

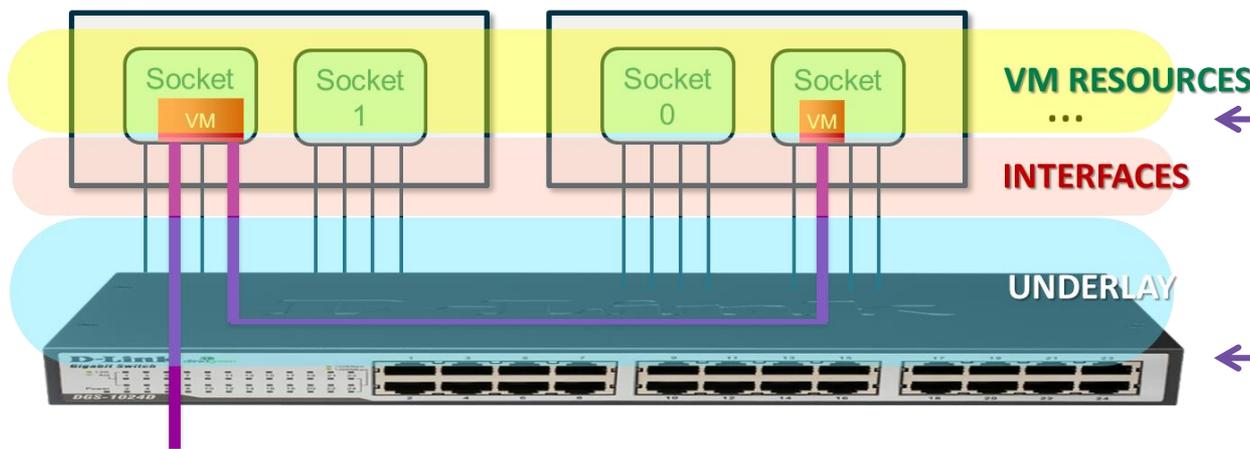
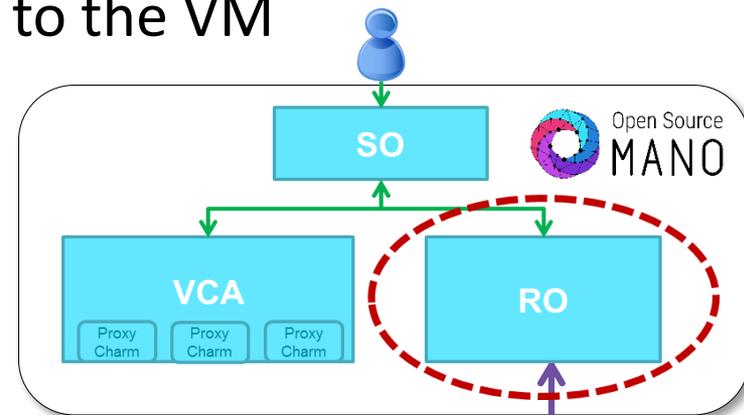
Open Source
MANO

OSM Hackfest – Session 4 Modeling EPA capabilities in VNF

Gerardo García (Telefónica)
Guillermo Calviño (Altran)

EPA support combined with SDN Assist enables chaining of high performance VNFs

1. Accurate assignment of resources at VM level
2. Proper assignment of I/O interfaces to the VM
3. **SDN gives the ability to create underlay L2 connections**
 - Interconnecting VMs
 - Attaching external traffic sources



- **EPA features** like use of large hugepages memory, dedicated CPUs, strict NUMA node placement, and the use of passthrough and SR-IOV interfaces, **can be used in OSM's VNF descriptors since Rel Zero.**
- If your VIM supports EPA, then you don't need to do anything extra to use it from OSM. VIM connectors in OSM take advantage of EPA capabilities if the VIM supports it. All you need to do is build your descriptors and deploy.
- Openstack configuration for EPA (reference guide):
 - [https://osm.etsi.org/wikipub/index.php/Openstack_configuration_\(Release_THREE\)#Configure_Openstack_for_OSM_.28EPA.29](https://osm.etsi.org/wikipub/index.php/Openstack_configuration_(Release_THREE)#Configure_Openstack_for_OSM_.28EPA.29)

