



Service Function Chaining (SFC)

Release brahmaputra.1.0 (28c87f4)

OPNFV

February 25, 2016

1	Introduction	1
2	Definitions	3
3	Abbreviations	5
4	Use Cases	7
5	Architecture	9
5.1	Service Functions	9
5.2	Service Function Forwarders	9
5.3	Service Chains	9
5.4	Service Chaining Encapsulation	9
5.5	Classifiers	10
5.6	VNF Manager	10
5.7	OPNFV SFC Network Topology	11
5.8	OVS NSH patch workaround	11
6	Requirements	15
6.1	Minimal Viable Requirement	15
6.2	Detailed Requirements	15
6.3	Long Term Requirements	15

INTRODUCTION

Note: This is the working documentation for the SFC project.

The [OPNFV Service Function Chaining \(SFC\)](#) project aims to provide the ability to define an ordered list of a network services (e.g. firewalls, NAT, QoS). These service are then “stitched” together in the network to create a service chain. This project provides the infrastructure to install the upstream ODL SFC implementation project in an NFV environment.

DEFINITIONS

Definitions of most terms used here are provided in the [IETF SFC Architecture RFC](#). Additional terms specific to the OPNFV SFC project are defined below.

ABBREVIATIONS

Table 3.1: Abbreviations

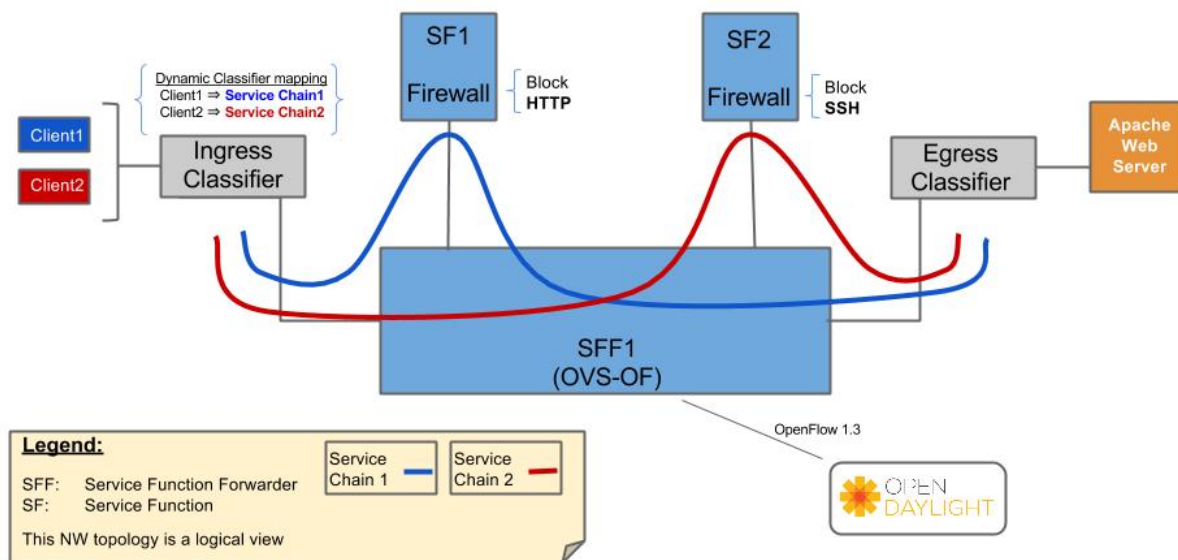
Abbreviation	Term
NS	Network Service
NFVO	Network Function Virtualization Orchestrator
NF	Network Function
NSH	Network Services Header (Service chaining encapsulation)
ODL	OpenDaylight SDN Controller
RSP	Rendered Service Path
SDN	Software Defined Networking
SF	Service Function
SFC	Service Function Chain(ing)
SFF	Service Function Forwarder
SFP	Service Function Path
VNF	Virtual Network Function
VNFM	Virtual Network Function Manager
VNF-FG	Virtual Network Function Forwarding Graph
VIM	Virtual Infrastructure Manager

USE CASES

This section outlines the Brahmaputra use cases driving the initial OPNFV SFC implementation.

