OVN Fabric Integration

Agenda

1. Current External Connection
2. Target Architecture
3. Current Work

Current External Connection

Example Openstack project

With “normal“ internet connectivity

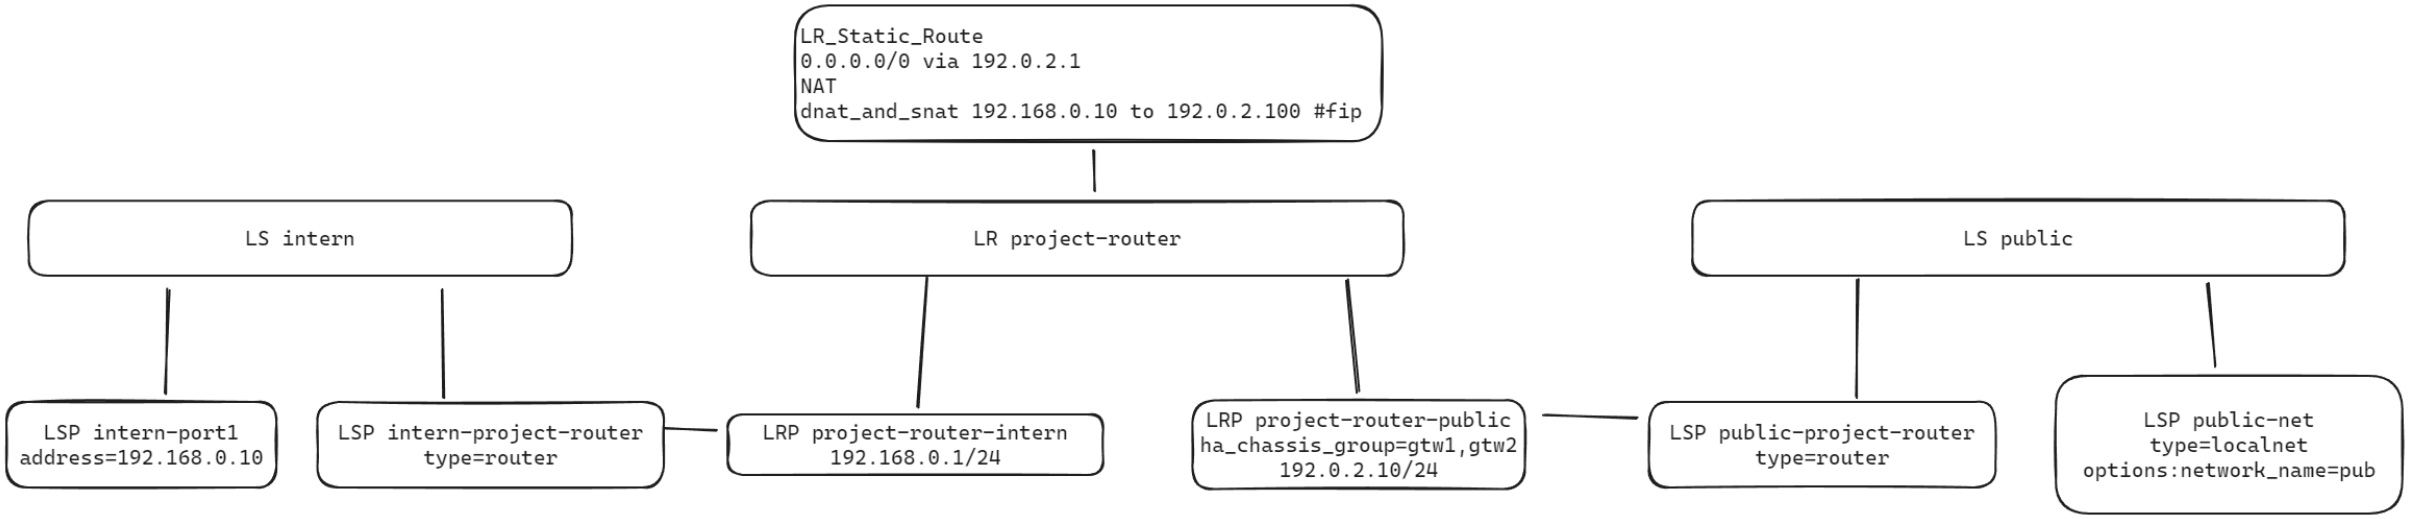

What happens on the localnet side

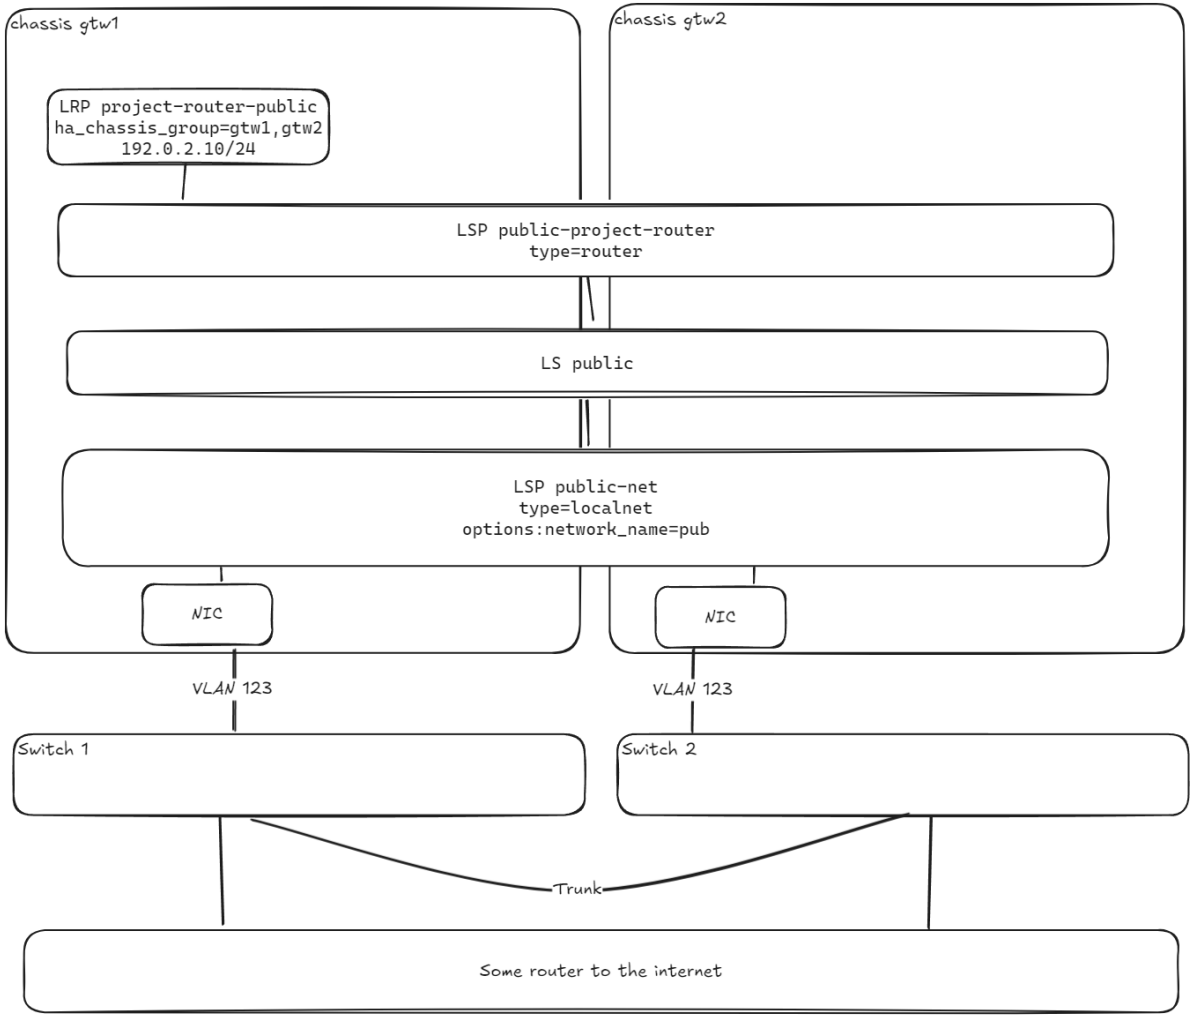

Drawbacks

- **Requires stretched L2 across all NICs bound to a localnet port**
  - All the problems of large scale stretched L2 networks
- **On failover MAC Address moves between chassis -> potential for MAC flapping**
- **Dynamic routes need to be handled outside of OVN (ovn-bgp-agent)**
- **Gateway chassis with multiple NICs need to use bonding**
  - Connection to multiple switches would need MLAG

