RIP Restart Extensions

Saravana Kumar, Janardhan Kulkarni

Why RIP GR ??

Existing Behavior of RIP in a Restart..

the good part

- RIP does not maintain a stateful neighbor relationship.
- Routes have a comfortable aging time of 180s. “No news is good news”.
- On Initialization RIP sends a full table request message to all its peers to receive the updates.

Existing Behavior.. the issues

- Routes are pre-maturely updated to forwarding table and advertised to peers as and when they are learnt.
- Topology changes cannot be handled. If any upstream routes go down during restart the downstream peers will not remove those routes for 180s[1].
- Nexthop flapping due to ECMP route selection.
- Vulnerable to replay attacks.

Note 1 :
-The topology change issue is because the route withdraws for non-existing routes are never propagated to peers. Non-existing because RIP might've had those routes before the restart but after the restart it cannot know for sure if the received withdraws are for the previously active routes or inactive routes.

-Even with a static timer workaround as a replacement for RIP GR the above issues will be there.
-Note : the premature updating of routes WILL NOT cause any route flapping, only the nexthops will flag as based on the order of receiving updates a less preferred nexthop could be chosen, only to be replaced by a more preferred nexthop later on.

Basic Restart Process

Restarting Router - actions in brief

- Request helpers on all its interfaces.
- Wait for route updates from all its peers (both helpers and non-helpers)
- Do not update the routing table or send out any route updates during the waiting period.

Helper Router - actions in brief

- Send full routing table update to the restarting router.
- Notify the restarting router once the full routing update is sent by sending a Update Complete Notification msg.
- Track topology changes.

RIP Extension Message

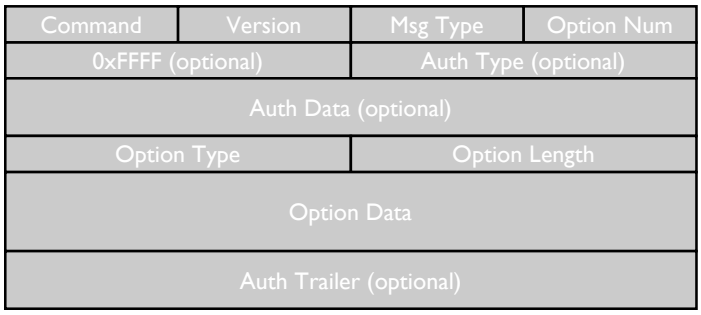

- Flexible for future extensions, ignored by legacy routers.
- New command type : 7
- 254 possible new message types, each having multiple options
- Options in TLV format.

```
 Different Message Types :
  1 - Restart Notification (RN)
   2 - Restart Acknowledgment (RA)
   3 - Update Complete Notification (UCN)
   4 - Update Complete Acknowledgement (UCA)
   5 - Restart Abort Notification (RAN)
   6 - Restart Abort Acknowledgement (RAA)
   7 - Restart Complete Notification (RCN)
   8 - Restart Complete Acknowledgement (RCA)
   9 - Sequence Number Request
  10 - Sequence Number Reply

Different Options :
  1 - Neighbor-List Option
  2 - Delayed-Restart Option
  3 - Sequence Number Option
```

Step 1: Helper Solicitation

- On initialization Restarting Router requests its peers to Help
- Using Restart Notification (RN) and Restart Acknowledgment (RA) messages
- Acknowledging peers are marked as **Restart-Capable** and the others as **Restart-InCapable.**
- Along with the RN, a Full-Table request message (always with sequence number zero[1]) is also sent to request the Restart-Incapable routers to send their route updates

Note 1:

-The full-table request is always sent with a cryptographic sequence number of zero because it achieves two things

-The Restart-Incapable routers will process the full table request only if it has sequence number zero, otherwise as per the MD5 RFC they will drop the packet (this is assuming that the RA is sent with sequence number zero and full table request with sequence number 1)

-The Restart-Capable capable peers (assuming they have Sequence Number Re-acqusition mechanisms enabled) will ignore the packet as it will not pass the sequence number test.

