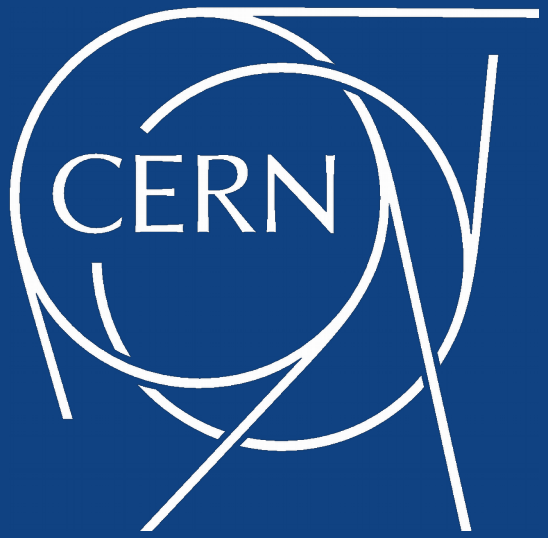
CERN


CERN

OpenStack @ CERN: Status update

Spyros Trigazis

CERN OpenStack Infrastructure

Production since 2013
> 6000 hypervisors ~ 190.000 cores ~4 million vms created ~200 vms per hour

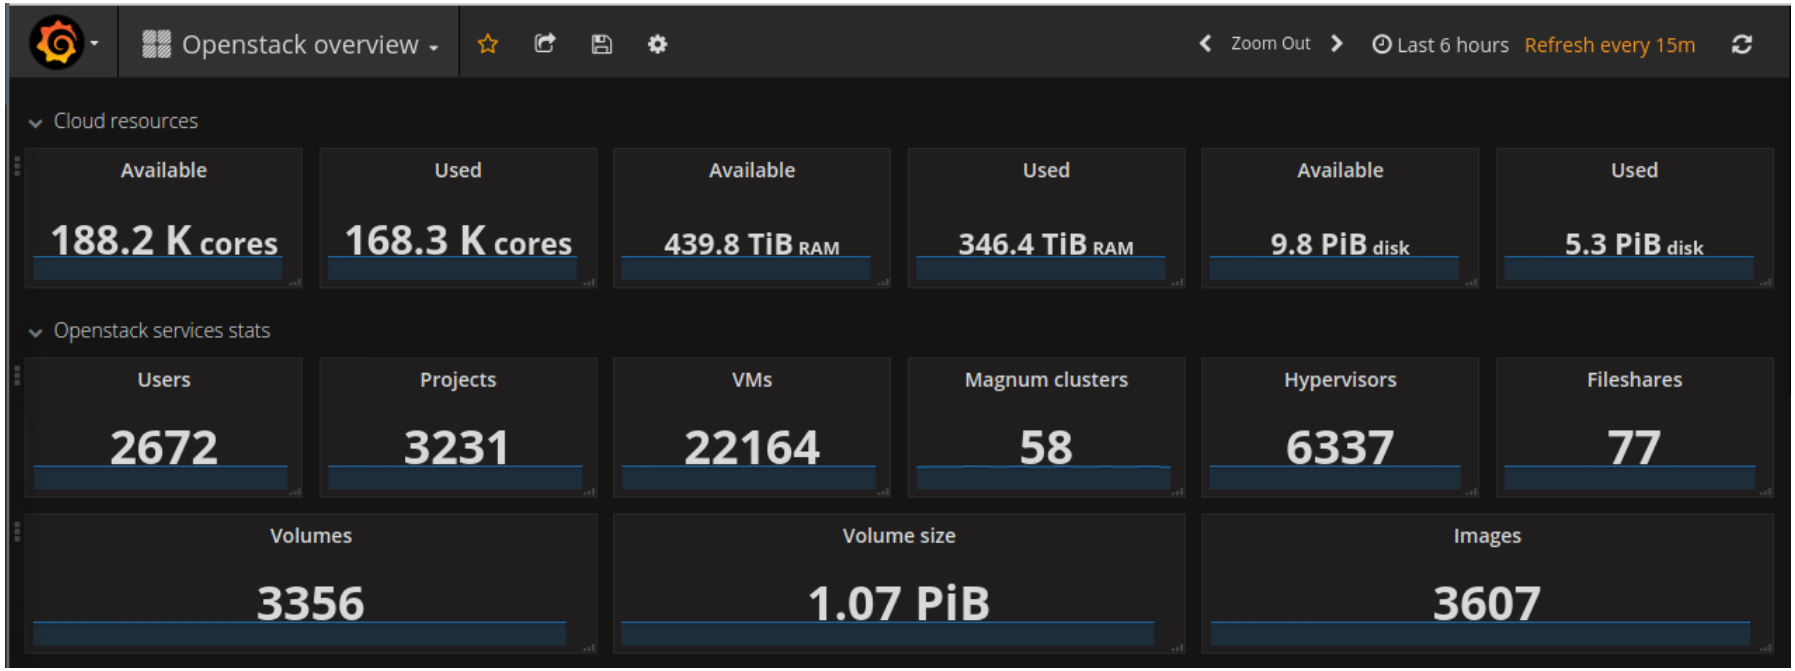

OpenStack@CERN Status

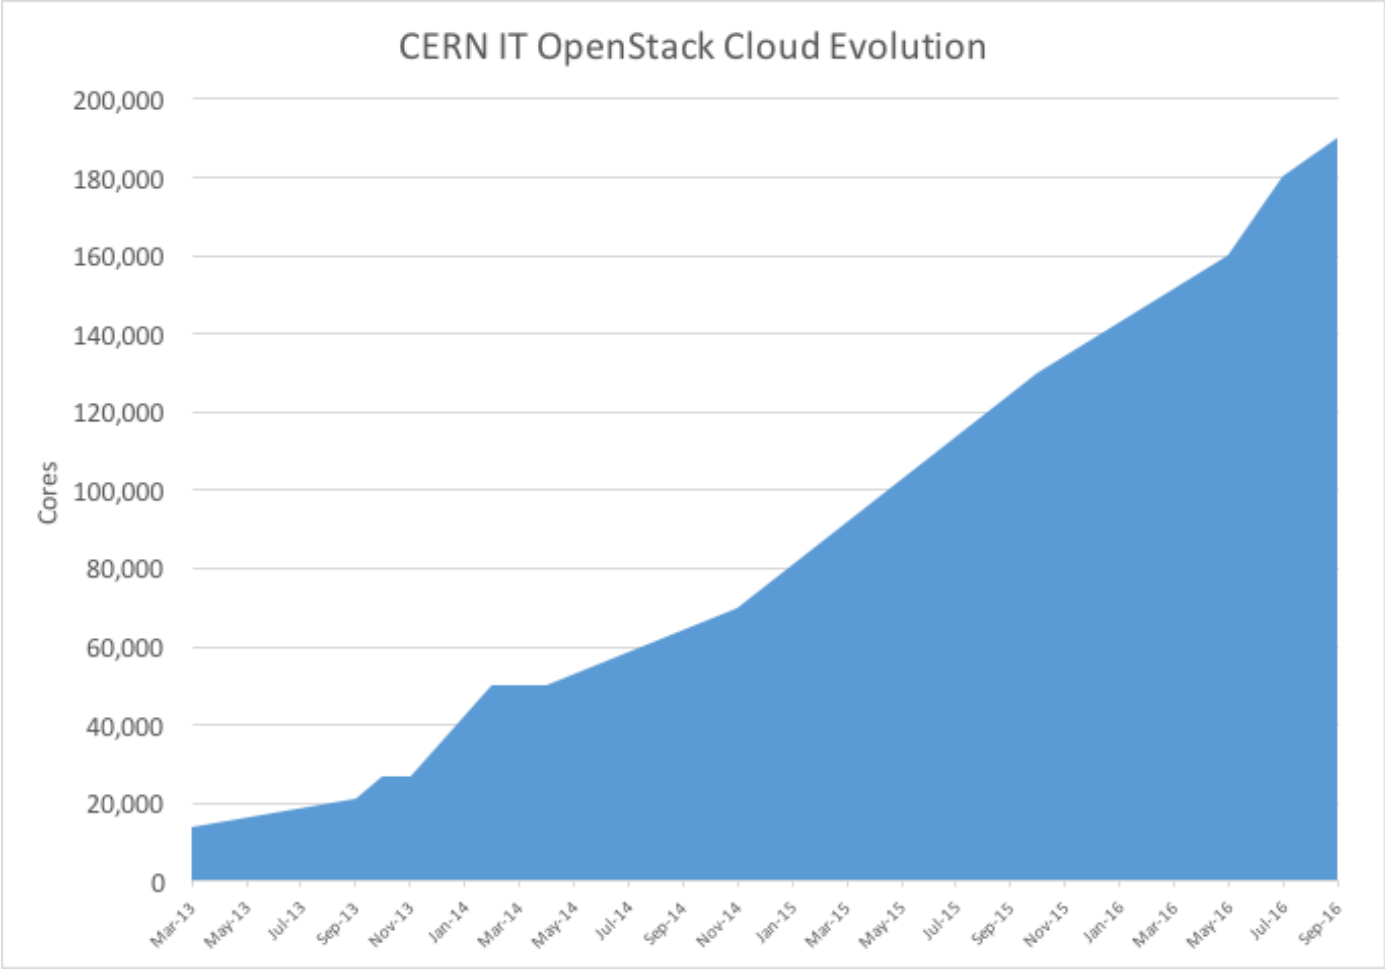

In production:

