

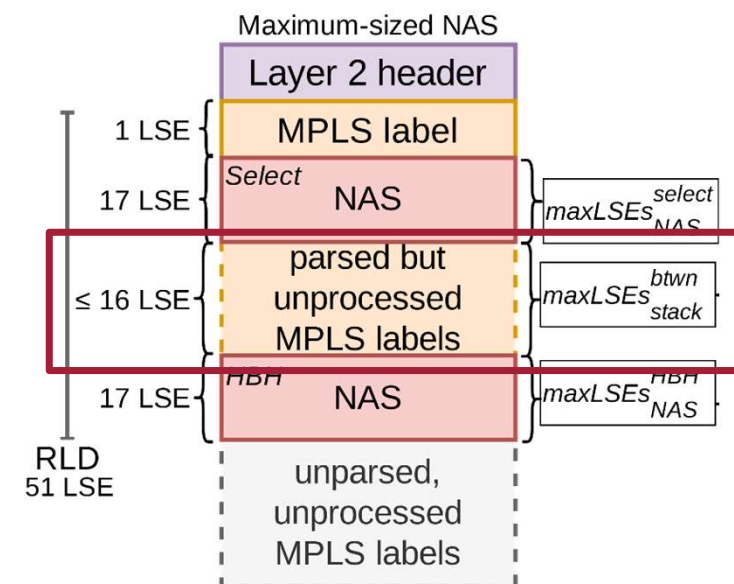


MPLS Network Action for Stack Management

Fabian Ihle <fabian.ihle@uni-tuebingen.de>, Michael Menth <michael.menth@uni-tuebingen.de>
draft-ihle-mpls-mna-stack-management-00

<http://kn.inf.uni-tuebingen.de>

- ▶ With MNA, the required RLD in LSRs is very large
- ▶ Why?
 1. 34 LSEs for NAS from IETF draft
 2. We must parse the „in-between stack“
 - HBH-scoped NAS is located deeper in the stack
 - Parse labels until the HBH NAS is found
 - Result of the protocol design
- ▶ LSEs in the in-between stack are not relevant for the current node
 - Resources must be allocated → increases required RLD
- ▶ Eliminate the in-between stack?
 - Place the HBH-scoped NAS at the top
 - ⚡ It will be popped
 - ⚡ Many HBH NAS copies required



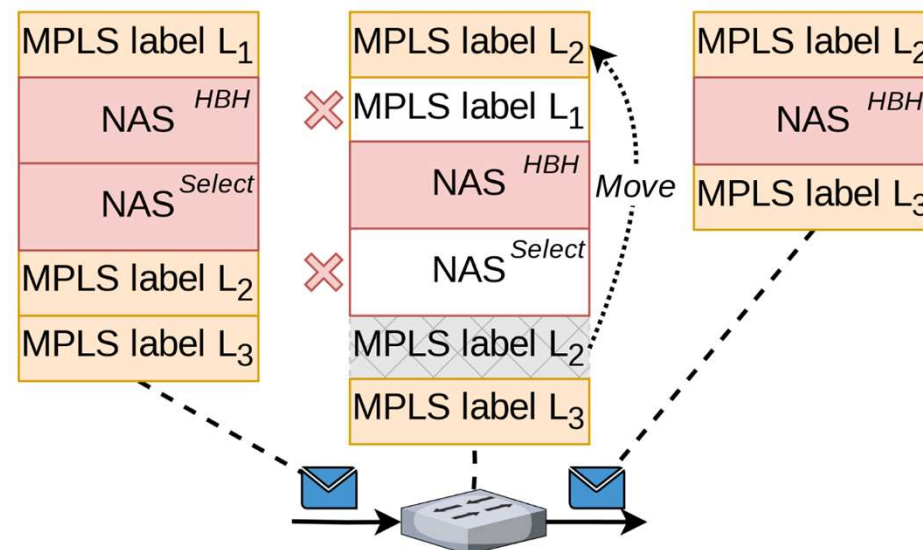
The HBH Preservation Mechanism

→ It would be better to keep the HBH NAS at the top

- Avoid parsing irrelevant LSEs
- But: The HBH NAS must be preserved

▶ Constrain the location of the HBH NAS and reorder the stack during forwarding

1. Always place the HBH NAS directly below the top-of-stack forwarding label
2. If top label is popped: Bring the label from below the NAS to the top
 - NAS is not exposed 😊
 - No HBH copies required 😊



The Stack Management Network Action (I)

- ▶ LSRs are required to perform this operation
 - How to indicate this?

