




OpenStack on FreeBSD Work in Progress

03/30/2023

Chih-Hsin Chang


Hello

- Chih-Hsin (Zespre) Chang
- Harvester HCI, SUSE
- Experiences
 - FreeBSD user since 2011
 - OpenStack user since 2013



Project Information

- Basic Info
 - Co-work with Li-Wen Hsu (lwhsu)
 - Kicked off in Jan. 2022
 - Sponsored by the FreeBSD Foundation, started in Jul. 2022
- Links
 - Weekly-based discussion
 - <https://hackmd.io/@lwhsu/rkgmgSlht>
 - GitHub organization: OpenStack on FreeBSD
 - <https://github.com/openstack-on-freebsd>
 - <https://github.com/openstack-on-freebsd/docs>
 - FreeBSD quarterly status report
 - https://www.freebsd.org/status/report-2022-07-2022-09/#_openstack_on_freebsd
 - https://www.freebsd.org/status/report-2022-10-2022-12/#_openstack_on_freebsd

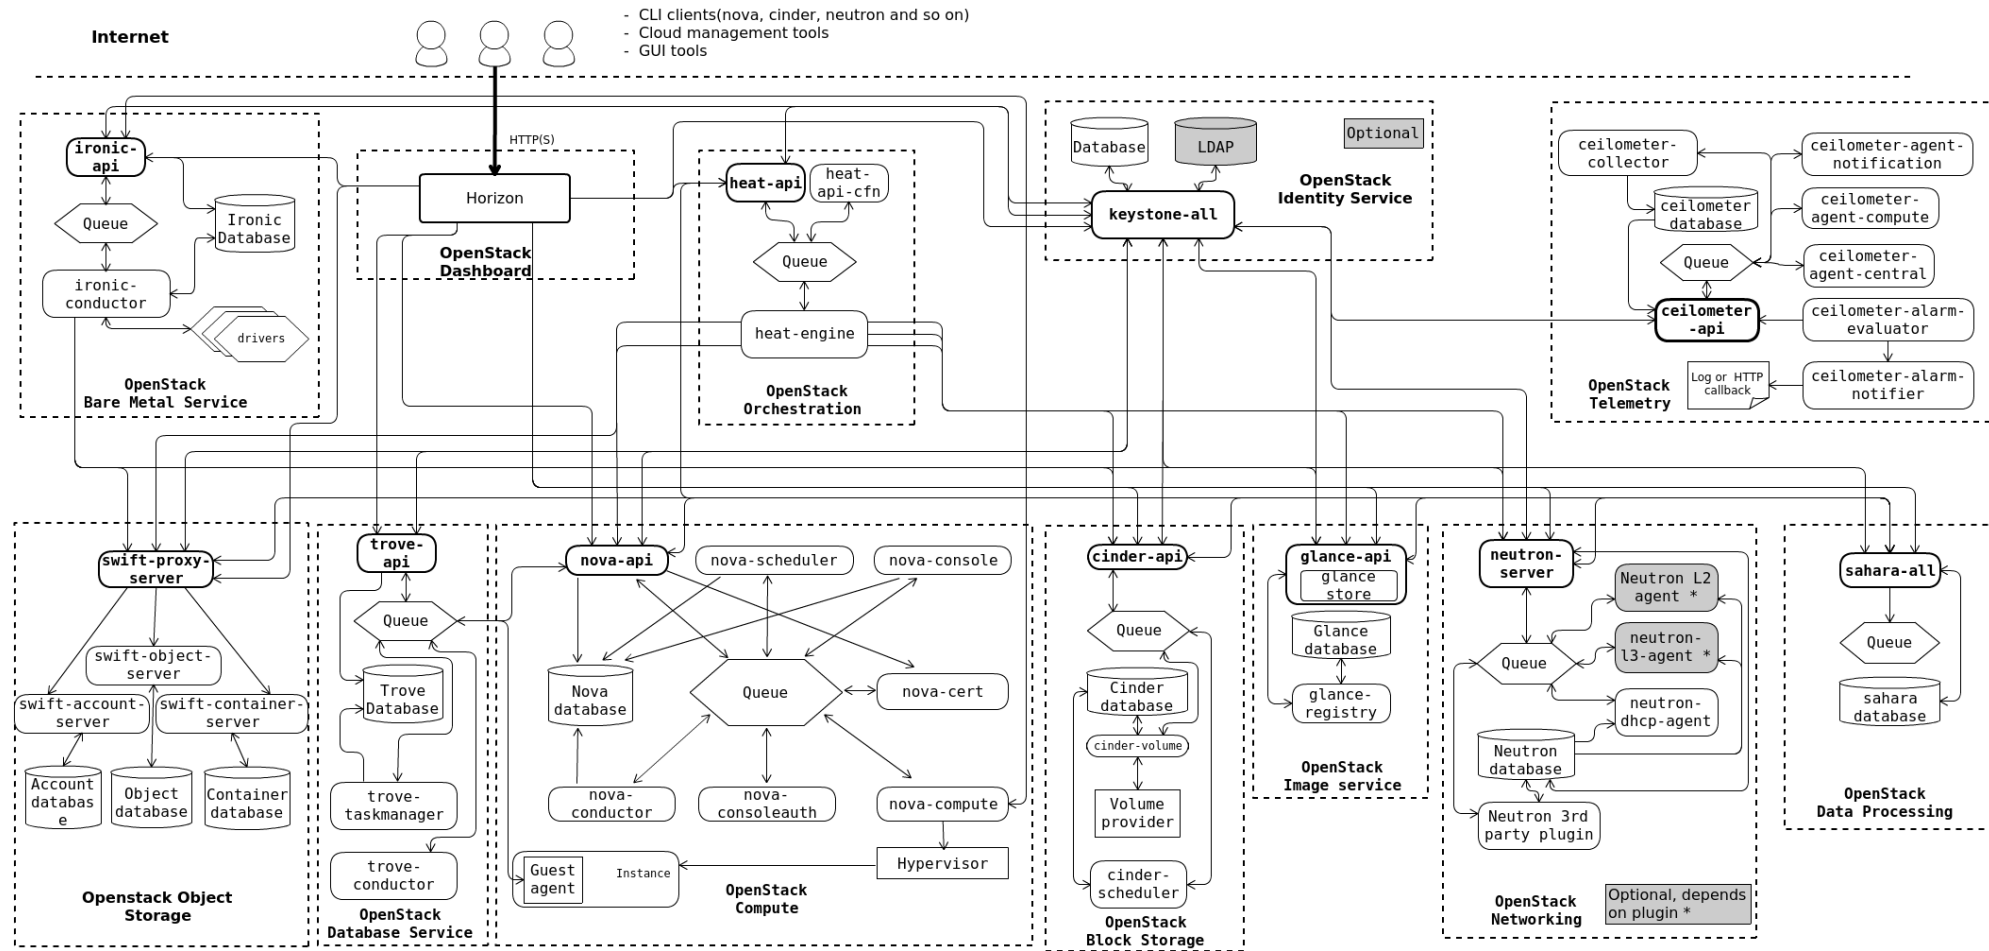


Why OpenStack?

