

# Scalability Considerations for Network Resource Partition (NRP)

*draft-ietf-teas-nrp-scalability-02*  
Jie Dong, Zhenbin Li @Huawei

Liyan Gong, Fengwei Qin @China Mobile

Guangming Yang @China Telecom

James Guichard @Futurewei

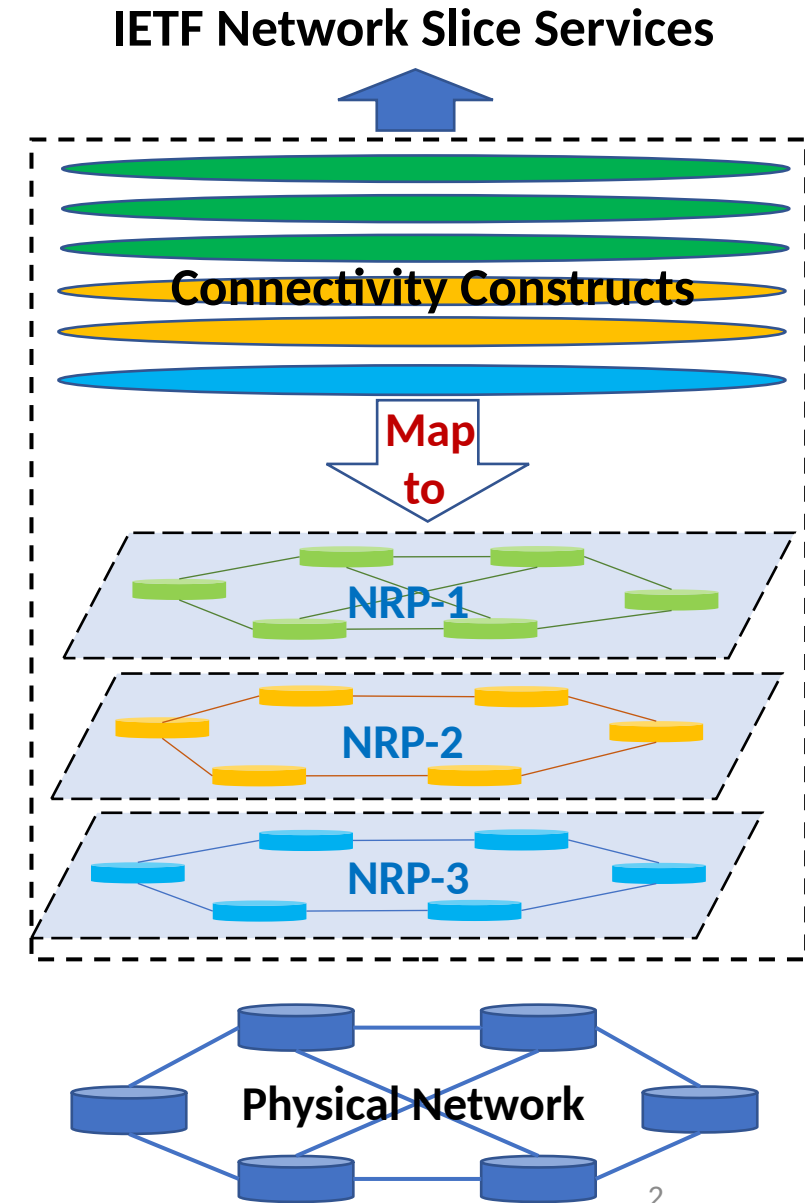
Gyan Mishra @Verizon

Tarek Saad @Cisco

Vishnu Beeram @Juniper

# Recap of Network Slicing and NRP

- The concept and general framework of IETF network slice is described in draft-ietf-teas-ietf-network-slices
  - IETF network slice services can be realized by mapping network slice connectivity constructs to the underlay NRPs
- 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 Control Plane Scalability Concerns

- The NRP scalability impact to distributed control plane may be related to the following aspects:
  - The number of control protocol instances maintained on each node
  - The number of control protocol sessions maintained on each link
  - The number of control messages advertised by each node
  - The amount of attributes associated with each message
  - The number of computations (e.g. SPF computation) executed by each node
- The scalability of centralized controller may also be a concern
  - The processing burden of global computation or optimization for multiple NRPs
  - The load on the communication channels for real time NRP information update from network nodes due to some network events (e.g. link or node failure)

