

R1  
R2 <  
R3

R1 and R3 are 2-way ECMP'd with R2 each sending same 500k routes.

At 0sec, shut R1 interface  
At 8sec, unshut R1 interface  
At 16sec, shut R2 interface  
At 24sec, unshut R2 interface  
At 32sec, shut R1 interface  
At 40sec, unshut R1 interface  
At 48sec, shut R1 interface  
At 54sec, unshut R2 interface  
At 60sec, shut R1 and R2 interface  
At 70sec, shut R1 and R2 interface

Enabled FPM as well to keep zebra do twice the amount of work (kernel and FPM)

Observation:

- In baseline, with 500K routes from R1 and R3 initially,  
we see 288K streams taking 22MB in zebra which after the above trigger shoots upto  
773K streams with 61MB

- With the fix, and same initial condition as well as after triggers, memory allocated by zebra for  
fifo has reduced a lot..2K streams, 2MB

##### BASELINE #####

Initial state with 500K routes (2way ECMP):

```
root@r2:~# ip route show | wc -l  
1500004
```

```
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 2000006

RIB entries 1000003, using 92 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	359	338	2000006	0	0	00:04:24	500000 500002	N/A
192.168.2.3	4	65003	222	197	2000006	0	0	00:00:10	500002 500002	N/A

Total number of neighbors 2

```
root@r2:~# vtysh -c "sh mem zebra" | grep "Stream\|Total"
Total heap allocated: 335 MiB
Type          : Current# Size   Total   Max# MaxBytes
Stream        :    13 variable 1522544 288362 22310576
```

After Trigger:

```
root@r2:~# ip route show | wc -l
1500004
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 7569058  
RIB entries 1000003, using 92 MiB of memory  
Peers 2, using 40 KiB of memory

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt
192.168.1.1	4	65001	555	624	7569058	0	0	00:00:45	500000 500002	N/A
192.168.2.3	4	65003	415	480	7569058	0	0	00:00:46	500002 500002	N/A

Total number of neighbors 2

```
root@r2:~# vtysh -c "sh mem zebra" | grep "Stream\|Total"
Total heap allocated: 939 MiB
Type          : Current# Size   Total   Max# MaxBytes
Stream        :    13 variable 1522544 773837 61662112
```

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

Initial State with 500K routes(2-way ECMP):

```
root@r2:~# ip route show | wc -l
1500004
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 3028150  
RIB entries 1000003, using 92 MiB of memory

Peers 2, using 40 KiB of memory

Neighbor Desc	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt
192.168.1.1	4	65001	2601	2604	3028150	0	0	00:00:42	500000	500002 N/A
192.168.2.3	4	65003	2776	2725	3028150	0	0	00:00:12	500002	500002 N/A

Total number of neighbors 2

```
root@r2:~# vtysh -c "sh mem zebra" | grep "Stream\|Total"
```

Total heap allocated: 835 MiB

Type	Current#	Size	Total	Max#	MaxBytes
Stream	:	13 variable	1522544	1936	2051072

After Trigger:

```
root@r2:~# ip route show | wc -l  
1500004
```

```
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 8728162

RIB entries 1000003, using 92 MiB of memory

Peers 2, using 40 KiB of memory

Neighbor Desc	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt
192.168.1.1	4	65001	2799	2893	8728162	0	0	00:00:50	500000	500002 N/A
192.168.2.3	4	65003	2973	3053	8728162	0	0	00:00:51	500002	500002 N/A

Total number of neighbors 2

```
root@r2:~# vtysh -c "sh mem zebra" | grep "Stream\|Total"
```

Total heap allocated: 1156 MiB

Type	Current#	Size	Total	Max#	MaxBytes
Stream	:	13 variable	1522544	1936	2051072

Example Debugs:

```
2024/02/23 23:48:29.976698 ZEBRA: [ZQ6GS-2TVH6] Read 99 packets from client: bgp. Current  
ibuf fifo count: 99
```

2024/02/23 23:48:29.977308 ZEBRA: [ZQ6GS-2TVH6] Read 187 packets from client: bgp.  
Current ibuf fifo count: 187

2024/02/23 23:48:29.977380 ZEBRA: [ZQ6GS-2TVH6] Read 32 packets from client: bgp. Current  
ibuf fifo count: 32

2024/02/23 23:48:29.977460 ZEBRA: [ZQ6GS-2TVH6] Read 29 packets from client: bgp. Current  
ibuf fifo count: 61

2024/02/23 23:48:29.977518 ZEBRA: [ZQ6GS-2TVH6] Read 21 packets from client: bgp. Current  
ibuf fifo count: 82

2024/02/23 23:48:29.977563 ZEBRA: [ZQ6GS-2TVH6] Read 16 packets from client: bgp. Current  
ibuf fifo count: 98

2024/02/23 23:48:29.977588 ZEBRA: [ZQ6GS-2TVH6] Read 9 packets from client: bgp. Current  
ibuf fifo count: 107