


Large scale MC simulation in Helix Nebula commercial cloud


C. Cordeiro, A. Di Girolamo, L. Field, L. Villazon
D. Giordano
(CERN IT-SDC)

Helix Nebula: Public & Private Partnership




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 LHC experiments

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

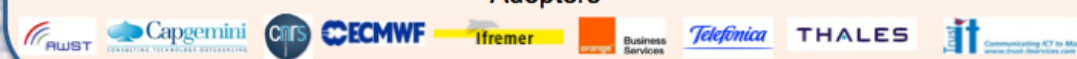


To improve the speed and quality of research for finding surrogate biomarkers based on brain images

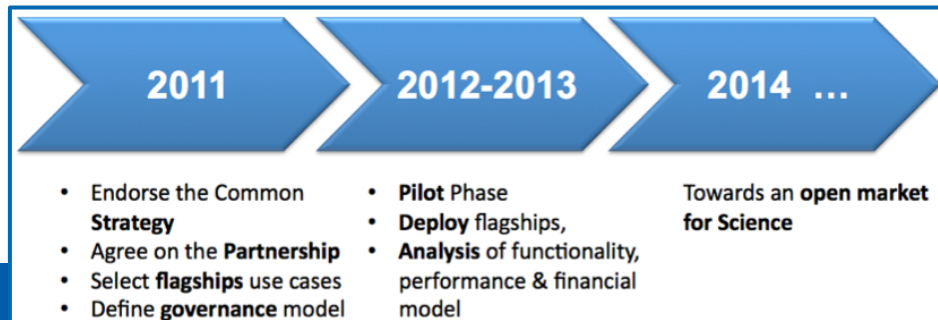
Suppliers



Adopters



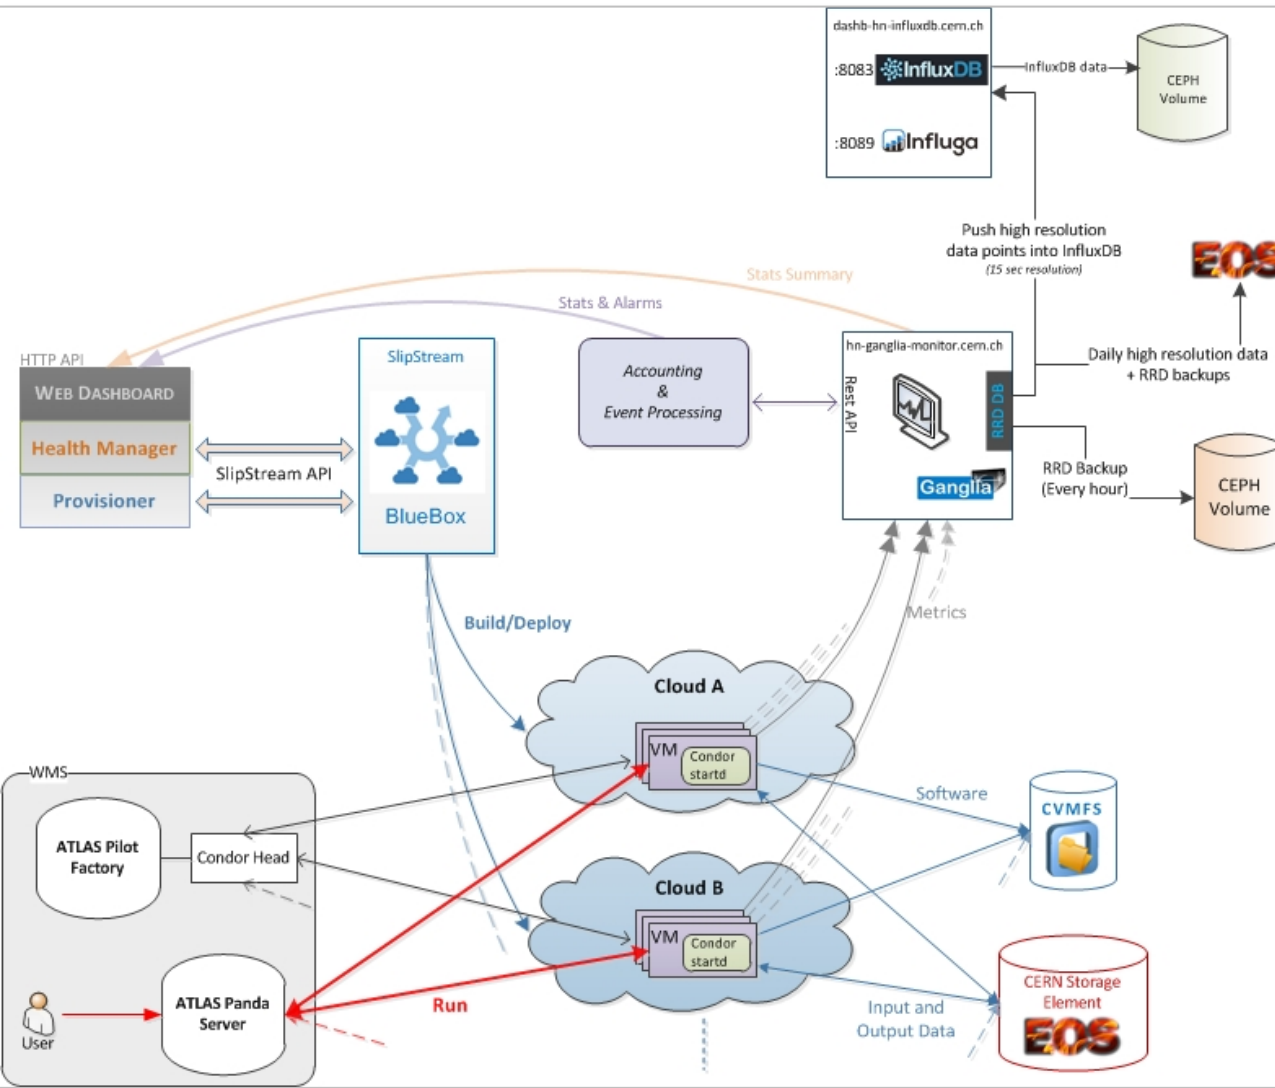
Project phases



Monte Carlo production in HN: Milestones

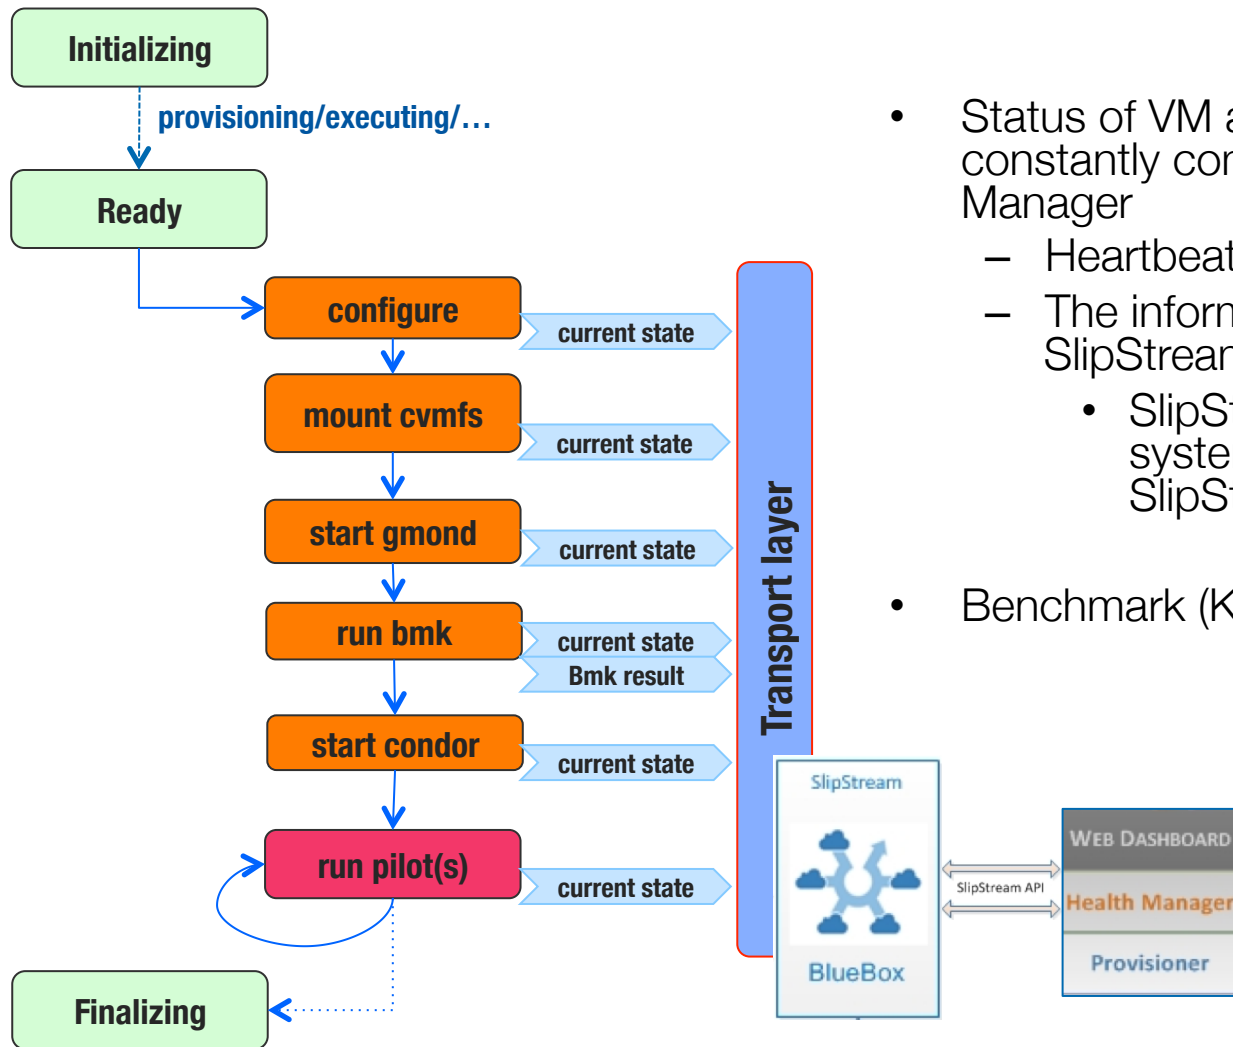
- Nov. 2014 CERN price enquiry
 - Up to 2,000 VMs for 45 days (from mid Feb. to end March)
 - VM specs: 1 vCPU, 2 GB memory, 20 GB free HD, public IP. OS: CentOS6
 - Run ATLAS Geant4 Simulation: high CPU/wall time jobs, reduced I/O
 - **Atos** made best offer. Successful bid confirmed early Jan 2015
- Preparation phase to get the IaaS ready