- ▶ If only we had a framework to instruct LSRs with network actions ...

→ The Stack Management Network Action


The Stack Management Network Action (II)

- ▶ Opcode with two operations indicated in AD
 - MOVE-N-LSE: Move LSEs from below NAS to top-of-stack
 - HBH preservation
 - POP-N-LSE: Pop LSEs below the NAS
 - Stateless MNA-based Egress Protection (SMEP), see next talk


▶ Not tied to those use cases

▶ Can be encoded in Format B or Format C

- 4 bit for MOVE-N and POP-N



MOVE-N operation.

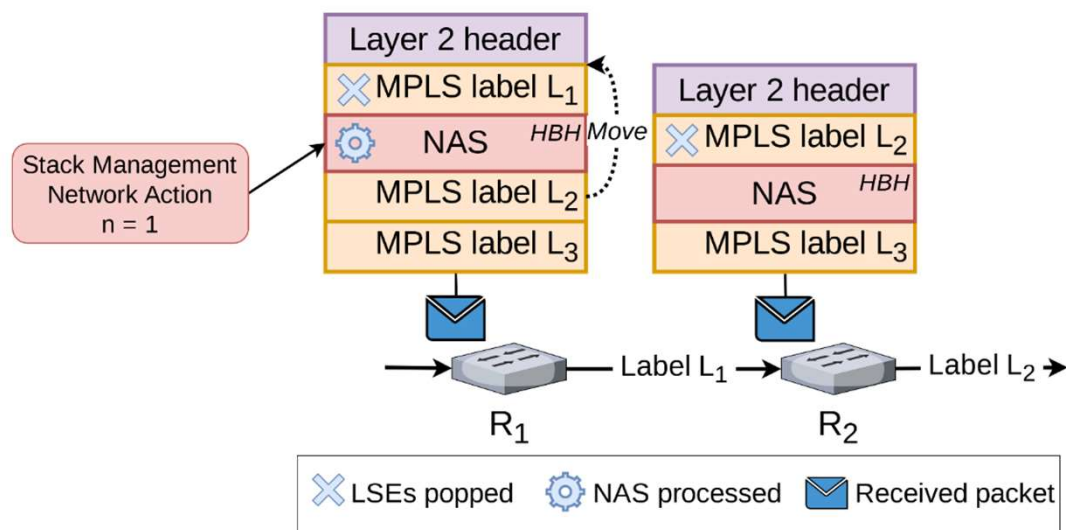


POP-N operation.

Bit →	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
FMT	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
B	Opcode = MGMT				Reserved				POP-N				MOVE-N				R	IHS				S	U	NASL = 0				NAL = 0				
C	Opcode = MGMT				Reserved								POP-N				MOVE-N				S	U	Reserved				NAL = 0					

- ▶ Stack management network action is used to MOVE-N LSEs to the top
 - Added by the ingress LER to the HBH-scoped NAS with MOVE-N = 1