- >190K cores
- >6000 hypervisors

~100,000 additional cores being installed in next 6 months

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

CERN OpenStack Project

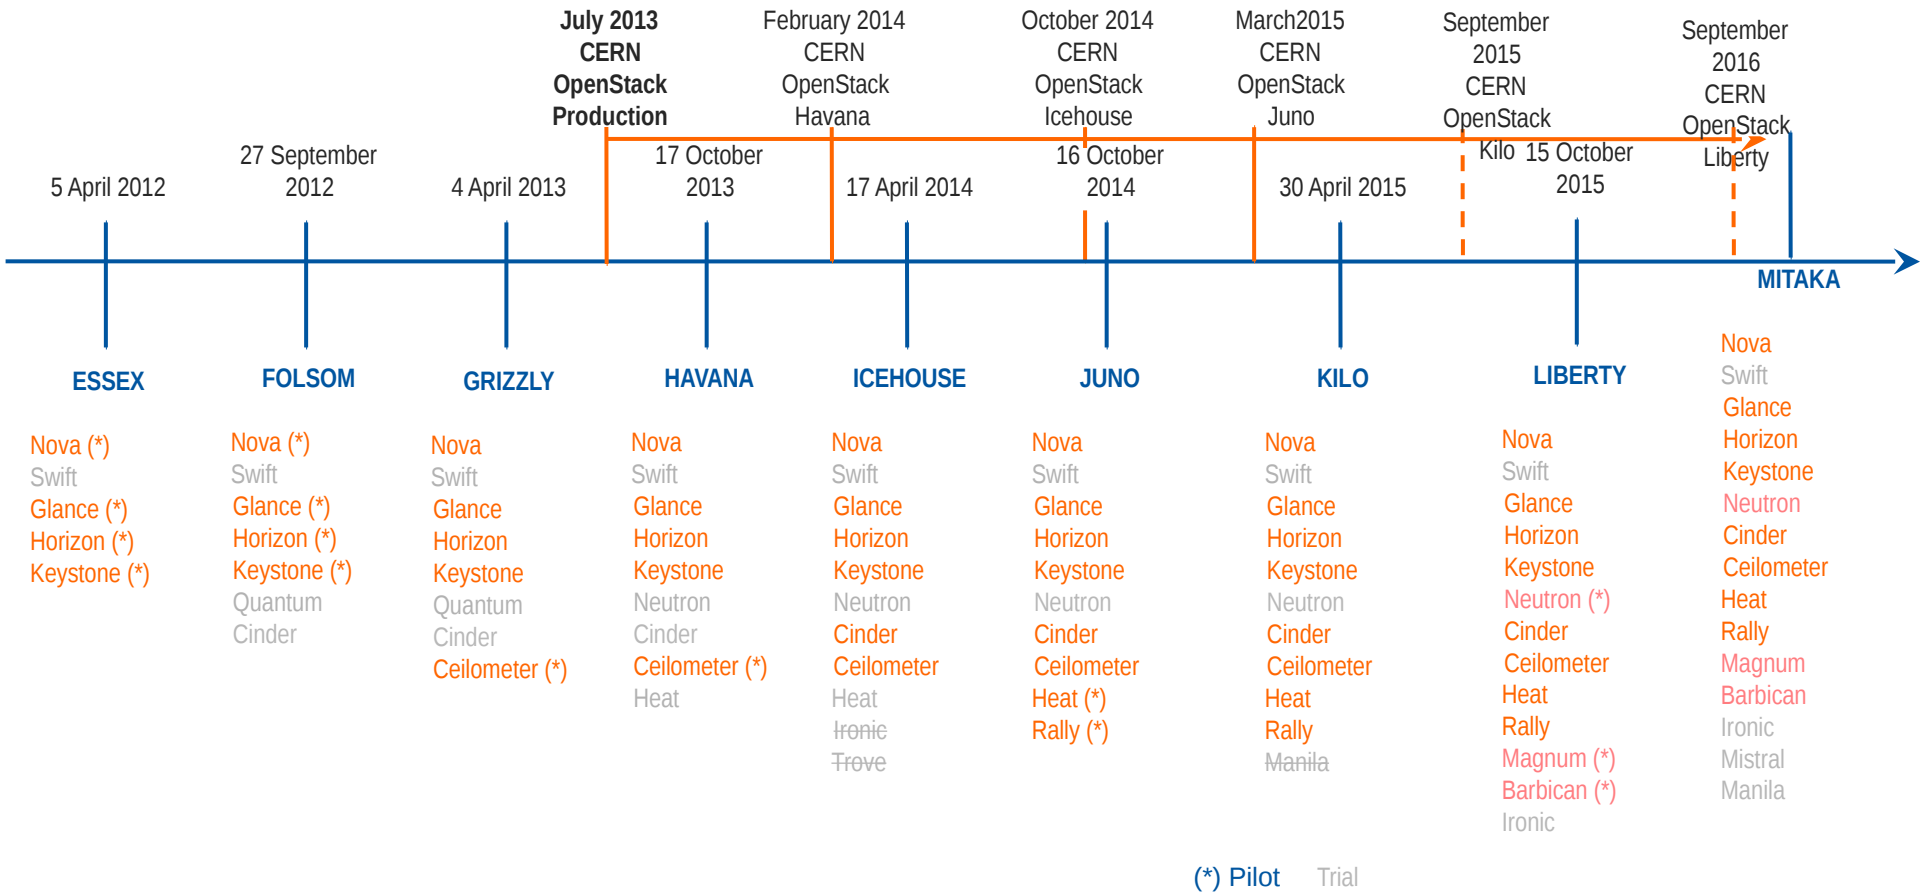

Compute Service - Nova

- Started in 2012 and moved to production in 2013
- Several Patches for nova-network
- Heterogeneous setup with KVM and Hyper-V
- Challenges
  - Upgrades
  - Migrate from SLC6 to Centos 7 (complete)
    - pre-req to upgrade to Liberty
  - Migrate to KVM only (on-going)
- Liberty release

