

THE MPLS “FIRST NIBBLE” AND POST-STACK DATA REGISTRY

Kireeti Kompella

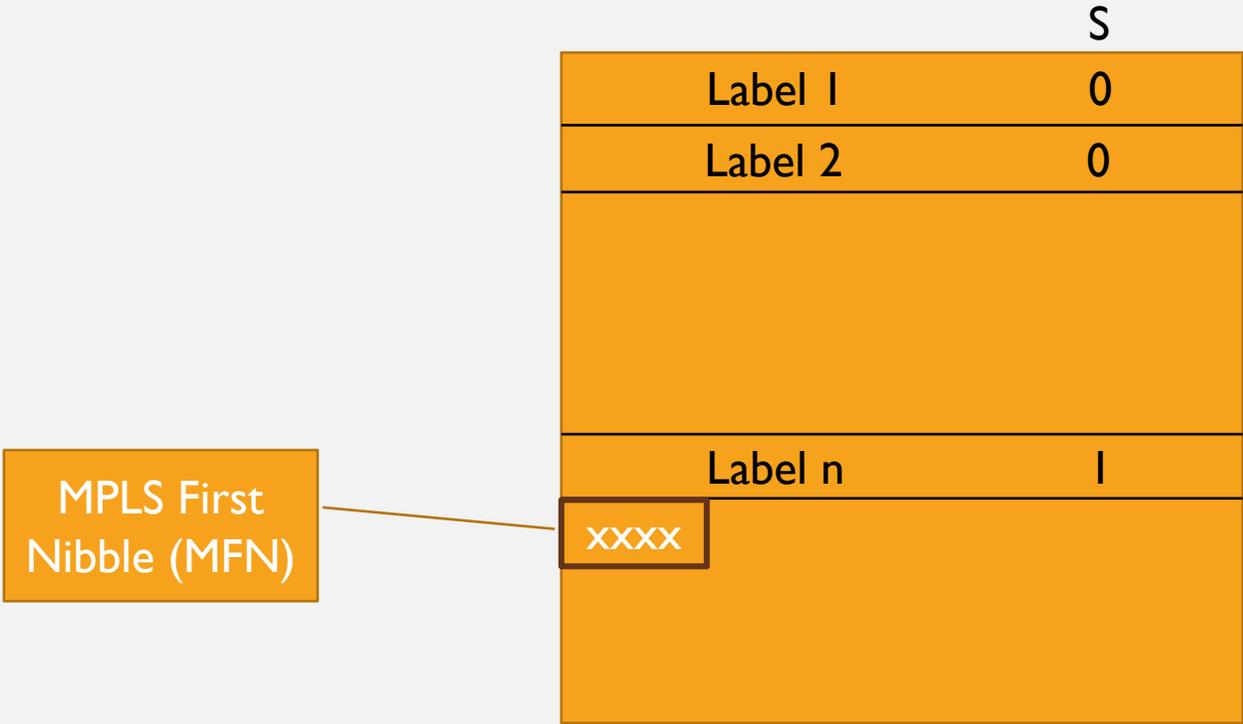
IETF 112

draft-kbbma-mpls-1st-nibble

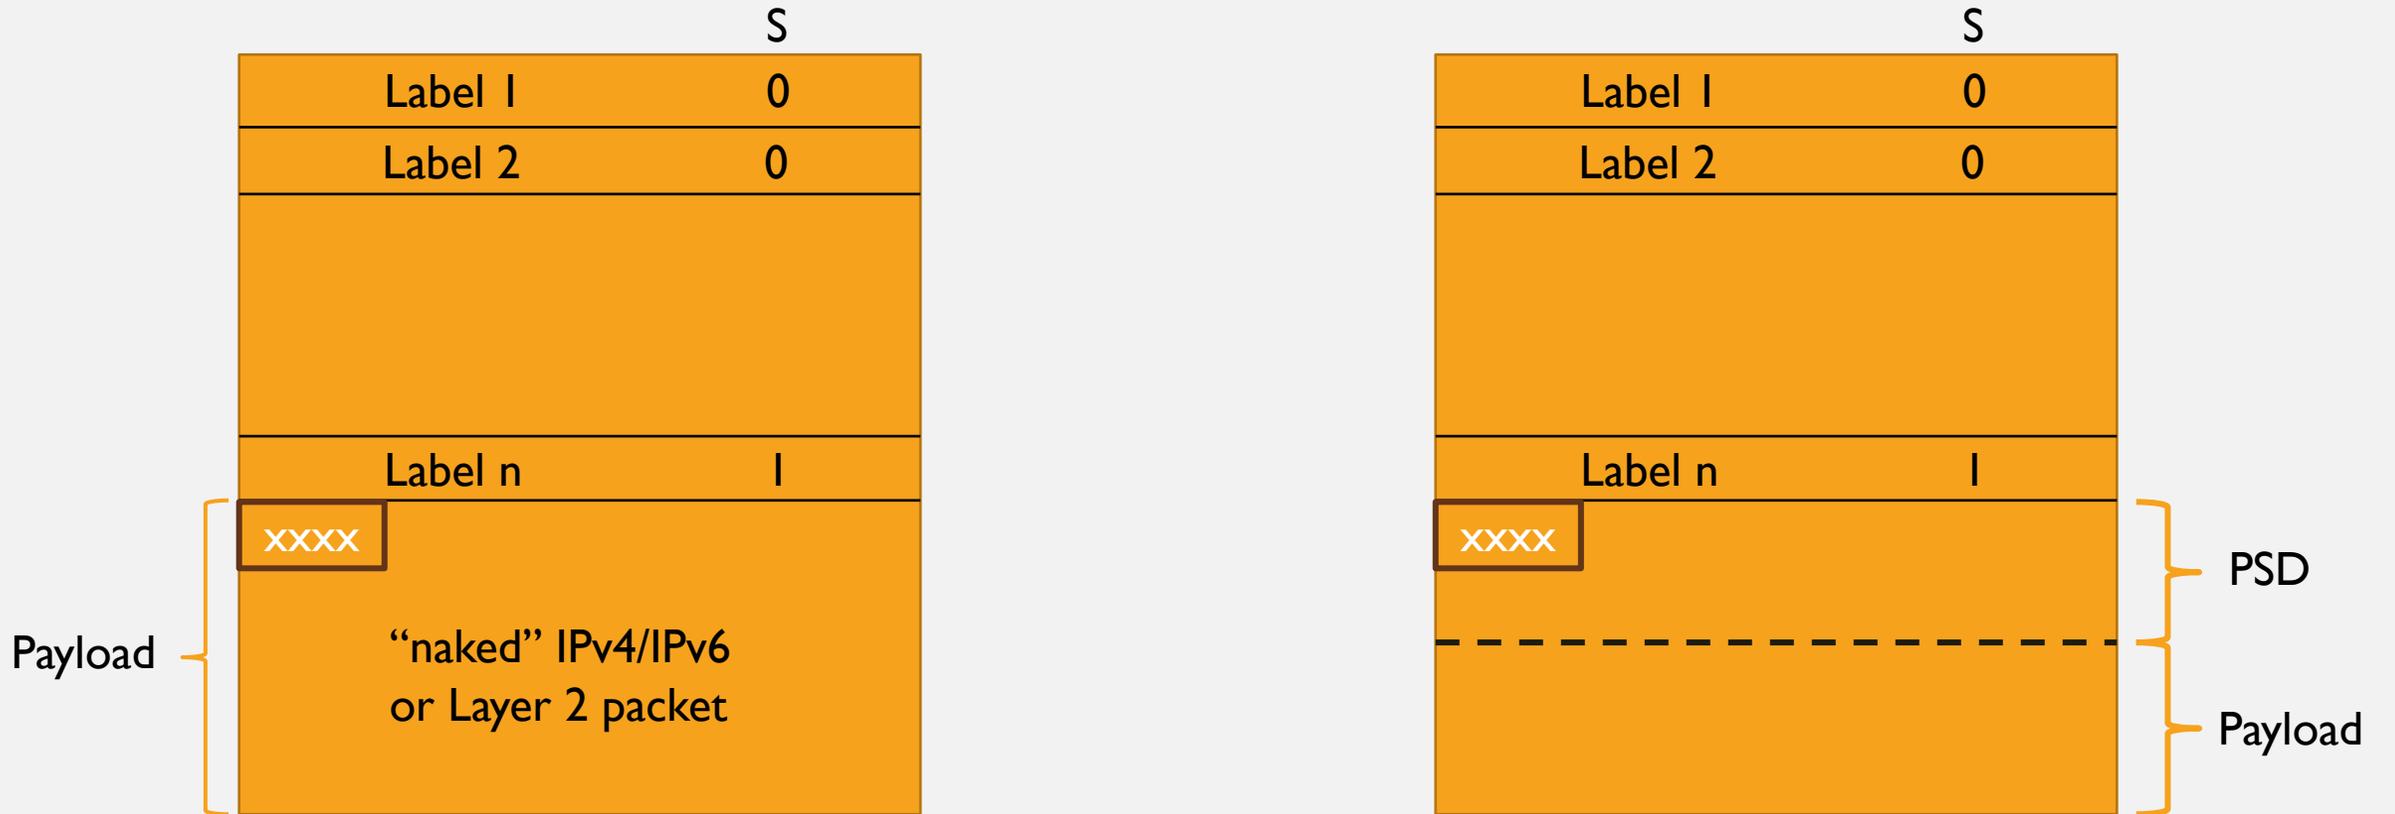
BACKGROUND

- The MPLS “first nibble” is the high-order 4-bit field (nibble) of the first octet after the last label in the stack (i.e., the one with the BoS bit set)
- This nibble has been used for various things in the past
 - and is still being used currently
- This draft explains the uses of this nibble, and proposes:
 1. A Requirement and a Recommendation on its use
 2. A registry for the first nibble values
 3. A registry for “post-stack data”

TERMINOLOGY



TWO CASES



USES OF THE MFN

1. As a heuristic that the payload is an IP packet and thereby extract fields for load balancing
2. As an indication of the type of “post-stack data” (PSD)

PSD today can be a pseudowire or detnet control word, or a BIER header, or a GACh, or ...

Note: PSD is often signaled in the control plane (e.g., it is known ahead of time that the packet will carry a PW control word)

