

Neutron Port Binding and Impact of Unbound Ports on DVR Routers with Floating IPs

Presenters

Brian Haley - Red Hat

Swaminathan Vasudevan - SUSE

Agenda

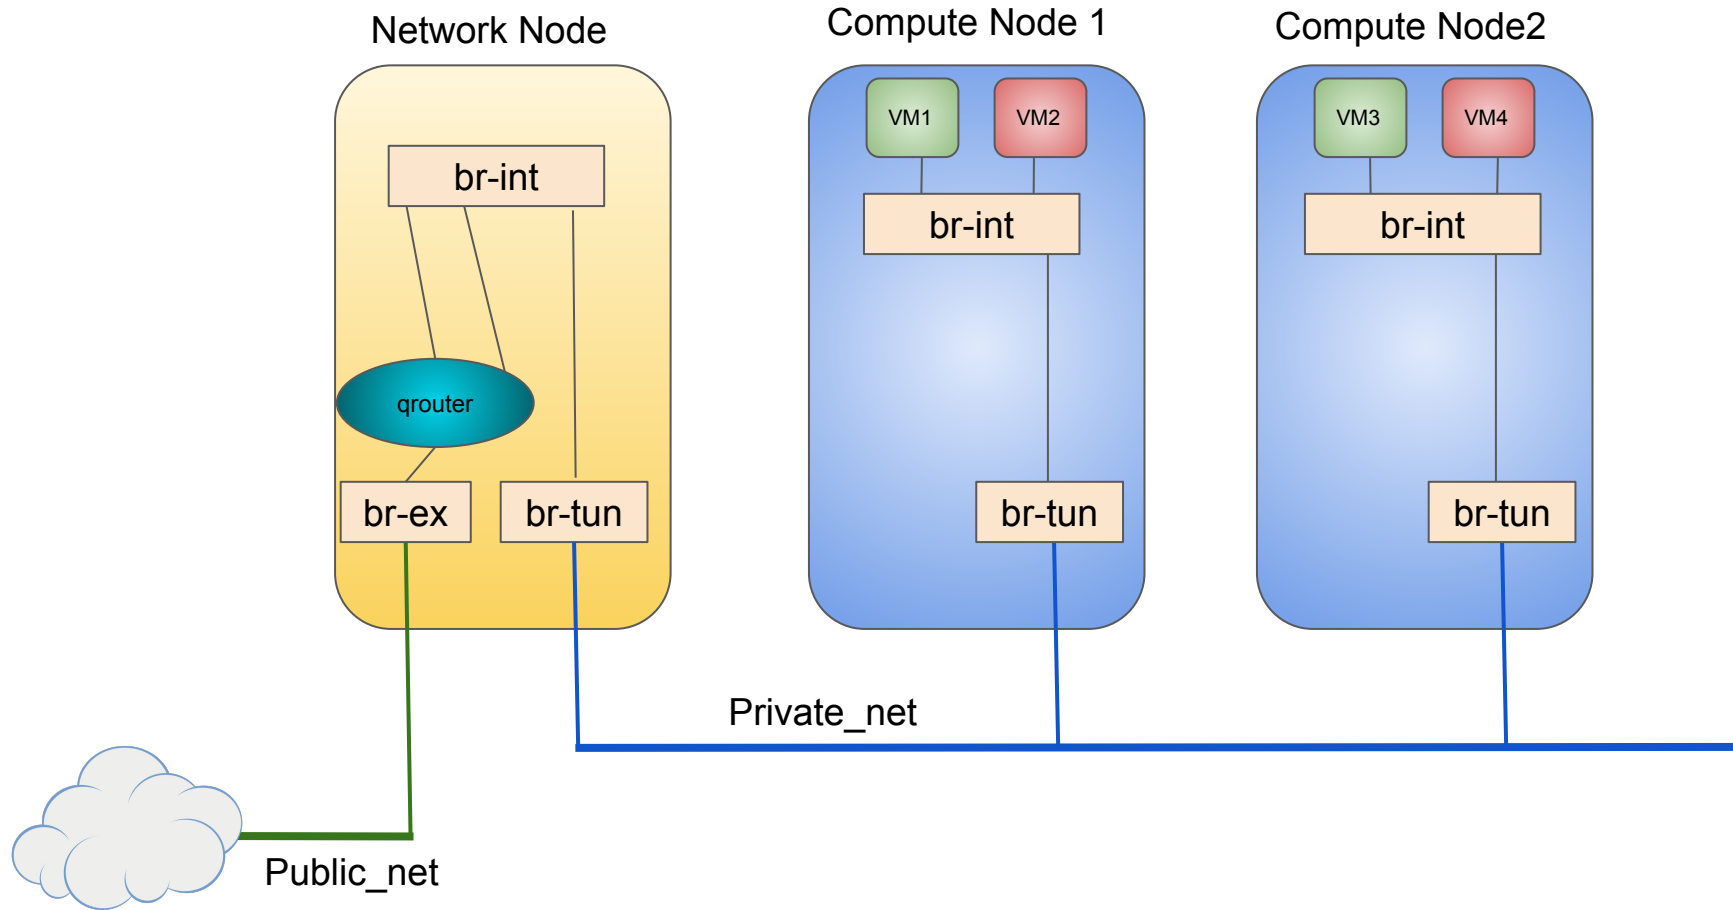
- Introduction
- Neutron Port Binding
- DVR Router scheduling
- Handling Floating IPs with DVR
- DVR Limitation on Unbound ports
(Current issues with Octavia and VRRP use case)
- Proposed Approach to handle the Unbound ports for DVR
- Design considerations for unbound ports with FIP
- Future plans for scheduler change for Floating IPs with DVR

