BGP Graceful Restart-Test Cases for FRR community

Purpose

BGP Graceful Restart functional Testing in FRR-Master

Setup / Topology

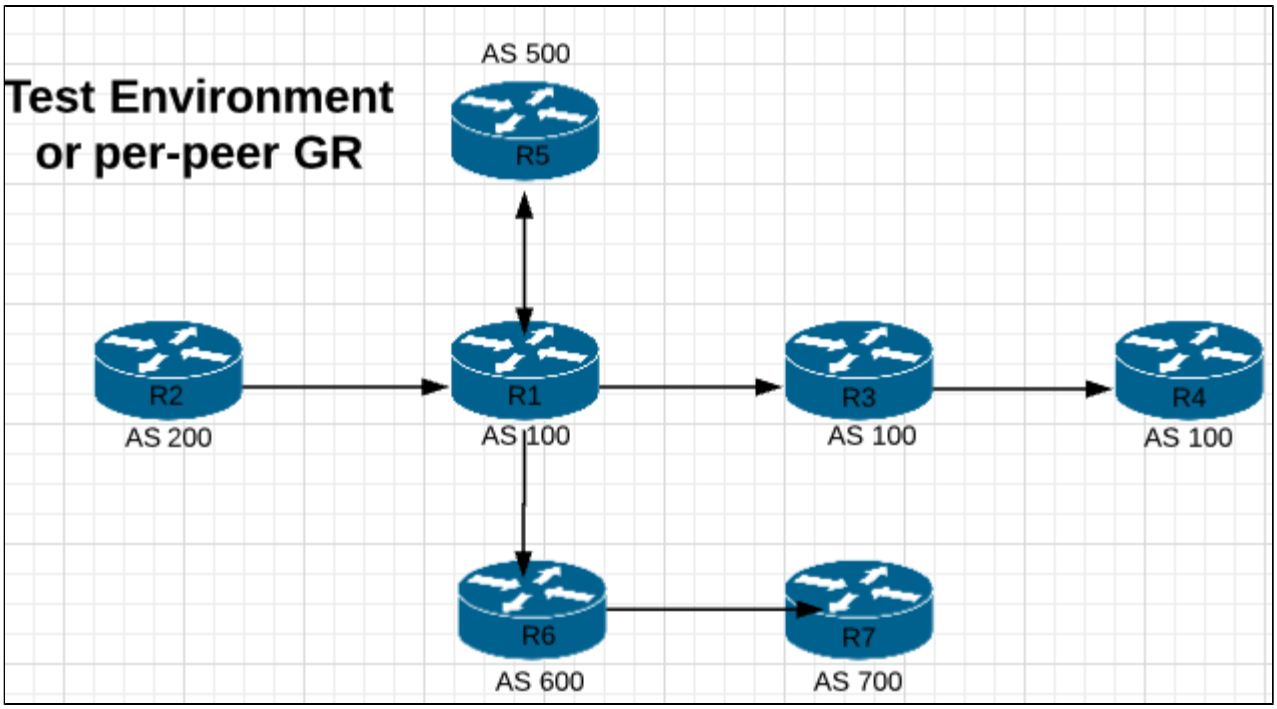

Fig-1

1. **Configuration-1**:

