



MPLS Network Actions: A Technological Overview

Fabian Ihle, Michael Menth

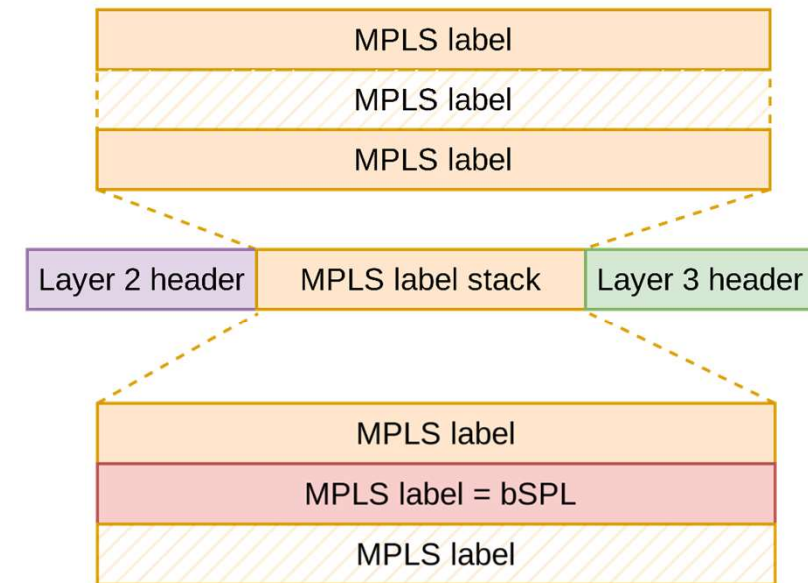
<http://kn.inf.uni-tuebingen.de>

- ▶ Motivation
- ▶ The MNA Framework
 - Encoding
 - The Network Action sub-stack
 - Scopes and placement
 - The readable label depth
- ▶ Use Cases for MNA
 - Link-specific packet loss measurement
 - Network slicing
 - Stateless MNA-based egress protection
- ▶ P4-MNA Prototype
- ▶ Conclusion



Motivation

- ▶ MPLS is a predominant WAN technology
 - Packets are switched based on assigned labels
- ▶ New applications require not only forwarding information in labels
 - Entropy labels for ECMP, OAM, Service Function Chaining (SFC), ...
- ▶ → Special Purpose Labels (SPLs) are well-established to provide special processing
 - Reserved label range
 - Available label space is limited ⚡
- ▶ The MPLS Network Actions (MNA) framework introduces a mechanism for transmitting and processing predefined network actions
 - Facilitates extensions to MPLS
 - MNA is to MPLS what IPv6 extension headers are to IPv6
 - Under standardization by the IETF MPLS WG
 - Disclaimer: the presented concepts are not yet finalized and may change



<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-fwk-15.html>

The MNA Framework

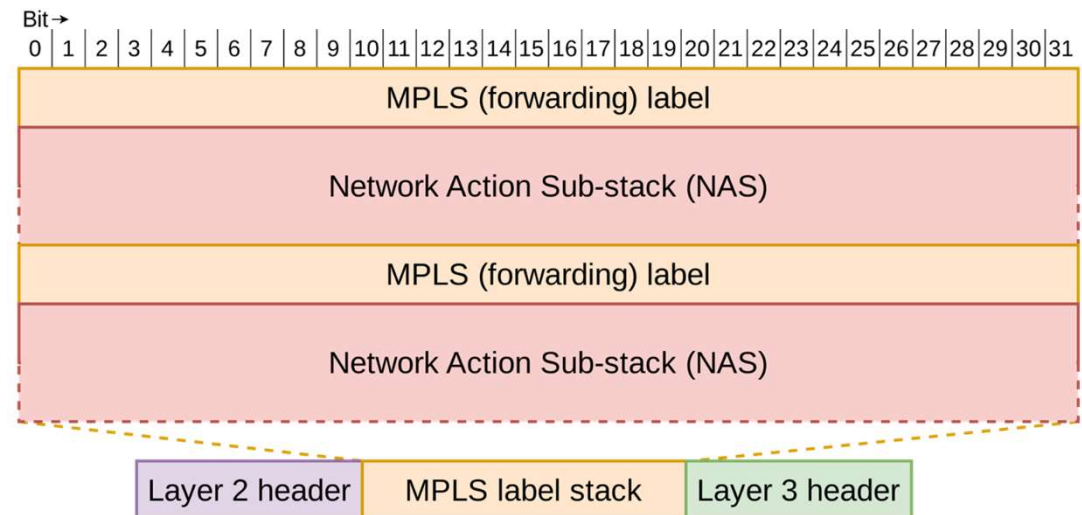


- ▶ New MPLS header encoding to specify network actions and their data
 - Encoded into label stack entries (LSE)

