


LRZ-LMU report – A&C research compound meeting – Feb 2023

- Pledges & Usage
- Person-power
- Procurement plans
- Power costs
- Additional activities

LMU-Team:
O. Biebel
G. Duckeck
A. Lory
C. Mitterer
R. Walker

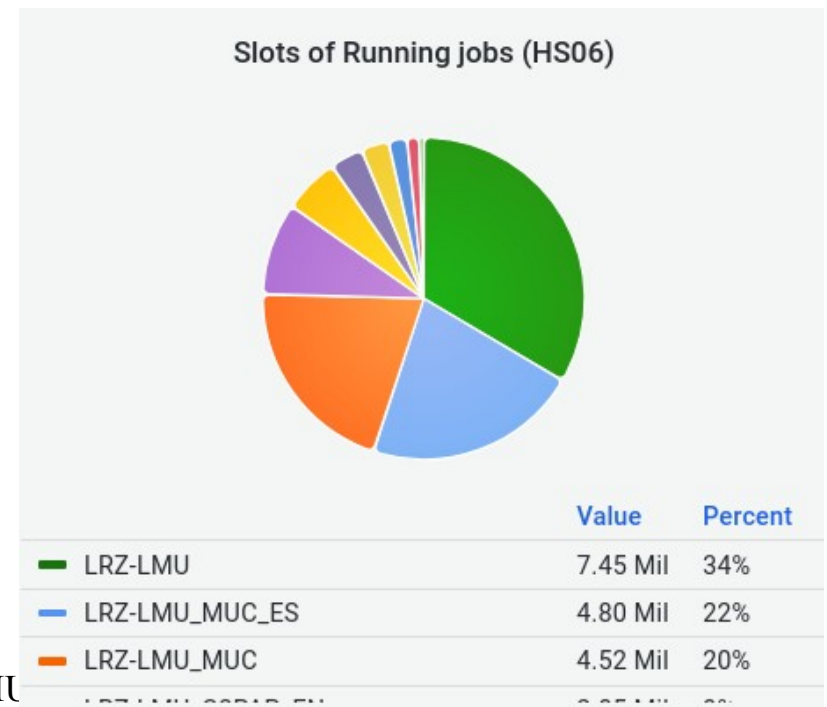
LRZ-LMU: pledges & usage



LRZ-LMU

MPPMU

- Federation w/ partner site MPCDF/MPP
 - Std ATLAS-DE T2 2023 pledges:
29 kHS06 CPU, 2800 TB Disk
- Substantial addtl **opportunistic** CPU resources:
 - T2 (40%), **SuperMuc (42%), C2PAP (12%), LRZ-cloud (3%), ...**
 - Last year CPU: ~2.2 x pledge



LRZ-LMU: person-power

- G. D. (coordination, procurement, ...) – LMU position
- C. Mitterer (dCache & T2 services) – LMU position
- R. Walker (T2 CE, opportunistic resources, ATLAS DPA, ADC technical coordinator) – BMBF core computing
- A. Lory (HammerCloud, HC coordinator since fall 22) – BMBF DE computing support (full position since 03/2022)
 - HammerCloud – coordination, operational support, auto-exclusion check & optimization
- LRZ personnel: sys-admin, hardware install, networking, ...
 - LRZ positions

LRZ-LMU: hardware status

- CPU:
Purchase usually via LRZ Rahmenvertrag for current HPC architecture
 - 44 Lenovo nodes, Xeon E5-2697 from 2015-2017
 - SuperMuc/CoolMuc-2 architecture, ~25 kHS06
 - 18 Megware nodes, Xeon 6230, 2020
 - CoolMuc-3 architecture, ~18 kHS06
- Disk:
Usually self-organized purchase, standard tender procedure
 - 18 HP servers a 10x8TB HDD from 2017
 - ~1400 TB
 - 19 HP servers a 12x16 TB HDD from 2021
 - ~3600 TB

LRZ-LMU: procurement plans

- No procurement so far in current BMBF-FP
- 2023: CPU :
 - Replace Lenovo systems (6-8 y old) + increase: ~40 kHS06
- 2024 Disk :
 - replace old HP + increase: ~3000 TB

LRZ-LMU: power costs

- Power contracts handled by LRZ
 - Long-term contract – only modest increase 2022 → 2023 (<10%)
 - **Though cost 2024+ completely unclear**
 - Cost covered by LMU administration (so far)
- Side-note:
 - From 2016→ 2022 power consumption **halved** & resources **doubled**
- Active project by Rod Walker to reduce CPU clock rate in periods of power shortage/high-cost/high-CO2 (→ [presentation WLCG 2022](#))
 - Implemented on CPU cluster and functionality demonstrated
 - **Rough estimate: for 10% CPU loss one could save up to 25% power-cost**
 - If there would be a dynamic power tariff – so far only dry-run

LRZ-LMU: add-ons

- Substantial use of opportunistic resources since many years
 - SuperMuc-NG: Rather special requirements → dedicated setup
 - Multi-node jobs, pre-empt, GlobusOnline for data IO, ATLAS Harvester, ...
 - Rather variable availability, >10 k cores sometimes
 - C2PAP (CPU cluster for ORIGINS exc-cluster)
 - Purely opportunistic usage, IO via T2 dCache, mix of regular and preempt
 - LRZ openstack cloud and K8s test cluster: ~100 cores
- Storage site for HEPHY-UIBK (Innsbruck)
 - OpenStack cluster ~1500 cores, IO via LRZ-LMU dCache
- XCache: several active servers
 - Integrated in ATLAS VP (virtual placement) for analysis: ~300 slots
 - Regular operation since ~2020

LRZ-LMU: add-ons - 2

- Fidium TA-3 – analysis technologies
 - Ongoing activity related to Iris-HEP Analysis Grand Challenge (D. Koch)
 - Perform standard analysis test-suite (Coffea/Dask) on different cluster environments in DE
 - Tests in progress on
 - Vispa/RWTH, Desy-NAF/Uni-HH,
 - LRZ-LMU: dCache, HDD XCache, SSD XCache, ...
- Considerations to setup OpenStack/K8s cluster
 - Substantial CPU procurement in 2023 → K8s cluster?
 - Similar plans & experience at MPCDF partner site
 - Strong interest & activity on K8s analysis cluster by new ATLAS group at TUM (L. Heinrich)
- Belle 2
 - LRZ-LMU dCache integrated in Belle 2 as resource for local group
 - Using phased-out storage servers