

# A BGP-Based Control Plane for Service Function Chaining

draft-mackie-bess-nsh-bgp-control-plane-01

Adrian Farrel (adrian@olddog.co.uk)  
J. Drake (jdrake@juniper.net)  
E. Rosen (erosen@juniper.net)  
J. Uttaro (ju1738@att.com)  
L. Jalil (luay.jalil@verizon.com)

# Introduction

- SFC is (obviously) chartered to work on Service Function Chaining
  - BUT...  
All protocol extension work resulting from these requirements should be carried out in the working group responsible for the protocol being modified in coordination with this working group, but may be done in this working group under a revised charter
- BESS is chartered to work on BGP extensions to support SFC
- So clearly...
  - This draft is targeted at BESS
  - SFC must be given amply opportunity to review it
  - Any proposed changes to the architecture or forwarding plane must come and be worked on in SFC
- Please be nice!
  - The SFC chairs asked me to come and present this in SFC
  - This document was first posted on October 14th
  - The BESS WG hasn't even discussed it yet

# Objectives

- Use BGP to
  - Discover SFFs
  - Learn what SFs are supported by each SFF
  - Distribute information about complete SFPs
- Make those two functions separable
- Re-use BGP VPN methodology and lessons learned
- Fully support SFC architecture and NSH protocol
  - RFC 7665
  - draft-ietf-sfc-nsh
- Be guided by draft-ietf-sfc-control-plane
- Allow flexible, high-function implementations and deployments
- Support multiple SFC overlay networks on a common underlay

# SFF and SF Discovery

- Each SFF speaks BGP
  - Probably to a Route Reflector
- Each SFF supports access to one or more SFs
  - It issues an advertisement for each SF
    - Indication of to which overlay network the SF belongs (the BGP Route Target)
    - A type (from an IANA registry) for the SF (the SFT)
    - A unique identifier for this SFI advertisement (the BGP route distinguisher)
- This allows:
  - A controller to know about the existence of all SFIs in the overlay network
  - A controller to build an SFP listing the desired SFIs
  - Each SFF to know how to route to any SFI on an SFP
    - I.e., by knowing which SFF to route to reach an SFI

# Service Function Instance Route (SFIR)

- Each SFF advertises for each SFI to which it provides access
  - Allows other SFFs to know how to route to the advertised SFI
  - Allows controller to see all available SFIs
- Advertisement contains
  - Route Target
    - Identifies the overlay network
    - Other nodes only import when the RT matched
  - Route Distinguisher (SFIR-RD)
    - Identifies this SFI advertisement
  - SF Type (SFT)
    - From the FCFS IANA registry
- The combination SFIR-RD/SFT uniquely identifies a specific SFI

# SFP Advertisement

- When a Controller has constructed an SFP it must tell the SFFs
  - That they are on the SFP
  - Where to forward packets for any given SI
- For each SFP the Controller advertises:
  - Indication of to which overlay network the SFP belongs (the BGP RT again)
  - A unique identifier for this advertisement of the SFP (the BGP RD)
  - The SPI to use for this SFP in the NSH
  - A series of hops
- Each hop in the advertised SFP contains
  - SFT of the SF to be executed
  - The SFI to use (indicated by the RD of its advertisement)
  - The SI for the hop and for use in the NSH