THINGS ARE BROKEN

- We should recognize this
- The heuristic can fail badly
- We must allow current implementations to work, with the above big caveats
- We cannot continue to confuse IP version numbers with the MFN
- We have new and better ways to load balance all packets, not just IPv4/v6

MUST, MUST, MUST

- Allow current implementations to continue working, bad heuristics notwithstanding
- Lay the groundwork for better, more efficient implementations
- Lay the groundwork for easier, self-contained handling of PSD

PROBLEMS AND WORKAROUNDS

1. In case of an Ethernet payload, the MFN can take any value
 - There are RFCs recommending that Ethernet packets SHOULD have a control word
2. Current PSD types stay away from the values of 4 and 6
 - This is to prevent confusion between PSD and IPv4 and IPv6 packets which would have MFN values of 4 and 6, respectively
 - This leads to possible confusion between the IP version number registry and MFN values – suggestion to “merge” them
3. Recognition and parsing of the PSD is dependent on signaling and/or the presence of heralding labels (e.g., BIER label, GAL label)

LONGER TERM SUGGESTIONS

1. Mandate that, GOING FORWARD, the only “naked” payload types are IPv4 and IPv6
 - Any other version of IP, and any other type of packet MUST have a PSD
 - This decouples MFN from IP version numbers (except 4 and 6)
 - **Could go further and mandate a PSD for all payloads**
2. RECOMMEND that, for load balancing purposes, an Entropy or a FAT PW Label SHOULD be used
3. Keep recognition and parsing of the PSD completely in the data plane
4. Create an MFN registry that will serve simultaneously as a PSD registry

not trying to “backdoor”
an MPLS protocol field

EXISTING IMPLEMENTATIONS

- The load balancing heuristic can continue to be used
 - with the caveats already mentioned
- PW and DetNet CWs work as is; no changes are to be made to them
- GAch PSD continues as is
- BIER PSD continues as is
- Any new ideas for PSD MUST use a new MFN codepoint

The MFN space is very small; accommodations for expansion must be made

PROPOSAL

(not in draft)

MFN (4) | subtype | total length of PSD | rest of PSD