Compute Service - Nova

Top level cell
- Runs API service
- Top cell scheduler

Child cells run
- Compute nodes
- Scheduler
- Conductor
- 40+ cells

Version 2 coming
- Default for all

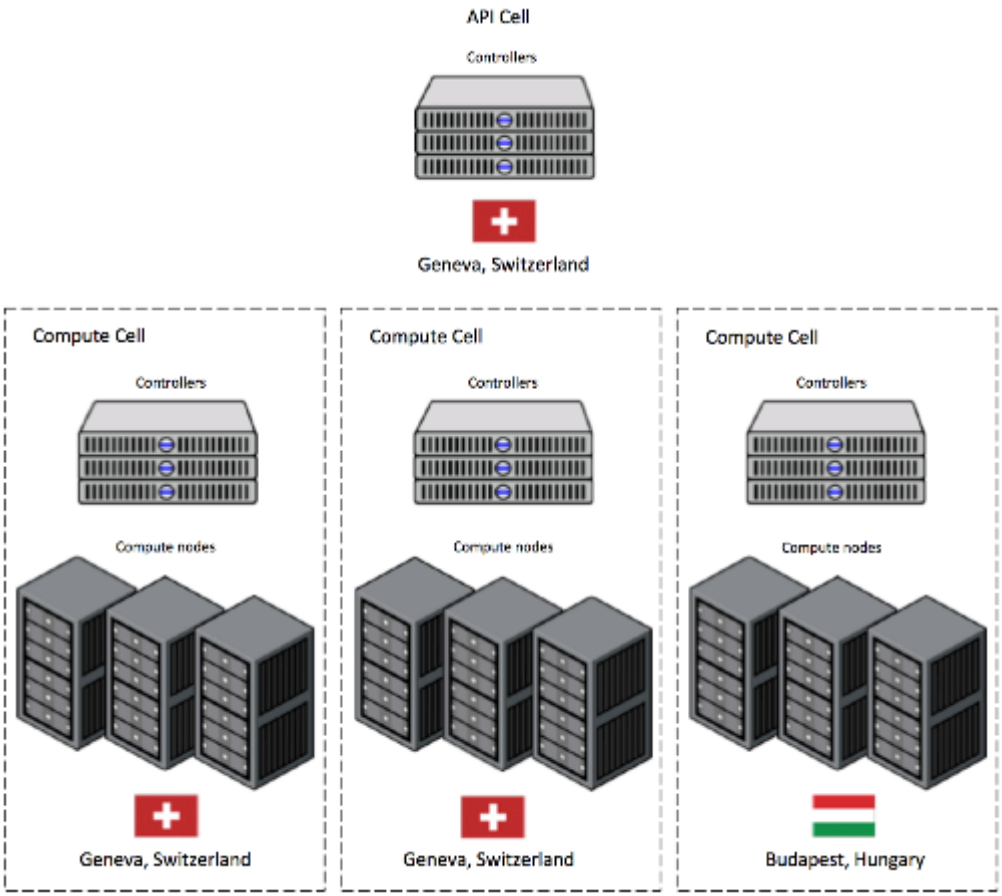

Networking Service - Neutron

- Started as a pilot service in Q3 2015
- Deployed only in one cell
- Deployed with a custom driver, based on the linux bridge driver, for integration with CERN's network database
- Migration from deprecated nova-network to Neutron networks
- Liberty release, depends on Nova

Block-Storage Service - Cinder and Ceph

- Block storage
  - Predominantly Ceph
  - Some NetApp
- Ceph is backend for Glance images
- CephFS with Manila is under investigation for user clusters
- Mitaka release