 - **Network**: acquire enough IP addresses and configure network
 - **Storage**: configure NetApp mapping to volumes and storage LUNs, tune # of vFiler images to cope with ISCSI protocol specs
 - **Hypervisor**: balance KVM load, speed up LUNs' scanning
 - **Brokerage**: improve robustness and fault tolerance
- Production phase (March 3rd to 31st)
 - Resources available: **200 KVM, 16 cores** each
 - Capacity acquired: started with ~2,000 running VMs, reached 3,000 VMs
 - Several improvements introduced during the production phase
 - **auto-scaling**, revised **provisioning approach**, **reduced** operational overhead

Provisioning & monitoring chain




- VM management agent
 - Define, achieve, maintain needed capacity
- WMS
 - Dedicated PanDA resource: HELIX_NEBULA_ATOS
 - APF: aipanda02/09/13
- Key role of VM monitoring
 - Real-time monitoring
 - Alarming
 - Accounting
 - Benchmarking
- Strategy
 - Ganglia data preserved with **15s** time resolution
 - **Benchmark** each VM at startup

VM execution progress




- Status of VM and running services is constantly communicated to the Resource Manager
 - Heartbeats and sensors
 - The information is retrieved from SlipStream dashboard
 - SlipStream includes a messaging system between the VMs and SlipStream server
- Benchmark (KV) runs before starting condor

Results in a nutshell



- Up to **3,000** concurrent running VMs
 - 4 (+1) weeks of production
 - **~1.2 million** CPU hours of processing
- ATLAS GEANT4 Simulation of $t\bar{t}$ events
 - **~11.5 million** events processed \Leftrightarrow $\sim 80,000$ jobs
 - **~93%** CPU/Wall time ratio
 - ~ 9 hours single job duration
 - **~97%** job wall time used for successful runs
 - Lost heartbeat is the main source of failures ($\sim 81\%$)




Compared with other ATLAS cloud sites (March)




NEvents Processed in MEvents (Million Events) (Sum: 23.00)

HELIX_NEbULA_ATOS - 38.78%



- Significant contribution compared with other ATLAS cloud sites running simulation
 - **Largest** # of processed events
 - **Longest** wall time consumption
 - **High** wall time efficiency