- This feature allows to use an external controller to create the underlying connectivity
- Wiki page:
 - https://osm.etsi.org/wikipub/index.php/EPA_and_SDN_assist
- Requirements:
 - A dataplane switch with Openflow capabilities that will connect the physical interfaces of the VIM compute nodes.
 - An external SDN controller controlling the previous dataplane switch.
 - The mapping between the switch ports (identified by name) and the compute node interfaces (identified by host-id and PCI address)
 - Some VIMs as Openstack requires admin credentials in order to be able to get the physical place of the SRIOV/passthrough VM interfaces

SDN assist. The way it works

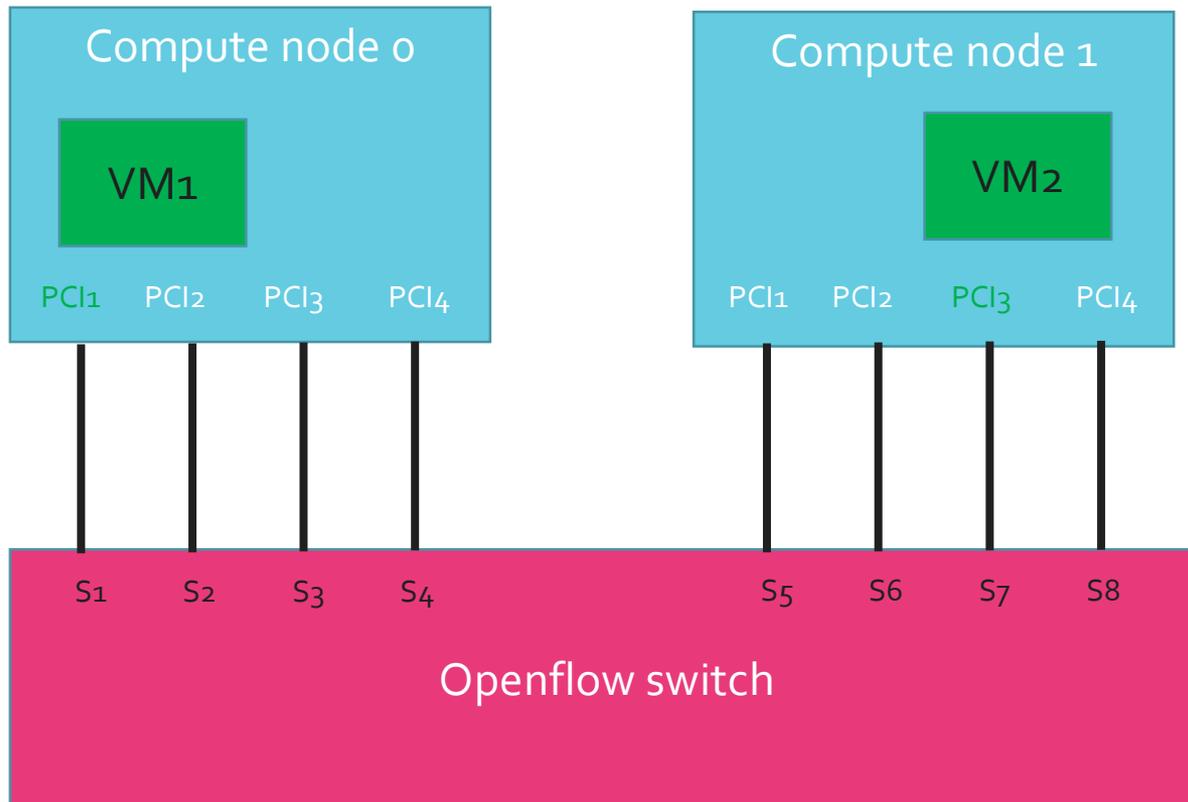


- OSM will deploy the VMs of a NS with Passthrough and/or SRIOV interfaces
- OSM will get from the VIM (in your case, Openstack) the compute node where the VM was deployed and the physical interface assigned to the VM (identified by its PCI address).
- OSM will map those interfaces to Openflow ports in the switch making use of the mapping that you should have introduced in the system
- OSM will create the dataplane networks by talking to the SDN controller and connecting the appropriate Openflow ports to the same network.

