Graceful Restart Extensions for BGP

draft-keyupate-idr-bgp-gr-extensions-00

Keyur Patel, Rex Fernando, John Scudder, Jeff Haas

IETF 80, March 2011, Prague, Czech Republic

Motivation

- Current BGP Graceful Restart mechanism limits its usage to BGP messages other than BGP Notification messages
- Any error within BGP results into generation of BGP Notification messages
  - Without GR support, BGP Notification messages are disruptive. They result in removal of routes from RIB and FIB
- Extending GR support for BGP Notification messages would result in avoiding unwanted routing flaps within routers and across the network

BGP GR Extensions

- New AF based GR capability Extension flag “BGP Graceful Notification (N) bit” defined:

```
0 1  2 3 4 5  6  7
+-+-+-+-+-+-+-+-+
|F|N| Reserved |
+-+-+-+-+-+-+-+-+
```

- Upon a successful GR capability negotiation with N bit set, both the peers transition into GR mode when BGP Notification messages are generated or received
- New BGP Notification type Cease message subcode known as “Hard Reset” defined to explicitly prevent peers from getting into GR mode: Analogous to neighbor shutdown command without GR Support

BGP GR Extensions (Cont'd)

- Both the peers are supposed to preserve their RIB and FIBs and follow normal GR procedures for Notification messages
- Consecutive restarts should not result in purging of RIB and FIB with this extension

Questions?