

# Export of MPLS-SR Label Type Information in IPFIX

Enabling insights in MPLS-SR forwarding plane  
by adding Segment Routing dimensions

thomas.graf@swisscom.com

18. February 2021

# MPLS-SR @ IPFIX

## Vendor Status

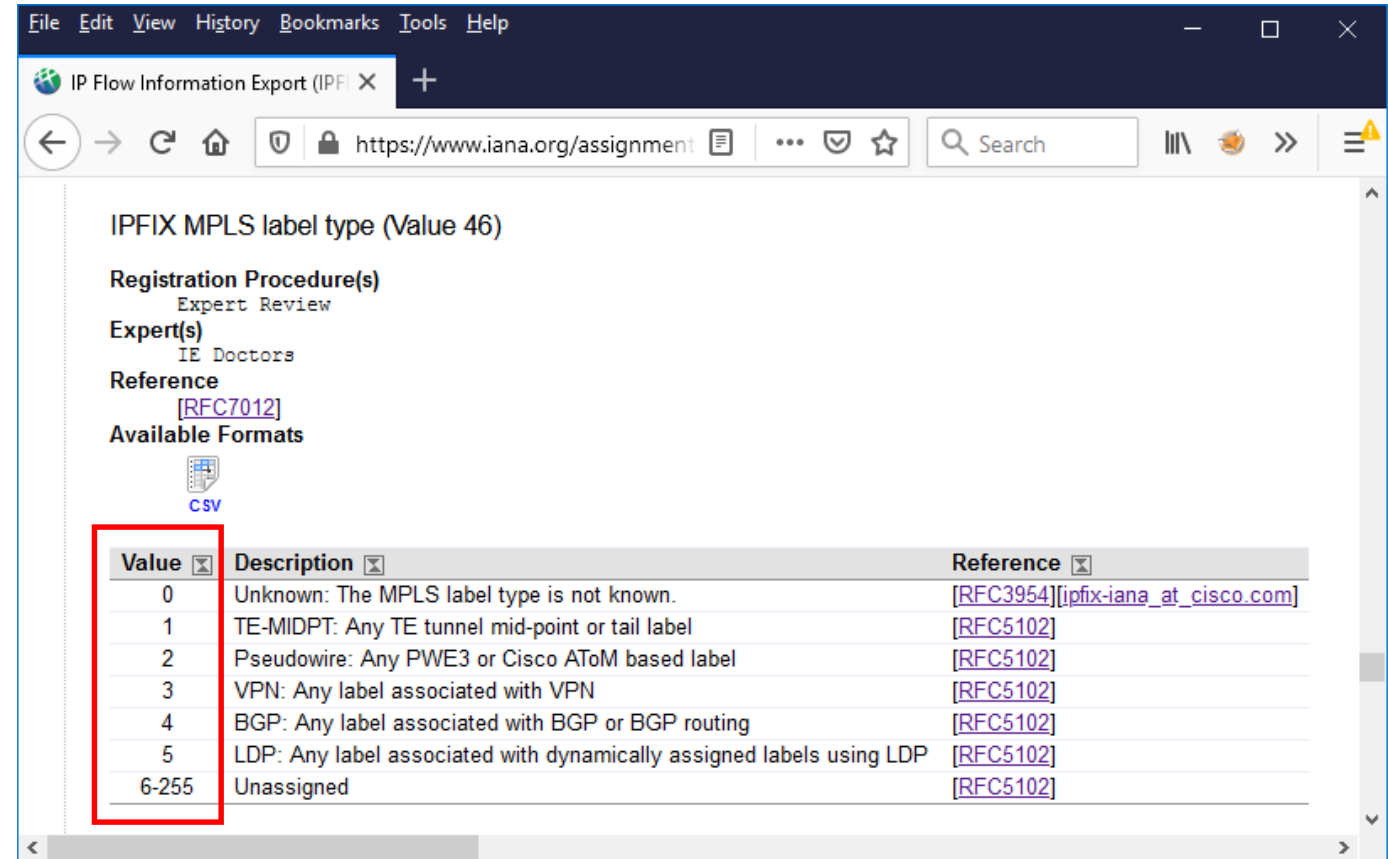
- MPLS-SR uses the existing MPLS data plane.
- Therefore, looking how IPFIX metrics are exposed at a current MPLS-SR vendor implementation we see **not much of a difference to classical MPLS**.
- Looking more deeply, **we notice "not much" is pretty much what is missing**.
- mplsTopLabelType is referencing LDP even though there isn't any LDP anymore. -> Funny

```
> Frame 527: 182 bytes on wire (1456 bits), 182 bytes captured (1456 bits) on interface 0
> Ethernet II, Src: Cisco_ea:ad:1c (00:32:17:ea:ad:1c), Dst: Vmware_21:95:d2 (00:0c:29:21:95:d2)
> Internet Protocol Version 4, Src: 138.187.57.63, Dst: 138.187.58.13
> User Datagram Protocol, Src Port: 44542, Dst Port: 9991
Cisco NetFlow/IPFIX
  Version: 9
  Count: 1
  SysUptime: 516154.381000000 seconds
  > Timestamp: Feb 23, 2020 13:57:18.000000000 W. Europe Standard Time
  FlowSequence: 23685
  SourceId: 0
  > FlowSet 1 [id=313] (1 flows)
    FlowSet Id: (Data) (313)
    FlowSet Length: 120
    [Template Frame: 9]
    > Flow 1
      > MPLS-Label1: 17002 exp-bits: 0
      > MPLS-Label2: 24622 exp-bits: 0 bottom-of-stack
      > MPLS-Label3: 0 exp-bits: 0
      > MPLS-Label4: 0 exp-bits: 0
      > MPLS-Label5: 0 exp-bits: 0
      > MPLS-Label6: 0 exp-bits: 0
      InputInt: 87
      OutputInt: 111
      Octets: 216000
      Packets: 2000
      > [Duration: 5.753000000 seconds (switched)]
      TopLabelAddr: 138.187.57.13
      SrcAddr: ::
      DstAddr: ::
      ipv6FlowLabel: 0
      IPv6 Extension Headers: 0x00000000
      SrcAddr: 10.248.4.236
      DstAddr: 10.248.4.222
      SrcPort: 0
      DstPort: 2048
      MPLS Top Label Prefix Length: 32
      TopLabelType: LDP (5)
      > Forwarding Status
      Direction: Ingress (0)
      IP ToS: 0x00
      Protocol: ICMP (1)
      > TCP Flags: 0x00
      SamplerID: 1
      Ingress VRFID: 1610612736
      Egress VRFID: 1610612736
      Padding: 0000
```