- ▶ The Network Action Sub-stack (NAS)

- A stack of related LSE
 - containing network actions and their data
- Inserted into the MPLS stack
 - in-stack data, ISD
- Added after the MPLS stack
 - post-stack data, PSD

- ▶ A NAS must never be at the top-of-stack
 - Ensures backward compatibility
 - NAS at the top-of-stack must be popped



<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-hdr-11.html>



The Network Action Sub-stack (NAS)

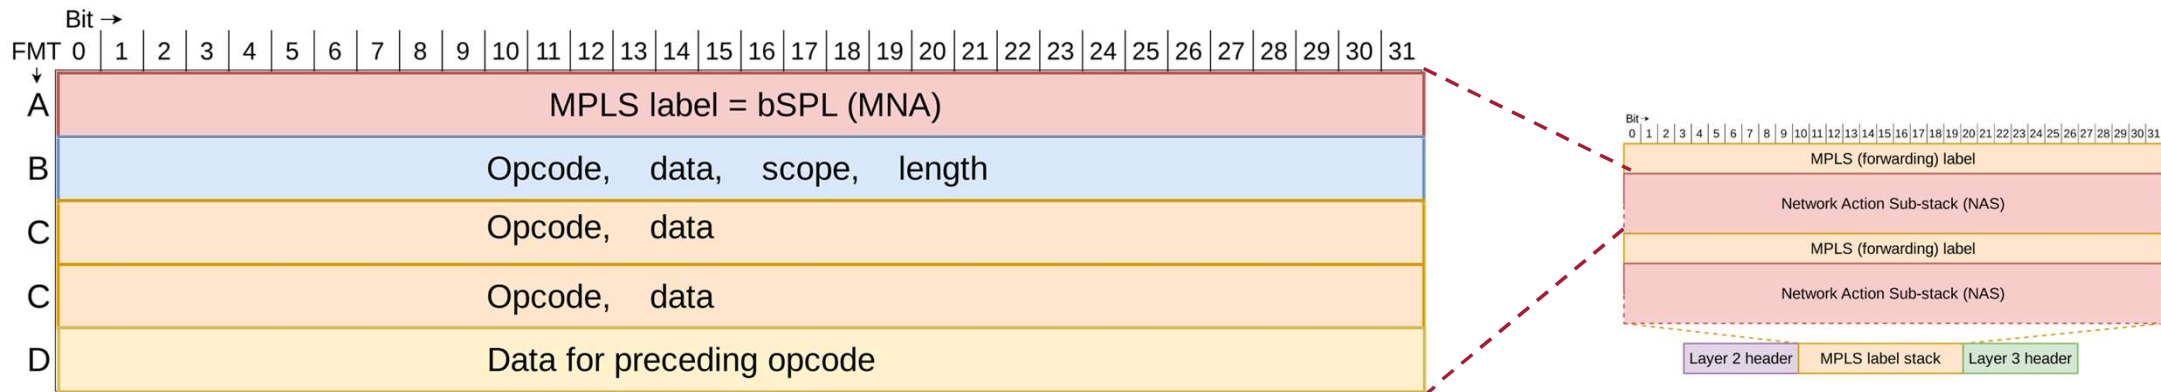
► What does a NAS look like?

- → Special label value (bSPL) followed by network actions and their data
 - = NAS indicator

► LSE encoding redefined for network actions and ancillary data in a NAS

- Initial opcode (mandatory, Format B)
- Subsequent opcode (optional, Format C)
- Ancillary data (optional, Format D)

