Annual Meeting of the BMBF-funded Research Compound - "Föderiertes Computing für die ATLAS- und CMS-Experimente am Large Hadron Collider in Run-3"

26–27 Mar 2024
Europe/Zurich timezone

Enter your search term

Overview
Timetable
Contribution List

Timetable

Tue 26/03 | Wed 27/03 | All days

Print | PDF | Full screen | Detailed view | Filter

14:00	**Welcome and Introduction**	*Markus Schumacher et al.* 14:00 - 14:10
	Aachen Report	*Thomas Kress* 14:10 - 14:35

My first meeting in Freiburg

Aachen Tier-2 Report

Thomas Kreß for the Aachen Grid-Team
(from the institutes Ib, IIIa, IIIb)

IDIUM
GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
CMS Germany
III. Physikalisches Institut B
RWTH AACHEN UNIVERSITY

Team and Technical Setup

Team:

Alexander Jung (Fidium)
Thomas Kreß
Martin Lipinski (Fed. A&C Computing LHC)
Andreas Nowack (Fed. A&C Computing LHC)
Alexander Schmidt (replaced Achim Stahl as project leader)
… plus sometimes local WiHis

Presently T2+T3:

identical T2 and T3 technical setup
now only DELL server hardware and Nexus/CISCO switches
8784 logical cores = 12.338 HepScore23
- last procurement fall 2023: AMD Bergamo Zen4 with 2*112 cores + 384 GB RAM / node -> HT off
8.2 PiB (net) disks, users' data sets are mirrored
2*40 Gbps LHCOne WAN, internally a mixture of 10G and 1G links

OS, dCache, CEs:

moved from Quattor to Foreman+Puppet fabric tools
still SL 7 Linux, will move to Alma 9 in May, (connected) desktop cluster almost done
dCache version 8.2.37, will update to next golden release soon
HTCondor batch v9 -> version 23.x soon
two HTCondor CEs located and maintained (very well !) by KIT

Site Realibility

WLCG pledges German CMS Tier-2 federation:

2/3 DESY Hamburg, 1/3 RWTH Aachen

pledges always fully available to CMS plus substantial opportunistic (“T3”) resources

CMS‘s 14-day “site readiness” metric:

very reliable, occasionally problems which need local expert intervention, then typically unavailable for 1-2 working days