Bug 1583694


<https://bugs.launchpad.net/neutron/+bug/1583694>

DVR support for Allowed_address_pair port that is bound to multiple ACTIVE VM ports

Centralized Virtual Router (CVR)



Distributed Virtual Router (DVR)



Floating IP

- How to access an instance from a public network
 - DNAT from public IP to private fixed IP address
- With Centralized Routers, floating IPs are configured on the network node
- With Distributed Routers, they are configured on the compute node
 - Direct access to the external network

Modifying port association is much different between the two!

Allowed Address Pairs

- Allowed address pairs is a port attribute extension that allows an arbitrary MAC address/IP address pair to pass through a port regardless of the subnet associated with the network
- Needed because by default neutron ports only allow traffic through that matches the MAC address and fixed-IP fields on a port (which is done to enforce anti-spoofing)

Virtual Router Redundancy Protocol

- The Virtual Router Redundancy Protocol (VRRP) is a computer networking protocol that provides for automatic assignment of available Internet Protocol (IP) routers to participating hosts ¹
- VRRP works by having hosts participate in a group which share a configured virtual IP address. One node is elected master and is the only node that will respond on that given address. The other nodes are slaves and just monitor the current master node periodically to ensure that it is still running.

1. https://en.wikipedia.org/wiki/Virtual_Router_Redundancy_Protocol

Neutron Port Binding

Port Binding

- ML2 Port Binding for Legacy Routers.
- ML2 Port Binding for DVR Routers

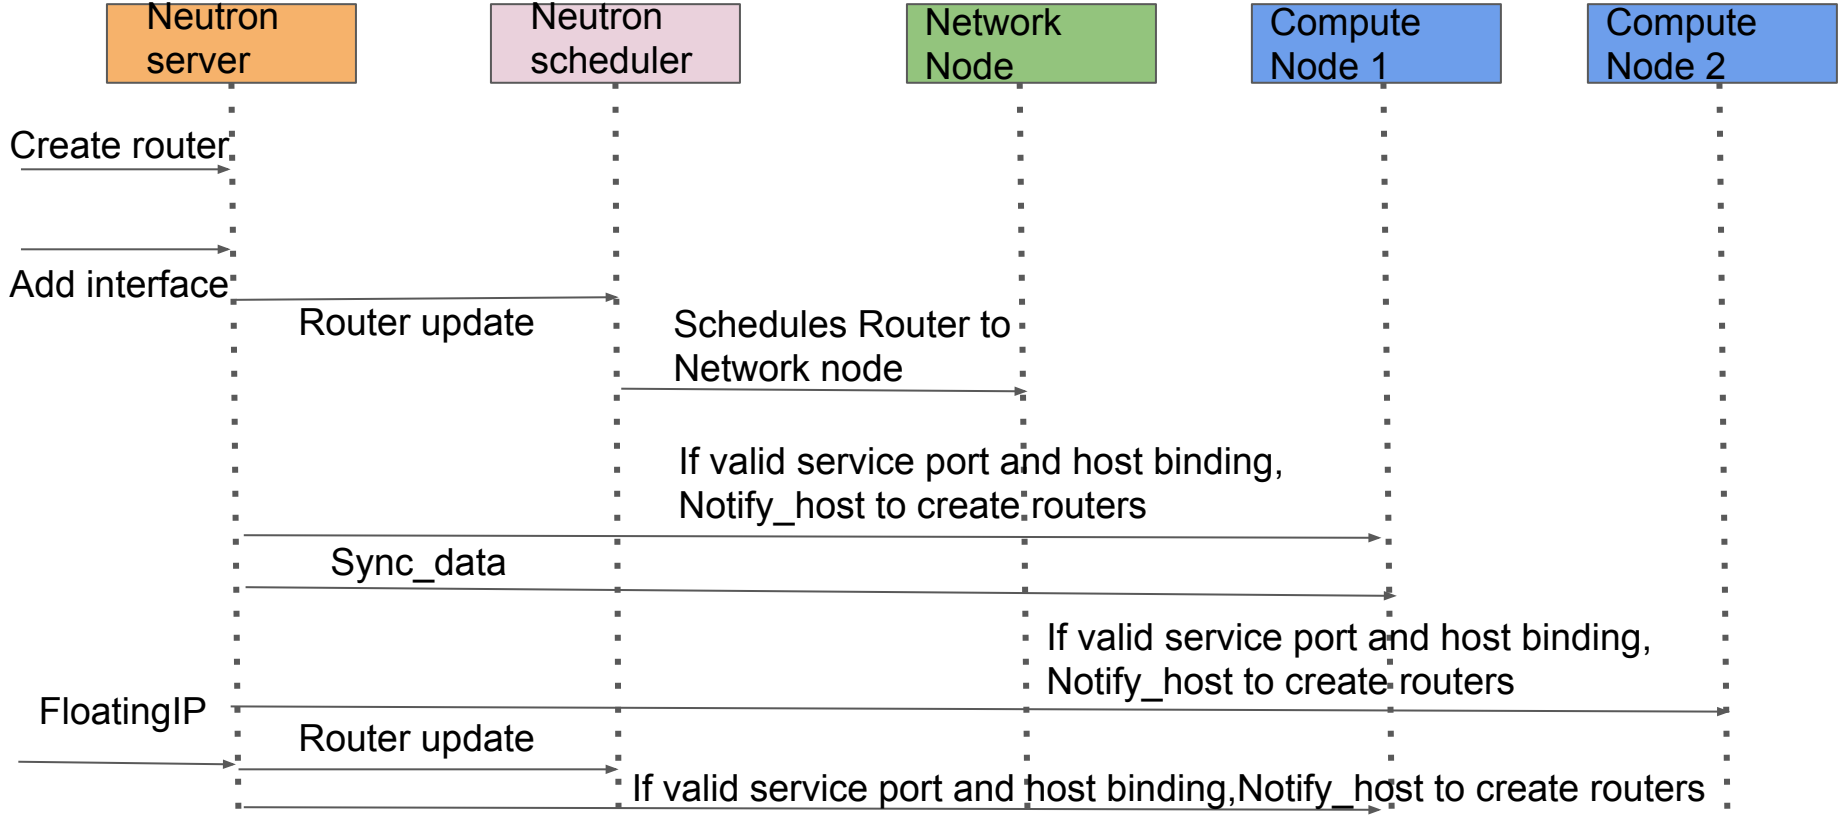
Port Binding attributes

- `device_owner`
- `binding:hostid`

DVR Router Scheduling

- DVR Routers are always scheduled to the 'dvr_snat' agents. (Network node)
- Update (router_update) notifications (Create/Update and Delete) are sent to the hosts that have valid dvr service ports based on host binding.

