Scalability Considerations for Network Resource Partition (NRP)

draft-ietf-teas-nrp-scalability-03

Jie Dong, Zhenbin Li @Huawei

Liyan Gong, Fengwei Qin @China Mobile

Guangming Yang @China Telecom

James Guichard @Futurewei

Gyan Mishra @Verizon

Tarek Saad @Cisco

Vishnu Beeram @Juniper

TEAS WG IETF 118 Meeting Nov. 2023

Recap of Network Slice and NRP

- The concept and general framework of “IETF network slice ” are described in draft-ietf-teas-ietf-network-slices
  - IETF network slice services can be realized by mapping network slice connectivity constructs to NRPs in the underlay
- An NRP consists of a set of dedicated or shared 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 widely deployment of IETF network slices
  - This document provides scalability considerations of NRP in both the control plane and data plane
  - Some optimization mechanisms are also investigated

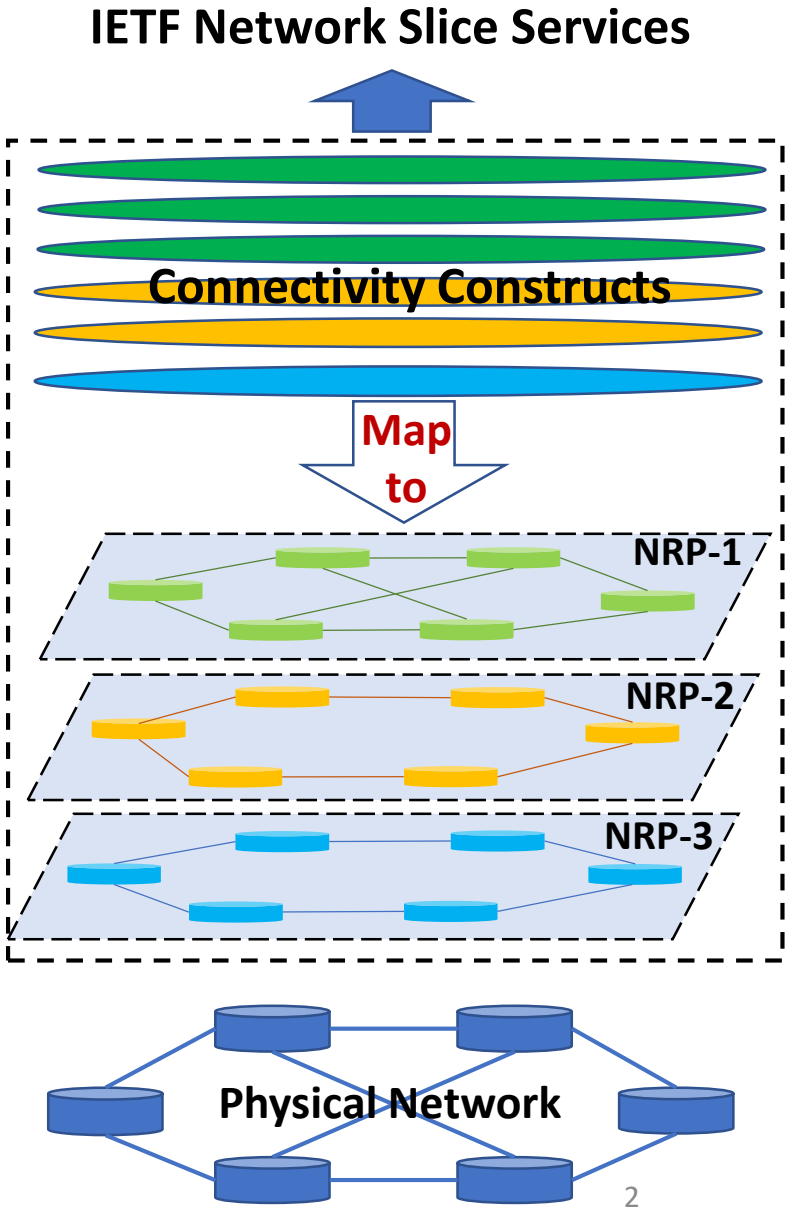

NRP Scalability Considerations

- The scalability of distributed control plane is related to:
  - The number of control protocol instances
  - The number of control protocol sessions
  - The number of control messages advertised by each node
  - The amount of information carried in each message
  - The number of computations (e.g. SPF computation) executed by each node
- The scalability of data plane is related to:
  - The pros and cons of reusing existing IDs/fields VS introducing new IDs/fields for NRP identification
  - The scalability implication of different encapsulation options

