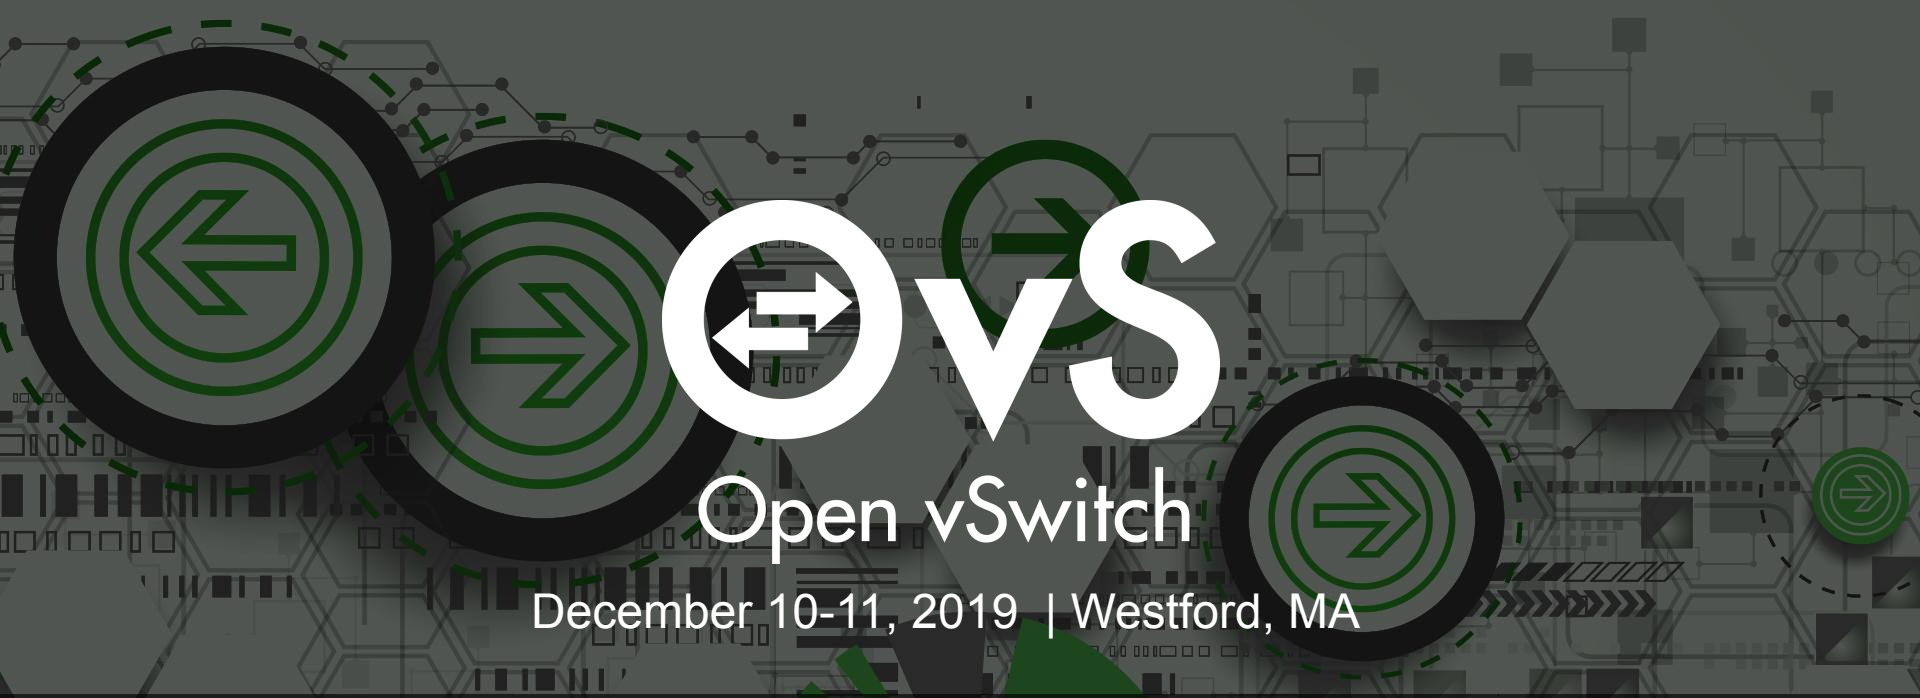

December 10-11, 2019 | Westford, MA

IP Multicast in OVN
IGMP Snooping and Relay

Dumitru Ceara, Red Hat