DVR Scheduling



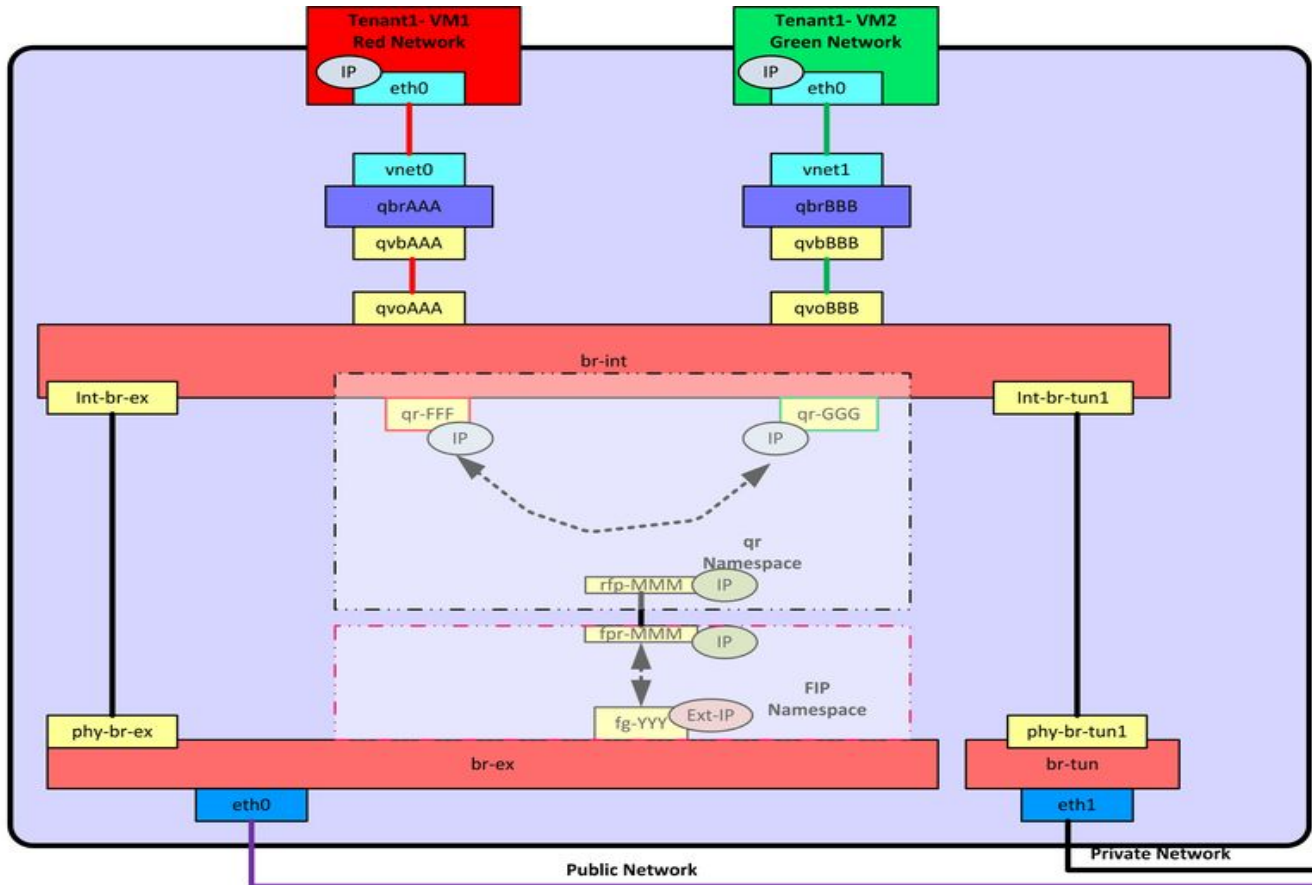
DVR Limitation on Unbound Ports

- DVR routers have a tight dependency on ports host binding to be scheduled in compute hosts.
- Unbound Ports are untouched or unscheduled by DVR routers
- VRRP ports or the Allowed-address-pair ports are not bound to any permanent host.
- When Floating IP's are configured on an unbound port, in the case of a VRRP port or Allowed-address-pair port, it is untouched by the agent, since the server does not send a notification message to the agents.

