

Sep 5, 2023 |

## 📅 Community meeting: ovn-heater for OVN scale testing with ...

Attendees: Ihar Hrachyshka Felix Huettner Daniel Alvarez Sanchez  
Itomasbo@redhat.com Ilya Maximets Mark Michelson tiago.pires@luizalabs.com  
xsimonar@redhat.com jishi@redhat.com Terry Wilson nusiddiq@redhat.com  
hakhande@redhat.com lbiancon@redhat.com Billy Olsen aasmith@redhat.com  
Frode Nordahl Martin Kalcok mheib@redhat.com Ales Musil  
roberto.acosta@luizalabs.com rjarry@redhat.com spk@redhat.com Slawomir Kaplonski  
~~Dumitru Ceara Rashid Khan james.troup@luizalabs.com bcfarel@redhat.com~~

### Notes

- Response data:
  - <https://docs.google.com/spreadsheets/d/1HMAhy0zQDhhD59PNJWeS3PWCrvmNkAkwo-BrqDay4FU>
  - Not a lot of responses
  - Discussion about the difference between network types for “provider” and “tenant” networks, where the former have most of their ports outside OVN
  - It’s interesting that non-public cases have higher density of ports per compute
  -
- How to proceed?
  - We will need to simulate SB clients
    - Ovn-controller, ovn-metadata-agent, ovn-agent?
      - There may be a unified agent in the future, but we don’t know if and when
    - Write a fake client to simulate (python-ovs?)
  - We need a mapping between Neutron API and OVN NB
    - We can probably largely ignore SB since it’s mostly R/O
    - There is no Neutron doc for this ATM
      - We may want to have one? (Ihar)
    - Reusing ml2 code is not ideal because OVN devs should be able to work on issues without learning how Neutron works
- 

### Action items



Jul 25, 2023 |

## 📅 Community meeting: ovn-heater for OVN scale testing with ...

Attendees: Ihar Hrachyshka Felix Huettner james.troup@luizalabs.com  
hakhande@redhat.com Daniel Alvarez Sanchez lbiancon@redhat.com  
bcafarel@redhat.com Dumitru Ceara aasmith@redhat.com Frode Nordahl  
Ilya Maximets Martin Kalcok mheib@redhat.com Ales Musil  
roberto.acosta@luizalabs.com rjarry@redhat.com tiago.pires@luizalabs.com  
xsimonar@redhat.com Terry Wilson nusiddiq@redhat.com jishi@redhat.com  
spk@redhat.com Mark Michelson ~~ltomasbo@redhat.com~~ ~~Billy Olsen~~ ~~Rashid Khan~~

### Notes

- Questionnaire
  - <https://docs.google.com/document/d/151saO5a5PmCt7clZQ7DkgvU755od76BIN2bKjXc1n08/edit>
  - Discussion over who's the audience for the questionnaire - are they technical?
  - Suggested that we should talk less about low level / OVN details in the questions and more about OpenStack facing
  - (Ihar) suggest to do another swipe through the formulas used in the questions and simplify them
  - How do we aggregate answers?
    - (Frode) Start with a manual approach, then consider better tools to track (spreadsheet?)
  - (Dumitru) at some point we may want to move this questionnaire under ovn-heater to make it more public / better tracked
  - (Frode) What's the timeline for when we'll have answers?
    - (Ihar) no answer but I will follow up with higher ups on what's the best way to reach out customers for answers
  - (Ihar) how do we settle on the final version of the questionnaire?
    - Spend a day or two flushing out / polishing the wording, then consider it done and start reaching out / collecting answers
  - Do we need a chat platform to discuss it if needed?
    - Let's reuse the upstream OVN meeting in IRC (Thus)
- (Frode) Questionnaire is great, there are some steps that we should / can take regardless of results
  - (Robin) answers will give sizing info, but the implementation may be independent of size. E.g. metadata connection simulation.
  - (Frode) ovn-heater has util functions simulating kubernetes workloads. Should we have the same for openstack?