R1 (DUT)	R2	R3	R4	R5	R6	R7
ip addr flush ens192 ip addr flush ens161 ip addr flush ens224 ip addr flush ens256 ifconfig ens192 13.1.1.1 netmask 255.255.255.0 ifconfig ens161 12.1.1.1 netmask 255.255.255.0 ifconfig ens224 15.1.1.1 netmask 255.255.255.0 ifconfig ens256 16.1.1.1 netmask 255.255.255.0 ifconfig ens161 inet6 add 12::1/120 ifconfig ens192 inet6 add 13::1/120 ifconfig ens224 inet6 add 15::1/120 ifconfig ens256 inet6 add 16::1/120 router bgp 100 bgp graceful-restart bgp graceful-restart preserve-fw-state coalesce-time 1000 neighbor 12.1.1.2 remote-as 200 neighbor 13.1.1.3 remote-as 100 neighbor 15.1.1.5 remote-as 500 neighbor 16.1.1.6 remote-as 600 neighbor 12::2 remote-as 200 neighbor 13::3 remote-as 100 neighbor 15::5 remote-as 500 neighbor 16::6 remote-as 600 ! address-family ipv4 unicast network 1.1.1.1/32 network 11.1.1.1/32 network 11.11.1.1/32 network 11.11.11.1/32 network 11.11.11.11/32 neighbor 13.1.1.3 next-hop-self no neighbor 12::2 activate no neighbor 13::3 activate no neighbor 15::5 activate no neighbor 16::6 activate exit-address-family ! address-family ipv6 unicast network 1::1/128 network 11::1/128 network 11::11/128 network 11:11::1/128 network 11:11::11/128 neighbor 12::2 activate neighbor 13::3 activate neighbor 13::3 next-hop-self neighbor 15::5 activate neighbor 16::6 activate exit-address-family !	ip addr flush ens224 ip addr flush ens192 ifconfig ens224 12.1.1.2 netmask 255.255.255.0 ifconfig ens224 inet6 add 12::2/120 ifconfig ens192 125.1.1.2 netmask 255.255.255.0 ifconfig ens192 inet6 add 125::2/120 ! router bgp 200 bgp router-id 20.1.1.1 coalesce-time 1000 neighbor 12.1.1.1 remote-as 100 neighbor 125.1.1.1 remote-as 1000000 neighbor 12::1 remote-as 100 neighbor 125::1 remote-as 1000000 ! address-family ipv4 unicast network 2.2.2.2/32 network 22.2.2.2/32 network 22.22.2.2/32 network 22.22.22.2/32 network 22.22.22.22/32 no neighbor 12::1 activate no neighbor 125::1 activate exit-address-family ! address-family ipv6 unicast network 2::2/128 network 22::2/128 network 22::22/128 network 22:22::2/128 network 22:22::22/128 neighbor 12::1 activate neighbor 125::1 activate exit-address-family !	ip addr flush ens224 ip addr flush ens256 ifconfig ens224 13.1.1.3 netmask 255.255.255.0 ifconfig ens224 inet6 add 13::3/120 ifconfig ens256 34.1.1.3 netmask 255.255.255.0 ifconfig ens256 inet6 add 34::3/120 ! router bgp 100 bgp router-id 30.1.1.1 coalesce-time 1000 neighbor 13.1.1.1 remote-as 100 neighbor 34.1.1.4 remote-as 400 neighbor 13::1 remote-as 100 neighbor 34::4 remote-as 400 ! address-family ipv4 unicast network 3.3.3.3/32 network 33.3.3.3/32 network 33.33.3.3/32 network 33.33.33.3/32 network 33.33.33.33/32 no neighbor 13::1 activate no neighbor 34::4 activate exit-address-family ! address-family ipv6 unicast network 3::3/128 network 33::3/128 network 33::33/128 network 33:33::3/128 network 33:33::33/128 neighbor 13::1 activate neighbor 34::4 activate exit-address-family !	ip addr flush ens256 ifconfig ens256 34.1.1.4 netmask 255.255.255.0 ifconfig ens256 inet6 add 34::4/120 router bgp 400 bgp router-id 40.1.1.1 coalesce-time 1000 neighbor 34.1.1.3 remote-as 100 neighbor 34::3 remote-as 100 ! address-family ipv4 unicast network 4.4.4.4/32 network 44.4.4.4/32 network 44.44.4.4/32 network 44.44.44.4/32 network 44.44.44.44/32 no neighbor 34::3 activate exit-address-family ! address-family ipv6 unicast network 4::4/128 network 44::4/128 network 44::44/128 network 44:44::4/128 network 44:44::44/128 neighbor 34::3 activate exit-address-family !	ip addr flush ens224 ifconfig ens224 15.1.1.5 netmask 255.255.255.0 ifconfig ens224 inet6 add 15::5/120 router bgp 500 bgp router-id 50.1.1.1 coalesce-time 1000 neighbor 15.1.1.1 remote-as 100 neighbor 15::1 remote-as 100 ! address-family ipv4 unicast network 5.5.5.5/32 network 55.5.5.5/32 network 55.55.5.5/32 network 55.55.55.5/32 network 55.55.55.55/32 no neighbor 15::1 activate exit-address-family ! address-family ipv6 unicast network 5::5/128 network 55::5/128 network 55::55/128 network 55:55::5/128 network 55:55::55/128 neighbor 15::1 activate exit-address-family !	ip addr flush ens193 ip addr flush ens256 ifconfig ens193 67.1.1.6 netmask 255.255.255.0 ifconfig ens193 inet6 add 67::6/120 ifconfig ens256 16.1.1.6 netmask 255.255.255.0 ifconfig ens256 inet6 add 16::6/120 ! router bgp 600 bgp router-id 60.1.1.1 coalesce-time 1000 neighbor 16.1.1.1 remote-as 100 neighbor 67.1.1.7 remote-as 700 neighbor 16::1 remote-as 100 neighbor 67::7 remote-as 700 ! address-family ipv4 unicast network 6.6.6.6/32 network 66.6.6.6/32 network 66.66.6.6/32 network 66.66.66.6/32 network 66.66.66.66/32 no neighbor 16::1 activate no neighbor 67::7 activate exit-address-family ! address-family ipv6 unicast network 6::6/128 network 66::6/128 network 66::66/128 network 66:66::6/128 network 66:66::66/128 neighbor 16::1 activate neighbor 67::7 activate exit-address-family !	ip addr flush ens224 ifconfig ens224 67.1.1.7 netmask 255.255.255.0 ifconfig ens224 inet6 add 67::7/120 rrouter bgp 700 bgp router-id 70.1.1.1 coalesce-time 1000 neighbor 67.1.1.6 remote-as 600 neighbor 67::6 remote-as 600 ! address-family ipv4 unicast network 7.7.7.7/32 network 77.7.7.7/32 network 77.77.7.7/32 network 77.77.77.7/32 network 77.77.77.77/32 no neighbor 67::6 activate exit-address-family ! address-family ipv6 unicast network 7::7/128 network 77::7/128 network 77::77/128 network 77:77::7/128 network 77:77::77/128 neighbor 67::6 activate exit-address-family !

