

Introduction into HTCondor in Zeuthen

A guide to success ;-)

Andreas Haupt – DV –
Data Science Seminar, Zeuthen

Bye, bye

... the end of an era



- Grid Engine (incl. its predecessor CODINE) has been in use for more than 25 years at DESY, Zeuthen
 - Our institute was one of the first customers, and it looks like it is one of the last, as well ...
- Grid Engine was recently purchased by Altair
 - But no real user community any longer
 - More or less all other HEP institutes moved to HTCondor during the last years
- Plan is to switch off Grid Engine in Zeuthen by the end of the year
 - If you really need it a few weeks longer, talk to us ...

Welcome



... the beginning of a new era?

- HTCondor has been driving the DESY grid clusters for some years already
- HTCondor-9 as the new driver of the interactive batch was announced in February
 - with an „invitation to test“ to all ...
 - some minor issues were found so far by a few number of testers
 - a BIG thanks to them!
- This introduction is (shamelessly) stolen from the last HTCondor week:
 - <https://agenda.hep.wisc.edu/event/1733/timetable/#20220523>
 - You will find more user-specific introduction talks there ...
 - Use them, they are really good!
- HTCondor team even provides a Youtube channel, if you prefer watching over reading ;-)
 - <https://www.youtube.com/channel/UCd1UBXmZlgB4p85t2tu-gLw/playlists>

HTCondor week user introduction

Differences to Grid Engine

- A physical (WGS) submit host is provided: **htc-submit.zeuthen.desy.de**
 - remote job submission & interaction from all other WGS is working, too
 - DAG jobs need to be submitted directly on the submit host!
- Authentication via Kerberos by default
 - kinit / klist
 - no interactions with batch system without authentication
- For DAG jobs this can/will be a problem
 - use IDTOKENS instead
 - example script on our official docs
- Also use DV service documentation at:
 - https://dv-zeuthen.desy.de/services/htcondor/job_submission/

Default settings

... if not specified differently

- Complete overview on DV service documentation web site:
https://dv-zeuthen.desy.de/services/htcondor/job_submission/

