

draft-jaufeerally-bgp-lg-cap

Rayhaan Jaufeerally <ietf@rayhaan.ch>

2021-07-27

GROW WG | IETF 111

Motivation

- When turning up a new BGP session, how do you know if your routes are accepted?
- Anycast makes this harder, since a simple ping can go to another PoP.
- What if you could programatically check that the new PoP is online and healthy?

Current solutions & limitations

- Looking glasses are useful for humans:
 - You can ping and traceroute,
 - You can lookup some details about the routes the peer sees.
- A human operator can manually reason about what reachability should look like.
- This process is manual, and one off.

Administrative message (S)AFI

- Re-use multiprotocol extensions to carry an administrative message in MP_REACH and MP_UNREACH path attributes.
- TLV payload in the administrative payload to let it be used for other purposes in the future.
- One type of message can be the looking glass, which carries:
 - ASN,
 - Looking glass URL conforming to RFC8522.

Using the path attribute

- Support to peers is advertised with standard multiprotocol capability in the OPEN message,
- Direct peers can send the admin message with looking glass address,
- When a direct peer forwards an announcement to another peer, which it knows the LG address for, it can send that peers LG address too in the admin message,

Example

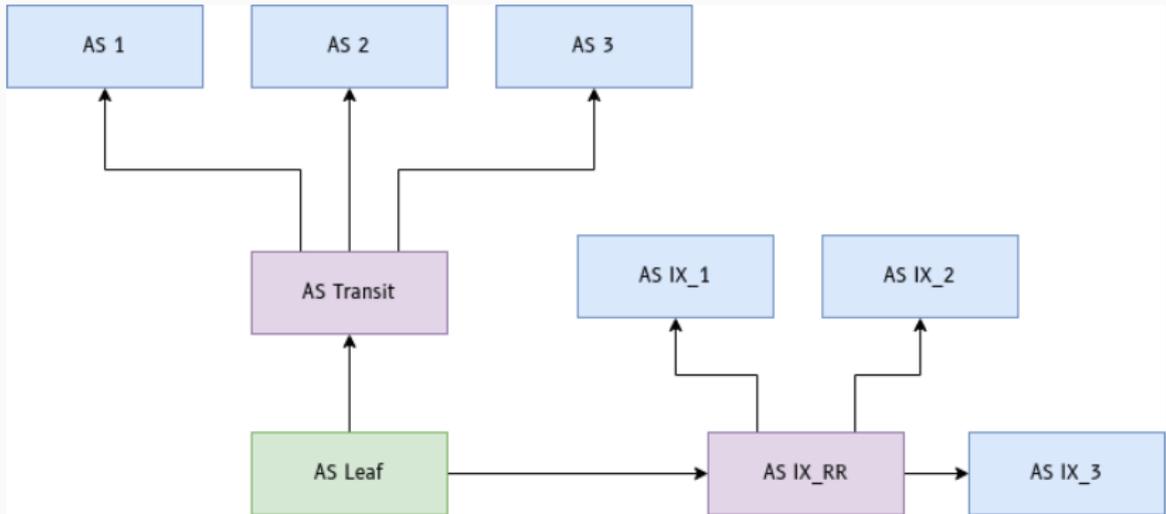


Figure 1: An example topology

- RFC8522 does not really specify a defined machine readable format, it outputs freeform text,
- Maybe we should use DNS instead for this, discussion on the mailing list about .bgp TLD,
- Unclear about what the looking glass will show with respect to the RIB of the ASBR,

End

Discussion