Problems with IP Multicast in OVN

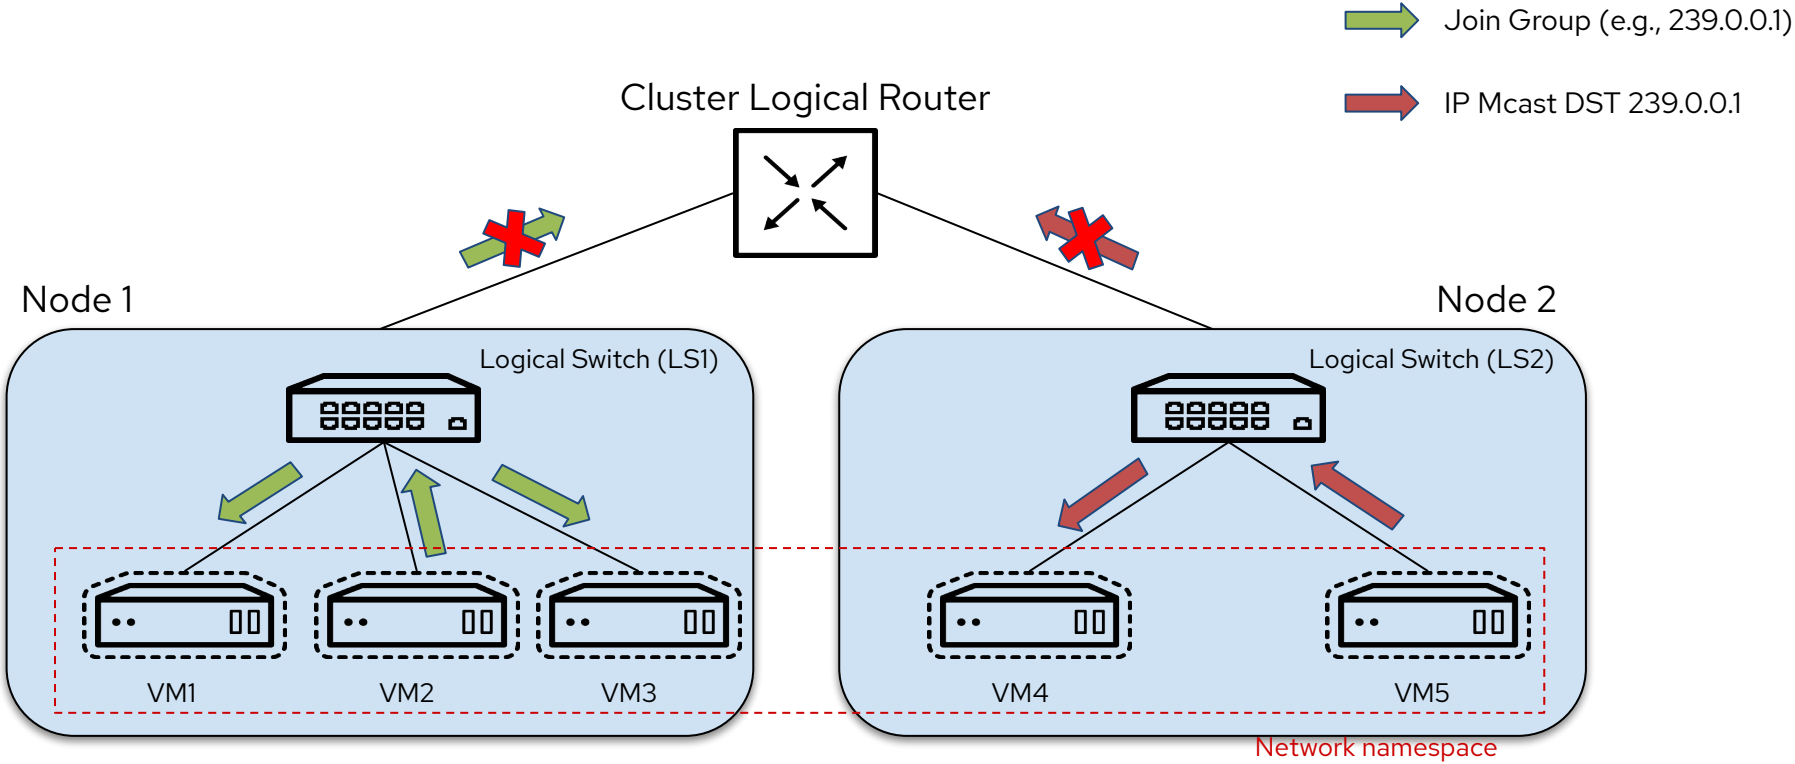

"Traditional" networking use case

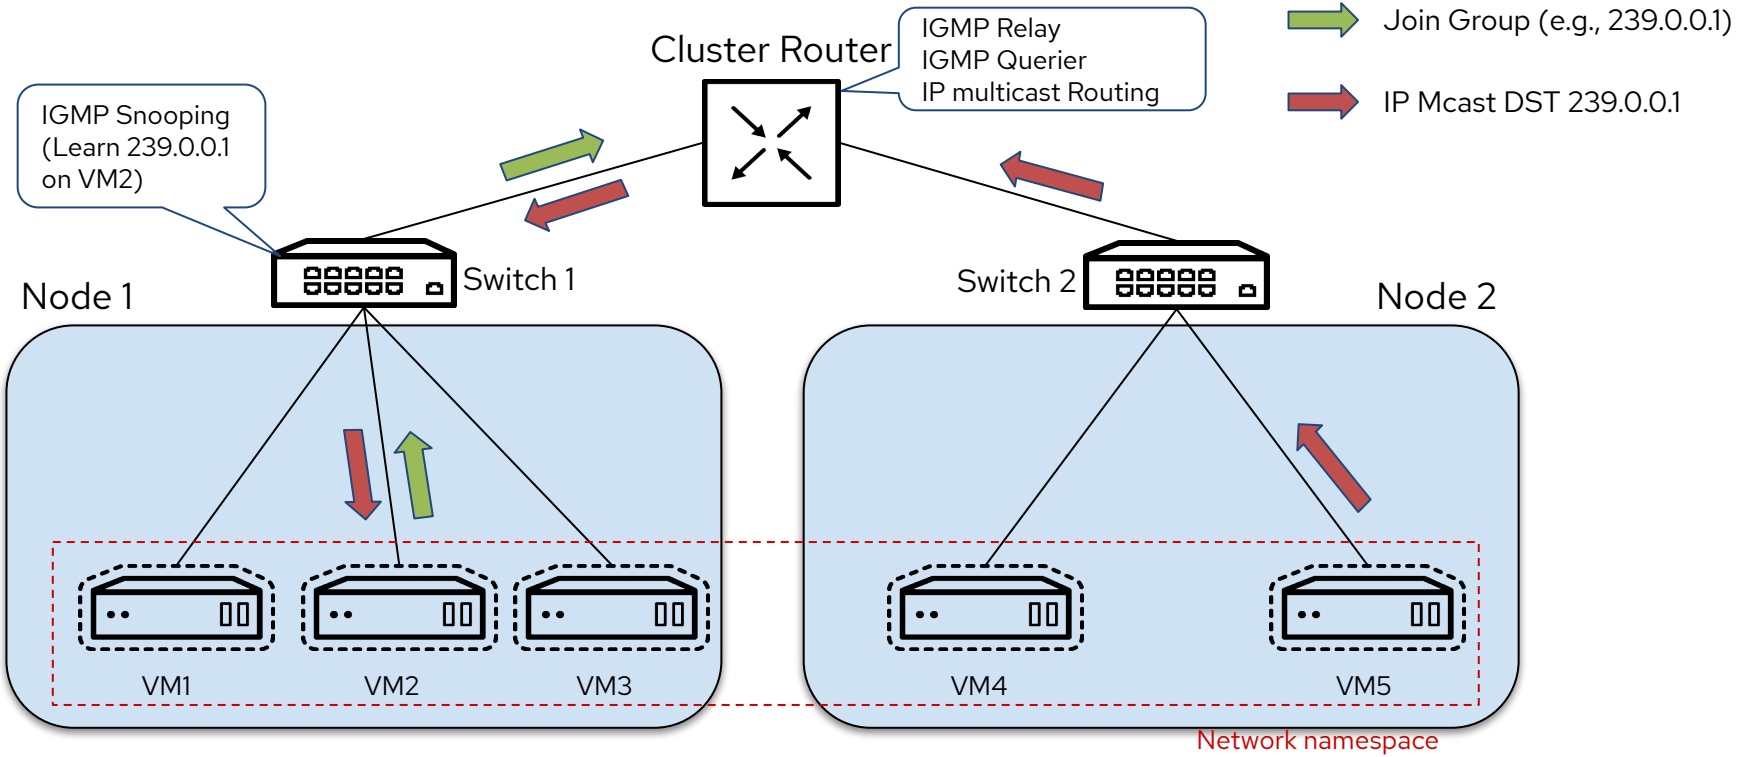

