



SFC with NSH and OVS

Thomas Graf (@tgraf__)

Uri Elzur (@ElzurUri)

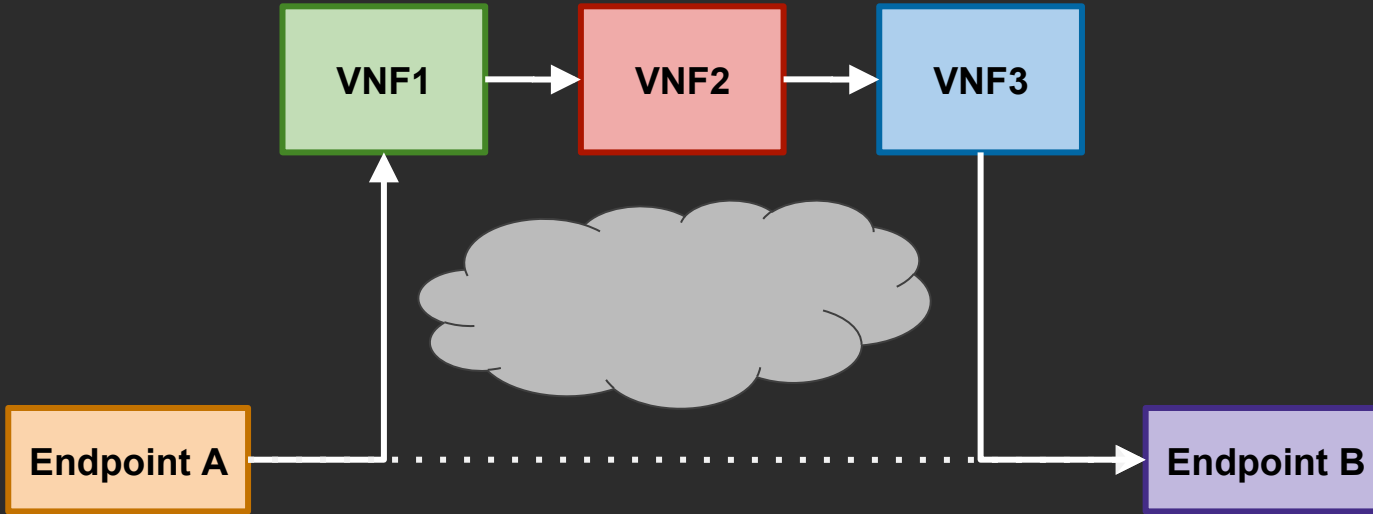
Russell Bryant (@russellbryant)

Danny Zhou

November 16th, 2015

What is SFC?

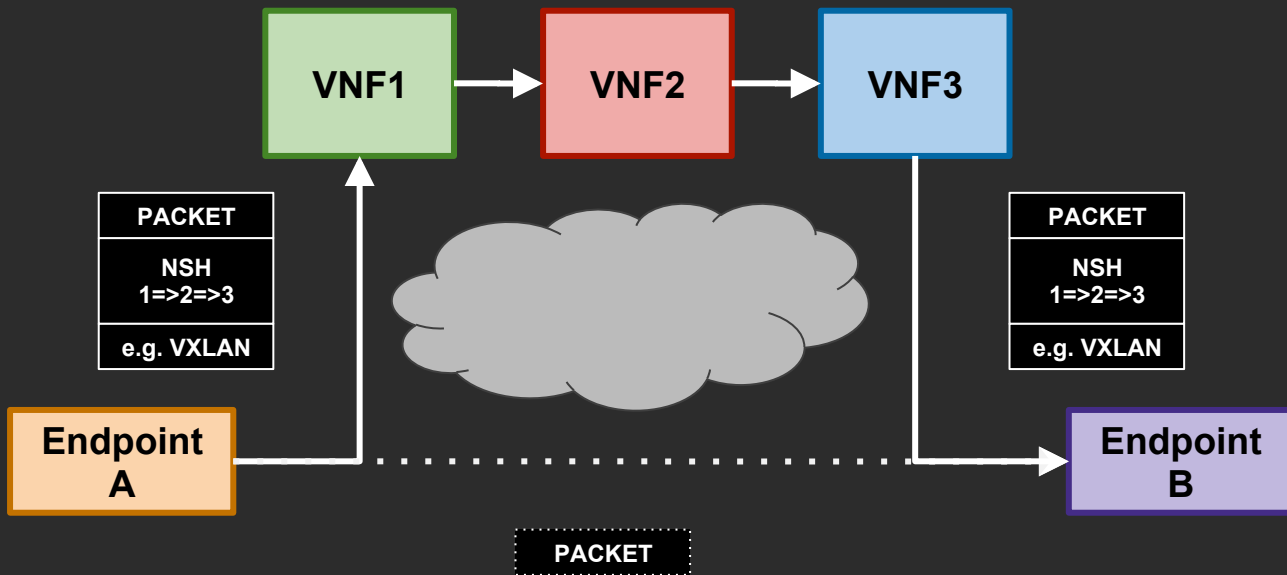
Programmable chains of Network Functions