SDN assist. The way it works

1) VMs are deployed

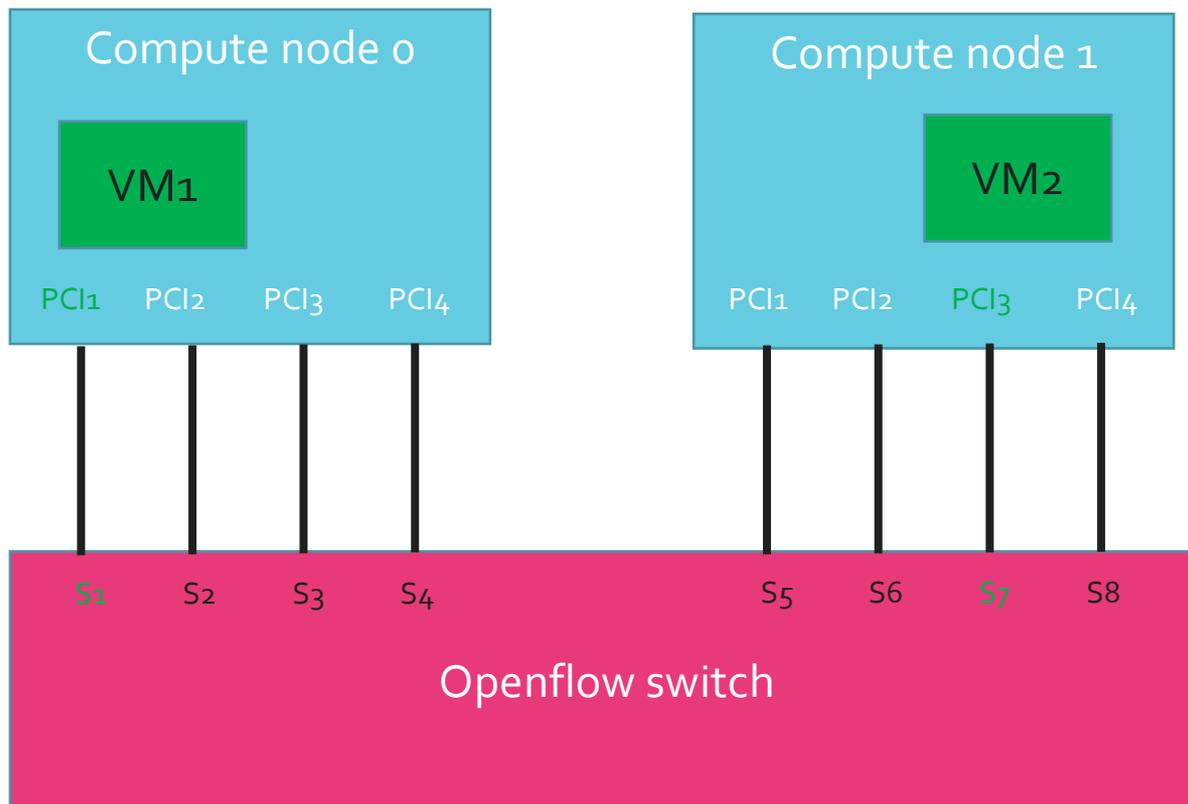
VM1 will be e.g. deployed in compute node 0 and will consume interface identified by PCI address PCI1 and MAC MACX



VM2 will be e.g. deployed in compute node 1 and will consume interface identified by PCI address PCI3 and MAC MACY