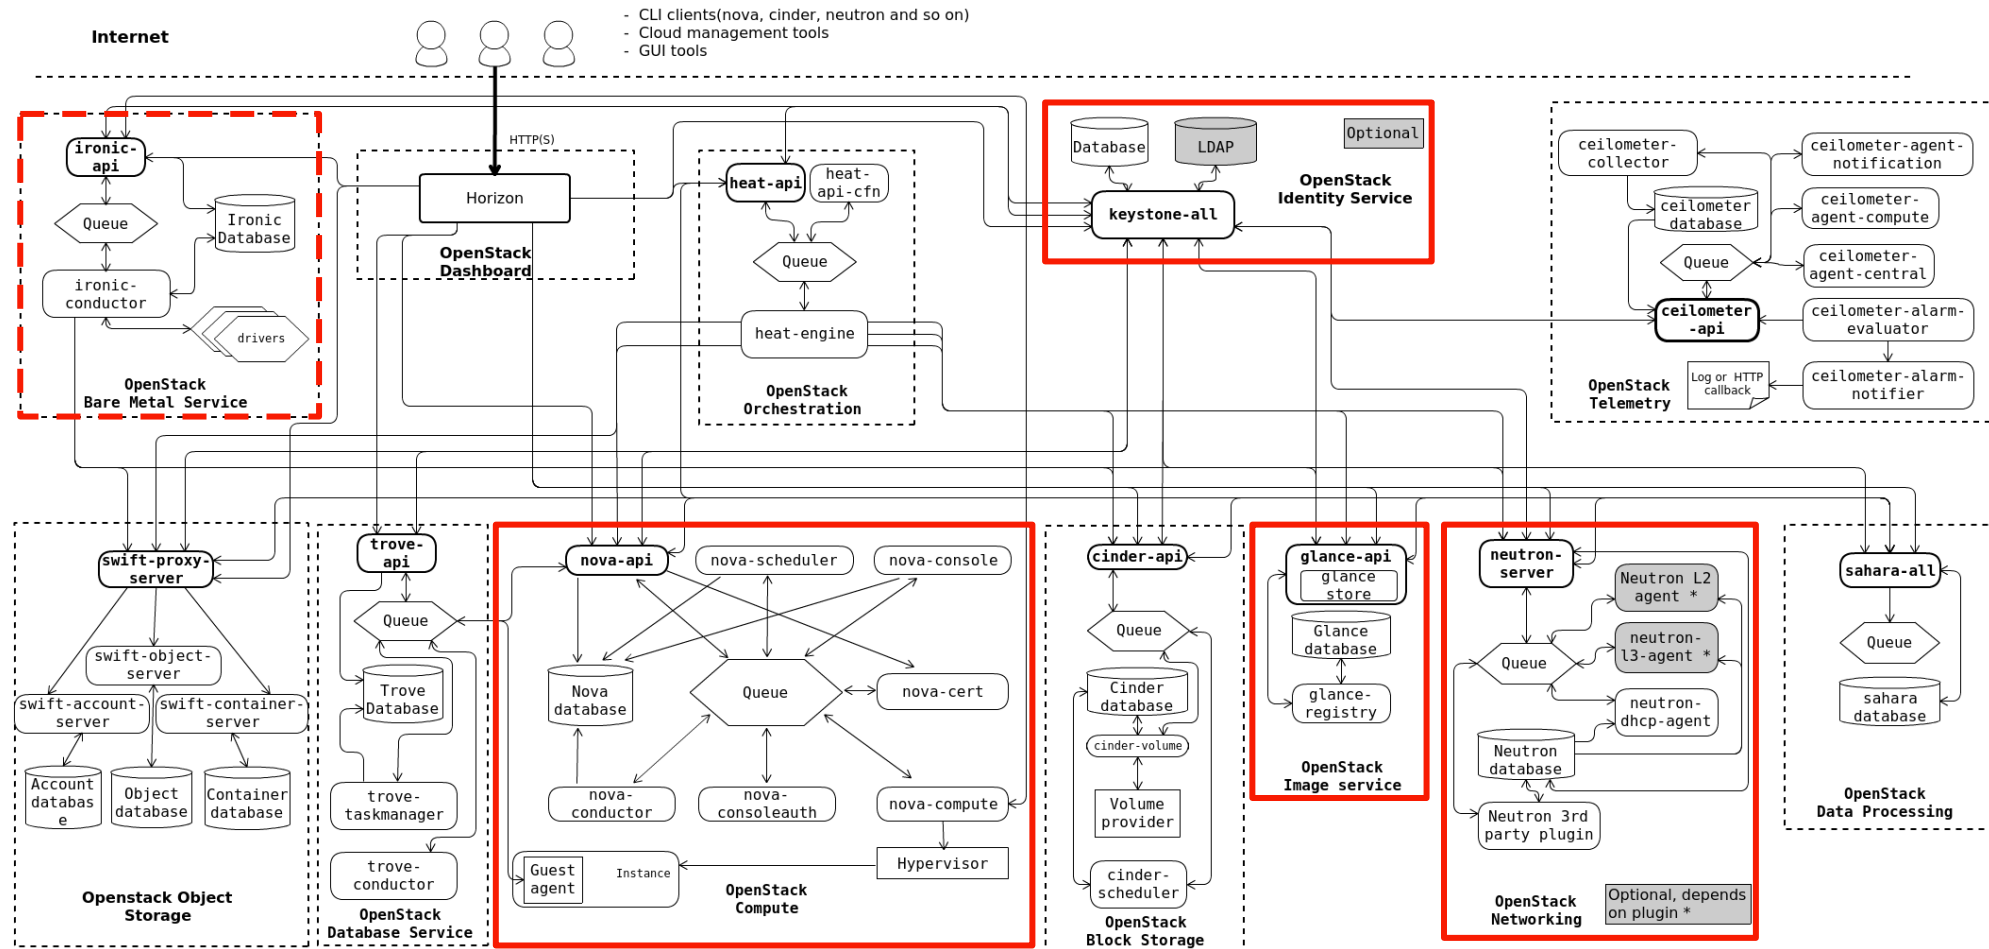
- It is a cloud orchestration platform for almost all kinds of resources
 - Compute
 - Networking
 - Storage
 - ...
- FreeBSD OS is only available as a guest OS
- Excellent opportunity for the FreeBSD community
 - To promote the capability of FreeBSD OS
 - To communicate with another big/mature community