# NRP Data Plane Selector Scalability Considerations

- The identification of NRP needs to be carried in data packet
  - Allow network nodes to determine the subset of network resources of the NRP and the associated topology for packet processing and forwarding
  - Different options of carrying NRP identification information in packet have different scalability implications
- Option 1: reuse existing fields in data packet to identify NRPs
  - Examples: SR SIDs, IP addresses, MPLS forwarding labels
  - Pros: No need of introducing new encapsulation to data packet
  - Cons: May bring scalability issues to the existing functions, may require change in semantics and processing behavior, may require control protocol extensions
- Option 2: introduce a dedicated data plane NRP ID
  - Examples: IPv6 extension headers, special purpose MPLS labels, MNA, etc.
  - Pros: Avoids the impacts to the existing fields, better scalability with network-wide NRP IDs
  - Cons: Cost of new data plane encapsulation and processing, may require control protocol extensions

# Proposed Scalability Optimizations

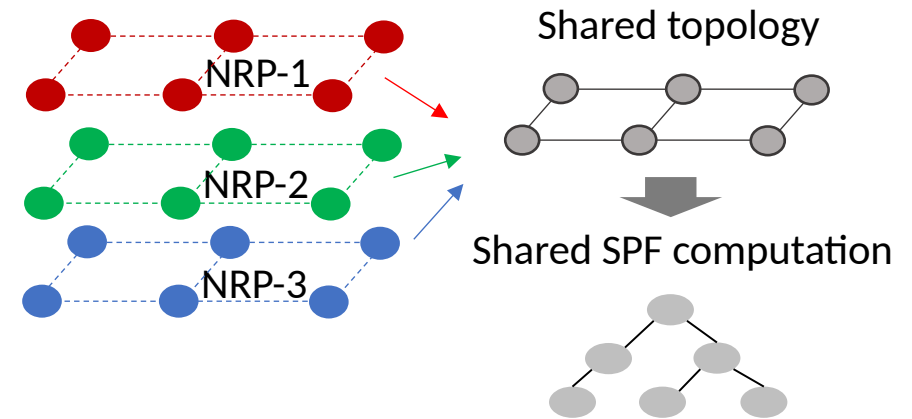
- Control plane scalability optimization

- Shared control protocol instances/sessions among multiple NRPs
- Shared topology specific computation among a set of NRPs
- Hybrid control plane with the help of centralized controller

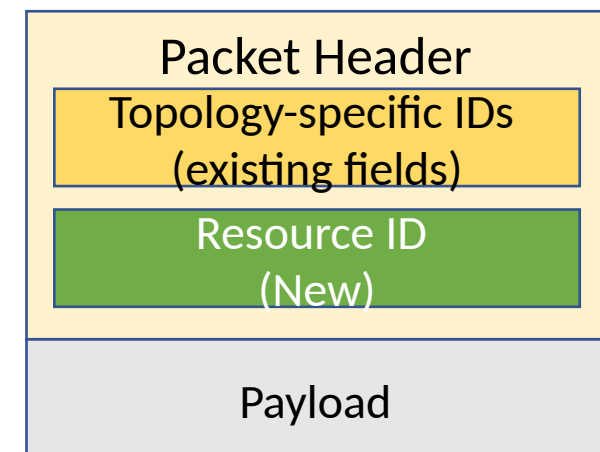
- Data plane scalability optimization

- Introduce a dedicated resource ID in data packet
- Decouple the resource ID from the topology-specific IDs in packet forwarding

- Shared topology and SPF computation between multiple NRPs



- Dedicated data plane resource ID



# Updates after WG Adoption

- Version -01:
  - Comprehensive editorial changes and clarifications according to comments received during and after the WG adoption
  - Many thanks to Donald, Kenichi, Med, Christian and Kiran for the review comments and suggestions
- Version -02:
  - Refresh with minor updates

# Recent Discussion on Design Principles of Scalable NRP

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

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

# Recent Discussion on Design Principles of Scalable NRP (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), or even experiments like CRH, 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.
11. Different packet transport mechanisms may use different means to carry the NRP ID. The writeup we need at this stage is agnostic to that.

# Recent Discussion on Solutions with Different Scales (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 plane is unaffected by this.
- There is always a trade-off between optimal solutions and scalable solution:
- 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
- The optimal solutions are good for specific environments, but
  - Might not work in other environments
  - May have scalability issues
- We should allow for both approaches, but we need to be clear of the costs and benefits in all cases
- 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

# Next Steps

- Have further discussion in the WG on NRP scalability
- Reach consensus on the design principles of scalable NRP solution
- Incorporate additional considerations and principles into this document
- All solution documents **MUST** include a section on scalability considerations

Thank You