



# CSCS Tier2 Status LHCb

*Roland Bernet*

*Universität Zürich*




- 
- Status: Generally smooth running for both clusters (Phoenix and Piz Daint).
- Tickets: In general I do not see CSCS tickets as I am not involved. Tickets are dealt with centrally. I get only involved, when CSCS and LHCb cannot solve them or something has to be discussed and/or a ticket is not the right channel.
- Monitoring: CSCS monitoring:  
Some plots are not updated anymore since end of last year, e.g. disk usage.
- LHCb monitoring:  
We have split the site CSCS into LCG.CSCS.ch (Phoenix) and LCG.CSCS-HPC.ch (Piz Daint) and have now the same monitoring for both clusters and we can therefore create the same type of plots.
- Concerns: Multicore Jobs:  
There is no big benefit for LHCb to run multicore jobs, it just adds an additional layer which can fail. We are however rewriting our framework for different reasons and multicore jobs will be supported, but still fare in the future.
- Stability:  
My impression is, that we had more problems with the service like BDII, ARC or SLURM than with the hardware of the CRAY to run LHCb jobs. There might be some issues accessing the storage from Diz Daint, but the statistic is still low.



SAM Availability: 01/05/2017 - 28/09/2017

From 2017/05/01 to 2017/09/28

Phoenix ARCs




From 2017/05/01 to 2017/09/28

Phoenix SRM



From 2017/05/01 to 2017/09/28

Piz Daint ARC






SAM Reliability: 01/05/2017 - 28/09/2017

From 2017/05/01 to 2017/09/28

Phoenix ARCs




From 2017/05/01 to 2017/09/28

Phoenix SRM



From 2017/05/01 to 2017/09/28

Piz Daint ARC







Pilot Efficiency:

01/04/2017 - 28/09/2017

Phoenix

Piz Daint







Pilot Efficiency per CE:

01/04/2017 - 28/09/2017


arc01:




arc02:



arc03:



arc04:







Job Efficiency:

01/04/2017 - 28/09/2017

Phoenix



Piz Daint




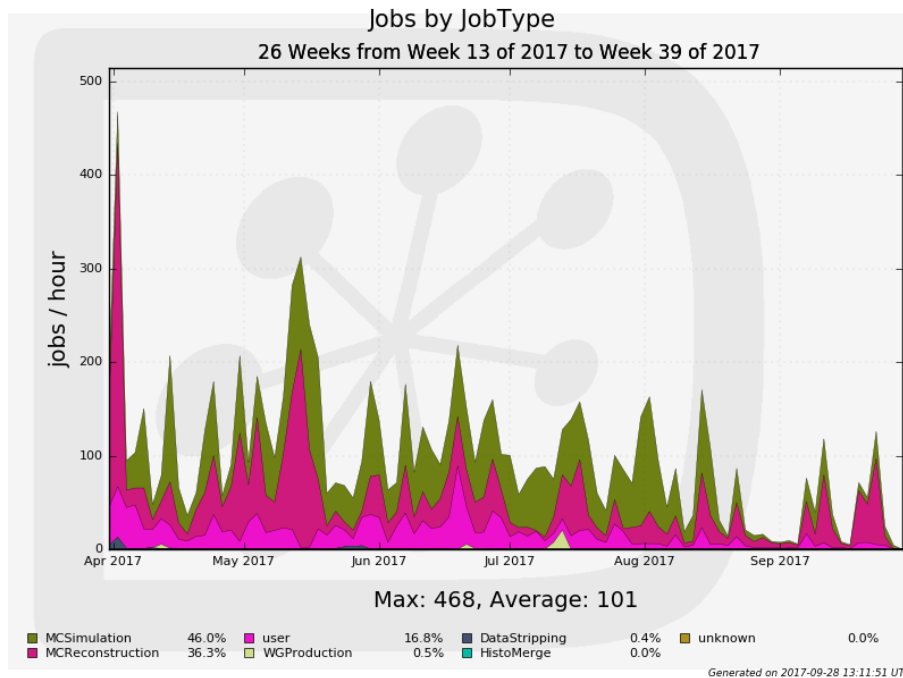


Job Types:

01/04/2017 - 28/09/2017

Phoenix

Piz Daint









CPU Efficiency:

01/04/2017 - 28/09/2017

Phoenix


Piz Daint






## Piz Daint Performance:

### Job Statistic on Piz Daint




### Job Types on Piz Daint





Piz Daint Performance:

## Pilot Statistic on Piz Daint






Job Efficiency per Type:

01/09/2017 - 28/09/2017


Simulation:




Reconstruction



WGProduction:



User:





Status: LHCb Zürich maintains a local simulation and analysis cluster, which is integrated into the institute Linux cluster. We are not part of WLCG but are part of the LHCb DIRAC framework, which allows us to run LHCb Grid jobs on idle CPUs.

Hardware: - 220 CPU cores (ca. 3700 HS06)  
- 250 TB disk space

Development: Recently we started to use the UZH ScienceCloud, an OpenStack multi-purpose compute and storage infrastructure, for our needs. Instead of replacing old hardware in our cluster, we are using CPUs in the ScienceCloud running our own worker node images. This reduces the maintenance burden for us and should overall be cheaper for everybody. Currently around 40% of the CPU power is delivered by the ScienceCloud. We also started to migrate the system scratch area to the ScienceCloud.  
The experience with the ScienceCloud and the collaboration with S3IT is very encouraging, next step is to sort out the costs / payment model.

Usage: last week (21-28 September 2017):  
dirac: LHCb LCG jobs  
others: local user jobs