OpenStack Overview

- **Crucial components**
 - keystone
 - glance
 - placement-api
 - neutron
 - nova
- **Middleware/dependencies**
 - RabbitMQ
 - memcached
- **Auxiliary**
 - ironic



OpenStack Big Picture



Porting Area

Project Timeline


- Milestone 01
 - Porting keystone and ironic
 - CHERI project
 - FreeBSD netperf cluster
- Milestone 02
 - Porting nova and neutron
 - FreeBSD.org cluster
- Milestone 03
 - Tenant-aware networking
 - OpenStack upstream contribution

Current Progress

- **Up and running**
 - keystone
 - glance
 - placement-api
 - neutron
 - neutron-server
 - neutron-metadata-agent
 - neutron-dhcp-agent
 - neutron-openvswitch-agent
 - nova
 - nova-api
 - nova-scheduler
 - nova-conductor
 - nova-compute
- **Integration**
 - Basic VM life-cycle management
 - Limited functions
 - Raw images
 - FreeBSD instance only
 - Network connectivity issue

Porting Strategy

- **Challenges in general**
 - Linux-based software and framework
 - Comparing behaviors between FreeBSD and Linux environments
- **Target**
 - Minimal viable instance
- **Strategies**
 - Tweaking OpenStack provided configurations
 - Patching source code directly
 - oslo library
 - Leveraging existing architecture to write new drivers/plugins
 - nova
 - neutron