# MPLS-SR @ IPFIX

## IANA code points

- Looking at IANA makes it clear, there is **NO** mplsTopLabelType code point for BGP, IS-IS, OSPFv2 and OSPFv3 Segment Routing.



IPFIX MPLS label type (Value 46)

Registration Procedure(s)  
Expert Review

Expert(s)  
IE Doctors

Reference  
[\[RFC7012\]](#)

Available Formats

[CSV](#)

Value	Description	Reference
0	Unknown: The MPLS label type is not known.	<a href="#">[RFC3954]</a> <a href="#">[ipfix-iana_at_cisco.com]</a>
1	TE-MIDPT: Any TE tunnel mid-point or tail label	<a href="#">[RFC5102]</a>
2	Pseudowire: Any PWE3 or Cisco AToM based label	<a href="#">[RFC5102]</a>
3	VPN: Any label associated with VPN	<a href="#">[RFC5102]</a>
4	BGP: Any label associated with BGP or BGP routing	<a href="#">[RFC5102]</a>
5	LDP: Any label associated with dynamically assigned labels using LDP	<a href="#">[RFC5102]</a>
6-255	Unassigned	<a href="#">[RFC5102]</a>

# MPLS-SR @ IPFIX

## IANA Reference vs. Requester

- Two IPFIX doctor reviews, revealed that requester column is hidden at IANA and content in reference column should be in requester column and vice versa.
- IANA and IE doctors confirmed that IE46 registry will be fixed soon so that new code points can be added with this draft.

IPFIX MPLS label type (Value 46)

Registration Procedure(s)  
Expert Review

Expert(s)  
IE Doctors

Reference  
[\[RFC7012\]](#)

Available Formats  
CSV

Value	Description	Reference
0	Unknown: The MPLS label type is not known.	<a href="#">[RFC3954]</a> <a href="#">[ipfix-iana_at_cisco.com]</a>
1	TE-MIDPT: Any TE tunnel mid-point or tail label	<a href="#">[RFC5102]</a>
2	Pseudowire: Any PWE3 or Cisco AToM based label	<a href="#">[RFC5102]</a>
3	VPN: Any label associated with VPN	<a href="#">[RFC5102]</a>
4	BGP: Any label associated with BGP or BGP routing	<a href="#">[RFC5102]</a>
5	LDP: Any label associated with dynamically assigned labels using LDP	<a href="#">[RFC5102]</a>
6-255	Unassigned	<a href="#">[RFC5102]</a>

Untitled - Notepad

Value	Description	Reference	Requester
0	Unknown: The MPLS label type is not known.	RFC3954	RFC3954
1	TE-MIDPT: Any TE tunnel mid-point or tail label	RFC4736	RFC5102
2	Pseudowire: Any PWE3 or Cisco AToM based label	RFC3985	RFC5102
3	VPN: Any label associated with VPN	RFC4364	RFC5102
4	BGP: Any label associated with BGP or BGP routing	RFC8277	RFC5102
5	LDP: Any label associated with dynamically assigned labels using LDP	RFC5036	RFC5102

