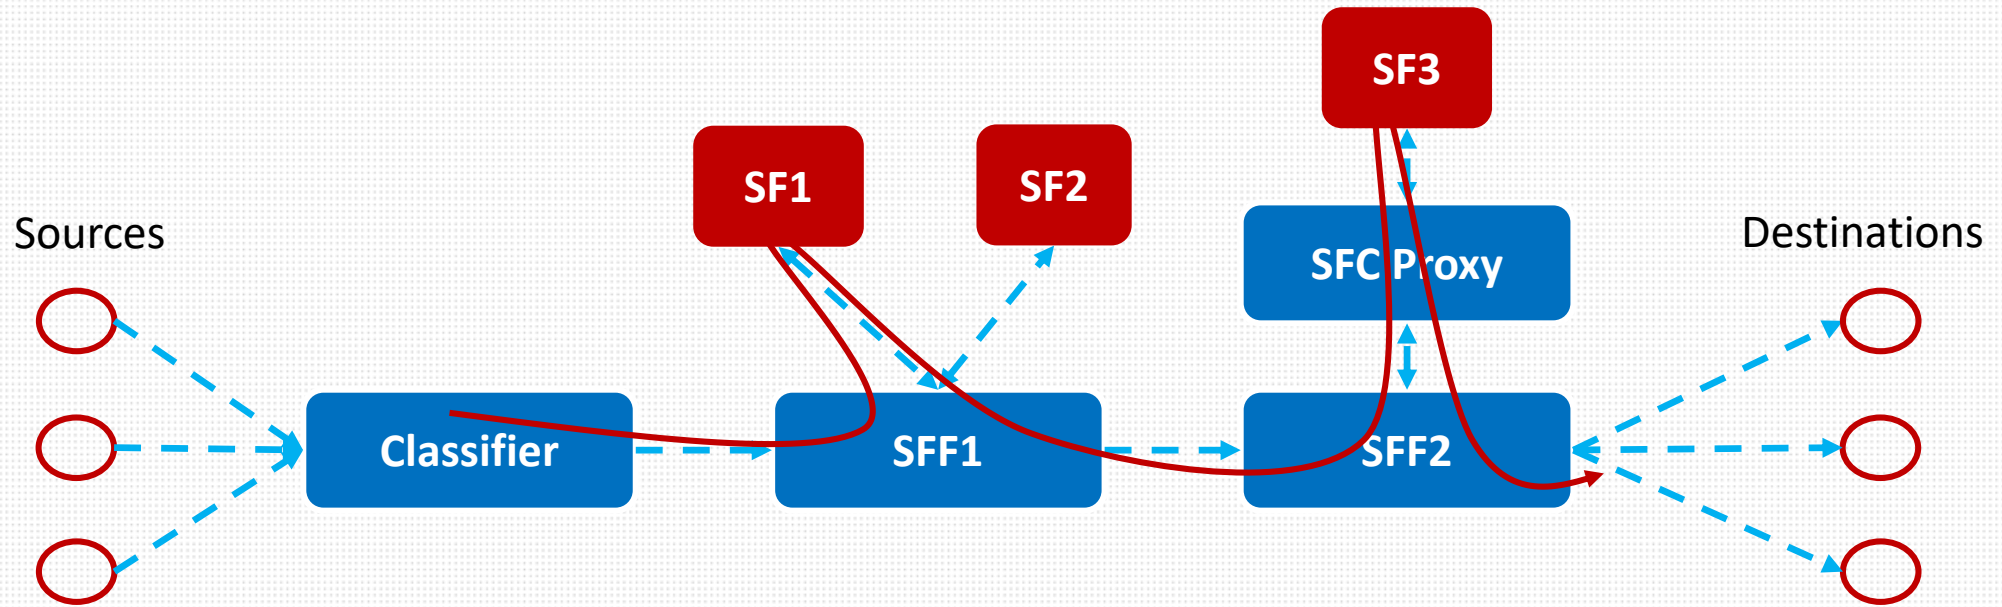




Leveraging SR to Realize Transport-Independent Service Function Chaining

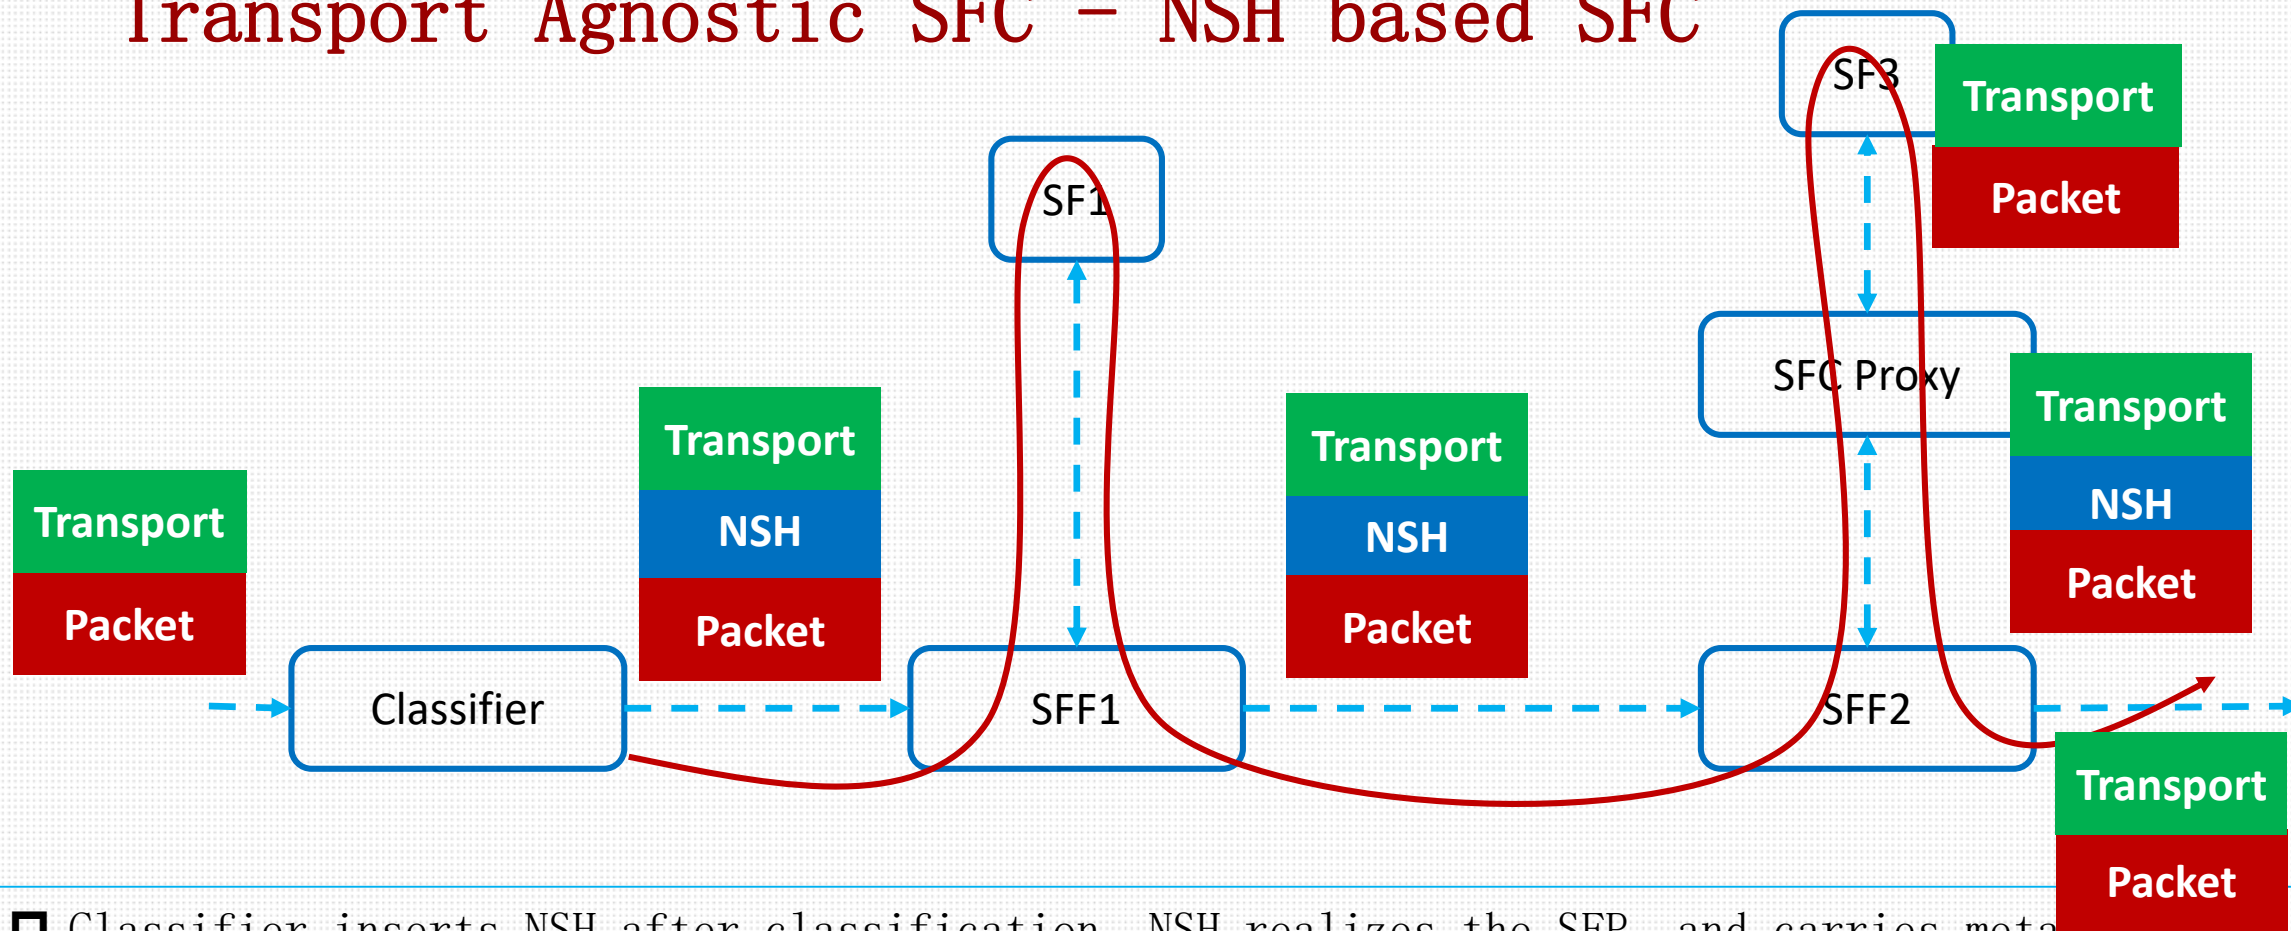
- Mach Chen
- mach.chen@huawei.com

