CERN

The CERN OpenStack Cloud

Compute Resource Provisioning for the Large Hadron Collider

Jan van Eldik

for the CERN Cloud Team

OpenStack Days Nordic

Stockholm

September 22, 2016

Agenda

- **Introduction to CERN**
- **Computing at CERN scale**
- **Cloud Service Overview**
- **Operations**
- **Performance**
- **Outlook**

CERN: home.cern

- European Organization for Nuclear Research (***C***onseil ***E***uropéen pour la ***R***echerche ***N***ucléaire)
  - Founded in 1954, today 22 member states
  - World's largest particle physics laboratory
  - Located at Franco-Swiss border near Geneva
  - ~2'300 staff members, >12'500 users
  - Budget: ~1000 MCHF (2016)

- CERN's mission
  - Answer fundamental questions of the universe
  - Advance the technology frontiers
  - Train the scientists and engineers of tomorrow
  - Bring nations together

The Large Hadron Collider (LHC)

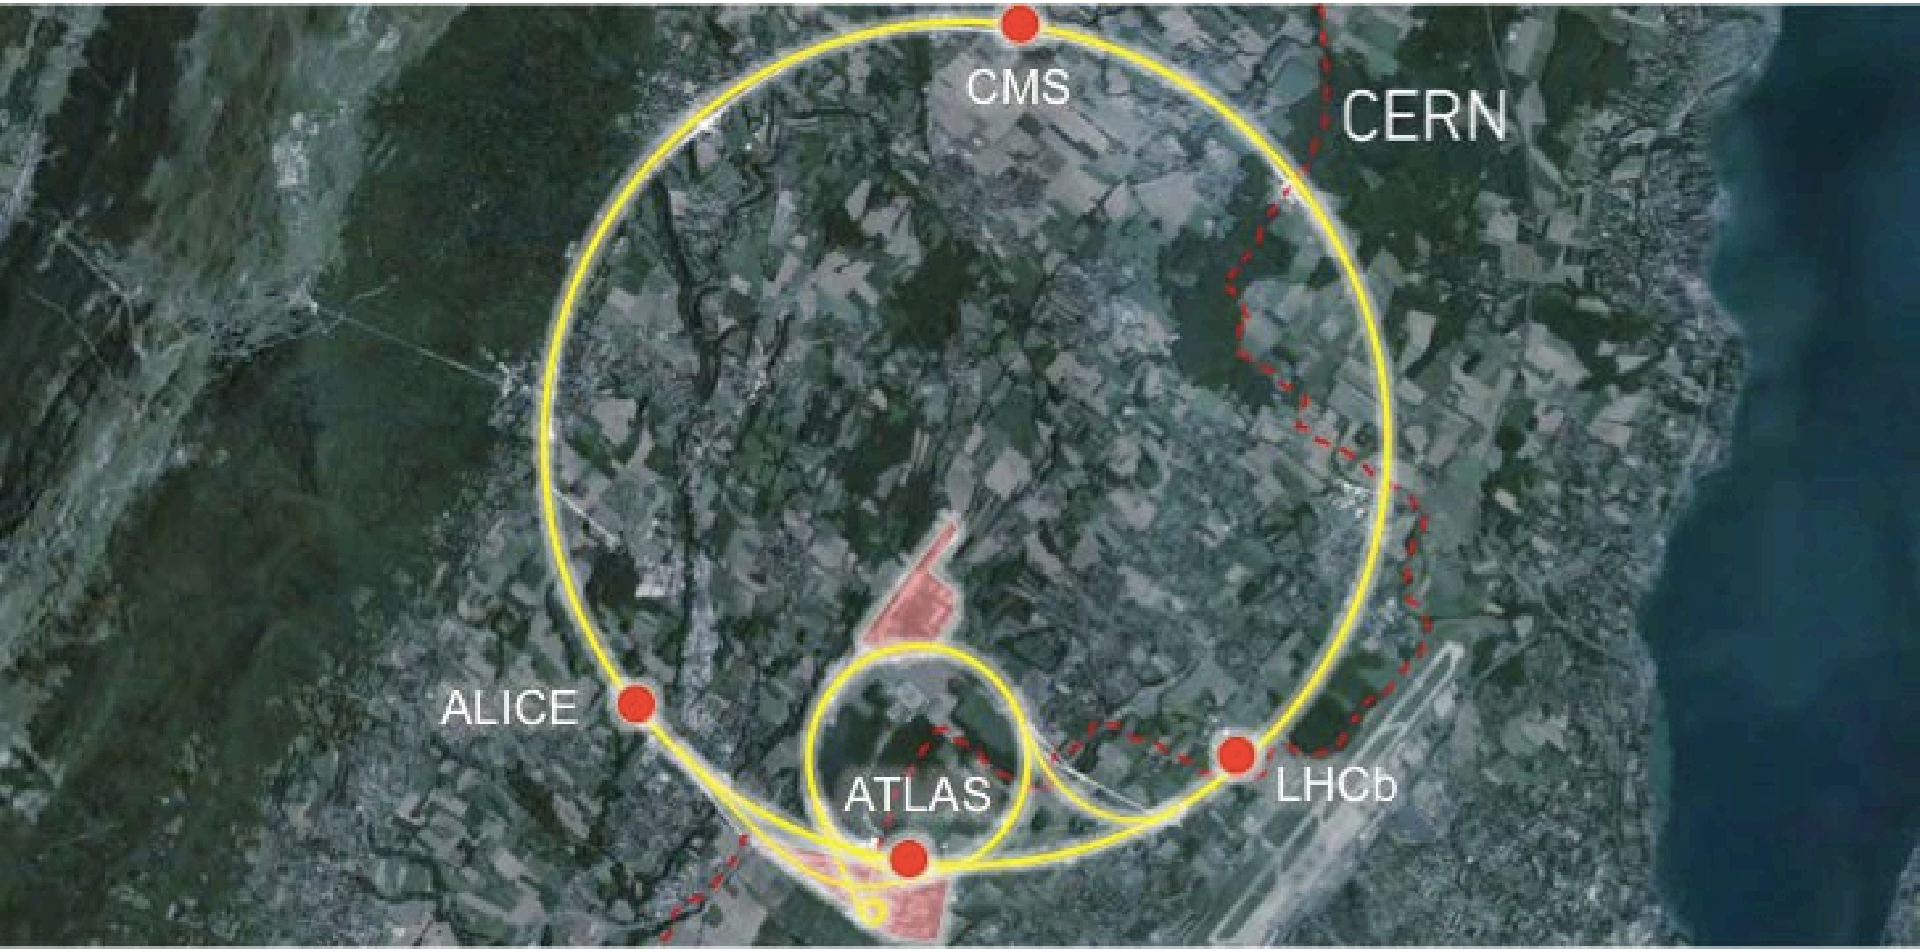

Largest machine on earth: 27km circumference

LHC: 9’600 Magnets for Beam Control

1232 superconducting dipoles for bending: 14m, 35t, 8.3T, 12kA

LHC: Coldest Temperature

World’s largest cryogenic system: colder than outer space (1.9K/2.7K), 120t of He

LHC: Highest Vacuum

Vacuum system: 104 km of pipes, 10^{-10}-10^{-11} mbar (comparable to the moon)