Environment Setup

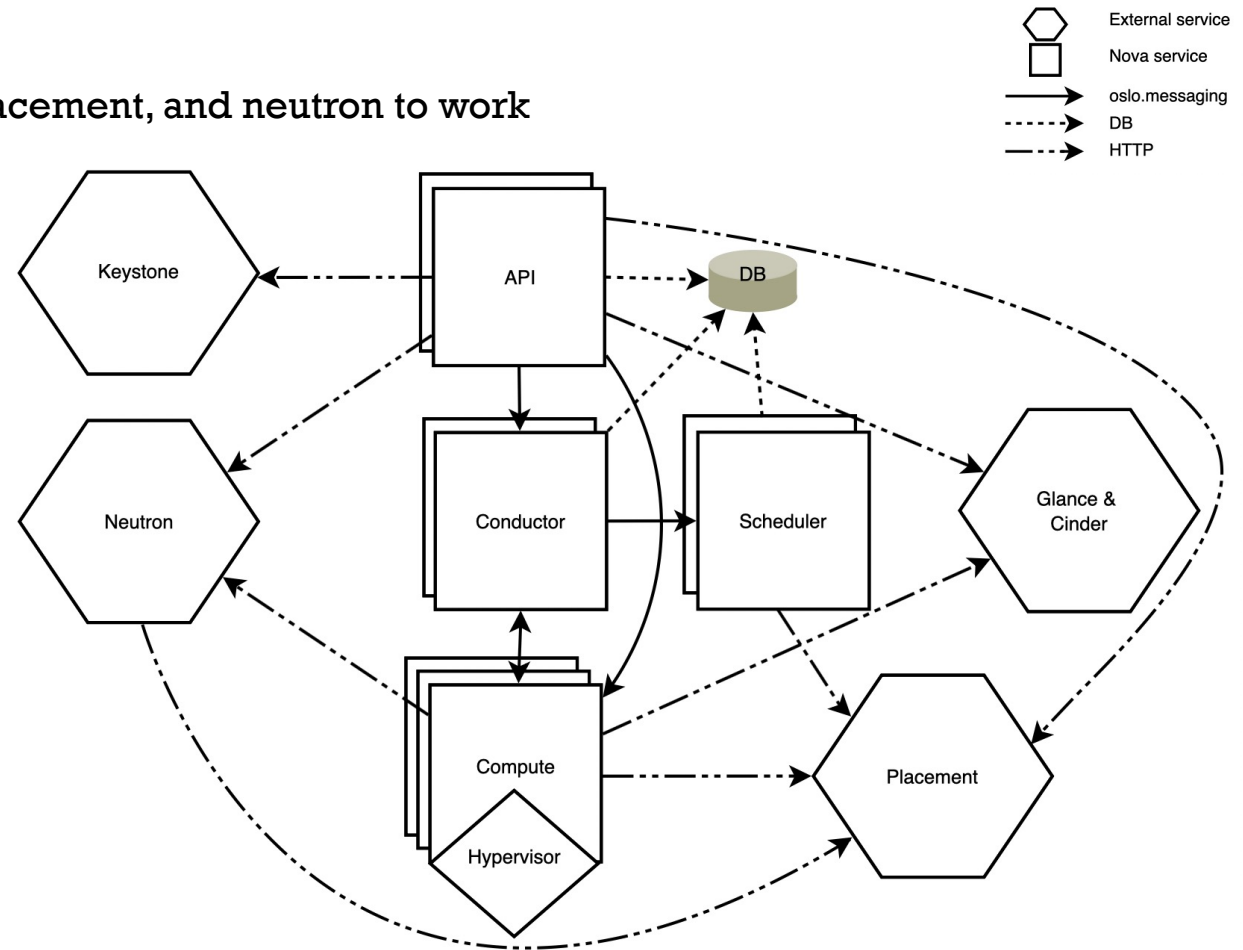
- **Hardware specs**
 - Intel Xeon E5-2680 v4
 - Supermicro X10DRL-i
 - 64 GB RAM
 - 1 TB SSD
- **FreeBSD 13.1-RELEASE**
- **OpenStack Xena**
- **Python 3.8**

Components – keystone, glance, and placement

- **Keystone**
 - Identity
 - Service catalog
- **glance**
 - Compute image repository
- **placement**
 - Tracking inventory

