

Beskar Cloud: Openstack deployment on top of Kubernetes

Adrian Rosinec
Cloud Engineer at CESNET

28.9.2023, IBERGRID 2023, Benasque

Agenda

- The context of the computing in Czechia
- Motivation for the new architecture
- OpenStack distribution overview
- Current status



Compute services at e-INFRA CZ



- Batch compute based on PBS (known as Metacentrum)
- **OpenStack IaaS cloud**
- SensitiveCloud - PaaS based on K8S
- Managed Kubernetes – PaaS based on K8S
- Karolina supercomputer



Cloud services as tool to support research



Cloud compute service

- e-INFRA CZ is national research e-infrastructure
 - 200 research/experiment oriented allocations
 - 600 users projects in “free tier” (treated as playground)
 - **50+ international projects (through EGI and ELIXIR)**
- 300 HV, 10K CPU, 200TB RAM
- Main focus on being HPC cloud
 - large flavors (up to 128 CPU), GPUs (NVIDIA A40), fast storage (local NVMEs) and networking
- Portion of resources/support dedicated to standard operation
 - Small VMs, databases + features like LBaaS, ...



Motivation for the new architecture

- End of life of GEN1 installation from 2016
 - Custom made solution “puppet-kolla” = not supported by community
- Enable Cloud as a service (to support specialized cloud deployments, BYOC)
- Improve cloud resiliency, frequent updates
- Tune current OpenStack cloud decisions, tidy up the cloud site
 - assignment public networks to projects
 - quota assignments, projects governance
 - improve flavor naming and unify functionality from UI and commandline
 - Improve various parameters of cloud (MTUs, storages, GUI, ...)
- Add second location in Czech Republic



Partnership with commercial partner



- Taikun Cloud, Czech Republic
- Main product Taikun.cloud
- Focus on DevOps automatization tools
 - To manage OpenStack / Kubernetes clusters
- Members of Cloud native foundation
- Cooperation to create “set of scripts” to deploy OpenStack Cloud easily



OpenStack distribution as a result




- Based on open-source, popular and modern tools = supported
 - Ubuntu OS, MAAS, Ansible, Kubernetes, FluxCD, OpenStack-Helm
- Published on Github
 - <https://github.com/beskar-cloud/>
- Building an OpenStack community around it





Cloud Architecture

From HW to OpenStack services



Infrastructure as Code, GitOps



- Infrastructure management is done
 - **manually by administrators**
 - via set of **custom scripts**
- Problems:
 - **Configuration file duplicities**
 - **Lack of automation**
 - **Non-standard custom management approaches**
 - **Secret management**
 - **Manual life-cycle management**



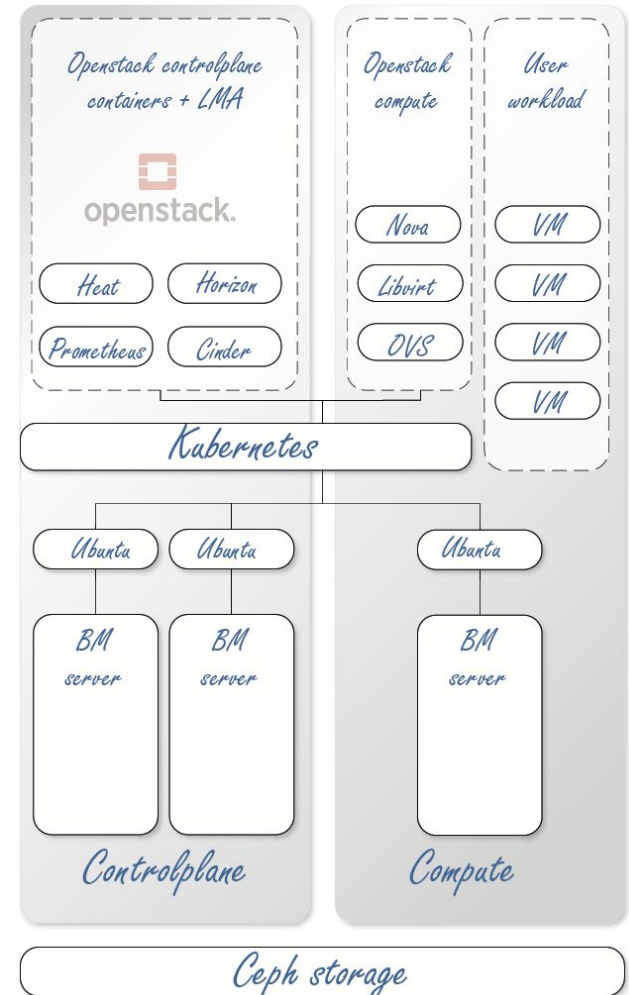
What's in the stack?

