

Open Source  
**MANO**


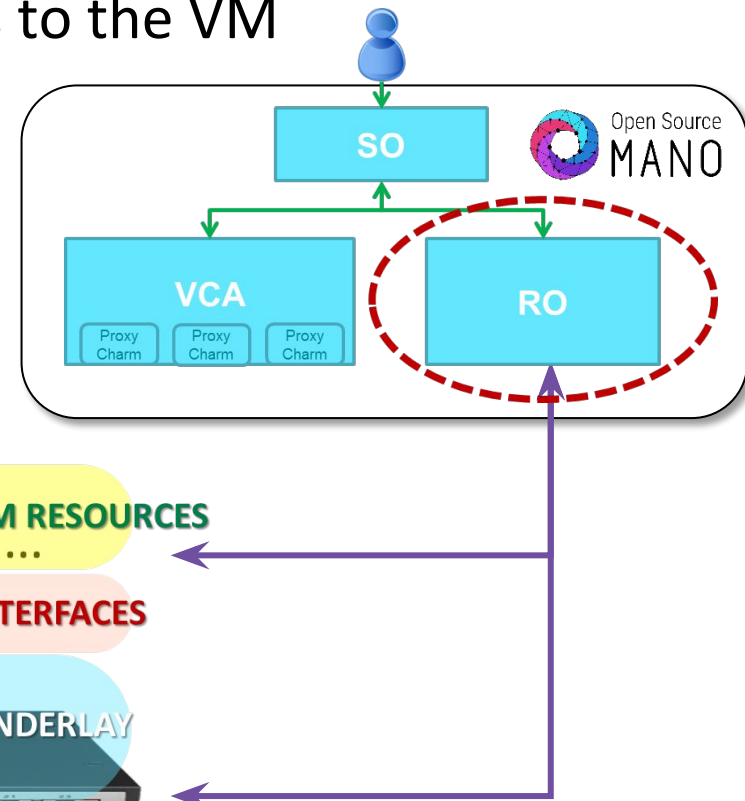
# OSM Hackfest – Session 5

## Modeling EPA capabilities in VNF

Eduardo Sousa (Whitestack)