- (Robin) someone with good knowledge of Neutron can list mapping between openstack and OVN resources / behaviors. Live checks, metadata agent, ...
- (Felix) Dump of northbound and southbound DB useful as starting point?
- (Terry) Internally, the Perf Team in Red Hat uses browbeat / rally to do some scale testing. We may check which scenarios they target and see if they may be useful for us.
- (Dumitru) For Kubernetes, we started with base topology. We can do the same to bring up the test environment, then consider targeting specific scenarios against it. Regardless of scenarios, we'll have to build the base environment.
  - (Numan) unless OCP, there's no hard wired logical topology to bring up. Chassis bring-up is enough.
- (Robin) how do we prepare the environment - over time or in an initial batch?
  - (Dumitru) for k8s, there are two steps. First we bring up a large number of initial resources, then second stage in which resources are added "one iteration" at a time (e.g., one "VM" at a time and measure latency to bring up networking).
  - (Robin) suggest targeting bringing up database with lots of resources instead of building it up over time through API.
    - (Terry) startup time is one of our big issues. A test of restarting all clients connected to ovsdb-server could be a good thing to simulate.
- Simulation of metadata client connection to SB.
  - Does it only read or write too?
    - A: It does write.
  - (Dumitru) we could define different kinds of nodes to simulate service distribution.
  - (Robin) about using the real metadata agent. There are probably bindings to the Neutron API that may not allow us to easily reuse a single component.
    - (Frode) +1 on this. Having all dependencies that OpenStack may bring into this project is not a good path.
    - (Dumitru) The goal is to reproduce the environment and allow OVN contributors to work with the environment and improve it for OpenStack but not using OpenStack.
- (Ihar) let's map out what should be done. How do we approach it?
  - GitHub should have some tooling to create Milestones / plan the project / split in pieces etc. Created under ovn-heater repo.
    - (Dumitru) will take a look at the tooling etc.
- Next meeting tentatively Aug 29th.
- (Q) Where do we put the results of the questionnaire?
  - (Robin) I can create a spreadsheet to put answers there.
  - (Frode) if Robin can take care of it, go for it.
  - (Robin) should be ready by Thu.

## Action items

- ☒ ~~(Everyone) Flush out the questionnaire in next day or two~~
  - ☐ (Everyone) Start figuring out how to reach customers with the questionnaire, collect answers
  - ☒ ~~(Robin) create spreadsheet for questionnaire answers, share the link during the next OVN upstream meeting in IRC (Thu).~~
  - ☐ (Dumitru) Take a look at GitHub tooling useful to capture the work.
  - ☐ (???) Map OpenStack resources / behaviors to OVN. Define work items needed for simulation.
  - ☐ (Frode) Share video from meeting
  - ☐ (Frode) Set up next video meeting
- 

Jul 11, 2023 |

## 📅 Community meeting: ovn-heater for OVN scale testing with ...

Attendees: Dumitru Ceara Felix Huettner Frode Nordahl hakhande@redhat.com  
Ilya Maximets Ihar Hrachyshka Martin Kalcok nusiddiq@redhat.com rjarry@redhat.com  
roberto.acosta@luizalabs.com Terry Wilson Maxim James Troup

## Notes

- Code of Conduct
- Introductions (companies represented: Red Hat, Canonical, StackIT, LuizaLabs - sorry if I missed someone, feel free to add)
- Ovn-heater brief intro (presentation by Dumitru)
  - <https://docs.google.com/presentation/d/1yf8lltZhA1YUr8VQ1ldxDiGoyZv0f3W42J0exS8XQsA/edit?usp=sharing>
- Discussion
  - Recreating CMS in ovn-heater has drawbacks (more work to do, the simulation may not be identical to CMS behavior)
    - That's why we should work closely with CMS teams to define scenarios
    - On the other hand, ovn-heater re-implementation may reveal improvements that can be applied in CMS when we stumble upon differences in behavior
  - Scenarios for non-gradual database population can reveal more issues than gradual population by a script
    - An existing CMS database dump may be used to capture a large database state to plug into fake-multinode environment
  - Data path testing - was it considered?

- The goal was to scale test control plane only, no plans to test data plane (yet)
  - It would be “nice” to do, but unclear how to do it at the same time as control plane is validated
  - Some of data plane testing could require SR-IOV like workloads, but other clusters don’t use it and could be validated with no special hardware
  - (amusil) ovn-heater packs lots of chassis on a single host, which may affect data plane performance too
- (terry) Pain point - how we do liveness checks in neutron agents through SB database
  - Agent API may not be a good idea in the first place, there are better tools to monitor services
- One aspect that may be missed in current ovn-heater simulation is that in neutron, we run multiple agents that maintain multiple connections to SB. Can we / should we simulate it in ovn-heater environment?
  - Can be simulated by starting an additional script in fake-multinode that would establish connections. Terry suggests that in his experience just running connections, regardless of specific load run through them, may affect the cluster significantly.
- Next steps
  - We’ll need a short list of scenarios that are most common / first order
  - We’ll need a questionnaire that we can bring to respective operators / companies and then sync the lists to produce the scenarios list
    - <https://mail.openvswitch.org/pipermail/ovs-dev/2023-May/405089.html> can be reused in questionnaire
  - Set up recurring bi-weekly sync meeting for the topic on Tues
  - Large-scale use case: Can we discuss how ovn-heater could measure the config time of openflow flows on a node so that a “VM” started via OpenStack could boot successfully? A common problem we face is related to cloud image metadata...

#### Action items

- ☐ [IN-PROGRESS] Write down a questionnaire / discuss it in ML ([draft document](#))
- ☒ ~~Set up a cadence for the sync meeting: Next meeting in two weeks~~
- ☒ ~~Share recording~~