# Service Function Path Route (SFPR)

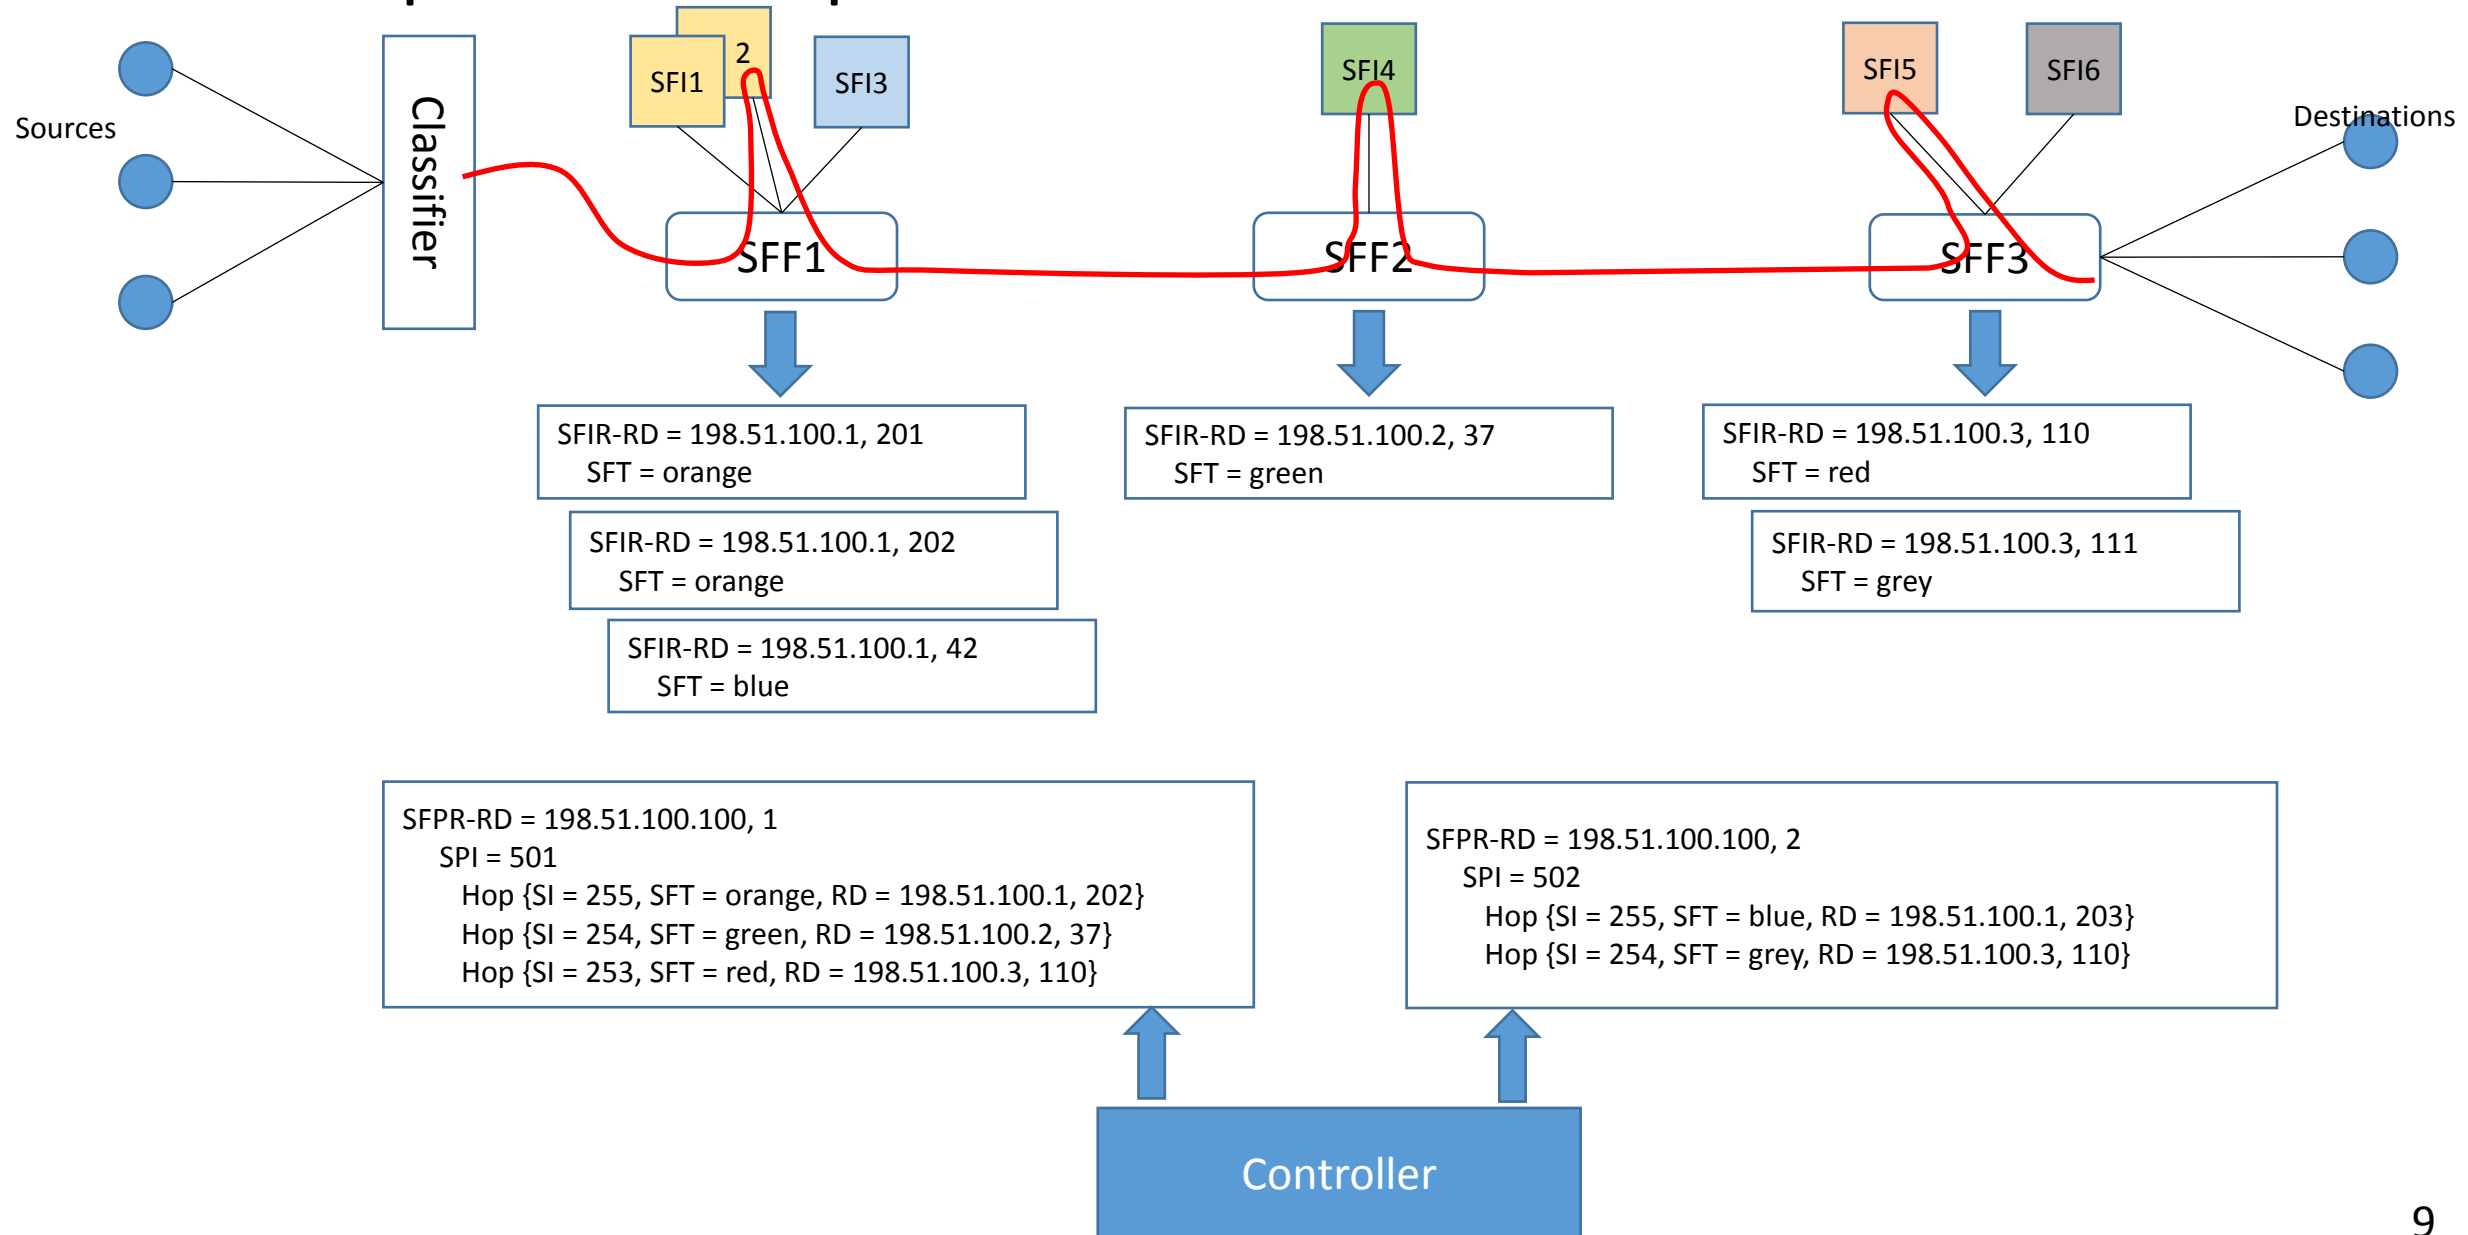
- Service Function Paths are constructed and advertised by controllers
- An SFP is a sequence of SFIs
- Advertisement contains:
  - Route Target
    - So only participating nodes need to import the advertisement
  - Route Distinguisher (SFPR-RD)
    - Identifies the SFP advertisement
  - Service Path Identifier (SPI)
    - Uniquely identifies the SFP
    - Used in the forwarding plane to identify this SFP
  - Series of hops in the path each encoded as a Hop TLV

