



CERN Cloud Infrastructure Report

Arne Wiebalck
for the CERN Cloud Team

HEPiX Autumn Meeting
Lincoln, Nebraska, U.S.
Oct 17, 2014

Numbers
Operations
Issues
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 two rolling upgrades since, now on "Icehouse"
 - Nova, Glance, Keystone, Horizon, Cinder, Ceilometer



CERN Cloud in Numbers (1)

- 3'000 hypervisors at the moment
 - Vast majority qemu/kvm on SLC6 (~100 Hyper-V hosts)
 - 550 HVs at Wigner in Hungary (so far only for the batch service)
 - 220 HVs on critical power (currently being deployed)
 - 2'000 HVs used by batch, rest shared by users, services, experiments
 - Additional 2'250 hypervisors will be added early 2015
- 8'000 VMs
 - Batch: 2'000
- 64k Cores
- 128 TB RAM






CERN Cloud in Numbers (2)

- 1'100 images/snapshots
 - Backed by Ceph
 - Finger trouble led to loss of 246 images (44 restored from caches)

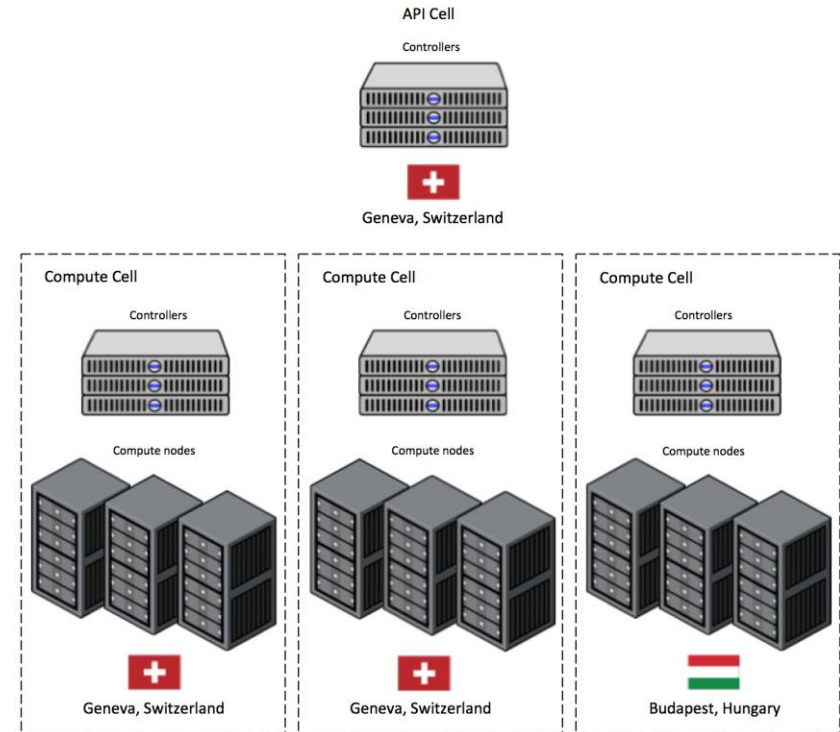
- 600 volumes
 - Backed by Ceph
 - Increasingly popular with users: Space, Tunable IOPS, Attach/Detach

- 700 active projects
 - Personal and shared



Operations: Cells

- Relying on “cell” feature
 - Structures our deployment
 - Needed for scale-out
 - 7 cells, size range: 4 ... 1500 nodes
 - Inter-cell consistency
- Top level controllers
 - Run multiple components
 - Currently being split to allow for easier per-component scale-out and independent upgrades
- H/W allocation non-trivial
 - Shared vs. dedicated
 - Physical location
 - Critical power
 - Hardware models: (10)GbE, SSDs, different generations ...



Operations: Updates

- Rolling upgrade to 'Icehouse' just finished!



