IPv6 BGP Identifier Capability for BGP-4

draft-fan-idr-ipv6-bgp-id-00

Peng Fan, Zhenqiang Li

Motivation

- The Identifier of a BGP speaker was specified as a valid IPv4 host address assigned to the BGP speaker in RFC4271; RFC6286 relaxed the definition to be a 4-octet, unsigned, non-zero and AS-wide unique integer.
- BGP Identifiers in a real network are often configured in the form of an IPv4 address to help network maintenance.
- The 4-octec integer Identifier in IPv6-only network requires additional configuration and planning consideration to guarantee uniqueness within the AS.
- This document extends BGP to allow a BGP Identifier to be a valid IPv6 global unicast address assigned to the BGP speaker.

Protocol Extension

- A new BGP capability code, “IPv6 BGP Identifier Capability”, is defined to indicate the support for IPv6 address as a BGP Identifier.
- OPEN message: the BGP Identifier field is set to zero, indicating the actual BGP Identifier is in the Capability Optional Parameter.
- IPv6 BGP Identifier Capability: The Capability Length field of the is set to 16, and the Capability value field is set to an IPv6 global unicast address.

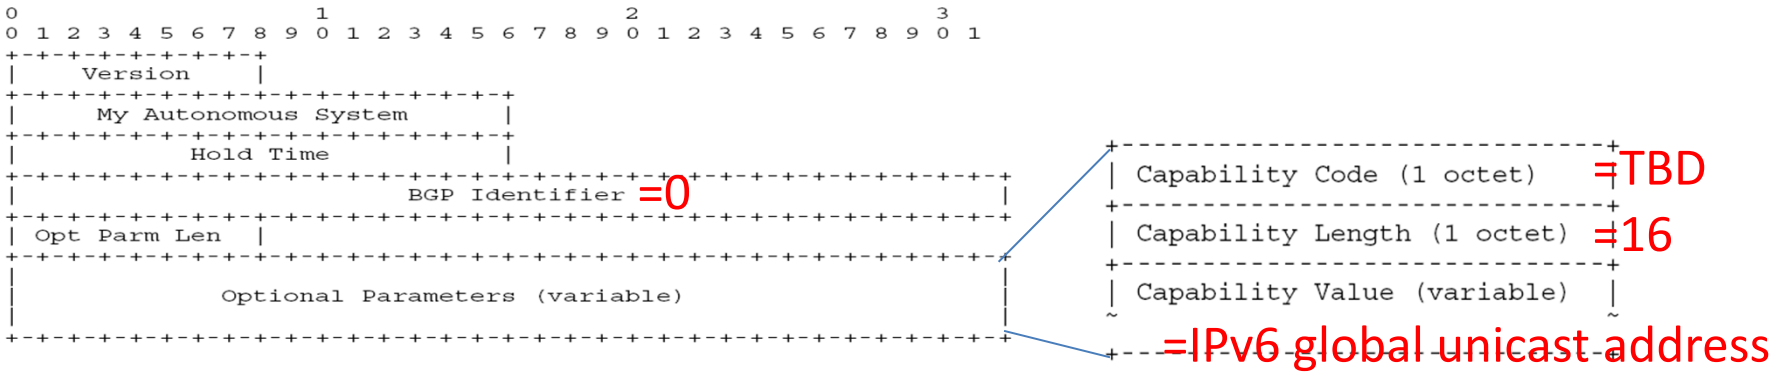

- AGGREGATOR attribute: set accordingly; the BGP Identifier carried in the attribute is encoded as a 16-octet entity.

Operation

- Processing received OPEN messages:
  - If the BGP Identifier field is not zero: process in the way of the message that does not contain IPv6 BGP Identifier, and any IPv6 BGP Identifier Capability is ignored.
  - If the BGP Identifier field is zero, then check if any IPv6 BGP Identifier Capability is carried. If there is no IPv6 BGP ID Capability, or the capability value is not a valid IPv6 global unicast address, then a Notification message is generated, with Error Code set to 2 (OPEN Message Error) and Error subcode set to 3 (Bad BGP Identifier).
- Connection collision detection:
  - The BGP Identifiers of the peers involved in the collision are compared and only the connection initiated by the BGP speaker with the higher-valued BGP Identifier is retained.
- Route selection decision:
  - If a route is advertised by an IPv4 BGP speaker and an IPv6 BGP speaker respectively, then the route advertised by the IPv6 BGP speaker is selected.
  - If a route is advertised by two IPv6 BGP speakers respectively, then their IPv6 BGP IDs are compared, and the route advertised by the BGP speaker with the lower-valued BGP Identifier is selected.

Transition

- A BGP speaker supporting the IPv6 BGP Identifier must set a 128-bit Identifier.
- If the speaker is not aware of the capability of its peers, then a 32-bit Identifier is assigned for backup purpose.
- The speaker tries the 128-bit identifier first; if the peer does not support the new 128-bit ID capability, then a “bad bgp identifier” error message is generated (Identifier field of OPEN message received is zero).
- The speaker initiates a second connection using 32-bit identifier in the old way, and the connection falls back using 32-bit identifier.

Consideration

- Pretty much discussion on the list
- Advantages of extending the length of identifier
  - Help identify the location, e.g. for diagnosis and troubleshooting
  - Can be autoconfigured
- Other suggestion on the list:
  - Use other separated mapping system, e.g. DNS, text, v4-v6 addr mapping. (Extra record keeping adds more work for OAM)

- Do we update ID to convey more information or just keep the 32 random bits?