LHC: Detectors

Four main experiments to study the fundamental properties of the universe

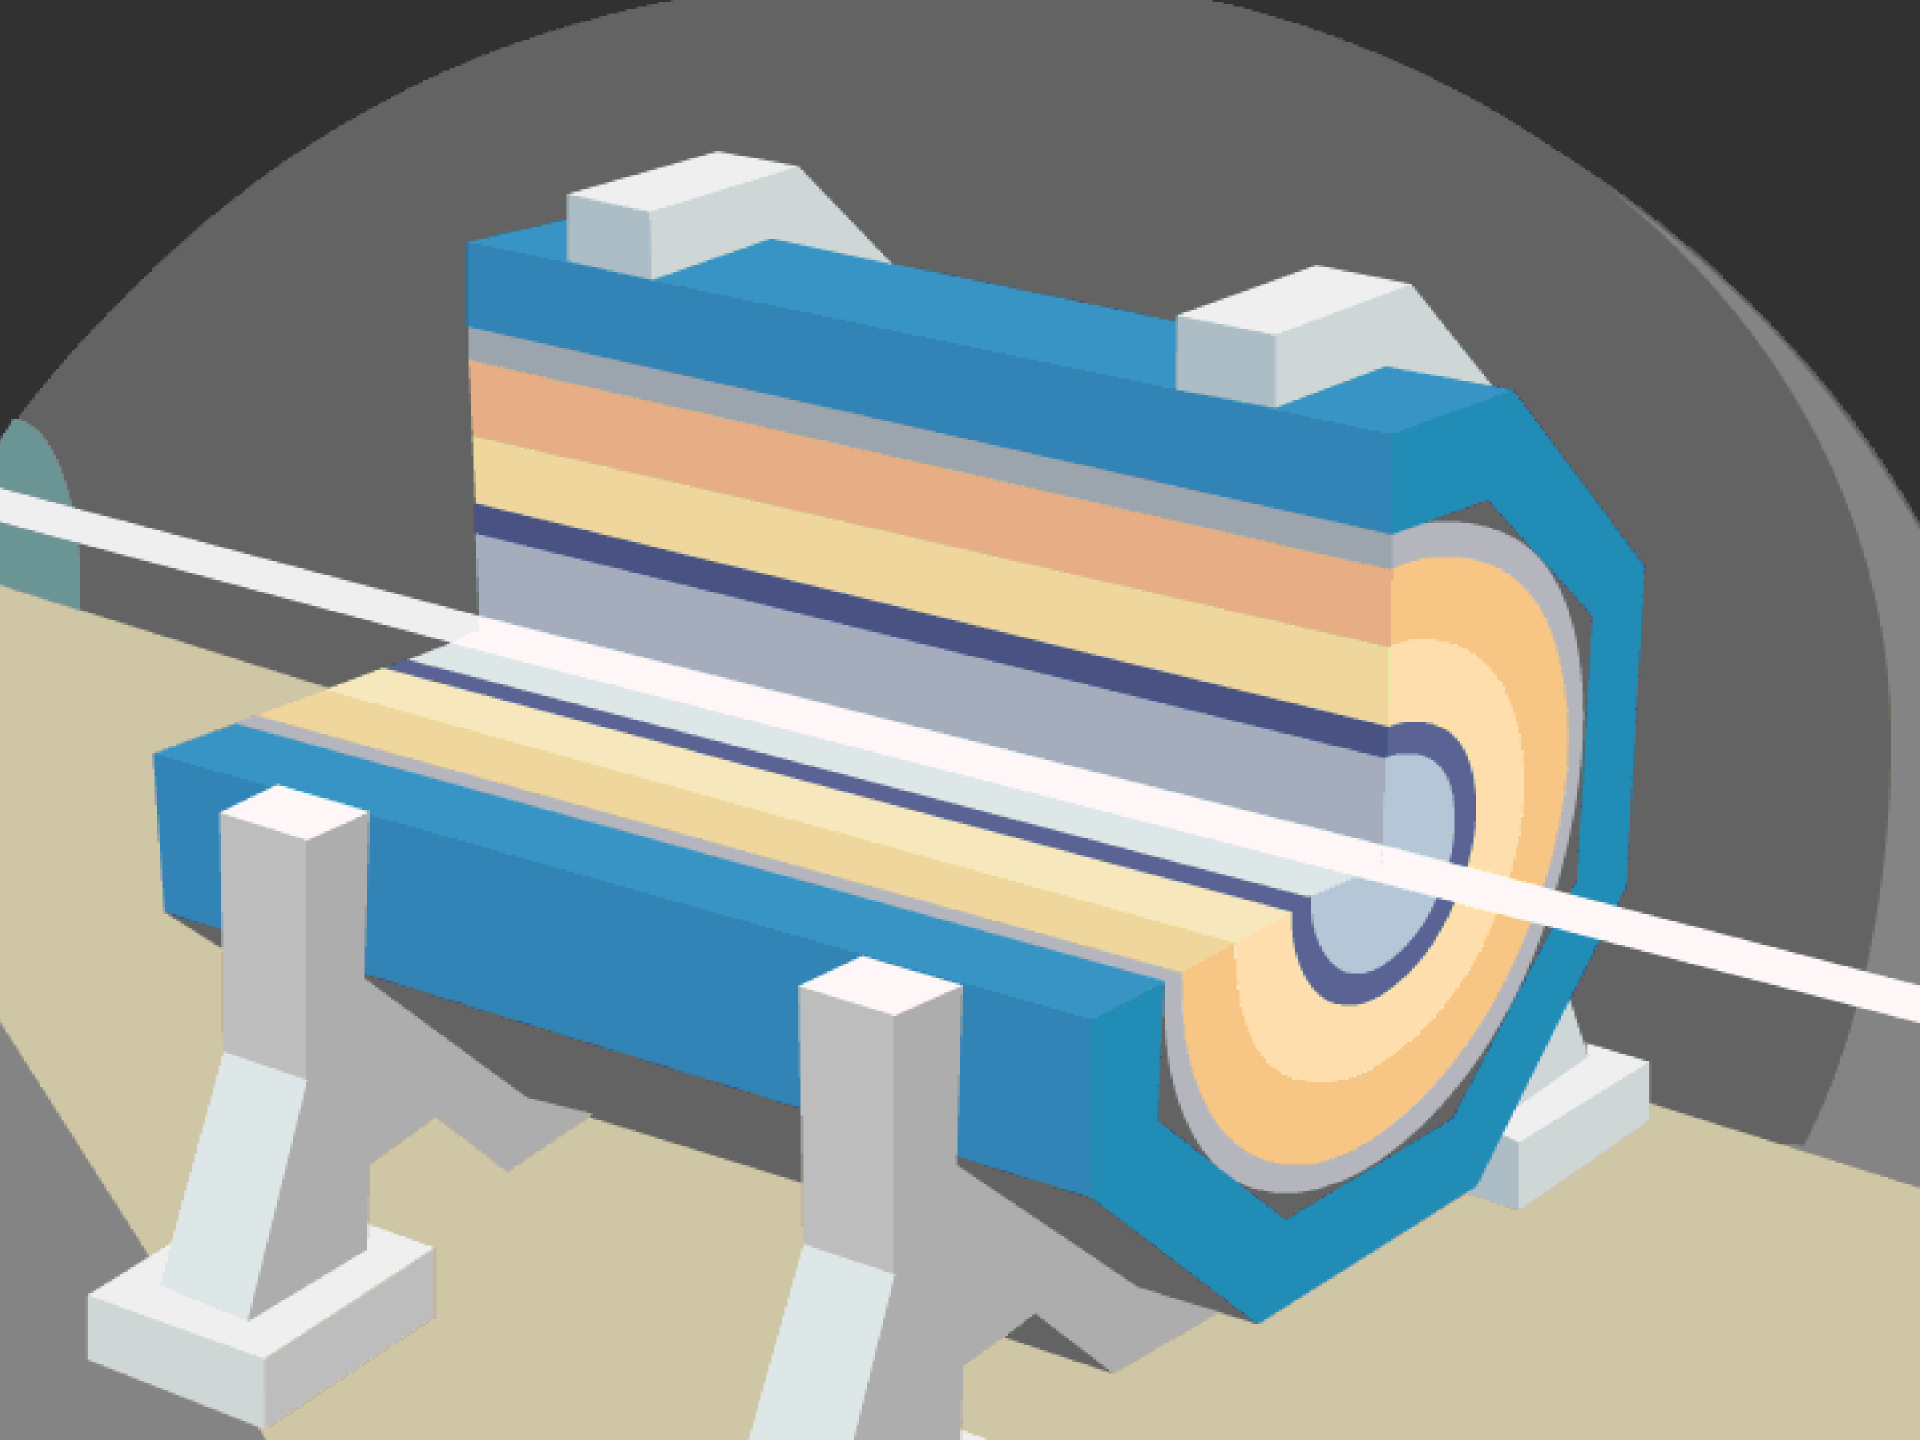

