

MINT

MPLS Inband Network Telemetry

Kireeti Kompella

Jai Kumar

IETF-104

Motivation

Provide simple, protocol agnostic, flexible, inband mechanism to track and collect per flow per packet network telemetry data with low cost and high performance of HW

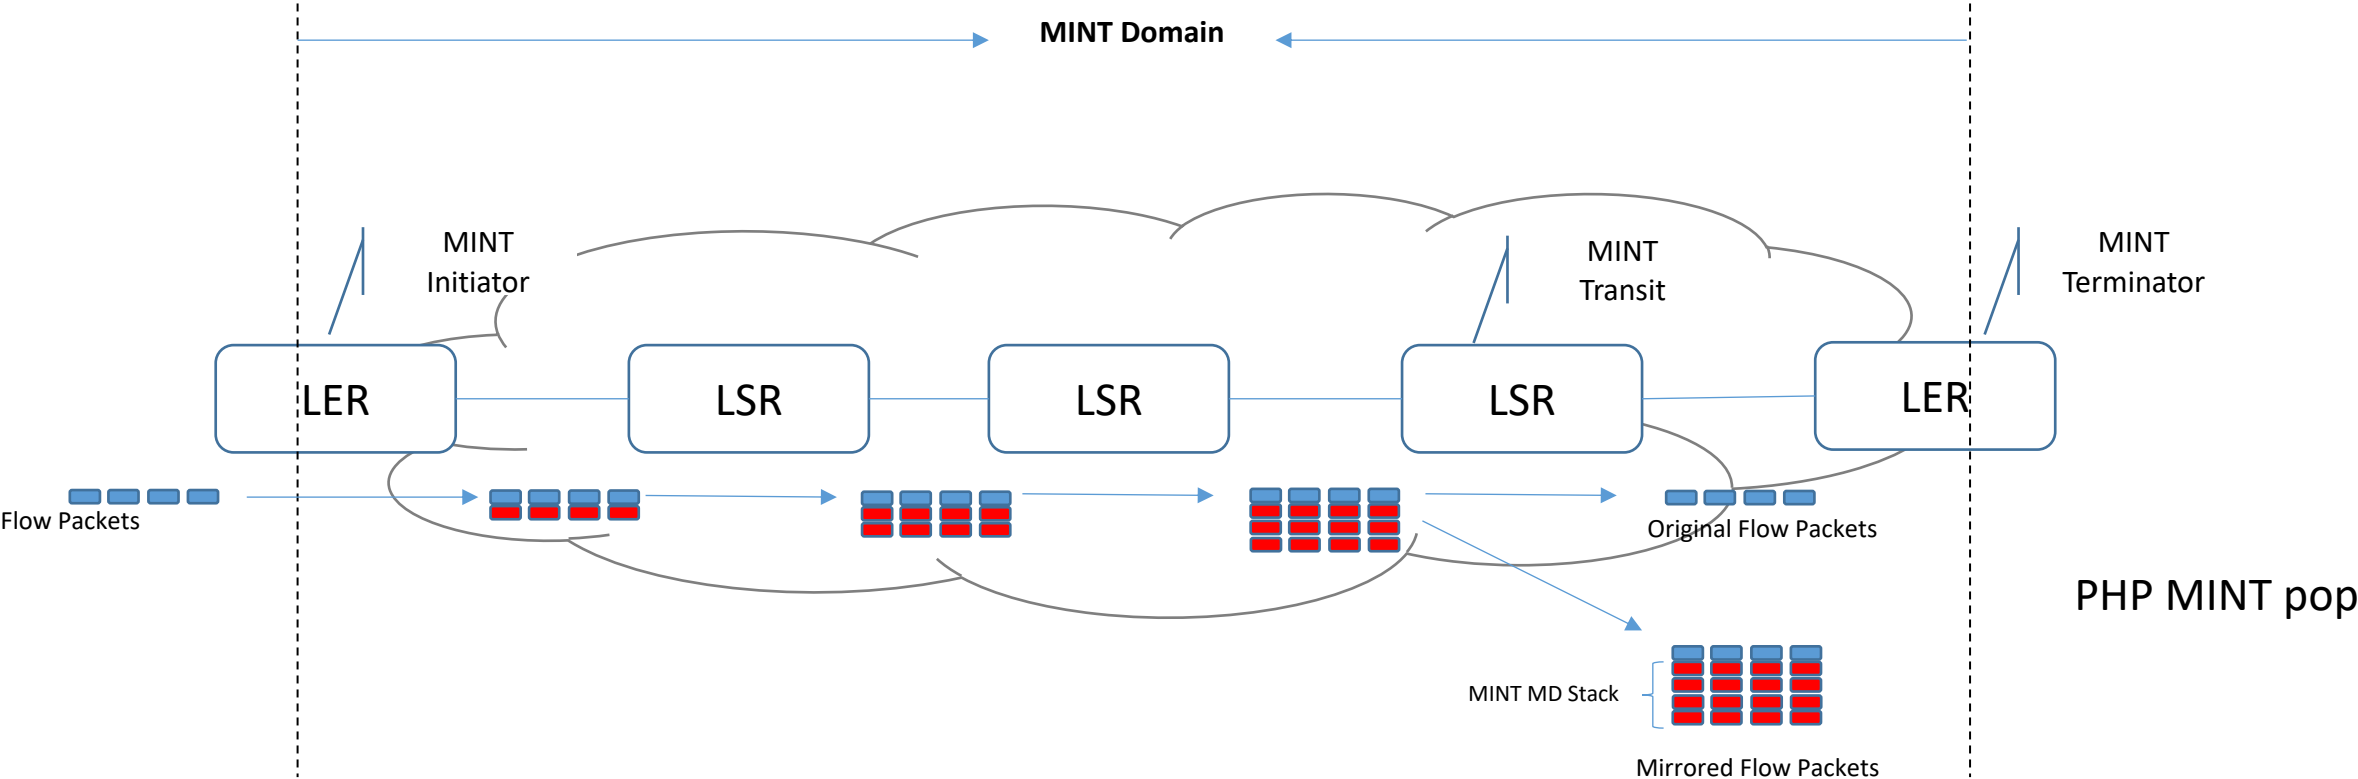
Requirements

- Collect per hop per packet metadata for a given flow
- Metadata is for live packet and instantaneous representation of network state for a given flow
- Use live metadata to solve interesting problems for a given flow
 - Flow E2E latency
 - Congestion on any hop
 - Mapping of Flow to LSP to underlay
 - Per hop resident time
 - Arrive rate
 - Queue congestion
 - Drops

