CERN

Tier-0 status and plans

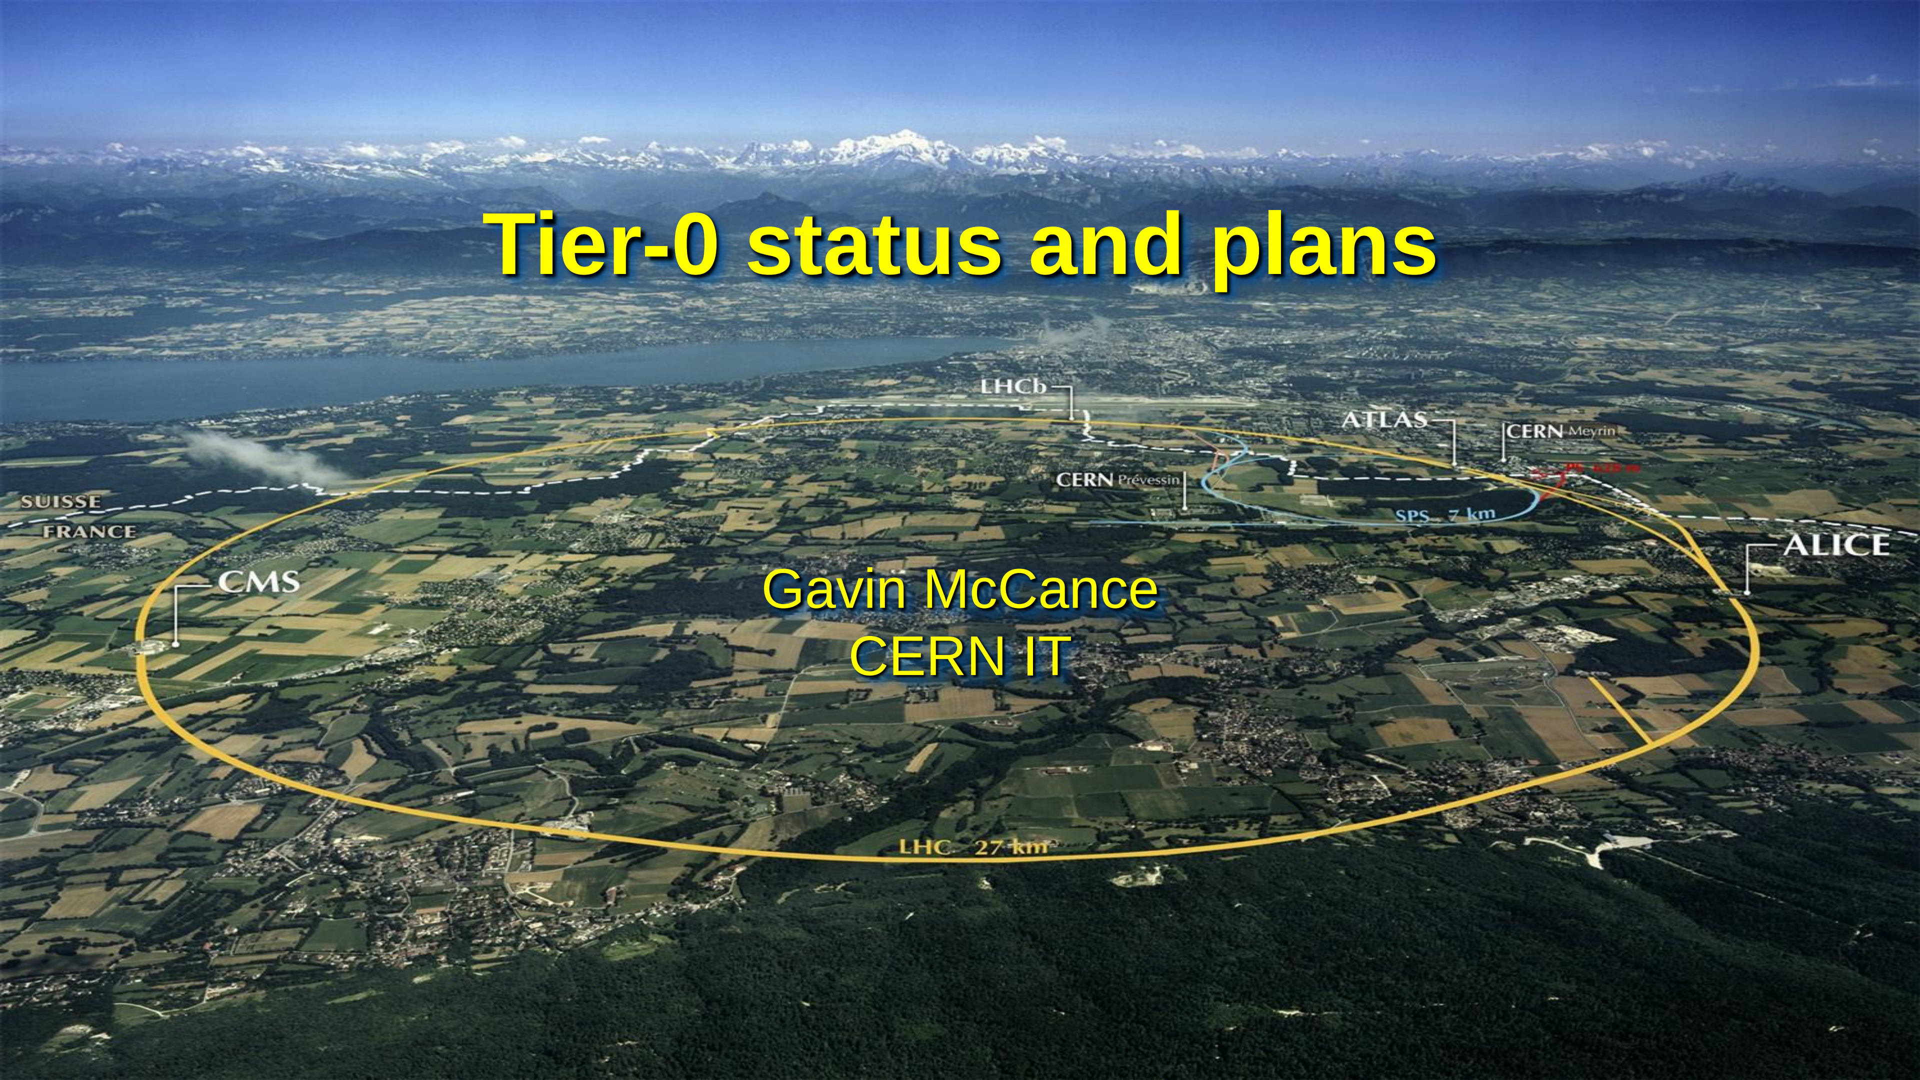

Gavin McCance
CERN IT

Outline

- Cloud and batch
- Opportunistic, mid-SLA, external cloud
- New computer centre
- Storage and databases

