NRP Selector – Options

Tarek Saad

Drafts of Relevance:

draft-ietf-teas-ns-ip-mpls

draft-ietf-teas-nrp-scalability

Network Resource Partition (NRP)

[RFC9543] defines:

- ***Network Resource Partition***: the subset of the buffer/queuing/scheduling resources and associated policies on the set of links in the underlay network associated with the NRP
- ***Connectivity Construct***: the set of Service Demarcation Points (SDPs) that together with a communication type defines how traffic flows between them

One or more connectivity constructs from one or more IETF Network Slices can be mapped to one NRP

The traffic flows in a connectivity construct that are mapped to an NRP are assigned the resources of that NRP

Network Resource Partition (NRP) Policy

[I-D.ietf-teas-ns-ip-mpls] defines:

- ***NRP Policy***: as a policy construct that enables instantiation of mechanisms in support of service specific control and data plane behaviors on topological elements associated with the NRP
  - The ***NRP*** is instantiated on the associated ***link/nodes*** using the ***NRP Policy***
- ***NRP ID***: an identifier that is globally unique within an NRP domain and that can be used in the control or management plane to identify the resources associated with the NRP
  - The ***NRP ID*** may be used to distinguish different NRPs – e.g. in the management plane
  - The ***NRP ID*** is a key identifier of the ***NRP Policy***

NRP Selector

[I-D.ietf-teas-ns-ip-mpls] also defines:

- ***NRP Selector:*** one or more fields (markings) within the packets of traffic streams that are inspected and used to map them to an NRP
  - The NRP Selector is used to select the NRP traffic streams and provide them with corresponding forwarding treatment
- A ***single NRP*** may have multiple ***NRP Selectors***
  - For example, multiple types of traffic streams (IPv6, MPLS, etc.) may be mapped to the same NRP
  - Each data-plane encapsulation may have different NRP selector(s) to select NRP packets

Options for Carrying NRP Selector

- ***Overload*** existing fields in the packet - for example:
  - Assign a unique MPLS forwarding label per NRP for each FEC
  - Assign a unique IPv6 (source or destination) address or SRv6 SID per NRP
- Assign a ***dedicated*** NRP selector identifier to be embedded in packets that belong to an NRP traffic stream (independent of forwarding)
  - Piggyback on an existing field of the packet:
    - Part of the IPv6 address or SRv6 SID
    - Part of the IPv6 Flow label
    - Part of an MPLS label (e.g. entropy label (EL))
  - Add a new field in the packet:
    - Within or post the MPLS label stack
    - Inside the IPv6 Hop-by-Hop Options header

Overloaded Forwarding Identifier as NRP Selector

- A unique forwarding address or MPLS forwarding label is assigned for each FEC per NRP – e.g. I-D.ietf-spring-resource-aware-segments
  - The per NRP forwarding address or forwarding labels needs to be distributed to nodes in the network
  - The forwarding plane needs to store per NRP forwarding state
- A node that is enabled for *'M'* NRPs, and has reachability to *'N'* other nodes will have to store state corresponding to *'MxN'* forwarding addresses/identifiers in its forwarding and control planes

Overloaded Service Identifier as NRP Selector

- A service identifier that is carried in a packet may be overloaded as an NRP Selector to map traffic to an NRP
  - For example, a VPN service label may be inspected by nodes to map traffic to an NRP
- The VPN service label(s) need to be distributed to NRP nodes, Options include:
  - Carrying in the NRP Policy
  - Signaling

Distribution of per NRP State

- Control plane as distribution mechanism
  - The signaling and/or routing protocols can be extended to carry the per NRP state
  - The distribution of the additional per NRP state, however, may impact the stability and scalability of existing routing protocols and/or signaling protocols
  - The WG discussed this approach in detail and produced a set of recommendations against it (documented in I-D.ietf-teas-nrp-scalability)
  - Rev-03 of I-D.ietf-teas-nrp-scalability:

The 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.

- Other NRP state distribution mechanisms
  - ***NRP Policy*** *via management plane*

Dedicated NRP Selector Identifier

- A dedicated NRP selector identifier can be used as NRP Selector that is carried in packets associated with the NRP
- The NRP Policy carries the dedicated NRP selector identifier is used as distribution to nodes
- No additional per NRP forwarding state on nodes
- No routing protocol extensions required
- Hardware needs to be upgraded to support parsing the new field

NRP Selector vs NRP-ID

- NRP Selector may or may not have a dedicated identifier
- The NRP-ID is a key identifier that identifies an NRP
  - Currently modelled as a 32-bit opaque identifier
  - The NRP-ID need not be the same as the dedicated NRP selector identifier
  - In certain scenarios, the NRP-ID may be encoded in the dedicated NRP selector identifier field (provided the NRP-ID fits in the NRP selector identifier space)

Open Questions

- NRP selector options:
  - Overloading existing field, or assigning a dedicated NRP selector field? Or, both?
- NRP scale:
  - How many NRPs can be expected in a given domain?
  - How many bits are needed for encoding the NRP-ID?
  - How many bits are needed for encoding the dedicated NRP selector identifier?
    - Should there be flexibility in the number of bits allowed?
    - Should there be multiple encoding options per dataplane?