An Overview of SFC Architecture



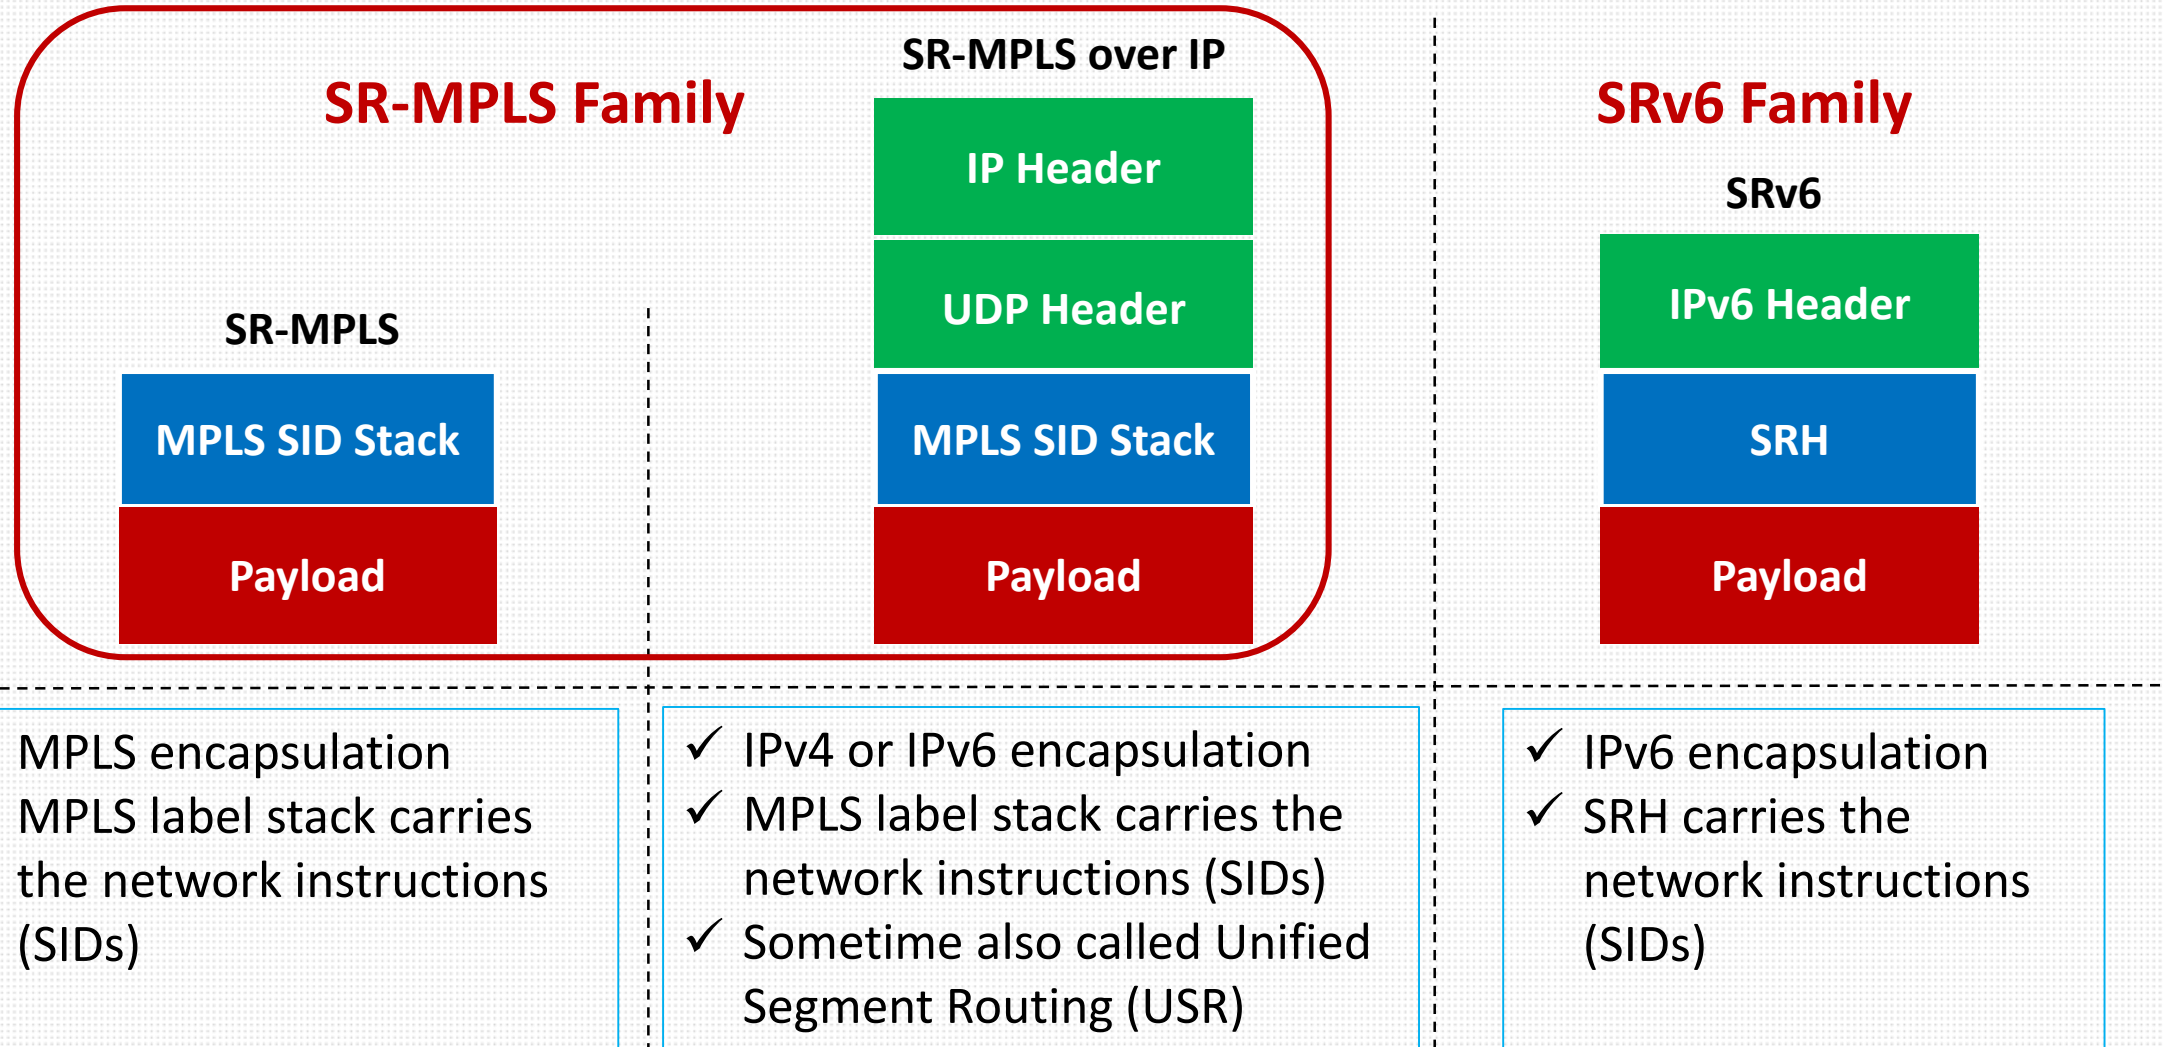
- ❑ Given an SFC
 - ✓ SFC: <SF1, SF3>
- ❑ Packets are classified onto a Service Function Path (SFP)
 - ✓ SFP: <SFF1, SF1, SFF2, SF3>
- ❑ Service Function Forwarder (SFF) determines packets to SF or to next SFF
- ❑ SFC proxy may be placed between SFF and SF to support SFC-unaware SF