OpenStack@CERN Status

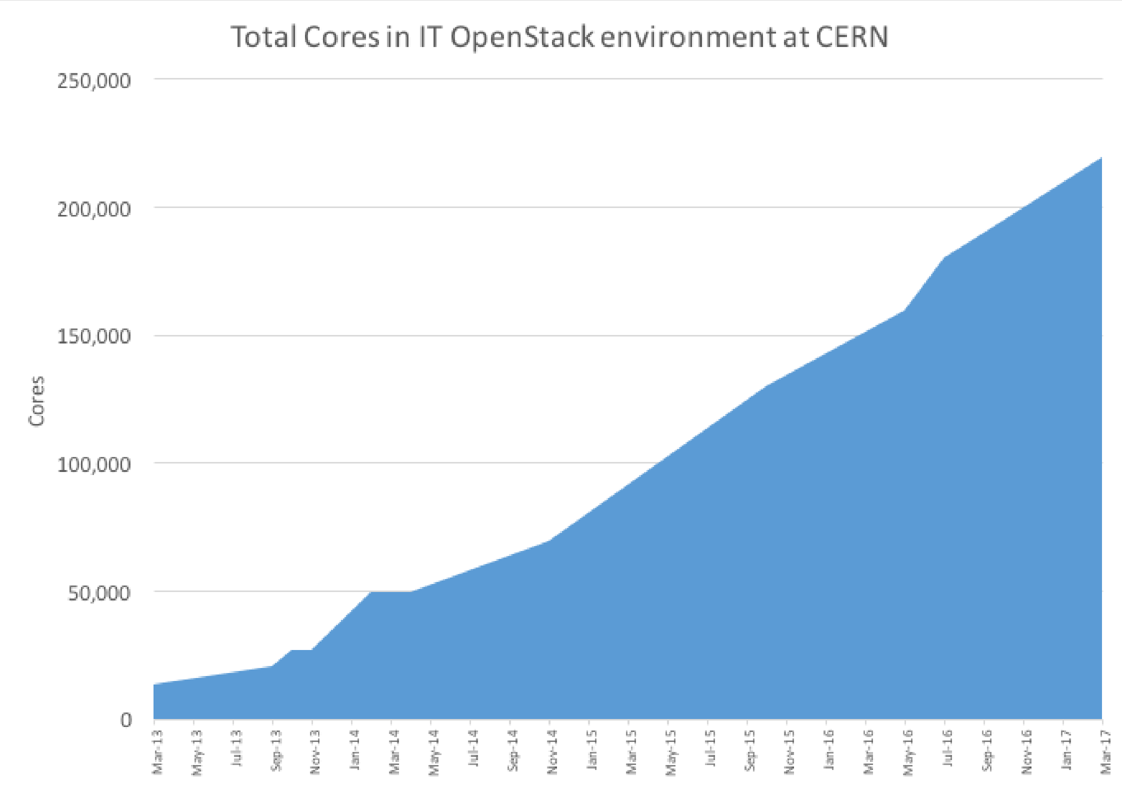

In production:

- >220K cores
- >7000 hypervisors

~86,000 additional cores being installed in next 6 months

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

Service / Compute split

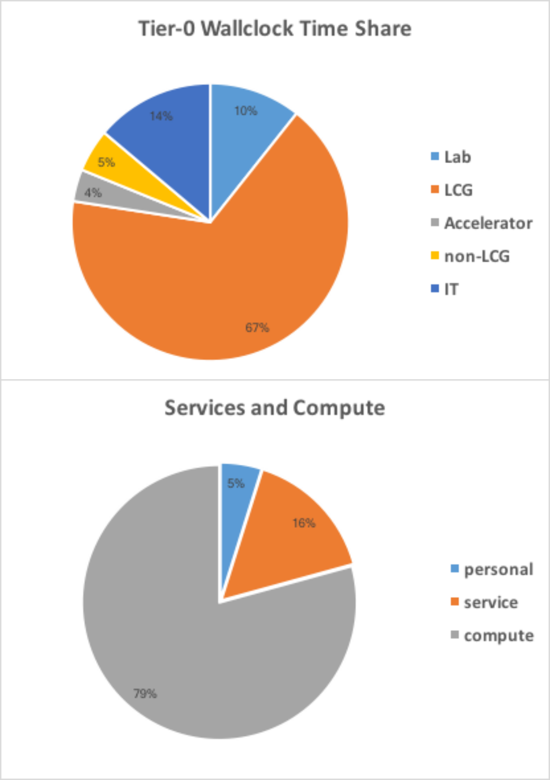

- Single resource pool for the lab infrastructure and the physics
- Majority of the resources are allocated for compute workloads
- However, a significant share of resources are used for “services” either IT or experiment

Openstack

- Still very happy with Openstack after 3 years
- Upgrades ~smooth and routine now
- Scaling out further soon > 280k
- Functional increments:
  - Live migrations
  - Container management (Magnum)
  - Bare-metal management (Ironic)
  - SDN (Neutron)

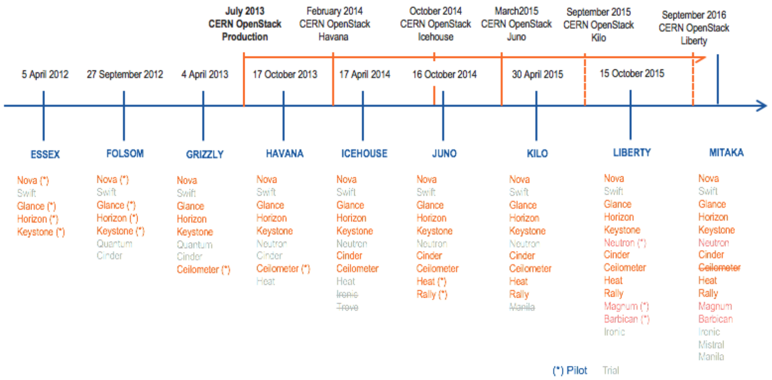

- We maintain involvement with upstream and actively contribute back to the project

Gavin.McCance@cern.ch

Filer: Manilla and CephFS

- We're evaluating CephFS with Openstack Manilla for the replacement of the NFS Filer service (currently Netapp)
  - Horizontally scalable on Ceph
  - Tunable IOPs
  - Easy to attach/detach as needed
  - Looking very positive…

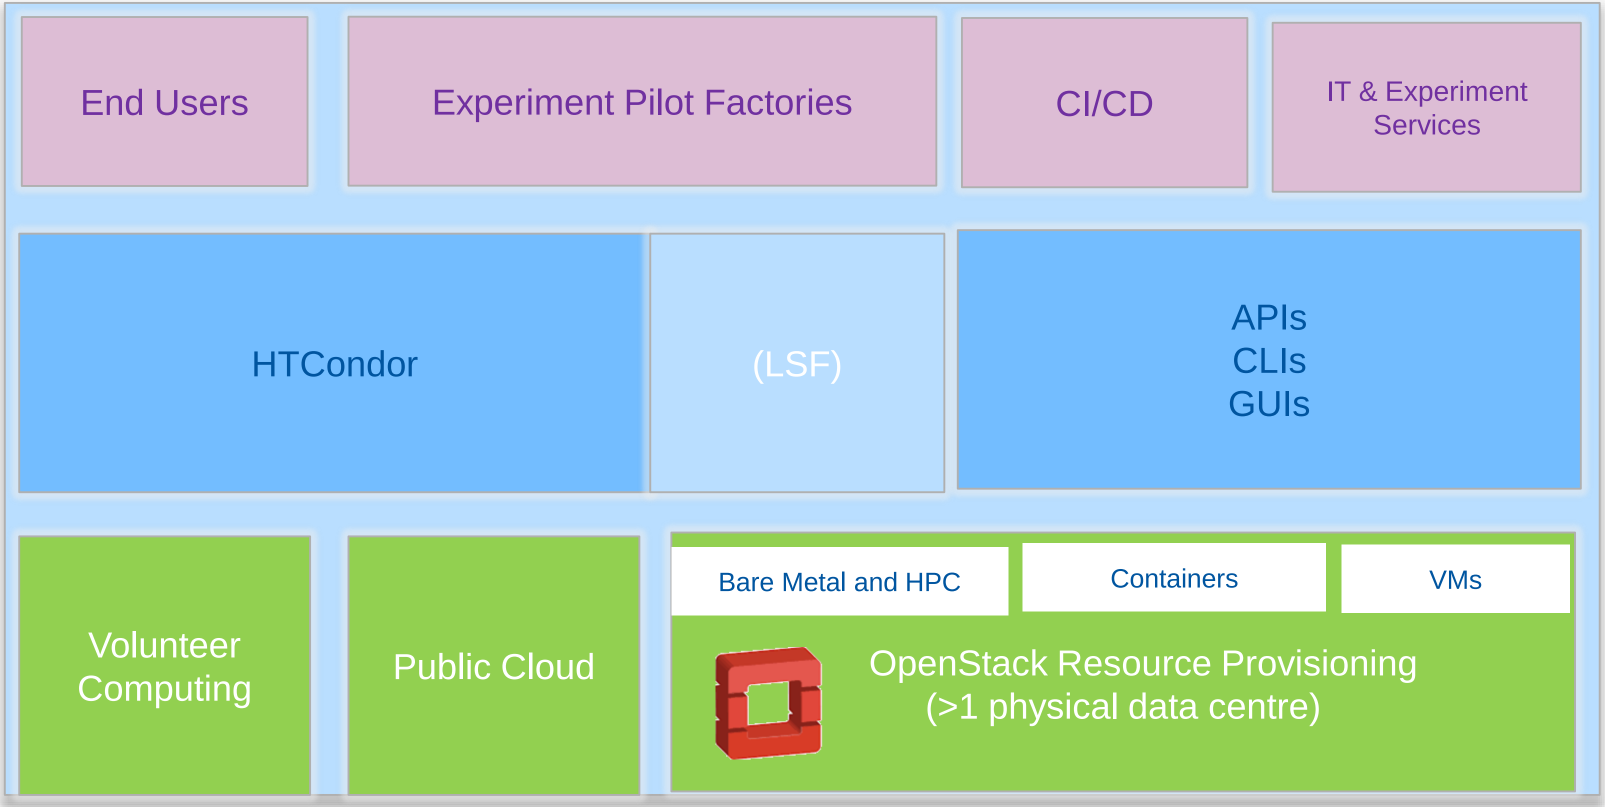
End Users
Experiment Pilot Factories
CI/CD
IT & Experiment Services
HTCondor
(LSF)
APIs
CLIs
GUIs
Volunteer Computing
Public Cloud
Bare Metal and HPC
Containers
VMs
OpenStack Resource Provisioning
(>1 physical data centre)