Use Cases

The Graceful Restart mechanism, which is specified for BGP in RFC-4724, is used to limit the prefix unavailability due to the BGP process restarting on a router. On a BGP interconnection between two peers, the Graceful Restart capacity declaration is used to keep the packet transfer during the BGP process restart for one of the two routers. The transfer is carried out during a limited time beyond which the routes used are deleted. Once the restart has been performed, the router selects the best routes among the ones sent by its peers and updates its RIB and FIB. This feature is essential in larger topologies, like in ISP environment, where a single routing domain has 100s or maybe more BGP speakers. If the BGP process restarts on a router it would be localised to it's directly connected peers only, other BGP speakers in that domain would not go through control-plane recalculation, provided the restarting router comes back up online within specified timeframe.

In order to localise the impact of a BGP speakers' restart, per peer BGP-GR adds more value to the feature by giving us the flexibility to chose a GR helper or restarting node based on our Network Designs.

Testing Gaps

N/A

Feature Tests/Test Case Details:

Test case ID	1
Priority	P0
Tag	Restarting
Automated	Yes
Objective	Verify that EOR message is sent out only after initial convergence.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Enable debugs for BGP update messages on R1 and log the session. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. Verify in the debugging logs that R1 first receives EOR from all helper routers(R2,R3,R5,R6) and only then send out an EOR to helpers.

Test case ID	2
Priority	P0
Tag	Restarting
Automated	Yes
Objective	Verify whether EOR message is received from all the peers after restart.

Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Enable debugs for BGP update messages on R1 and log the session. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. R1 should mandatorily receive EOR from all helper routers(R2,R3,R5,R6).

Test case ID	3
Priority	P0
Tag	Restarting
Automated	Yes
Objective	Verify that BGP restarting router downloads routes to RIB and send updates to neighbours only after it receives EOR from all of its peers.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Enable debugs for BGP update messages on R1 and log the session. 3. Stop R3 from sending an EOR 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. - 4. R1 should wait to receive an EOR from all helper routers(R2,R3,R5,R6). As it doesn't receive it from R3, restarting node R1 should not install routes into RIB/FIB until selection deferral timer expires.

Test case ID	4
Priority	P1
Tag	Restarting
Automated	Yes
Objective	Verify that the restarting node sets "R" bit while sending the BGP open messages after the node restart, only if GR is enabled.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify that helper node R3 receives R-bit set in OPEN message.

Test case ID	5_1
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify if restarting node sets R bit in BGP open message during normal BGP session flaps as well, when GR restarting mode is enabled.
Precondition and Setup	Fig-1
Steps	1. Change Hello/Dead interval to 5/15 seconds. 2. Bring down the connected interface on restarting node R1. 3. Wait for 16 seconds and bring up the connected interface again.
Expected Result	1. - 2. - 3. Verify that helper node R3 receives R-bit set in OPEN message during the re-negotiation of peering.

Test case ID	5_2
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify if restarting node sets R bit in BGP open message during normal BGP session flaps as well, when GR restarting mode is enabled.
Precondition and Setup	Fig-1
Steps	1. Change Hello/Dead interval to 5/15 seconds. 2. Restart BGPd on router R1.
Expected Result	1. - 2. Verify that helper node R3 receives R-bit set in OPEN message during the re-negotiation of peering.

Test case ID	6_1
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify if restarting node resets R bit in BGP open message during normal BGP session flaps when GR is disabled.

Precondition and Setup	Fig-1
Steps	1. Change Hello/Dead interval to 5/15 seconds on R1 and R3. 2. Change the GR mode on R1 from restarting to disabled. 3. Reset the BGP session between R1 and R3. 4. Bring down the connected interface on restarting node R1. 5. Wait for 16 seconds and bring up the connected interface again.
Expected Result	1. - 2. - 3. Verify the GR modes and hello/dead timers on both routers. 4. - 5. Verify that helper node R3 doesn't receives R-bit in OPEN message from R1.

Test case ID	6_2
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify if restarting node resets R bit in BGP open message during normal BGP session flaps when GR is disabled.
Precondition and Setup	Fig-1
Steps	1. Change Hello/Dead interval to 5/15 seconds on R1 and R3. 2. Change the GR mode on R1 from restarting to disabled. 3. Reset the BGP session between R1 and R3. 4. Restart BGPd on router R1.
Expected Result	1. - 2. - 3. Verify the GR modes and hello/dead timers on both routers. 4. Verify that helper node R3 doesn't receives R-bit in OPEN message from R1 after the restart.

Test case ID	7
Priority	P1
Tag	Restarting
Automated	Yes
Objective	Verify that BGP restarting node deletes all the routes of a neighbour (that were received before the restart) if BGP Graceful capability is not present in BGP Open message from the helper node.
Precondition and Setup	Fig-1
Steps	1. Restarting node R1 is receiving 5 prefixes from helper node R3. 2. Kill BGPd daemon on R1. 3. Change GR mode on R3 from helper to disabled. 4. Remove the advertized 5 networks from R3's BGP config. 5. Bring up the BGPd daemon on R1 again.