- Bare-metal provisioning
 - **MAAS + Ansible** automated in **CI/CD** pipeline
- Infrastructure management
 - **Kubernetes**
 - stable orchestrator / workload distributor
 - **HELM + Flux CD**
 - App configuration converted into HELM values
 - App deployment described declaratively
- OpenStack entity management
 - **Terraform**



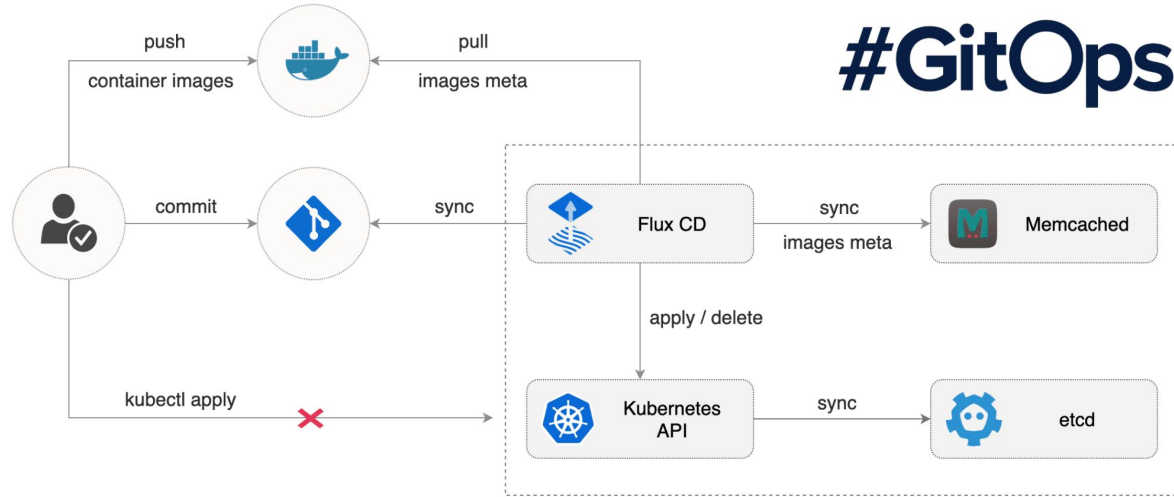
Architecture

- Kubernetes installed on every hypervisor
 - 2 types of hypervisor workload:
 - Compute
 - Running **K8S worker** with *Nova, OVS, ...*
 - Virtualisation service to run VMs via *Nova*
 - Control plane
 - Running **K8S controller** with *Horizon, Heat, Cinder, Keystone, Prometheus, ...*
 - No user VMs



Managing the site

- Infrastructure is declaratively described in git
- Repository is continuously watched by Flux CD and deployed (server-side) to Kubernetes



downstream ▾

ceph-openstack-lma / apps / prod-ostrava / 03-openstack / 03-openstack.base / + ▾