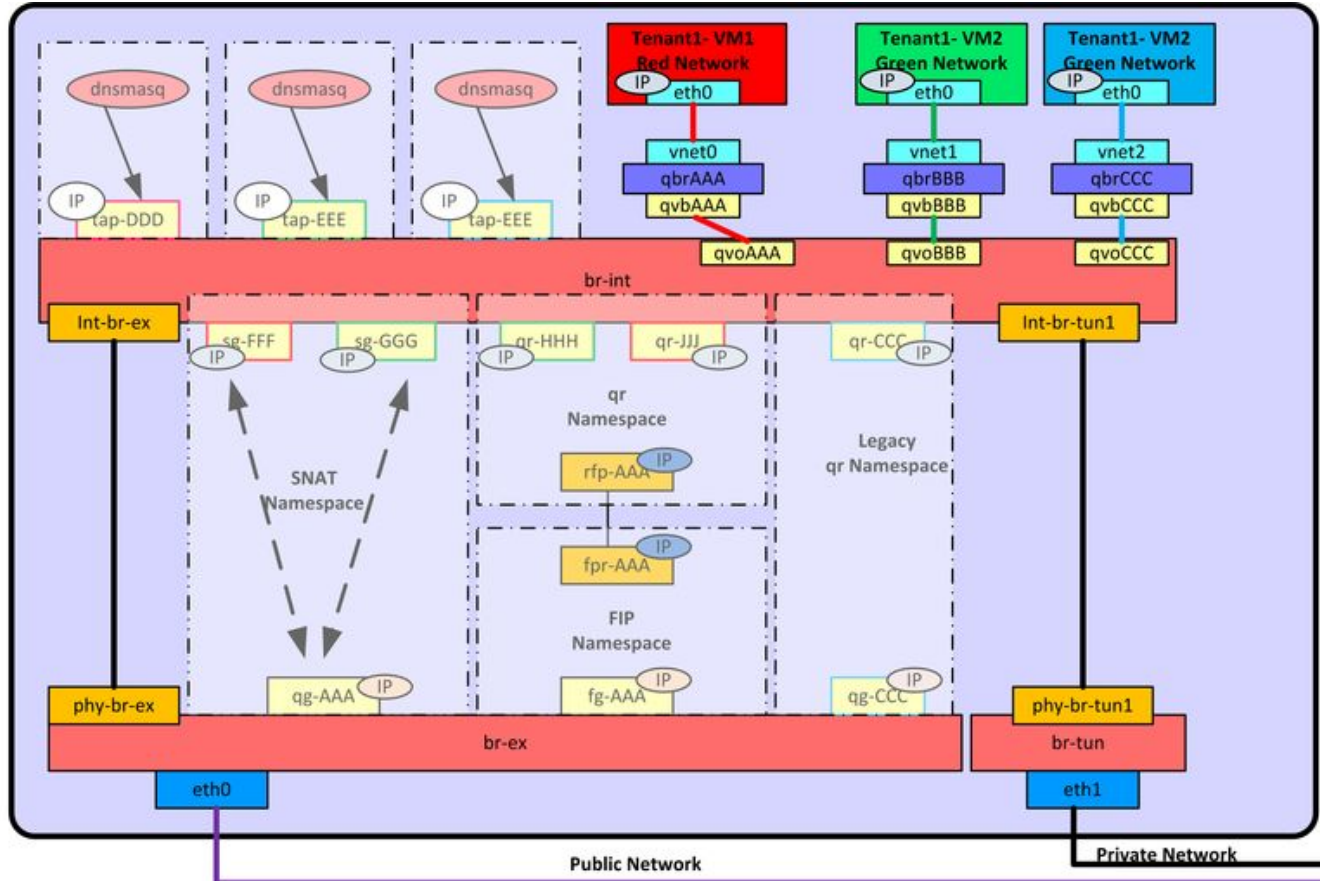
Handling Floating IPs with DVR

- Floating IPs are host bound and created when required
- Floating IPs are implemented in the shared FIP Namespaces
- Floating IPs require an agent gateway port on every compute node
- Floating IPs are also created on the SNAT node's FIP Namespace if nova-compute is enabled on SNAT node (Unique test only case).

Handling Floating IPs with DVR



Handling Floating IPs with DVR