# 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)

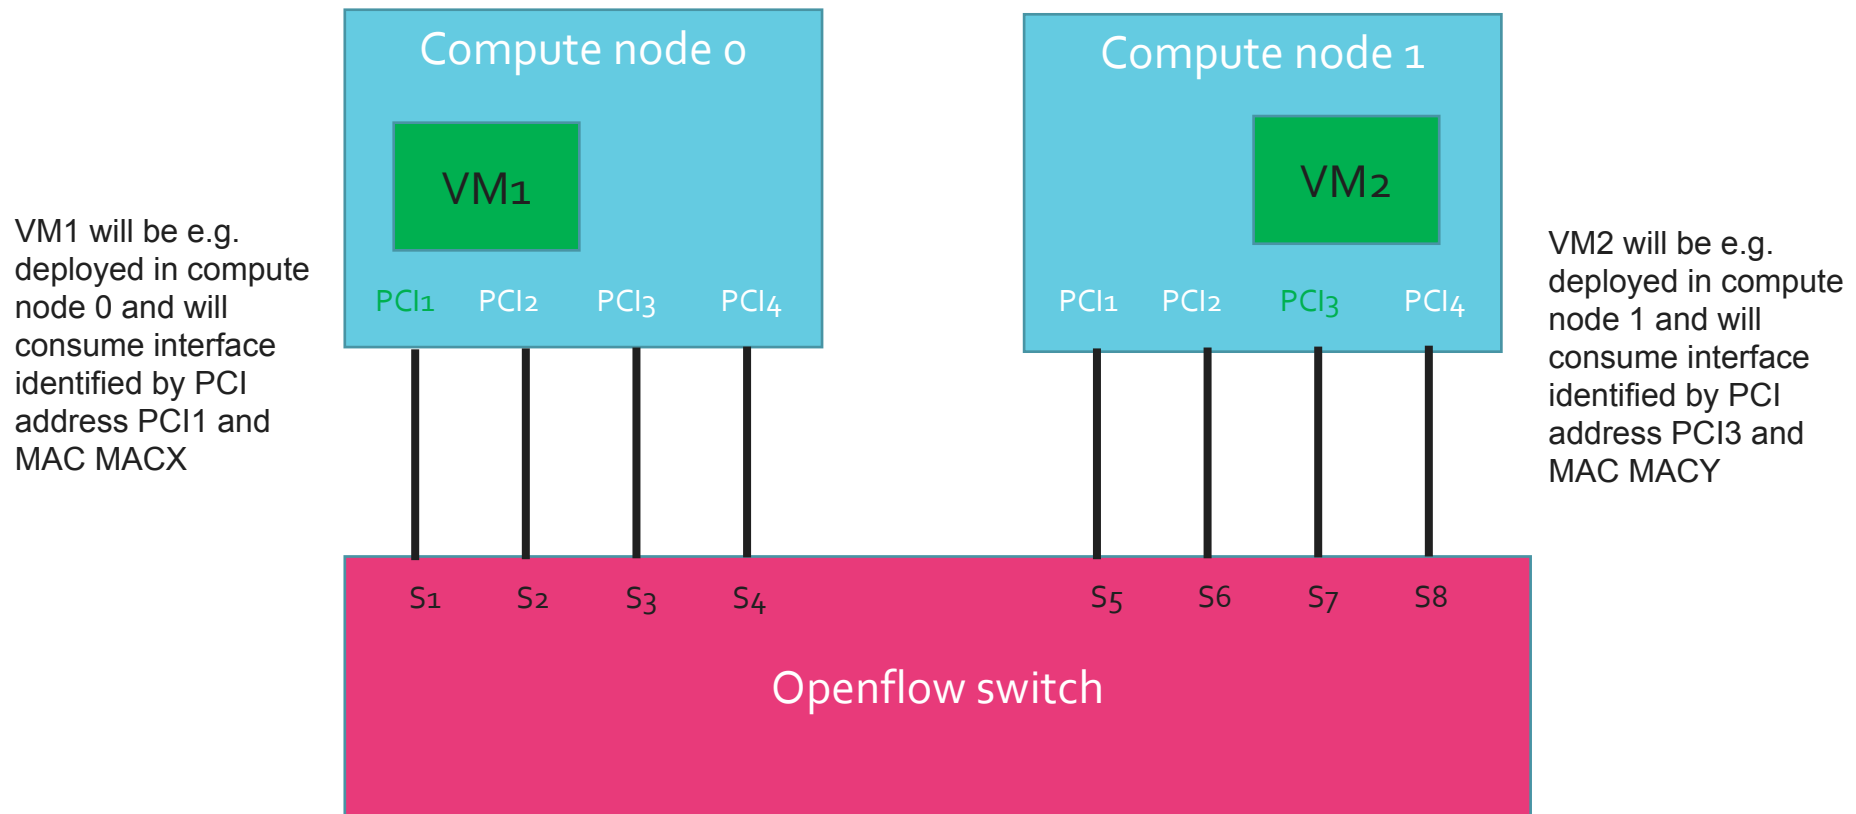
- This feature allows to use an external controller to create the underlying connectivity
- Wiki page:
  - [https://osm.etsi.org/wikipub/index.php/EPA\\_and\\_SDN\\_assist](https://osm.etsi.org/wikipub/index.php/EPA_and_SDN_assist)
- Requirements:
  - A dataplane switch with Openflow capabilities that 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.

