October 2018
IDR Working Group – Interim Meeting

draft-xu-idr-neighbour-autodiscovery-10

Xiaohu Xu, Chao Huang, Guixin Bao (Alibaba)
Ketan Talaulikar, Satya Mohanty (Cisco Systems)
Kunyang Bi, Shunwan Zhuang (Huawei)
Jeff Tantsura, Nikos Triantafillis (Apstra)
Jinghui Liu (Ruijie Networks)
Zhichun Jiang (Tencent)
Shaowen Ma (Juniper Networks)

Update

- Draft was presented at IETF 102 and actively discussed/debated at the mike and on the mailers
- Chairs express the need to have requirements documented before work can be progressed

- Draft now updated to include
  - Applicability
  - Requirements
- Scale improvements introduced
- Adjacency FSM and Procedures specified in further details

Applicability

- Where BGP is used as a hop-by-hop routing protocol e.g. in DCs based on RFC7938 design (or its variations)
- Not applicable for generic BGP deployments

Key Requirements

- Discovery of directly connected BGP neighbors over IPv4 and/or IPv6; agnostic to link layer
- Automatic bootstrapping of BGP TCP session by learning peering addresses; no change to BGP TCP FSM or BGP routing procedures
- Automatically setup reachability to neighbors peering address (e.g. loopback) when necessary
- Exchange link attributes & parameters between neighbors for topology discovery/advertisement and validation of link & peering policy
- Leverage existing fast-detection mechanisms (e.g. BFD, FEF, etc.)
- Security

Simplicity = do only what is needed & missing; leverage what exists

Updated Hello Message

- Periodic Hello : skip most TLVs (except security) for better scalability and simplified processing
- State Change Hello : carry TLVs and used only during initial discovery and subsequently when there is any change in state

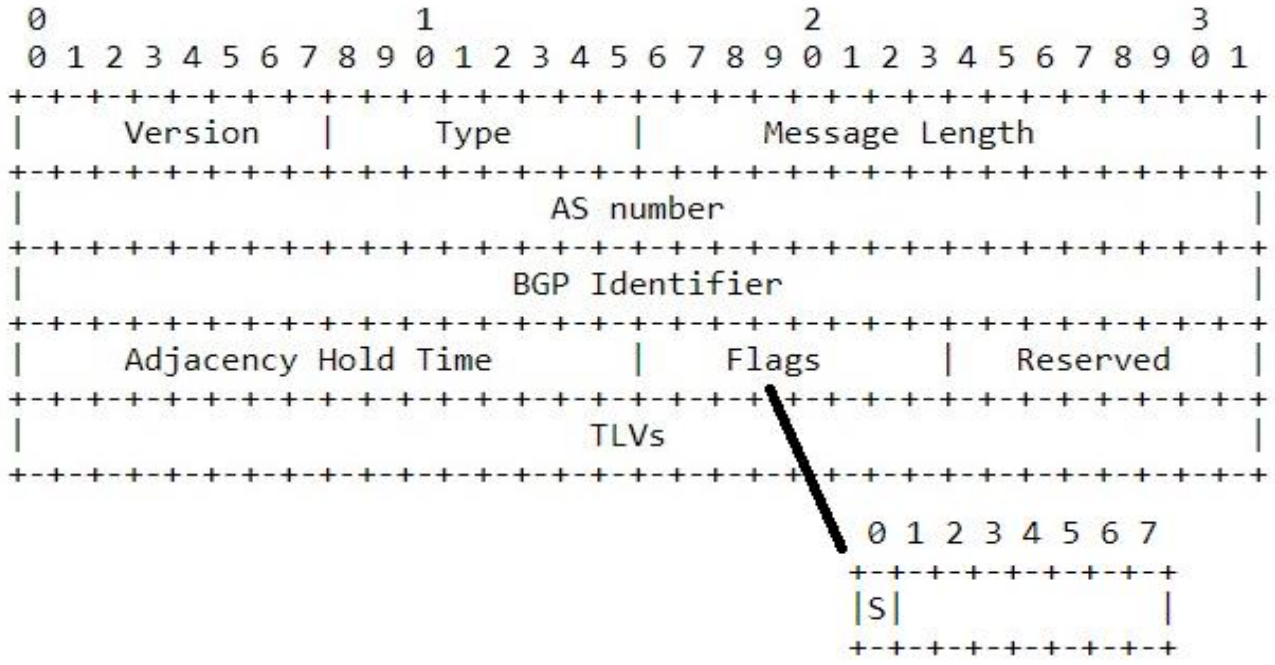

Updated Adjacency State Machine

- Down
  - Transient terminal state after which adjacency is deleted
- Initial
  - Transient initial state when adjacency is created for new neighbor
- 1-way
  - When router detects a neighbor and includes it in its own hello message, but the neighbor has not yet detected it
- 2-way
  - When the router and its neighbor have both detected each other's hello messages
- Adj-Reject
  - When the router rejects its neighbour due to failure of some validation checks based on local policy
- Adj-OK
  - When the router has accepted its neighbour after validation against local policy
- Accepted
  - When both router and its neighbour have accepted each other; and the BGP TCP Peering is initiated

Procedures

- Draft specifies the Adjacency FSM state transitions and events/triggers in detail
- Also specifies the procedures for handling of hello messages (both state change and periodic)
- Includes clarifications and details on interactions with the BGP TCP peering session
- Updated text to describe the handling of the Adjacency Route (i.e. the locally installed route for reachability to neighbor’s loopback)

Next Steps …

- Active interest for implementation
- Solicit WG review and comments/inputs/feedback for the updated revision
- Re-do WG adoption call ?