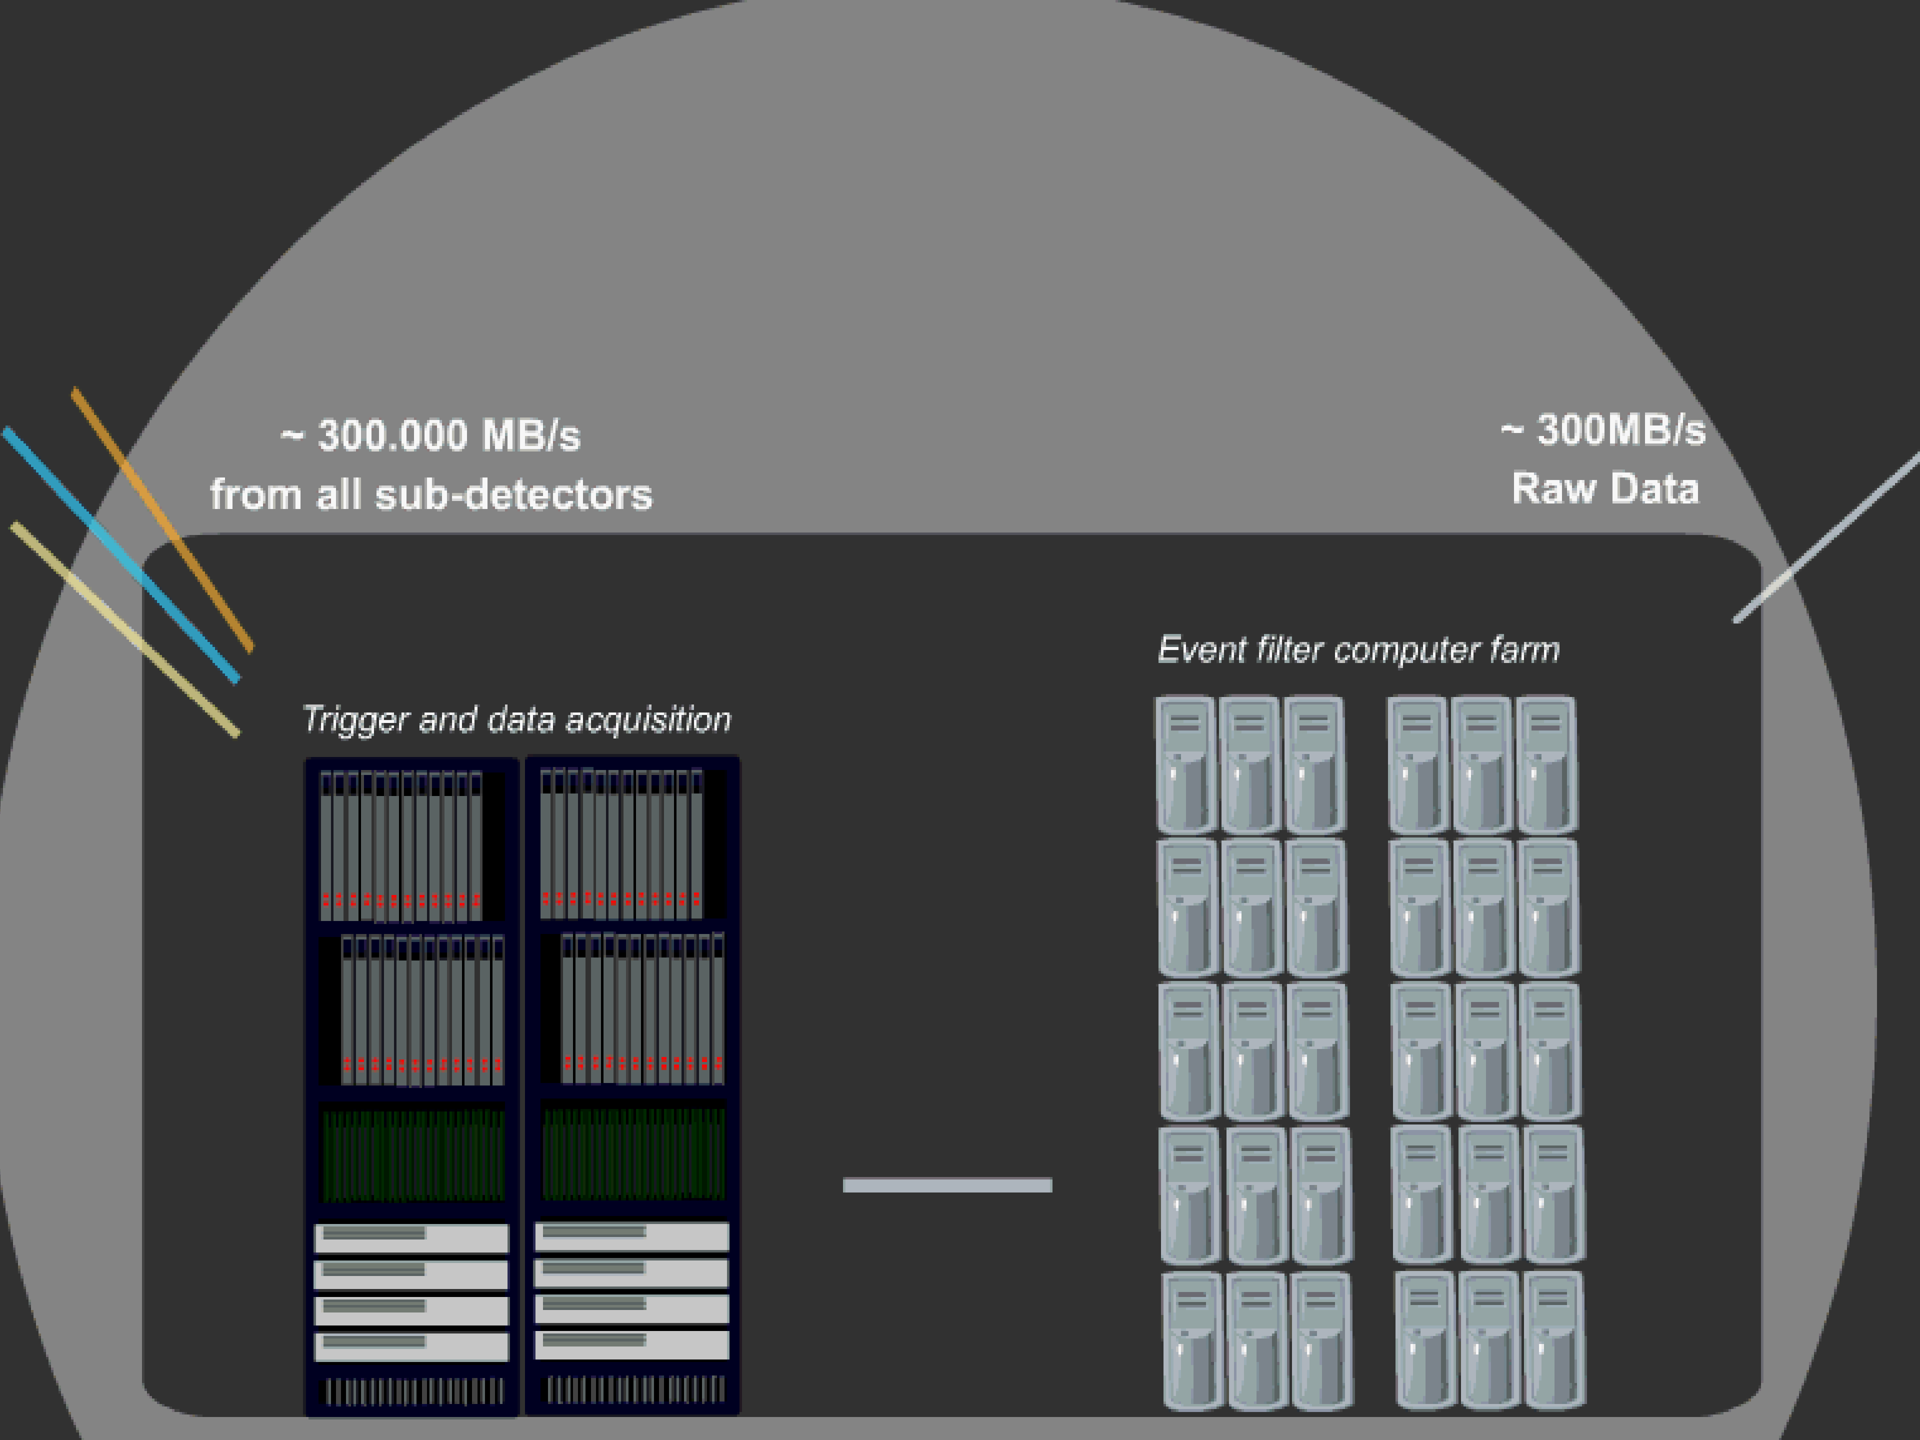
~ 300.000 MB/s
from all sub-detectors
~ 300MB/s
Raw Data
Event filter computer farm
Trigger and data acquisition

