Scalability Considerations for Network Resource Partition (NRP)

draft-ietf-teas-nrp-scalability-05

Jie Dong, Zhenbin Li @Huawei
Liyan Gong @China Mobile
Guangming Yang @China Telecom
Gyan Mishra @Verizon

TEAS WG   IETF 120 Meeting   July 2024

Recap of Network Slice and NRP

- The concept and general framework of “IETF network slice ” are described in RFC 9543
  - IETF network slice services can be realized by mapping network slice connectivity constructs to NRPs in the underlay network
- An NRP consists of a set of network resources, and is associated with a (filtered) topology
  - Can be used to support one or a group of network slice services
- The scalability of NRP is important for the widely deployment of IETF network slices
  - This document provides the design principles and considerations about scalability in both the control plane and data plane
  - Some scalability optimization mechanisms are suggested

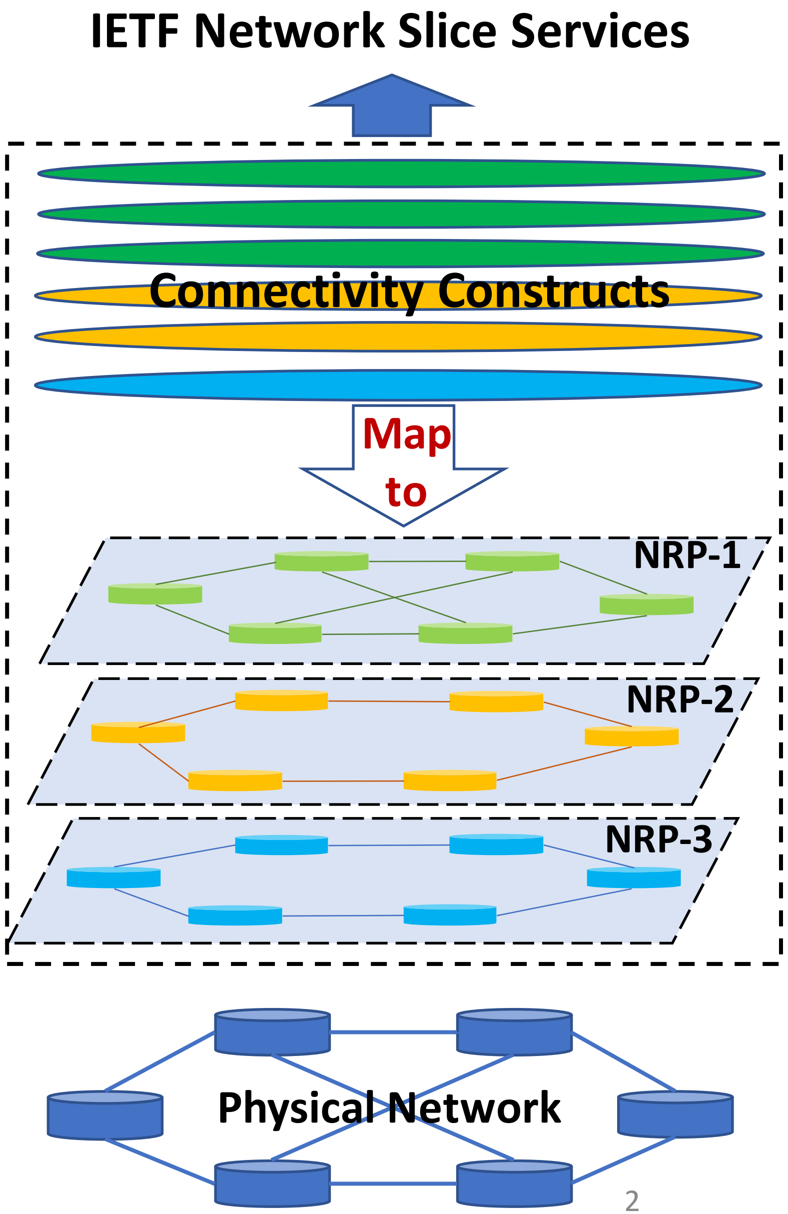

Updates in -04/-05 Version

- The NRP scalability design principles in section 3 are polished according to the feedbacks received
- Some examples of network slice realizations and their scalability implications are discussed in the appendix section
- A series of editorial changes and reference update