Transport Agnostic SFC – NSH based SFC

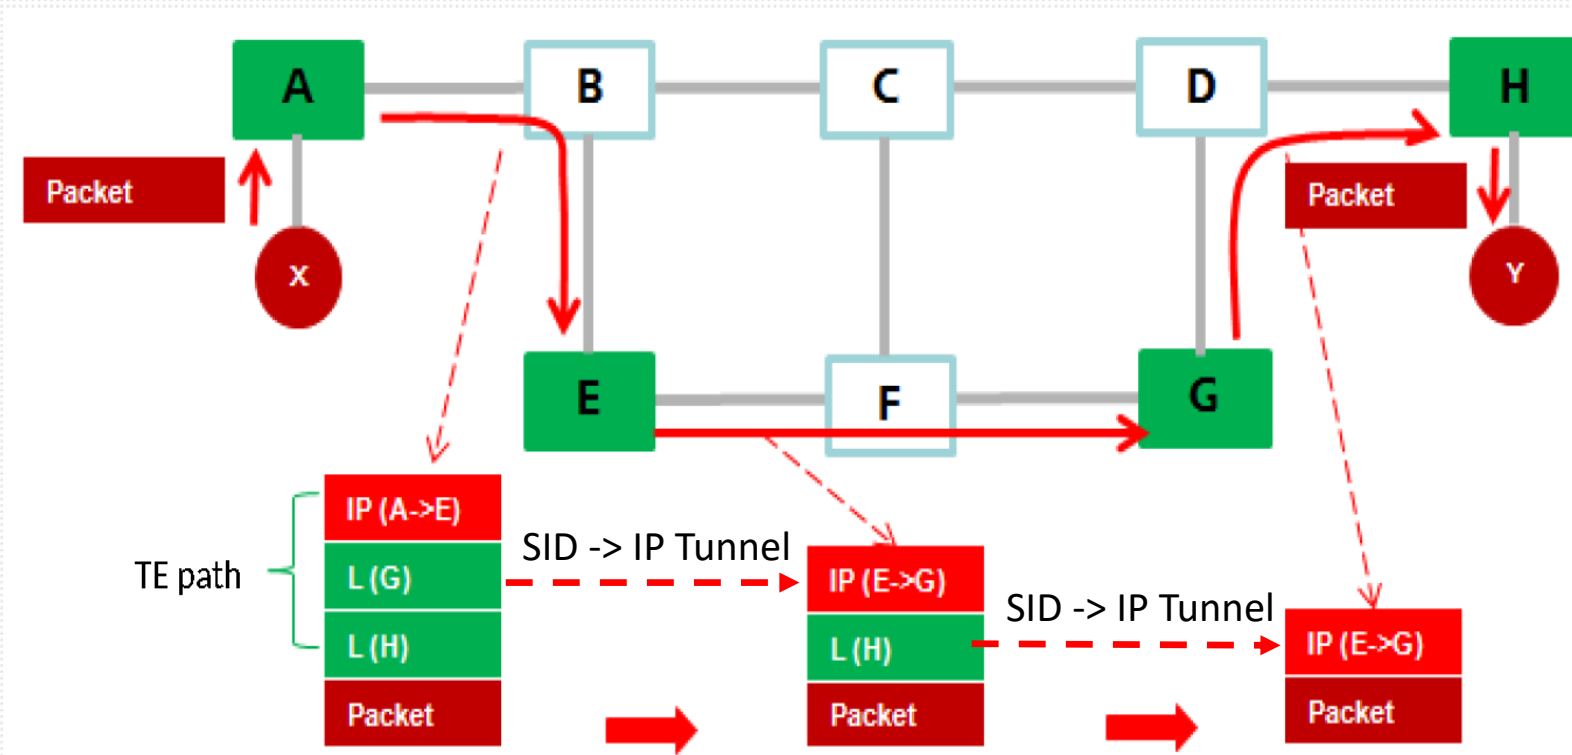


- ❑ Classifier inserts NSH after classification, NSH realizes the SFP and carries metadata
 - ✓ SFP: <SFF1, SF1, SFF2, SF3>
- ❑ SFFs maintain per-SFP states and determine (based NSH) to deliver packet to SF or to next SFF;
- ❑ SFC Proxy strips NSH when sending to NSH-unaware SF, puts the NSH back when sending back to SFF
- ❑ NSH based SFC is transport agnostic, the transport can be Ethernet, GRE, VXLAN-GPE, SR, etc.

A Recall of Segment Routing Family



What's SR-MPLS over IP?



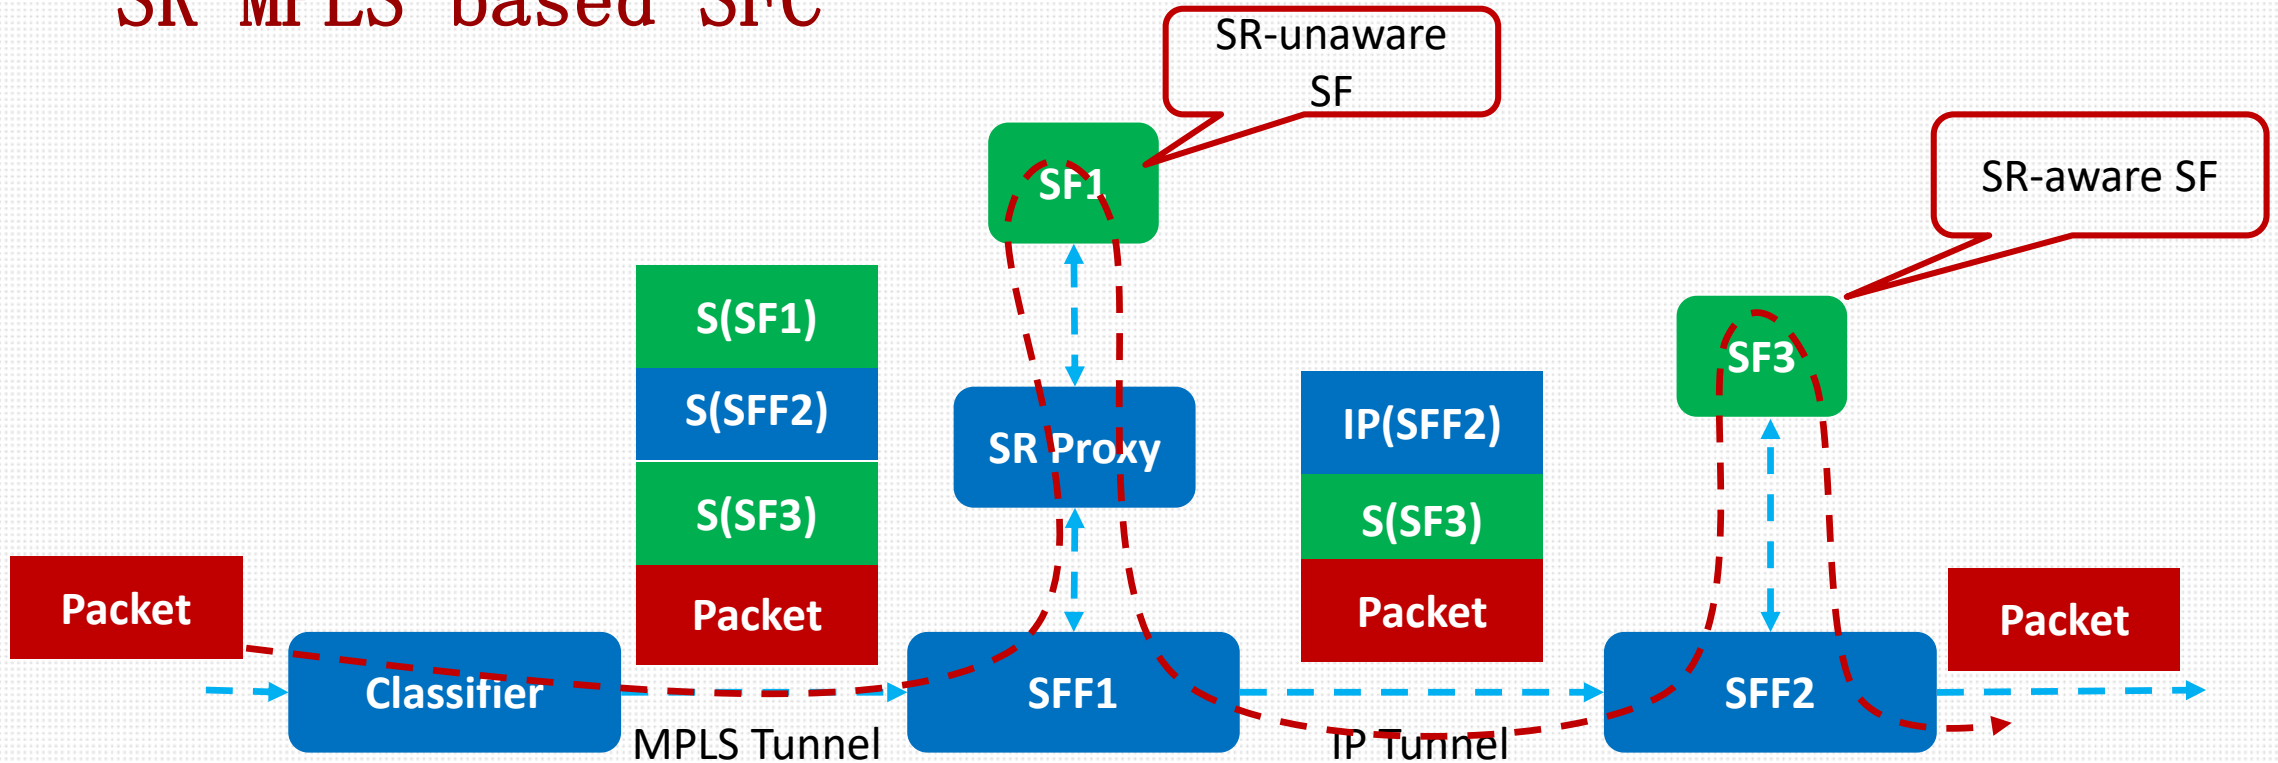
- ✓ Encapsulate a “normal” SR-MPLS packet in UDP in IP
- ✓ Each SID of the MPLS SID Stack maps to an IP tunnel
- ✓ Packets tunneled to next SR-capable node in the SR path by the IP tunnel
- ✓ UDP destination port indicates “MPLS below”
- ✓ UDP source port can be used to carry entropy for ECMP purpose

