



# BGP SendHoldTimer

---

draft-ietf-idr-bgp-sendholdtimer

Job Snijders, Ben Cartwright-Cox, Yingzhen Qu

# Problem Statement

---

- If the remote peer signals a TCP Receive Window of 0 for a period of time, local system cannot deliver KEEPALIVE messages, and if the (broken) remote peer doesn't close the connection because of the missing KEEPALIVES, on many systems the connection stays open, leading to operational problems.
- The local system should not solely rely on the remote system for session closure, and it SHOULD close the BGP connection when the remote system is not processing BGP messages for certain time (SendHoldTime).
- Implementation of the SendHoldTimer helps overcome the situation and terminates the BGP session after the local system detects that the remote system is not processing BGP messages for SendHoldTime.

# Summary of Changes

---

The modifications to the BGP FSM are trivial:

The SendHoldTimer ONLY works in the Established state, but it gets started when the FSM transits from the OpenConfirm state to the Established state.

# Update to the FSM - 1

---

- Two new mandatory session attributes:
  - 9) SendHoldTimer
  - 10) SendHoldTime (default value of 8 minutes, may be configured independently up to the implementation)
- Timer Event: SendHoldTimer\_Expires
  - Event XX1: SendHoldTimer\_Expires
  - Definition: An event generated when the SendHoldTimer expires.
  - Status: Mandatory

# Update to the FSM - 2

---

- Start the SendHoldTimer

In “OpenConfirm State”, if the local system receives a KEEPALIVE message, the local system sets the SendHoldTimer to the default for configured value.

- Reset the SendHoldTimer

Each time the local system successfully sends a BGP message, such as KEEPALIVE, UPDATE, and/or NOTIFICATION message, it resets the SendHoldTimer.

# Update to the FSM - 3

---

- If the `SendHoldTimer_Expires` (Event XX1), the local system:
  - sends a NOTIFICATION message with the BGP Error Code "Send Hold Timer Expired",
  - logs an error message in the local system with the BGP Error Code "Send Hold Timer Expired",
  - releases all BGP resources,
  - sets the `ConnectRetryTimer` to zero,
  - drops the TCP connection,
  - increments the `ConnectRetryCounter` by 1,
  - (optionally) performs peer oscillation damping if the `DampPeerOscillations` attribute is set to TRUE, and
  - changes its state to Idle.

# Implementation Status

---

- OpenBGPD
- FRRouting
- BIRD
- neo-bgp
- ... you ... ? :-)

# Next Steps

---

- Feedback and comments are welcome
- Request early IANA code point allocation for BGP Error Code "Send Hold Timer Expired"
- Request WGLC

*THANKS!*