

ErUM Data: COBaID / TARDIS Workshop COBaID / TARDIS @ KIT

Matthias Schnepf, René Caspart, R. Florian von Cube, Max Fischer, Eileen Kühn and Manuel Giffels

Karlsruhe Institute of Technology (KIT), SCC/ETP




■ GridKa cloud instance

- cloud provider Exoscale: 300 CPU cores (during HNSciCloud project)
- HPC Cluster ForHLR II: 400 CPU cores
- high throughput Tier-3 cluster: 462 CPU cores


■ ETP

- HPC Cluster NEMO @ Freiburg: 8000 CPU cores
- HPC Cluster ForHLR II | KIT: 260 CPU cores
- GridKa School Instance: 800 CPU cores



Current Setup


- entry point: `pps-arc-1-kit`
- HTCondor instance for opportunistic resources on `cloud-htcondor`
- resource manager TARDIS allocate and integrate resources via generalized pilots
so-called drones (batch job, container, VM)



Collaborations


Current Setup

- entry point: `pps-arc-1-kit`
- HTCondor instance for opportunistic resources on `cloud-htcondor`
- resource manager TARDIS allocate and integrate resources via generalized pilots so-called drones (batch job, container, VM)

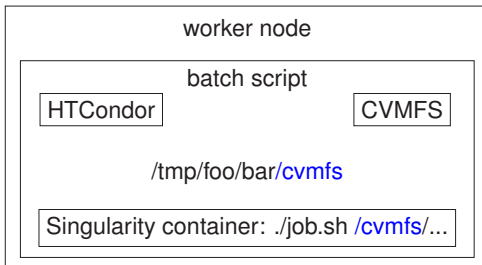


Collaborations

- request resources via batch job
- worker node runs a batch script
 - HTCondor worker node
 - CVMFS
- HTCondor starts jobs in singularity container to provide CERN CentOS 7 environment





- request resources via batch job
- worker node runs a batch script
 - HTCondor worker node
 - CVMFS
- HTCondor starts jobs in singularity container to provide CERN CentOS 7 environment



Current State


- integrate ForHLR II and GridKa Tier-3 resources
- management via TARDIS






Institute of Experimental Particle Physics

- schedule jobs to resources within an overlay batch system (OBS)
- users interact only with the OBS
- integrate resources from various resource providers on demand into a cloud



- schedule jobs to resources within an overlay batch system (OBS)
- users interact only with the OBS
- integrate resources from various resource providers on demand into a cloud



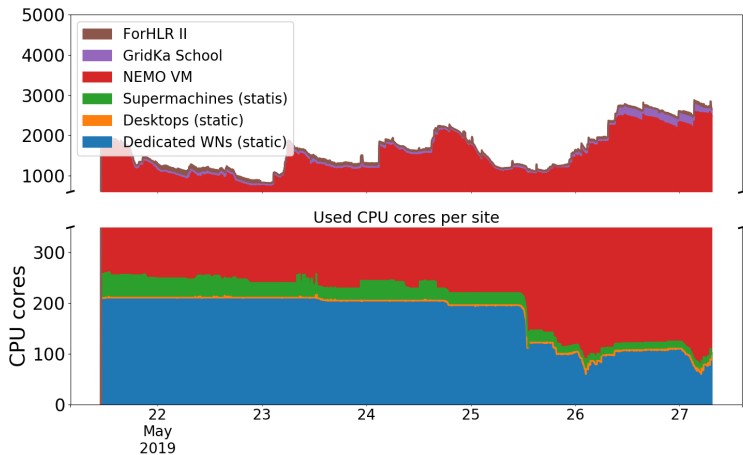
- static
 - usual worker nodes: 264 CPU cores
 - high throughput worker nodes: 84 CPU cores
 - desktop PCs (cycle stealing): 90 - 270 CPU cores
- included via flocking (connected HTCCondor pools)
 - GridKa Tier-3: 462 CPU cores
- managed via TARDIS
 - OpenTelekomCloud: 500 CPU cores (during HNSciCloud Project)
 - HPC Cluster ForHLR II @ KIT: 260 CPU cores
 - GridKa School Instance: 800 CPU cores
 - HPC Cluster NEMO @ Freiburg: 8000 CPU cores



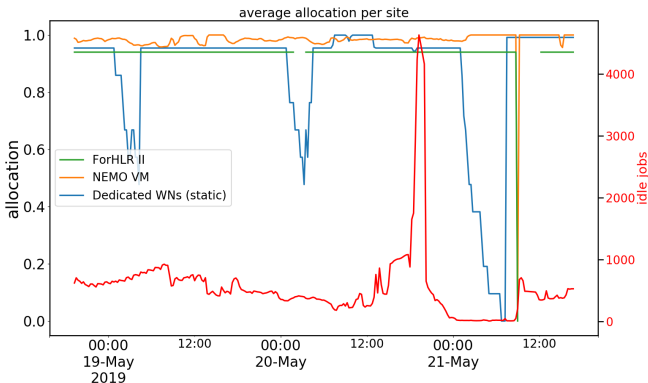
- static
 - usual worker nodes: 264 CPU cores
 - high throughput worker nodes: 84 CPU cores
 - desktop PCs (cycle stealing): 90 - 270 CPU cores
- included via flocking (connected HTCondor pools)
 - GridKa Tier-3: 462 CPU cores
- managed via TARDIS
 - OpenTelekomCloud: 500 CPU cores (during HNSciCloud Project)
 - HPC Cluster ForHLR II @ KIT: 260 CPU cores
 - GridKa School Instance: 800 CPU cores
 - HPC Cluster NEMO @ Freiburg: 8000 CPU cores



Resource Usage at ETP

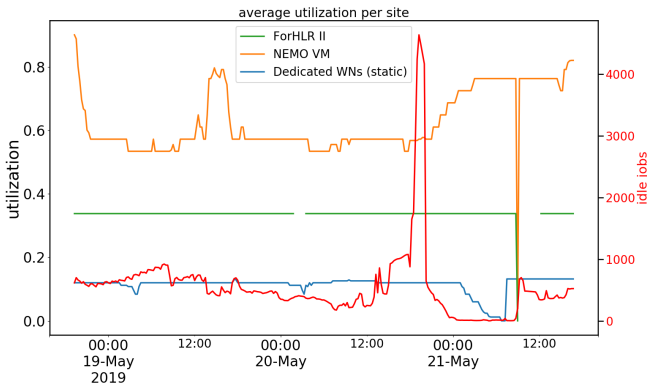


Allocation per Site at ETP



- demand and allocation (CPU, RAM, disk) are connected
- dynamic resources are more allocated than static resources with defragmentation

Utilization per Site at ETP



- low utilization (CPU, RAM, disk) \Rightarrow a lot of resources are unused

- **Telegraf** service queries batch system information via python scripts github.com and send to InfluxDB service
- **Grafana** presents information from InfluxDB