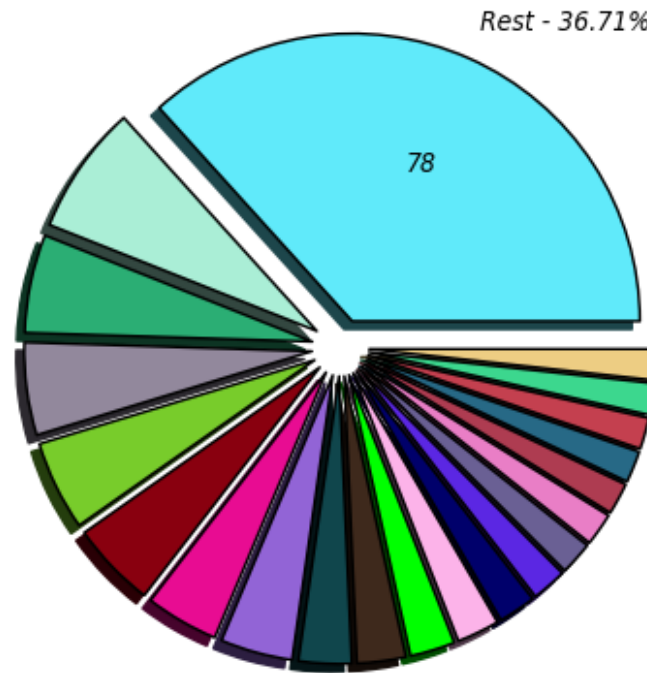
(*) CERN-PROD_CLOUD: 90% of processed events belongs to fast workloads (~24 s/evt)



Comparison with all ATLAS single core simulations in March



NEvents Processed in MEvents (Million Events) (Sum: 214.00)