Lock



















History

Find file

Edit ▾

↓ ▾

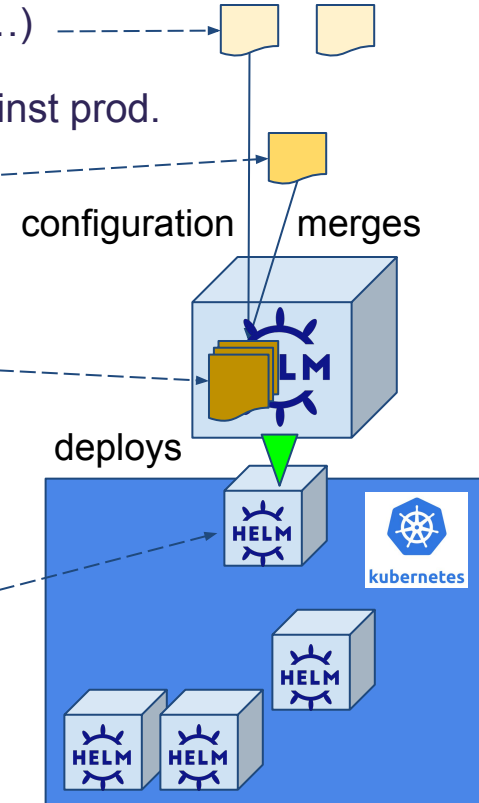
Clone ▾

Name	Last commit	Last update
..		
 00-common-configmap.yaml	refactor: rename cinder backend and re-enable	3 months ago
 00-common-encryptedsecret.yaml	#1097 introducing mariadb-backup	2 months ago
 00-namespace.yaml	feat: move namespace outside from infrastructure ks close to apps definition	3 months ago
 01-ingress-controller.yaml	refactor: yoga upgrade ostack/00-05 components	3 months ago
 02-ceph-cluster-config.yaml	refactor: yoga upgrade ostack/00-05 components	3 months ago
 02-ceph-rgw-users-encryptedsecret.yaml	fix: re-add ostack	5 months ago
 03-mariadb-backup.yaml	#1097 introducing mariadb-backup	2 months ago
 03-mariadb.yaml	refactor: upgrade monitoring/mariadb	3 months ago
 04-memcached.yaml	Update apps/prod-ostrava/03-openstack/03-openstack.base/04-memcache...	3 months ago
 05-rabbitmq.yaml	Update apps/prod-ostrava/03-openstack/03-openstack.base/04-memcache...	3 months ago
 06-keystone-apache-oidc-metadata-encryptedsecret.yaml	fix: re-add ostack	5 months ago
 06-keystone.yaml	refactor: disable keystone httpd debug logging	1 day ago
 07-radosgw-openstack.yaml	add back the glance, radosgw disabled atm, updated to last yoga, specify re...	3 months ago
 08-ceph-client-glance-key-images-rbd-keyring-encryptedse...	fix: re-add ostack	5 months ago
 08-glance.yaml	feat: add empty glance RBAC rules, add cinder policy rules	2 months ago
 09-ceph-client-cinder-backup-rbd-keyring-encryptedsecre...	fix: re-add ostack	5 months ago
 09-ceph-client-cinder-volume-rbd-keyring-encryptedsecre...	fix: re-add ostack	5 months ago
 09-cinder.yaml	refactor: drop unneeded cinder.conf [nova] entries	1 month ago

FRA
CZ

Customization of site

- Environment specific application configuration (production, staging, ...)
 - Short YAML document
 - Production secrets only | others secrets & differences against prod.
- Common cloud application configuration
 - Longer YAML document
 - Equivalent to production configuration
- HELM chart package with default application configuration
 - Long YAML document stating all possible config. options
- Configured HELM chart as application deployment via Flux CD



downstream ▾

ceph-openstack-lma / apps / prod-brno / kustomization.yaml

Find file

Blame

History

Permalink

 **kustomization.yaml**  1.05 KiB

Edit ▾

Lock

Replace

Delete



```
1 apiVersion: kustomize.config.k8s.io/v1beta1
2 kind: Kustomization
3 resources:
4   - ../00-kube-vip-controller
5   - ../base/02-rook-ceph
6   - ../base/03-openstack
7   - ../base/04-monitoring
8 patchesStrategicMerge:
9   - 02-rook-ceph/02-rook-ceph-cluster.yaml
10  - 03-openstack/01-osh-ingress-ingress-controller.yaml
11  # - 03-openstack/02-rook-ceph-client-config.yaml
12  - 03-openstack/03-openstack-mariadb.yaml
13  - 03-openstack/04-openstack-memcached.yaml
14  - 03-openstack/05-openstack-rabbitmq.yaml
15  - 03-openstack/06-openstack-keystone.yaml
16  - 03-openstack/07-openstack-radosgw-openstack.yaml
17  - 03-openstack/08-openstack-glance.yaml
18  - 03-openstack/09-openstack-cinder.yaml
19  - 03-openstack/10-openstack-openvswitch.yaml
20  - 03-openstack/11-openstack-libvirt.yaml
21  - 03-openstack/12-openstack-nova.yaml
22  - 03-openstack/13-openstack-placement.yaml
23  - 03-openstack/14-openstack-neutron.yaml
24  - 03-openstack/15-openstack-heat.yaml
25  - 03-openstack/16-openstack-horizon.yaml
26  - 03-openstack/17-openstack-barbican.yaml
27  - 03-openstack/18-openstack-prometheus-openstack-exporter.yaml
28
```


feat: enable FWaaS and VPNaaS csubcomponents



parent [c14fb575](#)

Branches > [Branches containing commit](#)

2 merge requests [1160](#) feat: enable FWaaS and VPNaaS csubcomponents, [12](#) Update...

