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)

Objectives

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

How it Works

- BGP used to advertise using a new AFI/SAFI with two route types
  - Service Function Forwarder and Service Function discovery
  - Service Function Path composition
- SFC Overlay Networks
  - The SFFs are connected together by tunnels crossing underlay networks
  - The SFFs form an overlay network
  - We allow multiple overlay networks at once and distinguish them using Route Targets
- Service Function Types
  - New registry of type indicators for service functions
- SFF/SF discovery
  - Each SFF advertises a Service Function Instance Route (SFIR) for each SFI it supports
- SFPR for SFP distribution
  - The controller advertises the whole SFP so that nodes on the path know about it

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
    - And the information to build tunnels across the underlay
  - 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 the SFI advertisement
  - SF Type (SFT)
    - From the FCFS IANA registry
- The combination SFIR-RD/SFT uniquely identifies a specific SFI

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

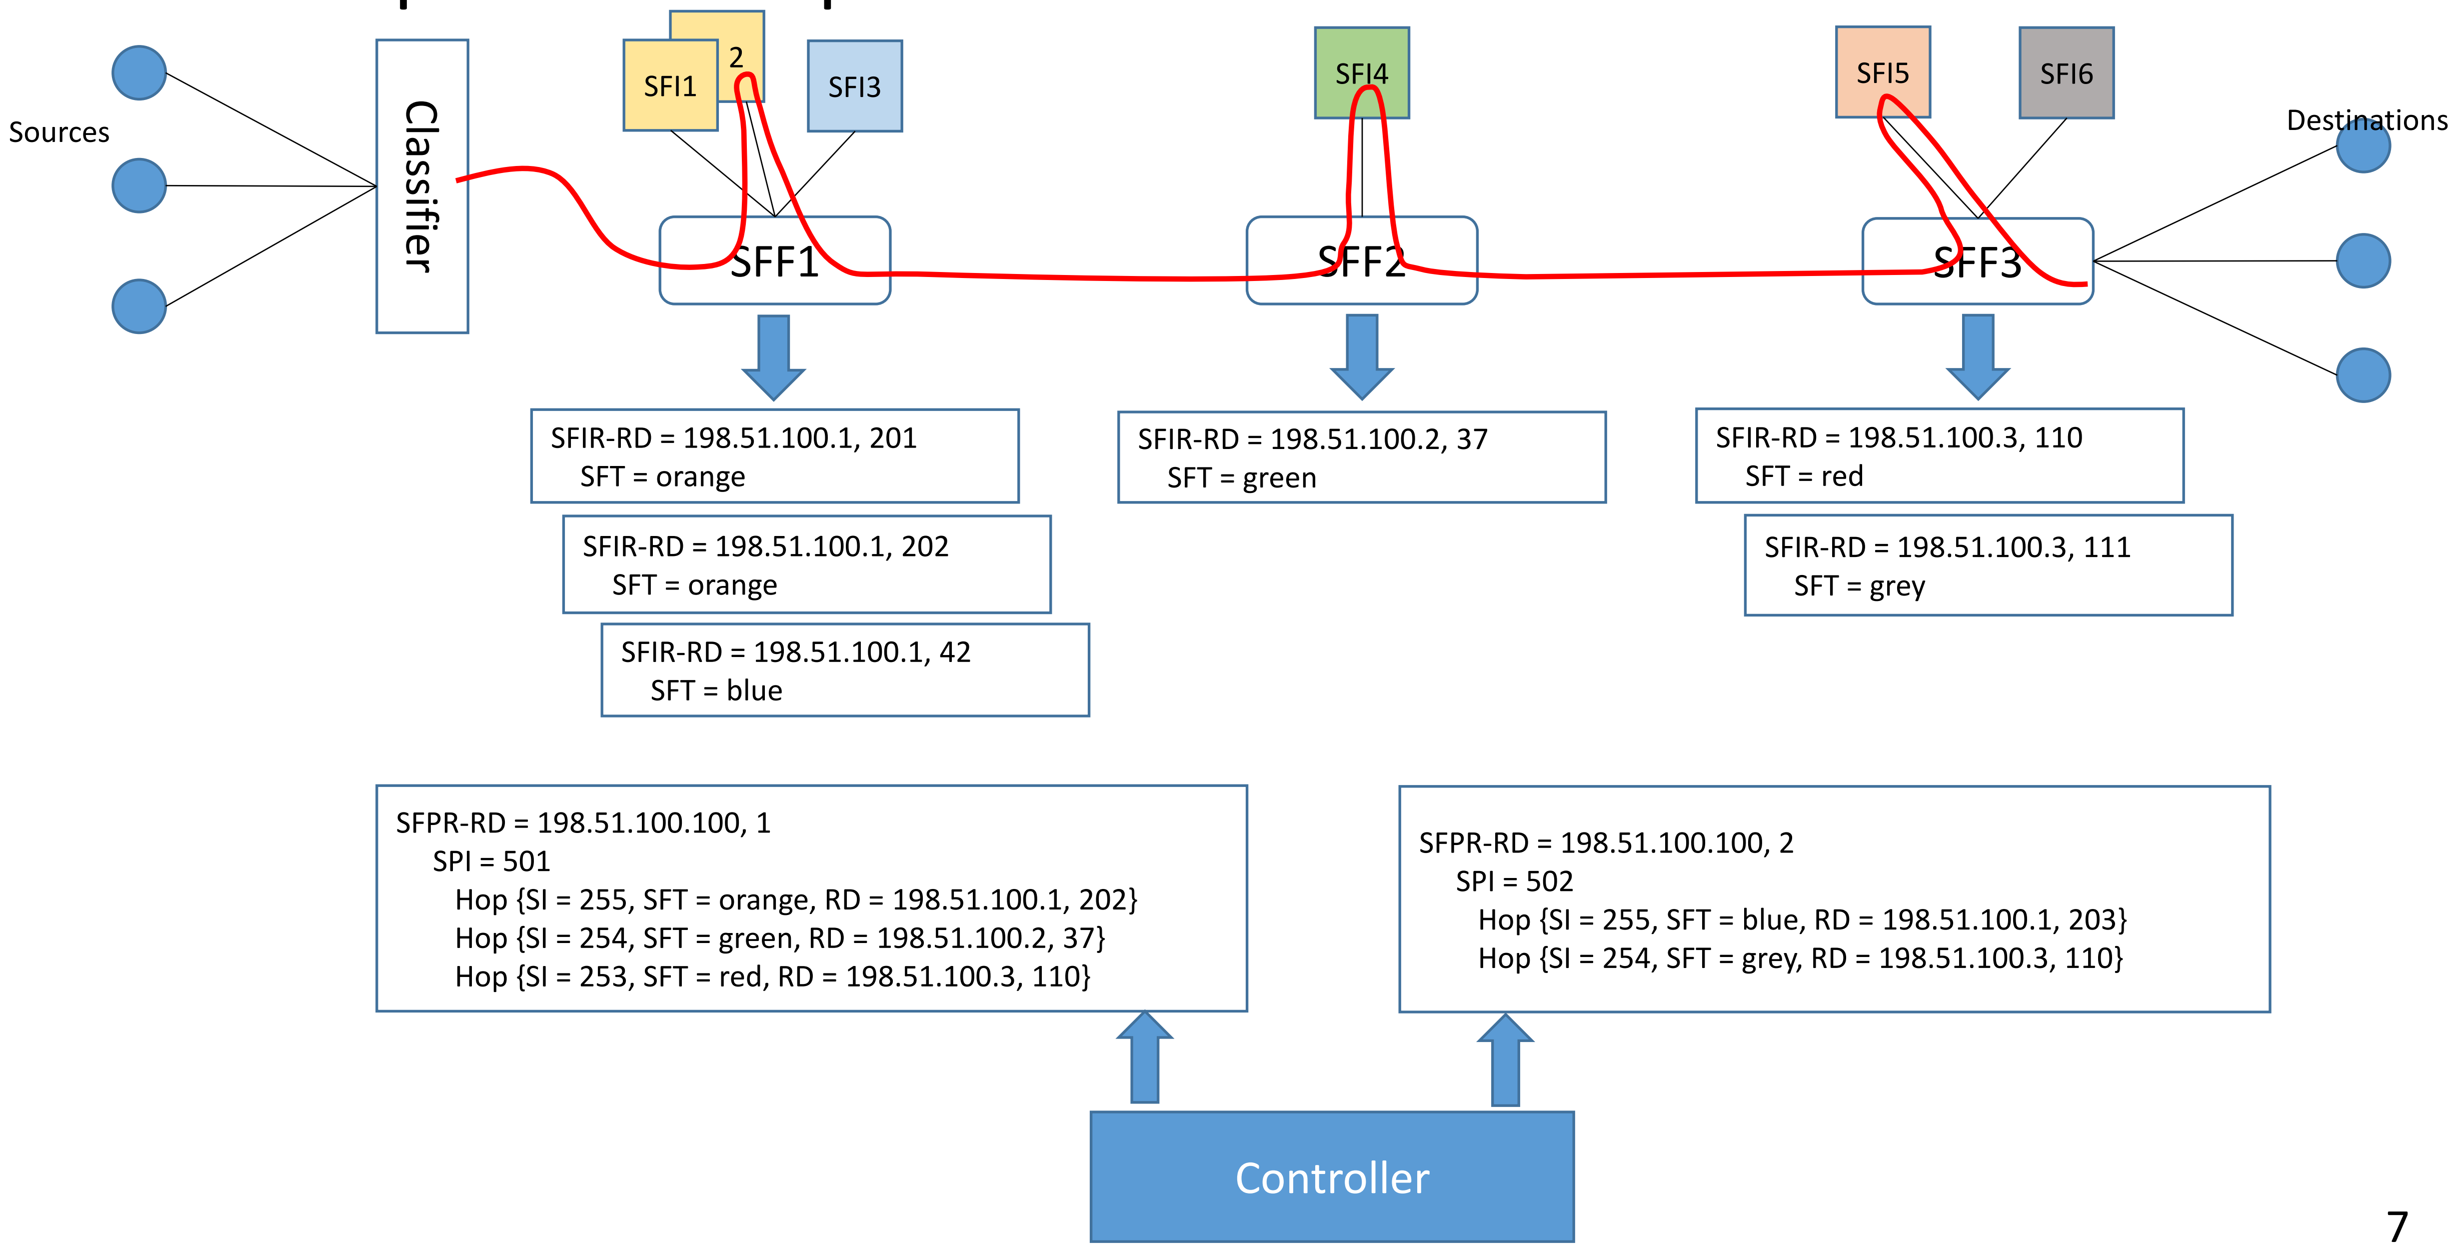

Advanced Function

- 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

- Is this work for BESS or SFC?
  - In charter at BESS, out of charter at SFC
  - BUT, MUST socialize to SFC
- Support for looping, jumping, branching, spiralling
  - Yes, we support all of them
  - There is a danger inherent in the SFC architecture of infinite loops caused by looping and branching
    - SFC WG needs to think about a solution
- Whose job is it to decrement SI?
  - Out of scope of this document
  - This control plane solution supports anyone decrementing SI
- Does “decrement SI” mean “decrement SI by one”?
  - Out of scope of this document
  - This control plane solution supports any decrement of SI
- When can re-classification happen
  - A re-classifier can be co-resident with SF or SFF, or in between per RFC 7665
  - This document supports any of these options
  - Our work allows the definition of choices in the SFP
    - Whenever a choice is made, this is “re-classification”. Also known as “local policy”.
- How does this relate to draft-ietf-sfc-control-plane?
  - Compatible with forwarding requirements in that draft