

CERN Cloud Infrastructure Report

Arne Wiebalck
for the CERN Cloud Team

HEPiX Spring Meeting
DESY, Zeuthen, Germany
Apr 19, 2019

Numbers

Operations

What's new

WIP

CERN Cloud Recap

- CERN Cloud Service one of the three major components in IT's AI project
 - Policy: Servers in CERN IT shall be virtual



- Based on OpenStack

- Production service since July 2013
- Performed (almost) 4 rolling upgrades since
- Currently in transition from Kilo to **Liberty**
- Nova, **Glance**, **Keystone**, Horizon, Cinder, **Ceilometer**, **Heat**, **Neutron**



CERN Cloud Architecture (1)

- Two data centers

- 1 region (1 API), 36 cells (+10!)
- Cells map use cases
hardware, hypervisor type, location, users, ...



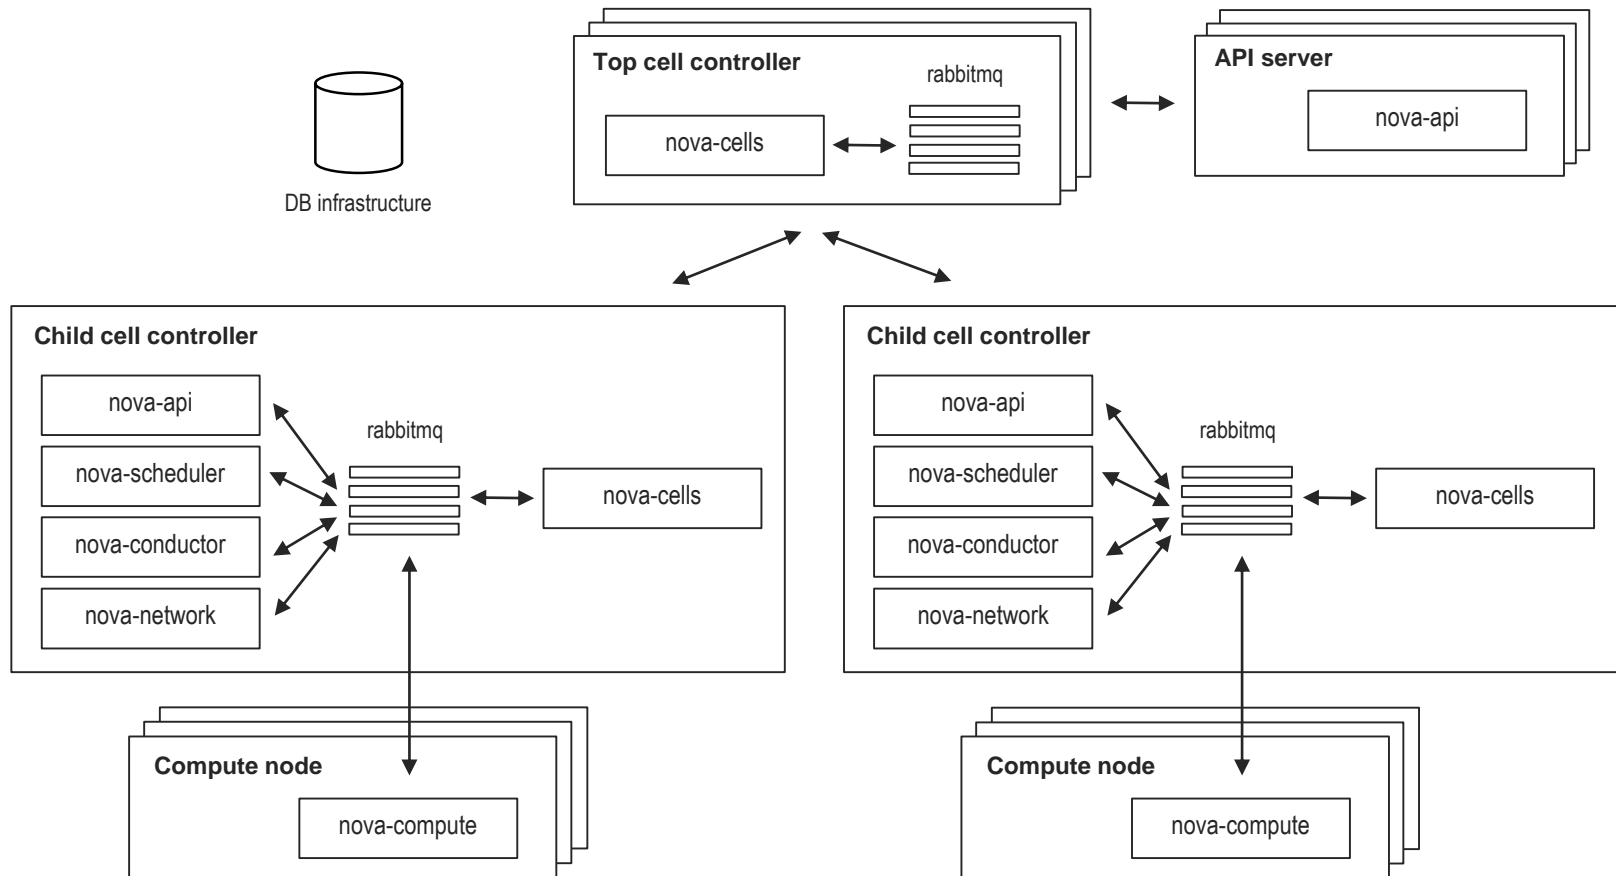
- Top cell on several physical nodes in HA