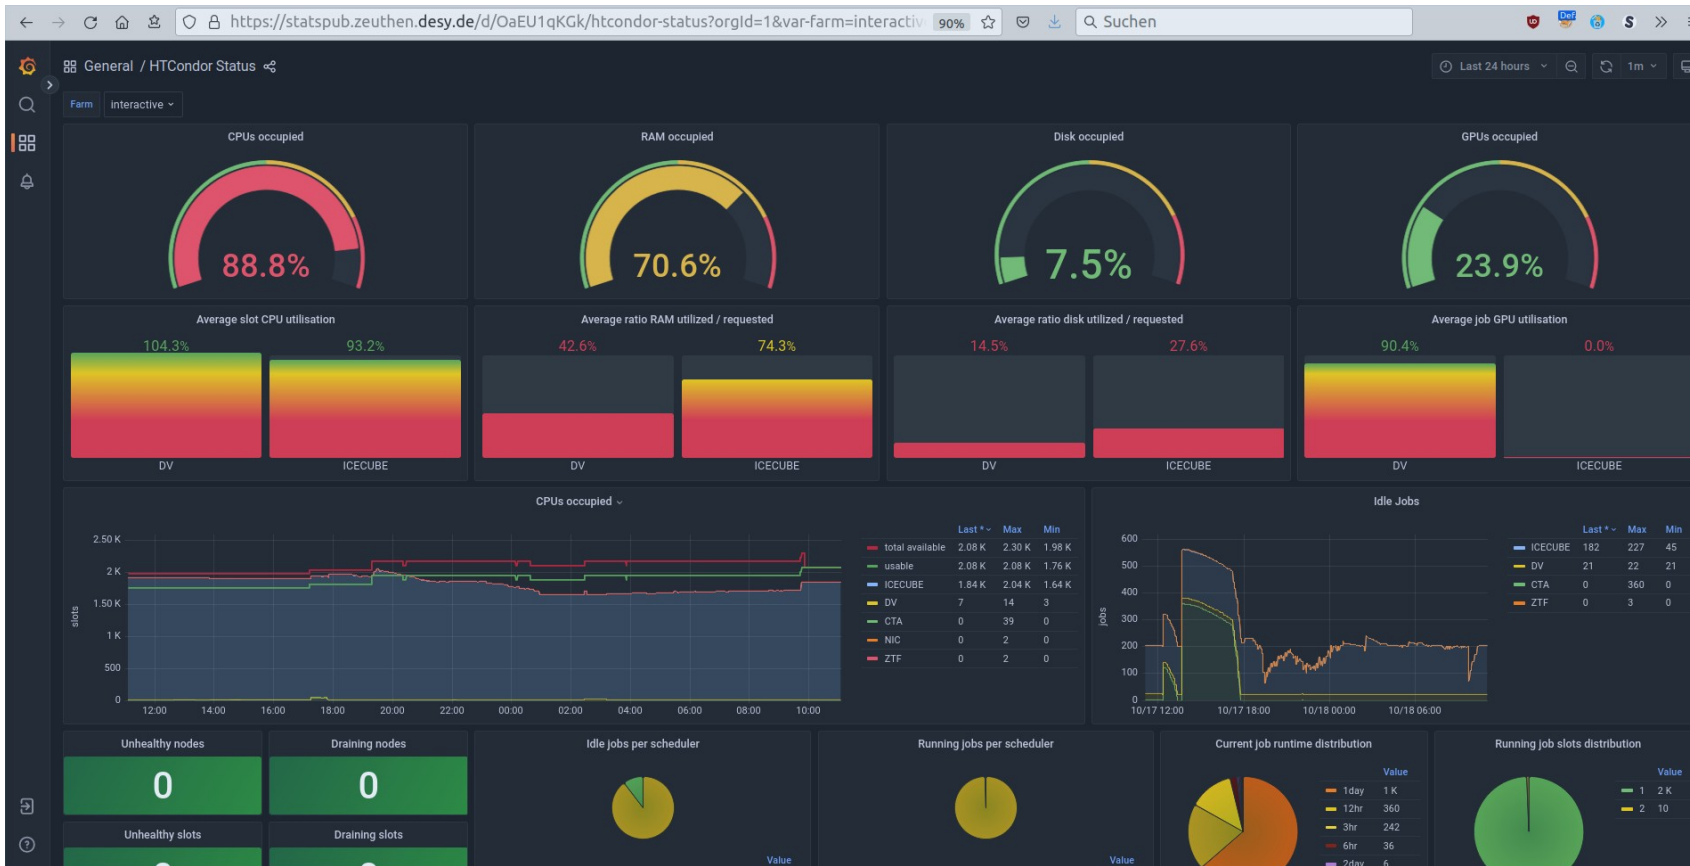
Resource	default	comment
Runtime	48 hours	up to 7 days possible, but unsupported job will be put on hold once exceeded
RAM (RSS)	1GB	job will be put on hold if exceeded by factor 2
Disk space in \$TMPDIR	1GB	no enforcement atm., subject to change if problematic!

- Jobs being on hold for more than a week will be (silently) removed from queue

A batch system dashboard ...

<https://statspub.zeuthen.desy.de/d/OaEU1qKGk/htcondor-status?orgId=1&var-farm=interactive&refresh=1m>

- Keep an eye on CPU & Memory utilisation
 - Farm is full if one of them is close to 100%



Dark times for computing

... due to European energy crisis

- Starting January 2023 (at the latest!) we need to heavily save electricity due to ridiculous price increase
- How much it affects available computing capacity is not yet 100% clear ...
 - However, be prepared for the worst
 - A reduction to 20%-40% of the current compute resources is most likely
 - ... but even worse scenarios are not unlikely
- It is not a local problem, but a European
 - Due to Zeuthen-specific energy distribution mix (~50% goes into computing) our site is more heavily affected than others, though

That's it folks ...



- Please also read the docs:
 - HTCondor in Zeuthen: https://dv-zeuthen.desy.de/services/htcondor/job_submission/
 - Upstream HTC docs: <https://htcondor.readthedocs.io/en/latest/users-manual/index.html>