An MPLS-Based Forwarding Plane for Service Function Chaining

draft-farrel-mpls-sfc-04/05

Stewart Bryant <stewart.bryant@gmail.com>
Adrian Farrel <afarrel@juniper.net>
John Drake <jdrake@juniper.net>

IETF-101, London, March 2018

Agenda

- Overview and (Non-)Objectives of the Design
- Issues and Changes
  - Nits and Editorial
  - Removal of Discussion of Segment Routing
  - Purpose and Intent
  - Transport Independence
  - SFC-Awareness and SFC Proxies
  - Metadata
  - Control Planes
- Future Plans

Recall the SFC Architecture

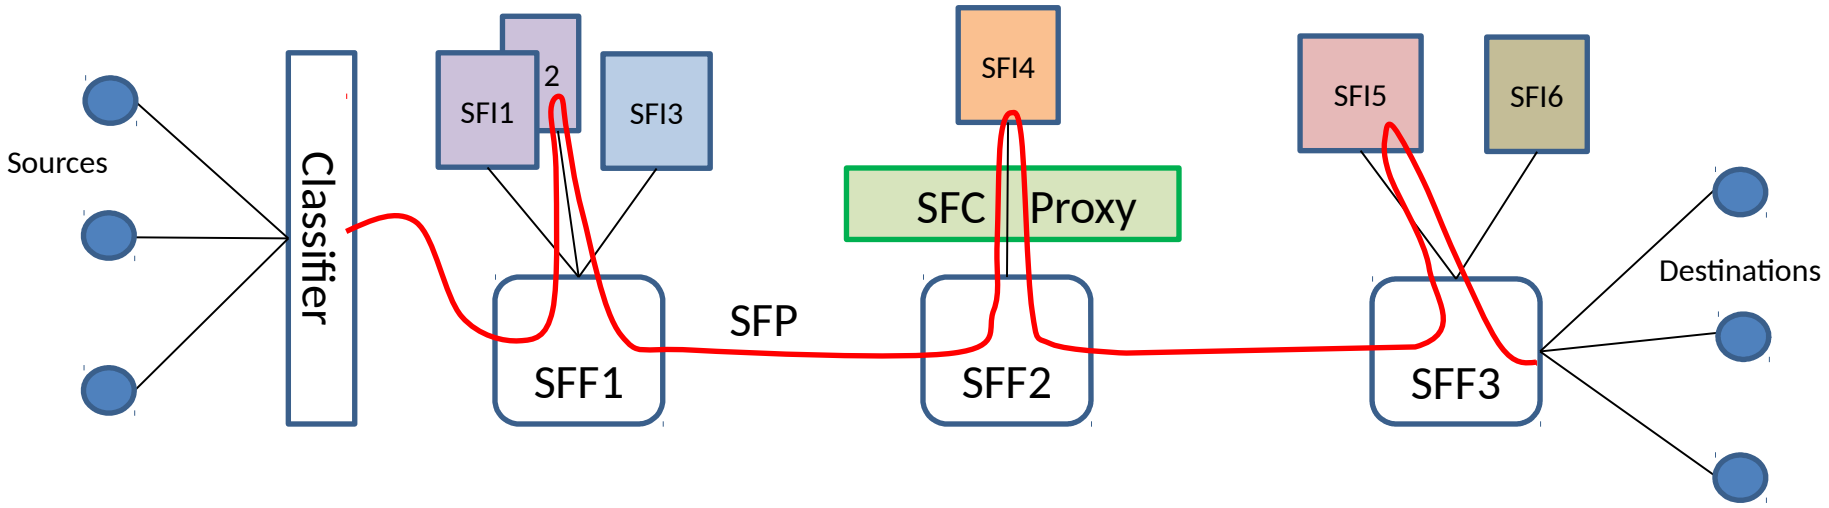

- Packets flow from source to destination
- Packets are classified onto a Service Function Path (SFP)
- SFP traverses a series of Service Function Forwarders (SFFs)
- Each SFF delivers packets on the SFP to a specific Service Function Instance (SFI)
- SFC Proxy may be placed between SFF and SFI

Objectives / Non-objectives

1. Not trying to replace or obsolete NSH
2. Looking at a specific environment where deployed MPLS routers can serve as SFFs
   - No change to forwarding plane
   - Work using existing MPLS forwarding operations (push/pop/swap)
   - Able to forward SFC packets “at line speed”
3. Aim to get high level of SFC functionality
   - Possible that some features will be sacrificed in compromise with desire to achieve objectives
   - Must support SFC architecture (RFC 7665)
   - Should support metadata
   - Try to integrate with control plane solutions that work with NSH
     - draft-ietf-bess-nsh-bgp-control-plane