The use cases targeted in the OPNFV SFC Brahmaputra release focus on creating simple Service Chains using Firewall Service Functions. As can be seen in the following diagram, 2 service chains are created, each through a different Service Function Firewall. Service Chain 1 will block HTTP, while Service Chain 2 will block SSH.

Initial Use Case Logical Overview



ARCHITECTURE

This section describes the architectural approach to incorporating the upstream OpenDaylight (ODL) SFC project into the OPNFV Brahma Putra platform.

5.1 Service Functions

A Service Function (SF) is a Function that provides services to flows traversing a Service Chain. Examples of typical SFs include: Firewall, NAT, QoS, and DPI. In the context of OPNFV, the SF will be a Virtual Network Function. The SFs receive data packets from a Service Function Forwarder.

5.2 Service Function Forwarders

The Service Function Forwarder (SFF) is the core element used in Service Chaining. It is an OpenFlow switch that, in the context of OPNFV, is hosted in an OVS bridge. In OPNFV there will be one SFF per Compute Node that will be hosted in the “br-int” OpenStack OVS bridge.

The responsibility of the SFF is to steer incoming packets to the corresponding Service Function, or to the SFF in the next compute node. The flows in the SFF are programmed by the OpenDaylight SFC SDN Controller.

5.3 Service Chains

Service Chains are defined in the OpenDaylight SFC Controller using the following constructs:

SFC A Service Function Chain (SFC) is an ordered list of abstract SF types.

SFP A Service Function Path (SFP) references an SFC, and optionally provides concrete information about the SFC, like concrete SF instances. If SF instances are not supplied, then the RSP will choose them.

RSP A Rendered Service Path (RSP) is the actual Service Chain. An RSP references an SFP, and effectively merges the information from the SFP and SFC to create the Service Chain. If concrete SF details were not provided in the SFP, then SF selection algorithms are used to choose one. When the RSP is created, the OpenFlows will be programmed and written to the SFF(s).

5.4 Service Chaining Encapsulation

Service Chaining Encapsulation encapsulates traffic sent through the Service Chaining domain to facilitate easier steering of packets through Service Chains. If no Service Chaining Encapsulation is used, then packets must be classified at every hop of the chain, which would be slow and would not scale well.

In ODL SFC, Network Service Headers (NSH) is used for Service Chaining encapsulation. NSH is an IETF specification that uses 2 main header fields to facilitate packet steering, namely:

NSP (NSH Path) The NSP is the Service Path ID.

NSI (NSH Index) The NSI is the Hop in the Service Chain. The NSI starts at 255 and is decremented by every SF. If the NSI reaches 0, then the packet is dropped which avoids loop detections.

NSH also has metadata fields, but that's beyond the scope of this architecture.

In ODL SFC, NSH packets are encapsulated in VXLAN-GPE.

5.5 Classifiers

A classifier is the entry point into Service Chaining. The role of the classifier is to map incoming traffic to Service Chains. In ODL SFC, this mapping is performed by matching the packets and encapsulating the packets in a VXLAN-GPE NSH tunnel.

The packet matching is specific to the classifier implementation, but can be as simple as an ACL, or can be more complex by using PCRF information or DPI.