- Clustered RabbitMQ with mirrored queues
- API servers are VMs in various child cells

- Child cell controllers are OpenStack VMs


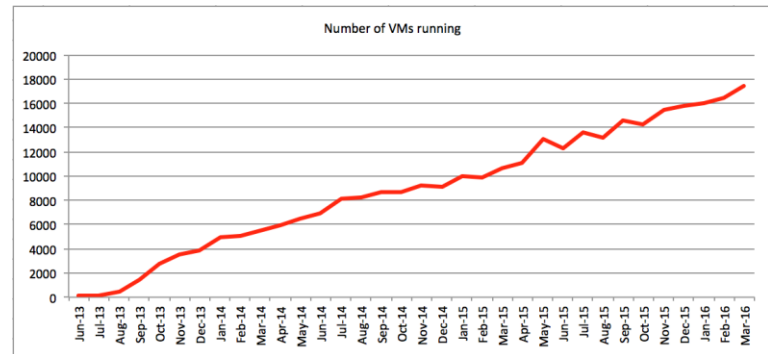
- **One** controller per cell
- Tradeoff between complexity and failure impact

CERN Cloud Architecture (2)

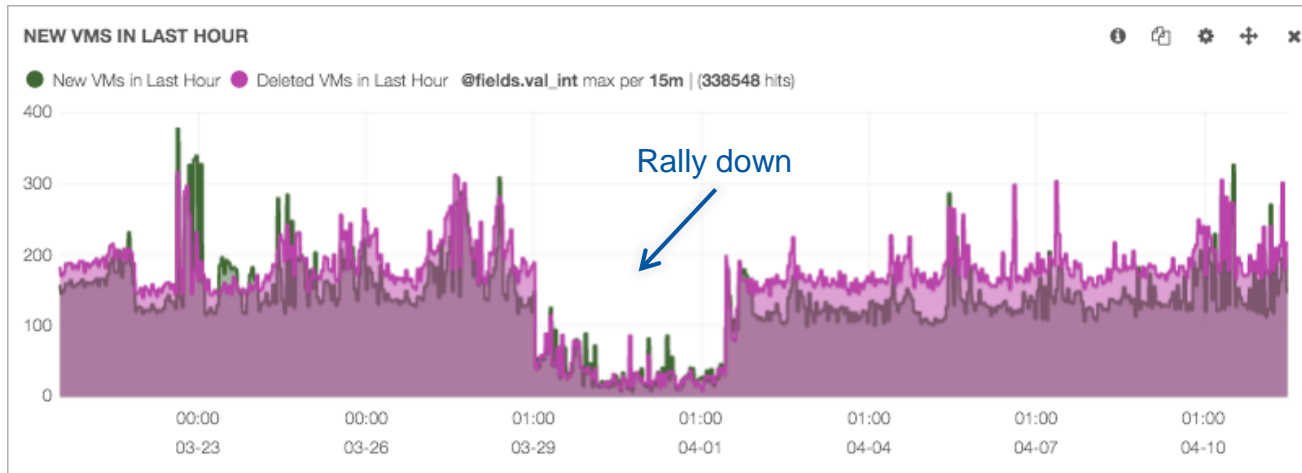


CERN Cloud in Numbers (1)

- 5'800 hypervisors in production (6m ago: +25%)
 - Majority qemu/kvm now on CC7 (~150 Hyper-V hosts)
 - ~2'100 HVs at Wigner in Hungary (batch, compute, services)
 - 370 HVs on critical power (+50%)
- 155k Cores (+30k)
- ~350 TB RAM (+100TB)
- ~18'000 VMs (+3'000)
- To be increased in 2016!
 - +57k cores in spring
 - +400kHS06 in autumn



CERN Cloud in Numbers (2)



Every 10s a VM gets created or deleted in our cloud!