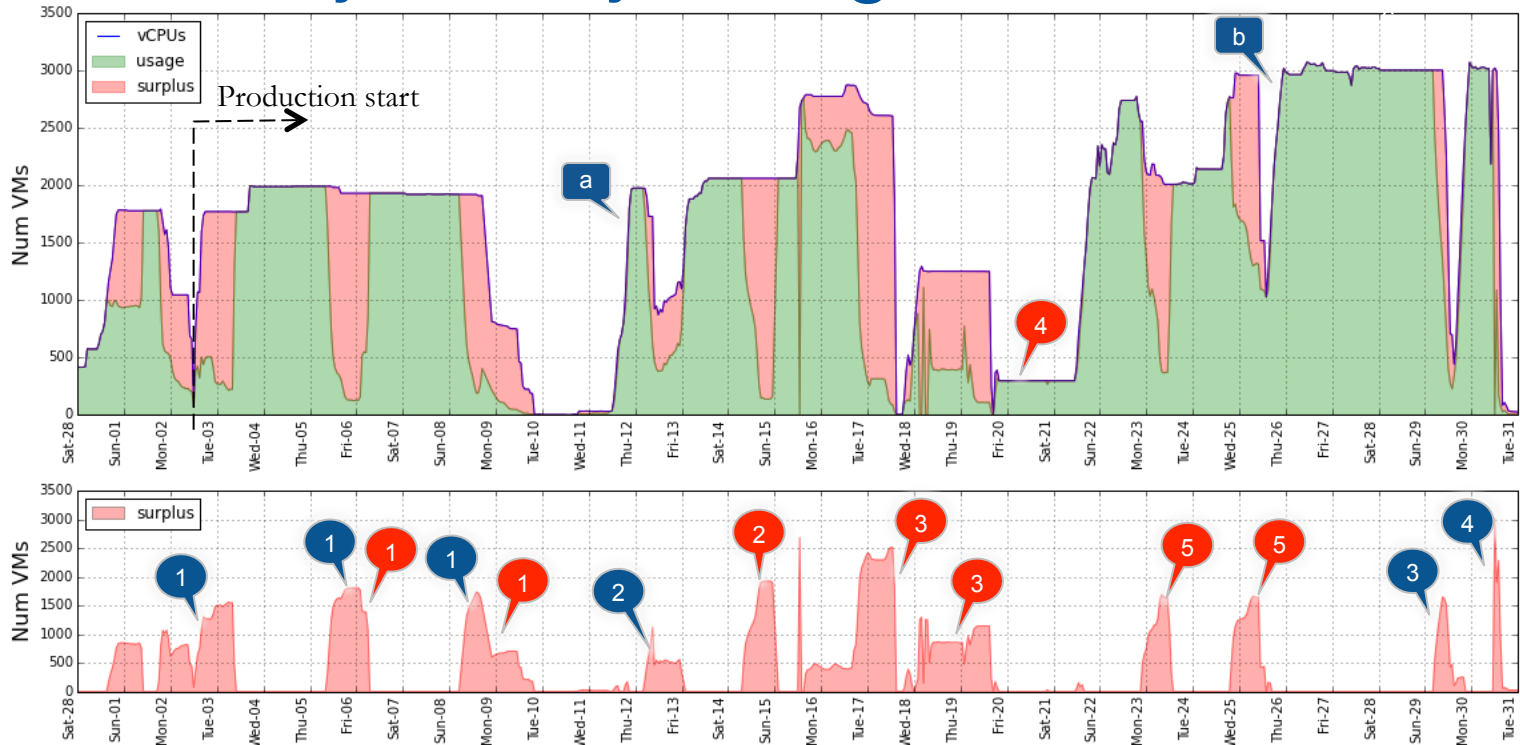
<http://cern.ch/go/M17V>

Rest - 36.71% (79.00)	TRIUMF - 7.27% (16.00)
RAL-LCG2_SL6 - 5.53% (12.00)	IN2P3-CC - 5.20% (11.00)
UKI-NORTHGRID-LANCS-HEP_SL6 - 5.04% (11.00)	BNL_PROD - 4.82% (10.00)
UKI-SCOTGRID-GLASGOW_SL6 - 4.11% (9.00)	HELIX_NEBULA_ATOS - 4.06% (9.00)
CONNECT - 2.96% (6.00)	UKI-SOUTHGRID-OX-HEP_SL6 - 2.74% (6.00)
CERN-PROD - 2.53% (5.00)	UKI-SOUTHGRID-RALPP_SL6 - 2.40% (5.00)
BU_ATLAS_Tier2_SL6 - 2.11% (5.00)	FZK-LCG2 - 1.97% (4.00)
UKI-NORTHGRID-MAN-HEP_SL6 - 1.87% (4.00)	INFN-T1 - 1.85% (4.00)
CERN-PROD_CLOUD - 1.78% (4.00)	UKI-LT2-Brunel_SL6 - 1.78% (4.00)
UKI-LT2-OMIL_SL6 - 1.77% (4.00)	

plus 2 more

Helix Nebula has been the 7th ATLAS resource in terms of sim events processed in March by ATLAS single cores

The monthly activity: Ganglia source



CERN

- 1 Task completed (no more jobs)
- 2 CERN network issues
- 3 Agent auth. cache not renewed
- 4 Task abruptly terminated

- a Improvement: auto-scaling (up/down) based on load
- b Improvement: Orchestrator-less single-VM runs

Surplus: amount of running VMs not effectively used for production

Reasons


-
-

IaaS

- 1 VMs stuck in provisioning
- 2 Missing cloud layer report
- 3 Read-only file system
- 4 Stuck orchestrators
- 5 Stuck deployments



The monthly activity: Ganglia source



CERN

- 1 Task completed (no more jobs)
- 2 CERN network issues
- 3 Agent auth. cache not renewed
- 4 Task abruptly terminated

- a Improvement: auto-scaling (up/down) based on load
- b Improvement: Orchestrator-less single-VM runs

Resource usage:

- Effective: **77%**
 - Surplus causes 10% CERN, 13% IaaS
- After improvement “b”:
 - Effective: **93%**


IaaS

- VMs stuck in provisioning 1
- Missing cloud layer report 2
- Read-only file system 3
- Stuck orchestrators 4
- Stuck deployments 5




