

CERN and HelixNebula, the Science Cloud

Fernando Barreiro Megino (CERN IT)

Ramón Medrano Llamas (CERN IT)

Dan van der Ster (CERN IT)

Rodney Walker (LMU Munich)



HELIX NEBULA THE SCIENCE CLOUD


Strategic Plan

- ▶ Establish multi-tenant, multi-provider cloud infrastructure
- ▶ Identify and adopt policies for trust, security and privacy
- ▶ Create governance structure
- ▶ Define funding schemes



To support the computing capacity needs for the ATLAS experiment

EMBL



Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity

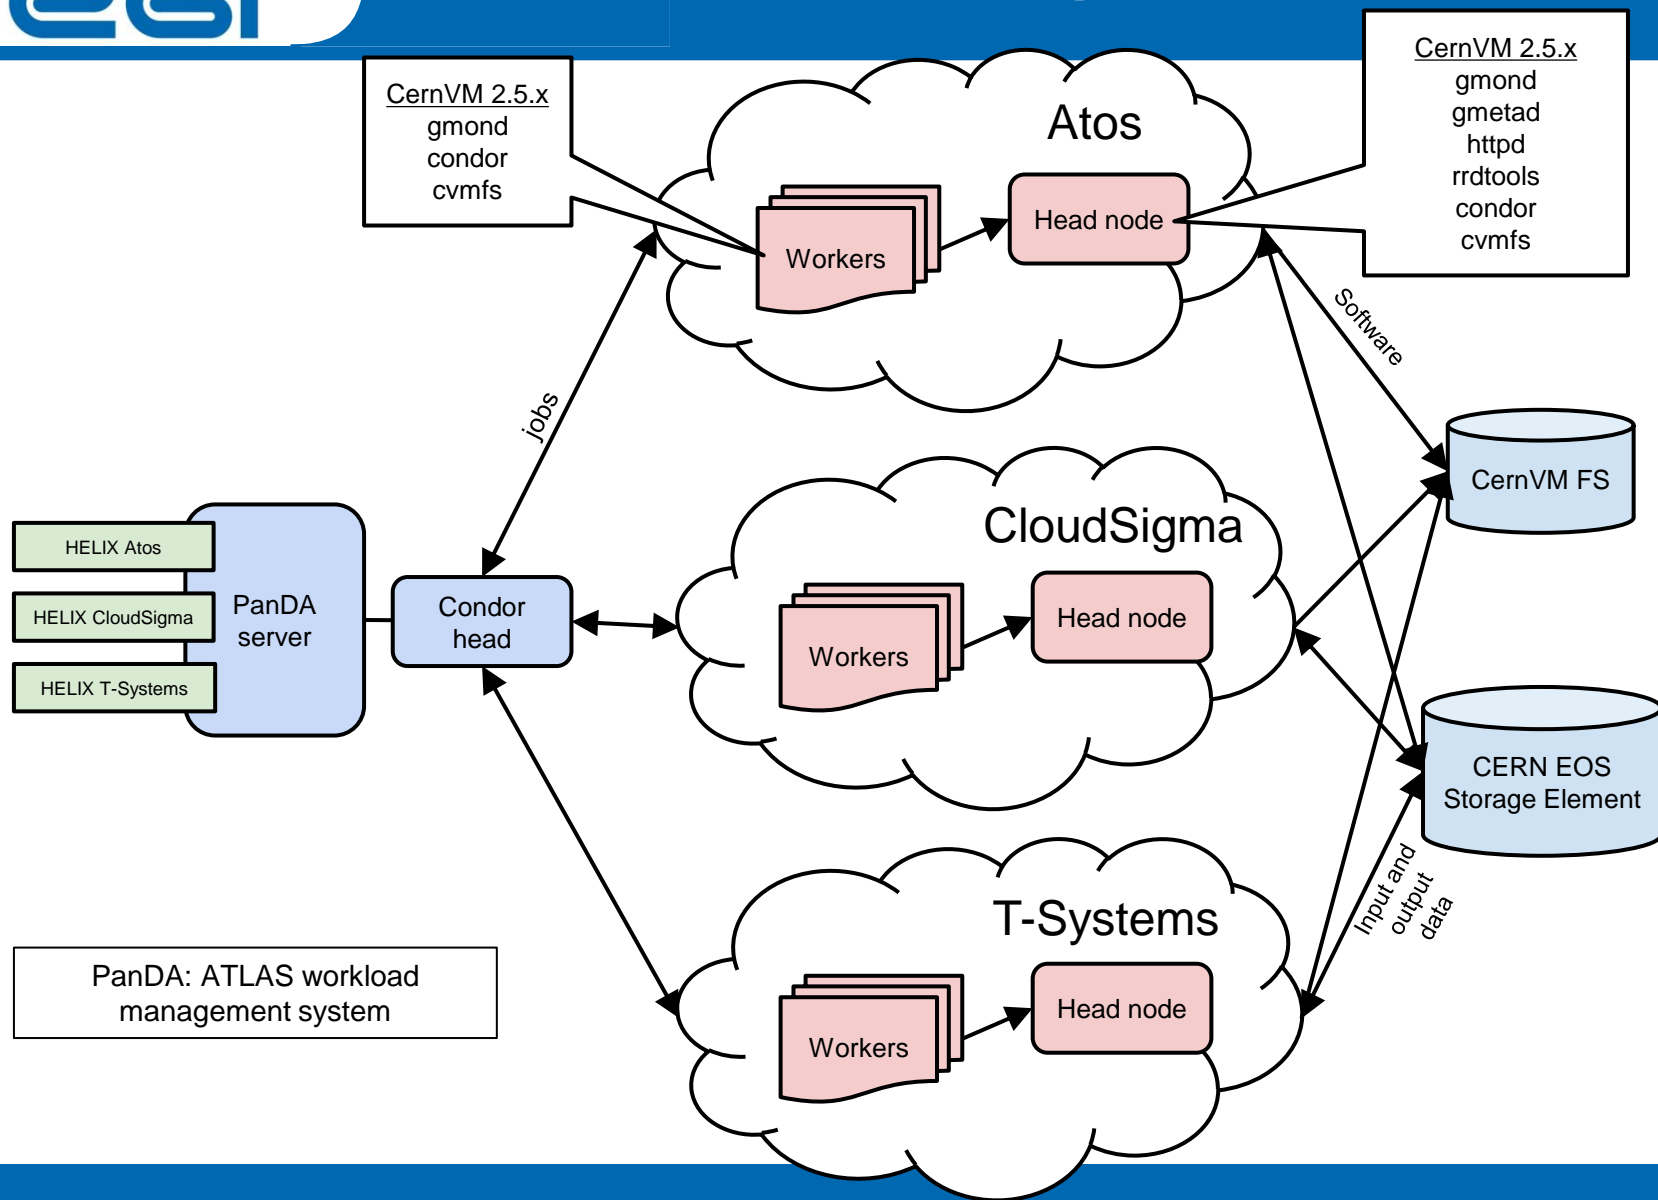


To create an Earth Observation platform, focusing on earthquake and volcano research



- ***Can ATLAS jobs run on cloud resources?***
- Run **Monte Carlo** jobs on 3 commercial providers (Atos, CloudSigma, T-Systems):
 - Very simple use case **without any data requirements**
 - 10s MB in/out
 - CPU intensive jobs with durations of ~6-12 hours/job

Configuration



- Each cloud offering was different
- Different concepts of IaaS
 - Persistent VMs: cloning of full disk to boot new instance