► Up to 17 LSEs large

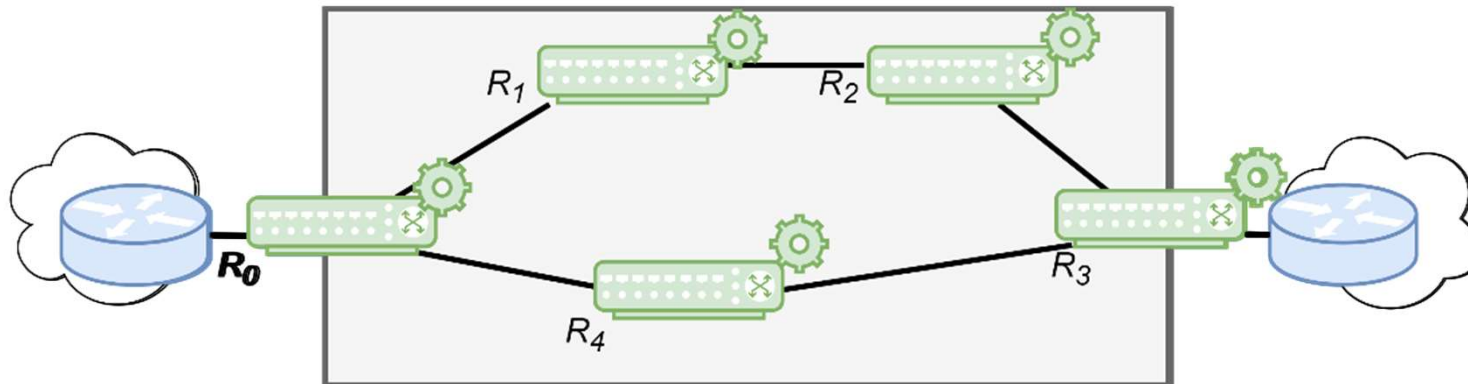


<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-hdr-11.html>



Scopes in the MNA Framework

- ▶ On which nodes in a network should network actions be processed?
 - Only on **selected** nodes
 - Only on the egress node (Ingress-to-Egress, **I2E**)
 - On all nodes (hop-by-hop, **HBH**)
- ▶ Indicated in Format B LSE of NAS
- ▶ MPLS stack including actions is pushed by ingress node





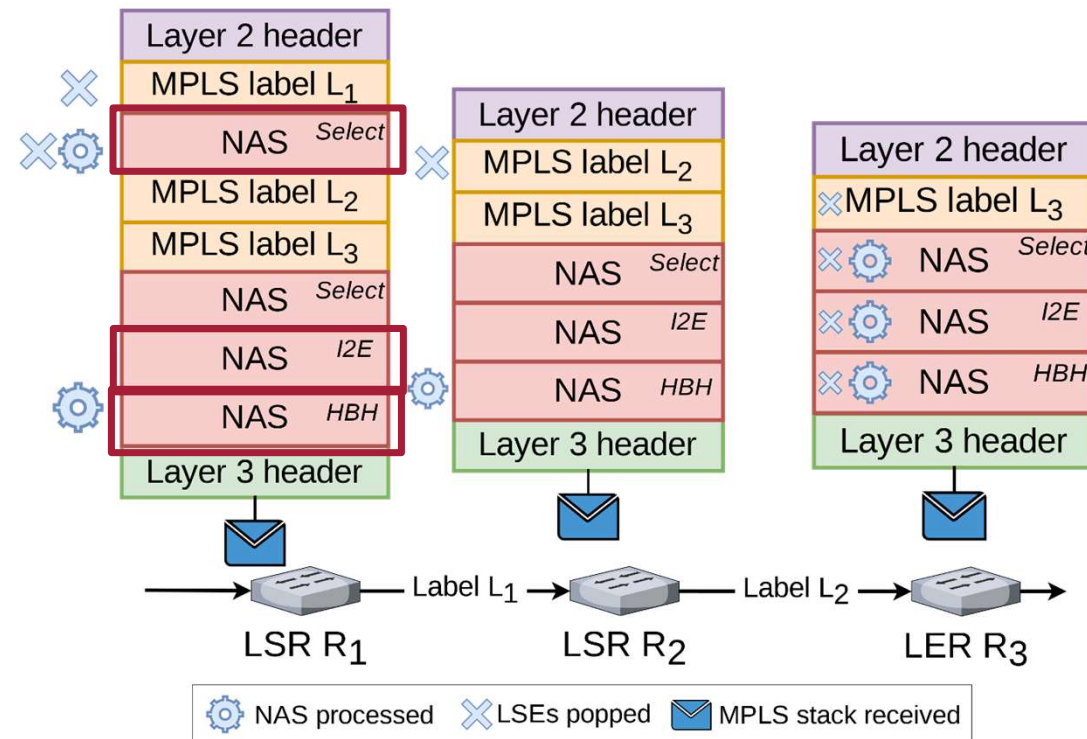
Placement of NAS in the Stack

► Placement per scope

- *Select*-scoped NAS are located below the forwarding label for the node
- *I2E*-scoped NAS are placed at the bottom of stack
- *HBH*-scoped NAS must be readable by each node
 - Here: Located at the bottom of stack

► A transit node must...

