



Initial state- r1 and r3 ECMP'd with r2 each sending 600k routes.

Trigger

every 8 secs, do the below in two iterations..

shutdown r1 interface

no shutdown r1 interface

shutdown r2 interface

no shutdown r2 interface

shutdown r1 and r2 interface

no shutdown r1 and r2 interface

OBSERVATION:

Without code changes Buffer data on this simple trigger jumped from 700KB to 78MB

With code changes Buffer data on this simple trigger jumped from 12KB to 8.4MB

#####

Without changes code

#####

-----  
Initial Boot state

-----  
root@r2:~# ip route show | wc -l  
1800004

root@r2:~# vtysh -c "sh ip bgp sum"

IPv4 Unicast Summary:

BGP router identifier 192.168.7.2, local AS number 65002 VRF default vrf-id 0

BGP table version 1200002

RIB entries 1200003, using 110 MiB of memory

Peers 2, using 40 KiB of memory

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
State/PfxRcd	PfxSnt	Desc						
192.168.1.1	4	65001	158	48	1200002	0	0	00:02:13
600000	1 N/A							
192.168.2.3	4	65003	169	48	1200002	0	0	00:02:13
600002	1 N/A							

Total number of neighbors 2

root@r2:~# vtysh -c "sh mem bgpd"

System allocator statistics:

Total heap allocated: 476 MiB

Holding block headers: 18 MiB

Used small blocks: 0 bytes

Used ordinary blocks: 475 MiB

Free small blocks: 1888 bytes

Free ordinary blocks: 508 KiB

Ordinary blocks: 1220

Small blocks: 55

Holding blocks: 4

(see system documentation for 'mallinfo' for meaning)

--- qmem libfrr ---

Type	Current#	Size	Total	Max#	MaxBytes
Buffer	5	24	120	5	120

Buffer data	:	1	4120	4120	174	716880
Host config	:	6	variable	368	6	368
Command Tokens	:	13416	72	966400	13422	966832
Command Token Text	:	9638	variable	337008	9644	337232
Command Token Help	:	9638	variable	232560	9644	232704
Command Argument	:	2	variable	48	8	192
Command Argument Name	:	2340	variable	56320	2344	56416
Lexer token (temporary)	:	0	variable	0	2	64
RCU thread	:	2	136	272	2	272
RCU sequence barrier	:	1	32	40	2	80
FRR POSIX Thread	:	4	variable	352	4	352
POSIX sync primitives	:	4	48	256	4	256
Graph	:	42	8	1008	43	1032
Graph Node	:	15782	32	632672	15784	632752
Hash	:	195	variable	9944	195	9944
Hash Bucket	:	2526	24	60704	2526	60704
Hash Index	:	98	variable	385120	98	385120
ID Number Allocator	:	1	2072	2072	1	2072
ID Number Allocator Name	:	1	50	56	1	56
ID Number Allocator Directory	:	1	1024	1032	1	1032
ID Number Allocator Subdirectory:	:	10	1024	10320	10	10320
ID Number Allocator Page	:	1172	144	178144	1172	178144
Interface	:	7	280	1960	7	1960
Connected	:	6	64	432	6	432
Link List	:	51	40	2120	56	2320
Link Node	:	162	24	4016	291	7112
Temporary memory	:	17	variable	408	28	1024
Bitfield memory	:	2	variable	10256	2	10256
Nexthop	:	2	152	304	2	304
Northbound Node	:	261	1208	315288	261	315288
Northbound Configuration	:	2	16	48	2	48
Northbound Configuration Entry:	:	2	1032	2064	2	2064
Prefix List	:	1	88	88	1	88
Prefix List Str	:	1	9	24	1	24
Prefix List Entry	:	1	136	136	1	136
Prefix List Trie Table	:	2	4096	8208	2	8208
Prefix	:	6	56	336	6	336
Privilege information	:	3	variable	152	3	152
Ring buffer	:	6	variable	1978440	8	2637920
Route map name	:	0	9	0	1	24
Skip List	:	2	56	112	2	112
Skip Node	:	2	160	336	2	336
Skiplist Counters	:	2	68	144	2	144
Socket union	:	4	128	544	4	560
Stream	:	9	variable	436424	12	567576
Stream FIFO	:	6	72	432	8	576
Route table	:	100	56	5664	100	5664
Route node	:	1200008	120	144001552	1200008	144001552
Thread	:	21	168	3528	21	3528
Thread master	:	12	variable	787680	12	787680

-----  
After Trigger  
-----

```
root@r2:~# ip route show | wc -l
1800004
```

```
root@r2:~# vtysh -c "sh ip bgp sum"
```

IPv4 Unicast Summary:

BGP router identifier 192.168.7.2, local AS number 65002 VRF default vrf-id 0  
BGP table version 9200018  
RIB entries 1200003, using 110 MiB of memory  
Peers 2, using 40 KiB of memory

Neighbor State/PfxRcd	V PfxSnt	AS Desc	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
192.168.1.1 600000	4 1 N/A	65001	410	132	9200018	0	0	00:00:43
192.168.2.3 600002	4 1 N/A	65003	425	129	9200018	0	0	00:00:43