- ▶ If a node pops the top-of-stack label → move one label to the top





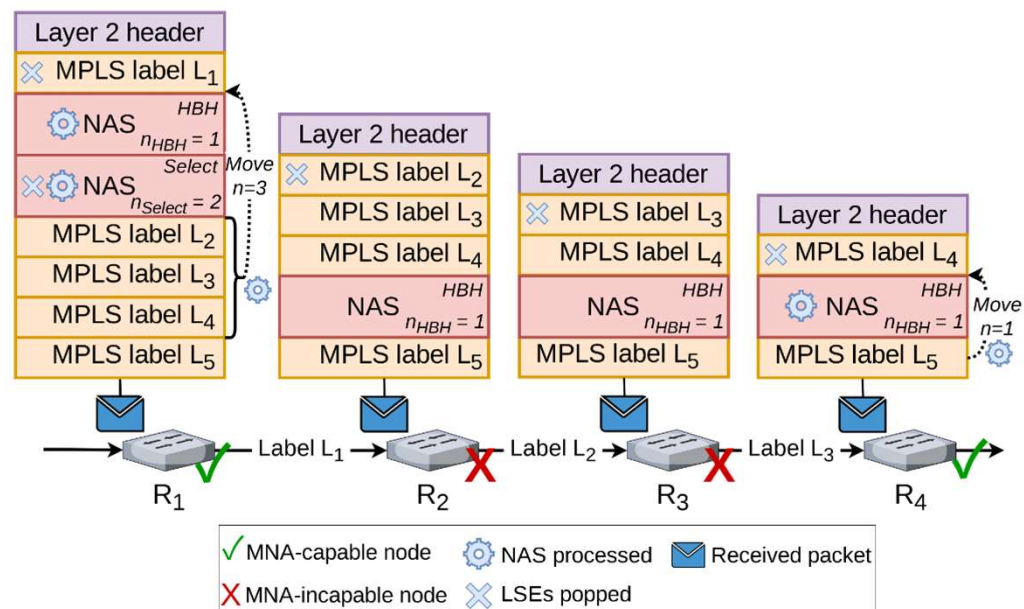
HBH Preservation – Backward Compatibility

► ↯ MNA-incapable nodes may expose the HBH NAS to the top

► The preceding LSR of an MNA-incapable node can bring more than one label to the top

- As many LSEs as popped by the MNA-incapable nodes
- Using a select-scoped NAS

→ MNA-incapable nodes do not expose the NAS to the top

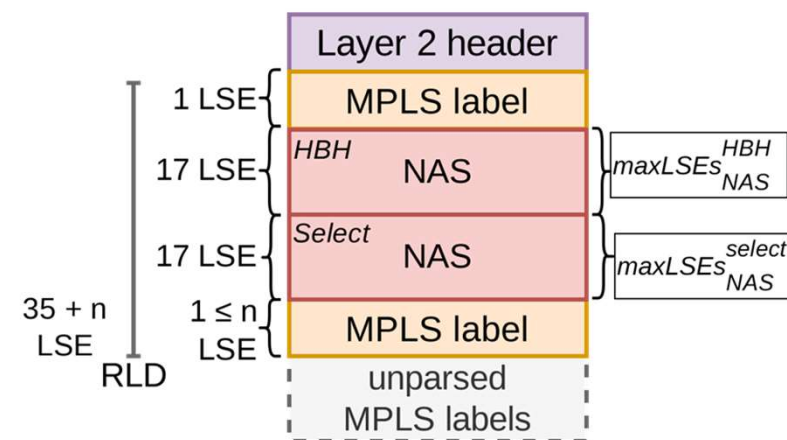


► Proof of Concept implementation in P4 on switching ASIC

- Line rate processing of the stack management operation at 400 Gb/s per port
 - Tested with network topology from previous example
- Does not work on legacy (push, pop, swap) hardware
 - → Mechanism for backward compatibility

► Effect on the required RLD

- No more „in-between“ stack
 - No resources wasted for parsing
- No more HBH copies
 - Reduces the overall stack size
- Minimum required RLD is $35 + n$ LSEs
 - $n > 0$: number of LSE that needs to be brought to the top at once
- Nodes with a smaller RLD can be integrated in MNA



- ▶ The security issues discussed in [[I-D.ietf-mpls-mna-hdr](#)] apply to this document.

- ▶ This document requests that IANA allocates a new codepoint with the name “Stack Management” in the “Network Action Opcodes Registry”.