1. Process and pop a select-scoped NAS if it is exposed to the top
2. Search through the stack to find a HBH-scoped NAS



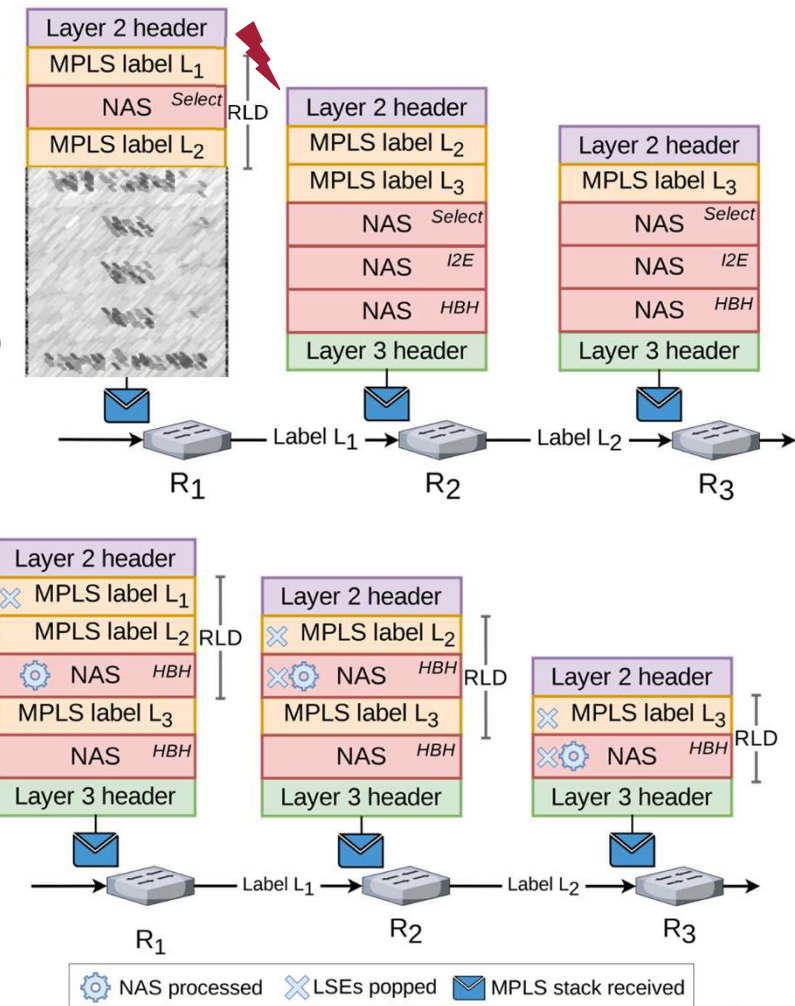
<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-hdr-11.html>



The Readable Label Depth (RLD)

- ▶ Hardware devices can only read a limited number of LSEs
 - = The readable label depth (RLD)
 - Network actions must be within RLD to process them
- ▶ Difficult for HBH-scoped NAS: must be readable by every node
 - Placing the NAS at the bottom works if the entire MPLS stack is in RLD
 - This is typically not the case
 - A single NAS may be up to 17 LSEs large
- ▶ Ingress node must ensure that a NAS is readable once it reaches a node
 - All nodes signal their RLD to the ingress node
 - Places copies of the HBH NAS in the stack
 - A node only processes the top copy of a HBH NAS
- ▶ → The RLD is a critical parameter

<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-hdr-11.html>



Use Cases

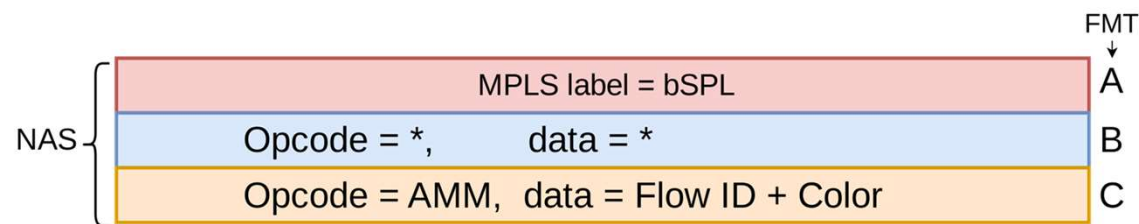


► Performance measurement using the alternate marking method (AMM)

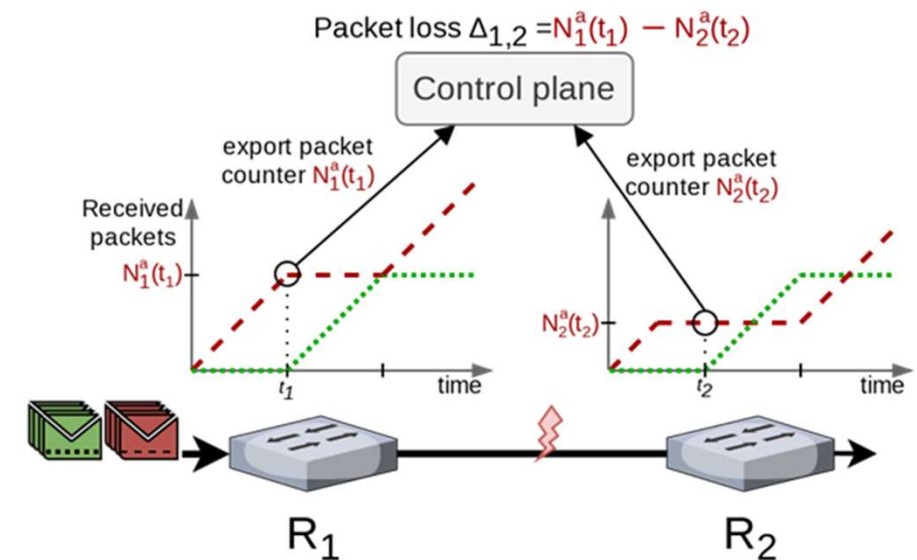
- Measure packet loss between two points without sequence numbers