Why SFC for OVS?

- Service functions are being virtualized, become more scalable and are migrated to live in VMs
- OVS is usually the first forwarding point to redirect packets to a service graph

What is NSH?



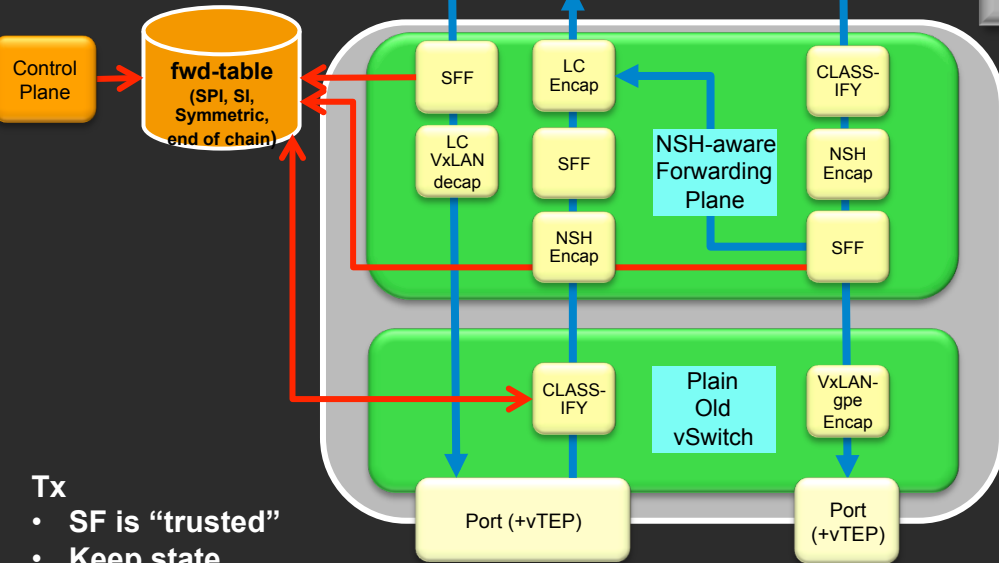
- Carries service graph + service metadata in additional header
- Allows VNFs to exchange metadata

Is NSH just another encap for OVS?

- Almost, but an NSH vport won't work:
 - Guests (VNFs) require to see the NSH headers
 - A push/pop action like VLAN is a better comparison
 - Lightweight Tunnel (LWT) in datapath requires some work to feed encapsulated packets back into OVS for forwarding to VNF

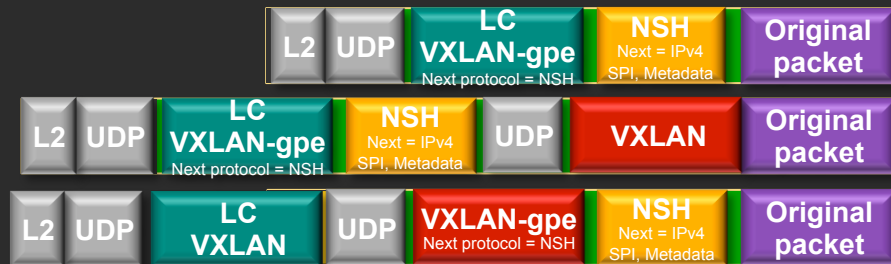
NSH-Based SFC in OvS

Pass metadata / change metadata / add external headers as metadata



Tx


- SF is “trusted”
- Keep state
 - E.g. NAT, SI
- Decap/Encap from local to external
- Shared VNI – VxLAN and gpe?