- Patches needed to be ported, mostly for Nova
 - Careful testing of each component, non-trivial e.g. for Nova
 - Risk-management: Sequential, component-wise updates
 - Service incidents
 - Shift of priorities, e.g. allocation of new resources
-
- Juno on RDO RHEL6? Not planned at the moment ...
 - In touch with the community to see what can be done
 - CERN CentOS 7 in testing already

Operations: Consultancy

- “I cannot use virtual servers for my service.”
 - Most worries around VM IO performance
 - Explain service offering, suggest tests, tune
 - See my talk on the “Ixplus problem” later
- Requests for special flavors
 - Bigger disks, more RAM per core, a lot of RAM
 - Make VM packing a hard problem
- Requests for special services
 - Virtualization service vs Cloud service misunderstanding
 - Not everything that is technically doable should be done

“Users learn how to use the service while we learn how to run it.”

Before we come to some issues ...

OpenStack is a solid cloud management product.

The CERN Cloud is a stable production service.

No VMs were harmed during any of the following incidents!


Issues: Rabbit



- Standard problem: unconsumed messages
 - Compute node stops acting
 - Watcher logic now in place
 - Has been improved with recent upgrades
- Incident spread over several days:
 - Problem on Hadoop cluster
 - Rabbit kept data in memory
 - Flushed to disk at some point
 - Filled up disks (but stopped “in time”)
 - Something else filled the disk
 - Rabbit crashed, restart did not bring back queues
 - Restart of the whole service
 - Queues found corrupted, queues needed deletion
 - Restart of (all) OpenStack services eventually ...
- Messaging (so far) was a time-consuming component.

Issues: smartd vs. mdadm

- Symptom: VMs spend large fraction in IOwait
 - Hypervisor seemed more or less OK



CPU Utilization on an ATLAS PanDa VM during the incident

- Started during weekend ...
- Various VMs affected ...
- Not at the same hour ...
- Not on all machines ...
- Regular SMART tests introduced some days before! (spread out over the day)
- Mdadm scrubs starts by default on Sunday night at 1 am!
- Both at the same time plus multiple VMs on the HV break things ...

Issues: Unsolved so far ...

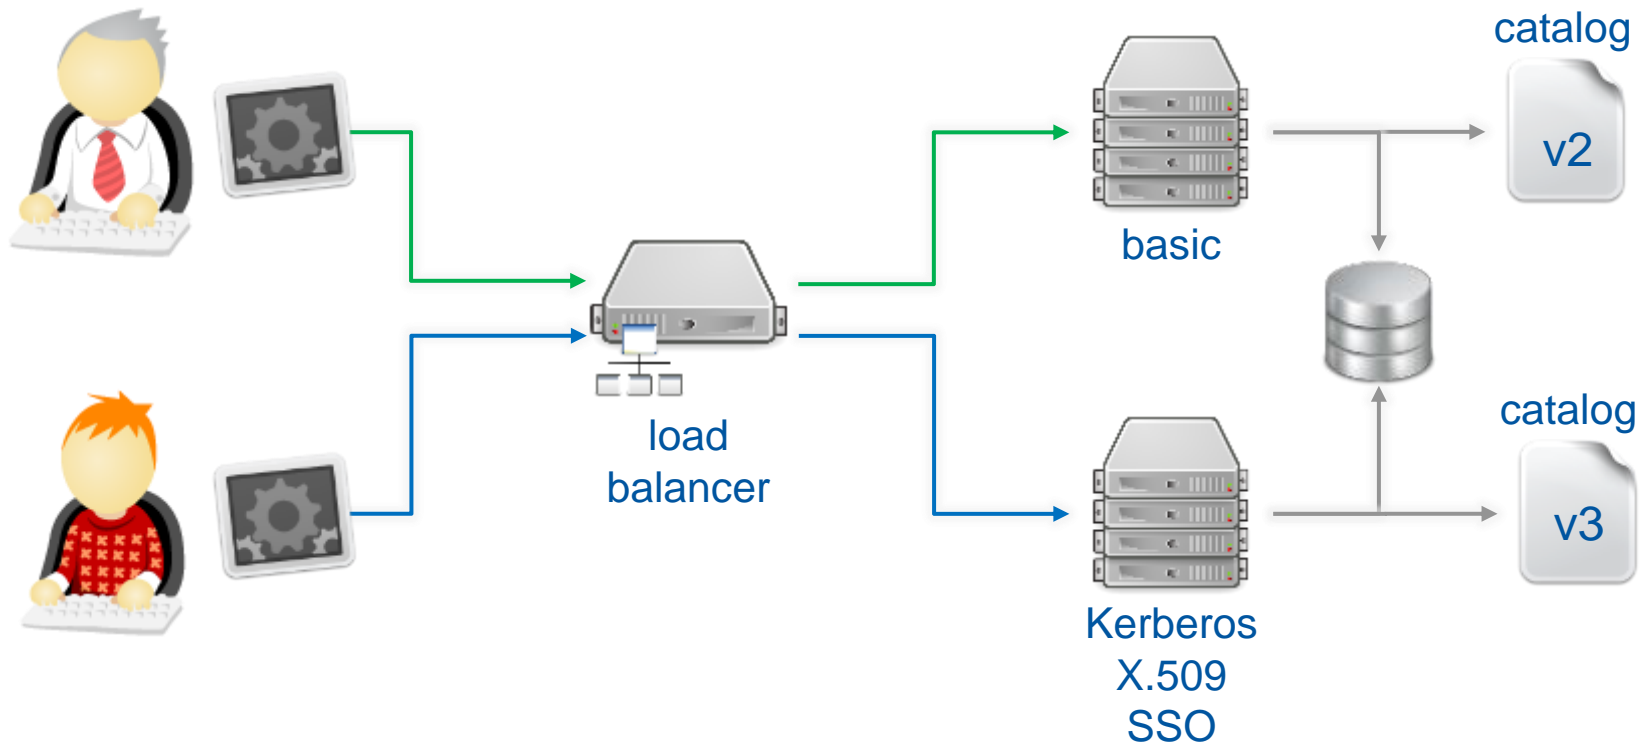
- Spontaneous VM shutdowns
 - For OpenStack, these look like normal shutdowns
- Unkillable libvirtd processes
 - In state 'R'
- qemu/kvm crashes of two-volume VMs
 - No usable dump yet
- ksmd creating high load
 - Long standing or new issue?
- ...