► Measuring per-link packet loss

1. Send packets in colored batches
 - Color is contained in network action
2. Alternate color between batches
3. Count packets per color per batch at nodes
4. On color change, export per-node counters to control plane
5. → Calculate packet loss



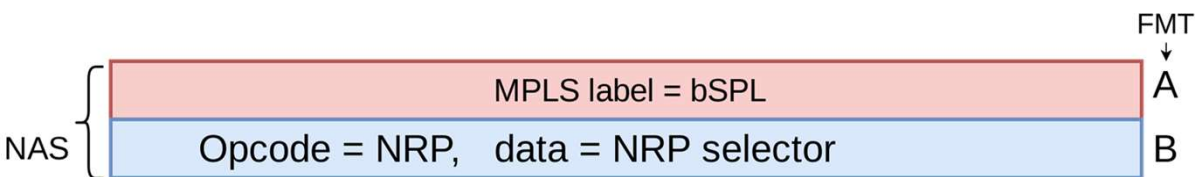
<https://datatracker.ietf.org/doc/draft-cx-mpls-mna-inband-pm/>



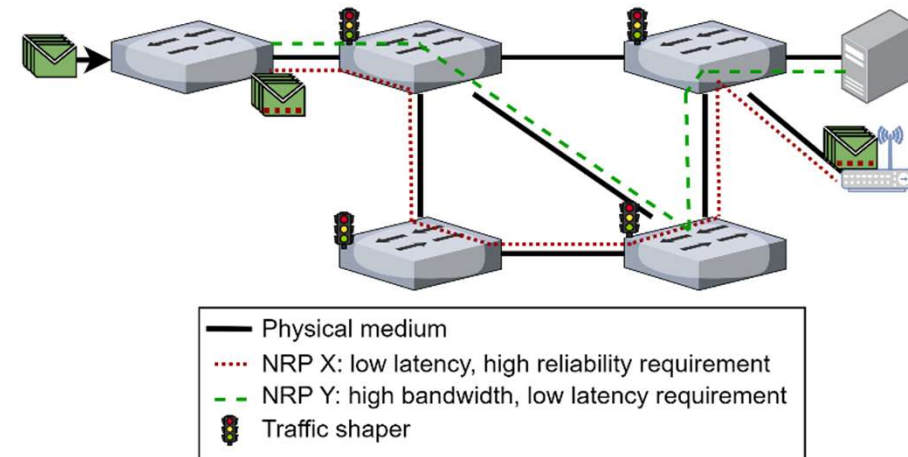


Use Cases – Network Slicing

- ▶ Network slicing
 - End-to-end logical network over a shared physical infrastructure
 - Connectivity coupled with a set of network resource commitments
 - Bandwidth, latency, ...
 - = Network Resource Partition (NRP)
- ▶ NRP selector identifies to which NRP a packet belongs to
 - → Indicated in network action
- ▶ Nodes apply traffic shaping based on configured NRP and NRP selector in packet