Tunnel SR-MPLS over an IP network, enable SR in legacy networks by tactically introducing SR-capable nodes at strategic points in the network.

Leverage SR to Realize SFC

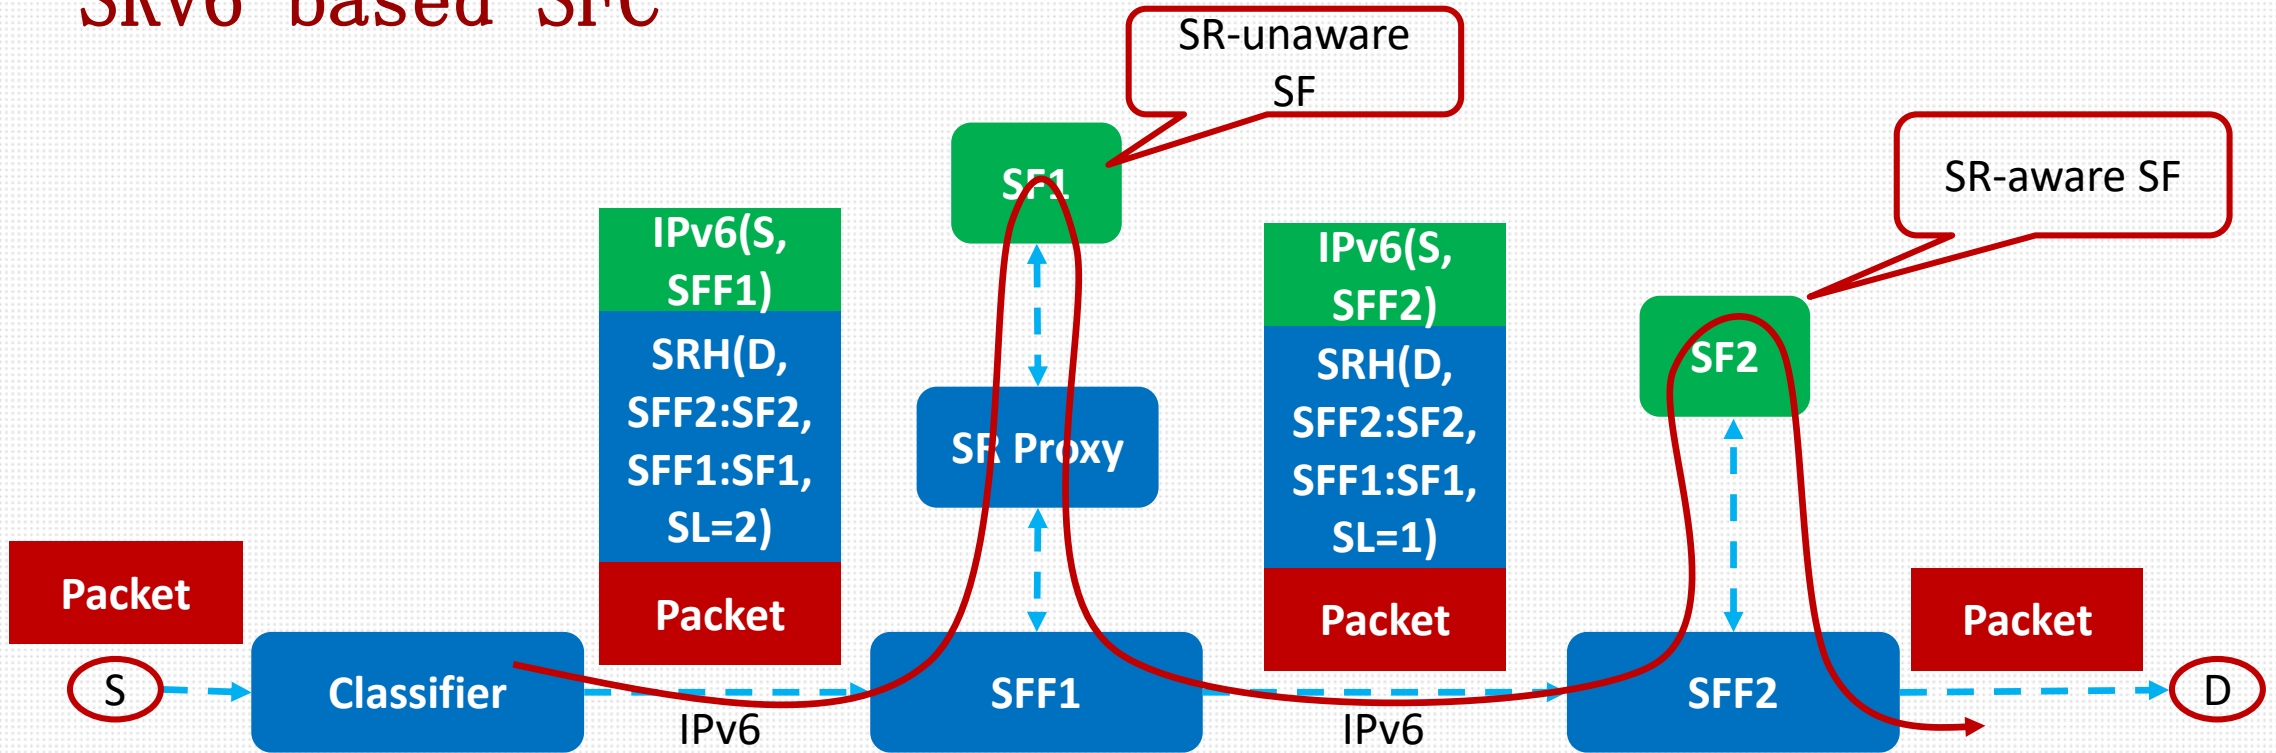
- **SR-MPLS based SFC**
 - › Services encoded as stack of MPLS labels
 - › No good way to carry metadata
- **SRv6 based SFC**
 - › Services encoded as SRH list
 - › SRH TLV to carry metadata
- **NSH-based SFC with SR as the transport tunnel**
 - › SR tunnels connect SFF
 - › SFF uses NSH to determine packets to SF or to next SFF
- **SR-based SFC with integrated NSH service plane**
 - › SR is responsible for steering traffic through the necessary SFFs
 - › NSH is responsible for maintaining the service plane, and holding the SFC instance context and associated metadata