Tier 0 at CERN: Acquisition, First pass reconstruction, Storage & Distribution

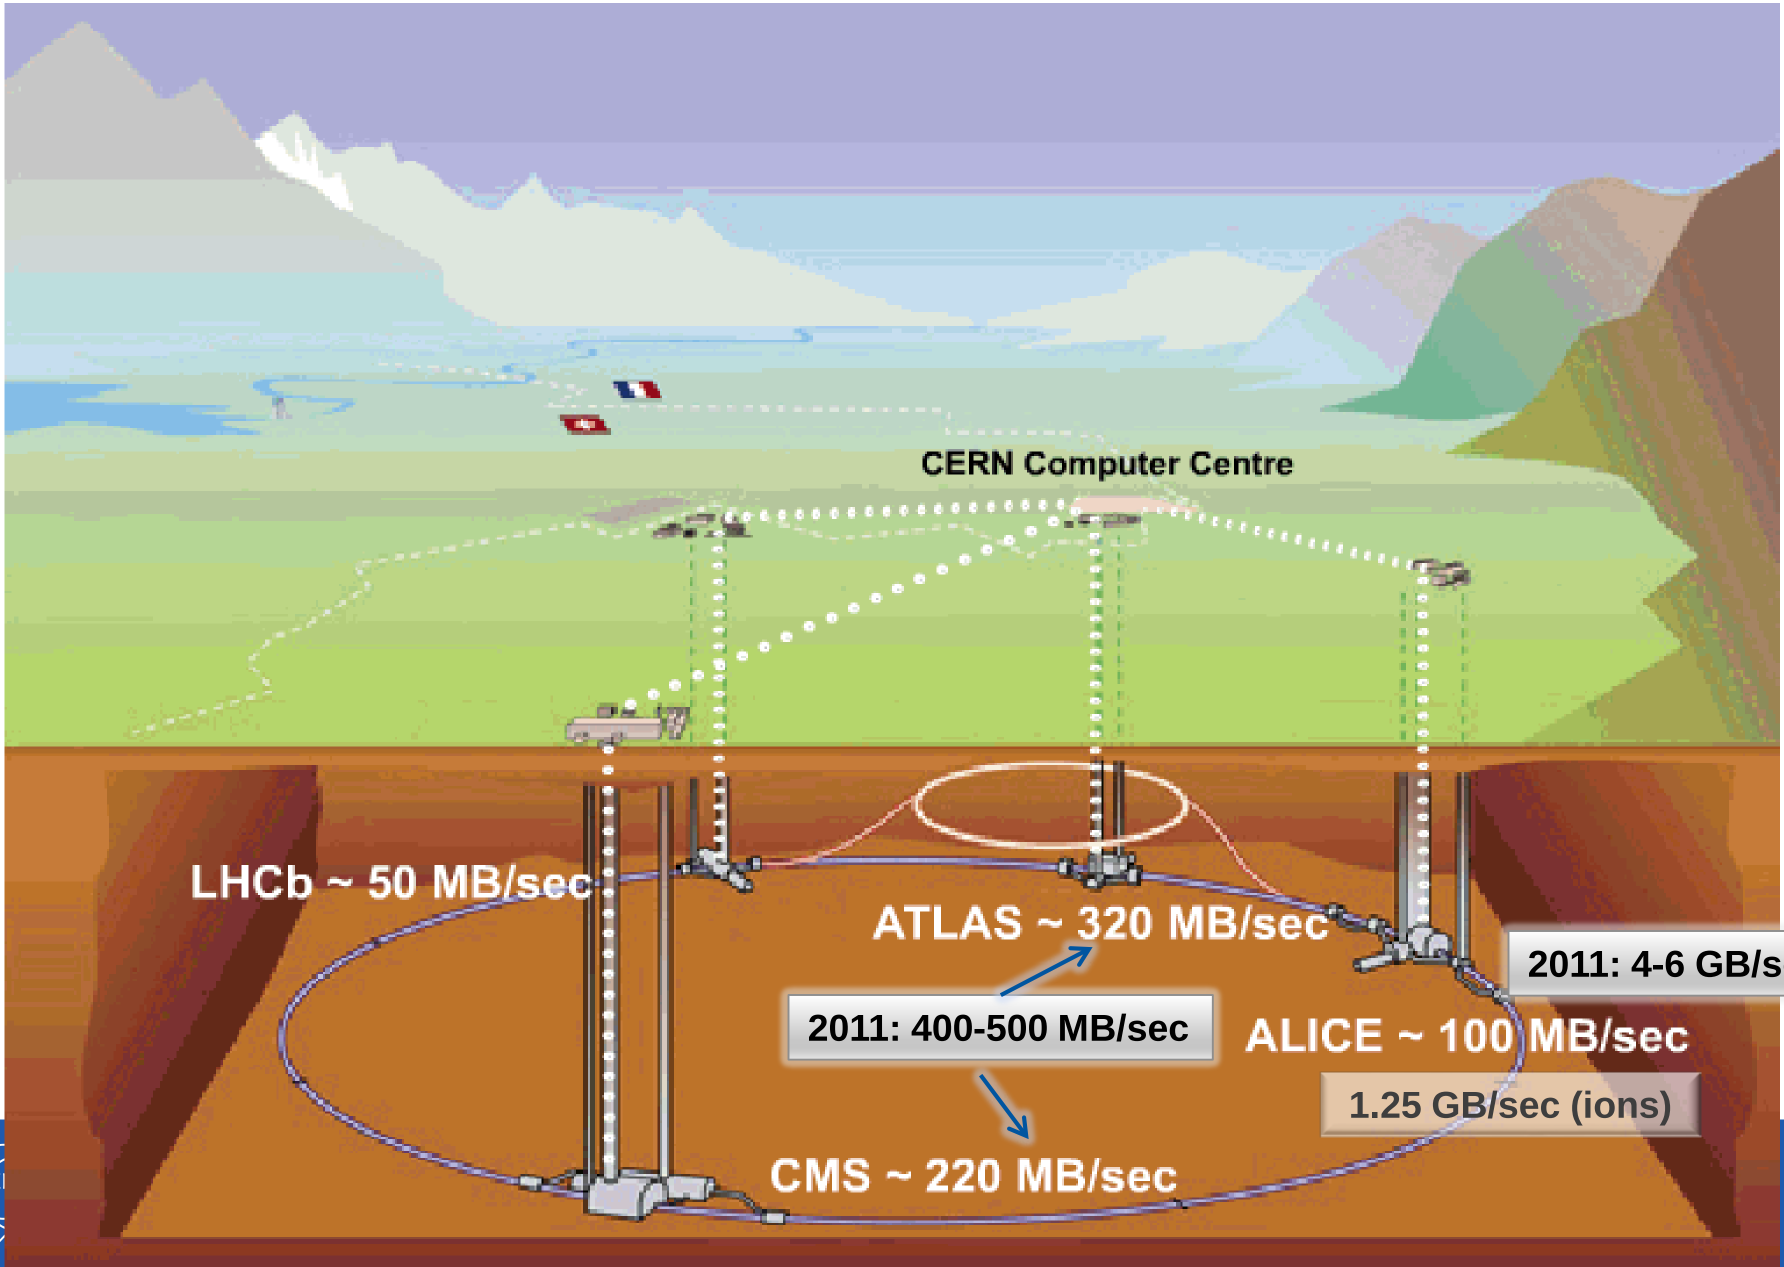

Solution: the Grid

- Use the Grid to unite computing resources of particle physics institutes around the world

The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe

LHC: World-wide Computing Grid

TIER-0 (CERN):
data recording, reconstruction and distribution

TIER-1:
permanent storage, re-processing, analysis

TIER-2:
simulation, end-user analysis

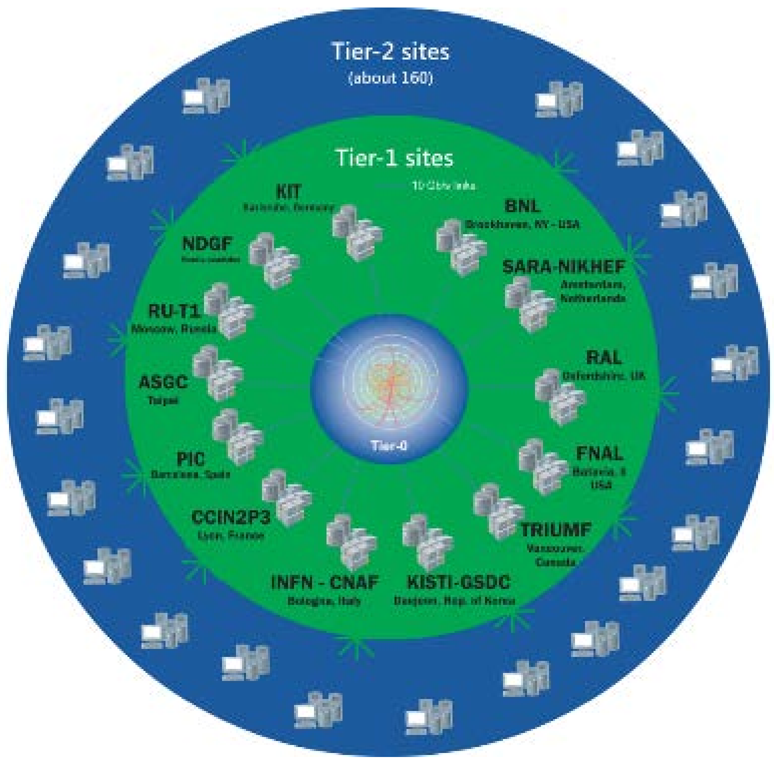

nearly 170 sites, 40 countries

~350’000 cores

500 PB of storage

> 2 million jobs/day

10-100 Gb links

LHC: Data & Compute Growth

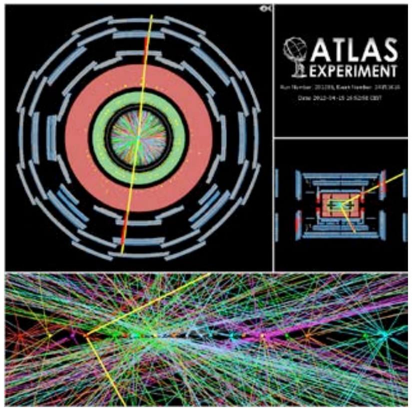

Collisions produce ~1 PB/s

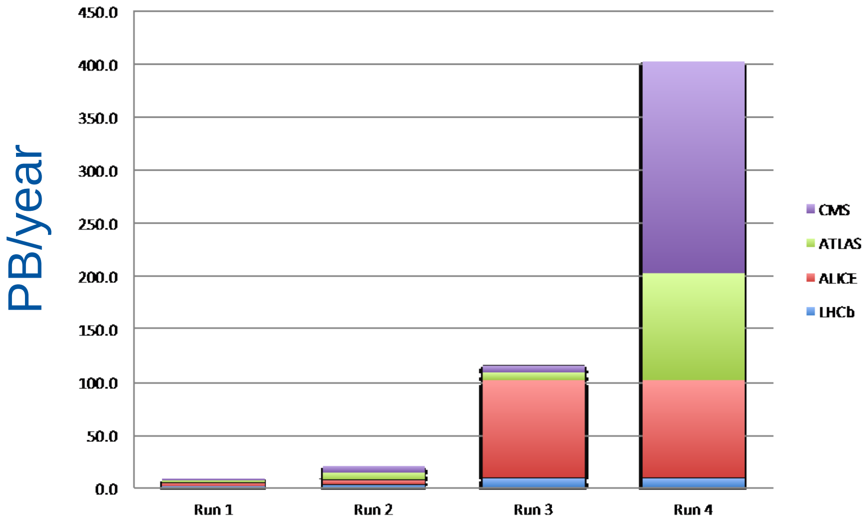

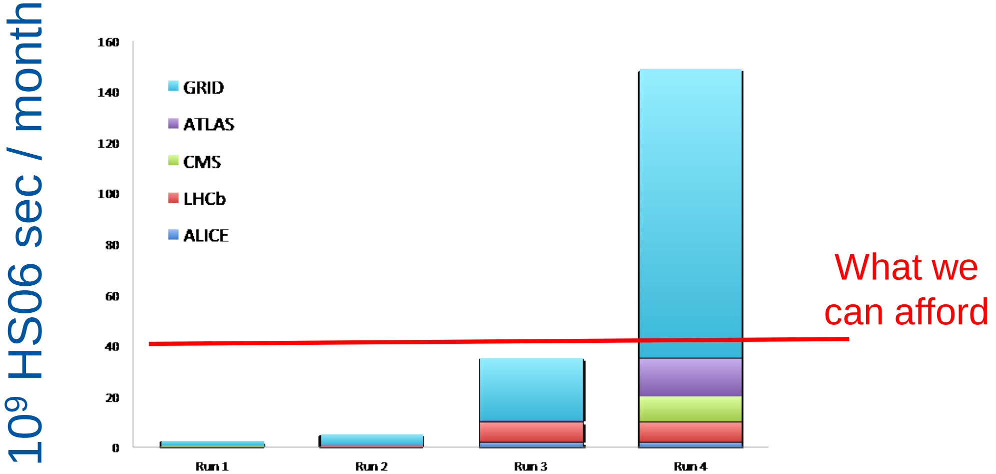

2012: Enter the cloud

- Aim: virtualize all the machines
  - Unless really, really, really not possible
- Offer Cloud endpoints to users
- Scale horizontally
- Consolidate server provisioning
  - Yes, use the private cloud for server consolidation usecases as well