Target Architecture

Rearchitecting for dynamic routing

Project insides stay the same

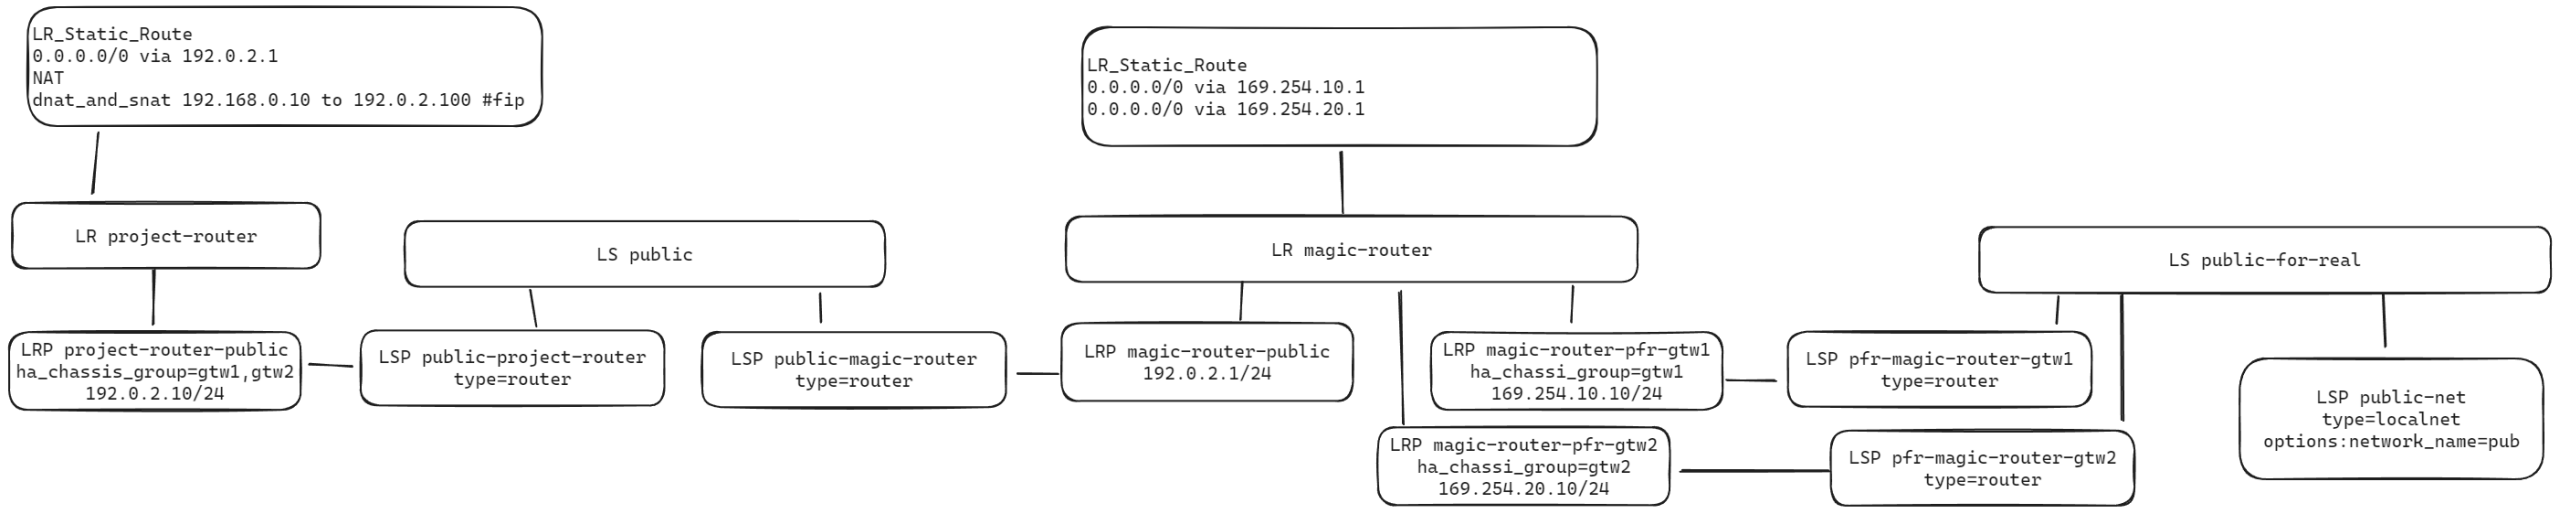

Rearchitecting for dynamic routing

- **Introduce additional „magic router“**
  - Translation from one overlay l2 network to per chassis p2p networks
  - „magic router“ is not located on any chassis, just egress traffic needs to use one specific chassis
  - ECMP for egress traffic
  - Ingress traffic can enter any of the „magic router“ LRPs