HTC Batch service

- Users:
  - LHC: Our share of WLCG Grid quota
  - LHC: Prompt Tier-0 calibration / hot events
  - LHC: Tier-0 bulk reconstruction
  - All other CERN experiments (e.g. Compass)
  - Local CERN departments (EP groups, Theory, Beams and Engineering)
  - Associated experiments (e.g. AMS)

  - About 50% Grid, 50% local

HTC Batch service

- Currently we have around 130k cores
  - Around 650k jobs per day
  - Split 50/50 LSF and HTCondor
  - On track for migration of major users and all LHC by Q3 this year

- Scaling to over 200k cores in batch by end of 2017
- 2018 – 2020 … funding is being discussed…

Batch deployment model

- Vast majority deployed as long-lived VMs on Openstack using HTCondor vanilla universe
- For HTCondor we're aiming for uniform 8-core node standard workers, 30-50% as multicore

- Configuration maintained with Puppet
- Normal drain/reboot life-cycle
  - but if it fails, shoot it and get another one

- Small high-memory (~1TB) facility to be provided this year for special cases (CAF merging, theory, electronics, …)

Containers

- Plans for Batch Service:
  - We'll soon deploy Singularity for experiments (pilot isolation)
  - We'll (likely) roll-out HTCondor Docker universe for job isolation
    - Condor-managed containers
    - Standard SLC6/CC7 image runners
    - …and choose-your-own-image
    - CVMFS / EOS mounts, no AFS
  - Once AFS is gone, we'd like this to be our default for the future

- N.B. Also containers for services:
  - Also now running Openstack Magnum with DCOS/Kubernetes for better deployment of our *services*)
  - Need to work out reliable way of configuring the container's contents

HTCondor experience

- Very good experience, very good support
- Migrating all CEs to HTCondorCE
- “Free” monitoring from Fifemon (FNAL)
- Contribute where we can
  - e.g. improvements to “haggis” share management tool
- User migration fairly easy
  - Same pattern
  - Consultancy.. some wrinkles.. user limits, etc.
  - Hope to be done by end of Run 2

HPC facilities

- Apps that don’t fit normal HTC pattern
- MPI, shared memory across nodes, infiniband
  - Lattice QCD Theory simulations, Beam / plasma, fluid dynamics applications (fire safety, cryo), engineering simulations (civil and electronic)

- New theory cluster now working well (72 nodes)
- Beams clusters (2*72) – funding being discussed for Q4 ‘17

- SLURM batch system being deployed for this (~5k cores)
  - ...will backfill via HTCondor / SLURM interface

Volunteer

- Significant resources for beam simulation and LHC Monte Carlo, now integrated via HTCondor

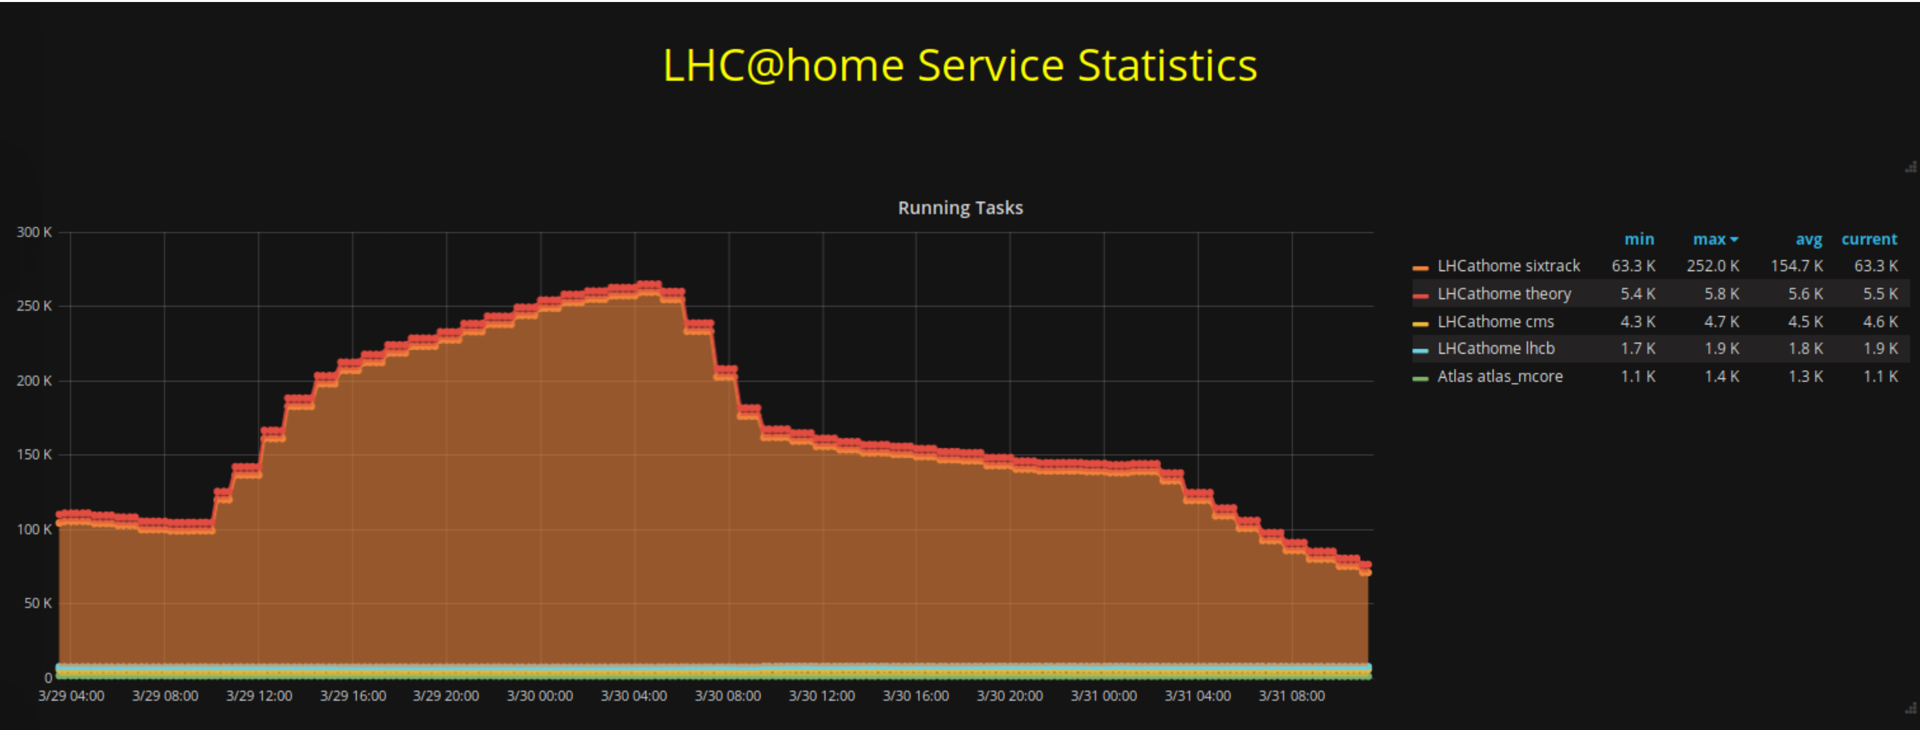

Making better use

- Identifying how to exploit poorly used resources
- Significant potential extra capacity @CERN, if we are able to relax the normal HEP SLA a bit

