



# ATLAS COMPUTING ON THE SWISS CLOUD SWITCHEngines

## BACKGROUND

Switzerland contributes with up to 4% of the ATLAS T2 CPU capacity with four standard linux clusters. We investigate usage of HPC Crays and academic cloud resources for more compute for same money. Since 2015 the Swiss National Research and Education Network (NREN) offers hours on an OpenStack infrastructure for academic usage. Since 2016 ATLAS runs on the infrastructure.

## IaaS

SWITCHEngines ([www.switch.ch/engines](http://www.switch.ch/engines)) is an Infrastructure as a Service (IaaS) based on OpenStack. In 2016 ATLAS has access to testbed quota with 304 logical cores on Intel x86\_64, 2 GB RAM and 2.5 GB disk per core, and ethernet interconnect.

## CLUSTER

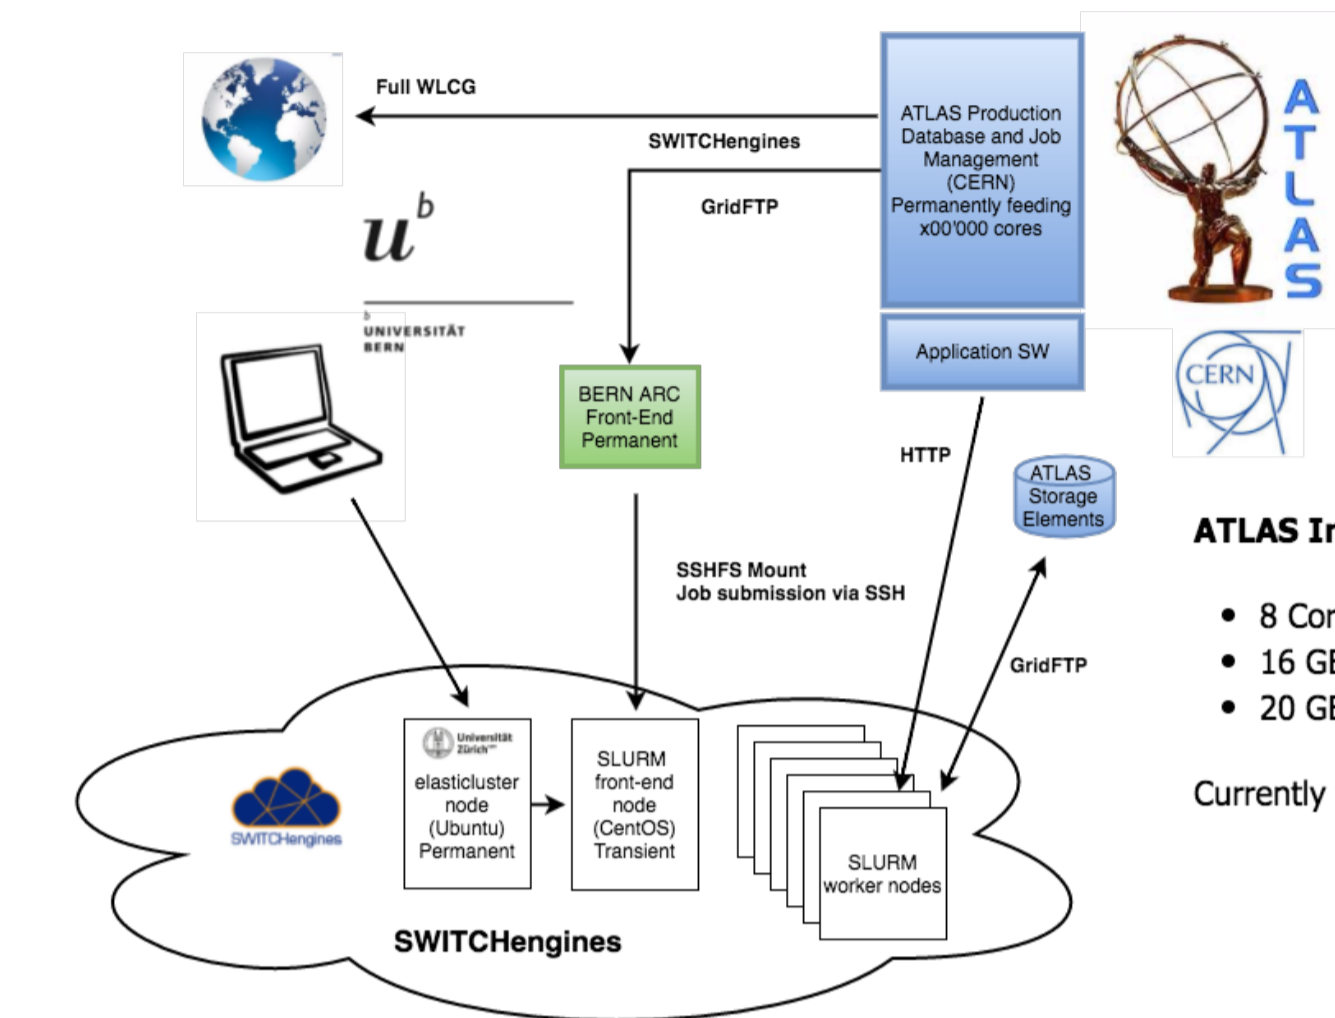
The Elasticcluster software (<http://elasticcluster.readthedocs.io/>) developed and maintained at the University of Zurich is used for setting up and resizing a SLURM cluster with CentOS nodes. This SW runs on a separate management node. With a minimal configuration, basically telling number of nodes, flavour of nodes, and image, it fires up a 300 nodes cluster on the IaaS within 30 minutes.