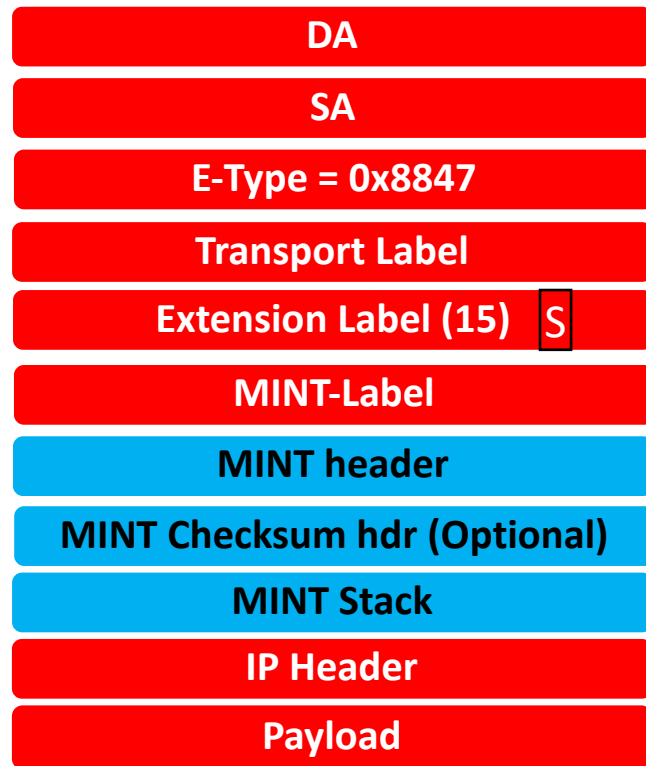
Terminology

- MINT Domain
 - *MPLS domain with all nodes participating in MINT*
- MINT Node
 - *Single node in MINT domain*
- MINT Label
 - *IANA allocated extension labels*
- MINT metadata
 - *Information inserted by each MINT node*
- MINT Initiator
 - *MINT nodes at the edge of MINT domain*
 - *Responsible for inserting the Extension-Label, MINT-Label and MINT metadata*
- MINT Transit
 - *MINT nodes within the MINT domain*
 - *Responsible for detecting the MINT-Label and inserting MINT metadata*
- MINT Terminator
 - *MINT nodes at the edge of MINT domain*
 - *Responsible for detecting the MINT-Label*
 - *Stripping the MINT metadata and forwarding the packet*

MINT Deployment



MINT Packet Format

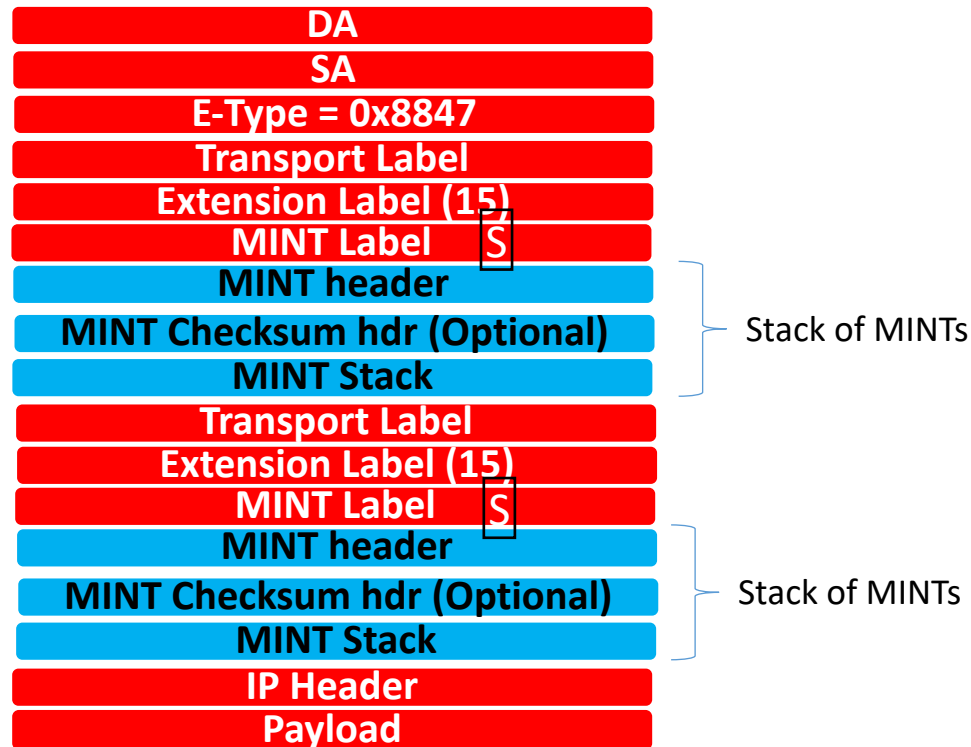


- MINT presence is indicated by MINT label present after the Extension Label (15)
- MINT label is always the last label in a label stack with S (end of stack) bit set
- MINT header contains MINT length and other checks
- MINT checksum is for MINT stack and header only
- MINT header will provide next protocol field and can be used to create a next header chain
- Entropy label can be used for load balancing
- Each Tunnel will have its own MINT stack
 - RSVP over LDP can have RSVP MINT stack over LDP MINT stack