Ensuring Reliability of RN

- On multi-access links send RN multiple times irrespective of the number of RAs received[1].
- On P2P links send RN till a RA is received.
- Each RN should contain a list of routers which have already responded with a RA on that link.
- If the peer sees itself in the RN message, it should ignore the message (prevents redundant processing)

Note 1 : Sending RN multiple times is necessary as there may be many routers on the broadcast link and there is no way for the restarting router to know how many there are.
-We considered having the Helper router include a list of its peers in the RA message but it wont work due to issues like split horizon and selective filtering (one of the routers in the network maybe configured Not to send any updates to the Helper but it may be configured to send updates to the restarting router, so relying on the helper to find out the list of routers on a link cannot work).
-Anyway it is simpler this way and the redundant processing in the helper is eliminated by the self check in the RN message)
-On P2P links also we can consider the same behavior as it can keep things simple

Step 2: Route Synchronization

- Restarting router waits for the route updates to be received from all its peers.
  - For Restart-Capable Helpers wait till a Update Complete Notification message is received from the helper.
  - For Restart-Incapable peers wait for a fixed amount of time (45 to 60s)
- No updating the forwarding table or sending out any periodic or triggered updates during this time.

Step 3: Restart Completion

- The restarting router sends a Restart Complete notification message to all its peers.
- Using the received route updates, add the received routes to the forwarding table and clear any existing stale routes
- Start sending periodic updates to all its peers.
- Helper routers exit the Helper mode only after receiving the Restart complete notification.

Extras

Handling Topology Changes

- If the Helper detects a topology change (restarting router is not an existing peer, local router restarting, local config change, triggered updates required etc). It sends a Restart Abort message to the restarting router[1].
- Restarting router exits the restart mode on receiving an abort msg.
- After an abort the restarting router should propagate received withdraws for non-existing destinations (for 180s)[2].
- For planned restarts, the Restart Notification can be sent with the 'Delayed-Restart' option causing peer to enter the helper mode even before the restart.

Note 1: The Helper should send the restart abort message only when it knows that the restarting router is active. I.e only after it receives the restarting notification message without the 'Delayed Restart' option.

Note 2: This will cause route flapping for a short duration, but better than keeping unreachable routes for 180s

Handling ECMP routes

- Equal cost routes are normally selected on a first come first serve basis
- The set of selected ECMP routes can be different before and after a restart.
- Make the ECMP selection deterministic by preferring routes associated with the lowest neighbor and nexthop addresses.

Security Consideration

- RIP susceptible to replay attacks (specs allow the receipt of messages with sequence number zero from existing neighbors)
- Restart Notification message can be captured and replayed.
- To avoid it, use the optional sequence number re-acquisition (SNR) mechanism[1] after a restart.
- With SNR enabled, messages with sequence number zero SHOULD NOT be accepted from existing peers.

Note 1: The SNR mechanism can be used independent of RIP GR for added security.

SNR mechanism (optional)

- SNR must be supported on all the routers over a link for it to succeed.
- It is performed whenever RIP is enabled on an interface.
- Restarting router sends a sequence number request message (containing a random session id value).
- Peers reply with a sequence number reply message containing the last sequence number received from the requesting router. The highest among the received numbers is picked.
- Using the random session id value prevents replay attacks.

-Restarting router uses this sequence number to send its other 'normal' messages (RN, route updates etc)
-If the peer router has no existing neighbor relationship with the restarting router then it should respond with a sequence number zero.

Issues

Issues

- RIP route updates are not reliable, if any of the updates are lost during the route synchronization phase then it cannot be handled.
- Topology change handling is effective only for route withdraws, for route addition or modification its advantage is negligible and can actually be disruptive to the network.
- Multiple peers undergoing restart is not supported currently (can be done though by relaxing the route update propagation rules for restarting peers)

Q & A