Updates in -03 Version

- The output of previous scalability discussion are listed in the draft as a new section: Scalability Design Principles
    - The text still needs further review and feedback from the WG
- Replaced the term “IETF Network Slice” with “RFC XXXX network slice”
    - Align with the decision made on draft-ietf-teas-ietf-network-slices
- A series of editorial changes

Scalability Design Principles

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
   - Running SPF per filtered topology is not a high burden
3. Multiple NRPs can run on a single filtered topology
   - Meaning that the NRPs can be associated with the same filtered topology and use that topology's SPF computation results.
4. Three separate things need to be identified by information carried within a packet:

   a. Path

   b. NRP

   c. Topology (i.e., filtered topology)

   How this information is encoded (separate fields, same field, overloading existing fields) forms part of the solution work.
5. NRP IDs should have domain-wide scope, and must be unique within a filtered topology
6. Configuration mechanisms are used to set up packet/resource treatments on nodes
7. Configuration mechanisms (such as southbound protocols from a controller) are used to install bindings on network nodes between domain-wide resource treatment identifiers (NRP IDs) and configured packet treatment

Scalability Design Principles (2)

8. The path selection performed by or within a traffic engineering process, within or external to the head end node, (in particular the topology selection and path computation within that topology) may consider the characteristics of the filtered topology and the attributes of the NRP
   - But is agnostic to the resource treatment that the packets will receive within the network.
   - Ensuring that the selected components of the path that are configured are capable of supporting the resource treatments identified by the NRP ID, is a separate matter.
9. The selected path is indicated in the packets using existing or new mechanisms.
   - Whether that is SR-Policy (for some variety of SR), flex-algo (for whatever flex-algo expression you like), is something out of scope for now, but it will obviously form part of the full set of solution specifications.
10. The components or mechanisms that are responsible for deciding what path to select, for deciding how to mark the packets to follow the selected path, and for determining what resource treatment identifier (NRP ID) to apply to packets are also responsible for ensuring sufficient consistency so that the whole solution works.

Scalability Design Principles (3)

- Different operators can choose to deploy things at different scales, and while we may have opinions about what scales are sensible / workable / desirable, we do not have to get WG agreement on that aspect.
- Routing protocols (IGP or BGP) do not need to be involved in any of these points, and it is important to isolate them from these aspects in order that there is no impact on scaling or stability. Furthermore, the complexity of SPF in the control plan is unaffected by this.
- There is always a trade-off between optimal solutions and scalable solutions.
  - We need to achieve a scalable solution that can be deployed in all circumstances.
    - We may need some extensions to the data/control/management plane to achieve this result
    - The scalable solution might not be optimal everywhere
  - We must understand that optimal solutions are good for specific environments, but
    - Might not work in other environments
    - May have scalability issues.
- We should allow for both of these approaches, but we need to be clear of the costs and benefits
- In particular, we should be open to the use of approaches that do not require control plane extensions and that can be applied to deployments with limited scope, such as:
  - Resource-aware SIDs
  - VPNs

How to Coordinate with Other WGs?

- There are on-going work in other WGs on protocol extensions/application for network slice realization
- Some solutions are targeted at specific network scenarios and limited scales, and may not be applicable to scenarios with large number of slices/NRPs
  - Suggests to make it clear in the scalability considerations of those documents
- Some protocols may be more sensitive than others to scalability issues
  - So far the major concern is about the scalability of IGP
  - Suggest to make it clear in the scalability considerations of those documents
- Some other solutions aims to provide highly scalable solutions for wide network slice deployment (e.g. 1000 or more NRPs)
  - The design principles in this document can be used as guidance for the solution design

Suggestions & Next Steps

- Initiate further discussion on NRP scalability in the WG
- Reach consensus on the scalability design principles
- Add an appendix with several examples showing different possible NRP realizations and their scaling properties
- Coordinate with relevant WGs on the solutions which are targeted at better scalability

Thank You