5.6 VNF Manager

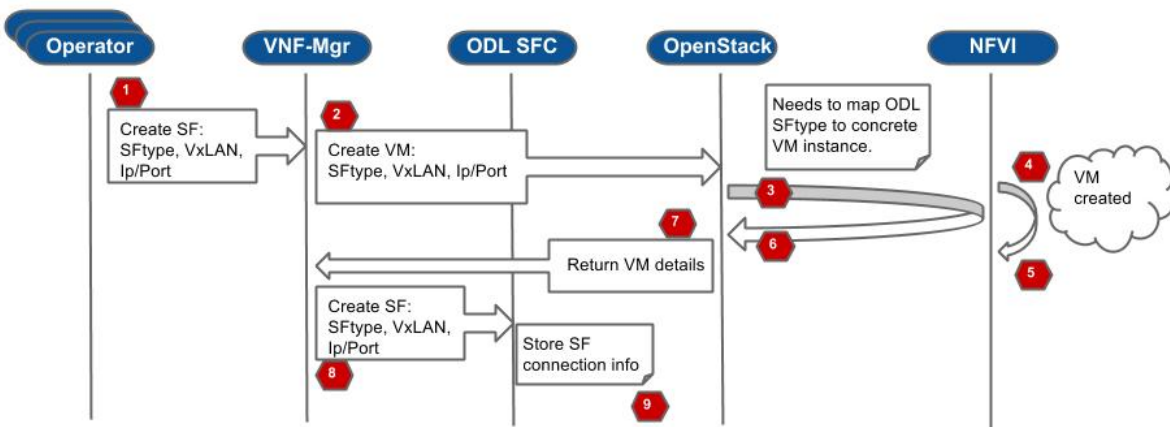
In OPNFV SFC, a VNF Manager is needed to spin-up VMs for Service Functions. It has been decided to use the OpenStack Tacker VNF Mgr to spin-up and manage the life cycle of the SFs. Tacker will receive the ODL SFC configuration, manage the SF VMs, and forward the configuration to ODL SFC. The following sequence diagram details the interactions with the VNF Mgr:

Sequence diagram: SF creation

In this scenario, the VNF-Mgr is responsible for:

- driving configuration into the ODL SFC
- requesting VM creation from NFVI

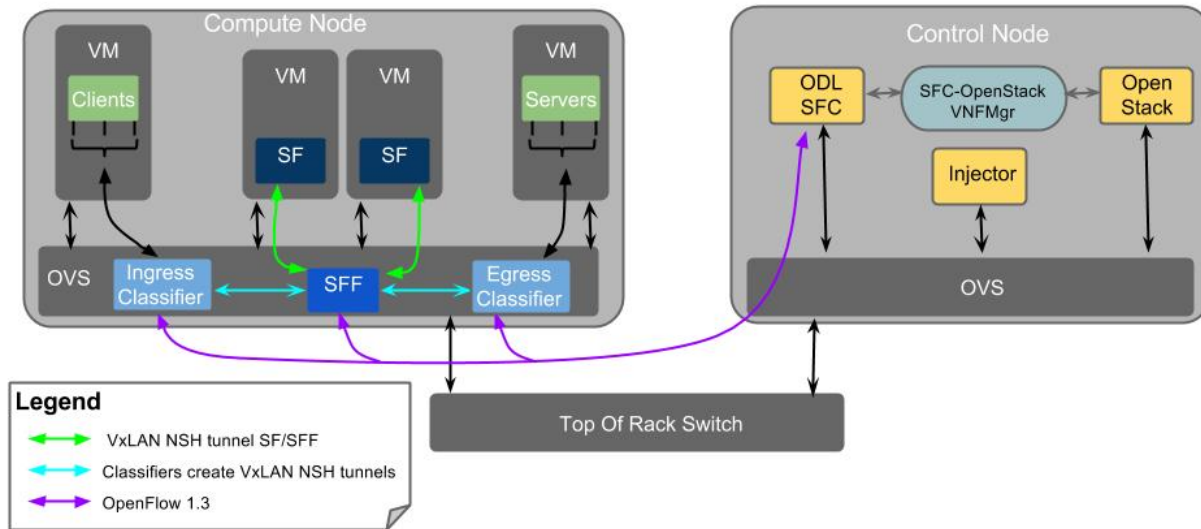
Operator is any generic trigger



5.7 OPNFV SFC Network Topology

The following image details the Network Topology used in OPNFV Brahmaputra SFC:

OPNFV SFC Initial NW Topology

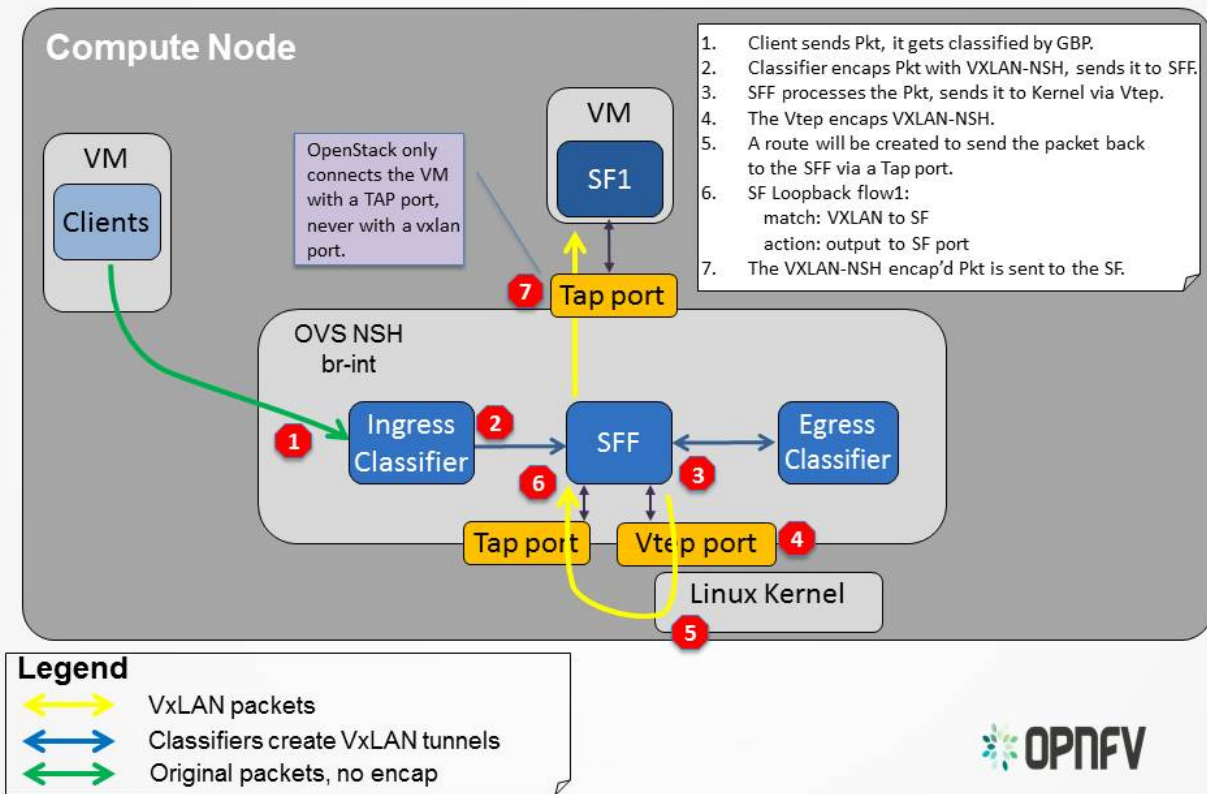


5.8 OVS NSH patch workaround

When using NSH with VXLAN tunnels, it's important that the VXLAN tunnel is terminated in the SF VM. This allows the SF to see the NSH header, allowing it to decrement the NSI and also to use the NSH metadata. When using VXLAN with OpenStack, the tunnels are not terminated in the VM, but in the "br-int" OVS bridge. There is work ongoing in the upstream OVS community to implement NSH encapsulation. To get around the way OpenStack handles VXLAN tunnels, the OVS work will also include the ability to encapsulate/decapsulate VXLAN tunnels from OpenFlow rules, instead of relying on the Vtep ports. The ongoing upstream OVS work will probably not be finished by the time OPNFV Brahmaputra is released, so a work-around has been created. This work-around will use a private branch of OVS that has a preliminary version of NSH implemented.