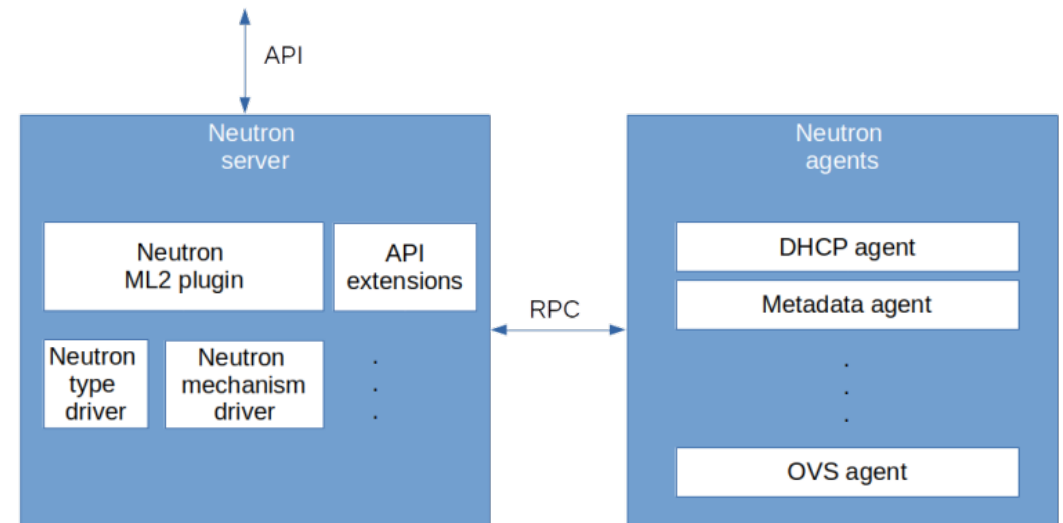
Components – nova

- Requires keystone, glance, placement, and neutron to work
- Control nodes
 - Scheduling
 - Orchestration
 - Database proxy
- Compute nodes
 - Hypervisor
 - Network interface
- Compute drivers
 - Libvirt
 - Hyperv
 - Vmware
 - Xen



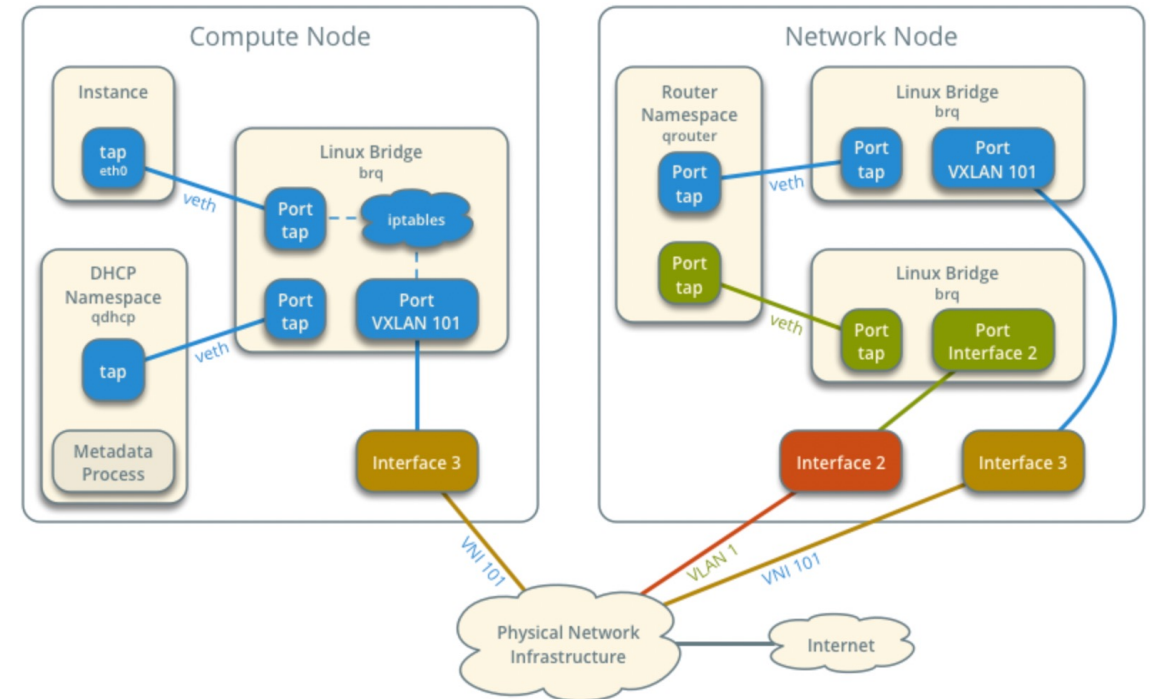
Components – neutron

- API server: neutron-server
- Agents
 - Routing: l3-agent
 - DHCP: dhcp-agent
 - Metadata: metadata-agent
- ML2 (Modular Layer 2) plugin
 - Type driver
 - Flat
 - VLAN
 - VXLAN
 - Mechanism driver
 - Open vSwitch
 - Linux bridge