 - Ephemeral VMs: can be accidentally lost
- Requirement of different image formats of CernVM
- Proprietary interfaces
- Possibility of user contextualization only straightforward in one provider
 - Otherwise we followed the “golden disk” model


Consequences:

- Needed to start from scratch with each provider
- Set up and configuration was time consuming and limited the amount of testing


- ATM we are the only institute to finish successfully the 3 PoCs
- In total ran ~40.000 CPU days
 - Few hundred cores per provider for several weeks
 - CloudSigma generously allowed us to run on 100 cores for several months beyond the PoC
- Running was smooth, most errors were related to network limitations
 - Bandwidth was generally tight
 - One provider required VPN to CERN
- Besides MonteCarlo simulation, we executed HammerCloud jobs:
 - Standardized type of jobs
 - Production Functional Test
 - All sites running the same workload
 - Performance measuring and comparison

Total job completion time


Wallclock time at CERN-PROD (s)



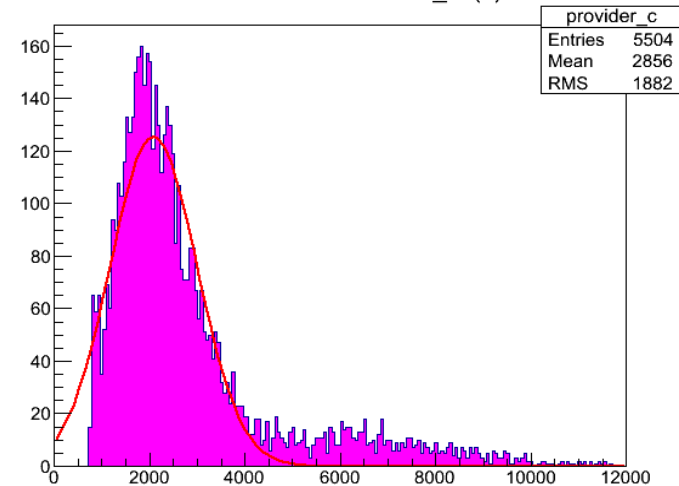
Wallclock time at HELIX_A (s)