Expected Result	1. - 2. Verify if all 5 prefixes received from R3 are marked as stale on R1. 3. - 4. Restarting node R1 should not receive GR capability from R3, as this is now GR disabled. 5. R1 should delete all the stale entries from it's database without waiting for stale path timer's expiry.

Test case ID	8
Priority	P1
Tag	Restarting
Automated	Yes
Objective	Verify that restarting nodes set "F" bit while sending the BGP open messages after it restarts, only when BGP GR is enabled.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify that helper router R3 receives F bit from restarting node R1.

Test case ID	9
Priority	P1
Tag	Restarting
Automated	Yes
Objective	Verify that restarting nodes reset "F" bit while sending the BGP open messages after it's restarts, when BGP GR is **NOT** enabled.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Change the GR mode from restarting to disabled on R1 for peer router R3. 3. Configure the CLI for preserving the forwarding state on R1. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. - 4. Verify that helper router R3 doesn't receive F bit from restarting node R1.

Test case ID	10

Priority	P2
Tag	Restarting
Objective	Test GR scenarios by enabling Graceful Restart for multiple address families.
Automated	Yes
Precondition and Setup	Fig-1
Steps	1. Send IPv6 routes over IPv4 peering and vice versa. 2. Reset BGP session between R1 and R3 3. Kill BGPd daemon on R1. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify that restarting node R1 sends both AFIs in OPEN message. 3. Verify that helper router R3 keeps the stale entries for both AFIs. 4. -

Test case ID	11
Priority	P0
Tag	Restarting
Automated	Yes
Objective	Verify that selection-deferral timer sets the maximum time to avoid deadlock during which the best-path selection process is deferred, after a peer session was restarted.
Precondition and Setup	Fig-1
Steps	1. Change the mode on R1 be a restarting node on global level. 2. Kill BGPd daemon on R1. 3. Stop the helper router R3 from sending EOR message using CLI. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. Verify the selection deferral timer is set to 360 seconds by default. 2. - 3. - 4. Verify that R1 sends EOR to other helper routers only after 360 seconds.

Test case ID	12
Priority	P1
Tag	Restarting
Automated	No
Objective	Verify the behaviour of restarting node when connected with multiple peers consisting helper nodes, restarting nodes, GR disabled nodes.
Precondition and Setup	Fig-1

Steps	1. Change the mode on R1 be a restarting node on global level. 2. Configure R2 as GR helper, R3 as GR restating and R6 as GR disabled nodes. 3. Kill BGPd daemon on R1. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify the GR mode of operation for peer routers based on received capabilities from R2 R3 and R6. 3. Verify that R1 keeps the stale entries for R2 and R3 only, during restart. 4. -

~~**Test case ID**~~	~~13~~
~~**Priority**~~	~~P2~~
~~**Tag**~~	Restarting
~~**Automated**~~	~~Yes~~
~~**Objective**~~	~~Verify that Restarting node doesn't advertise any GR capability to peer router if transitioned to GR disabled.~~
~~**Precondition and Setup**~~	~~Fig-1~~
~~**Steps**~~	1. ~~Change GR mode from restating to disabled on R1.~~ 2. ~~Reset the session between R1 and R3.~~
~~**Expected Result**~~	1. ~~Verify the changed mode on R1.~~ 2. ~~Verify on R3 that after re-negotiation of BGP peering, R1 doesn't advertise any GR capability to R3.~~

Test case ID	14
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify that Restarting node deletes the stale routes as it goes through restart, if BGP process doesn't come up within defined stale-timer.
Precondition and Setup	Fig-1
Steps	1. Change the timer for BGP RIB using "**bgp graceful-restart rib-stale-time XXX**" command on R1. 2. Kill BGPd daemon on R1. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. verify the routing table using "show ip route command". These routes should be kept in the database for configured time stamp only. 3. -

Test case ID	15
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify BGP-GR feature when restarting node is a transit router for it's eBGP peers.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Bring up the BGPd daemon on R1 again.
Expected Result	1. Check if R6 has all the prefixes originated from R2 and R1 in it's Rib/FIB.

Test case ID	16
Priority	P2
Tag	Restarting
Automated	Yes
Objective	Verify BGP-GR feature when restarting node is a transit router for it's iBGP peers.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Bring up the BGPd daemon on R1 again.
Expected Result	1. Check if R3 has all the prefixes originated from R2 and R1 in it's Rib/FIB.

Test case ID	17
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that only GR helper routers keep the stale route entries, not any GR disabled router.
Precondition and Setup	Fig-1
Steps	1. Configure R5 as the BGP GR disabled router. 2. Reset the BGP session between R1<---> R5 and R1<---> R3. 3. Kill BGPd daemon on R1. 4. Bring up the BGPd daemon on R1 again.