<https://www.ietf.org/archive/id/draft-li-mpls-mna-nrp-selector-02.html>

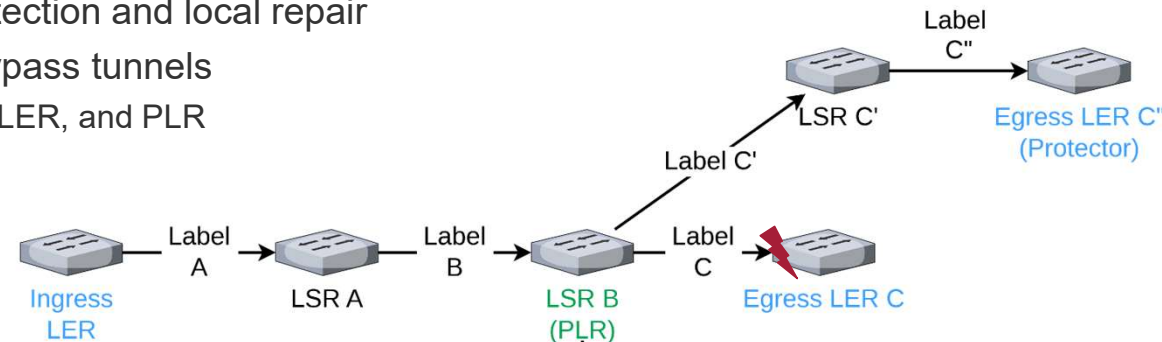




Use Cases – SMEP

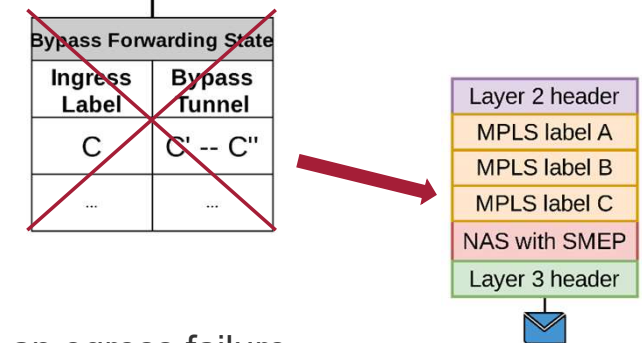
► MPLS Egress Protection Framework (RFC 8679) establishes bypass tunnels for egress routers on an egress failure

- Relies on Point of Local Repair (PLR) for failure detection and local repair
- PLR maintains a mapping of transport tunnels to bypass tunnels
 - Signaled between egress LER, Protector LER, ingress LER, and PLR



► Stateless MNA-based Egress Protection (SMEP)

- Egress bypass tunnels are carried in a network action
 - No mapping from label to bypass tunnel in PLR required → „stateless“
 - Signaling with the PLR is not required anymore
- PLR uses the bypass MPLS label from the network action to reroute traffic on an egress failure



<https://datatracker.ietf.org/doc/draft-ihle-mpls-mna-stateless-egress-protection/>



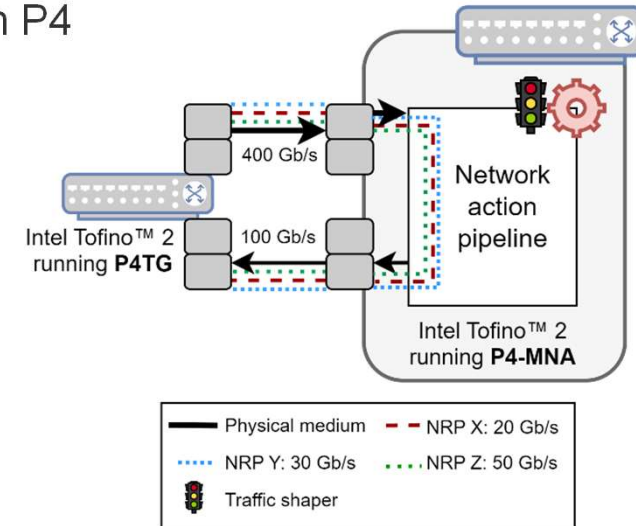
P4-MNA Prototype

► We implemented the MNA framework on the Intel Tofino 2 switching ASIC in P4

- In-stack data
- RLD of 51 LSEs
- Up to 32 network actions in a packet supported
- Line rate of 400 Gb/s

► Implemented use cases

- Packet loss measurement using AMM
- Bandwidth reservation using network slicing



► MPLS Network Actions: Technological Overview and P4-Based Implementation on a High-Speed Switching ASIC”, accepted for IEEE Open Journal of the Communications Society (OJ-COMS), 2025 ([preprint](#))

►  <https://github.com/uni-tue-kn/P4-MNA>

- Source code
- Wireshark dissector for MNA traffic
- Python library for building MNA packets





- ▶ The MPLS Network Actions (MNA) framework introduces a mechanism for transmitting network actions
 - MNA is to MPLS what IPv6 extension headers are to IPv6
 - The MPLS encoding is repurposed to hold network actions in a container (NAS)
 - Network actions can be processed on different nodes indicated by the scope
- ▶ The readable label depth (RLD) is a critical parameter when deploying MNA
- ▶ Use cases include...
 - Per-link packet loss measurement
 - Network slicing
 - Egress protection
- ▶ More information: MPLS Network Actions: Technological Overview and P4-Based Implementation on a High-Speed Switching ASIC”, accepted for IEEE Open Journal of the Communications Society (OJ-COMS), 2025 ([preprint](#))



Any Questions?