Wallclock time at HELIX_B (s)

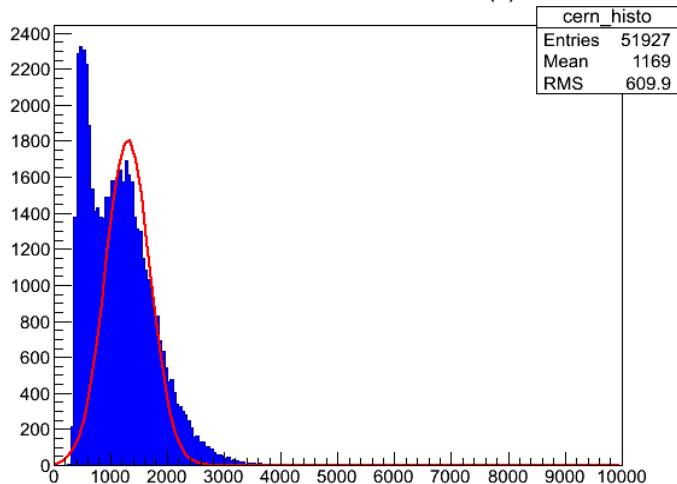


Wallclock time at HELIX_C (s)




Athena execution time


Athena time at CERN-PROD (s)



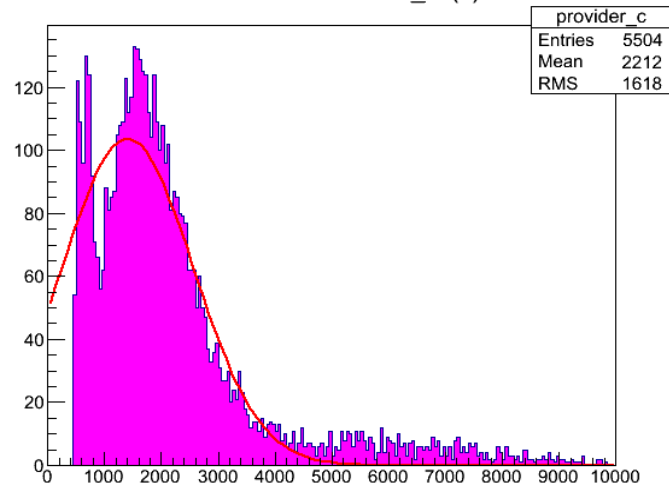
Athena time at HELIX_A (s)



Athena time at HELIX_B (s)




Athena time at HELIX_C (s)

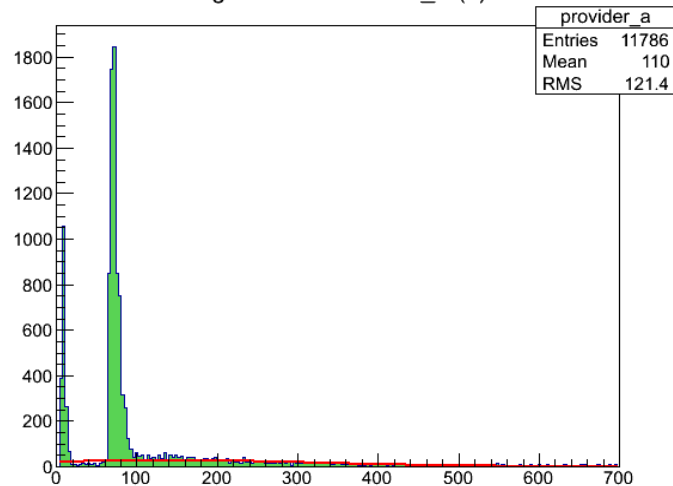


Stage in time

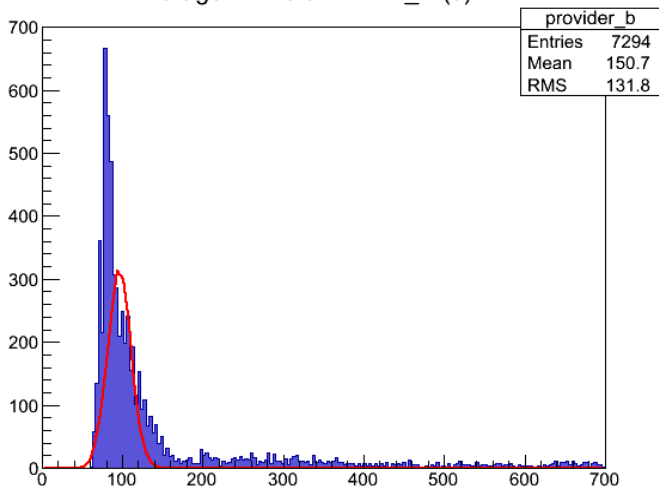
Stagein time at CERN-PROD (s)