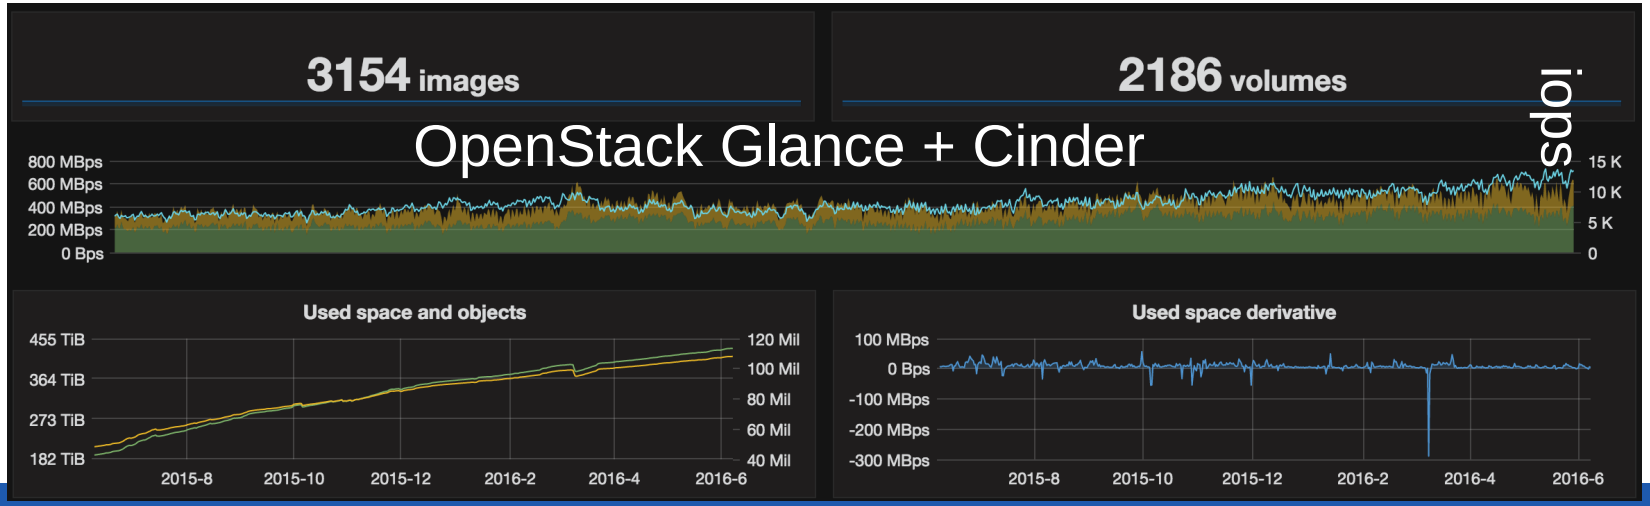

OpenStack Rally benchmarks

- Deployed at 2015
- Deploys m1.tiny VMs every hour in selected cells to check the status of all services
- Used to stress-test the orchestration service and consequently all other services
- Patched to not run as admin user
- newton release with cherry-picks

Container Infrastructure - Magnum

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

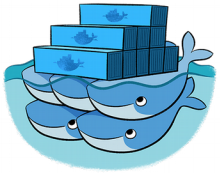

- Production service since Q4 2016
  - Support for Docker Swarm, Kubernetes, Mesos
  - Storage drivers for (CERN-specific) EOS, CVMFS

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

- Newton release, with cherry-picks

Upcoming Services

- Baremetal service Ironic
  - API server and Conductor already deployed
  - First node deployed this month
- Workflow service Mistral
  - Will simplify operations, create users, clean up resources
  - Deployed and testing prototype workflows
- FileShare service Manila
  - Pilot since Q4 2016
  - Share configuration, certificates, etc

Operations

Humans in the loop

- Multiple CERN teams
  - Openstack team: service managers and service developers
  - Procurement, hardware mgmt. teams
  - LinuxSoft, Ceph, DBoD teams
- “Upstream”
  - OpenStack superusers ☺
    - Development, large deployment, board representative
  - RDO