Expected Result	1. - 2. Verify that R5 no longer advertises GR capabilities to R1. 3. R5 should immediately remove all the routes learned from R1, without keeping any stale entries. While R3 retains the stale routes. 4. -

Test case ID	18
Priority	P2
Tag	Helper
Automated	Yes
Objective	Verify that GR helper router deletes stale routes received from restarting node, if GR capability is not present in restarting node's OPEN message.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Change the GR mode from restarting to disabled on R1. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. As soon as BGP peering comes up, helper routers R2, R3 and R6 should delete the stale entries immediately after receiving Open message.

Test case ID	19
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that GR routers keeps all the routes received from restarting node if both the routers are configured as GR restarting node.
Precondition and Setup	Fig-1
Steps	1. Change the mode on R3 from helper to GR restarting node. 2. Reset BGP session between R1 and R3. 3. Kill BGPd daemon on R1. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify that both routers are acting as GR restarting nodes. 3. R3 should keep the stale entries and support R1 during it's restart. 4. -

Test case ID	20
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that GR routers delete all the routes received from a node if both the routers are configured as GR helper node
Precondition and Setup	Fig-1
Steps	1. Change GR mode on R1 from restarting to helper. 2. Reset the BGP session between R1 and R3 3. Kill BGPd daemon on R1. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify that both routers are acting as GR helper nodes. 3. R3 should delete the stale entries for R1. 4. -

Test case ID	21
Priority	P2
Tag	Helper
Automated	Yes
Objective	Verify BGP-GR feature when helper node is a transit router for it's eBGP peers.
Precondition and Setup	Fig-1
Steps	1. Change R2 as a GR restarting node globally. 2. Change R1 as helper node only for R2 and disabled for R6. 3. Kill BGPd daemon on R2. 4. Bring up the BGPd daemon on R2 again.
Expected Result	1. - 2. - 3. Check if R6 has all the prefixes originated from R2 and R1 in it's Rib/FIB, pointing next-hop to helper router R1.

Test case ID	22
Priority	P2
Tag	Helper
Automated	Yes
Objective	Verify BGP-GR feature when helper node is a transit router for it's iBGP peers.
Precondition and Setup	Fig-1

Steps	1. Change R2 as a GR restarting node globally. 2. Change R1 as helper node only for R2 and disabled for R3. 3. Kill BGPd daemon on R2. 4. Bring up the BGPd daemon on R2 again.
Expected Result	1. - 2. - 3. Check if R3 has all the prefixes originated from R2 and R1 in it's Rib/FIB, pointing next-hop to helper router R1.

Test case ID	23
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that helper routers are deleting stale routes after stale route timer's expiry. if all the routes are not received from restating node after restart.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Remove 3 prefixes from BGP, originated via network command on R1. 3. Configure R1 to prevent from sending EOR to R3. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. - 4. Verify that R3 waits for an EOR until stale path timer's expiry and remove the stale entries for withdrawn prefixes from R1.

Test case ID	24
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that helper routers are deleting stale routes after restart timer's expiry, in case they don't receive an OPEN message from the restarting node.
Precondition and Setup	Fig-1
Steps	1. Change restart timer on R1 to 150 seconds. 2. Reset the BGP session. 3. Kill BGPd daemon on R1 and wait for 160 seconds. 4. Bring up the BGPd daemon on R1 again.

Expected Result	1. - 2. Verify that R1 advertises GR-restart time as 150 seconds. 3. Verify that R3 keeps stale routes only for 150 seconds and delete after that, as R1 doesn't send an OPEN message again. 4. -

Test case ID	25
Priority	P1
Tag	Helper
Automated	Yes
Objective	Verify that BGP helper routers are deleting the stale routes after receiving the EOR messages from the restarting node.
Precondition and Setup	Fig-1
Steps	1. Change the mode on R1 to be a restarting node on global level. 2. Kill BGPd daemon on R1. 3. Wait for 20 seconds. 4. Bring up the BGPd daemon on R1 again. 5. Run the debugging on R3 and R6 for BGP update messages to see when they receive EOR message.
Expected Result	1. - 2. Stale entries should be present in RIB of helper routers R3 and R6 during these 20 seconds. 3. - 4. Verify if both R3 and R6 deletes stale entries as soon as they receive EOR from restarting node R1.

Test case ID	26
Priority	P2
Tag	Helper
Automated	Yes
Objective	Test GR scenarios on helper router by enabling Graceful Restart for multiple address families.
Precondition and Setup	Fig-1
Steps	1. Send IPv6 routes over IPv4 peering and vice versa. 2. Kill BGPd daemon on R1. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. Verify that helper node R3 receives both AFIs in OPEN message from R1.

Test case ID	27

Priority	P1
Tag	Helper
Automated	No
Objective	Verify the behaviour of helper node when connected with multiple peers consisting helper nodes, restarting nodes, GR disabled nodes.
Precondition and Setup	Fig-1
Steps	1. Change the mode on R1 be a helper node on global level. 2. Configure R2 as GR helper, R3 as GR restating and R6 as GR disabled nodes. 3. Kill BGPd daemon on R2, R3 & R6. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. Verify the GR mode of operation for peer routers based on received capabilities from R2 R3 and R6. 3. Verify that R1 keeps the stale entries for R2 and R3 only, during their restart and R5 doesn't see the routes originated from R6. 4. -

Test case ID	28
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Verify if helper node goes down before restarting node comes up online, helper node sets the R-bit to avoid dead-lock till SDT expiry.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Kill BGPd daemon on R3. 3. Bring up the BGPd daemon on R1 again. 4. Bring up the BGPd daemon on R3 again.
Expected Result	1. - 2. - 3. - 4. Verify that R3 sets R-bit in it's OPEN message to R1.