OpenStack at CERN

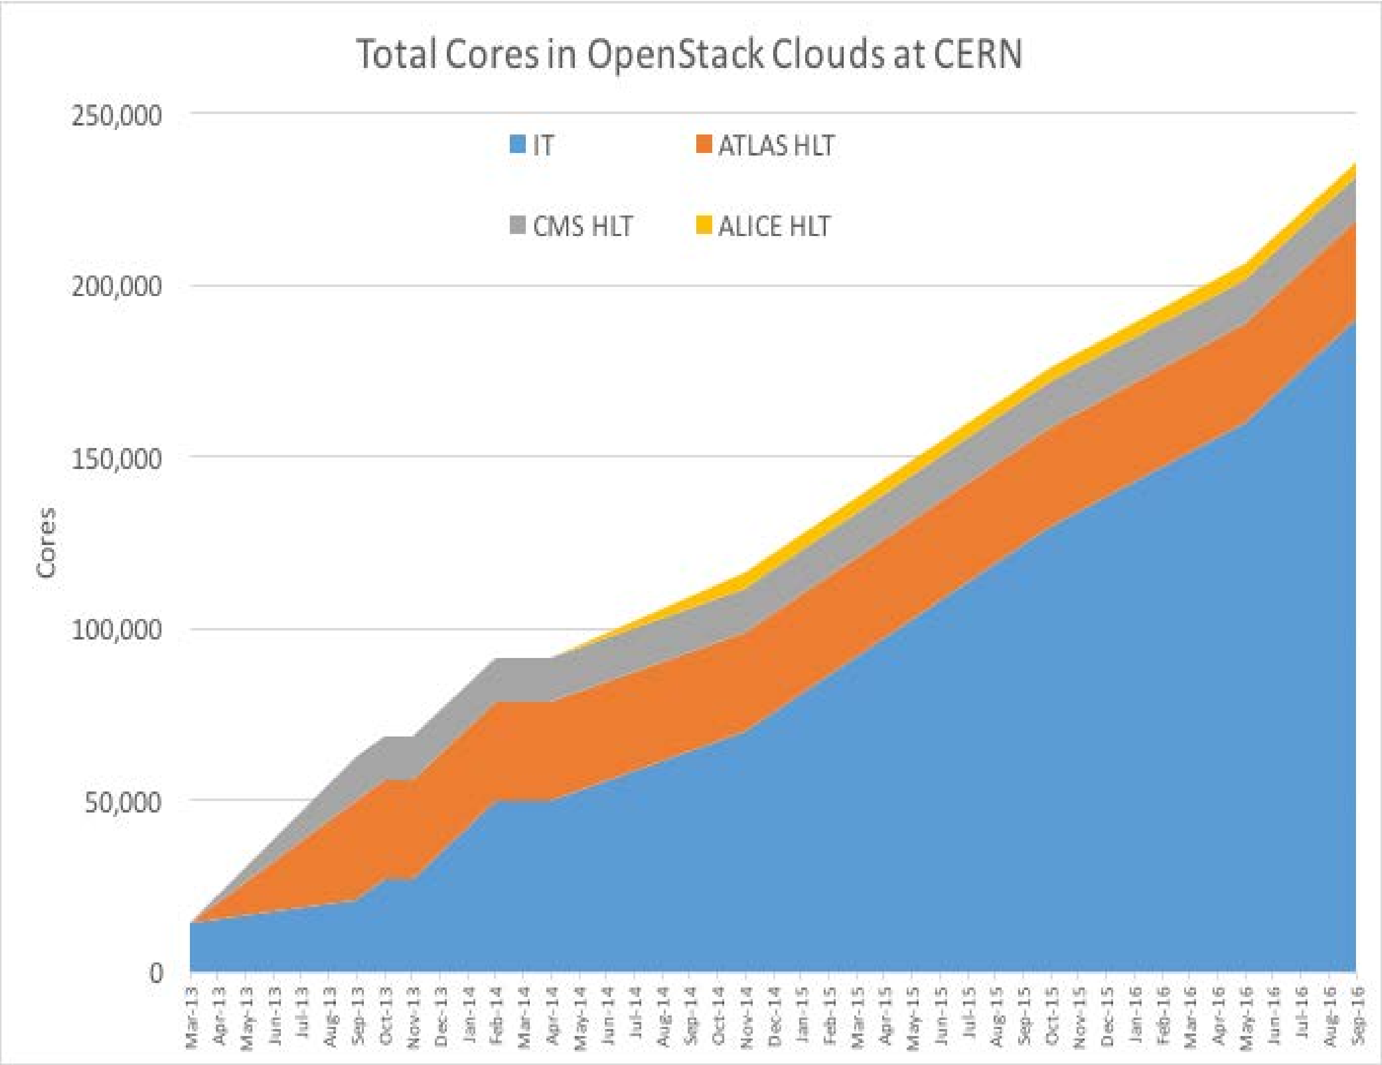

In production:

- 4 clouds
- >200K cores
- >8,000 hypervisors

90% of CERN's compute resources are now delivered on top of OpenStack

openstack™

Cloud Service Context

- CERN IT to enable the laboratory to fulfill its mission
  - Main data center on the Geneva site
  - Wigner data center, Budapest, 23ms distance
  - Connected via two dedicated 100Gbs links

http://goo.gl/maps/K5SoG

- 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 4 rolling upgrades since
  - Currently in transition from Liberty to Mitaka
    - Nova, Glance, Keystone, Horizon, Cinder, Ceilometer, Heat, Neutron, Magnum, Barbican

LIBERTY
THE TWELFTH RELEASE OF OPENSTACK 12

OPENSTACK mitaka 13

CERN Cloud Architecture (1)

- Deployment spans our two data centers
  - 1 region (to have 1 API), ~40 cells
  - Cells map use cases
    hardware, hypervisor type, location, users, …

- Top cell on physical and virtual 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)

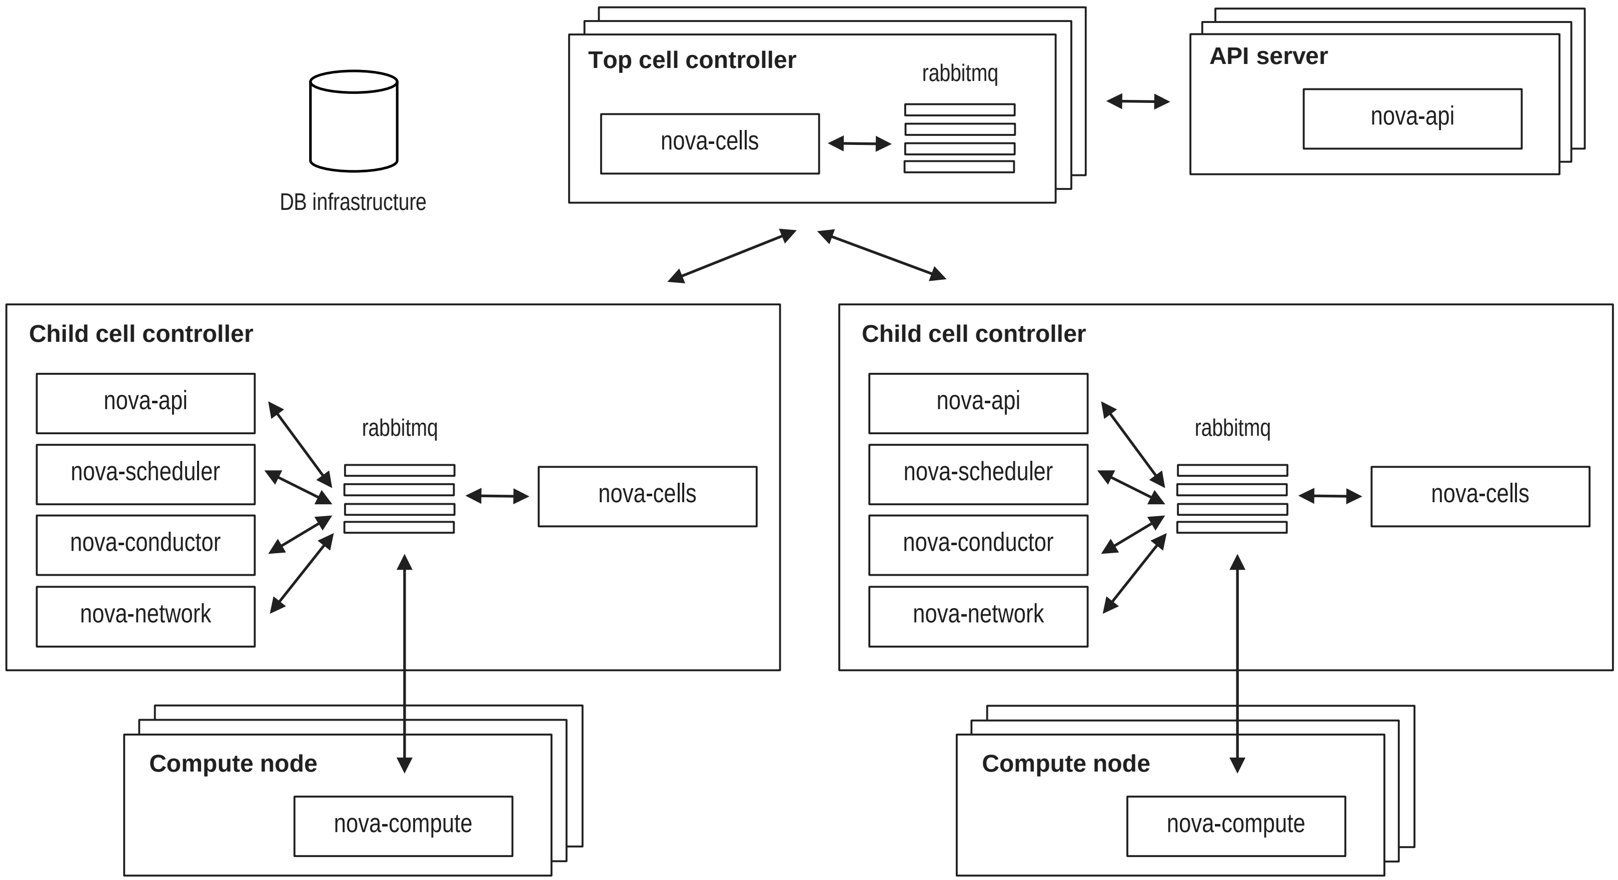

CERN Cloud in Numbers (1)

