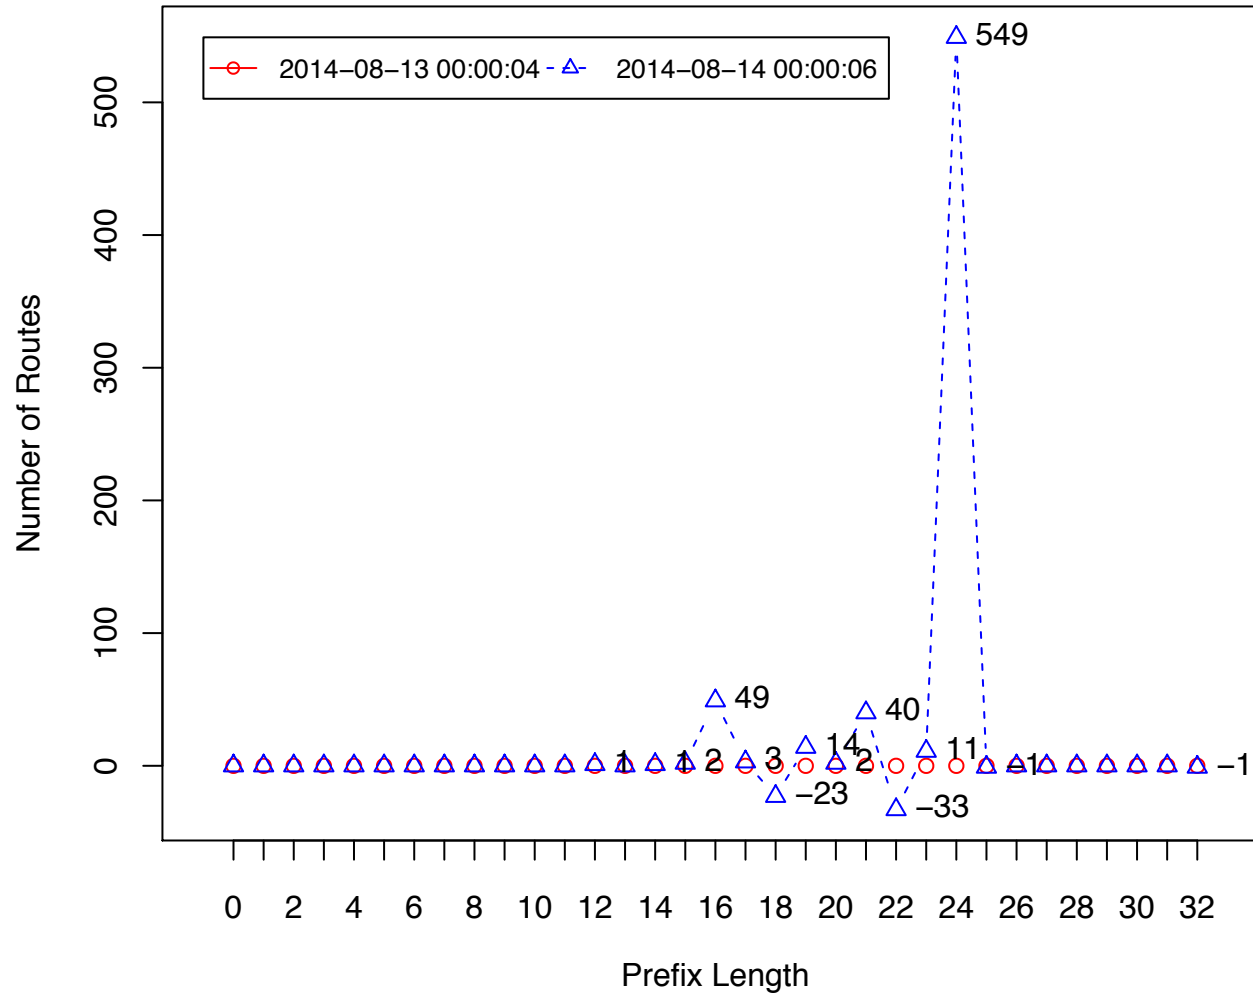


Distribution by Prefix Length

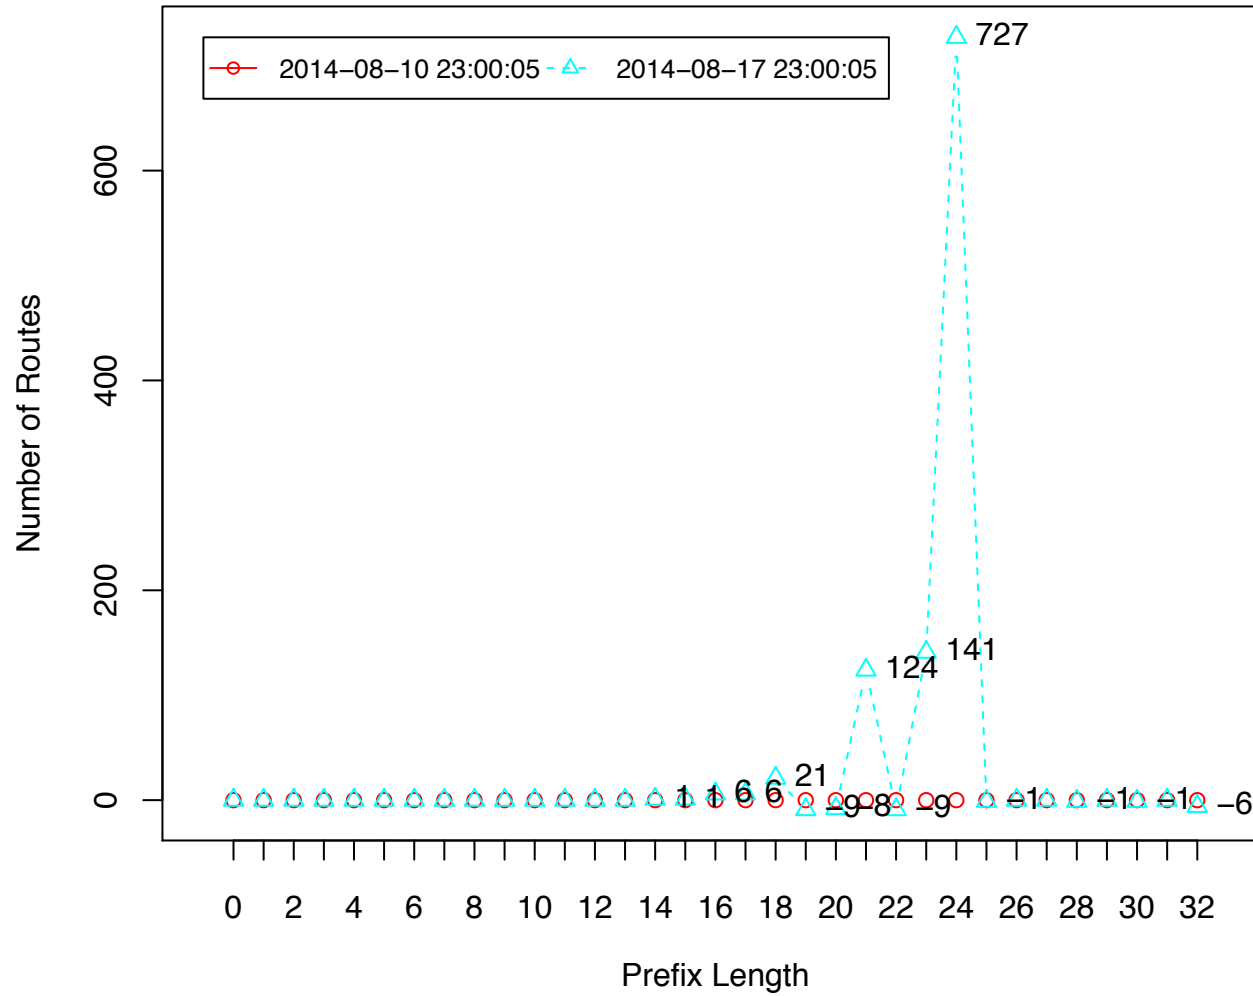


Increase in the number of route.
Disbribution per prefix length.
2014-08-12 00:00 – 2014-08-13 00:00

Distribution by Prefix Length



Distribution by Prefix Length



Discussion

- Peer table is changing in the RIB.
- What is the thing BGP full route routing table ? Even 2914's in Oregon and Tokyo are different.
- Do we need to calculate the common set ?
- Can we evaluate ISP's quality based on the BGP routes it has ? (I guess we can but don't know how)
- Does this tool contribute to debugging the BGP routes ?

Features to be added

- REGEXP in AS_PATH
- match (RE?) on NEXT_HOP

Wrap up

- An open-source tool is newly developed to analyze BGP dump: bgpdump2.
- currently support only RIB dump and only bz2
- PATRICIA-based routing lookup (longest-match) is supported.