Test case ID	29
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Change timers on the fly, and verify if it takes immediate effect.

Precondition and Setup	Fig-1
Steps	1. Change GR restating time to 150 on R1 without resetting the session. 2. Kill BGPd daemon on R1. 3. Wait for 120 seconds. 4. Bring up the BGPd daemon on R1 again.
Expected Result	1. Verify the advertised restarting time from R1 on R3 2. Verify that restarting time doesn't change in received capability on R3. 3. Verify that R3 holds stale routes till default restarting time(120 seconds). 4. -

Test case ID	30
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Restarting node removes stale routes from Zebra after receiving an EOR from helper router.
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting and R3 as helper node. 2. Kill BGPd daemon on R1. 3. Withdraw advertised prefixes from R3. 4. Bring up BGPd on R1 again.
Expected Result	1. Verify the GR mode on both the routers. 2. - 3. Verify that R3 is no longer advertising the prefixes to R1 in BGP process. 4. Once the BGP session is re-established, R1 should remove stale entries immediately after receiving an empty EOR from R3.

Test case ID	31_1
Priority	P1
Tag	Chaos
Automated	Yes
Objective	After BGP neighborship is established and GR capability is exchanged, transition restarting router to disabled state.
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting and R3 as helper node. 2. Change the config on R1 from GR restarting to disabled mode. 3. Reset the session from R3. 4. Kill the BGPd daemon on R1.

Expected Result	1. Verify the GR mode on both the routers. 2. - 3. Verify that R1 change GR mode from restarting to disabled mode. 4. R1 should not retain routes in Zebra using "show ip/ipv6 route" command.

Test case ID	31_2
Priority	P1
Tag	Chaos
Automated	Yes
Objective	After BGP neighborship is established and GR capability is exchanged, transition disabled router to restarting state and vice versa.
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR disabled and R3 as helper node. 2. Change the config on R1 from GR disabled to restarting mode. 3. Reset the session from R3. 4. Kill the BGPd on R1.
Expected Result	1. Verify the GR mode on both the routers. 2. - 3. Verify that R1 change GR mode from restarting to disabled mode. 4. R1 should retain routes in Zebra using "show ip/ipv6 route" command.

Test case ID	32
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Restarting node is connected to multiple helper nodes, one of them does't send EOR to restarting router. Verify that only after SDT restarting node send EOR to all helper peers excluding the prefixes originated by faulty router.
Precondition and Setup	Fig-1
Steps	1. Change the mode on R1 be a restarting node on global level. 2. Kill BGPd daemon on R1. 3. Withdraw all the advertised prefixes from R5. 4. Stop the helper router R5 from sending EOR message using CLI. 5. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. Verify that R5 is no more advertising any prefixes in BGP 4. - 5. Verify that all helper routers delete stale entries (originated from R5), as well as those prefixes are no longer appear in BGP routing table.

Test case ID	33
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Helper router receives same prefixes from two different routers (GR-restarting and GR-disabled). Keeps the stale entry only for GR-restarting node(next-hop is correct).
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting and R3 as helper node and R4 as GR disabled mode. 2. Advertise same networks from R1 and R4. 3. Kill BGPd process on R1 and R4. 4. Bring up the BGPd daemon on R1 and R4 again.
Expected Result	1. Verify the GR mode on all the 3 routers. 2. Verify on R3 that it receives same prefixes from 2 next-hops (R1 & R4) 3. Verify that R3 keeps the stale entries only for R1 only. 4. -

Test case ID	34_1
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Restarting node doesn't preserve forwarding state, helper router should not keep the stale entries.
Precondition and Setup	Fig-1
Steps	1. Remove the "preserve-fw-state" command from restarting node R1's config. 2. Reset the session between R1 and R3. 3. Kill the BGPd daemon on restarting node R1.
Expected Result	1. - 2. Verify that R3 doesn't receive F-bit in newly negotionated GR capabilities. 3. Verify that R3 deletes all the stale entries.

Test case ID	34_2
Priority	P1
Tag	Chaos
Automated	Yes

Objective	Restarting node doesn't preserve the forwarding state after restart verify the behaviour on helper node, if it still keeps the stale routes.
Precondition and Setup	Fig-1
Steps	1. Configure restarting router R1 to prevent sending an EOR. 2. Reset the session between R1 and R3. 3. Kill BGPd daemon on R1. 4. Withdraw/delete the prefixes originated from R1. 5. Remove the CLI from R1's config to set the F-bit. 6. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. - 4. - 5. - 6. Verify that after restart when F-bit is not advertised, helper router R3 deletes stale routes(originated from R1) immediately.