Design considerations for centralized vs distributing the unbound port (VRRP/Allowed-address-pair) Floating IPs

Distributed Model:

- Slow failover
 - Compute Node -> Compute Node teardown/setup takes time.
- Router namespace should receive the ARP update for the VRRP port
- ARP poll for the IP in the router-namespace.
- Trigger the GARP for the floating IP from the current host.
 - Lost gratuitous ARP can cause packet loss.

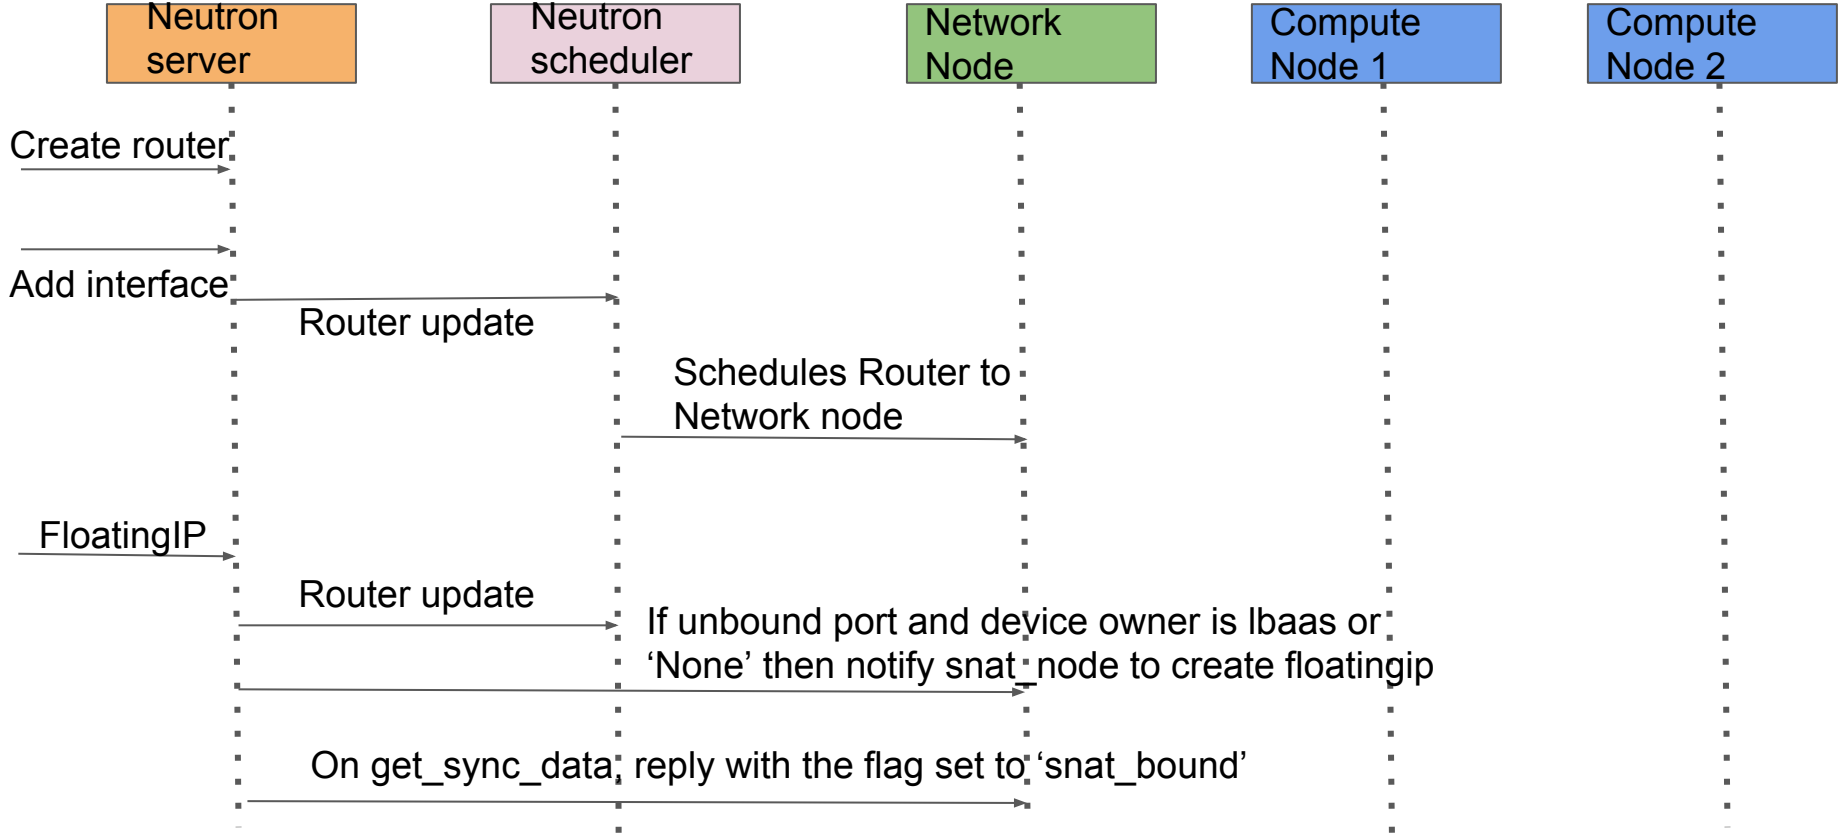
Centralized Model:

Fast failover, reliable, uses existing gateway port and no single point of failure or packet loss. Can take advantage of the DVR SNAT high availability.

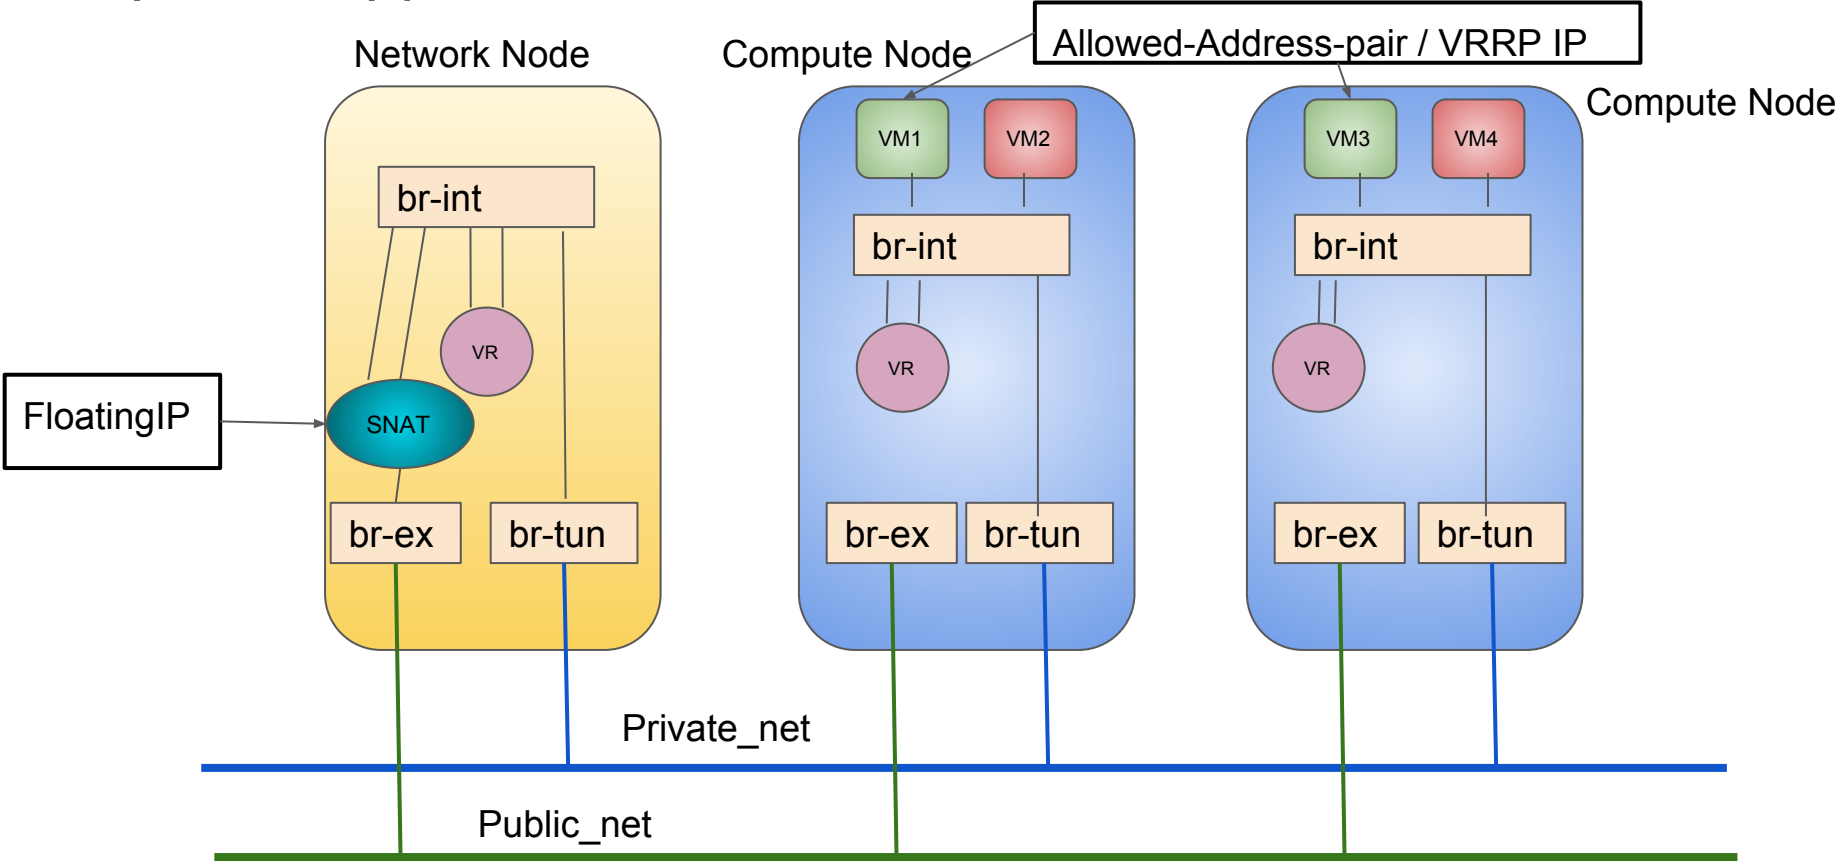
Proposed Approach to Handle Unbound Ports with DVR