SR-MPLS based SFC



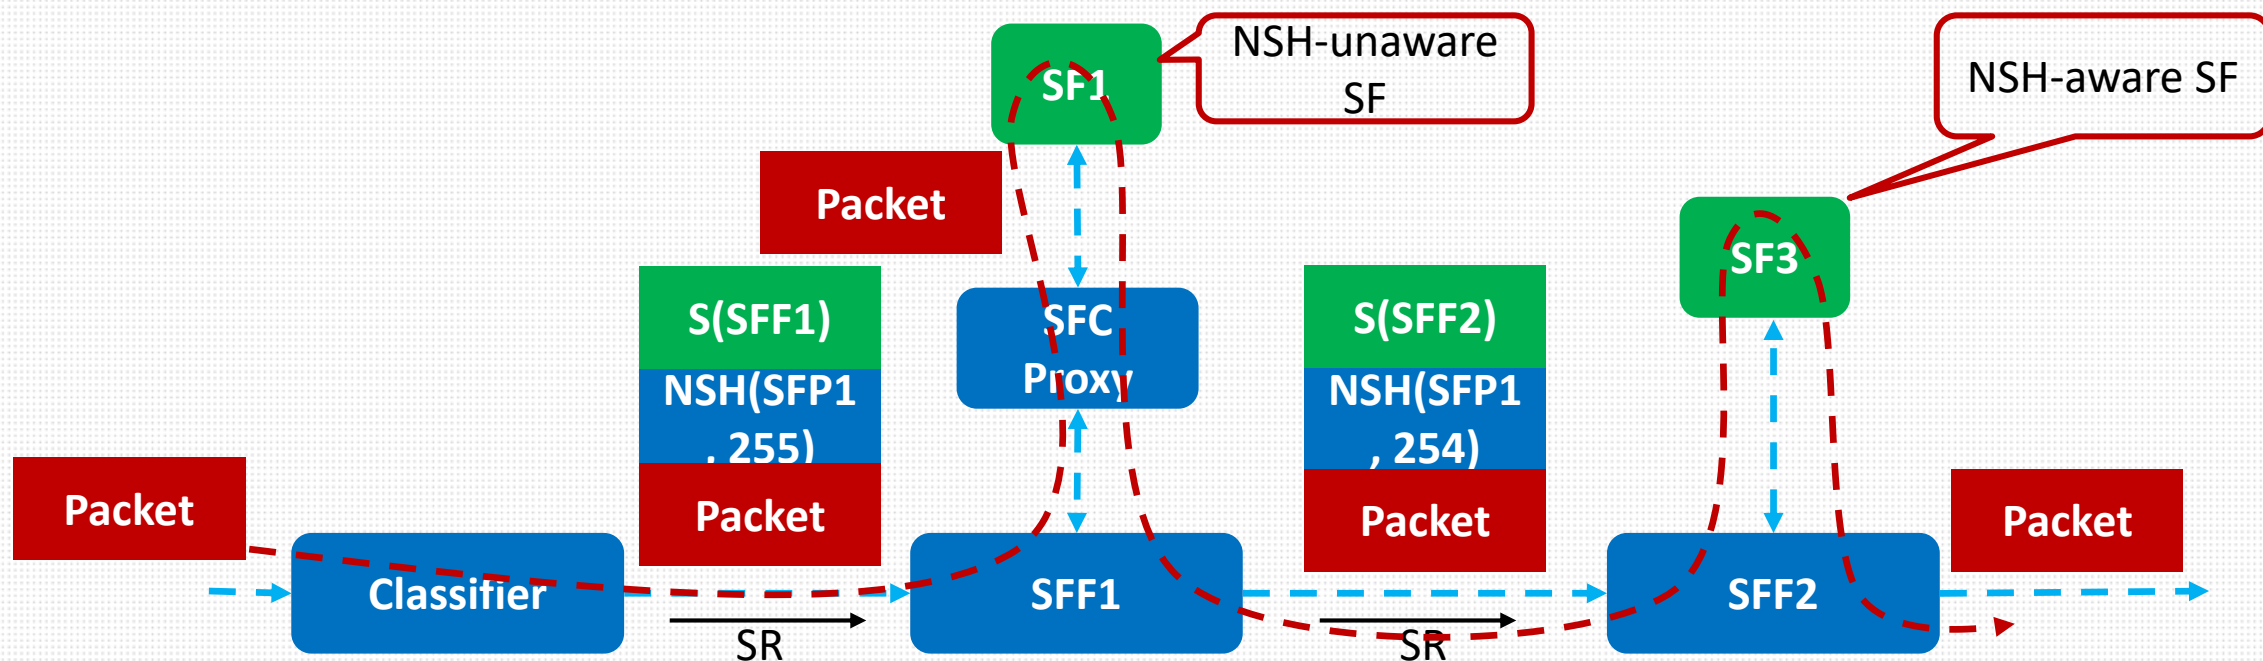
- ✓ Each SFF and SF will be allocated an SID, an SFP is realized by a SID list (MPLS label stack)
- ✓ SF can be either SR-aware or SR unaware, SR unaware needs the help of an SR Proxy
- ✓ SR Proxy strips the SR header when send to SF, and put the SR header back when send packet to SFF
- ✓ Underneath tunnel could be MPLS, IPv4 or IPv6;
- ✓ No per-path state maintained at SFF;
- ✓ No good way to carry metadata;