SDN assist. The way it works

2) OSM uses port mapping to identify the ports in the switch

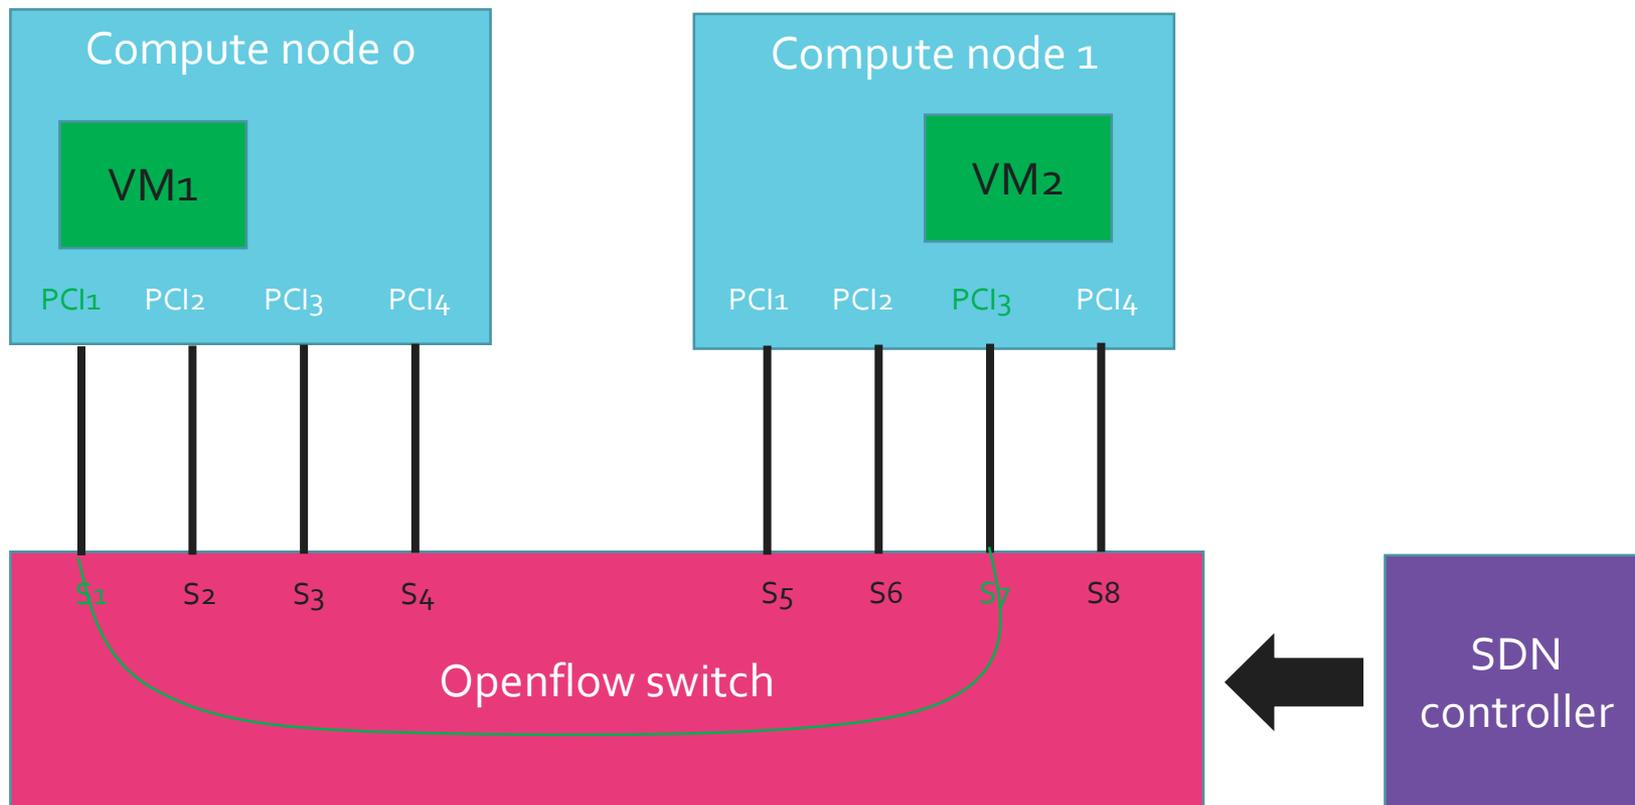


The interface at PCI address PCI1 of compute node 0 is connected to port S1 in the switch

The interface at PCI address PCI3 of compute node 1 is connected to port S7 in the switch

SDN assist. The way it works

3) OSM talks to the SDN controller and connect the ports to the same network.



Example of rules for E-LINE and PCI passthrough interfaces:

IN: Port S1 → OUT: Port S7

IN: Port S7 → OUT: Port S1

Device role tagging

- Whenever a new NS is deployed, and if the VIM allows it, VM interfaces will be tagged with the interface name specified in the descriptor
- This allows proper identification of interfaces in the VM
- A service and a script have to be pre-installed in the VM:
 - RedHat-based VMs: https://github.com/oglok/udev_data
 - Ubuntu-based VMs: https://github.com/gcalvino/udev_data