Rx

- Control of vSwitch Port vTEP decap action?
- OvS recognize & pass VxLAN-gpe and NSH (NSH patch/s...)
- Local Circuit header required to form a “legal” NSH packet
- Ability to pass LC tunnel + NSH headers to SF
 - Header passing exclusively to SFs?
 - How to identify an SF?
- Multi-tenant SF:
 - External VXLAN SHOULD be sent to SF
 - Authority of SFF or OvS?

3 packet format examples



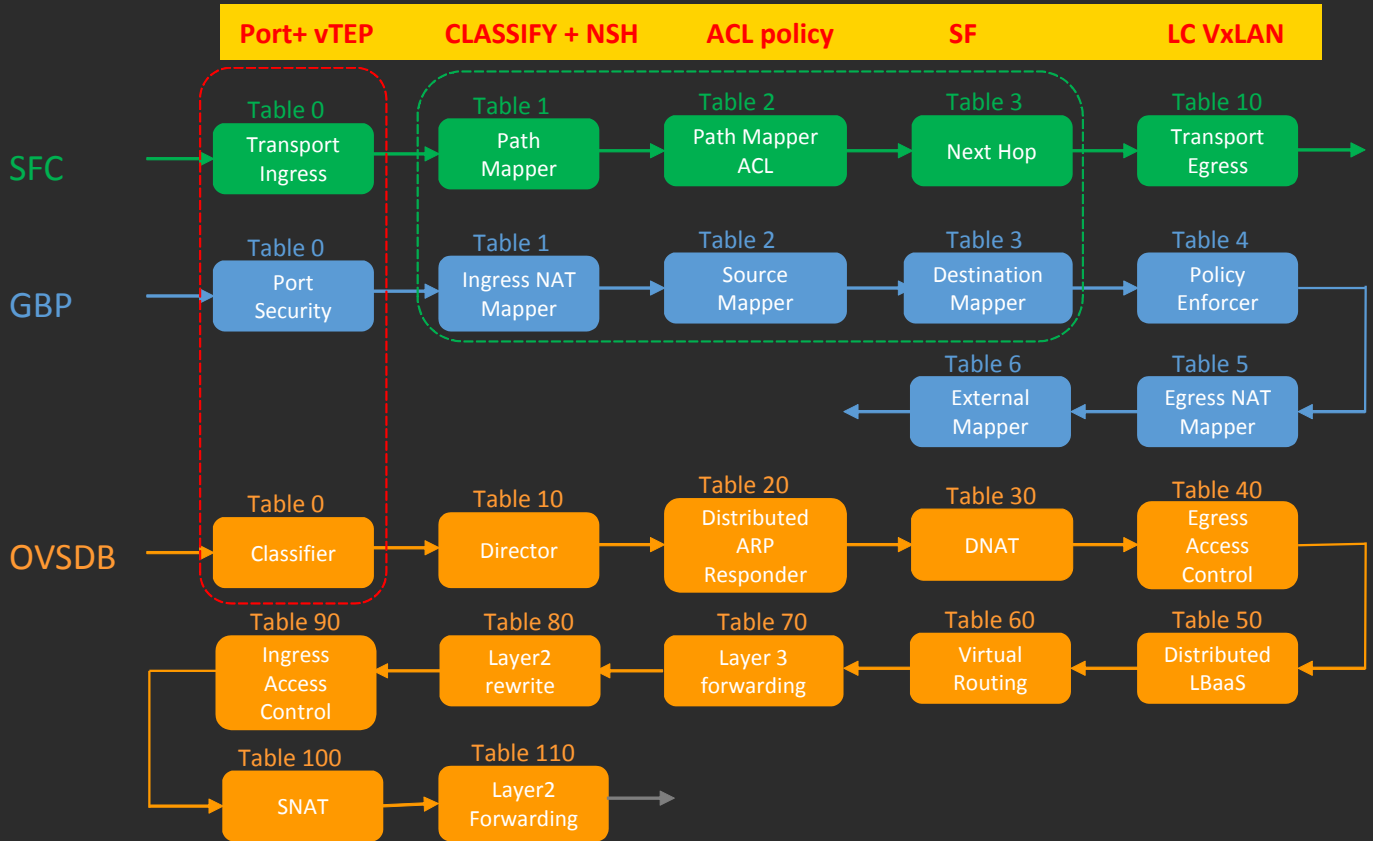
SFC, ODL and OvS OF Tables

Classify impacts
Table 0

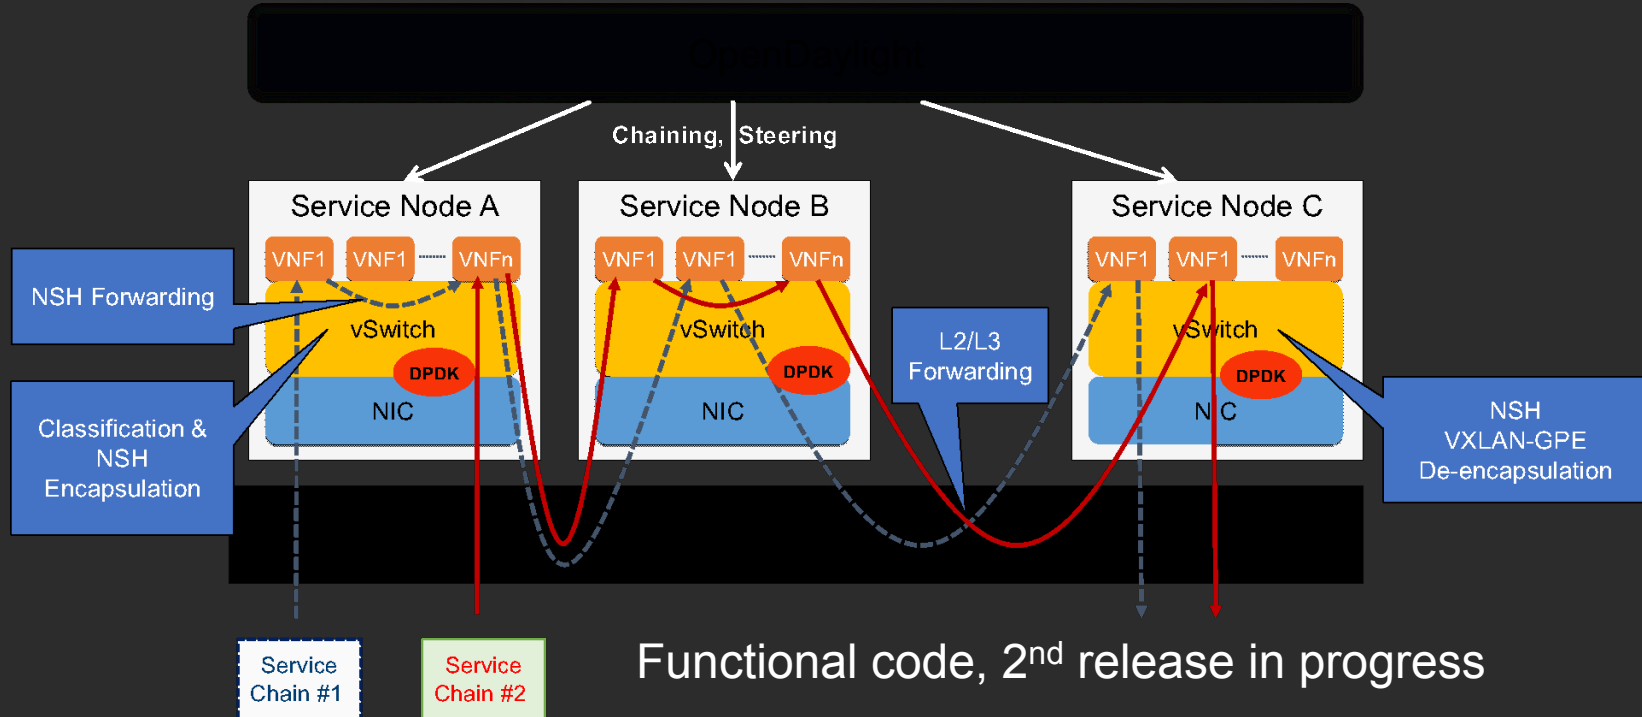
OvS table vs.
ODL
assumptions?

Assumes a single
Controller per
OVS!