- Examples:
  - disk-server CPU
  - spare service “headroom” on cloud, choppy cloud compute capacity, external cloud spot
  - HPC backfill, pre-empt by prompt work (Tier-0/CAF),
  - intervention draining

Mid-SLA

- Standard-kit, access to local storage
- Potentially resource squeezed (notably IOPS)
- No notice termination if the real owner wants it back
  - Possible batch-system pre-empt in some cases

- Will be exploring over next ~year with experiments how to make better use of these

Commercial clouds

- Various projects and exercises with experiments to understand how we could use them
  - HNSciCloud currently

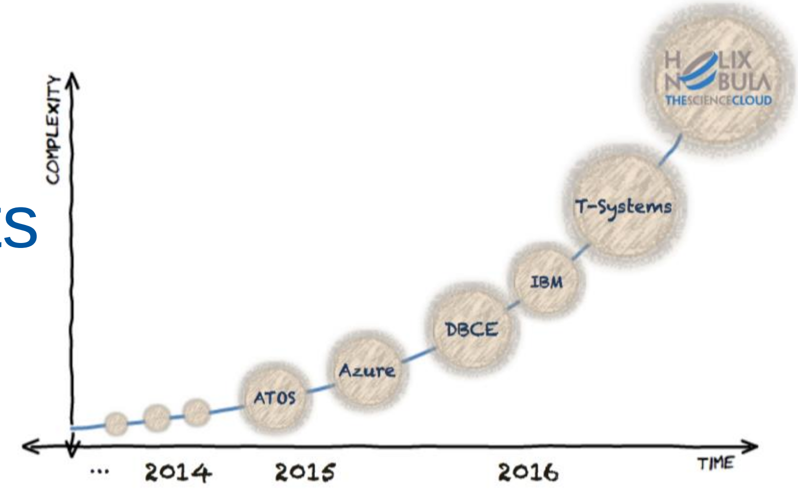

- Strategy:
  - Configure them as we do here (same Puppet)
  - Flat capacity, long-lived worker nodes exposed via the CERN HTCondor pool (special route)

Commercial clouds

- Most workflows tried
- Data is hard with current HEPpy protocols
    - Would need some evolution of experiment data models
- Can do reco from CERN, but heavy on network
- Strategy for mid-term, “cheap and cheerful”:
    - Keep the capability and be able to spin up on any cloud with no extra effort using our standard tools (Terraform + Puppet)
    - Procure for standard tools support
    - Focus on simulation when we do
    - Related to mid-SLA, explore with experiments “spot” models

Strategy

Consolidate HTC around HTCondor interface

- Local
- Public cloud
- Volunteer
- Opportunistic
- Backfill

SLURM for HPC

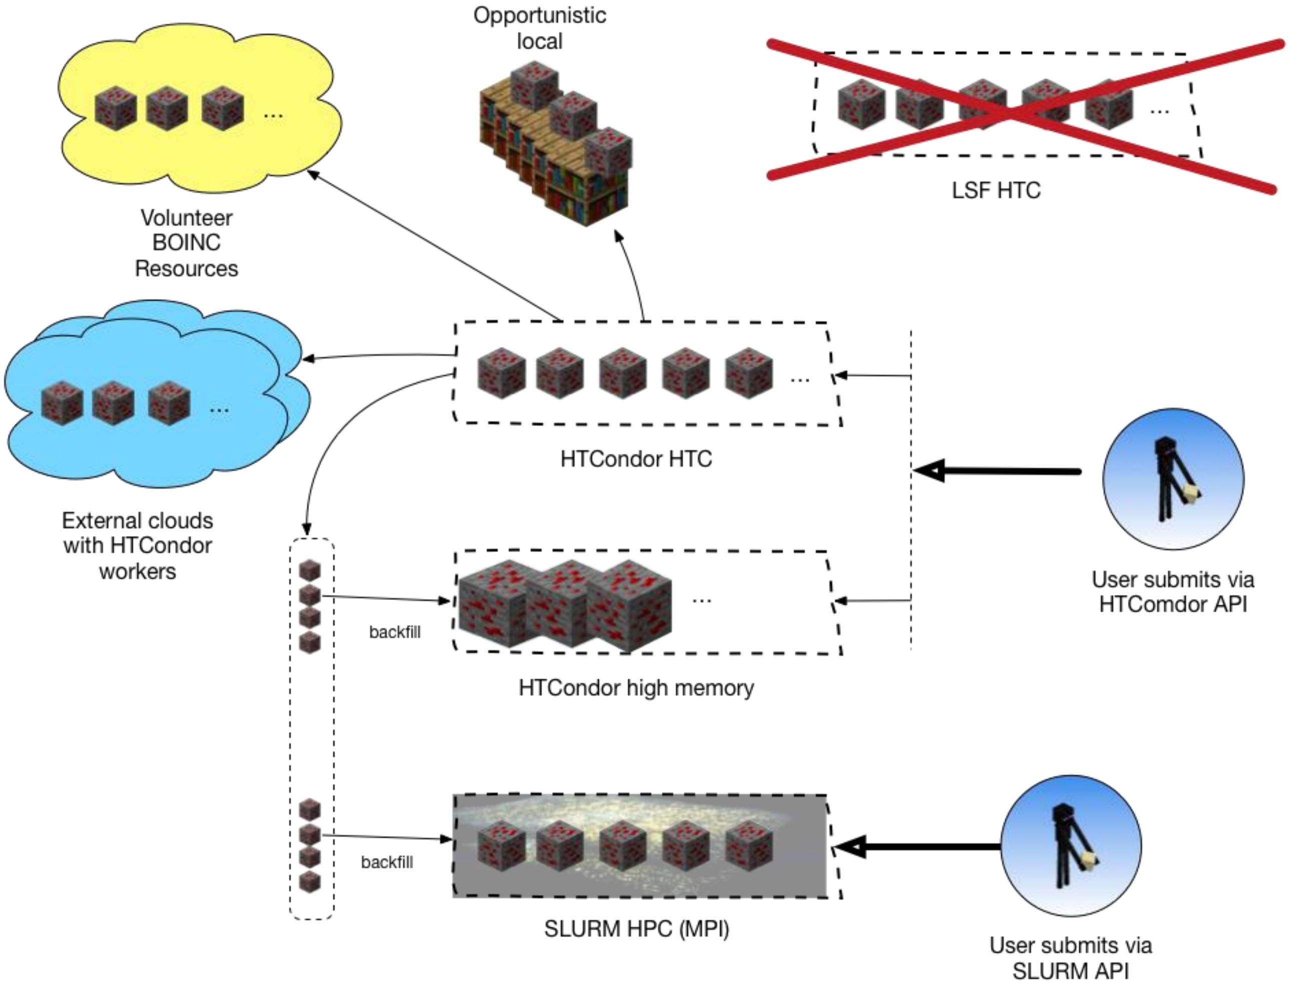

New computer centre – why?