MNA Opcode	Description	Reference
TBA1	Stack Management Network Action	This document



▶ MPLS network action for stack management

- Features a MOVE-N and a POP-N operation
 - MOVE-N moves a number of LSEs from below the NAS to the top
 - POP-N pops a number of LSEs below the NAS
 - More operations may be added

▶ Not tied to a specific use case

- Can be used for various applications
 - MOVE-N: HBH preservation
 - POP-N: Stateless MNA-based Egress protection (SMEP), see next talk

▶ HBH preservation reduces required RLD using the proposed action

- MPLS stack is reorganized
- No changes in MNA drafts required
- P4 implementation with 400 Gb/s per port available
- More information tomorrow at [ACM/IRTF Applied Networking Research Workshop](#)
 - El Prado, 17:00: <https://www.irtf.org/anrw/2025/program.html>



F. Ihle, M. Menth: [Stack Management for MPLS Network Actions: Integration of Nodes with Limited Hardware Capabilities](#), in [ACM/IRTF Applied Networking Research Workshop \(ANRW\)](#), Jul. 2025, Madrid, Spain



▶ Discussion

- More use cases?
- More operations?
- Thoughts about the HBH preservation?
- Precedence of operations?
 - MOVE-N before POP-N?

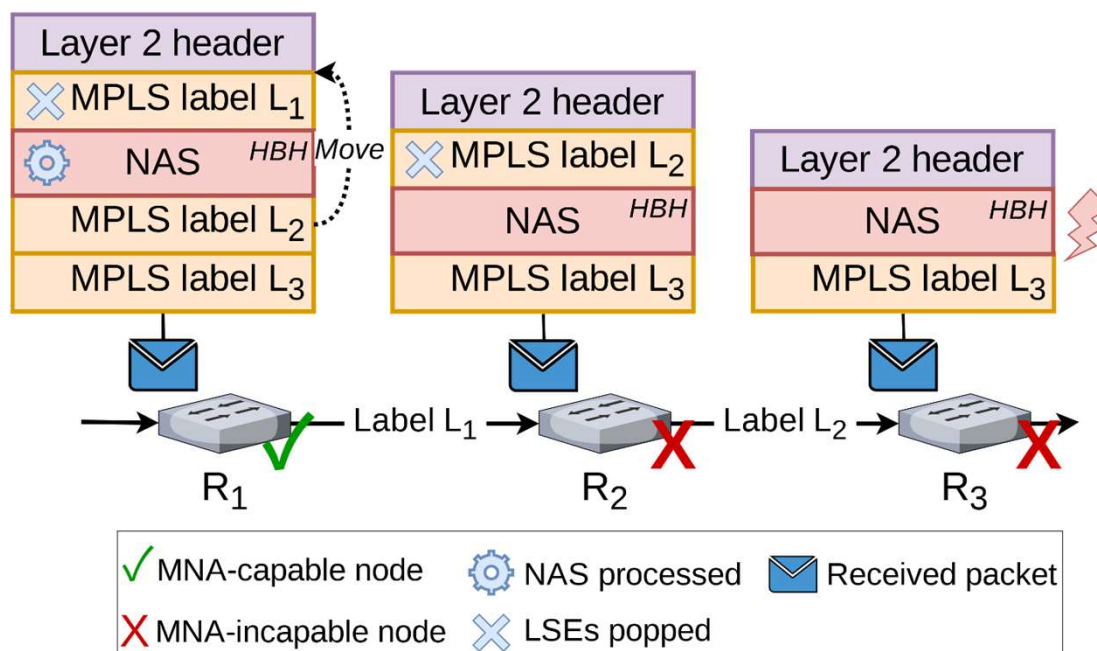
▶ Looking for collaborators

▶ Feedback welcome!



THANK YOU

- ▶ MPLS networks contain many legacy nodes
- ▶ ⚡ The mechanism is not compatible with brownfield networks
 - ⚡ An MNA-incapable node exposes the NAS to the top without popping it
 - The next node will drop such a packet



- ▶ The preceding node of a MNA-incapable node brings more than one label to the top
 - As many labels as popped by the MNA-incapable nodes

→ An MNA-incapable node that pops the top-of-stack forwarding label does not expose a NAS to the top