Test case ID	35
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Verify that if helper node removes an AFIs from GR capability before the restarting node comes back up online, restating node deletes all stale routes for that AFI.
Precondition and Setup	Fig-1
Steps	1. Kill BGPd daemon on R1. 2. Remove IPv6 AFIs from GR capability on R3. 3. Bring up the BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. Verify that restarting node R1 doesn't keep stale entries for that specific AFI soon after receiving EOR.

Test case ID	36
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Restarting node restarts again for the second time, before BGP route convergence after first restart.
Precondition and Setup	Fig-1

Steps	1. Kill BGPd daemon on R1. 2. Configure R1 to prevent sending EOR. 3. Bring up the BGPd daemon on R1 for 30 seconds. 4. Kill BGPd daemon on R1. 5. Bring up the BGPd daemon on R1.
Expected Result	1. - 2. - 3. - 4. Helper router should reset the received restart timer to hold the stale entries. 5. -

Test case ID	37
Priority	P1
Tag	Chaos
Automated	Yes
Objective	Verify if helper node restarts before sending the EOR message, restarting node doesn't wait until stale path timer expiry to do the best path selection and sends an EOR.
Precondition and Setup	Fig-1
Steps	1. Stop helper router R3 from sending the EOR using the CLI. 2. Reset the session between R1 and R3. 3. Kill BGPd daemon on restating node R1. 4. Bring up the BGPd daemon on R1 again. 5. Kill BGPd daemon on R3. 6. Bring up the BGPd daemon on R3 again.
Expected Result	1. - 2. - 3. - 4. - 5. Verify that R3 sets the R-bit in OPEN message and R1 doesn't wait until stale path timer expiry to do the best path selection and sends an EOR.

Test case ID	38
Priority	P1
Tag	Scale
Automated	No
Objective	Helper router is connected to 320 restarting nodes, all of these 320 restarts at the same time.
Precondition and Setup	Fig-1

Steps	1. Change GR mode from restarting to helper on R1(at global level). 2. Establish 320 neighborships between IXIA and R1 using unique link addresses. 3. Advertise GR-restarting capability from IXIA side for each BGP neighborship with R1. 4. Advertise 10 unique prefixes from IXIA side for each BGP neighborship with R1. 5. Reset BGP session between IXIA and R1 to re-negotiate GR capabilities. 6. Shutdown BGP peering from IXIA side. 7. Bring up BGP peering from IXIA side.
Expected Result	1. - 2. - 3. - 4. - 5. Once the session is reset, please verify that R1 is acting as helper for all neighborships. 6. Verify that R1 keeps stale entries for all (10X320 prefixes) pointing to correct next-hop addresses. 7. Verify that R1 deletes the stale entries after the neighborships are up again.

Test case ID	39
Priority	P1
Tag	Scale
Automated	No
Objective	Restarting router is connected to 320 helper nodes, and goes for a restart.
Precondition and Setup	Fig-1
Steps	1. Configure GR mode as restarting on R1(at global level). 2. Change the RIB stale-timer to 200 Seconds on R1 using "**bgp graceful-restart rib-stale-time**" command 3. Establish 320 neighborships between IXIA and R1 using unique link addresses. 4. Advertise GR-helper capability from IXIA side for each BGP neighborship with R1. 5. Advertise 10 unique prefixes from IXIA side for each BGP neighborship with R1. 6. Reset BGP session between IXIA and R1 to re-negotiate GR capabilities. 7. Kill BGPd daemon on R1. 8. Wait for 201 seconds. 9. Bring up BGPd daemon on R1 again.
Expected Result	1. - 2. - 3. - 4. - 5. - 6. Once the session is reset, please verify that R1 is acting as restarting node for all neighborships. 7. Verify that R1 maintains route entries for all (10X320) prefixes in FIB pointing to correct next-hop addresses. 8. Verify that R1 deletes the route entries after RIB stale-timer expires. 9. Verify that R1 re-learns all the entries from associated neighborship again.

Test case ID	40
Priority	P1
Tag	Scale

Automated	No
Objective	Restating node has 320 helper node and 100 out of those 320 restarted before sending EOR message to restarting node.
Precondition and Setup	Fig-1
Steps	1. Configure GR mode as restarting on R1(at global level). 2. Establish 320 neighborships between IXIA and R1 using unique link addresses. 3. Advertise GR-helper capability from IXIA side for each BGP neighborship with R1. 4. Advertise 10 unique prefixes from IXIA side for each BGP neighborship with R1. 5. Reset BGP session between IXIA and R1 to re-negotiate GR capabilities. 6. Kill BGPd daemon on R1. 7. Uncheck the flag on IXIA to send EOR for 100 peers. 8. Bring up BGPd daemon on R1 again. 9. Restart BGP process on IXIA for those specific 100 neighbors, where we have prevented EOR flag.
Expected Result	1. - 2. - 3. - 4. - 5. Once the session is reset, please verify that R1 is acting as restarting node for all neighborships. 6. Verify that R1 maintains route entries for all (10X320) prefixes in FIB pointing to correct next-hop addresses. 7. - 8. Verify that R1 doesn't receive EOR from those specific 100 peers. 9. Verify that R1 receives R-bit from all specific 100 peers(helpers). After that R1 doesn't wait selection deferral timer to expire for route selection(sends EOR out to all 320 peers).