- 2'700 images/snapshots (+700)
 - Glance on Ceph
- 2'200 volumes (+700, uptake doubled)
 - Cinder on Ceph (& NetApp) in GVA & Wigner




Operations: NUMA/THP Recap (1)

- The HS06 rating of full-node VMs was about 20% lower than the one of the underlying host
 - Smaller VMs much better
- Investigated various tuning options
 - KSM, EPT, PAE, Pinning, ... +hardware type dependencies
- Comparison with Hyper-V: no general issue
 - Loss w/o tuning ~3% (full-node), <1% for small VMs
 - NUMA-awareness!




Operations: NUMA/THP Recap (2)

- NUMA-awareness identified as most efficient setting
 - Full node VMs have ~3% overhead in HS06
- “EPT-off” side-effect
 - Small number of hosts, but very visible there
- Use 2MB Huge Pages
 - Keep the “EPT off” performance gain with “EPT on”
- All details in [Arne Wiebalck’s talk at BNL](#)



Operations: NUMA/THP Roll-out

- Rolled out on ~2'000 batch hypervisors (~6'000 VMs)
 - HP allocation as boot parameter → reboot
 - VM NUMA awareness as flavor metadata → delete/recreate
- Cell-by-cell (~200 hosts):
 - Queue-reshuffle to minimize resource impact
 - Draining & deletion of batch VMs
 - Hypervisor reconfiguration (Puppet) & reboot
 - Recreation of batch VMs
- Whole update took about 8 weeks
 - Organized between batch and cloud teams
 - No performance issue observed since



Operations: SSDs

- New h/w has SSDs only
 - To solve the recurrent I/O issues

```
# fio --name xyz --rw=randwrite --size=1G --direct=1
xyz: (g=0): rw=randwrite, bs=4K-4K/4K-4K/4K-4K, ioengine=sync, iodepth=1
fio-2.2.8
Starting 1 process
Jobs: 1 (f=1): [w(1)] [100.0% done] [0KB/41132KB/0KB /s] [0/10.3K/0 iops] [eta 00m:00s]
xyz: (groupid=0, jobs=1): err= 0: pid=3507: Mon Apr 11 16:09:01 2016
write: io=1024.0MB, bw=42442KB/s, iops=10610, runt= 24706msec
clat (usec): min=70, max=2200, avg=92.32, stdev=17.43
lat (usec): min=70, max=2201, avg=92.57, stdev=17.44
clat percentiles (usec):
| 1.00th=[ 73], 5.00th=[ 76], 10.00th=[ 78], 20.00th=[ 82],
| 30.00th=[ 84], 40.00th=[ 86], 50.00th=[ 88], 60.00th=[ 91],
| 70.00th=[ 95], 80.00th=[ 100], 90.00th=[ 112], 95.00th=[ 124],
| 99.00th=[ 141], 99.50th=[ 149], 99.90th=[ 171], 99.95th=[ 191],
| 99.99th=[ 644]
bw (KB /s): min=39192, max=46832, per=100.00%, avg=42449.80, stdev=1713.07
lat (usec): 100=79.44%, 250=20.53%, 500=0.01%, 750=0.01%, 1000=0.01%
lat (msec): 2=0.01%, 4=0.01%
cpu      : usr=2.15%, sys=11.96%, ctx=262146, majf=0, minf=32
IO depths: 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete: 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued  : total=r=0/w=262144/d=0, short-r=0/w=0/d=0, drop=r=0/w=0/d=0
latency  : target=0, window=0, percentile=100.00%, depth=1
```

```
Run status group 0 (all jobs):
WRITE: io=1024.0MB, aggrb=42442KB/s, minb=42442KB/s, maxb=42442KB/s, mint=24706msec,
maxt=24706msec
```

```
Disk stats (read/write):
vda: ios=0/261747, merge=0/0, ticks=0/22147, in_queue=22114, util=89.35%
```

- New flavors
 - To match smaller disks
- SSD caching ...
 - ZFS ZIL/l2arc in front of Cinder volumes?

bcache/dm-cache update:

- dm-cache for lxplus: worked OK
- we'll drop it none the less in favor of bcache (performance vs. operations)