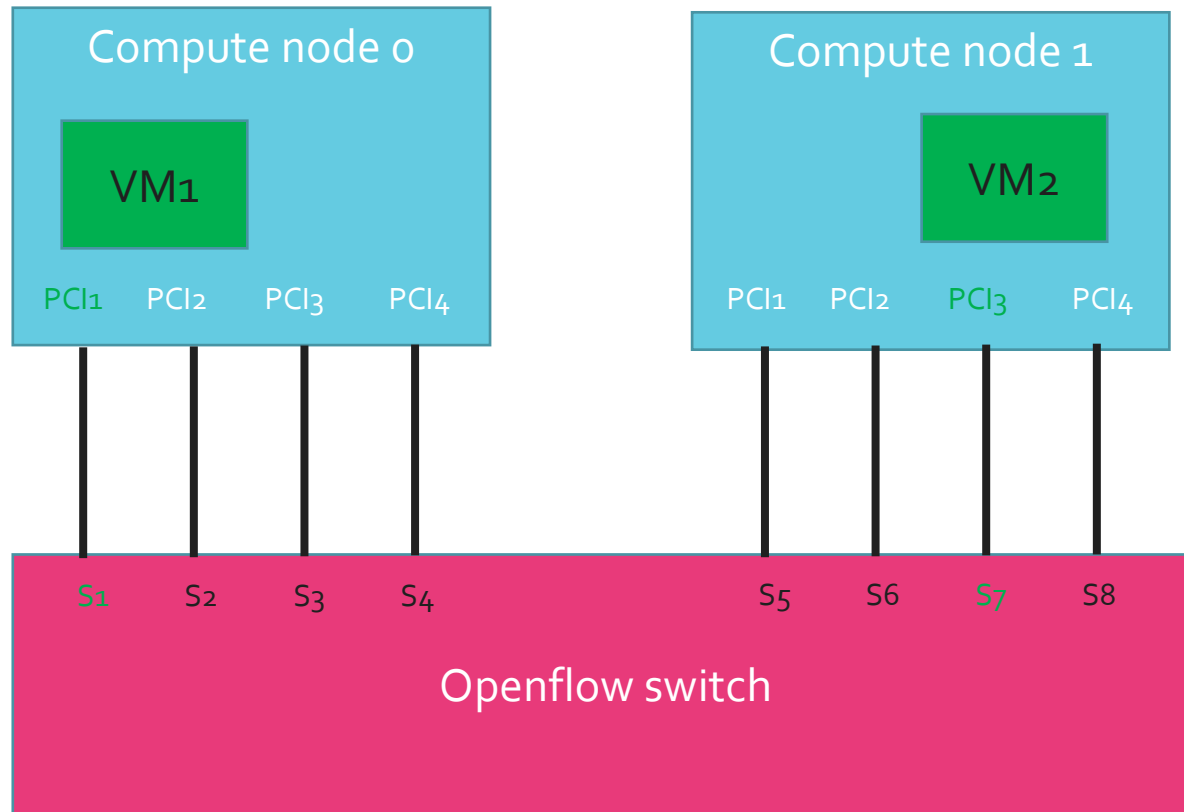
# SDN assist. The way it works

## 1) VMs are deployed



# SDN assist. The way it works

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

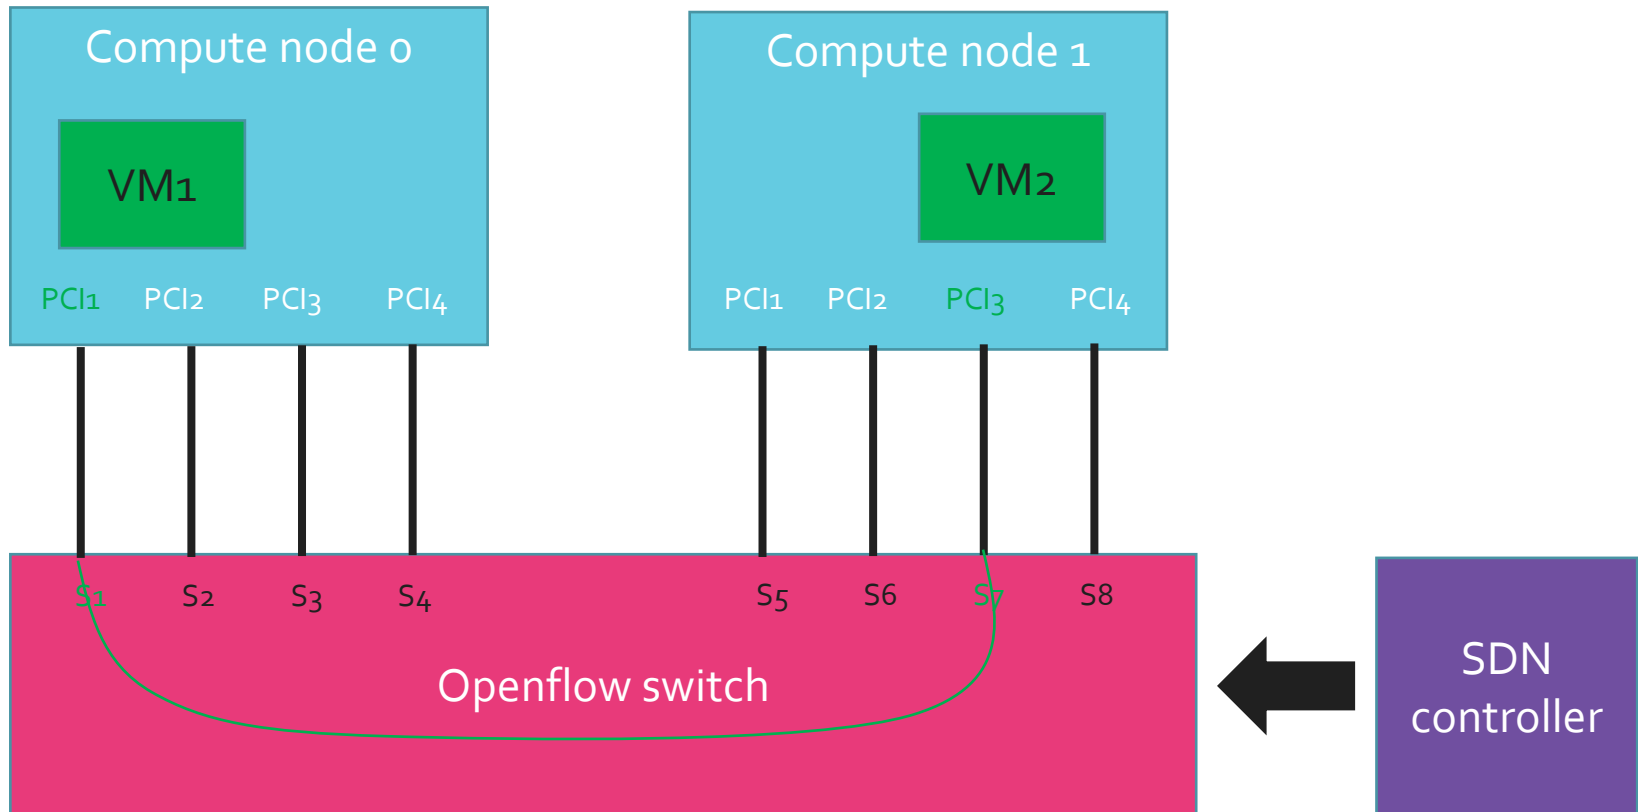


The interface at PCI address PCI1 of compute node 0 is connected to port S1 in the switch

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 S<sub>1</sub> → OUT: Port S<sub>7</sub>

IN: Port S<sub>7</sub> → OUT: Port S<sub>1</sub>



# 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](https://github.com/oglok/udev_data)
  - Ubuntu-based VMs: [https://github.com/gcalvino/udev\\_data](https://github.com/gcalvino/udev_data)



Open Source  
**MANO**

Find us at:  
[osm.etsi.org](http://osm.etsi.org)  
[osm.etsi.org/wikipub](http://osm.etsi.org/wikipub)