

VLAN Trunking with OVN

Karthik Chandrashekar,
Nutanix Inc.



WHY

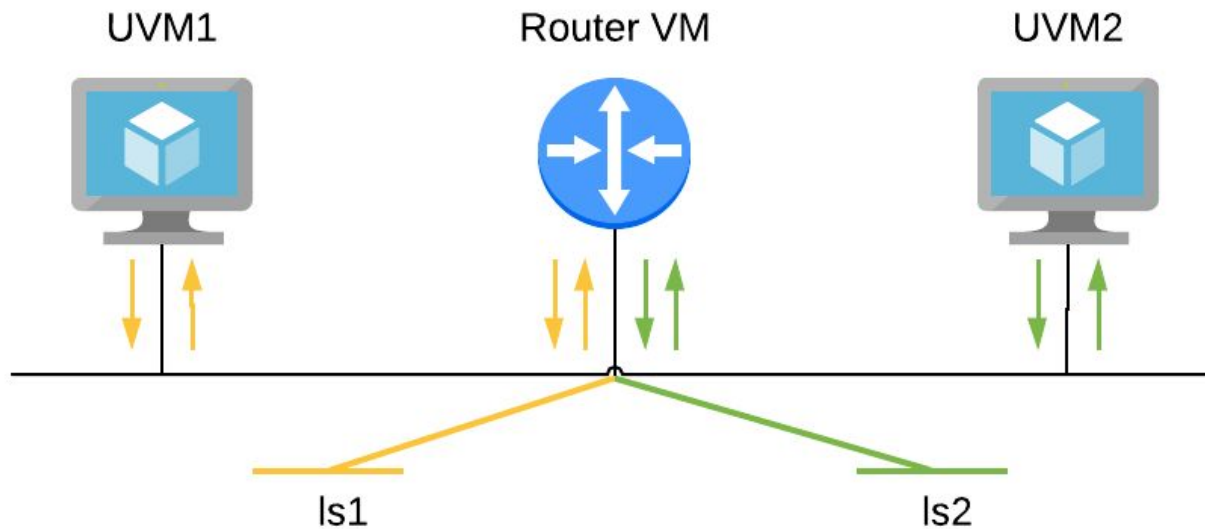


Figure 1: Network function virtualization - VM acting as a router for two different VLANs.

WHY

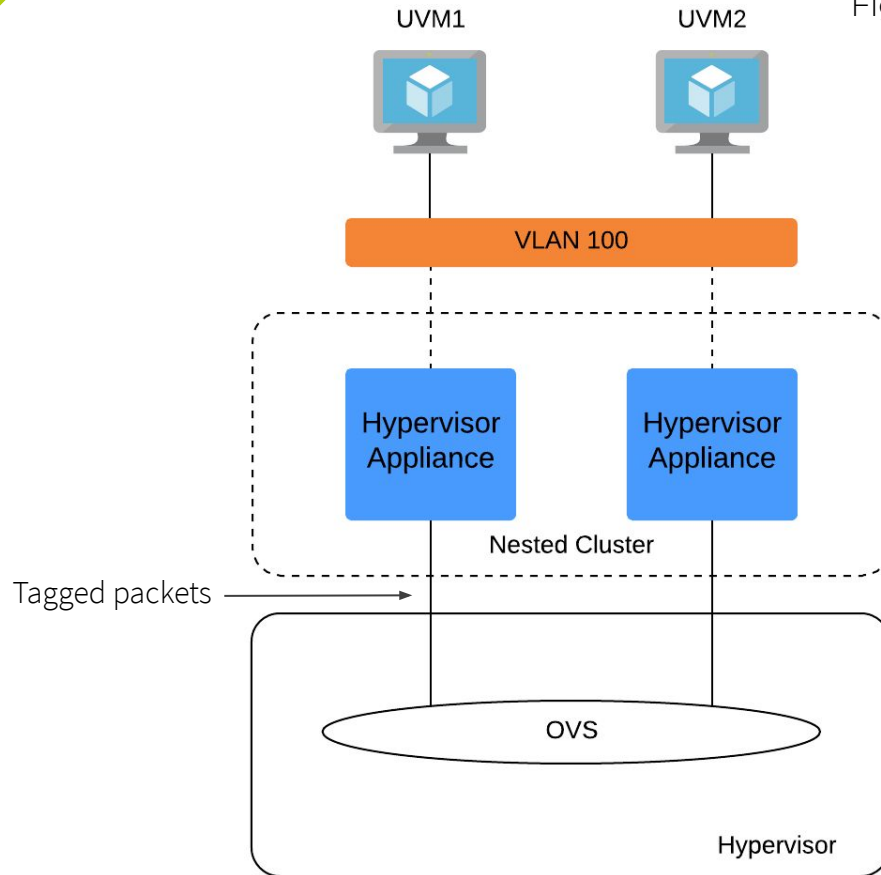


Figure 2: Nested hypervisor scenario.