Test case ID	41
Priority	P1
Tag	Scale
Automated	No
Objective	Verify if restarting router can hold upto 1 Million prefixes(50% IPv4 + 50% IPv6) as stale entries.
Precondition and Setup	Fig-1
Steps	1. Configure GR mode as restarting on R1 for IXIA. 2. Establish BGP neighborships between IXIA and R1 using unique link address. 3. Advertise GR-helper capability from IXIA side for BGP neighborship with R1. 4. Advertise 1M (50% IPv4 + 50% IPv6) unique prefixes from IXIA side for BGP peer R1. 5. Reset BGP session between IXIA and R1 to re-negotiate GR capabilities. 6. Kill BGPd daemon on R1. 7. Bring up BGPd daemon on R1.
Expected Result	1. - 2. - 3. - 4. - 5. Once the session is reset, please verify that R1 is acting as restarting node for this neighborship with IXIA. 6. Verify that R1 maintains route entries for all (1M) prefixes in FIB pointing to correct next-hop address. 7. -

Test case ID	42
Priority	P1
Tag	Scale
Automated	No
Objective	Verify if helper router can hold upto 1 Million prefixes(50% IPv4 + 50% IPv6) as stale entries.
Precondition and Setup	Fig-1
Steps	1. Configure GR mode as Helper on R1 for IXIA. 2. Establish BGP neighborships between IXIA and R1 using unique link address. 3. Advertise GR-restarting capability from IXIA side for BGP neighborship with R1. 4. Advertise 1M (50% IPv4 + 50% IPv6) unique prefixes from IXIA side for BGP peer R1. 5. Reset BGP session between IXIA and R1 to re-negotiate GR capabilities. 6. Kill BGPd daemon on R1. 7. Bring up BGPd daemon on R1.
Expected Result	1. - 2. - 3. - 4. - 5. Once the session is reset, please verify that R1 is acting as helper node for this neighborship with IXIA. 6. Verify that R1 maintains stale entries for all (1M) prefixes in FIB pointing to correct next-hop address. 7. -

==========

Below tests must be executed in sequential order as mentioned below:

==========

Test case ID	43
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Restarting router Peer level Config --->>> Helper router
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting node in global level. 2. Configure R1 as GR helper node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a helper node.

Test case ID	44
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Restarting router Peer level Config --->>> Restarting router
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting node in global level. 2. Configure R1 as GR restarting node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a restarting node.

Test case ID	45
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Restarting router Peer level Config --->>> Disabled router
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR restarting node in global level. 2. Configure R1 as GR disabled node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 does't advertise any GR capabilities.

Test case ID	46
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Restarting router Peer level Config --->>> No config
Precondition and Setup	Fig-1

Steps	1. Configure R1 as GR restarting node in global level. 2. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. Verify on R3 that R1 advertises GR capabilities as a restarting node based on inheritance.

Test case ID	47
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Helper router
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR helper node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a helper node.

Test case ID	48
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Restarting router
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR restarting node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a restarting node.

Test case ID	49
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Disabled router
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR disabled node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 does't advertise any GR capabilities.

Test case ID	50
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> No GR config
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Remove R1's GR config at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. Verify on R3 that R1 advertises GR capabilities as a helper node based on inheritance.

Test case ID	51
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Disabled router Peer level Config --->>> Helper router
Precondition and Setup	Fig-1

Steps	1. Configure R1 as GR disabled node in global level. 2. Configure R1 as GR helper node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a helper node.

Test case ID	52
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Disabled router Peer level Config --->>> Restarting router
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR disabled node in global level. 2. Configure R1 as GR restarting node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a restarting node.

Test case ID	53
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Disabled router Peer level Config --->>> Disabled router
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR disabled node in global level. 2. Configure R1 as GR disabled node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 does't advertise any GR capabilities.

Test case ID	54
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global Config --->>> Disabled router Peer level Config --->>> No config
Precondition and Setup	Fig-1
Steps	1. Configure R1 as GR disabled node in global level. 2. Remove R1's GR config at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 does't advertise any GR capabilities based on inheritance.

Test case ID	55
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Helper router
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR helper node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a helper node.

Test case ID	56
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Restarting router

Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR restarting node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 advertises GR capabilities as a restarting node.

Test case ID	57
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> Disabled router
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Configure R1 as GR disabled node at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. - 3. Verify on R3 that R1 does't advertise any GR capabilities.

Test case ID	58
Priority	P1
Tag	
Automated	Yes
Objective	Verify GR mode of operation when R1 is configured as below: Global level --->>> No GR config Peer level --->>> No GR config
Precondition and Setup	Fig-1
Steps	1. Remove R1's GR config from global level. 2. Remove R1's GR config at per peer level for R3. 3. Reset the BGP session between R1 and R3.
Expected Result	1. - 2. Verify on R3 that R1 advertises GR capabilities as a helper node based on inheritance.

Result : All the above test cases has passed.