



Contribution ID: 50

Type: not specified

## Operating a production HTCondor cluster: Seamlessly automating maintenance, OS and HTCondor updates with (almost) zero downtimes

*Tuesday, 13 October 2020 11:20 (20 minutes)*

Our HTC cluster using HTCondor has been set up at Bonn University in 2017/2018. All infrastructure is fully puppetised, including the HTCondor configuration.

OS updates are fully automated, and necessary reboots for security patches are scheduled in a staggered fashion backfilling all draining nodes with short jobs to maximize throughput. Additionally, draining can also be scheduled for planned maintenance periods (with optional backfilling) and tasks to be executed before a machine is rebooted or shutdown can be queued. This is combined with a series of automated health checks with large coverage of temporary and long-term machines failures or overloads, and monitoring performed using Zabbix.

In the last year, heterogeneous resources with different I/O capabilities have been integrated and MPI support has been added. All jobs run inside Singularity containers allowing also for interactive, graphical sessions with GPU access.

Combining increasingly heterogeneous resources and different data centre locations in one cluster allows operations with almost zero (full) downtime. During this talk, some examples will be presented on how the automations can be leveraged for different interventions and how the cluster the impact on users and cluster CPU efficiency is minimized.

**Primary author:** Dr FREYERMUTH, Oliver (University of Bonn (DE))

**Co-author:** WIENEMANN, Peter (University of Bonn (DE))

**Presenter:** Dr FREYERMUTH, Oliver (University of Bonn (DE))

**Session Classification:** Computing and Batch Services

**Track Classification:** Computing & Batch Services