- LHCb and ALICE will be moving to SW triggers for RUN3 -> HLT farms need more capacity
  - Hard to accommodate this capacity at the pits
  - Probably also for ATLAS and CMS for RUN4
- Increased Tier-0 capacity will be required for RUN3 and beyond
- Other large computing needs are coming up, e.g. ProtoDune
- Commercial cloud computing not yet attractive for bulk

New computer centre

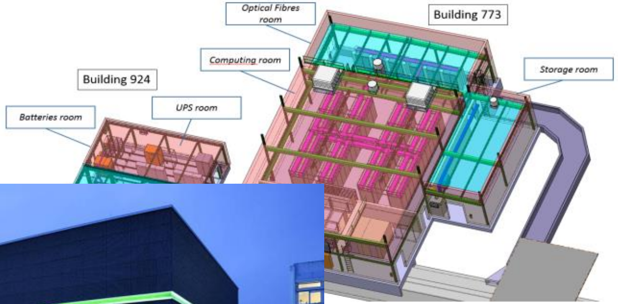

- GreenITCube @GSI
  - ...starting point for design call
  - Excellent PUE
  - Quick to build
- Current status:
  - ED will decide formally in May if we go ahead
    - In parallel, tenders out – review in June FC
    - In parallel, CERN groups studying potential computing and network architectures (e.g. links to expt. areas)
  - Contract signed summer (assuming approval)

Networking: LHCONE plans

- Continue to increase capacity of links
  - To handle LHC data growth
  - 100 Gbit T0-T1 where cost effective
- Expanding e.g. Asian links
- Continue deployment of perfSonar monitoring

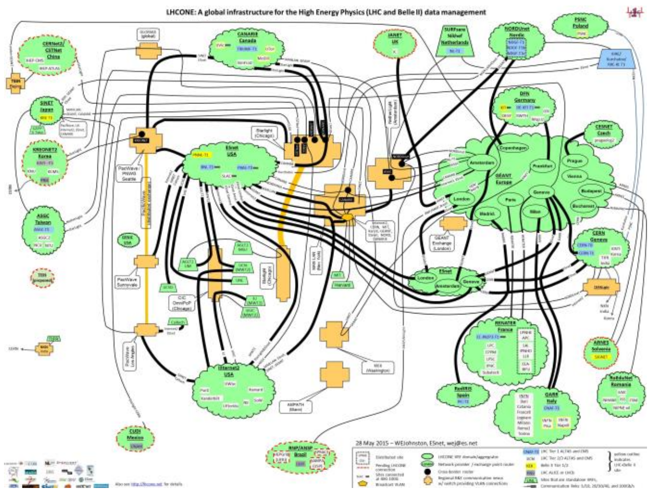

Storage stats

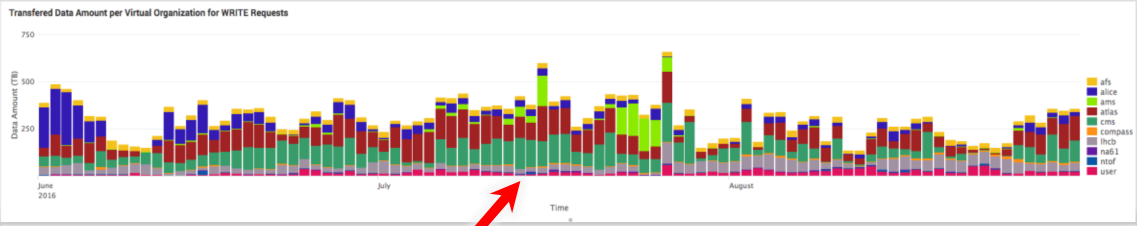

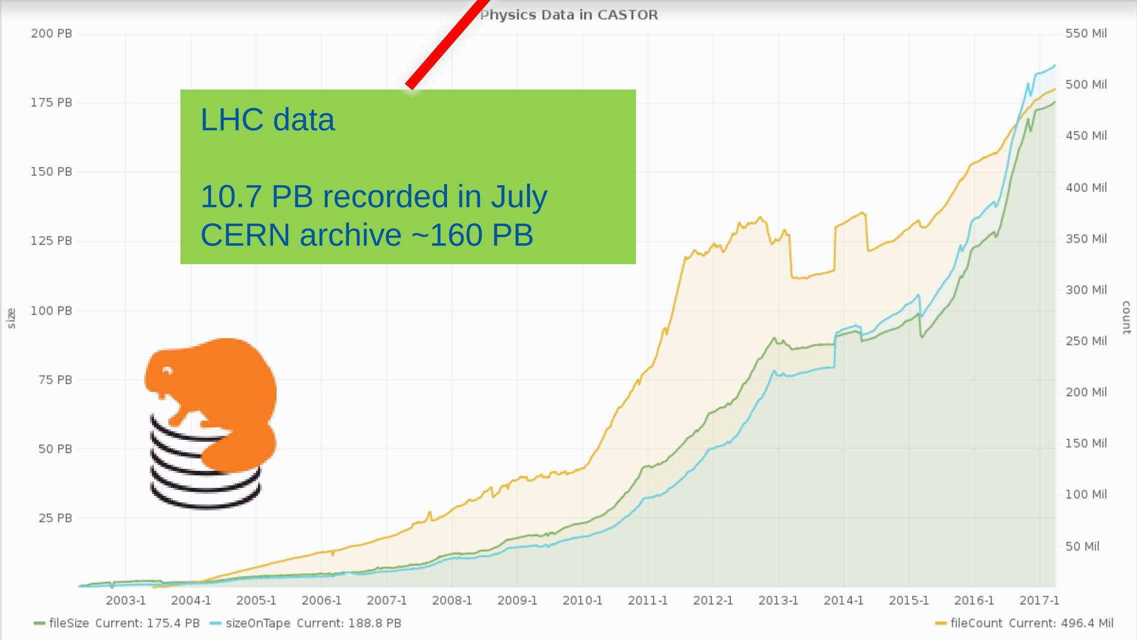

LHC data

10.7 PB recorded in July
CERN archive ~160 PB

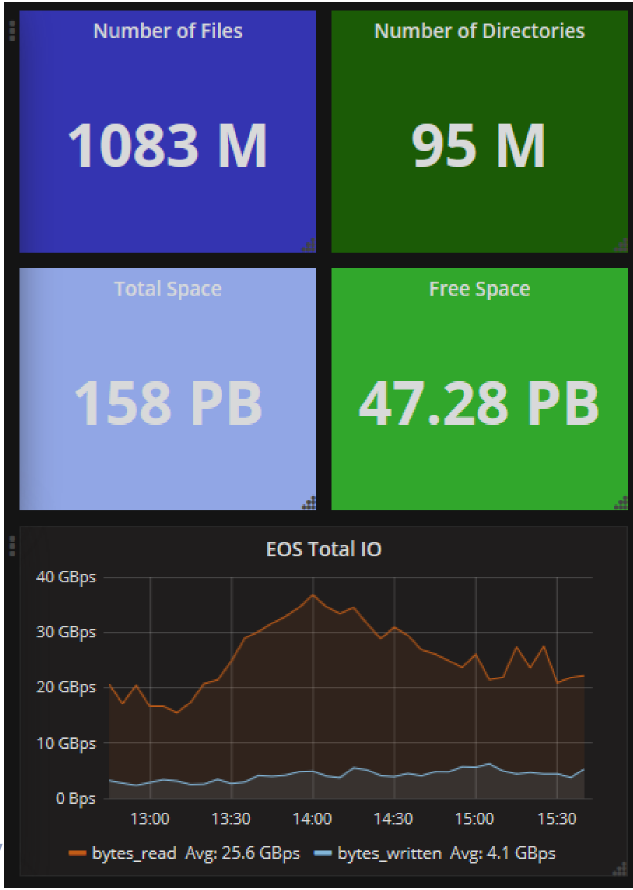

Gav

EOS and CERNBox

- Consolidation on EOS for storage
  - Designed for very high performance open/read; low latency; tunable QoS
  - JBOD commodity hardware, ignore failed disks
  - Expanding range of protocols
  - Future: be able to take advantage of whatever disk technology brings
- CERNBox provides cloud sync (à la Dropbox)
  - Synchronise EOS files (data at CERN) and offline data access