Adding new VIM account: openstack2-epa

- VIM:
 - openstack2-epa: 172.21.2.22
- Test VIM:
 - ping <IP>
 - curl http://<IP>:5000/v2.0
 - Load Openstack credentials:
 - export OS_AUTH_URL=http://172.21.2.22:5000/v2.0
 - export OS_USERNAME=xxx
 - export OS_TENANT_NAME=xxx
 - export OS_PASSWORD=xxx
 - Run some commands:
 - openstack image list
 - openstack network list
 - openstack flavor list
 - openstack server list

Adding new VIM account: openstack2-epa

- Add your second VIM 'openstack2-epa' with the OSM client:
 - `osm vim-create --name openstack2-epa --account_type openstack \`
`--auth_url http://172.21.2.22:5000/v2.0 \`
`--user xxx --password xxx --tenant xxx \`
`--description "ETSI openstack site 2, with EPA, with tenant xxx" \`
`--config '{dataplane_physical_net: physnet_sriov, microversion: 2.32}'`
 - `osm vim-list`
 - `osm vim-show openstack2-epa`
- Config options:
 - `dataplane_physical_net`:
 - Used to instantiate VMs with SR-IOV and Passthrough interfaces
 - Value: The physical network label used in Openstack both to identify SRIOV and passthrough interfaces (nova configuration) and also to specify the VLAN ranges used by SR-IOV interfaces (neutron configuration).
 - `microversion`:
 - Used for device role tagging
 - Value: 2.32

Adding an SDN controller

- Try 'osm' and look for an option to add SDN controller
 - No option will be found
- SDN controller has to be added through the RO
- Log into the RO container and run:
 - `openmano sdn-controller-create etsi_DSS9000_fl --ip 172.21.2.23 --port 8080 --dpid 00:01:64:00:6a:e6:b3:14 --type floodlight`
 - `openmano datacenter-edit -f openstack2-epa --sdn-controller etsi_DSS9000_fl`
 - `openmano datacenter-list openstack2-epa -vvv`