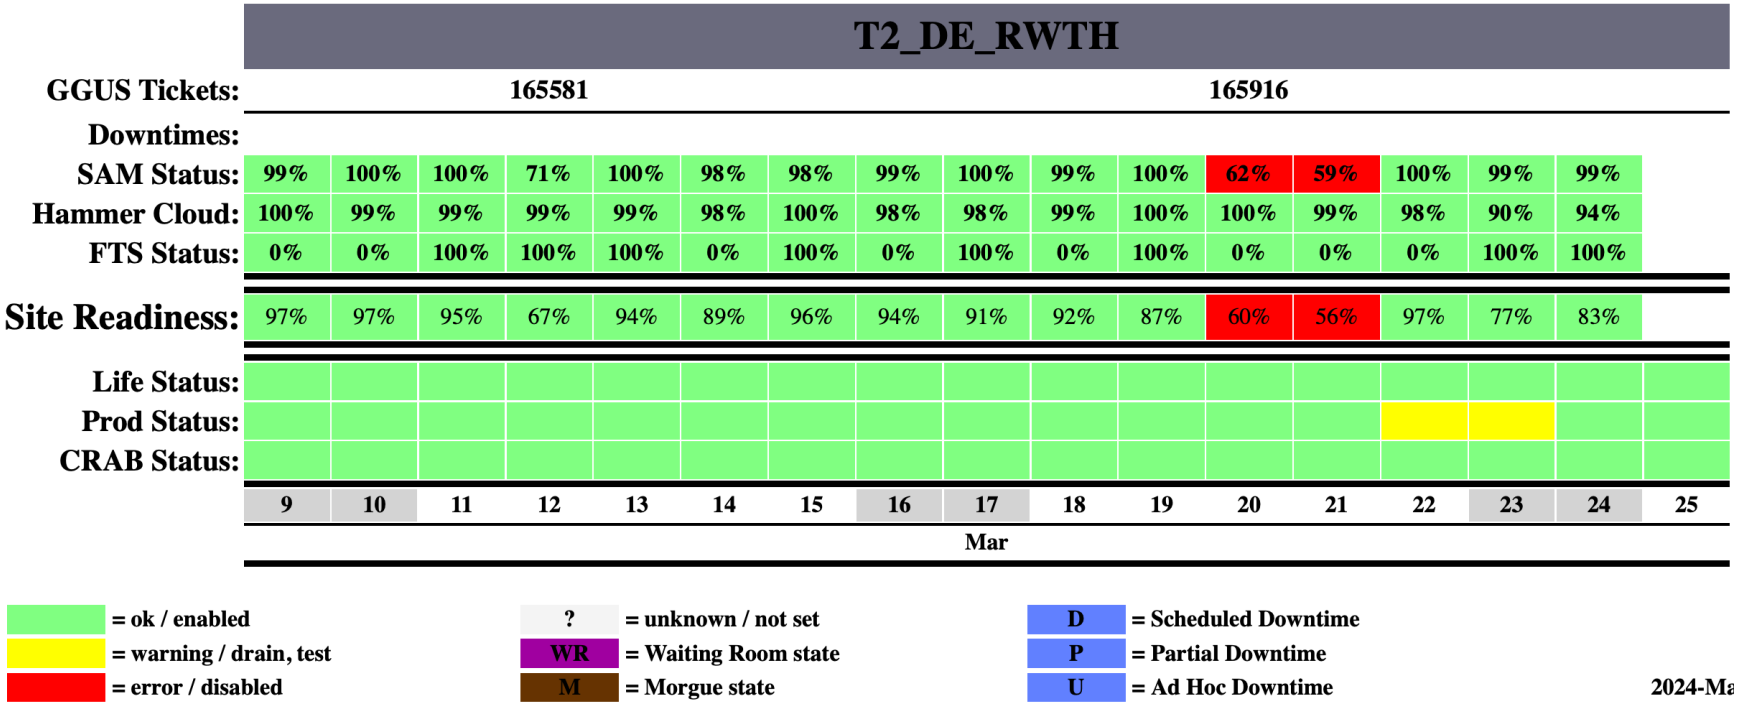

WLCG + Belle2 Data Challenge DC2024

Aim to test (integral):

- 25% of HL-LHC
- ATLAS tested their Tier mesh more heavily than CMS

No Aachen site problems found:

- up to 20 Gbps WAN data transfers sustained with even higher peaks
- experiments did not try to reach sites' limitations

General limitations for CMS:

- FTS servers overloaded
- deletions (not a relevant workflow) slow

Tokens instead of X.509 certs:

- mainly OK

After the DC is before the next ...