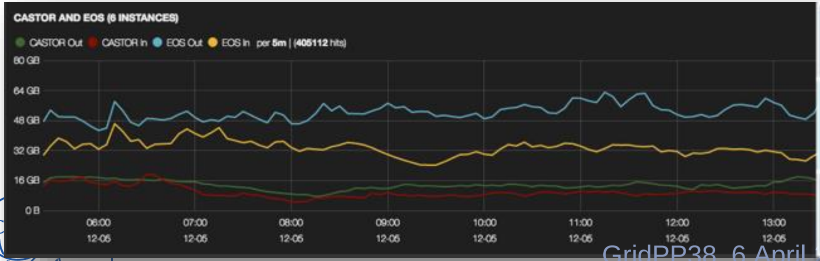

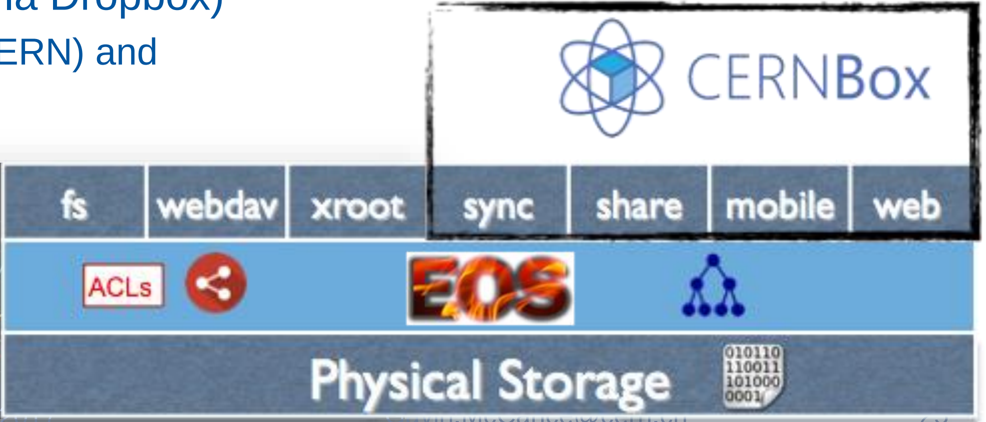

AFS replacement

- Driven by gradual demise of upstream
- CERN planning to phase out gradually
  - Large migration effort – aiming to be done during LS2

- Plan to consolidate on EOS for remaining use-cases
  - Significant steps to migrate already taken
  - Some developments ongoing inside EOS to support harder use-cases (e.g. homedirs)

SWAN Architecture –
Data Analysis as a Service

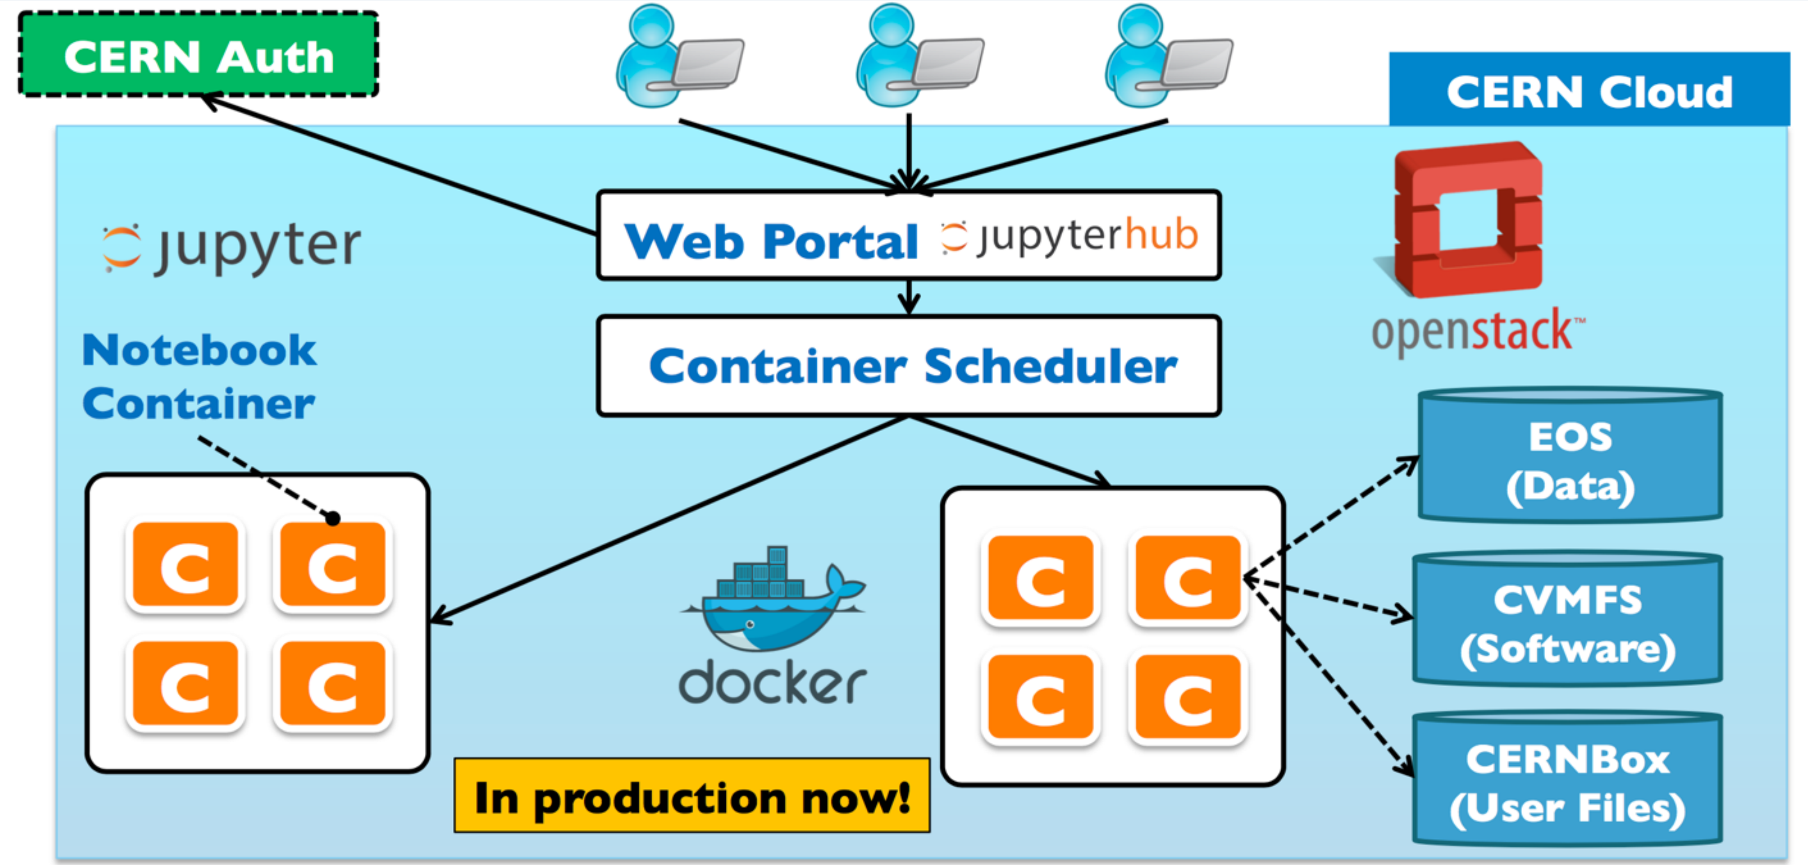

Databases

- Oracle
  - For critical, transactional load (~100)
    - Example of critical production DBs:
      - Quench Protection System 150.000 changes/s
      - LHC logging database ~550 TB, expected growth up to ~100 TB / year
  - Administered by DBA team
- DBaaS
  - Different database engines:
    - MySQL (340), PostgreSQL (89), InfluxDB (31)
  - Instance owners have the most of DBA rights
- Hadoop
  - For sequential load and analytics