Changes [2](#)

Showing [2](#) changed files with [26](#) additions and [1](#) deletion

Hide whitespace changes

Inline

Side-by-side

apps/prod-ostrava/03-openstack/03-openstack.base/14-neutron.yaml

+25 -0



View file @6c4132e0

```
... .. @@ -221,6 +221,10 @@ spec:
221 221     ml2_conf:
222 222         ml2_type_vlan:
223 223             network_vlan_ranges: provider
224 224     +     agent:
225 225     +         extensions: fwaas_v2
226 226     +         fwaas:
227 227     +             firewall_l2_driver: noop
228 228     neutron:
229 229     +     quotas:
230 230     +         quota_network: 1
... .. @@ -230,5 +234,26 @@ spec:
230 234     +     quota_floatingip: 1
231 235     +     quota_security_group: 10
232 236     +     quota_security_group_rule: 100
237 237     +     DEFAULT:
238 238     +         service_plugins: router,firewall_v2,vpnaaS
239 239     +         service_providers:
240 240     +             service_provider: FIREWALL_V2:fwaas_db:neutron_fwaas.services.firewall.service_drivers.agents.agents.FirewallAgentDriver:default
241 241     +     l3_agent:
242 242     +         AGENT:
243 243     +             extensions: fwaas_v2,vpnaaS
244 244     +             vpnagent:
245 245     +                 vpn_device_driver: neutron_vpnaaS.services.vpn.device_drivers.strongswan_ipsec.StrongSwanDriver
246 246     +             fwaas:
247 247     +                 agent_version: &fwaas_agent_version v2
248 248     +                 driver: &fwaas_agent_driver neutron_fwaas.services.firewall.service_drivers.agents.drivers.linux.iptables_fwaas_v2.IptablesFwaasDriver
249 249     +                 enabled: true
250 250     +             fwaas_driver:
251 251     +                 fwaas:
252 252     +                     agent_version: *fwaas_agent_version
253 253     +                     driver: *fwaas_agent_driver
254 254     +                     enabled: True
255 255     +             neutron_vpnaaS:
256 256     +                 service_providers:
257 257     +                     service_provider: VPN:strongswan:neutron_vpnaaS.services.vpn.service_drivers.ipsec.IPsecVPNDriver:default
258 258
259 259
```





feat: add basic openstack entities

František Řezníček authored 5 months ago

Code owners Assign users and groups as approvers for specific file changes. [Learn more.](#)

master ▾

openstack-entities / prod-ostrava /

+ ▾

Name	Last commit
..	
aggregate	feat: add basic openstack entities
flavors	feat: add basic openstack entities
network	feat: add basic openstack entities
main.tf	feat: add basic openstack entities
providers.tf	feat: add basic openstack entities





feat: add basic openstack entities

František Rezníček authored 5 months ago

afdf31bf



Code owners Assign users and groups as approvers for specific file changes. [Learn more.](#)

Manage branch rules

master ▾

openstack-entities / prod-ostrava / flavors / c2.16core-30ram-flavor.tf

Find file

Blame

History

Permalink



c2.16core-30ram-flavor.tf 632 B

Edit ▾

Lock

Replace

Delete



```
1 # OpenStack flavor c2.16core-30ram terraform declaration
2
3 resource "openstack_compute_flavor_v2" "c2_16core_30ram" {
4   name = "c2.16core-30ram"
5   ram = "30720"
6   vcpus = "16"
7   disk = "80"
8   is_public = false
9   extra_specs = {
10     "hw_rng:allowed" = "true",
11     "hw_rng:rate_bytes" = "2048",
12     "hw_rng:rate_period" = "1",
13     "quota:disk_total_bytes_sec" = "2097152000",
14     "quota:disk_total_iops_sec" = "1000",
15     "quota:vif_inbound_average" = "2560000",
16     "quota:vif_outbound_average" = "2560000",
17     "aggregate_instance_extra_specs:flavor" = "c2",
18   }
19 }
```



Beskar is deployed

- Test deployment in second datacenter of Czech Republic (in Ostrava)
- 30 HV, part of Karolina supercomputer cluster
- Performance testing in progress - for API, DB and internal storage
 - Spawning and deleting VMs, ...
- Migration of Brno site (the one with 300HV) is planned for Nov 2023 - Jan 24
 - as seamless as possible
 - without user interaction



This is the cloud way!



- You are welcome to join the community
- Deploy OpenStack cloud of any size using Beskar.cloud distribution
- Documentation and code:
 - <https://github.com/beskar-cloud/>



Thank you for your attention!

Questions?

Please contact us {adrian, moravcova}@cesnet.cz