Operations: Retirement Campaign

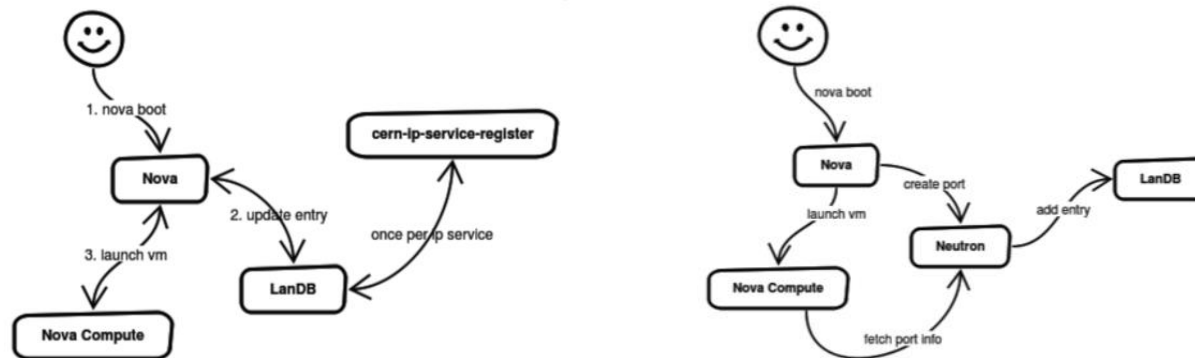
- About 1'600 nodes to retire from the service by 3Q16
 - ~1'200 compute (started), ~400 with services
- We have gained some experience with (manual) migration
 - Live and cold
 - Seems to work reliably (where it can)
- We have developed a tool that you can instruct to drain hypervisor (or simply live-migrate given VMs)
 - Main tasks are VM classification and progress monitoring
 - The nova scheduler will pick the target (nova patch)
- We will use the “IP service bridging” mentioned at BNL
 - See [Carles Kishimoto Bisbe's talk](#) for all details

Operations: Recent Issues

- **Unexpected VM shutdowns**
 - Symptoms: VMs shut down without reason, looks like 'shutdown -h now'
 - Rare! ~1-2 VMs / month in >15'000 VMs
 - SLC6, CC7, no core dumps (well almost ...), VMs w/ volumes
 - Seems we're hitting a Ceph Issue 6480: using non thread-safe libnss functions
- **libvirtd updates can freeze VMs**
 - Symptoms: Nova is blocked (b/c libvirt is stuck), stuck msgs in Rabbit, no operations get through
 - Work-around: `virsh destroy <offending_domain>`
- **Blocked I/Os on SSD-hosted VMs**
 - Symptoms: I/Os in the VMs completely blocked
 - Happened only on new SSD hypervisors
 - Fix: Update to recent `qemu-kvm` version

Operations: Neutron Migration

- We'll need to replace nova-network
 - It's going to be deprecated (really really really this time)



- We have a number of patches to adapt to the CERN network constraints
 - We patched nova for each release ...
 - ... neutron allows for out-of-tree plugins!
- New potential features (Tenant networks, LBaaS, Floating IPs, ...)

Operations: Neutron Status

- We have a working Neutron cell ...
 - Neutron control plane in Liberty (fully HA)
 - Bridge agent in Kilo (nova)
- ... but there is still quite some work to do.
 - IPv6, custom DHCP, network naming, ...
- “As mentioned, there is currently no way to cleanly migrate from nova-network to neutron.”
 - All efforts to establish a general migration path failed so far
 - Should be OK for us, various options (incl. in-place, w/ migration, ...)

What's new: Keystone

- Access to EduGAIN users via Horizon
 - Allow (limited) access given appropriate membership
- Additional project attributes
 - LanDB owner/responsible for extended access
 - More attributes being discussed (accounting groups)
- Introduced endpoint filtering
 - Allows access to features on a per tenant basis
 - Open service on demand (e.g. Magnum)




What's new: EC2 API Project

- EC2 API support in nova was deprecated with Kilo, will be removed with Mitaka
- New EC2API project
 - CERN provided initial packaging and Puppet modules
 - Currently being tested with first users at CERN
 - <https://github.com/openstack/ec2-api>

What's new: Container Integration

- **Magnum: OpenStack project to treat Container Orchestration Engines (COEs) as 1st class resources**
- **Pre-production service available**
 - Supporting Docker Swarm and Kubernetes for now
- **Many users interested, usage ramping up**
 - GitLab CI, Jupyter/Swan, FTS, ...
- **All details in Bertrand's talk on Thursday!**



Future Plans

- Investigate Ironic (Bare metal provisioning)
 - OpenStack as one interface for compute resource provisioning
 - Allow for complete accounting

- Replace Hyper-V by qemu/kvm?
 - Successfully done at other sites
 - Remove dependency on Windows expertise
 - Reduce complexity in service setup



Summary

- **Cloud service continues to grow and mature**
 - While experimental, good experience with cells for scaling
 - Experience gained helps with general resource provisioning
 - New features added (federation, containers)
 - Expansion planned (bare metal provisioning)
- **Major operational challenges ahead**
 - Transparent retirement of service hosts
 - Replacement of network layer
- <http://openstack-in-production.blogspot.com>

