

# Revised Error Handling for BGP Messages


<https://datatracker.ietf.org/doc/draft-wang-idr-bgp-error-enhance/>

H. Wang, M. Shen, J. Dong @Huawei

IETF 112

Nov. 2021

# BGP Error Handling Scenario



- R1~R3 are all connect to the RR
- R3 advertises an update message to the RR. When the RR found that the packet has an error, the RR will reset the connection.
- Later the session will re-established, and R3 will continue advertise that message. The session will goes down and up again and again.
- This may be occurred for each session. Such procedure do not help to solve the problem.
- Why this occurred? Perhaps R3 exist some software bugs or inconsistent understanding of standards.
- This procedure is defined in 4271 but is too strict, and it was revised by RFC7606, But there's still some not clear description of the situation.

# Proposal

- Our proposal is to reduce the impact of these errors on BGP's existing services, as this does not help resolve the problem. We want to describe some scenarios where the current definition is not clear enough and some are too strict.
- Following are some Suggestions:
- MP\_REACH\_NLRI/MP\_UNREACH\_NLRI:
  - Prefix Length Error. Recommend to discard the error Prefix and stop the current message's parsing. We may keep the correct Prefixes parsed before.
  - Appear more than once. Recommend to process with the last ( or first) one and ignore the others
  - Next-hop Length Error. Recommend to ignore current message
- Validation of the address
  - Address is indicate the BGP ID, according to RFC6286 only zero is invalid. Attribute AGGREGATE, AS4\_AGGREGATOR, ORIGINATOR\_ID, Cluster-List are the case.
  - Address is indicate the IP, mainly the NEXTHOP attribute. In most cases, we can assume that a valid unicast address can be used. So the Multicast Address(Class D), Reserved Address(Class E) and Broadcast Address(All One), and invalid address( such as 0.1.1.1) shouldn't be used. Also the special address 0.0.0.0 shouldn't be used either. But the use of some other special address such as localhost address is a service behavior, we may not focus on it here.
- Prefix\_SID: Used for create SR tunnel. If the attribute is incorrect, it means that we cannot use it to create SR tunnel, so we may withdraw the route.

# Received Comments

- Thanks for Robert's comments. Some comments have been incorporated in the previous description.
- Two general comments for discussion:
  - Comment #1: Issue a RFC7606bis or separate draft?
    - Answer: We may first discuss which problem scenarios need to be modified and supplemented
  - Comment #2 : BGP-4 protocol running on TCP is not bullet proof when it comes to handling bad implementations or malicious protocol attacks. Consider using new transport which no longer runs on TCP and essentially not only treats all SAFIs as fully independent streams but also cut's interdependencies between all NLRI even with given SAFI.
    - Answer: We suggest that we first consider polishing the existing BGP mechanism and also consider improving the BGP protocol mechanism

# Next steps

- Welcome more comments and discussion
- Revise the draft accordingly

Thank you!