- ~6700 hypervisors in production
  - Split over 40+ Nova cells
  - Vast majority qemu/kvm now on CERN CentOS 7 (~150 Hyper-V hosts)
  - ~2’100 HVs at Wigner in Hungary (batch, compute, services)
  - 370 HVs on critical power
- 190k Cores
- ~430 TB RAM
- ~20’000 VMs
- Big increase during 2016!
  - +57k cores in spring
  - +40k cores in autumn

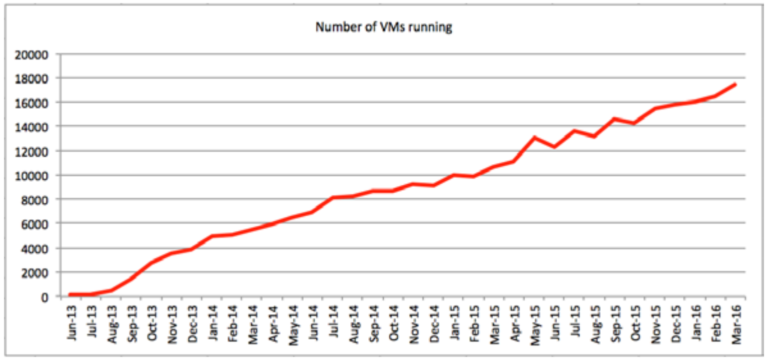

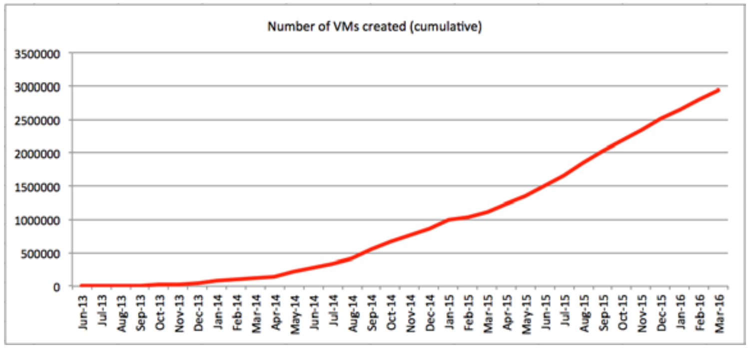

CERN Cloud in Numbers (2)

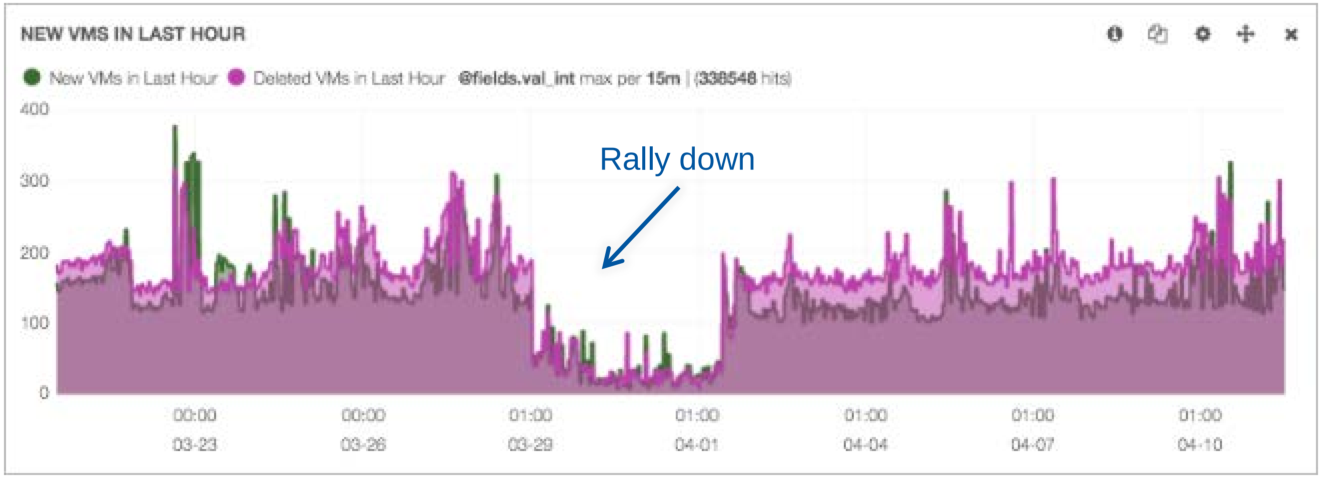

Every 10s a VM gets created or deleted in our cloud!

- 2'700 images/snapshots
  - Glance on Ceph
- 2'300 volumes
  - Cinder on Ceph (& NetApp) in GVA & Wigner

ceph

Only issue during 2 years in prod: Ceph Issue 6480

Software Deployment

- Deployment based on CentOS and RDO
  - Upstream, only patched where necessary
    (e.g. nova/neutron for CERN networks)
  - Some few customizations
  - Works well for us
- Puppet for config’ management
  - Introduced with the adoption of AI paradigm
- We submit upstream whenever possible
  - openstack, openstack-puppet, RDO, …
- Updates done service-by-service over several months
  - Running services on dedicated (virtual) servers helps
    (Exception: ceilometer and nova on compute nodes)
- Upgrade testing done with packstack and devstack
  - Depends on service: from simple DB upgrades to full shadow installations

‘dev’ environment

- Simulate the full CERN cloud environment on your laptop, even offline

- Docker containers in a Kubernetes cluster
  - Clone of central Puppet configuration
  - Mock-up of
    - our two Ceph clusters
    - our secret store
    - our network DB & DHCP
  - Central DB & Rabbit instances
  - One POD per service