SRv6 based SFC



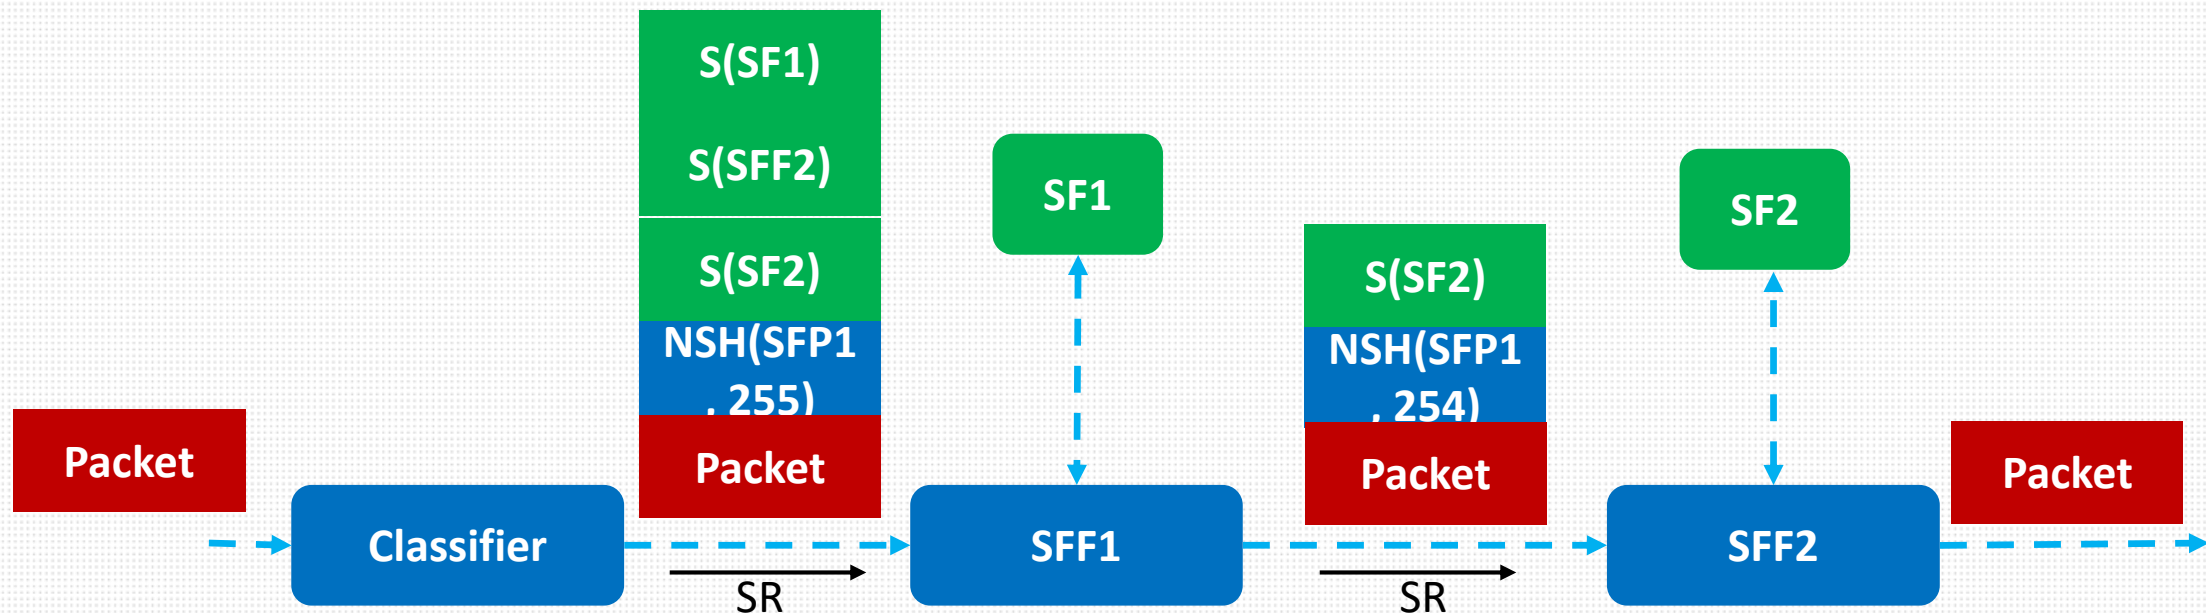
- ✓ Each SFF and SF will be allocated an SID, an SFP is realized by a SID list (SRH)
- ✓ SF can be either SR-aware or SR unaware, SR unaware needs the help of an SR Proxy
- ✓ SR Proxy strips the SR header when send to SF, and put the SR header back when send packet to SFF; SR proxy needs a to correlate a packet to an SRH (e.g., interface id, not perfect)
- ✓ Use SRH TLV to carry metadata;
- ✓ No per-path state maintained at SFF;