# The Hop TLV

- One instance of the Hop TLV for each hop in the path
- Each Hop TLV contains
  - Service Index
    - Used in the forwarding plane to identify this hop
  - A Service Function Type
    - The type of SF that must be executed
  - An SFIR-RD
    - The RD of the SFIR that advertised the SFI to be executed
- The uncomplicated case
  - SFPR is just a series of Hop TLVs each with one SFT/SFIR-RD



# A Simple Example



# Advanced Function

- Offering a choice of next hop
  - A Hop TLV can carry multiple SFI identifiers
    - Allows for load-balancing or other policy choices through re-classification
  - Choice may be between SFIs of same or different types
- Choice may be open
  - A Hop TLV indicates a specific SFT, but leaves choice of SFI open
  - Allows SFF to select “best” next hop considering load and underlay network
- Explicit control of next hop can be achieved using a “special purpose SFT”
  - Standards action range (1-31)
  - One value defined : “Change Sequence”
  - In this case the SFIR-RD is overloaded to contain SPI/SI of next hop
    - May be anywhere on the same SFP (“jumping”)
    - May be another SFP (“branching”)
- Encapsulation between SFFs
  - The SFIR can include a Tunnel Encapsulation attribute to tell other SFFs how to reach the SFI
- Association of SFPs
  - SFPR can include an Association TLV containing the SFPR-RD and SPI of an associated SFP
  - Allows creation of a bidirectional SFP
    - Opposite directions do not need to be co-routed

# Points of Contention (1 of 3)

1. Is this work for BESS or SFC?
  - In charter at BESS, out of charter at SFC
  - But, MUST socialize to SFC
2. The Terminology is not aligned
  - First version got confused
    - Mainly SFC versus SFP
  - Current version is better
    - Still some errors around the location of a Classifier
  - It's a work in progress!
3. Whose job is it to decrement SI?
  - Out of scope of this document
  - It is a question for the SFC WG to resolve
  - This control plane solution supports anyone decrementing SI
  - In particular
    - SF can decrement
    - Classifier can decrement
    - Classifier can be separate or co-resident

# Points of Contention (2 of 3)

4. When can re-classification happen
  - Out of scope of this document
  - It is a question for the SFC WG to resolve
  - A re-classifier can be co-resident with SF or SFF, or in between per RFC 7665
  - This document supports any of these options
5. Does “decrement SI” mean “decrement SI by one”?
  - Out of scope of this document
  - It is a question for the SFC WG to resolve
  - This control plane solution supports any decrement of SI
  - In particular
    - SF decrement by 1 is supported
    - Classifier can modify SI
6. Support for looping, jumping, branching, spiralling
  - Yes, we support all of them
  - There is a danger of infinite loops caused by re-classification
    - That is out of scope of this document
    - SFC WG may need to think about a solution, but it does not belong in this discussion
  - Our work allows the definition of explicit “change sequence” instructions in the SFP distributed by BGP
    - Allows “splicing” of SFPs

# Points of Contention (3 of 3)

## 7. Support for re-classification

- Yes, re-classification is really important
- The role of the Classifier is defined in SFC docs and is out of scope for this I-D
- Our work allows the definition of explicit “choice” in the SFP distributed by BGP to be acted on by re-classification function
  - Simple next hop choice such as for load balancing
  - Specific

## 8. Does “choice” work with OAM?

- It seems so
  - In band OAM will follow the same path as actual traffic
    - Just like in IP forwarding the secret is that choices must be consistent and stable
  - Out of band OAM should report on choices that are available
- Aim to be consistent with OAM discussed in RFC7665 and draft-ietf-sfc-nsh
- draft-ietf-sfc-oam-framework expired three months ago
  - What is the plan?

## 9. How does this relate to draft-ietf-sfc-control-plane?

- Compatible with forwarding requirements in that draft
- Need to do deeper analysis and could make references to relevant sections
- The relevant bullets from Section 3.1 are

The SFC control plane is responsible for the following:

- Build and monitor the service-aware topology. For example, this can be achieved by means of dynamic SF discovery techniques. Those means are out of scope of this document.
- Provision SFP Forwarding Policy Tables of involved SFFs and provide classifiers with traffic classification rules.