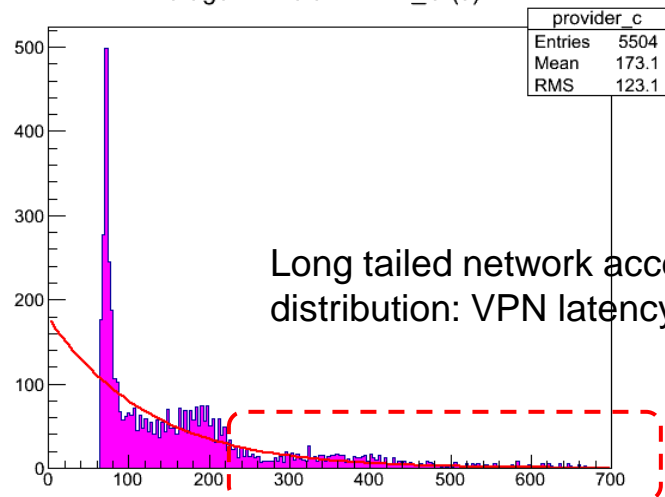
Stagein time at HELIX_A (s)



Stagein time at HELIX_B (s)



Stagein time at HELIX_C (s)

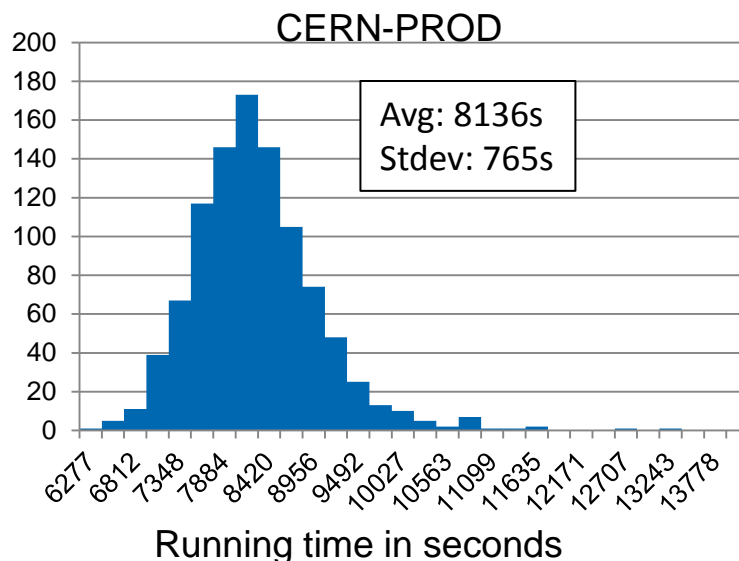




Long tailed network access distribution: VPN latency

- Finished 6 x 1000-job **MC** tasks over ~2 weeks
- 200 CPUs at HELIX_A used for production tasks
- We ran 1 identical MC simulation task at CERN to get reference numbers

	HELIX_A	CERN
Failures	265 failed, 6000 succeeded	36 failed, 1000 succeeded
Running times	16267s±7038s	8136s±765s

- Wall clock performance cannot be compared directly, since the HW is different
 - HELIX_A: ~1.5Ghz AMD Opteron 6174
 - CERN-PROD: ~2.3GHz Xeon L5640





- We demonstrated the feasibility of running ATLAS workload on the cloud
- We did not touch the storage and data management part
 - each job brought the data over the WAN
 - MC simulation is best suited for this setup
- Cloud computing is still a young technology and the adoption of standards would be welcome
- We do not have cost estimates
 - Best comparison would be CHF/event on the cloud and on the sites

- TechArch Group:
 - Working to federate the providers
 - Common API, Single-sign-on
 - Image & Data marketplace
 - Image Factory (transparent conversion)
 - Federated accounting & billing
 - Cloud Broker service
 - Currently finalizing requirements
 - Two implementation phases:
 - Phase 1: *lightweight* federation, common API
 - Phase 2: image management, brokerage
- Improve connectivity by joining cloud providers to GEANT network
 - only for non-commercial data
- Collaborate with EGI Federated Cloud Task Force
 - seek common grounds
- For ATLAS: Possibly repeat PoCs after Phase 1 implementations