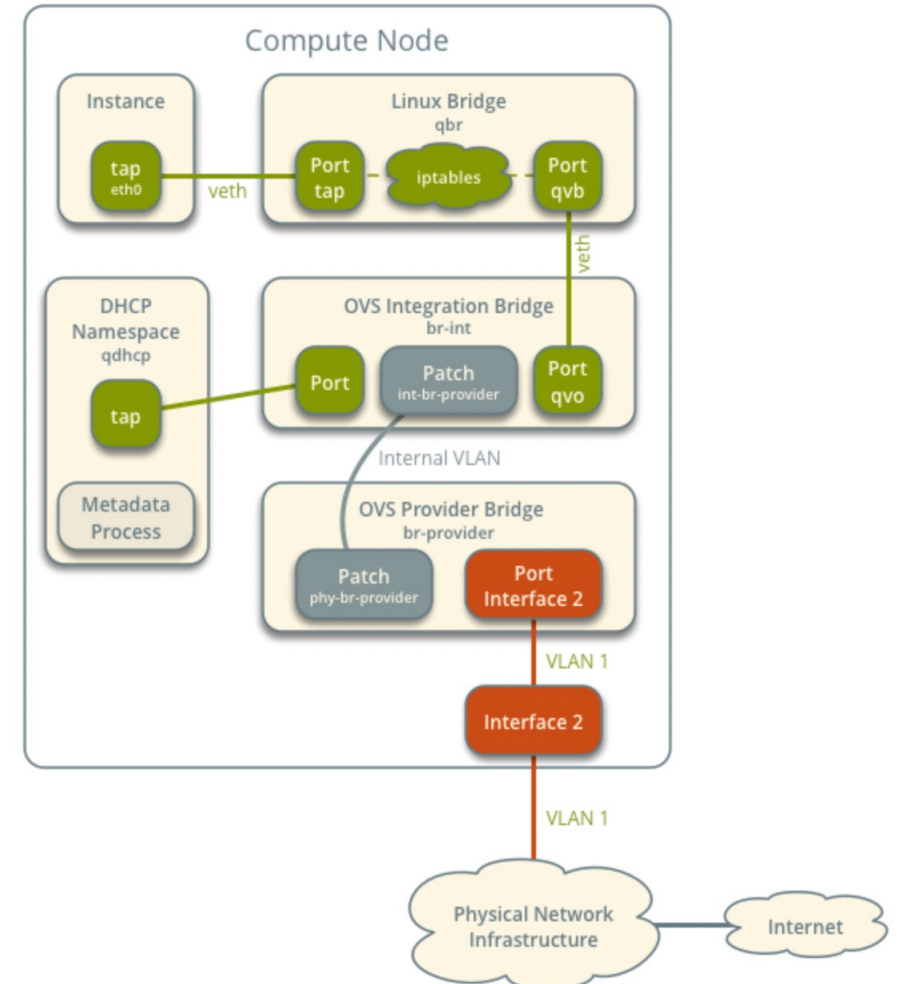
Typical Linux Bridge Self-service Networks

- **VXLAN + Linux bridge**
 - Users can create their networks
 - Isolated overlay networks



Challenges – Networking

- Flat + Open vSwitch
 - Users cannot create their networks
 - No network isolation
 - ovs-implement security group
- Why
 - Minimal changes
 - 0% vs. X% ($X > 0$)
- Eventually
 - FreeBSD bridge driver/agent
 - pf-implement security group



Challenges – Compute

- **Environment issue**
 - Nested virtualization not working (bhyve on KVM)
- **Virtualization-related**
 - Some functions not supported by libvirt + bhyve combination
 - New virt_type to handle bhhyve specific actions
 - libvirt-python might need some modification too
- **Image/storage**
 - Filesystem resizing
- **Network-related (access layer)**
 - `delegate_create: True -> False`

Challenges – Privilege Separation

- **oslo.privsep**
 - Run functions with more (or less) privileges than the rest of the code
 - Based on Linux capabilities
- **Workaround**
 - Fallback to oslo.rootwrap (sudo with filters)
 - Disable client_mode (Windows path)
- **Better solution**
 - Leverage a FreeBSD capability framework

```

from oslo_privsep import capabilities
from oslo_privsep import priv_context

sys_admin_pctxt = priv_context.PrivContext(
    'nova',
    cfg_section='nova_sys_admin',
    pypath=__name__ + '.sys_admin_pctxt',
    capabilities=[capabilities.CAP_CHOWN,
                  capabilities.CAP_DAC_OVERRIDE,
                  capabilities.CAP_DAC_READ_SEARCH,
                  capabilities.CAP_FOWNER,
                  capabilities.CAP_NET_ADMIN,
                  capabilities.CAP_SYS_ADMIN],
)

