

BGP Extensions for Network Resource Partition

draft-li-idr-bgp-nrp-01

Zhenqiang Li (China Mobile)
Ruiqian Hu(China Mobile)


IETF 127, IDR Working Group · Interim Meeting 2026

Problem

- Existing work *draft-ietf-idr-sr-policy-nrp* extends the BGP SR Policy protocol [RFC 9830] to deliver the NRP ID, and binds each NRP to one SR Policy on a one-to-one basis. See: <https://datatracker.ietf.org/doc/draft-ietf-idr-sr-policy-nrp/>
- As the number of NRPs grows, SR Policies grow linearly, each with its own BFD/SBFD monitoring session.
- This brings heavy configuration and operational overhead and weakens scalability.

Solution

- Extending BGP, rather than BGP SR Policy, to advertise NRP ID information in routing updates between network nodes.
- This solution decouples SR Policy from NRP, while coupling NRP with the services that demand the corresponding NRP for forwarding resources.
 - *Our key idea is that it is the service—not the SR Policy—that requires NRP.*
- Our solution allows multiple NRPs to share a common SR Policy path, thereby avoiding linear growth of SR Policies. This eliminates heavy OPEX burdens related to SR Policies, removes duplicated SR Policies, and avoids the need for associated BFD/SBFD sessions.



Benefits

Higher scalability for slice deployments

- Multiple NRPs share one SR Policy path; policies stop growing linearly with the slice count.
- Duplicated policies and their per-policy BFD/SBFD sessions are eliminated.

Lower operational complexity; covers both TE and BE

- NRP information is exchanged with conventional BGP route propagation; intermediate nodes perform no NRP-specific processing.
- Same mechanism applies to SR TE and SR Best Effort; no new security considerations are introduced (RFC 4360, RFC 9543, RFC 9012 apply).

One shared SR Policy path can serve multiple slices; NRP isolation and QoS are enforced per service at the resource layer.

Questions

- Do you also see the one-to-one NRP-to-SR-Policy binding as a scalability and operational gap as the number of slices grows?
- Does this approach create any issue with existing TEAS, SPRING, or IDR technology (SR Policy, extended communities, BGP route propagation)?
- Other feedback