Adding a port-mapping

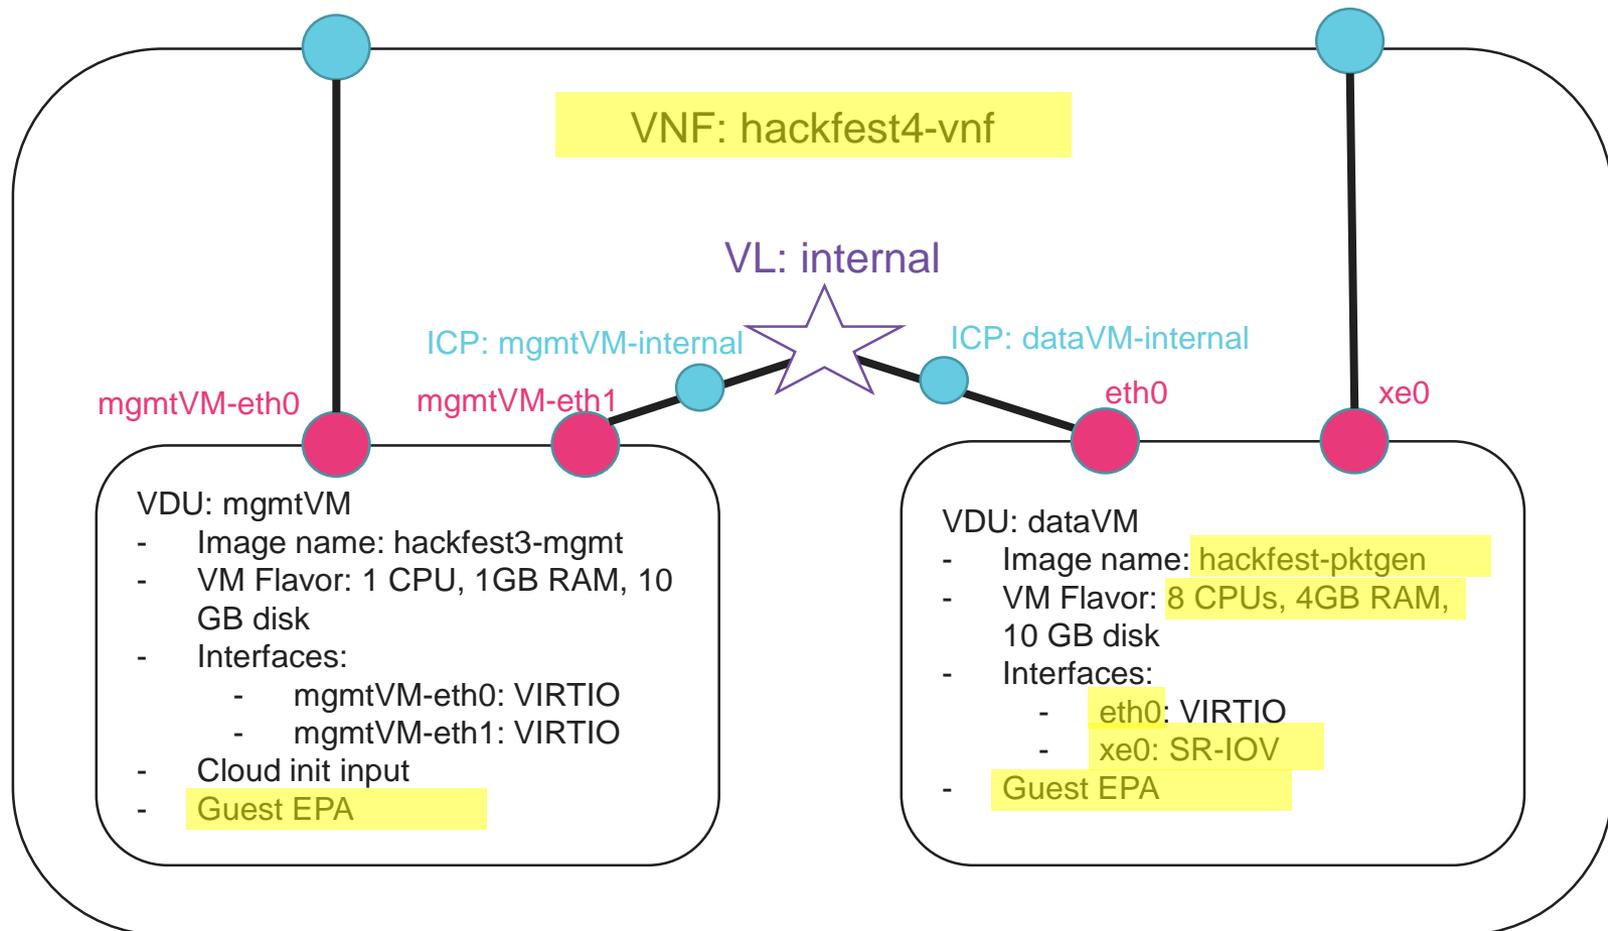
- Port-mapping has to be added through the RO
- Port mapping file:
 - <https://osm-download.etsi.org/ftp/osm-3.0-three/2nd-hackfest/other/port-mapping-openstack2.yaml>
- From the OSM host, download port-mapping and push it to RO container:
 - `wget https://osm-download.etsi.org/ftp/osm-3.0-three/2nd-hackfest/other/port-mapping-openstack2.yaml`
 - `lxc file push port-mapping-openstack2.yaml RO/root/port-mapping-openstack2.yaml`
- Log into the RO container and run:
 - `openmano datacenter-sdn-port-mapping-set openstack2-epa port-mapping-openstack2.yaml`
 - `openmano datacenter-sdn-port-mapping-list openstack2-epa`