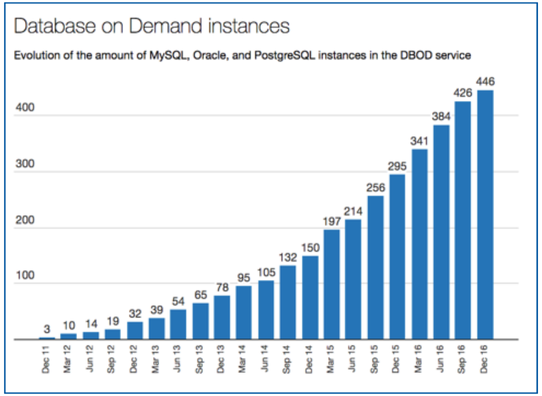

Hadoop and Analytics

- **New scalable analytics services**
  - Hadoop ecosystem
  - Spark and Kafka
  - Time Series databases
- Activities and objectives
  - CERN IT monitoring now moved to this
  - New CPU accounting developing on this infrastructure
  - Scalable platform, future evolution based on requirements

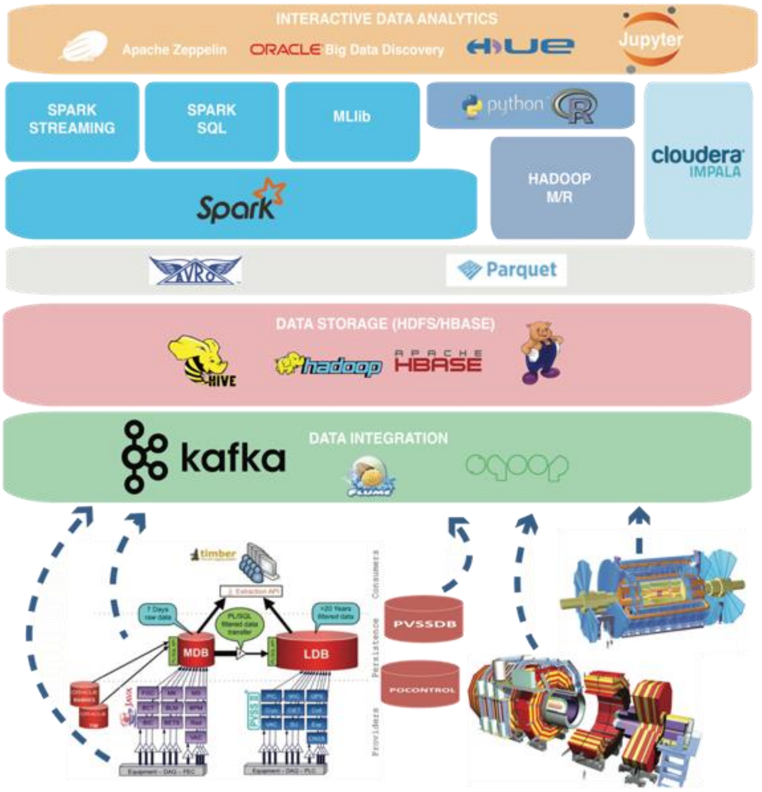

DB future work

- Hadoop / Spark / Kafka analytics service
  - Major commitments towards Kafka/Hadoop (LS2): WinCC OA (PVSS), ACCelerator LOGging (NXCALS), etc. require advanced service level
  - ATLAS EventIndex support and help for performance
  - Kafka pilot started to satisfy key projects (BE, IT, etc.)
  - Spark with ROOT/SWAN with EP-SFT

- “Preparation for Oracle database release 12.2 upgrade during LS2
- High Availability deployment for Database on Demand

- Database Futures” workshop 29th-30th May this year

Ops tools

- Open-source everywhere
  - Focusing on Puppet
  - Focusing on automation (e.g Rundeck)
  - Contribute, don’t build (or build within eco-system)
- Orchestration
  - Terraform (@hashicorp) for general cloud
  - OpenStack Magnum for containers, Heat for VMs
- Monitoring
  - Moving to collectd + analytics platform
    - Kafka to Elasticsearch (evaluating InfluxDB for time-series)
    - Kibana for detail, Grafana for dashboards
    - Spark/Hadoop for offline analysis

Tier-0 summary

- Use OpenStack base everywhere
- Use HTCondor for all compute
  - Efficiency, opportunistic mid-SLA, agility with standard tools
  - Cloud: no rush, maintain strategic capability
- New computing centre, to decide soon
- Consolidation on EOS for storage
- Databases DBoD further strengthening
- Big data / analytics tools in production
  - use-cases moving to the new analytics platform
- Ops tools – open-source, contribute
  - Focus on monitoring and automation

Backup

Some consequences

- Need to constrain costs
  - The main driver for WLCG is overall data volume (replication) & wide distribution
    - 2/3 of total global cost is in disk
- Infrastructure must no longer be (too) special
- Need to be able to use commercial & opportunistic resources
  - Including non-traditional for HTC: HPC, cloud, special architectures etc.
- Thus need significant agility and performance optimisation through software
- Learn from our experience and that of large internet companies
- Need flexibility/agility to changing markets – e.g. cost of commercial resources, obsoleted technologies (perhaps overnight)
- Must recognise and leverage opportunities of commonality at all levels
  - Between experiments, HEP, across disciplines, with industry, …