Rearchitecting on the localnet side

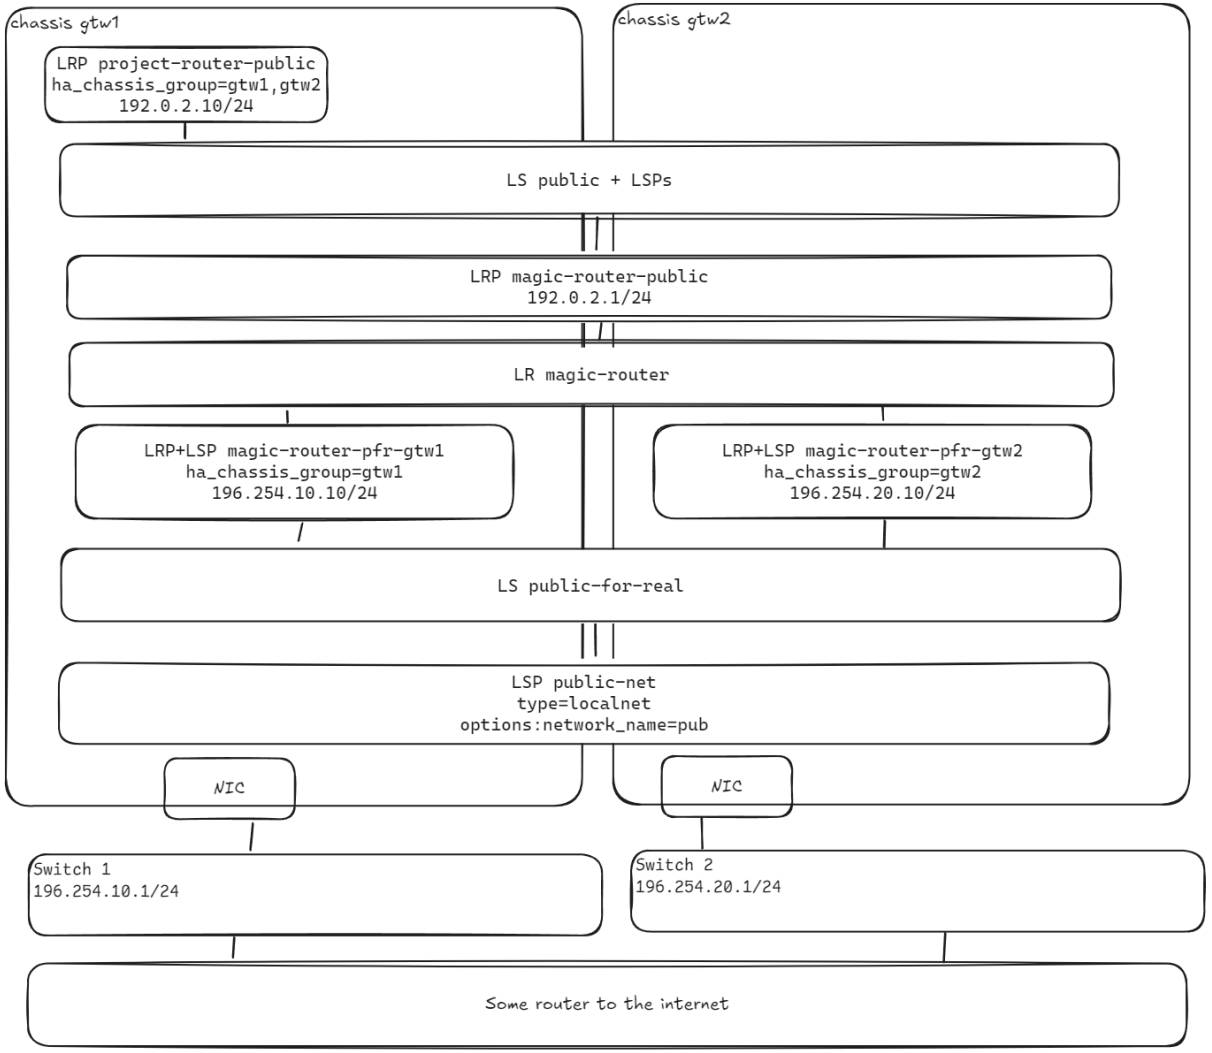

Rearchitecting on the localnet side

- **No stretched L2 network required anymore**
- **No bonding required anymore**
  - Individual NICs can be modeled as separate LRPs
- **Failover now works using moving dynamic routes instead of MAC moving**
  - Dynamic routing convergence does not limit failover speed
    - Outages can be detected with BFD and the Fabric will chose other paths
    - Planned Failover for maintenance does not need routing convergence

- **Future Optimization**
  - Traffic to project-router-public should prioritize the active gateway chassis
  - Traffic from project-router-public should prioritize the local gateway chassis

Current Work

Current Work

- Support exporting routes from LRs to Linux route tables
- Support importing routes from Linux route tables to LRs
- Automatically generate the LRPs/LSPs per chassis -> avoid changes needed in the CMS
- Optimize ingress and egress traffic based on LRP locality

- Patchset available at patchwork. Series: “OVN Fabric integration”