

Carrying NRP Information in IPv6 Extension Header

draft-ietf-6man-enhanced-vpn-vtn-id-06

Jie Dong, Zhenbin Li @Huawei

Chongfeng Xie, Chenhao Ma @China Telecom

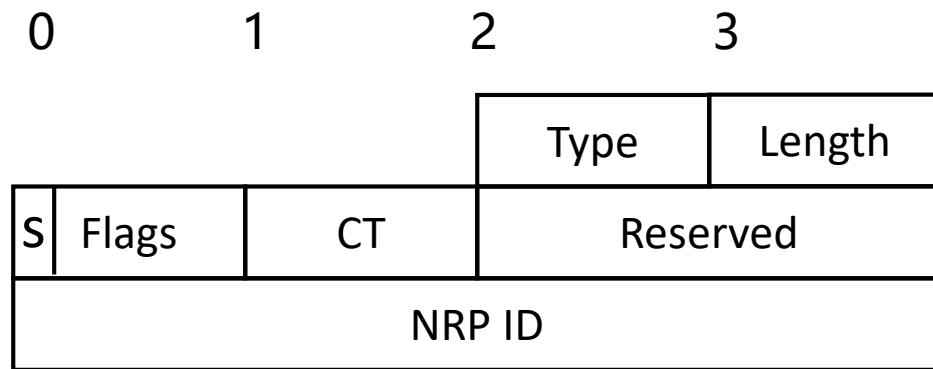
Gyan Mishra @Verizon

Background Recap

- Network Resource Partition (NRP) is a collection of network resources allocated on a set of links in the underlay network
 - The NRP concept is introduced in RFC 9543 as an underlay network resource construct to support network slice or enhanced VPN services
 - Packets belonging to an NRP needs to be processed and forwarded by network nodes using the set of resources allocated to the NRP
- This document introduces a new HBH option to carry NRP related information in IPv6 packets
 - According to the consensus in TEAS WG, the terminology has been aligned with NRP
 - The new HBH option can be used by the transit nodes to determine the NRP a packet belongs to
 - NRP-specific packet processing and forwarding can be performed
- There are interests in 6man to generalize the NRP option
 - So that this option could be extended for other related functions in the future
 - The encoding needs to meet the current network slicing use case and allow future extensions

NRP Option Encoding

- Optimized encoding for both network slicing and future extensions



- Option Type: TBA
 - Length: Length of the data fields in octets
 - **Flags:** the first bit is defined as Strict Match (S)
 - **S flag:** Indicate whether the NRP ID MUST be strictly matched for the processing of the packet
 - **Context Type (CT):** Indicate the semantics and length of the NRP ID
 - **CT=0:** The NRP ID is a 4-octet resource ID
 - Reserved field: leave for future extensions
 - **NRP ID:** Network-wide unique identifier of an NRP
- The data plane processing overhead is taken into consideration
 - For a specific CT, the length of the NRP ID is fixed
 - The format of the option carries all necessary information efficiently

Updates in -06 Version

- The terminologies and descriptions are aligned with IETF network slice framework and enhanced VPN framework draft
 - The term “NRP” is used consistently to refer to network resource constructs used to support network services with enhanced characteristics
 - The new HBH option and ID are renamed as “NRP option” and “NRP ID” respectively
- The usage of the S flag and its benefit are elaborated
 - The S flag allows fine granular control of the forwarding behavior for packets in an NRP, when the NRP ID in the packet cannot match with locally provisioned NRPs
- Editorial changes to improve readability

Discussion

- Information needed in the dedicated NRP ID encoding
 - Identifier: the basic information
 - Flags: used to control the forwarding conditions and behavior
 - Reserved field: for future extensibility
 - Context Type: for generalization

Discussion (Cont.)

- Further on terminology: NRP Selector and NRP ID
 - NRP selector: a general term to refer to all possible data plane fields which could be used to determine the NRP which a packet belongs to
 - Covers both derived IDs (address, label, SR SID, etc.) and a new dedicated ID
 - The usage of NRP selector needs to be clarified: on ingress node, intermediate nodes, or both?
 - NRP ID: the identifier of an NRP, which can be used in management plane, control plane and data plane
 - The length of NRP ID in management, control and data plane could be the same or different
 - The length of NRP ID in different data plane encapsulations could be the same or different
- Which term to use for the new IPv6 HBH option and identifier?

Next Steps

- The document is approaching WG last call in 6man WG
- Discuss whether (some of) the design in IPv6 NRP option could be applied to other data plane

Thank You

Discussion on the S Flag

- Comments from Med:
 - The usage of S flag is not clear, can it be replaced by local configurations?
- Replies from author:
 - The S Flag allows flexible and fine-granular processing policy (discard or best effort forwarding) for a subset of packets (e.g. OAM packets or specific service flows) of an unmatched NRP, which is not provisioned on the interface of a transit node
 - It would be burdensome to realize the same policy via local configurations, as it has to be done on all nodes along the forwarding path, and it would require additional configurations for NRPs which are not locally provisioned
- Comments from Ketan:
 - This seems as a local configuration of drop or fallback
- Replies from author:
 - The indication of drop or fallback is not a simple global knob, it needs to be per-NRP or at flexible and fine granularity , and it is for NRPs which are not provisioned on the node