Overview of Solution

- Basic building block is a two-label unit

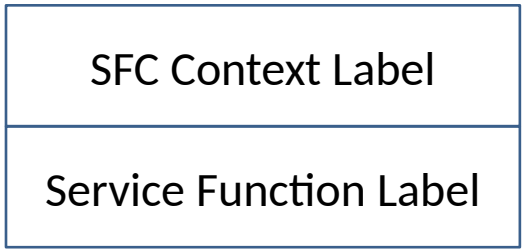

  - Labels included as Label Stack Entries
  - Neither of the labels can be in the range 0..15
    - I.e., must not overlap with Special Purpose Label values
  - An SFF uses top label to identify “path”
    - Local context
    - Select path to next SFF
  - An SFF uses second label to identify SF

Nits and Editorial

- Changes from -02 to -04
  - RFC 8300 published
  - RFC 8174 published
  - MPLS S-bit (not S-flag or S-field)
  - Abbreviation expansions
  - Add a section on Proxies (see later slide)
  - Clarify metadata usage is less functional than NSH
    - (see later slide)
  - Typos in examples

Discussion of Segment Routing

- Called out on mailing list
  - Resulted in Adrian's mea culpa email
- -05 will
  - Remove all discussion of SR (specifically MPLS-SR)
  - Talk only about the MPLS forwarding plane as already defined
    - push, pop, and swap
  - Not discuss control plane mechanisms in any detail
- Continue to discuss
  - Use of labels to encode information included in NSH
  - How to handle metadata with labels
- Where to discuss this draft?
  - Seems to leave the document as an MPLS draft
  - With necessary SFC review

Purpose and Intent

- As stated in objectives
  - SFFs built from existing (MPLS) routers
    - Able to forward packets at line speed
  - Functionally of 7665 and 8300
  - Authors think this will provide migration assist
    - Experience with SFC
    - Gateway to use of control plane and other tools
    - Easy way to introduce SFC to today's network
- Debate over whether such an SFF could exist
  - Should authors describe how to do that? Or is that secret-sauce?
- Non-objective
  - Obsolete or modify NSH
- -05 will
  - Add/clarify text on objectives and non-objectives

Transport Independence

- SFC transport means:
  - Between SFFs
    - NSH is transport independent
      - This draft shows MPLS as the transport
        - » This is the most likely use case for this work
  - Between SFF and SF
    - NSH and this document are transport independent
    - See also discussion of proxies on next slide
    - SFs are usually Ethernet/VxLAN/PW attached?
- Nothing proposed for -05
  - Anything needed?

SFC-Awareness and SFC Proxies

- “SFC-Aware” means “able to handle the SFC encapsulation”
- SFFs
  - Usually SFC-aware, but…
  - Might be programmed with label forwarding/operations
    - E.g. “pop and forward”, “incoming port maps to label imposition”
- SFs
  - Legacy VNFs and PNFs are not SFC-aware by definition
  - Must use an SFC Proxy
    - Strip encapsulation
    - Pass to SF
    - Receive from SF (on logical port)
    - Impose encapsulation
  - NSH and MPLS encodings have identical requirements and issues
- -04 introduced Section 8 on proxies
  - -05 clarifies and provides early pointer to Section 8

Metadata

- Document acknowledges it cannot do everything that NSH can do
  - Not carry metadata in user data packet
    - Cannot do per-packet metadata
  - Use same technique as draft-farrel-sfc-convent
    - (On RFC Editor queue with SFC WG consensus)
    - Send metadata in dedicated packets following the SFP
    - Point to metadata from packet using label
- Technique is not seamless
  - SFF can forward metadata just as user data
  - SFC Proxy must map metadata as SF is not MPLS-aware
    - But this is exactly how SFC Proxy must behave for all metadata
- Draft already includes explanation
  - No changes planned for -05 but welcome input

Control Planes

- This document does not depend on any control plane
  - But a control plane will probably be needed
- Want a YANG model?
  - Write one, probably in SFC WG
- Like SR?
  - Probably in SPRING where Xiaohu Xu has a draft
- Want to use BGP?
  - See draft-ietf-bess-nsh-bgp-control-plane
- Legacy world?
  - See draft-ietf-bess-service-chaining
    - This is a BGP VPN approach
    - Popular way to introduce the technology
    - SFP is achieved by programming SFFs (i.e., not following SFC WG)

Next Steps

- There are always things to polish, but…
  - This is now relative stable
  - Support for swapping and stacking in a common way took some effort, but has good benefits
- Fits with BESS control plane work
- To the authors approach seems “obvious”
  - What do other people think?
- The authors think this is in charter for MPLS WG
  - Use of special purpose labels belongs in MPLS
  - But **obviously** it needs review by SFC WG
- Actions for chairs
  - Decide where this belongs
  - Resolve adoption issues
- Actions for participants
  - Objective discussion of the design.

Backup Slides

Where To Have This Discussion?