Next Steps

- Reach consensus on the scalability design principles in this document
  - Other considerations on scalable NRP design could be added to this document
- Use this document as the guidelines for the data plane and control plane extensions of scalable NRP solutions
- Move this document towards WG LC

Scalability Design Principles (1)

1. A filtered topology is a subset of the underlying physical topology
2. It is not envisaged that there would be many filtered topologies active, so running SPF per filtered topology is not a high burden
3. Multiple NRPs can run on a single filtered topology

Topology sharing among NRPs to reduce SPF computation overhead

4. Three separate things need to be identified by information carried within a packet
   - Forwarding path (e.g. the next-hop)
   - NRP
   - Topology (i.e., filtered topology)

   How this information is encoded (using separate fields, the same field or overloading existing fields) forms part of the solution.
5. NRP IDs should have domain-wide scope, and must be unique within a filtered topology

Identifiers in data packets

6. Mechanisms such as local configuration or southbound protocols from a controller are used to set up resources and packet treatments of NRPs on the involved network nodes

NRP instantiation

7. The path computation or selection performed by a traffic engineering process may take the characteristics of the filtered topology and the attributes of the NRP into consideration

NRP-specific path computation

8. The selected path and the associated NRP can be indicated in the packets using existing or new mechanisms

NRP Identifier in data packet

9. The components or mechanisms that are responsible for deciding what path to select for specific network slice service packet, how to mark the packets to follow the selected path and what resource treatment identifier (NRP ID) to apply to packets are responsible for ensuring consistency for the solution to work

Steering services to NRP-specific paths

Scalability Design Principles (2)

- Different operators can choose to deploy network slices at different scales, we should not attempt to constrain operators in their deployment choices
- The routing protocols (IGP or BGP) does not have to be involved in any of these points, but when they need to, it is important to isolate information of network slices and NRPs from the existing routing information, so that there is no impact on scaling or stability. Furthermore, the overhead of SPF computation should not be impacted by the increasing number of network slices and NRPs.
- There is always a trade-off between optimal solutions and scalable solutions.
  - A solution that can be deployed at large scale needs to be provided.
    - Some extensions to the data/control/management plane may be needed to achieve this result
    - The scalable solution might not be optimal everywhere
  - A solution may be optimal for specific environments, but
    - Might not work in some environments
    - May have scalability issues in some other environments.
- Both of these approaches should be allowed, but we need to be clear of the costs and benefits
  - There is consensus on the efforts for better NRP scalability, the cost of necessary protocol extensions is acknowledged
  - Solutions which are suitable for specific environments are also allowed, as long as their limitations are documented
- In particular, we should be open to the use of approaches that do not require control plane or data plane extensions and that can be applied to deployments with limited scope, such as:
  - Resource-aware SIDs
  - VPNs

Network Slice Realization Examples and Scalability Implications

- VPNs with default NRP
  - Provides connectivity and simple SLA required by some network slice services
  - Scalability lies in the routing protocols used for the distribution of VPN routes
- Multiple routing instances for NRPs
  - Separate routing instance and separate set of links for each NRP
  - Scalability lies in the number of routing instances needed on each node
- Resource-aware segment routing based NRPs
  - Separate multi-topology or Flex-Algo with subset of network link resources for each NRP
  - Resource-aware segments are used to identify/distinguish the NRP resources in data plane
  - Scalability lies in the SPF computation overhead per NRP, and the amount of resource-aware SIDs that needs to be distributed and installed for each NRP
- MPLS-TE virtual networks
  - Virtual networks built with a set of resource-reserved TE-LSPs can be mapped to an NRP
  - Scalability lies in the overhead of LSP planning and per-LSP state maintenance for each NRP

Thank You