

MPLS Encapsulation for SFC NSH

draft-malis-mpls-sfc-encapsulation-01

Andrew G. Malis, Huawei

Stewart Bryant, Huawei

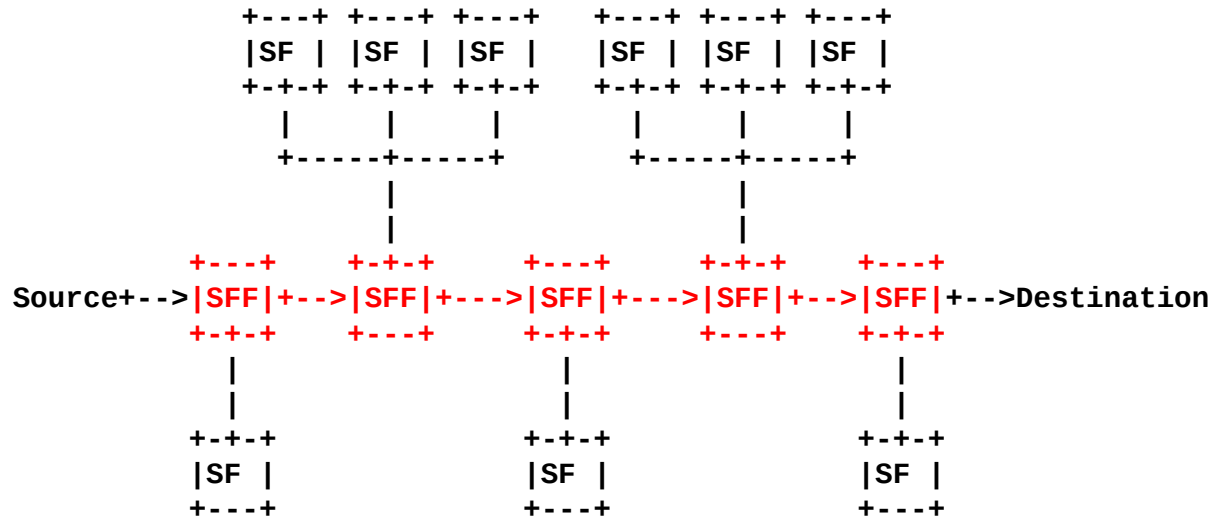
Joel Halpern, Ericsson

Wim Henderickx, Nokia

IETF 102, Montreal

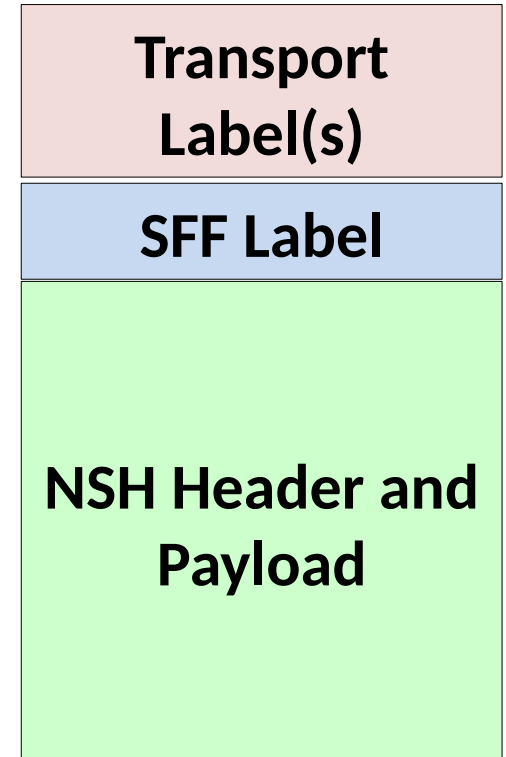
Basic Intention of the Draft

- Defines an encapsulation used to transport SFC packets that use the NSH from one SFF to the next SFF over an MPLS infrastructure
- Red indicates MPLS-enabled nodes (LSRs) and links in the diagram (adopted from RFC 7665)



Encapsulation Details

- SFF Label is very similar to a PW, VPN, or other service label
 - Like a VPN label, no control word
- SFF Label identifies the SFF instance at the downstream LSR
 - Allows more than one SFF instance at the downstream LSR
 - SFF label(s) advertised by downstream LSR to upstream LSR (standard MPLS label advertisement)
 - Label advertisement mechanism(s) could include LDP, RSVP, YANG, BGP, PCEP, etc. (see later slide)
- Because there can be multiple transport labels, works with MPLS-based Segment Routing (SR-MPLS)
 - Works with draft-guichard-sfc-nsh-sr



ECMP Considerations

- ECMP forwarding through the MPLS infrastructure may or may not be desirable for a particular SFC flow
 - Should be avoided for flows that require in-order delivery
- First nibble of the NSH provides protection to prevent unintended ECMP
 - Never equal to 0100 or 0110 to avoid appearing like an IP header to ECMP logic
- If ECMP is desired, MPLS has native mechanisms to provide entropy
 - Entropy label (RFC 6790)
 - Flow-aware transport label (RFC 6391)
- A recommendation between these options is for future study

OAM Considerations

- OAM at the SFC layer is handed by SFC-defined mechanisms (see RFC 8300)
- OAM may be required at the MPLS layer
- If so, standard MPLS-layer OAM mechanisms may be used, such as GAL (RFC 5586)

Comparison with draft-ietf-mpls-sfc

- draft-malis-mpls-sfc-encapsulation transports SFC packets with the NSH between SFFs over an MPLS infrastructure
 - Supports ALL SFC features, including per-packet metadata
- draft-ietf-mpls-sfc uses the MPLS label stack to “logically represent” the NSH for interim deployments in an MPLS infrastructure that doesn’t support the NSH
 - No NSH in packets
 - Encodes the SFC Service Path Indicator and Service Index as “labels” in the label stack
 - SPI and SI labels require processing different from normal label operations, see Section 6
 - No per-packet metadata, only per-flow metadata
 - Metadata requires control plane extensions or a new MPLS special purpose label that carries the metadata in a dedicated packet

Comparison with draft-xuclad-spring-sr-service-programming

- draft-malis-mpls-sfc-encapsulation transports SFC packets with the NSH between SFFs over a general MPLS infrastructure
 - Supports both traditional label swapping and SR-MPLS
 - The usual MPLS state (LIB, etc.) at every LSR when label swapping
 - Intended for SFC infrastructures; SFC NSH is present in every packet
- draft-xuclad-spring-sr-service-programming is intended to support generalized service programming in SR domains
 - Services are associated with SIDs
 - More general than SFC in that “services” could be more than just service functions as defined by the SFC WG
 - Works with both SR-MPLS and SRv6
 - Doesn’t support MPLS label swapping, no MPLS state in the routers
 - NSH is available, using the NSH Carrier TLV, if using standard SFC-defined SFs

Next Steps

- Progress the FFS items
 - ECMP recommendation
 - Control plane for SFF label advertisement
 - Already been in touch with authors of draft-ietf-bess-nsh-bgp-control-plane to discuss adding MPLS labels to the Service Function Instance Route to support this draft
 - Other control plane options?
- Start working towards adoption in the MPLS WG