

# Post-stack First Nibble

draft-ietf-mpls-1stnibble

Kireeti Kompella  
Stewart Bryant  
Matthew Bocci

Greg Mirsky  
Loa Andersson  
Jie Dong

IETF-119, March 2024

# Why Do We Need This?

- Post-stack First Nibble (PFN) is the first four bits immediately following the Label Stack Element (LSE) with the S bit set to 1
- When the payload of an MPLS-encoded packet is IP, the IP version number happens to fall in the field reserved for PFN for other applications
- IP Version Numbers are tracked in the IANA “IP Version Numbers” registry
- PFN in Post-Stack Header (PSH) doesn’t have an IANA registry. This draft proposes the creation of the Post-stack First Nibble registry
- Deprecation: regardless of how the deprecation is understood in other IETF documents, the interpretation in this document is that if a practice has been deprecated, that practice should not be included in the new implementations or deployed in deployments

# Use of the PFN in load-balancing

- Load-balancing techniques
  - Several possible techniques, but not in scope for the draft
  - Identify packets with non-IP payload, to avoid load-balancing (RFC 4928), e.g., Entropy Label (RFC 6790) or FAT PW Label (RFC 6391)
- Heuristic for load-balancing:
  - If the PFN value is 0100b or 0110b – there is a high probability that the payload is IPv4 or IPv6. If there are other indications in the label stack supporting this, load sharing based on payload information may be used
  - If the PFN value is anything else, then the payload is non-IP packet. Load-balancing based on payload information should not be used. Other load-balancing techniques may be used

# How to Use PFN Safely?

It is not possible to reliably determine what follows the LSE with the Bottom of Stack bit set through inspection of the first nibble following that LSE only.

Safety rule:

Correct interpretation of the PFN in PSH can be done only in the context of the LSE or a group of LSEs in the preceding label stack that characterize the type of the PSH, and that any attempt to rely on the value in any other context is unreliable.

# Updates of RFC 4928

Paragraph 3 in Section 3 of RFC 4928 states that:

It is REQUIRED, however, that applications depend upon in-order packet delivery restrict the first nibble values to 0x0 and 0x1.

This will ensure that their traffic flows will not be affected if some future routing equipment does similar snooping on some future version(s) of IP.

RFC 4928 is to be updated as follows:

Network equipment that complies with [I-D.ietf-mpls-1stnibble] MUST use a PSH (Post-Stack Header) with a PFN (Post-stack First Nibble) value that is neither 0x4 nor 0x6 in all cases when the MPLS payload is not an IP packet.

# Updates of RFC 4928 (cont.)

The recommendation proposes to update the paragraph 4 in Section 3 of RFC 4928:

This behavior implies that if in the future an IP version is defined with a version number of 0x0 or 0x1, then equipment complying with this BCP would be unable to look past one or more MPLS headers, and load-split traffic from a single LSP across multiple paths based on a hash of specific fields in the IPv0 or IPv1 headers. That is, IP traffic employing these version numbers would be safe from disturbances caused by inappropriate load-splitting but would also not be able to get the performance benefits.

**RFC 4928 is to be updated as follows:**

[I-D.ietf-mpls-1stnibble] deprecated the practice of deducing the payload type to avoid inaccurate load balancing based on the PFN value. This means that older implementations and deployments can continue to use that heuristic, while it must not be part of new implementations or deployments. The deprecation also means that concerns about load balancing for future IP versions with a version number of 0x0 or 0x1 are now moot.

A new document is to be published to obsolete MPLS encapsulations without PSH of non-IP payload when sufficient evidence exists that there are no marketed or deployed implementations using the heuristic practice.

**Also, will add in Terminology:**

PSH: Post-Stack Header

PFN: Post-stack First Nibble

# PFN IANA Registry

Code points allocated for BIER		
Value	PSH Type	Reference
0x5	BIER header - Normal traffic	RFC 8296

  

Code points allocated for DETNET		
Value	PSH Type	Reference
0x0	DetNet Control Word	RFC 8964
0x1	DetNet Associated Channel	draft-ietf-detnet-mpls-oam

  

Code points allocated for Network Service Header (NSH)		
Value	PSH Type	Reference
0x0	NSH Base Header, Payload	RFC 8300
0x2	NSH Base Header, OAM	RFC 8300

  

Code points allocated for Pseudowires (PW)		
Value	PSH Type	Reference
0x0	PW Control Word	RFC 4385
0x1	PW Associated Channel	RFC 4385

# PFN IANA Registry (Cont.)

Code points allocated for the MPLS Generic Associated Channel		
Value	PSH Type	Reference
0x1	MPLS G-ACh	RFC 5586
Reserved Code Points, not to be allocated		
Value	Usage	Reference
0x4	IPv4 Protocol Number	RFC 791
0x6	IPv6 Protocol Number	RFC 8200
Unassigned Code Points		
Value	PSH Type	Reference
0x3	-	-
0x7-0xF	-	-



# Takeaways

- Please consider this work (draft-ietf-mpls-1stnibble) in your updates of use of the Control Word in draft-ietf-bess-rfc7432bis
- Welcome your comments, questions
- Move towards obsolescence of non-PSH for non-IP payload in the MPLS networks – new document

Thank you!