MPLS Network Actions: A Technological Overview

Fabian Ihle
University of Tübingen
fabian.ihle@uni-tuebingen.de



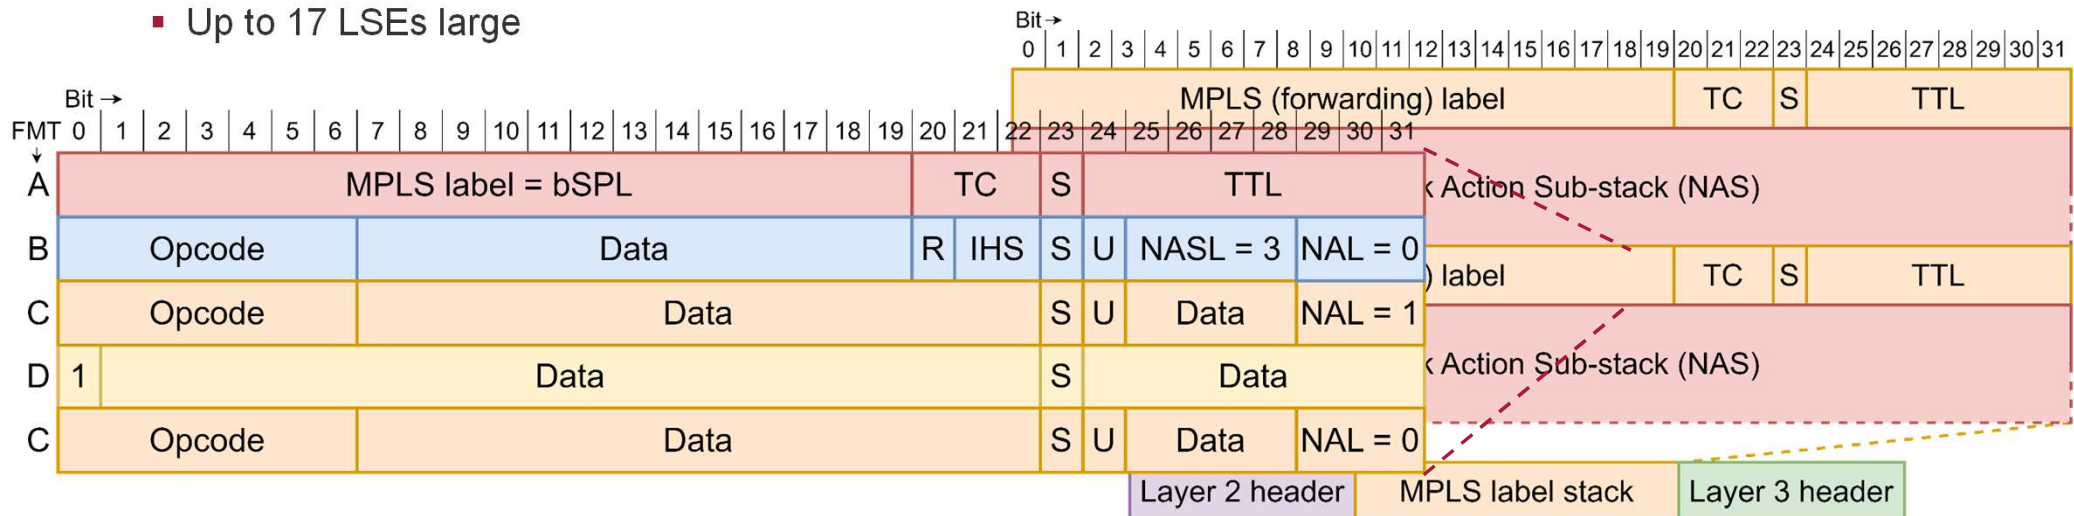
The Network Action Sub-stack (NAS)

► What does a NAS look like?

- Special label value (bSPL) followed by network actions and their data
 - = NAS indicator

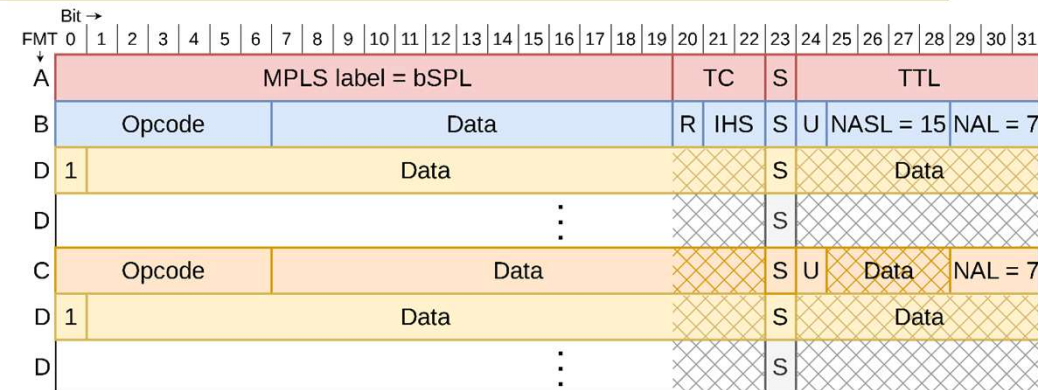
► LSE encoding redefined for network actions and ancillary data in a NAS