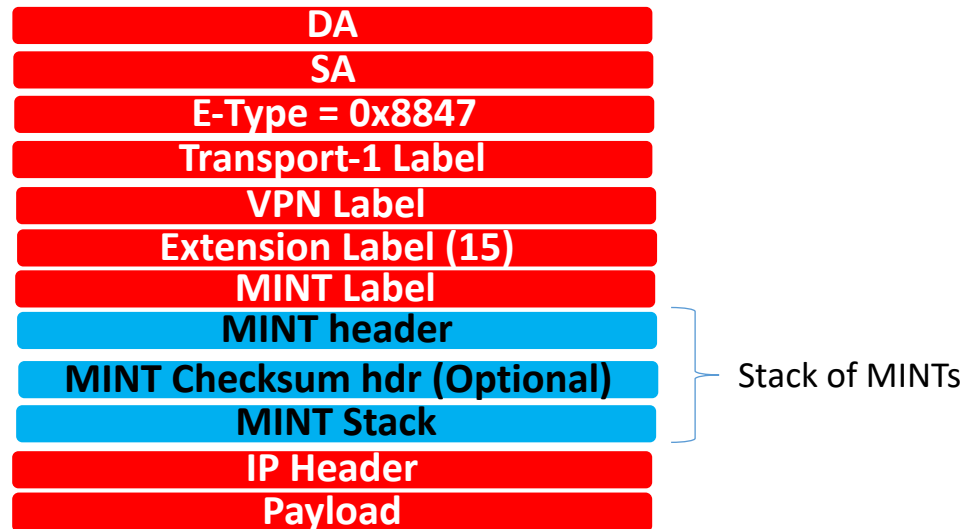
IANA Allocated Extension Labels

Tunnel Mode - MINT Packet Format for Tunnel



- Each tunnel can add Extension label and MINT label for tunnel scoped metadata collection
- Tunnel initiator must add Extension label and MINT label with S bit SET
- Tunnel terminator must dispose Extension label, MINT label and MINT stack
- Transit nodes parse the label stack for only top MINT label for MD insertion
- PHP can be supported for MINT pop so that UHP is not burdened with popping two MINT stacks

Transport Mode - MINT Packet Format for MPLS VPN



- No new Extension Label is inserted for new tunnels.
- Only when all labels are popped Extension label and MINT stack is popped
- Extension label and MINT MAY be popped based on end of MINT zone instead of transport domain termination
- There is a risk of leaking MINT data from one tunnel domain to other if MINT zones are not carefully defined and MINT stack is not popped when packet exits a MINT domain

TBAI Label 1 – MINT

MINT Header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Ver=2.0| GNS  | NextHdr = IP  |      Flags      | Max Length  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

MINT MD Header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Request Vector| Action Vector | Hop Limit.   | Curr Length. |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

MINT MD Fragment Header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                                     Packet ID           | MF ID |L|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```


MINT Metadata Stack:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                                     MINT Metadata Stack
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

NextHdr field reflects the payload type

TBAI Label 3 – RFC8321

Flow-ID Label Header:



- Flow ID: MPLS flow identifier
- L: Loss bit (for coloring)
- D: Delay bit (for coloring)

TBAI Label 1 – MINT chained with GPB

MINT Header:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Ver=2.0|  GNS  | NextHdr = GPB |           Flags           |  Max Length  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                                                                                                 ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|                                     MINT Metadata Stack
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

GPB Header:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|.          Group Policy ID           | RES |E|A|D|Ver| NextHdr = IP  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

- MINT header nexthdr field will point to GPB header

Status of I-D

draft-kumar-ippm-ifa-00

<https://tools.ietf.org/html/draft-kumar-mpls-mint-00>

Request

Would appreciate review and comments

WG Adoption

Thank You