- Unbound Floating IP ports are to be scheduled to Network Node
- Floating IP (DNAT) rules for unbound VRRP ports are to be deployed in SNAT Namespace at the Network Node
- The Floating IP for unbound VRRP ports utilizes the 'qg-interface' and there is no need for the 'fg-interface' similar to the compute hosts.
- Current patch under review:
- <https://review.openstack.org/#/c/437970/> (Server side patch)
- <https://review.openstack.org/#/c/437986/> (Agent side patch)
-

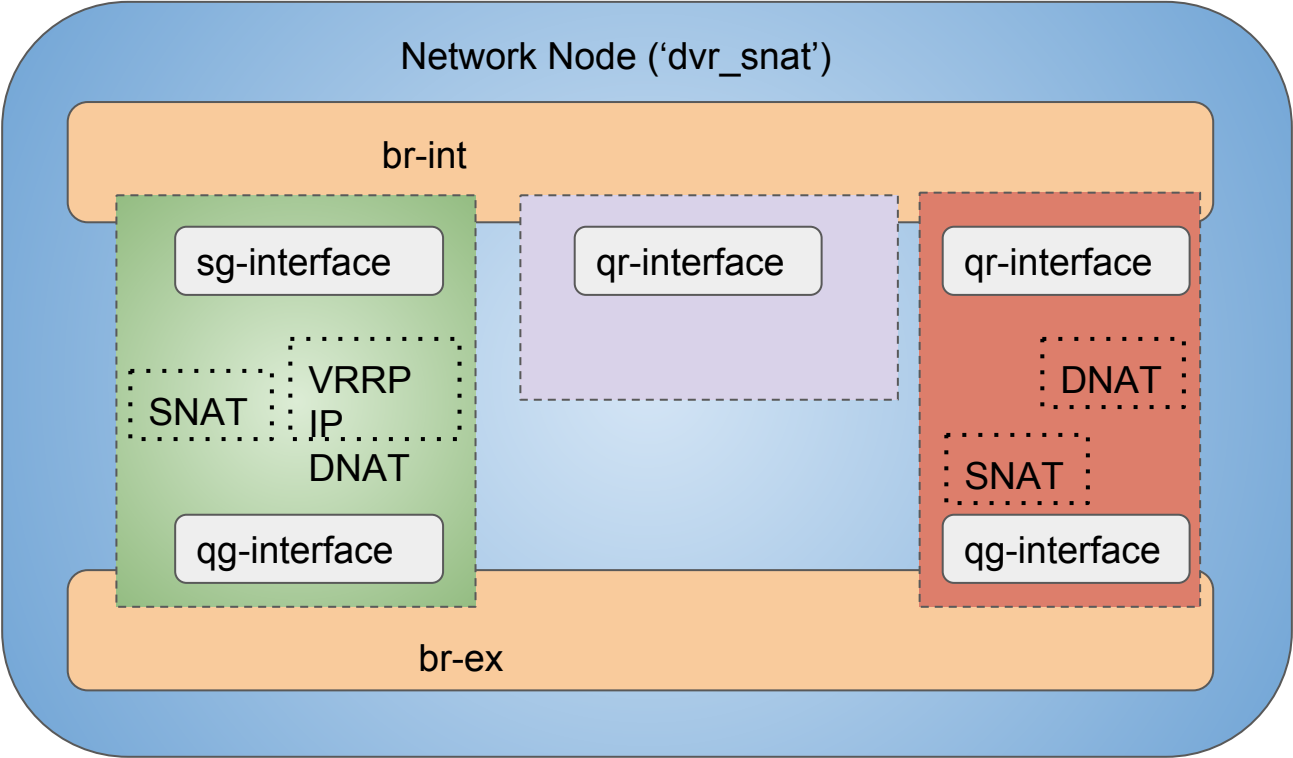
DVR Scheduling for unbound ports