➢ Failure to change will limit the scientific output

HEP Facility timescale

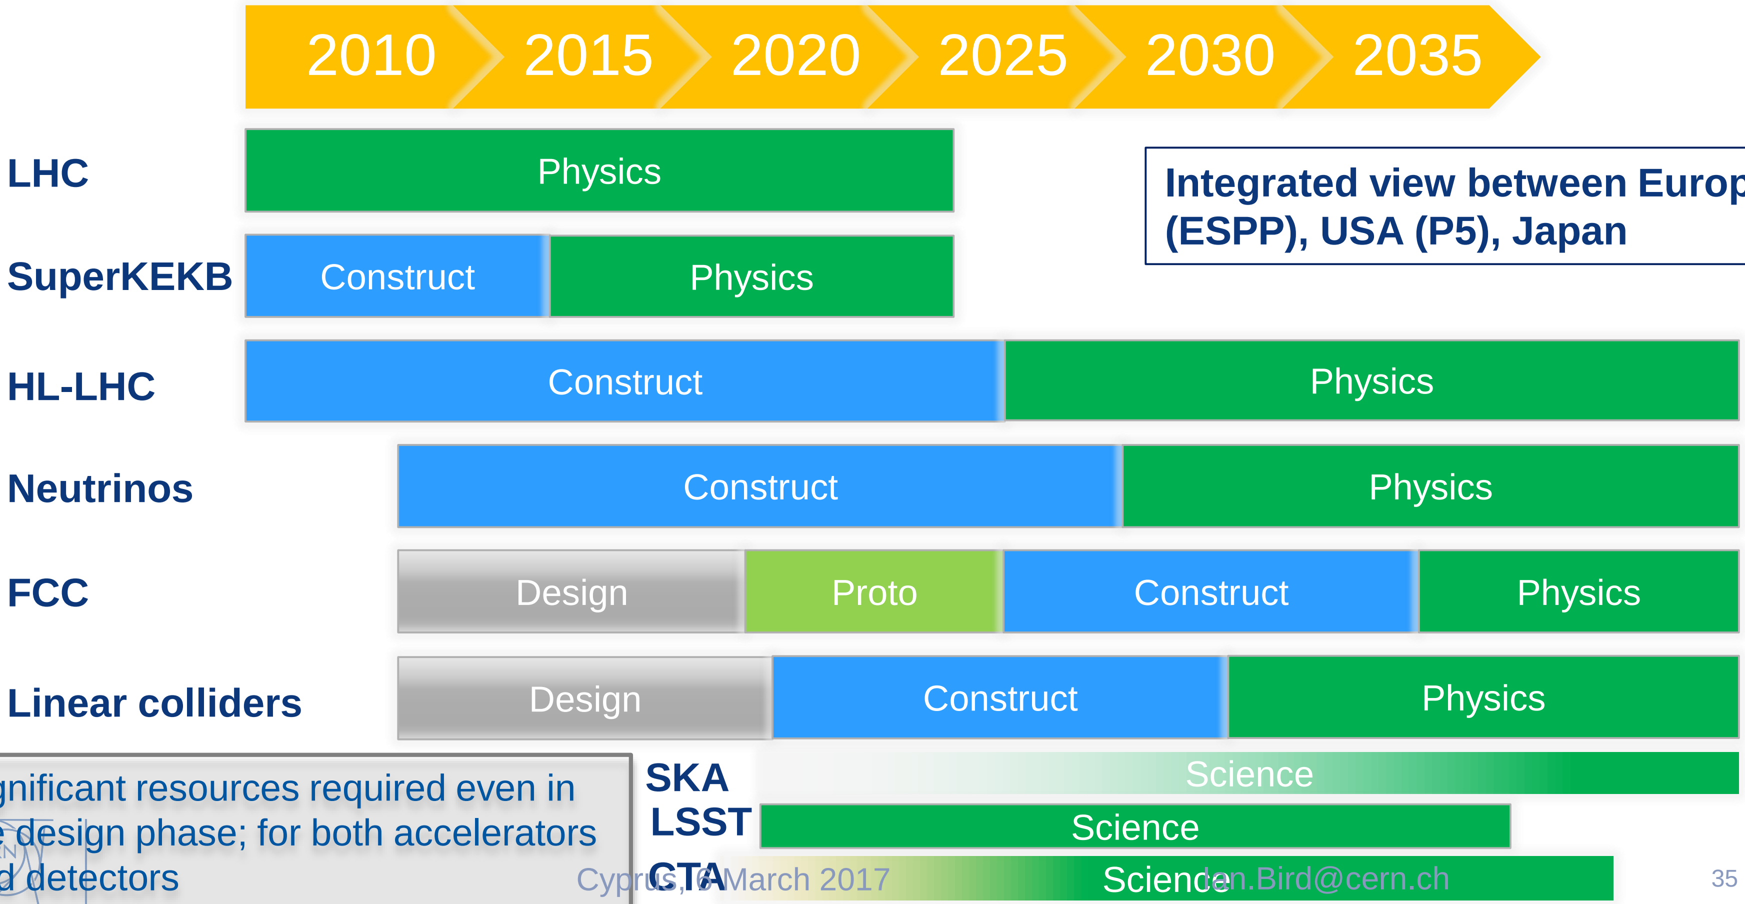

Integrated view between Europe (ESPP), USA (P5), Japan

Significant resources required even in the design phase; for both accelerators and detectors

Future Challenges

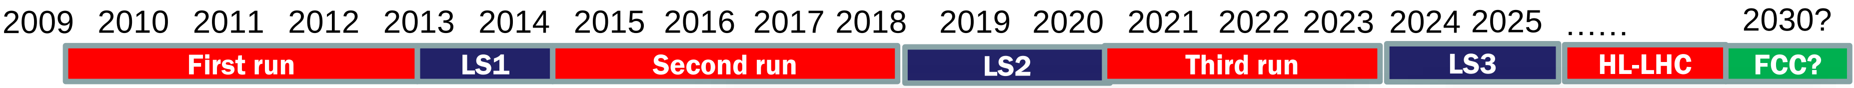

- Raw data volume for LHC increases exponentially and with it processing and analysis load
- Technology at ~20%/year will bring x6-10 in 10-11 years
- Estimates of resource needs at HL-LHC x10 above what is realistic to expect from technology with reasonably constant cost

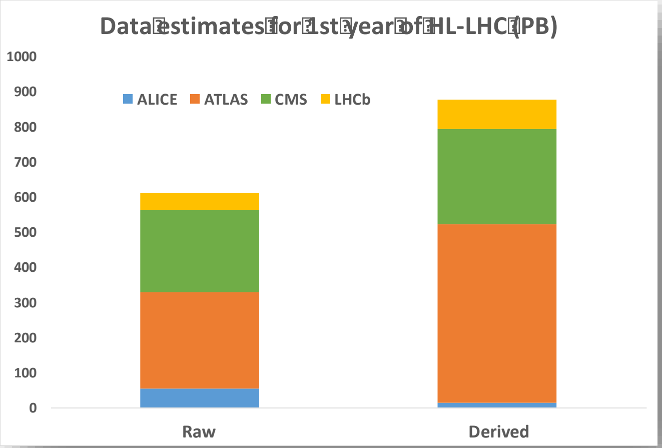

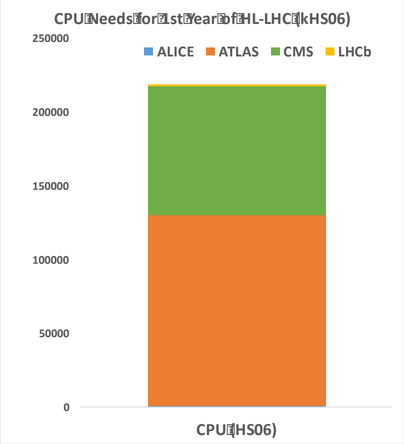

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU:
- x60 from 2016

Possible Model for future HEP computing infrastructure

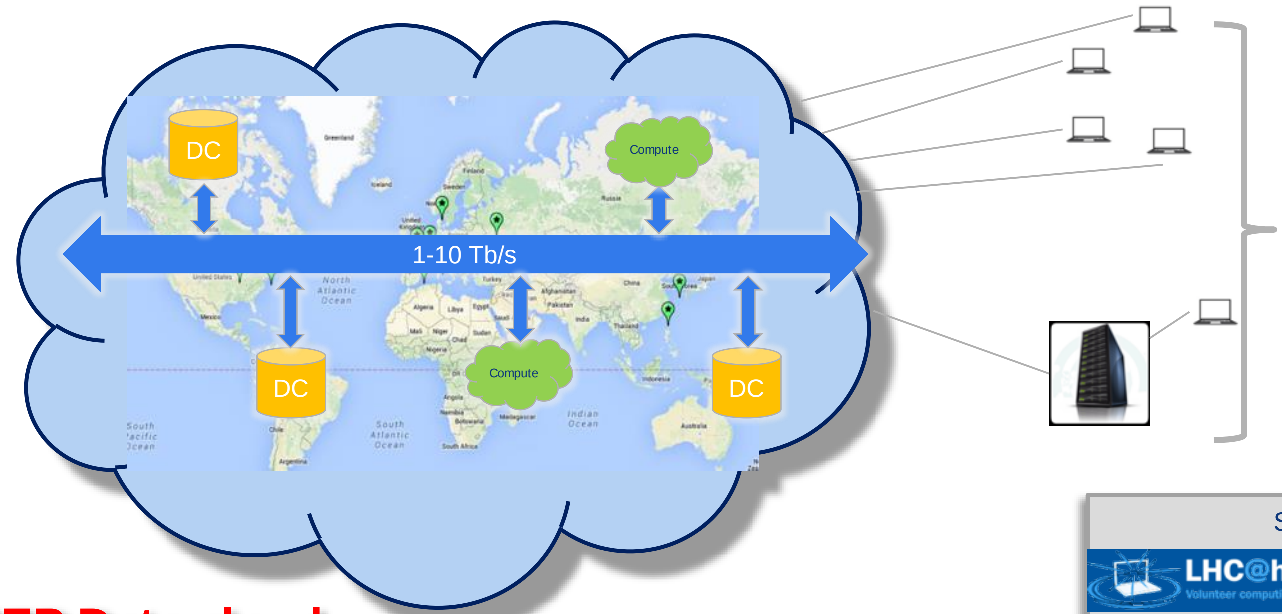

Cloud users:
Analysis

HEP Data cloud

Storage and compute