IP Multicast in an OVN logical switch (no snoop)

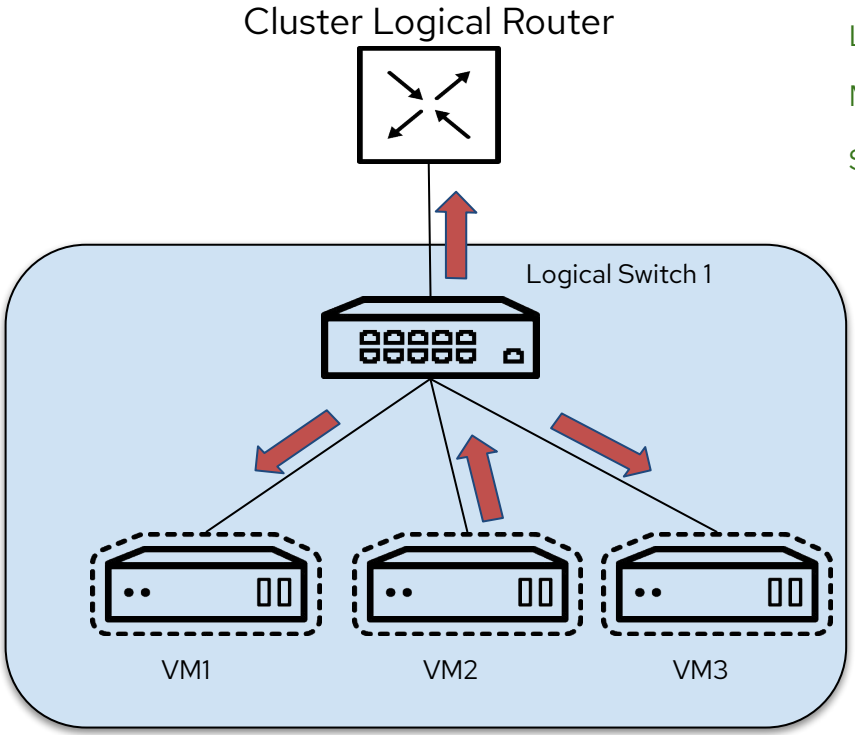

IP Mcast DST 239.0.0.1

Logical_Flow: **match=ip4.mcast**, **action=output(MC_FLOOD)**

Multicast_Group: **name=MC_FLOOD**, **ports={vm1,vm2,vm3,sw1-lrp}**

Simplified OVS Flows (in sequence):