- SFC WG has developed problem statement and architecture for SFC
  - We re-use these
- SFC chartered to work on “generic encapsulation” that is “agnostic to the
  layer at which it is applied”
  - Has developed the NSH
- This work is specific to an MPLS forwarding plane and uses an MPLS encapsulation
  - Need review from experts
    - Want to be sure MPLS parts work
    - Want to be sure SFC parts work
  - Some functions need specific MPLS extensions and codepoints
- Let the chairs and ADs work out where the work belongs

MPLS Label Swapping

- Tunnels between SFFs “as normal”
  - Of course, we are interested in MPLS as the transport
- SPI and SI used “as normal” for NSH
  - Some limitation as SPI is constrained here to 20 bits

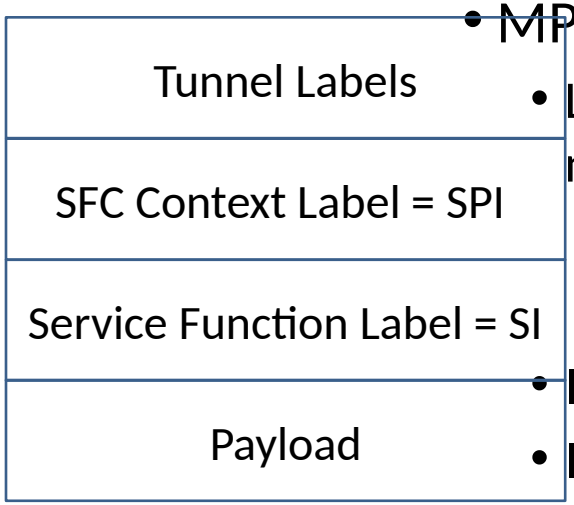

- MPLS-SFC processing…
- Labels are looked up and acted on by SFF to determine next hop
  - Maybe forward to SFI or SFC proxy
  - Maybe forward to next SFF
- In some cases action can be achieved simply through SPI
- In other cases need the two label context
- SI is updated before further forwarding (it’s a swap)
- SPI and SI set during classification
  - Potentially also during re-classification

Metadata

- MPLS encapsulation not well suited for carrying “arbitrary” metadata
- We define an Extended Special Purpose Label
  - This three-label sequence can be included at the bottom of the label stack
  - Metadata label is an index into a store of metadata
    - Must also not use 0..15
  - Store may be populated though management plane, control plane, or in-band (next slide)
    - This approach is not good for “per-packet metadata” (e.g., hashes)
    - Works fine for per-SFP or per-flow metadata

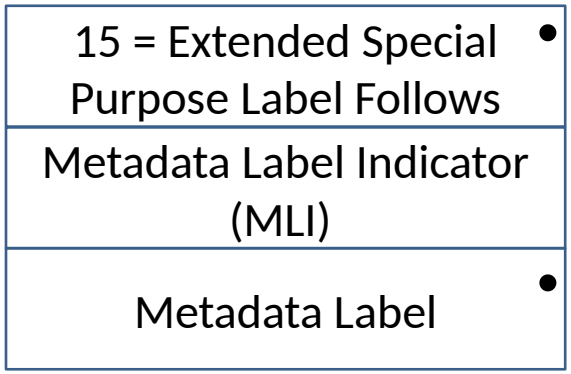

In-Band Metadata Distribution

- Consider draft-farrel-sfc-convent
  - Defines use of NSH with Next Protocol == None
  - Can be used to send NSH packets along an SFP without carrying payload (but still carrying metadata)
  - This draft defines how to do this in MPLS

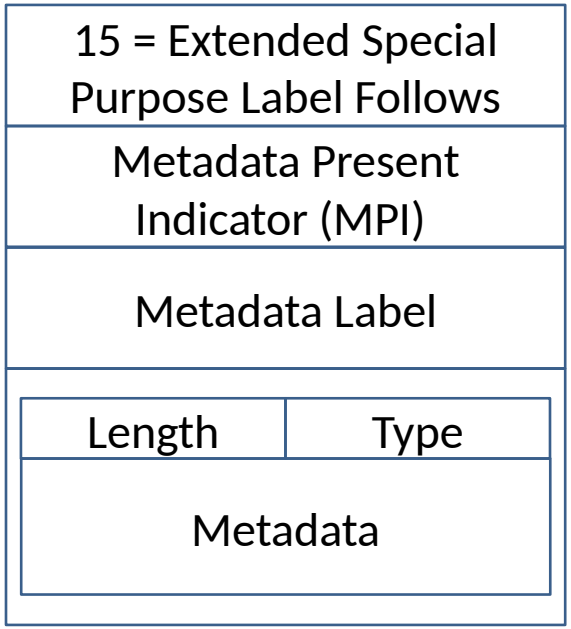

- Use an Extended Special Purpose Label
  - Hence, a three label sequence
- Placed at the bottom of the label stack
- Rest of stack exactly as for SFP
- Metadata carried as payload
  - Formatted as TLV
  - Type field defined by SFC WG for NSH
  - Metadata as defined by SFC WG