

MPLS Label Forwarding with No Swapping

draft-fang-mpls-label-forwarding-no-swap-01

Luyuan Fang, Microsoft

Fabio Chiussi

Barak Gafni, Mellanox

Nabil Bitar, Verizon

IETF 93, MPLS WG, July 23, 2015

Problem Statement

- Global / domain-wide labels are used more and more
 - Global labels, Source Routing, HSDN, etc. ...
- Global / domain-wide labels are not swapped in the switches
 - incoming label = outgoing label
 - From a forwarding behavior standpoint, there is no “label swapping”
- RFC3031 only defines swapping for forwarding action
 - This is for historical reasons: in 2001, labels were locally assigned

Label forwarding operation as defined in RFC3031

...

Section 3.10

The "Next Hop Label Forwarding Entry" (NHLFE) is used when forwarding a labeled packet. It contains the following information:

1. the packet's next hop
2. the operation to perform on the packet's label stack; this is one of the following operations:
 - a) replace the label at the top of the label stack with a specified new label
 - b) pop the label stack
 - c) replace the label at the top of the label stack with a specified new label, and then push one or more specified new labels onto the label stack.

...

Label Swapping defined in RFC 3031

...

3.13 Label Swapping

Label swapping is the use of the following procedures to forward a packet.

In order to forward a labeled packet, a LSR examines the label at the top of the label stack. It uses the ILM to map this label to an NHLFE. Using the information in the NHLFE, it determines where to forward the packet, and performs an operation on the packet's label stack. It then encodes the new label stack into the packet, and forwards the result.

...

Forwarding without Swapping is not current standard behavior

- “Swapping a label with another label taken from a table entry that happens to be equal” is not the same operation as no swapping
- An additional label forwarding operation referred to as No Swapping needs to be added to the current standard
- This is required for aligning the forwarding behavior with the needs of Global/Domain Wide label processing
- This is not a local optimization issue. This is an addition to the current standard to allow efficient implementation of domain wide label processing

Next Steps

- Encourage people to read RFC3031 again, carefully, and correctly interpret it from the perspective of “no swap”
- Continue discussion on the mailing list
- Welcome contributions