SetReg(LogOutPort=vm1),resubmit(egress_pipeline)
SetReg(LogOutPort=vm2),resubmit(egress_pipeline)
SetReg(LogOutPort=vm3),resubmit(egress_pipeline)
SetReg(LogOutPort=sw-lrp),resubmit(egress_pipeline)

Issues:

- Flood to all switch ports
- Multicast_group implementation requires sequentially executing the pipeline for each port in the flood group (4k resubmit limit)

IP Multicast in an OVN logical switch (with snoop)

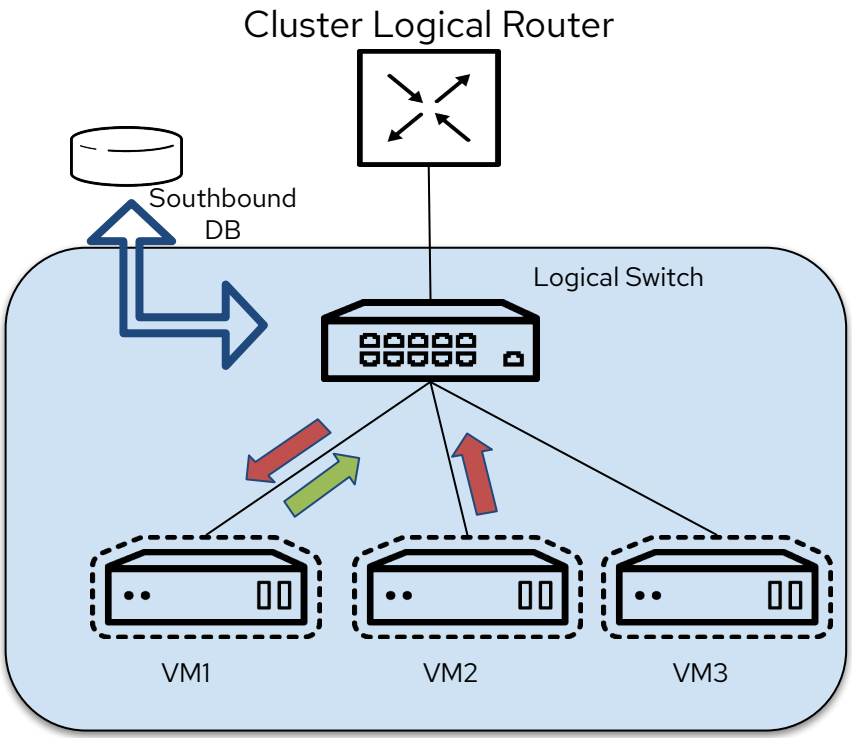

Join Group (e.g., 239.0.0.1)

Logical_Flow: **match=igmp,action=controller(...)**
IGMP_Group: **datapath=LS,ports={vm1},chassis=Node1**

ovn-northd (update Southbound DB):
Multicast_Group(MC-239.0.0.1, ports={vm1})
Logical_Flow(239.0.0.1, outport(MC-239.0.0.1)

IP Mcast DST 239.0.0.1

Logical_Flow: **match=ip4.dst==239.0.0.1,action=output(239.0.0.1)**

Multicast_Group: **name=239.0.0.1, ports={vm1}**

Simplified OVS Flows (in sequence):
SetReg(LogOutPort=vm1),resubmit(egress_pipeline)

Quite similar to “traditional” snooping!

Now let's route..