VNF diagram

Changes highlighted in yellow

External Connection point: vnf-mgmt

External Connection point: vnf-data



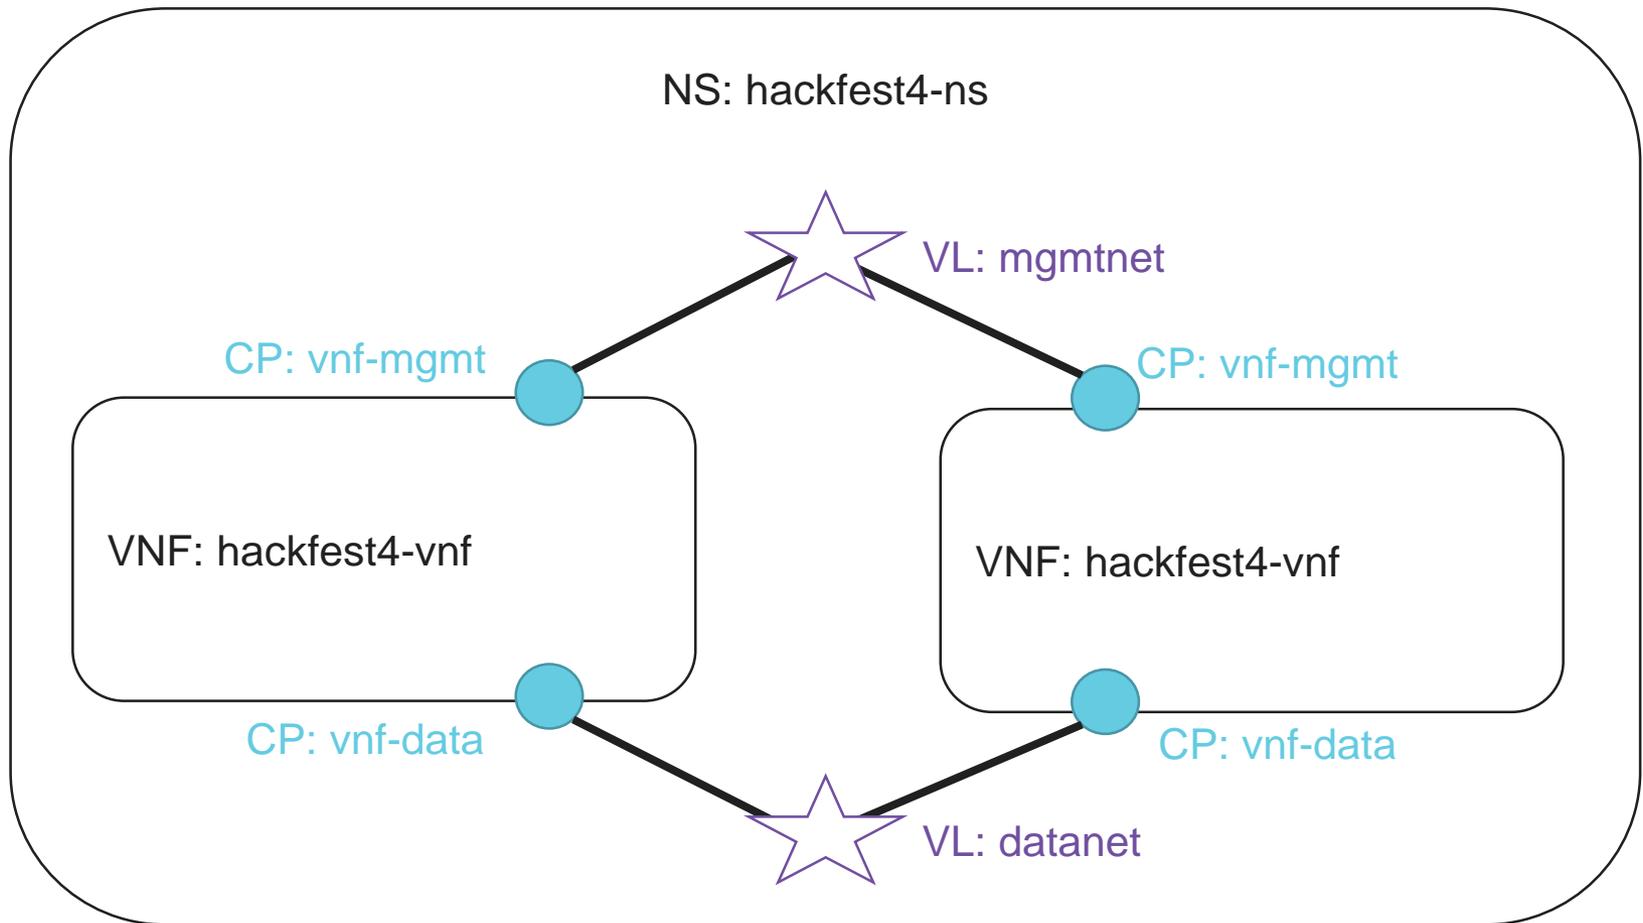
Creating the VNF in the UI (1/2)

- Go to the catalog
- Copy VNF hackfest3
- Modify the new VNF
 - Name: hackfest4-vnf
- Modify VDU mgmtVM:
 - Guest EPA
 - Mepage size: LARGE
 - CPU pinning policy: DEDICATED
 - CPU thread pinning policy: **PREFER (HW threads)** or ISOLATE (cores)
 - NUMA Policy: numa-aware
 - Node CNT: 1
 - Mem policy: STRICT

