



YANG DATA MODEL FOR NETWORK RESOURCE PARTITION POLICY

draft-bestbar-teas-yang-nrp-policy-00

Replaces draft-bestbar-teas-yang-slice-policy-02

Tarek Saad

Vishnu Pavan Beeram

Bin Wen

Daniele Ceccarelli

Shaofu Peng

Ran Chen

Luis M. Contreras

Xufeng Liu

Juniper Networks

Juniper Networks

Comcast

Ericsson

ZTE Corporation

ZTE Corporation

Telefonica

IBM Corporation

INTRODUCTION

- A Network Resource Partition (NRP) [[I-D.ietf-teas-ietf-network-slices](#)] is a collection of resources identified in the underlay network to support the IETF Network Slice service (or any other service that needs logical network structures with required characteristics to be created).
- An NRP Policy [[I-D.bestbar-teas-ns-packet](#)] is a policy construct that enables instantiation of mechanisms in support of service specific control and data plane behaviors on select topological elements associated with the NRP.
- Draft defines a YANG data model for the management of NRP policies on NRP capable nodes and controllers in IP/MPLS networks.
 - The latest (renamed) version is aligned with the terminology used in [[I-D.ietf-teas-ietf-network-slices](#)] and [[I-D.bestbar-teas-ns-packet](#)]

CATERING TO NRP POLICY MODES

- An NRP policy specifies the rules for determining the topology associated with the NRP and dictates how an NRP can be realized in IP/MPLS networks using one of three modes.
 - Partitioning of the shared network resources can be achieved in:
 - a) just the data plane or in
 - b) just the control plane or in
 - c) both the control and data planes.
 - The NRP policy modes (a) and (c):
 - Require the forwarding engine on each NRP capable node to identify the traffic belonging to a specific flow aggregate and to apply the corresponding Per-Hop Behavior (PHB).
 - Identification of the flow aggregate that the packet belongs to and the corresponding forwarding treatment that needs to be applied to the packet is dictated by the NRP policy.
 - The NRP policy modes (b) and (c):
 - Require the distributed/centralized resource reservation manager in the control plane to manage NRP resource reservation.
 - The provisions for enabling NRP state aware traffic engineering (NRP-TE) [[I-D.bestbar-teas-ns-packet](#)] are dictated by the NRP policy.
- The data model discussed in this document caters to all three NRP Policy modes.

MODEL STRUCTURE

- The top-level 'networks' container [RFC8435] is augmented with a set of NRP policies

```
module: ietf-nrp-policy
  augment /nw:networks:
    +--rw nrp-policies
      +--rw nrp-policy* [name]
        +--rw name string
        +--rw nrp-id? uint32
        +--rw resource-reservation
        | + .....
        +--rw flow-agg-selector
        | + .....
        +--rw phb? string
        +--rw topology
          +--rw filters
            +--rw filter* [filter-ref]
            + .....
            +--rw resource-reservation
            | + .....
            +--rw flow-agg-selector
            | + .....
            +--rw phb? string
```

MODULE IETF-NRP-POLICY: NRP POLICIES

- The 'nrp-policies' container carries a list of NRP policies.
- Each 'nrp-policy' entry is identified by a name and holds the set of attributes needed to instantiate the NRP.
- Each entry also carries an 'nrp-id' leaf which uniquely identifies the NRP created by the enforcement of this policy.
- Key elements of an NRP policy:
 - Resource Reservation
 - Flow-Aggregate Selector
 - Per-Hop-Behavior
 - Topology Filters

NRP POLICY: RESOURCE RESERVATION

- The 'resource-reservation' container carries data nodes that are used to support NRP state aware bandwidth engineering.
- The data nodes in this container facilitate preference-based preemption of NRP state aware TE paths, sharing of resources amongst a group of NRPs and backup path bandwidth protection.

```
+--rw resource-reservation
|  +--rw preference?                               uint16
|  +--rw (max-bw-type)?
|  |  +--:(bw-value)
|  |  |  +--rw maximum-bandwidth?                 uint64
|  |  |  +--:(bw-percentage)
|  |  |  |  +--rw maximum-bandwidth-percent?
|  |  |  |  |  rt-types:percentage
|  +--rw shared-resource-groups*                   uint32
|  +--rw protection
|  |  +--rw backup-nrp-id?                          uint32
|  |  +--rw (backup-bw-type)?
|  |  |  +--:(backup-bw-value)
|  |  |  |  +--rw backup-bandwidth?                uint64
|  |  |  |  +--:(backup-bw-percentage)
|  |  |  |  |  +--rw backup-bandwidth-percent?
|  |  |  |  |  |  rt-types:percentage
```

NRP POLICY: FLOW AGGREGATE SELECTOR

- The 'flow-agg-selector' container carries data nodes that specify the rules for identifying which packets belong to the flow aggregate that this NRP caters to.

```
+--rw flow-agg-selector
| +--rw mpls
| | +--rw (fas-type)?
| |   +--:(label)
| |   | +--rw (specification-type)?
| |   |   +--:(derived)
| |   |   | +--rw forwarding-label?          empty
| |   |   +--:(explicit)
| |   |   +--rw label?
| |   |   |   rt-types:mpls-label
| |   |   +--rw label-position?
| |   |   |   identityref
| |   |   +--rw label-position-offset?      uint8
| |   +--:(label-ranges)
| |     +--rw label-range* [index]
| |       +--rw index                      string
| |       +--rw start-label?
| |       |   rt-types:mpls-label
| |       +--rw end-label?
| |       |   rt-types:mpls-label
| |       +--rw label-position?            identityref
| |       +--rw label-position-offset?    uint8
| +--rw ipv4
| | +--rw destination-prefix*      inet:ipv4-prefix
| +--rw ipv6
| | +--rw (fas-type)?
| |   +--:(ipv6-destination)
| |   | +--rw destination-prefix*    inet:ipv6-prefix
| |   +--:(ipv6-hbh-eh)
| |     +--rw fas-hbh-eh*           uint32
| +--rw acl-ref*      nrp-policy-acl-ref
```

NRP POLICY: PER-HOP-BEHAVIOR

- The 'phb' leaf carries a name of a PHB profile available on the topological element where the policy is being enforced.

```
+--rw phb?                string
```


NRP POLICY: TOPOLOGY FILTERS

- The 'topology' container consists of a list of filters where each entry references a topology filter [I-D.bestbar-teas-yang-topology-filter].
 - The resultant topology from the union of these filters is referred to as the NRP topology.
- The topological elements that satisfy the membership criteria can optionally **override** the default resource-reservation, flow-agg-selector and phb specific leafs.

```
+--rw topology
  +--rw filters
    +--rw filter* [filter-ref]
      +--rw filter-ref
        |      nrp-policy-topo-filter-ref
      +--rw resource-reservation
        | + .....
      +--rw flow-agg-selector
        | + .....
      +--rw phb?                               string
```

NEXT STEPS

- In the next revision, the 'topology' container will be augmented to include a reference to the resultant NRP topology state.
- Request review and feedback.



THANK YOU

draft-bestbar-teas-yang-nrp-policy@ietf.org