WIP: External Authentication (1)

- **User requirements**
 - Secure way to authenticate (Kerberos, X.509)
 - Enable federated use-cases (SSO)
- **Service requirements**
 - Only available on Keystone API v3
 - Backwards compatible and transparent to end-users
 - Client only supports one API version in the Cloud
- **Community based**
 - Working actively on a solution
 - Involves several projects (CLIs and services)

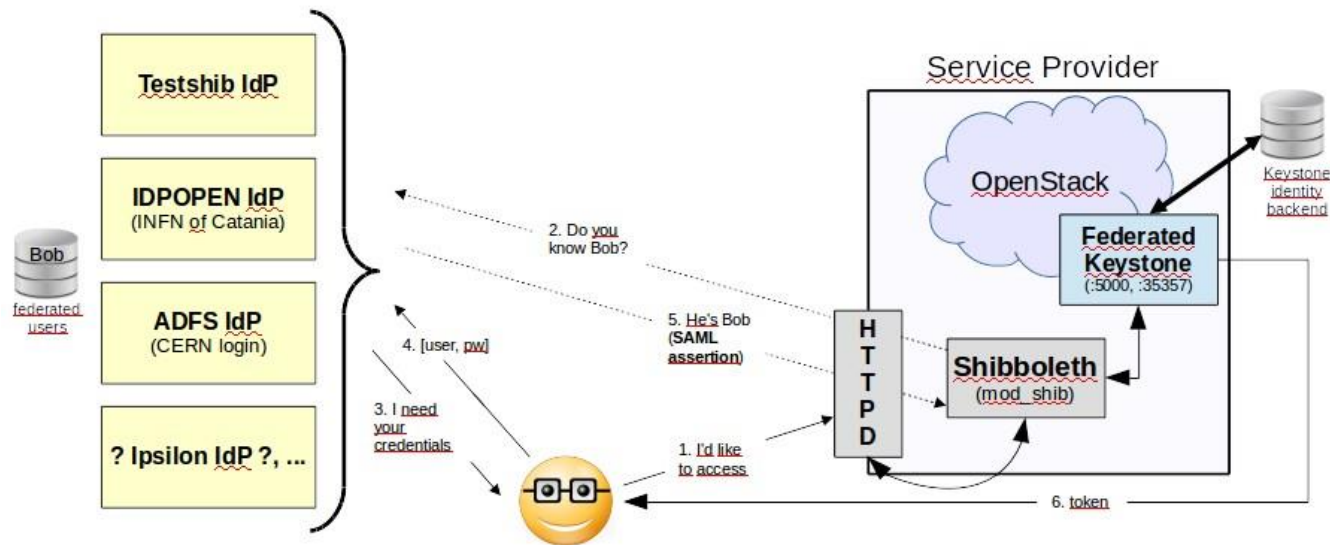
WIP: External Authentication (2)