OvS Match



SFC in NFV systems with OpenDaylight

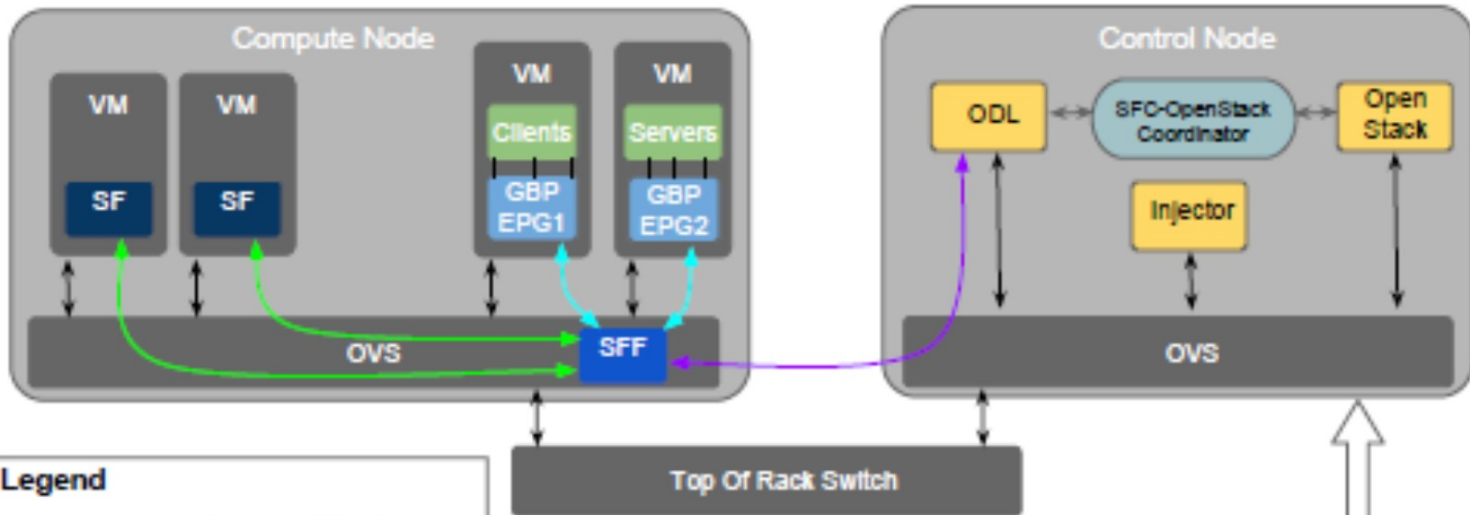


Functional code, 2nd release in progress

Used by OPNFV, allows the orchestrator to create Service Graph

SFC use by OPNFV

OPNFV SFC Initial NW Topology, V2



Legend

- VxLAN NSH tunnel SF/SFF
- GBP creates VxLAN NSH tunnel
- OpenFlow 1.3
- GBP EPG Group Based Policy End Point Group. Replaces SFC classifier

No ODL clustering
in first release

OpenStack SFC API

- Proposed SFC API for OpenStack Neutron -- “networking-sfc”
- <http://docs.openstack.org/developer/networking-sfc/api.html>
- API defines a service chain as:
 - Flow classifier - definition of what traffic enters the chain
 - An ordered list of Neutron ports that define the chain
 - Correlation type -- chain metadata encapsulation type

SFC in OVN

- Status: discussion and early prototyping
 - Prototype based on chaining logical ports on a single logical switch
- Seems like SFC will have a place in OVN
- Lots of questions to answer to come up with a design

SFC in OVN - Metadata

- Metadata in and out of the VM
 - networking-sfc defines use of an MPLS header
 - NSH seems to have the most interest

- Metadata between hosts
 - OVN uses Geneve today
 - Could use vxlan-gpe + NSH in the future

SFC in OVN -- Classifier

- OVN already exposes a nice traffic matching syntax
- We can reuse this for SFC
- See “match” for ACLs in OVN Northbound DB
- Examples:
 - HTTP: `ip && tcp && tcp.dst == {80,443}`
 - SIP: `ip && ((tcp && tcp.dst == {5060,5061}) || (udp && udp.dst == {5060,5061}))`

SFC in OVN's Northbound DB

- Could add as a new action to ACLs
 - Are priorities enough, or do we need separate stages for ACLs and chaining?
- Defining chains
 - Could be arguments to a chain() action
 - Can add new tables for structured chaining definition if needed

Conclusions and next steps

- Asks of OVS for SFC
 - NSH encap/decap, VxLAN-gpe encap/decap, VxLAN to VxLAN-gpe interop
 - Local Circuit and External, control Tunnel port actions!
 - Multi Tenant support – allow external headers and multi VNI to a multi-tenant SF
 - SF privileges vs. VM
 - Expose data plane / local capabilities to orchestrator for best SF placement
- Watch ovs-dev for discussion of OVN SFC design in coming weeks