# MPLS-SR @ IPFIX

## Draft Status

- Feedback collected from SPRING, MPLS, LSR and OPSAWG lists.
  - Called for adoption at IETF 108 OPSAWG. Piano feedback received with comment to review SrSidType and fix existing IE46 registry.
  - SrSidType removed from -05 version (backup slides) in favor of draft-ietf-spring-sr-yang-23#section-8.3
  - Two IE-DOCTOR reviews received and IE46 registry correction submitted to IANA and currently pending.
  - Presented at MPLS at IETF 109 and received positive feedback.
- > Call for adoption at OPSAWG at IETF 110 again**

thomas.graf@swisscom.com

18. February 2021

MPLS-SR @ IPFIX

Draft Status

# Backup

thomas.graf@swisscom.com

18. February 2021

# MPLS-SR @ IPFIX

## RFC 8402, SID's, SID's

- Segment Routing is all about SID's.
- An Adjacency-SID can be used by TI-LFA or uLoop avoidance to use a different path to the Prefix SID than what the routing protocol calculated as best path.
- Where are the SID's in IPFIX?  
-> **Nowhere!**

### Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">Link-State IGP Segments</a>	<a href="#">9</a>
<a href="#">3.1.</a>	<a href="#">IGP-Prefix Segment (Prefix-SID)</a>	<a href="#">9</a>
<a href="#">3.1.1.</a>	<a href="#">Prefix-SID Algorithm</a>	<a href="#">9</a>
<a href="#">3.1.2.</a>	<a href="#">SR-MPLS</a>	<a href="#">10</a>
<a href="#">3.1.3.</a>	<a href="#">SRv6</a>	<a href="#">12</a>
<a href="#">3.2.</a>	<a href="#">IGP-Node Segment (Node-SID)</a>	<a href="#">13</a>
<a href="#">3.3.</a>	<a href="#">IGP-Anycast Segment (Anycast-SID)</a>	<a href="#">13</a>
<a href="#">3.3.1.</a>	<a href="#">Anycast-SID in SR-MPLS</a>	<a href="#">13</a>
<a href="#">3.4.</a>	<a href="#">IGP-Adjacency Segment (Adj-SID)</a>	<a href="#">15</a>
<a href="#">3.4.1.</a>	<a href="#">Parallel Adjacencies</a>	<a href="#">17</a>
<a href="#">3.4.2.</a>	<a href="#">LAN Adjacency Segments</a>	<a href="#">18</a>
<a href="#">3.5.</a>	<a href="#">Inter-Area Considerations</a>	<a href="#">18</a>
<a href="#">4.</a>	<a href="#">BGP Segments</a>	<a href="#">19</a>
<a href="#">4.1.</a>	<a href="#">BGP-Prefix Segment</a>	<a href="#">19</a>
<a href="#">4.2.</a>	<a href="#">BGP Peering Segments</a>	<a href="#">20</a>
<a href="#">5.</a>	<a href="#">Binding Segment</a>	<a href="#">21</a>
<a href="#">5.1.</a>	<a href="#">IGP Mirroring Context Segment</a>	<a href="#">21</a>
<a href="#">6.</a>	<a href="#">Multicast</a>	<a href="#">22</a>
<a href="#">7.</a>	<a href="#">IANA Considerations</a>	<a href="#">22</a>
<a href="#">8.</a>	<a href="#">Security Considerations</a>	<a href="#">22</a>
<a href="#">8.1.</a>	<a href="#">SR-MPLS</a>	<a href="#">22</a>
<a href="#">8.2.</a>	<a href="#">SRv6</a>	<a href="#">24</a>
<a href="#">8.3.</a>	<a href="#">Congestion Control</a>	<a href="#">25</a>
<a href="#">9.</a>	<a href="#">Manageability Considerations</a>	<a href="#">25</a>
<a href="#">10.</a>	<a href="#">References</a>	<a href="#">26</a>
<a href="#">10.1.</a>	<a href="#">Normative References</a>	<a href="#">26</a>
<a href="#">10.2.</a>	<a href="#">Informative References</a>	<a href="#">27</a>
	<a href="#">Acknowledgements</a>	<a href="#">30</a>
	<a href="#">Contributors</a>	<a href="#">31</a>
	<a href="#">Authors' Addresses</a>	<a href="#">32</a>

# MPLS-SR @ IPFIX

## draft-tgraf-ipfix-mpls-sr-label-type

- Segment Routing adds the source routing paradigm to MPLS and enhances IGP routing protocol to carry label information.
- Let's bring visibility into **how Segment Routing applications change the MPLS forwarding plane**.
- "Show me all MPLS-SR controlled traffic where Adj-SID's were used, group by Label Stack, and show for each through which nodes and interfaces it was forwarded."
- Fill the missing gaps at IPFIX:
  - Update **mplsTopLabelType**
  - Introduce **SrSidType**