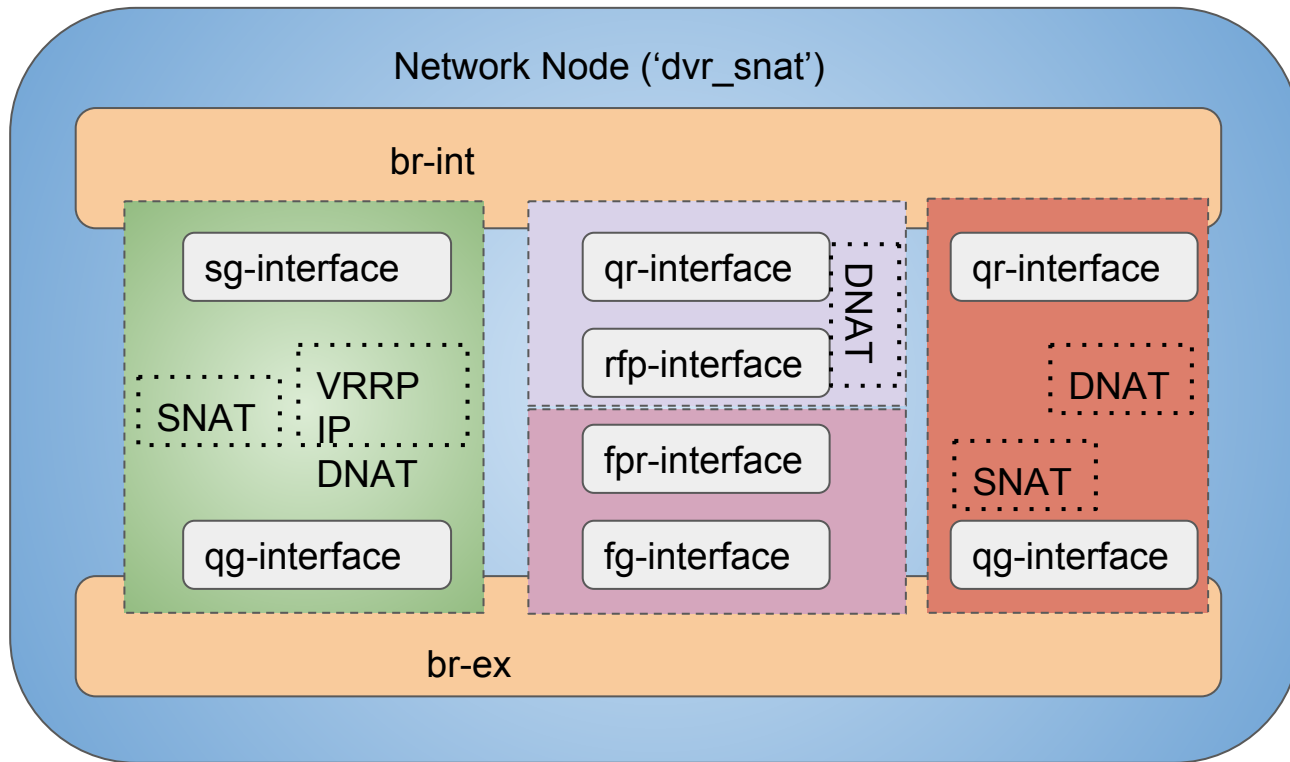
Proposed Approach to Handle the Unbound Ports with DVR



Network Node Namespace change for VRRP FIP



Network Node Namespace with unbound port FIP and bound port FIP



Future plans for scheduling the FIP (bound ports and unbound ports)

- Unbound ports always end up on the Network Node.
- Bound ports with Floating IPs can be distributed or centralized based on a config option or possibly API-based.
- <https://bugs.launchpad.net/neutron/+bug/1667877> (RFE)

Thank You and Questions