NSH-based SFC with SR as the Transport Tunnel



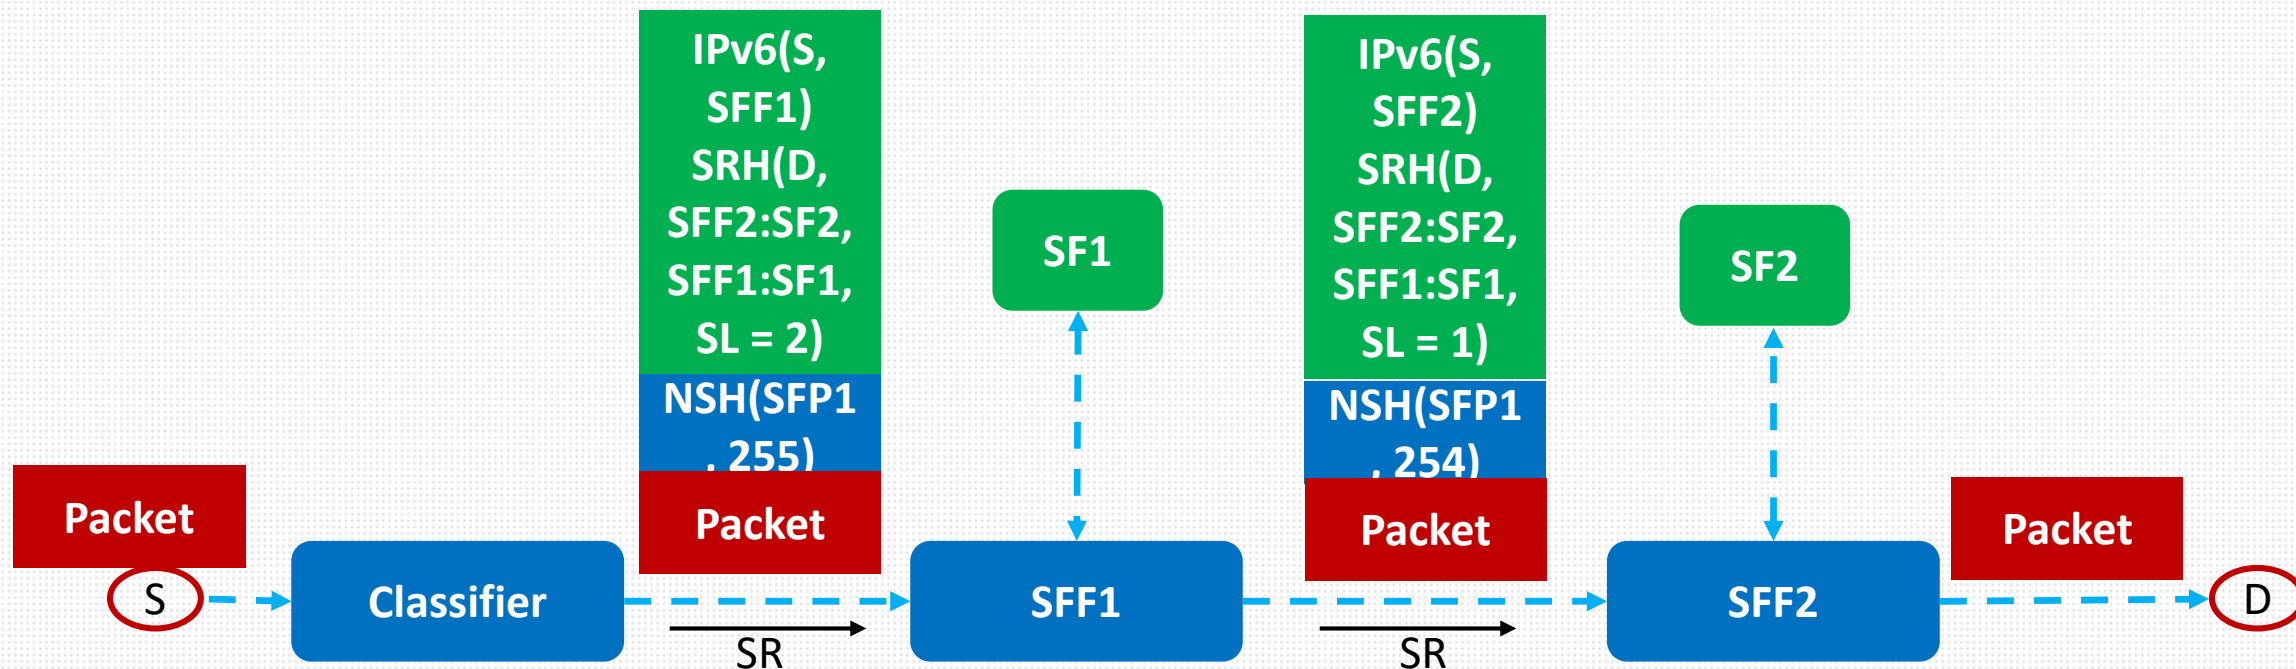
- ✓ Classifier inserts NSH header and determines (based on NSH) whether SR tunnel will be used;
- ✓ If so, put SR header and tunnel it to next SFF;
- ✓ When packet arrived SFF, strip the SR header, send the packet to SF;
- ✓ For NSH-unaware SF, an SFC proxy will be placed between SFF and SF, the proxy responsible for strip/add NSH header;
- ✓ Totally decouple SFC layer and the transport layer, keep all merits of NSH; it can also leverage SR for Traffic steering between SFFs

SR-MPLS based SFC with Integrated NSH Service Plane



- ✓ Classifier inserts NSH header and put the SR header that realize the SFP (e.g., <SFF1, SF1, SFF2, SF2>);
- ✓ SFF strips the SR header before sending the packet the SF;
- ✓ SR header cached at SFF by using SPI as the index;
- ✓ When received packet from SF, using the SPI to correlate the cached SR header and put the SR header back to the packet;
- ✓ Still transport agnostic, keep all the merits of NSH and SR, maintain fewer states at SFFs

SRv6 based SFC with Integrated NSH Service Plane



- ✓ Classifier inserts NSH header and put the SRH header that realize the SFP(e.g., <SFF1, SF1, SFF2, SF2>);
- ✓ SFF strips the SRH header before sending the packet the SF;
- ✓ SRH header cached at SFF by using SPI as the index;
- ✓ When received packet from SF, using the SPI to correlate the cached SRH header and put the SRH header back to the packet;
- ✓ Transport agnostic, keep all the merits of NSH and SRv6, maintain fewer states at

A Summary of SR based SFC

- SR based SFC has several options, each has its own pros and cons
- NSH is designed to be transport agnostic, NSH based SFC is more and more accepted by the industry.
- Segment Routing(SR) provides an efficient way to steering traffic without requiring intermediates to maintain per-flow states
- Combine with SR and NSH, a transport-independent SFC can be realized.
- A solution that keeps merits of both NSH and SR is attractive.



Thank you

www.huawei.com

Copyright©2014 Huawei Technologies Co., Ltd. All Rights Reserved.

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.