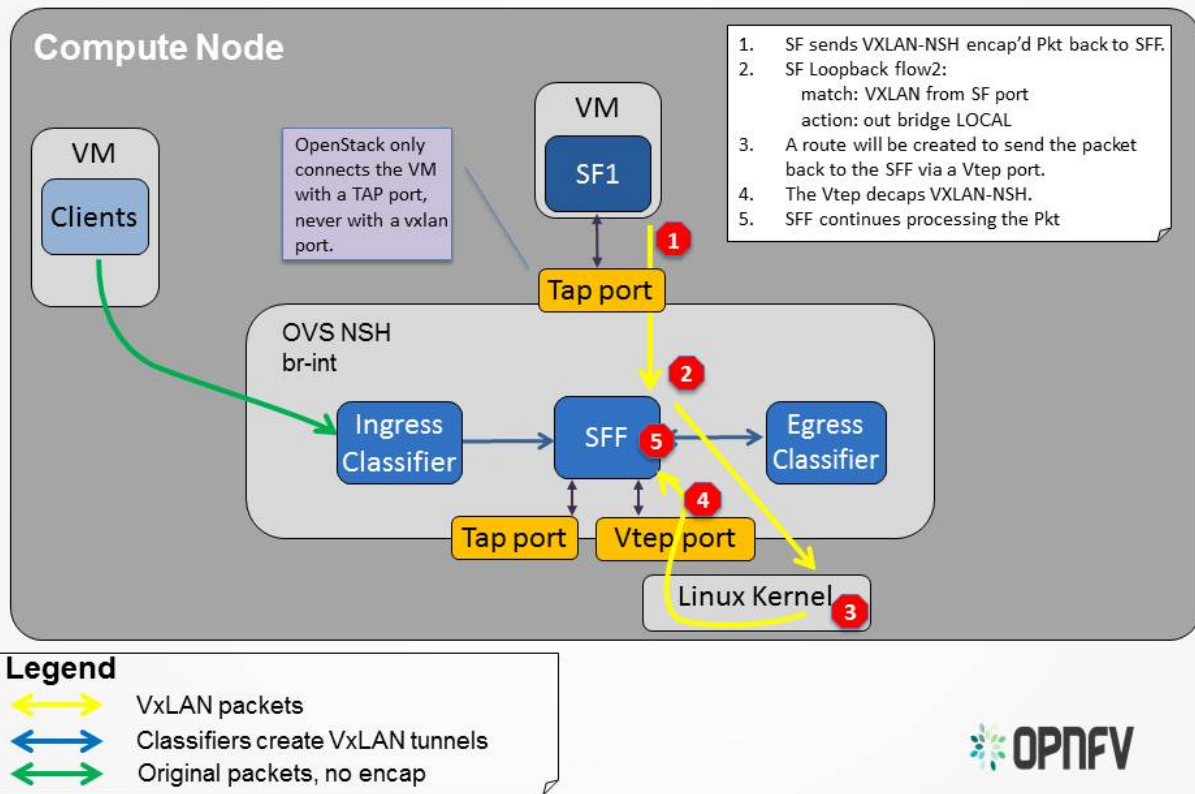
The following diagram illustrates how packets will be sent to an SF, when the SFF has processed the packet and wants to send it to the SF:

OPNFV SFC Loopback Workaround when sending packets to SF



The following diagram illustrates how packets will sent from an SF to an SFF, once the SF has processed a packet:

OPNFV SFC Loopback Workaround when receiving packets from SF



REQUIREMENTS

This section defines requirements for the initial OPNFV SFC implementation, including those requirements driving upstream project enhancements.

6.1 Minimal Viable Requirement

Deploy a complete SFC solution by integrating OpenDaylight SFC with OpenStack in an OPNFV environment.

6.2 Detailed Requirements

These are the Brahmaputra specific requirements:

- 1 Placement of SFs on only one Compute node will be supported.
- 2 The supported Service Chaining encapsulation will be NSH VXLAN-GPE.
- 3 The version of OVS used must support NSH.
- 4 The SF VM life cycle will be managed by the Tacker VNF Mgr.
- 5 The supported classifiers will be either ODL Netvirt or ODL GBP.
- 6 **ODL will be the OpenStack Neutron backend and will handle all networking** on the compute nodes.

6.3 Long Term Requirements

These requirements are out of the scope of the Brahmaputra release.

- 1 Placing SFs on multiple Compute nodes.
- 2 Dynamic movement of SFs across multiple Compute nodes.
- 3 Load Balancing across multiple SFs