WHY

WHAT

- Provide same **Layer 2** semantics like that of a physical switch.
- On a given VNIC accept untagged packets and process the packet on that VLAN.
 - This notion already exists today when we create a logical port in a logical switch.
- With trunking configured, a given (VNIC, VLAN) should map to a unique logical switch (broadcast domain) and should process the packet in that logical switch.
 - This notion can be achieved using child ports which exist today to support

Life-Cycle of Container Interface Inside a VM

WHY

WHAT

HOW

- Life-cycle of container interface inside a VM.
 - Multiple container interfaces map to same VIF.
 - Use VLAN tag to differentiate traffic from different containers.

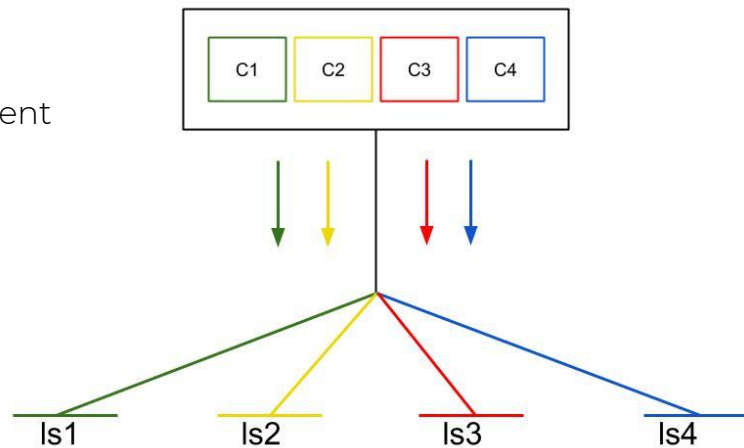


Figure 3: Containers (C1, C2, C3, C4) inside a VM.

WHY

WHAT

HOW

- Life-cycle of container interface inside a VM.
- Each logical switch is associated with a unique VLAN ID (broadcast domain).
- In order to trunk a VNIC on VLAN X, create a child port in logical switch for VLAN X with parent as VNIC and tag X.




Figure 4: VLAN Trunking with child ports.

WHY

WHAT

HOW

- Life-cycle of container interface inside a VM.
- Each logical switch is associated with a unique VLAN ID (broadcast domain).
- In order to trunk a VNIC on VLAN X, create a child port in logical switch for VLAN X with parent as VNIC and tag X.

Flow rule hit for ingress packet

```
cookie=0x0, . . . , table=0, . . . , priority=150, in_port=11, dl_vlan=30  
actions = . . . , strip_vlan, load:0x2713->OXM_OF_METADATA[] . . . ,  
resubmit(,8)
```

Flow rule hit for egress packet

```
cookie=0x0, . . . , table=65, . . . priority=100, metadata=0x2713  
actions=mod_vlan_vid:30, output:11
```

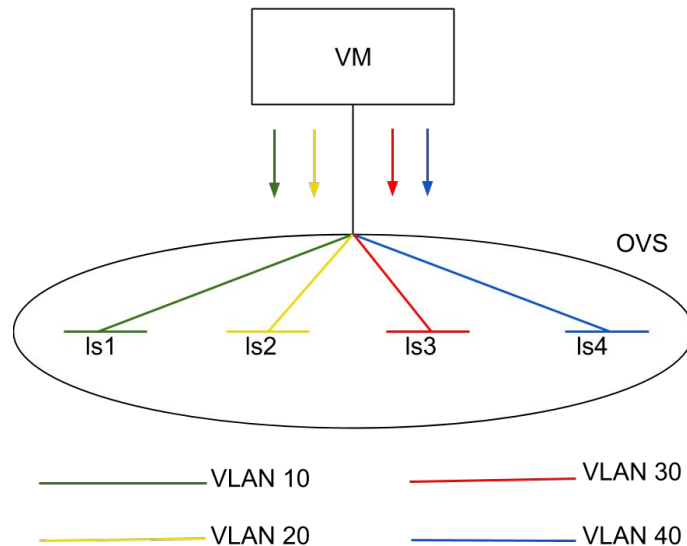


Figure 4: VLAN Trunking with child ports.

WHY

WHAT

HOW

- Life-cycle of container interface inside a VM.
- Each logical switch is associated with a unique VLAN ID (broadcast domain).
- In order to trunk a VNIC on VLAN X, create a child port in logical switch for VLAN X with parent as VNIC and tag X.
- Disable port security on trunked interface.




Figure 4: VLAN Trunking with child ports.

Thank you