Total number of neighbors 2

root@r2:~# vtysh -c "sh mem bgpd"

System allocator statistics:

Total heap allocated: 558 MiB  
Holding block headers: 644 KiB  
Used small blocks: 0 bytes  
Used ordinary blocks: 493 MiB  
Free small blocks: 2160 bytes  
Free ordinary blocks: 65 MiB  
Ordinary blocks: 10764  
Small blocks: 58  
Holding blocks: 1

(see system documentation for 'mallinfo' for meaning)

--- qmem libfrr ---

Type	Current#	Size	Total	Max#	MaxBytes
Buffer	5	24	120	5	120
Buffer data	1	4120	4120	19045	78465416
Host config	6	variable	368	6	368
Command Tokens	13416	72	966400	13422	966832
Command Token Text	9638	variable	337008	9644	337232
Command Token Help	9638	variable	232560	9644	232704
Command Argument	2	variable	48	8	192
Command Argument Name	2340	variable	56320	2344	56416
Lexer token (temporary)	0	variable	0	2	64
RCU thread	2	136	272	2	272
RCU sequence barrier	1	32	40	2	80
FRR POSIX Thread	4	variable	352	4	352
POSIX sync primitives	4	48	256	4	256
Graph	42	8	1008	43	1032
Graph Node	15782	32	632672	15784	632752
Hash	195	variable	9944	197	10040
Hash Bucket	2526	24	60704	2527	60776
Hash Index	98	variable	385120	99	387208
ID Number Allocator	1	2072	2072	1	2072
ID Number Allocator Name	1	50	56	1	56
ID Number Allocator Directory	1	1024	1032	1	1032
ID Number Allocator Subdirectory:	10	1024	10320	10	10320
ID Number Allocator Page	1172	144	178144	1172	178144
ID Number temporary holding pool entry:		0	16	0	600002
14402448					
Interface	7	280	1960	7	1960
Connected	6	64	432	6	432
Link List	51	40	2120	56	2320
Link Node	164	24	4064	293	7160
Temporary memory	17	variable	408	28	1024

```

Bitfield memory      :      2 variable    10256      2    10256
Nexthop              :      2    152        304      2      304
Northbound Node     :    261  1208    315288    261  315288
Northbound Configuration :      2    16         48      2      48
Northbound Configuration Entry:      2  1032    2064      2    2064
Prefix List         :      1    88         88      1      88
Prefix List Str     :      1     9         24      1      24
Prefix List Entry   :      1   136        136      1    136
Prefix List Trie Table :      2  4096    8208      2    8208
Prefix              :      6    56        336      6    336
Privilege information :      3 variable    152      3    152
Ring buffer         :      6 variable  1970264    8  2637920
Route map name      :      0     9           0      1      24
Skip List           :      2    56        112      2    112
Skip Node           :      2   160        336      2    336
Skiplist Counters  :      2    68        144      2    144
Socket union        :      4   128        576      6    896
Stream              :      9 variable  436424    105  6445560
Stream FIFO         :      6    72        432      8    576
Route table         :    100    56    5664    100    5664
Route node          :  1200008    120  144004896  1200008  144007120
Thread              :      23   168        3864    24    4048
Thread master       :      12 variable  787680    12  787680

```

```

#####
With changes code
#####

```

```

-----
Initial Boot state
-----

```

```

root@r2:~# ip route show | wc -l
1800004
root@r2:~# vtysh -c "sh ip bgp sum"

```

```

IPv4 Unicast Summary:
BGP router identifier 192.168.7.2, local AS number 65002 VRF default vrf-id 0
BGP table version 1200002
RIB entries 1200003, using 128 MiB of memory
Peers 2, using 40 KiB of memory

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
State/PfxRcd	PfxSnt	Desc						
192.168.1.1	4	65001	475	371	1200002	0	0	00:18:25
600000	1	N/A						
192.168.2.3	4	65003	490	371	1200002	0	0	00:18:24
600002	1	N/A						

```

Total number of neighbors 2
root@r2:~# vtysh -c "sh memory bgpd"
System allocator statistics:

```

```

Total heap allocated: 485 MiB
Holding block headers: 18 MiB
Used small blocks: 0 bytes
Used ordinary blocks: 485 MiB
Free small blocks: 1952 bytes
Free ordinary blocks: 474 KiB
Ordinary blocks: 5

```