```

```

import nova.privsep

@nova.privsep.sys_admin_pctxt.entrypoint
def update_motd(message):
    with open('/etc/motd', 'w') as f:
        f.write(message)

```

```

import nova.privsep.motd
...

nova.privsep.motd.update_motd('This node is currently idle')

```

Challenges – Keep up the pace

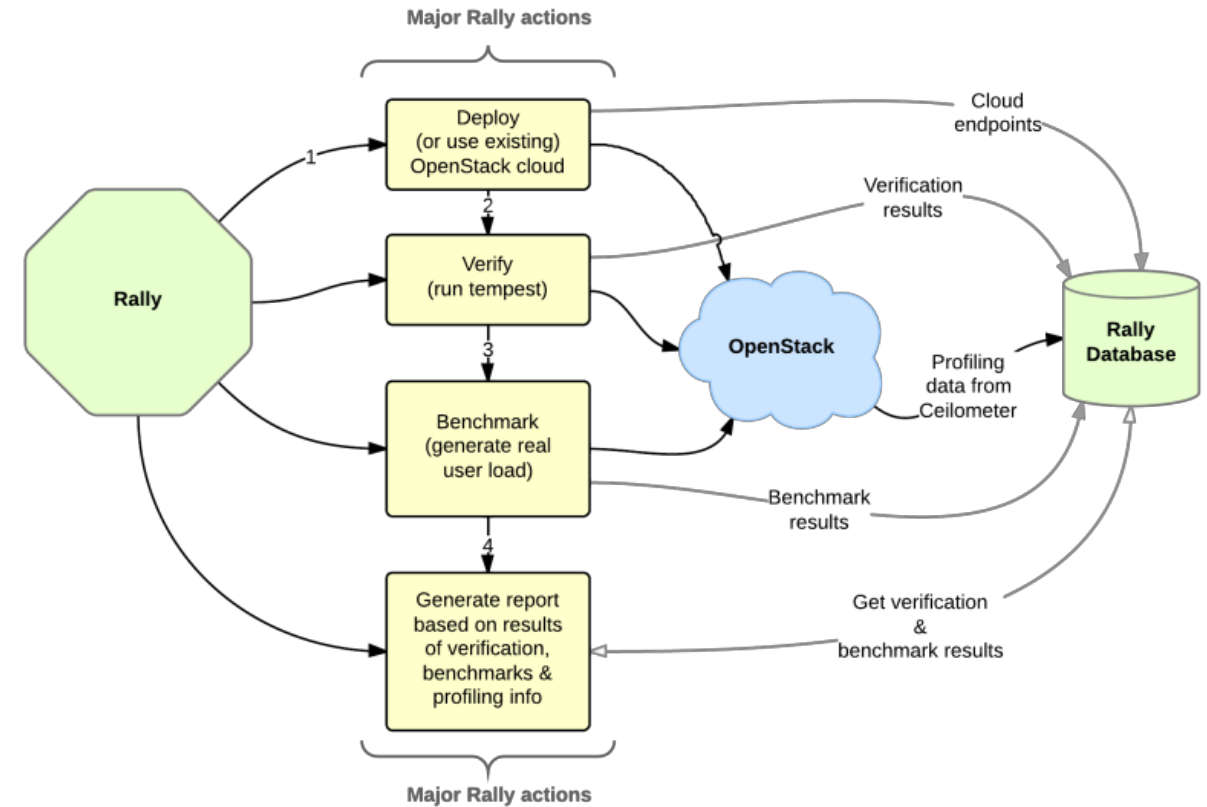
- **OpenStack Xena**
 - Released in Oct. 2021
 - Extended maintenance in Apr. 2023
- **Rebase onto 2023.1.Antelope**
- **Seeking to merge into OpenStack upstream**

Demo

- Upload images
- Create networks/subnets
- Create instances

Next

- Networking issue
- OpenStack Rally – testing
- Convert to FreeBSD Ports



- <https://github.com/openstack-on-freebsd/docs>



Thank You

Q & A