    - Packaging, testing

Operations

- User support
  - Tickets, documentation, etc.
- Deployment
  - RDO packages, with some customizations/selected patches
  - Puppetized configuration
- Upgrades
  - Component by component, twice a year
  - Operating systems
    - CentOS 7.2 -> 7.3 ongoing

Software Deployment

- Deployment based on CentOS and RDO
  - Upstream, only patched where necessary (e.g. nova/neutron for CERN networks)
  - 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

RDO

puppet labs

Custom Package management

- Production and QA tags
- Build with gitlab-ci and koji
  - VM builders and Docker builders
- Testing in:
  - a dedicated nova cell
  - in all-in-one VMs with devstack
  - cloud-dev environment on top of kubernetes

Package cluster drivers for Magnum

- Upstream challenges
  - Out of tree management
  - Manage common code between drivers (python and heat templates)
  - Integrate in openstack ci
  - development overhead
- CERN specific drivers
  - already packages but taken from in tree
  - qa tag, tag as production when we are happy with it

Summary

- OpenStack at CERN in production since 3.5 years
  - We're working closely with the various communities
  - OpenStack, RDO, Ceph, 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

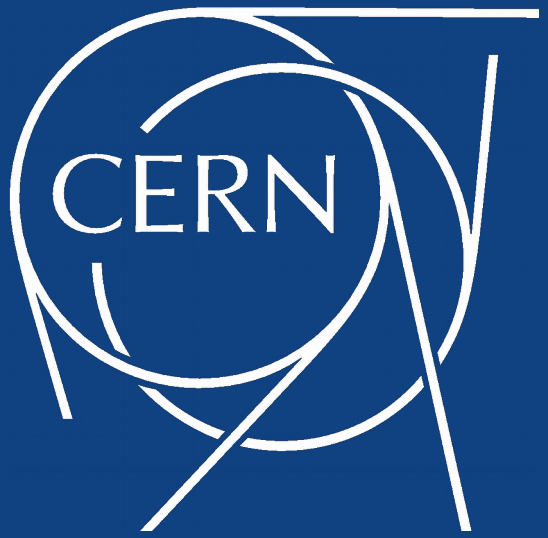
CERN


CERN