```

Small blocks:          53
Holding blocks:       4
(see system documentation for 'mallinfo' for meaning)
--- qmem libfrr ---

```

Type	Current#	Size	Total	Max#	MaxBytes
Buffer	5	24	120	5	120
Buffer data	1	4120	4120	3	12360
Host config	6	variable	384	6	384
Command Tokens	13416	72	966400	13422	966832
Command Token Text	9638	variable	337008	9644	337232
Command Token Help	9638	variable	232560	9644	232704
Command Argument	2	variable	48	8	192
Command Argument Name	2340	variable	56320	2344	56416
Lexer token (temporary)	0	variable	0	2	64
RCU thread	2	136	272	2	272
RCU sequence barrier	1	32	40	2	80
FRR POSIX Thread	4	variable	352	4	352
POSIX sync primitives	4	48	256	4	256
Graph	42	8	1008	43	1032
Graph Node	15782	32	632672	15784	632752
Hash	195	variable	9944	195	9944
Hash Bucket	2526	24	60704	2526	60704
Hash Index	98	variable	385120	98	385120
ID Number Allocator	1	2072	2072	1	2072
ID Number Allocator Name	1	50	56	1	56
ID Number Allocator Directory	1	1024	1032	1	1032
ID Number Allocator Subdirectory:	10	1024	10320	10	10320
ID Number Allocator Page	1172	144	178144	1172	178144
Interface	7	280	1960	7	1960
Connected	6	64	432	6	432
Link List	51	40	2120	56	2352
Link Node	162	24	4016	291	7224
Temporary memory	17	variable	408	28	1024
Bitfield memory	2	variable	10256	2	10256
Nexthop	2	152	304	2	304
Northbound Node	261	1208	315288	261	315288

```

-----
After Trigger
-----

```

```

root@r2:~# ip route show | wc -l
1800004

```

```

root@r2:~# vtysh -c "sh ip bgp sum"
IPv4 Unicast Summary:
BGP router identifier 192.168.7.2, local AS number 65002 VRF default vrf-id 0
BGP table version 9250018
RIB entries 1200003, using 128 MiB of memory
Peers 2, using 40 KiB of memory

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
State/PfxRcd	PfxSnt	Desc						
192.168.1.1	4	65001	763	490	9250018	0	0	00:01:04
600000	1	N/A						
192.168.2.3	4	65003	777	486	9250018	0	0	00:01:02
600002	1	N/A						

```

Total number of neighbors 2

```

```
root@r2:~# vtysh -c "sh memory bgpd"
```

```
System allocator statistics:
```

```
Total heap allocated: 506 MiB  
Holding block headers: 17 MiB  
Used small blocks: 0 bytes  
Used ordinary blocks: 486 MiB  
Free small blocks: 2096 bytes  
Free ordinary blocks: 20 MiB  
Ordinary blocks: 220382  
Small blocks: 57  
Holding blocks: 2
```

```
(see system documentation for 'mallinfo' for meaning)
```

```
--- qmem libfrr ---
```

Type	: Current#	Size	Total	Max#	MaxBytes
Buffer	: 5	24	120	5	120
Buffer data	: 1	4120	4120	2034	8380080
Host config	: 6	variable	384	6	384
Command Tokens	: 13416	72	966400	13422	966832
Command Token Text	: 9638	variable	337008	9644	337232
Command Token Help	: 9638	variable	232560	9644	232704
Command Argument	: 2	variable	48	8	192
Command Argument Name	: 2340	variable	56320	2344	56416
Lexer token (temporary)	: 0	variable	0	2	64
RCU thread	: 2	136	272	2	272
RCU sequence barrier	: 1	32	40	2	80
FRR POSIX Thread	: 4	variable	352	4	352
POSIX sync primitives	: 4	48	256	4	256
Graph	: 42	8	1008	43	1032
Graph Node	: 15782	32	632672	15784	632752
Hash	: 195	variable	9944	197	10040
Hash Bucket	: 2526	24	60704	2527	60744
Hash Index	: 98	variable	385120	99	387176
ID Number Allocator	: 1	2072	2072	1	2072
ID Number Allocator Name	: 1	50	56	1	56
ID Number Allocator Directory	: 1	1024	1032	1	1032
ID Number Allocator Subdirectory:	12	1024	12384	12	12384
ID Number Allocator Page	: 1422	144	216144	1422	216144
ID Number temporary holding pool entry:		2	16	48	600002
14400512					
Interface	: 7	280	1960	7	1960
Connected	: 6	64	432	6	432
Link List	: 51	40	2120	56	2352
Link Node	: 164	24	4064	293	7224
Temporary memory	: 17	variable	408	28	1024
Bitfield memory	: 2	variable	10256	2	10256
Nexthop	: 2	152	304	2	304
Northbound Node	: 261	1208	315288	261	315288
Northbound Configuration	: 2	16	48	2	48
Northbound Configuration Entry:	2	1032	2064	2	2064
Prefix List	: 1	88	88	1	88
Prefix List Str	: 1	9	24	1	24
Prefix List Entry	: 1	136	136	1	136
Prefix List Trie Table	: 2	4096	8208	2	8208
Prefix	: 6	56	368	6	368
Privilege information	: 3	variable	152	3	152
Ring buffer	: 6	variable	1970264	8	2637920
Route map name	: 0	9	0	1	24
Skip List	: 2	56	112	2	112

Skip Node	:	2	160	336	2	336
Skiplist Counters	:	2	68	144	2	144
Socket union	:	4	128	592	6	912
Stream	:	9	variable	436424	96	5855376
Stream FIFO	:	6	72	432	8	576
Route table	:	100	56	5664	100	5664
Route node	:	1200008	120	144001024	1200008	144001056
Thread	:	22	168	3696	22	3696