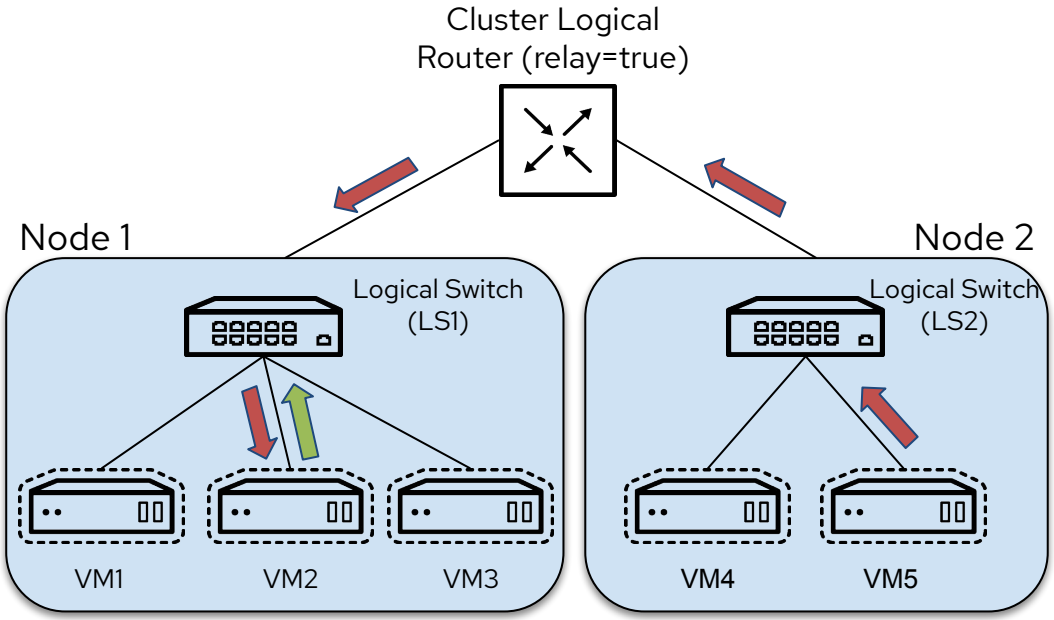

Join Group (e.g., 239.0.0.1)

Logical_Flow: **match=igmp,action=controller(...)**
IGMP_Group: **datapath=LS,ports={vm1},chassis=Node1**

ovn-northd (update Southbound DB):
Multicast_Group(LS1,MC-239.0.0.1,ports={vm1})
Logical_Flow(LS1,239.0.0.1,outport(MC-239.0.0.1)

Multicast_Group(LR, MC-239.0.0.1, ports={rtr-ls1})
Logical_Flows(239.0.0.1,outport(MC-239.0.0.1),dec_ttl, update_mac)

IP Mcast DST 239.0.0.1

Logical_Flow: **LS2,match=ip4.mcast, action=clone(MC_MROUTER_FLOOD)**
Multicast_Group: **LS2,name=MC_MROUTER_FLOOD, ports={sw2-lrp}**
Logical_Flow: **LR,match=ip4.dst==239.0.0.1, action=outport(MC-239.0.0.1)**
Multicast_Group: **LR,name=239.0.0.1, ports={lrp-sw1}**
Logical_Flow: **LS1,match=ip4.dst==239.0.0.1, action=outport(MC-239.0.0.1)**
Multicast_Group: **LS1,name=239.0.0.1, ports={vm2}**

Status, future work, potential issues

IGMP Snooping and Querier:

- available in OVN 2.12

IGMP Relay:

- available in OVN master

Static Multicast configuration:

- allows per port flood configuration for multicast
- enables connectivity to upstream multicast routers
- available in OVN master

Concerns:

- additional load on ovn-controller and ovn-northd
- the 4k resubmit limit is harder to reach for IP multicast traffic but still possible if many hosts join the same groups
- IPv6 support (MLD) to be implemented

Thanks!