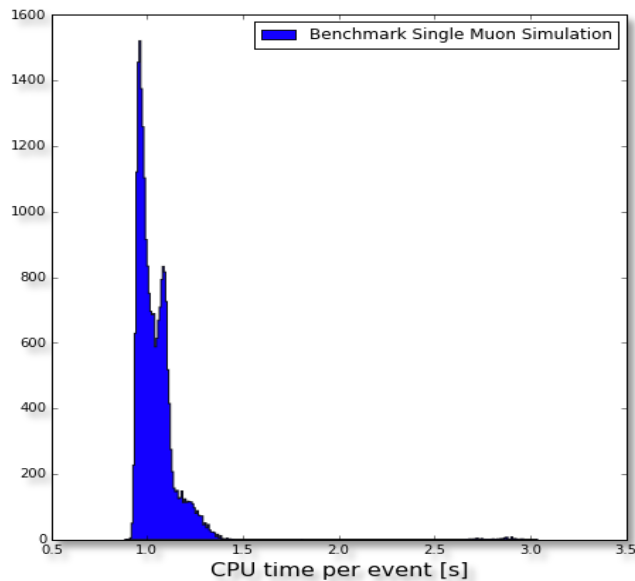
Consumer-side accounting

- Requirements
 - Validate the invoices received about the used resources
 - Monitor the effective usage of resources and identify inefficiencies
- CERN Ganglia monitoring as been the **reference** for used resources
 - IaaS accounting in agreement with CERN Ganglia report




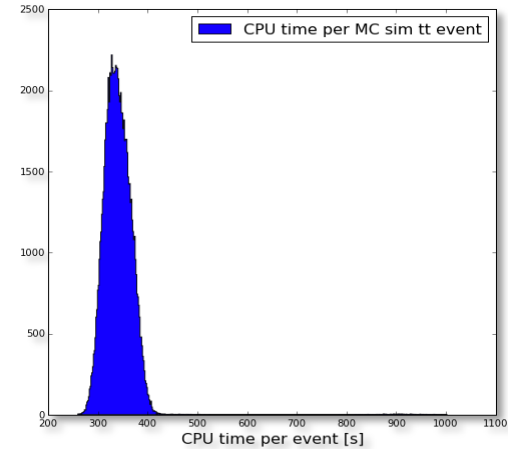
Benchmarking

- Each created VM has been benchmarked using ATLAS KitValidation
 - ~30,000 VM benchmark performed
 - 100 Single Muon events simulated (~2 min to run)
- Results
 - CPU performance uniform within 15% spread
 - Benchmark profile consistent over time




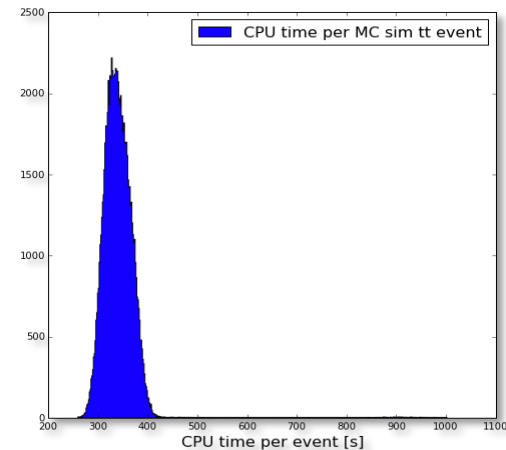

Benchmarks vs Job performance

- Consistent job CPU performance and benchmark




Benchmarks vs Job performance

- Consistent job CPU performance and benchmark
 - Correlated behavior
- Outliers detection
 - KV bmk (2') is a *prompt and effective solution* to identify VMs with poor performance




Activity organization

- CERN IT-SDC team: Cris, Luis, Alessandro, Domenico
 - Crucial support from ADC people: in particular P. Love, J. E. Garcia
- Activity tracking: [CERN Jira project](#)
 - Over **100** tickets
 - Issue tracking and consolidation work, daily monitoring, time tracking
 - ~30 FTE days (**240 hours**) over 74 days of overall activity
 - 3 areas: tool consolidation (~8 FTE days), pre-prod (~13 FTE days), prod (~9 FTE days)
- Regular **CERN-SixSQ-Atos meetings** (15 of ~1h each)
 - Comprehensive **notes** of Issues and Actions
 - Established ticketing system to notify issues with Atos



Lessons learned

- **Successful ATLAS production** with ~11.5 million events processed
- **Beneficial experience**
 - Managing VMs (and facing issues) in a “cloudy” commercial IaaS
 - Monitoring, accounting and benchmarking cloud resources
- Client-side accounting is mandatory in a commercial context
- Benchmarking of each single VM is possible
 - Adopted fast benchmark running at VM startup
 - Transparently used in other commercial cloud IaaS
- Results reported at CHEP’15
 - Two contributions: Helix Nebula, Ganglia
- Next CERN procurement initiatives [See R. Jones talk at previous GDB]
 - Continue evaluation & integration of commercial cloud IaaS





www.cern.ch