- Change & test in seconds

- Full upgrade testing

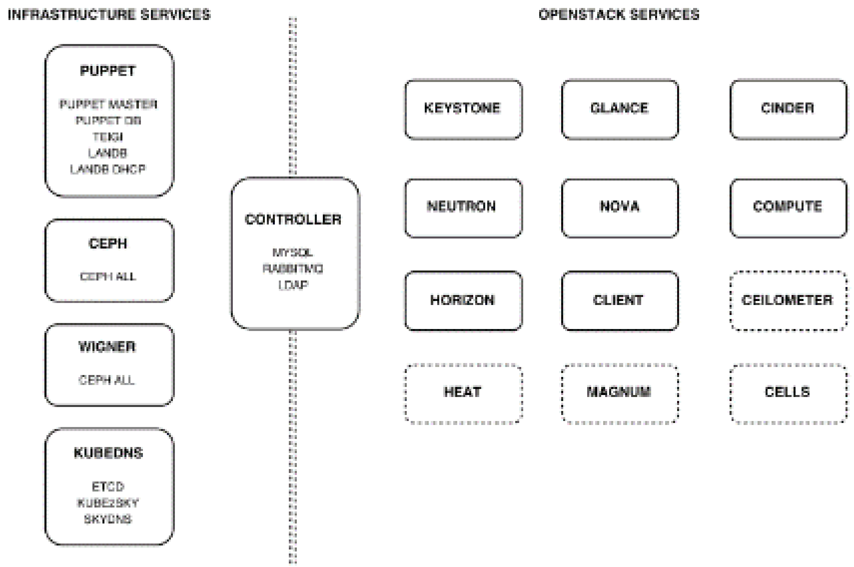

Cloud Service Release Evolution

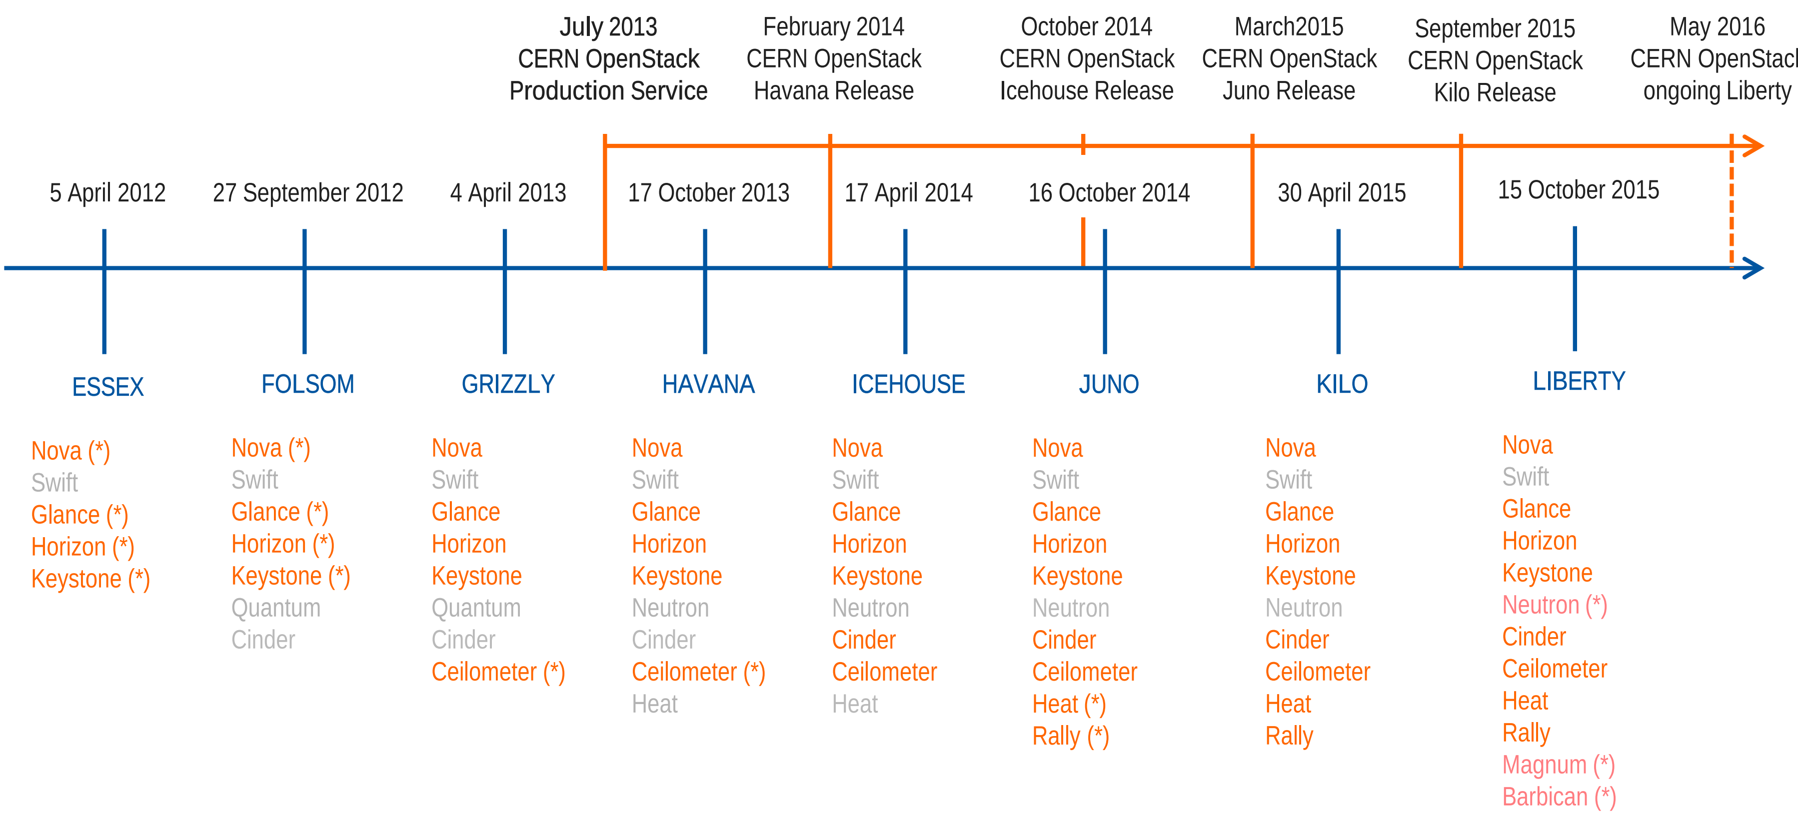

(*) Pilot

Rich Usage Spectrum …

- Batch service
  - Physics data analysis
- IT Services
  - Sometimes built on top of other virtualised services
- Experiment services
  - E.g. build machines
- Engineering services
  - E.g. micro-electronics/chip design
- Infrastructure services
  - E.g. hostel booking, car rental, …
- Personal VMs
  - Development

… rich requirement spectrum!

Usecases (1)

Server consolidation:

- Service nodes, dev boxes, Personal VMs, …
- Performance less important than “durability”
  - Live-migration is desirable
- Persistent block storage is required
- Linux VM @ KVM, Windows VMs @ HyperV
  - Starting to run Win VMs under KVM
- “Pets usecase”: 32K cores, 7500 VMs

Usecases (2)

- Compute workloads
  - Optimize for compute efficiency
    - CPU passthrough, NUMA aware flavours
  - Still, very different workloads
    - IT Batch: LSF and HTCondor, longlived VMs, 8 and 16-core VMs, “full-node” flavors
    - CMS Tier-0: medium-long, 8-core VMs
    - LHCb Vcycle: short-lived, single core VMs
  - Low-SLA, “cattle usecase”
  - 150K cores, 12500 VMs @ 6000 compute nodes

Wide Hardware Spectrum

- The ~6700 hypervisors differ in …
  - Processor architectures: AMD vs. Intel (av. features, NUMA, …)
  - Core-to-RAM ratio (1:2, 1:4, 1:1.5, ...)
  - Core-to-disk ratio (going down with SSDs!)
  - Disk layout (2 HDDs, 3 HDDs, 2 HDDs + 1 SSD, 2 SSDs, …)
  - Network (1GbE vs 10 GbE)
  - Critical vs physics power
  - Physical location (Geneva vs. Budapest)
  - Network domain (LCG vs. GPN vs. TN)
  - CERN CentOS 7, RHEL7, SLC6, Windows
  - …
- Variety reflected/accessible via instance types, cells, projects … variety not necessarily visible to users!
  - We try to keep things simple and hide some of the complexity
  - We can react to (most of the) requests with special needs

Basic Building Blocks: Instance Types

Name	vCPUs	RAM [GB]	Ephemeral [GB]
m2.small	1	1.875	10
m2.medium	2	3.750	20
m2.large	4	7.5	40
m2.xlarge	8	15	80
m2.2xlarge	16	30	160
m2.3xlarge	32	60	320

~80 flavors in total (swap, additional/large ephemerals, core-to-RAM, …)

Basic Building Blocks: Volume Types

Name	IOPS	b/w [MB/s]	feature	Location	Backend
standard	100	80		GVA	ceph
io1	500	120	fast	GVA	ceph
cp1	100	80	critical	GVA	ceph
cpio1	500	120	critical/fast	GVA	ceph
cp2	n.a.	120	critical/Windows	GVA	NetApp
wig-cp1	100	80	critical	WIG	ceph
wig-cpio1	500	120	critical/fast	WIG	ceph

m2.* flavor family plus volumes as basic building blocks for services

Automate provisioning with RUNDECK

Automate routine procedures

- Common place for workflows
- Clean web interface
- Scheduled jobs, cron-style
- Traceability and auditing
- Fine-grained access control
- …

Procedures for

- OpenStack project creation
- OpenStack quota changes
- Notifications of VM owners
- Usage and health reports
- …

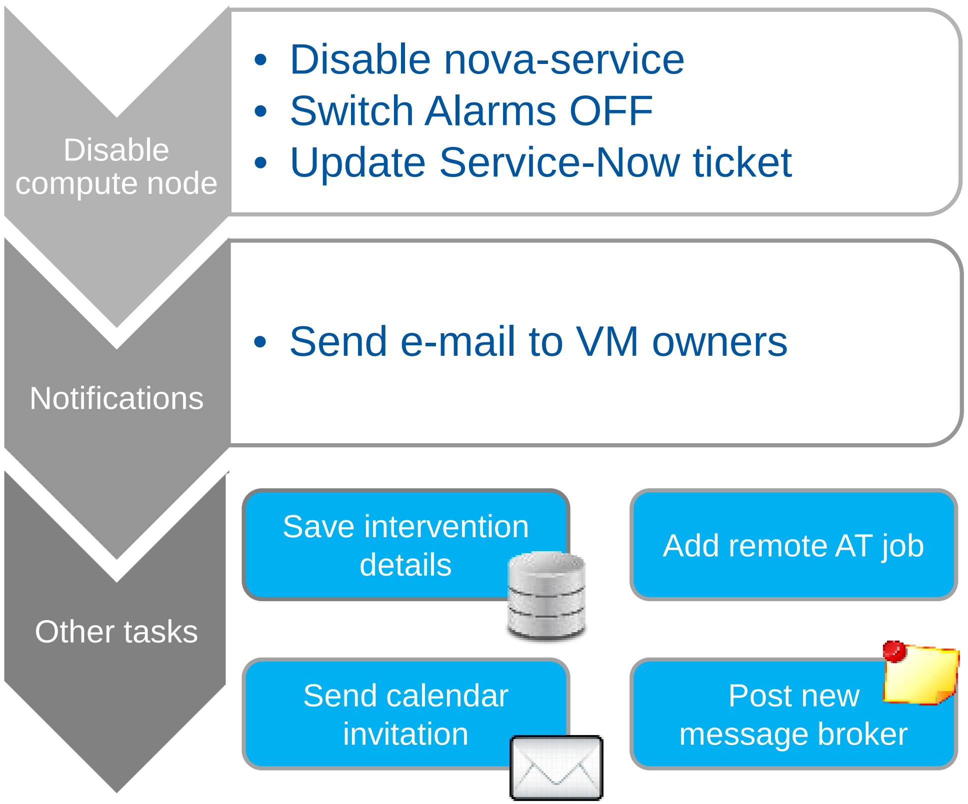

Operations: Retirement Campaign

- About 1’600 nodes to retire from the service by 3Q16
  - ~1’200 from compute, ~400 with services (hosting ~5000 VMs)
- We have gained quite some experience with (manual) migration
  - Live where possible and cold where necessary
  - Works reliably (where it can)
- We have developed a tool that you can instruct to drain hypervisor (or simply migrate given VMs)
  - Main tasks are VM classification and progress monitoring
  - The nova scheduler will pick the target (nova patch)
- We are using the “IP service bridging” to handle CERN network specifics

Operations: Network Migration

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

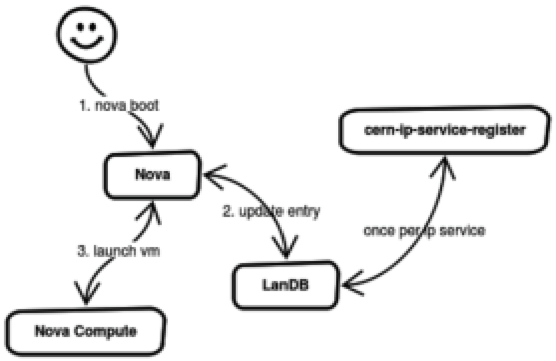

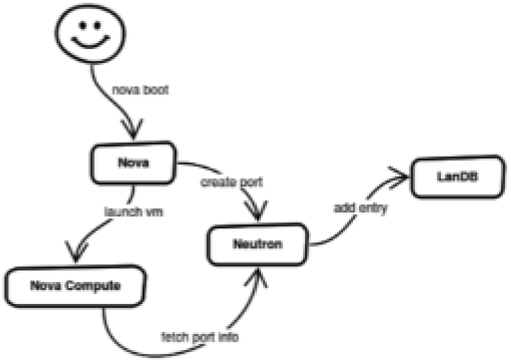

- New potential features (Tenant networks, LBaaS, Floating IPs, …)
- We have a number of patches to adapt to the CERN network constraints
  - We patched nova for each release …
  - … neutron allows for out-of-tree plugins!