- CERN Release planned for end of Oct 2014
 - Close to upcoming community solution



WIP: Cloud Federation 1

- Use several clouds with a federated identity
 - E.g.: combine resources of CERN IT's private cloud with experiments' clouds in the pit
 - Authentication done only against one IdP selected from a set



WIP: Cloud Federation 2

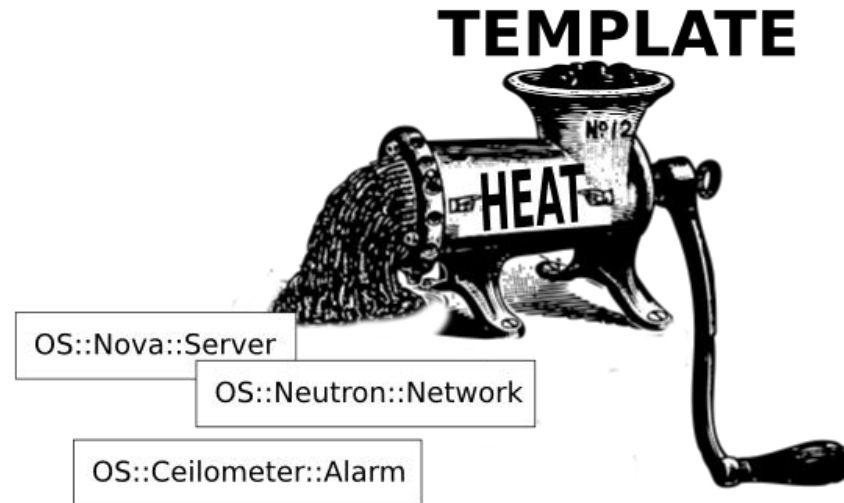
- OpenStack support for identity federation
 - Available with Icehouse
 - OpenStack Identity Service (Keystone) acts as a Service Provider mapping SAML assertions to roles
 - Support for SAML2 (OpenID and ABFAB to come)
- Cloud federation status at CERN
 - Collaboration with Rackspace
 - Successfully tested with INFN's IdP
 - CERN to join EduGAIN federation, providing cloud resources to other federation members
- Outlook
 - Native Web SSO support
 - Inter cloud image sharing, Inter cloud SDNs



CERNopenlab



WIP: Orchestration with Heat (1)



- Heat provides a mechanism for orchestrating OpenStack resources through templates
 - Analogous to AWS cloud formation
- Auto-scaling feature is main plus

WIP: Heat HOT template example



```
...
cpu_alarm_high:
  type: OS::Ceilometer::Alarm
  properties:
    description: Scale-up if the average CPU > 50% for 1 minute
    meter_name: cpu_util
    statistic: avg
    period: 60
    evaluation_periods: 1
    threshold: 50
    alarm_actions:
      - {get_attr: [web_server_scaleup_policy, alarm_url]}
    matching_metadata: {metadata.user_metadata.stack!: {get_param: "OS::stack_id"}}
    comparison_operator: gt
...
...
web_server_scaleup_policy:
  type: OS::Heat::ScalingPolicy
  properties:
    adjustment_type: change_in_capacity
    auto_scaling_group_id: {get_resource: web_server_group}
    cooldown: 60
    scaling_adjustment: 1
...
```

WIP: Orchestration with Heat (2)



- **Current Status**

- Heat test environment connected to production infrastructure available only for Cloud team
- Missing features (auto-scaling, load balancers) due to identity constraints

- **Deployment Plan**

- OpenStack Juno release will enable multiple identity drivers for different domains (LDAP, SQL)

Summary

- Cloud service at CERN still growing rapidly
 - +6'000 VMs in the past year
 - We will double the capacity in the next 6 months
- We're exploring options to address user expectations and requests
 - Performance of physical hardware is expected
 - Cloud-style of running services
- Clearly a learning process for users and operations team

Questions?