- Initial opcode (mandatory, Format B)
- Subsequent opcode (optional, Format C)
- Ancillary data (optional, Format D)
- Up to 17 LSEs large



Mutable Data in MNA

- ▶ Mutable data = data where the value changes during packet forwarding

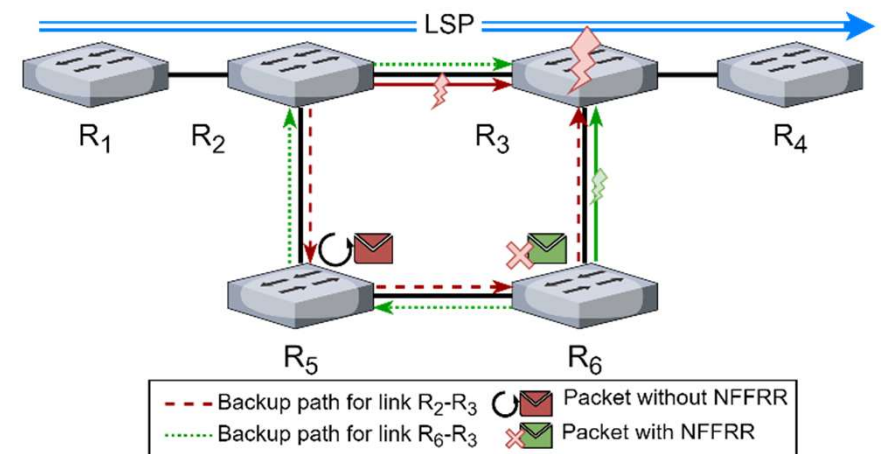


- ▶ Example: Collecting telemetry data along a path
 - Each node on the path adds information to the NAS, e.g., the node ID
 - That information must not be discarded on forwarding
- ▶ For backward compatibility, only last 12 bits are mutable
 - MPLS label may be hashed for ECMP by legacy nodes
- ▶ → A post-stack data implementation may be more feasible for mutable data



Use Cases – NFFRR

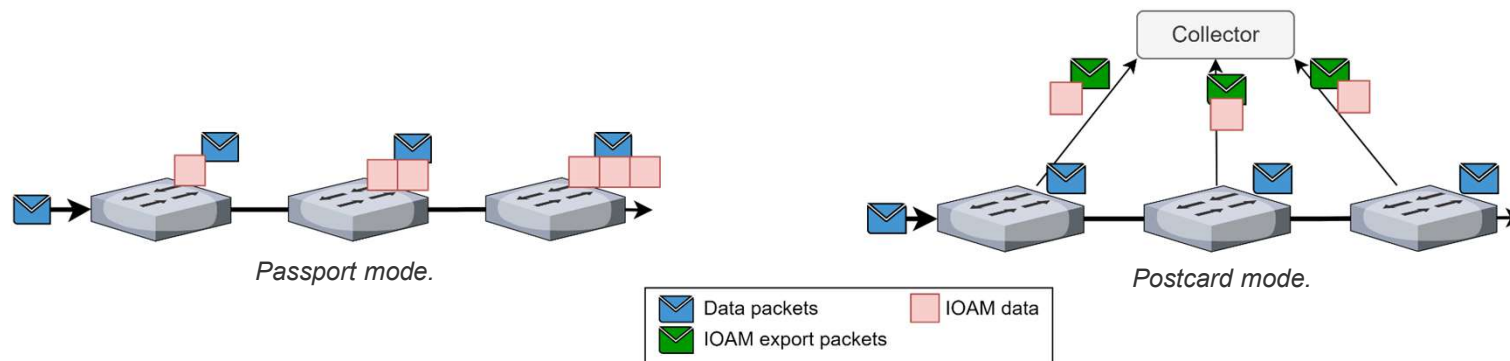
- ▶ No Further Fast Reroute (NFFRR)
 - Multiple failures may lead to looping packets in backup tunnels
 - → Congestion
- ▶ Network action indicates if a packet was rerouted through a backup tunnel
 - If a second reroute will be applied, they can be dropped



<https://datatracker.ietf.org/doc/draft-li-mpls-mna-nffrr/01/>



- ▶ In-situ operations, administration, and maintenance (IOAM)
 - Collect metrics (timestamps, interfaces, ...) along a path
- ▶ Two operation modes
 - Passport: Write metrics into data fields of network actions
 - Postcard: Trigger export to an entity in the network
- ▶ Network action contains mode and information on which metrics to collect
 - E.g., for path tracing



<https://datatracker.ietf.org/doc/draft-ietf-mpls-mna-ioam/>