Creating the VNF in the UI (2/2)

- Modify VDU dataVM:
 - Image name: hackfest-pktgen
 - Flavor:
 - 8 CPUs
 - 4096 MB RAM
 - 10 GB disk
 - Guest EPA
 - Mepage size: LARGE
 - CPU pinning policy: DEDICATED
 - CPU thread pinning policy: **PREFER (HW threads)** o ISOLATE (cores)
 - NUMA Policy: numa-aware
 - Node CNT: 1
 - Mem policy: STRICT
 - Modify interfaces of the VDU
 - Interface 1:
 - Name: eth0
 - Interface 2:
 - Name: xe0
 - Virtual-interface:
 - Type: SR-IOV
- Click on UPDATE to save your VNF

NS diagram



Creating the NS in the UI(1/2)

- Go to the catalog
 - Add NSD
 - Name: hackfest4-ns
 - Add 2 VNFs (hackfest4-vnf) by drag and drop
 - Add first VLDs:
 - VLD1:
 - name (optional): mgmtnet
 - TYPE: ELAN
 - MGMT NETWORK: True
 - INIT PARAMS
 - vim-network-ref
 - VIM NETWORK NAME: mgmt
- <- This is to have a default mapped VIM network name

Creating the NS in the UI(2/2)

- Add second VLD:
 - VLD2:
 - name (optional): datanet
 - TYPE:ELAN
 - MGMT NETWORK: False (default)
- Connect VNF Connection Points to the VLs:
 - vnf-mgmt to VLD:mgmtnet
 - vnf-data to VLD:datanet
- Click on UPDATE

Deploying NS in the UI (1/3)

- Go to Launchpad > Instantiate
- Select hackfest4-ns and click Next
- Complete the form
 - Add a name to the NS
 - Select the Datacenter where the NS will be deployed (openstack2-epa)
 - Add SSH key
- Go to the dashboard to see the instance and get the mgmt IP address of the VNF

Deploying NS in the UI (2/3)

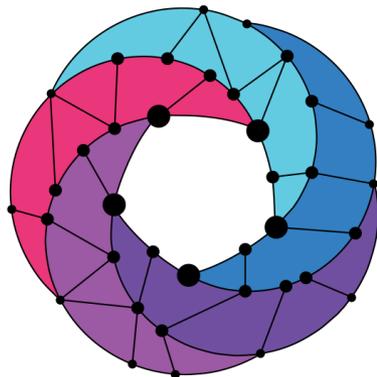
- SSH to the mgmtVM of one of the VNFs
- Check that cloud-init worked in mgmtVM
- Check interface names in mgmtVM
 - They must be different than the ones in the descriptor (no udev metadata service)
- Jump to dataVM:
 - `ssh -i test4.pem <IP_dataVM>`
- Check interface names in mgmtVM
 - They must be the same names as in the descriptor
- Run 'pktgen'

Deploying NS in the UI (3/3)

- SSH to the mgmtVM of the second VNF
- Check that cloud-init worked in mgmtVM
- Check interface names in mgmtVM
 - They must be different than the ones in the descriptor (no udev metadata service)
- Jump to dataVM:
 - `ssh -i test4.pem <IP_dataVM>`
- Check interface names in mgmtVM
 - They must be the same names as in the descriptor
- Run 'pktgen'

Instructions to run 'pktgen'

- Configure interfaces to use DPDK:
 - `cd /opt/dpdk-17.11`
 - `sudo -E ./usertools/dpdk-devbind.py --status`
 - `sudo -E ./usertools/dpdk-devbind.py --bind=igb_uio xe0`
- Run pktgen:
 - `cd /opt/pktgen-3.4.5/`
 - `sudo -E ./app/x86_64-native-linuxapp-gcc/pktgen -n 3 -l 2-7 -- -P -m "[4-7].0"`
- From pktgen console, set the destination MAC address to the one used by xe0 interface in the other VNF, and send traffic:
 - `set all dst mac xx:xx:xx:xx:xx:xx` <- destination MAC where to send traffic
 - `start all` <- to start traffic
 - `stop all` <- to stop traffic



Open Source
MANO

Find us at:

osm.etsi.org
osm.etsi.org/wikipub