Operations: Neutron Status

- We have ~5 Neutron cells in production
  - Neutron control plane in Liberty (fully HA)
  - Bridge agent in Kilo (nova)

- And it is very stable
  - All new cells will be Neutron cells

- “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, …)

Operations: Containers

- Magnum: OpenStack project to treat Container Orchestration Engines (COEs) as 1st class resources

- Pre-production service available
  - Support for Docker Swarm, Kubernetes, Mesos

- Many users interested, usage ramping up
  - GitLab CI, Jupyter/Swan, FTS, …

Operations: Federated Clouds

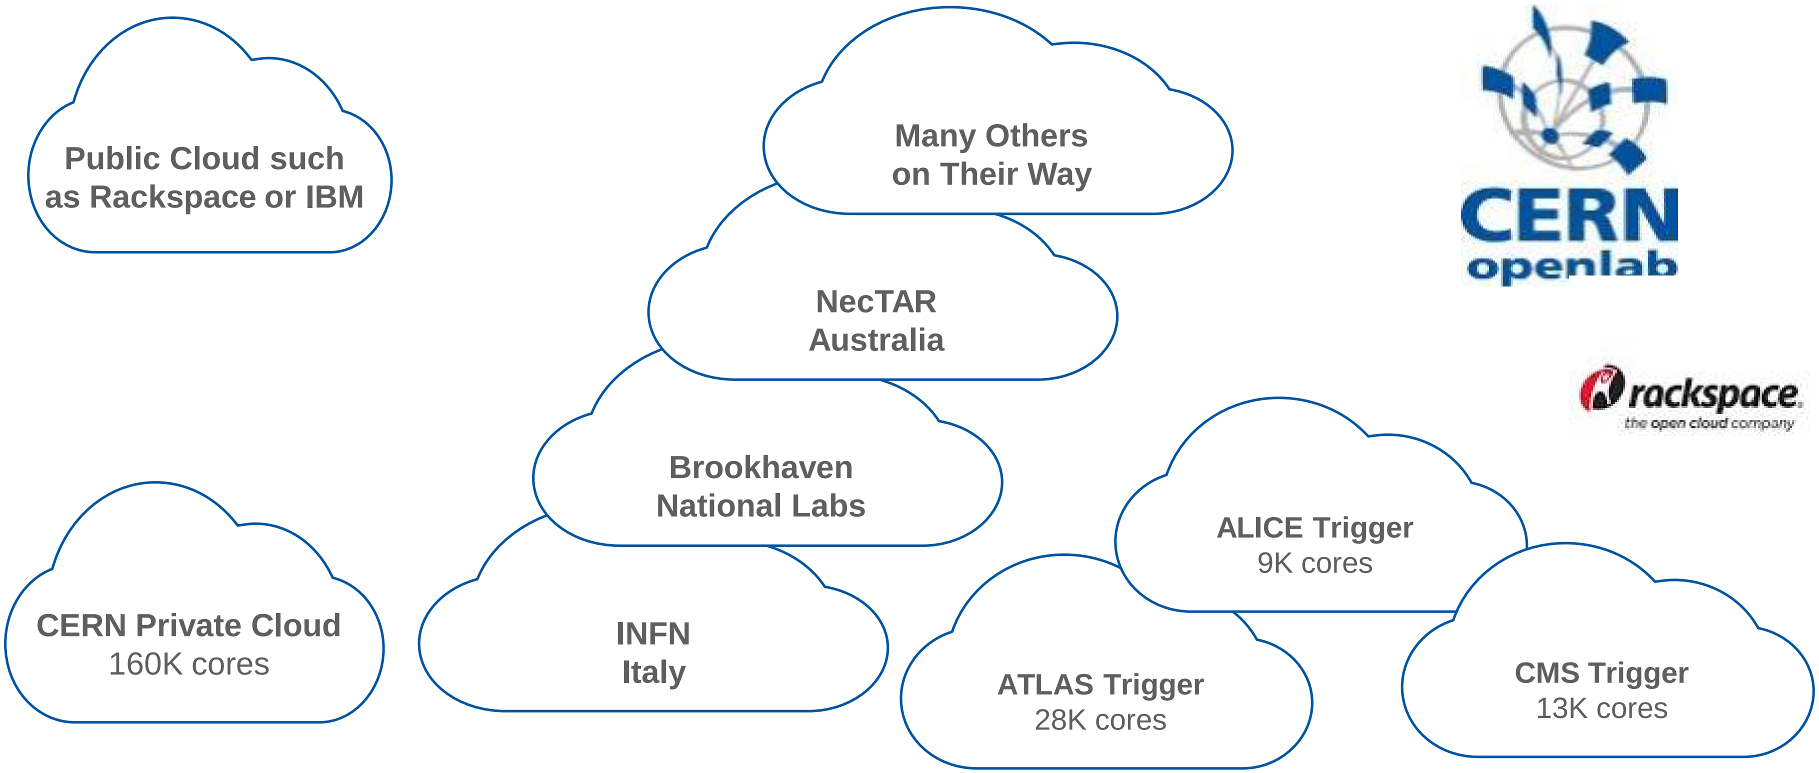

- Access to EduGAIN users via Horizon
  - Allow (limited) access given appropriate membership

CPU Performance Issues

- The benchmarks on 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
  - Discrepancy down to ~10% between virtual and physical
- Comparison with Hyper-V: no general issue
  - Loss w/o tuning ~3% (full-node), <1% for small VMs
  - … NUMA-awareness!

*KSM on/off: beware of memory reclaim! **EPT on/off: beware of expensive page table walks!

CPU: NUMA

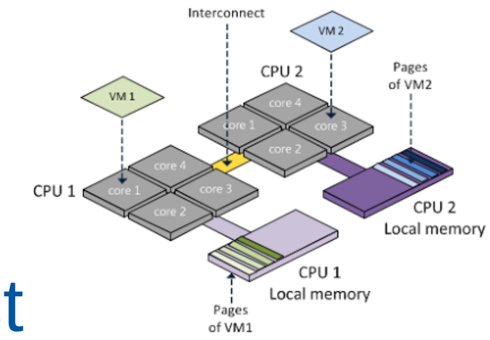

- NUMA-awareness identified as most efficient setting
  - Full node VMs have ~3% overhead in HS06

- “EPT-off” side-effect
  - Small number of hosts, but very visible there

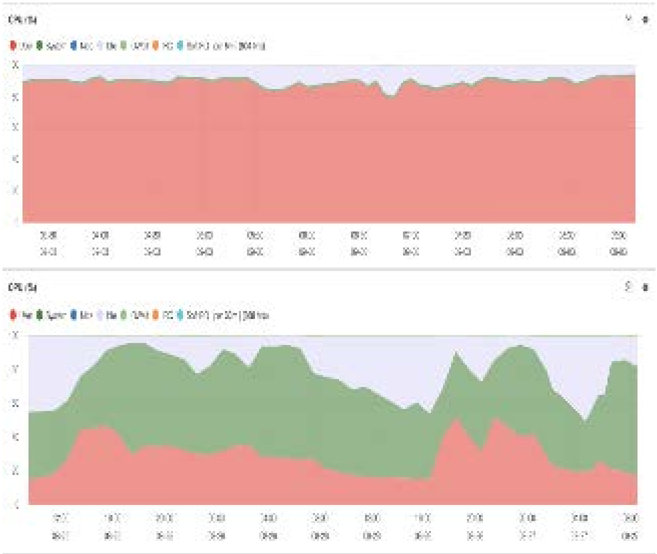

- Use 2MB Huge Pages
  - Keep the “EPT off” performance gain with “EPT on”

- More details in this talk

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

Future Plans

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

- Replace Hyper-V by qemu/kvm
  - Windows expertise is a scarce resource in our team
  - Reduce complexity in service setup

Summary

- OpenStack at CERN in production since 3 years
  - We’re working closely with the various communities
  - OpenStack, Ceph, RDO, Puppet, …
- Cloud service continues to grow and mature
  - While experimental, good experience with Nova cells for scaling
  - Experience gained helps with general resource provisioning
  - New features added (containers, identity federation)
  - Expansion planned (bare metal provisioning)
- Confronting some major operational challenges
  - Transparent retirement of service hosts
  - Replacement of network layer
- http://openstack-in-production.blogspot.com
  (read about our recent 2M req/s Magnum & Kubernetes!)

CERN