## GRID INTEGRATION

A remote Advanced Resource Connector Computing Element (ARC CE) is integrated within the ATLAS production system. The node ARC CE node mounts the SLURM head node via sshfs. Wrapper scripts for the ARC CE backend, formerly written and used by University of Bern to integrate HPC, are used for job submission and monitoring via ssh (Figure 1 and 2).

## PERFORMANCE

Running MC simulation only, the ATLAS accounted CPU time relative to the IaaS accounted PAY time is up to 90% (Figure 3). For comparison, relative to quota time this is up to 75%, on average lower.



**ATLAS Experiment at CERN is using SWITCHEngines 24/7 (we can set up a 1000 cores cluster for ATLAS in 1h)**

## ATLAS Instances on SWITCH

- 8 Cores
- 16 GB RAM
- 20 GB Disk

Currently testing 300 cores 24/7

**FIGURE 1 :** Integration of Slurm cluster on IaaS SWITCHEngines within the ATLAS production system. With the Elasticcluster software and ARC computing element an O(1000 cores) elastic compute resource can be initiated within an hour.

## ATLAS Grid Monitor

2016-08-30 CEST 13:27:44

Processes: ■ Grid ■ Local

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
Switzerland	ATLAS BOINC	98139	7894+6883	1571+4063
	ATLAS BOINC 3	98139	5815+8163	1253+4371
	ATLAS BOINC TEST	644	8+0	0+0
	Bern ce01 (UNIBE-LHEP)	1513	1848+0	156+0
	Bern ce02 (UNIBE-LHEP)	770	624+0	159+0
	Bern ce04 (UNIBE-LHEP)	304	384+0	192+0
	Bern UBELIX T3	4472	385+2822	208+2450
	CSCS BRISI Cray XC40	1500	576+8	154+0
	Geneva (UNIGE-DPNC)	720	168+349	169+0
	Lugano PHOENIX T2 arc>	1920	1526+4848	411+14
	Lugano PHOENIX T2 arc>	2240	2865+3584	391+4
	Lugano PHOENIX T2 arc>	2048	1864+3784	407+1
<b>TOTAL</b>	<b>12 sites</b>	<b>212409</b>	<b>22269 + 28665</b>	<b>5071 + 10903</b>

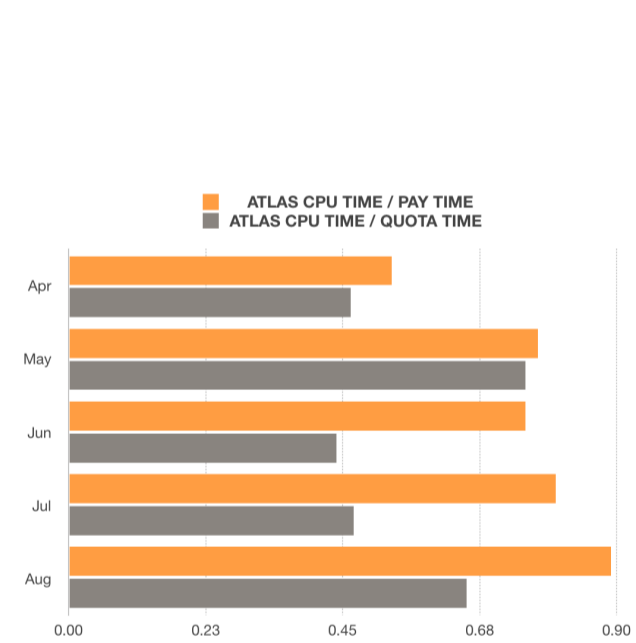
**FIGURE 2 :** Swiss compute resources serving ATLAS. In addition to standard linux clusters there are volunteer based resources (BOINC), high-end HPC (Cray) and recently IaaS (ce04) represented.

## PRICES

In 2016 regular price ([www.switch.ch/engines](http://www.switch.ch/engines)) for 1000 cores with 2GB RAM per core is 73 kCHF p.a. This is comparable to prices from other academic providers.

## CONCLUSIONS

- Cloud compute cluster is very convenient: Setup in half a day.
- Full elasticity may give about 40% gain relative to quota.
- Cost is competitive with other outsourcing options.



**FIGURE 3 :** CPU time seen by ATLAS relative to PAY time and QUOTA time seen by IaaS provider.



Elasticcluster



u<sup>b</sup>

UNIVERSITÄT BERN

AEC ALBERT EINSTEIN CENTER FOR FUNDAMENTAL PHYSICS

swissuniversities