

Restart Signalling for IS-IS

draft-ginsberg-isis-rfc5306bis-01

Les Ginsberg, Cisco Systems

Paul Wells, Cisco Systems

Existing Functionality

Allows a restarting router which maintains forwarding plane across a restart to hitlessly reacquire the LSPDB

```
 0  1  2  3  4  5  6  7
+--+--+--+--+--+--+--+--+
|   Reserved   |SA|RA|RR|
+--+--+--+--+--+--+--+--+
```

RR - Restart Request

RA - Restart Acknowledgement

SA - Suppress adjacency advertisement

Remaining holding time (in seconds)

Restarting Neighbor System ID (for sending RA on LANs)

Sent in Hellos

Neighbor initiates LSPDB sync when receiving RR

What is Lacking...

Current support works well when the control plane restart takes a very short amount of time (less than adjacency holdtime)

Useful for process restarts, redundant control planes lacking local checkpoint capability

Does not support non-redundant control planes which take a significant amount of time to reload (minutes)

Simply extending the holdtime prior to reload leaves neighbor unaware of the impending restart

New Functionality

```
 0  1  2  3  4  5  6  7
+--+--+--+--+--+--+--+
|Reserved|PA|PR|SA|RA|RR|
+--+--+--+--+--+--+--+
```

RR - Restart Request

RA - Restart Acknowledgement

SA - Suppress adjacency advertisement

PR - Restart is planned

PA - Planned restart acknowledgement

Remaining holding time (in seconds)

Restarting Neighbor System ID (for LANs)

Allows neighbor to be aware that a restart is imminent (PR) and to acknowledge (PA).

Neighbor Behavior on Receipt of PR

- Adjacency remains UP – marked in Planned Restart State
- Holdtime is updated (once only)
- PA is sent

Clearing Planned Restart State

- Receipt of RR IIH
- Receipt of IIH w/o Restart TLV or w/o RR or PR
- Holdtime expires

Neighbor Behavior In Planned Restart State

- If topology changes occur, neighbor MAY bring down adjacency (stale forwarding plane)
- If restarting system is DIS, adjacency SHOULD be brought down if topology changes occur
- On P2P circuits flooding of LSPs, xSNPs MAY be suppressed
- If BFD session fails and Control Plane Independent bit is NOT set, BFD failure can be ignored

Changes since V0

- Added Appendix documenting changes from RFC 5306

Next Step

WG adoption