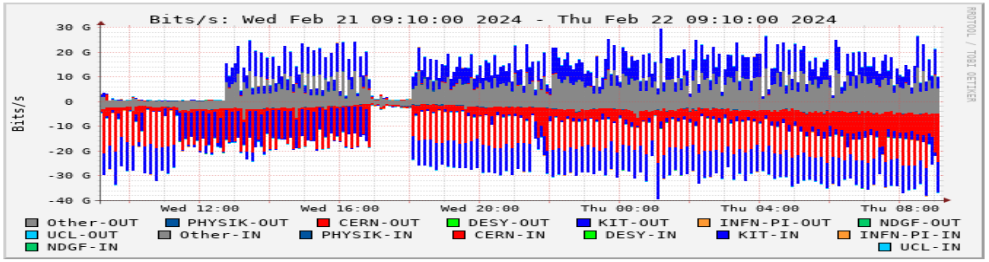

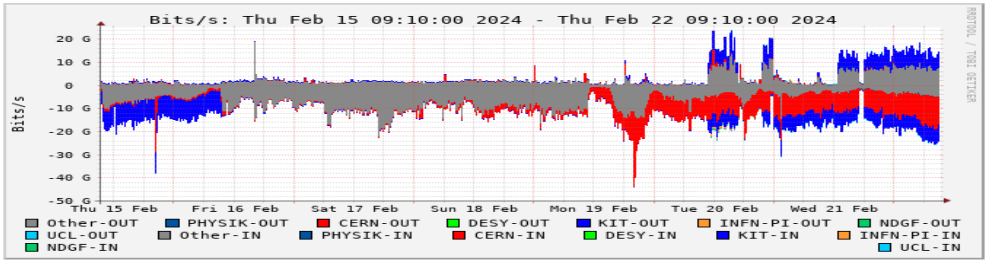

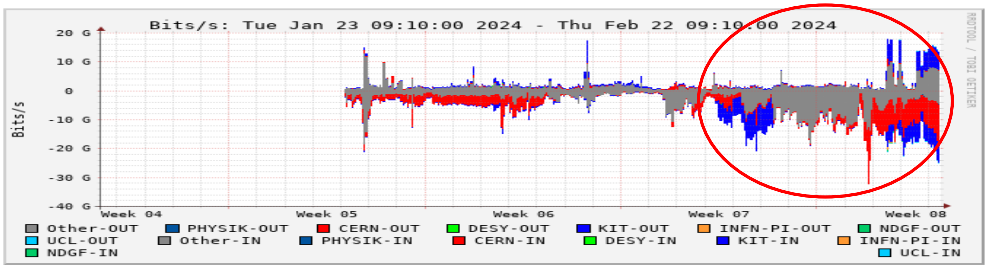

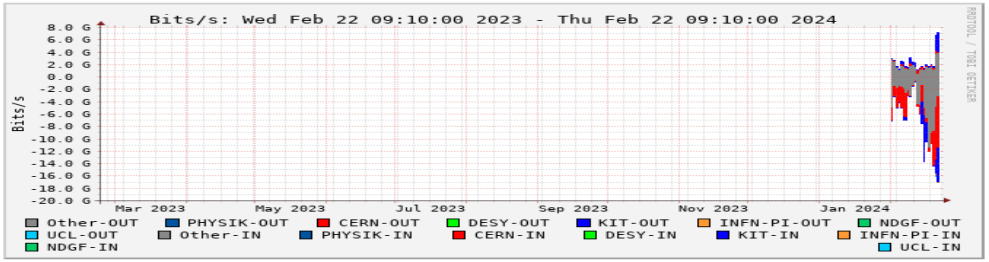

III. Physikalisches Institut B
RWTH AACHEN UNIVERSITY

Pledge and Local Resources' Extrapolation

Fall 2023:

last procurement in recent BMBF funding period

5 years amortization period for T2 hardware resources

From 2025 on +15%" pledge scenario (could/will become higher):

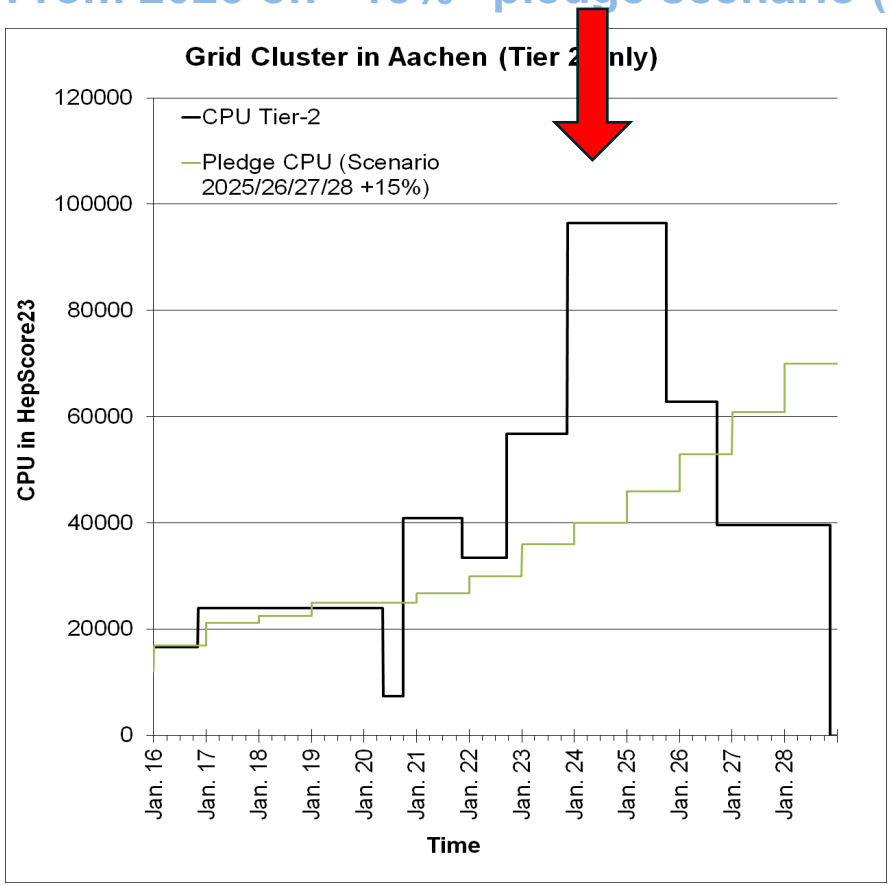

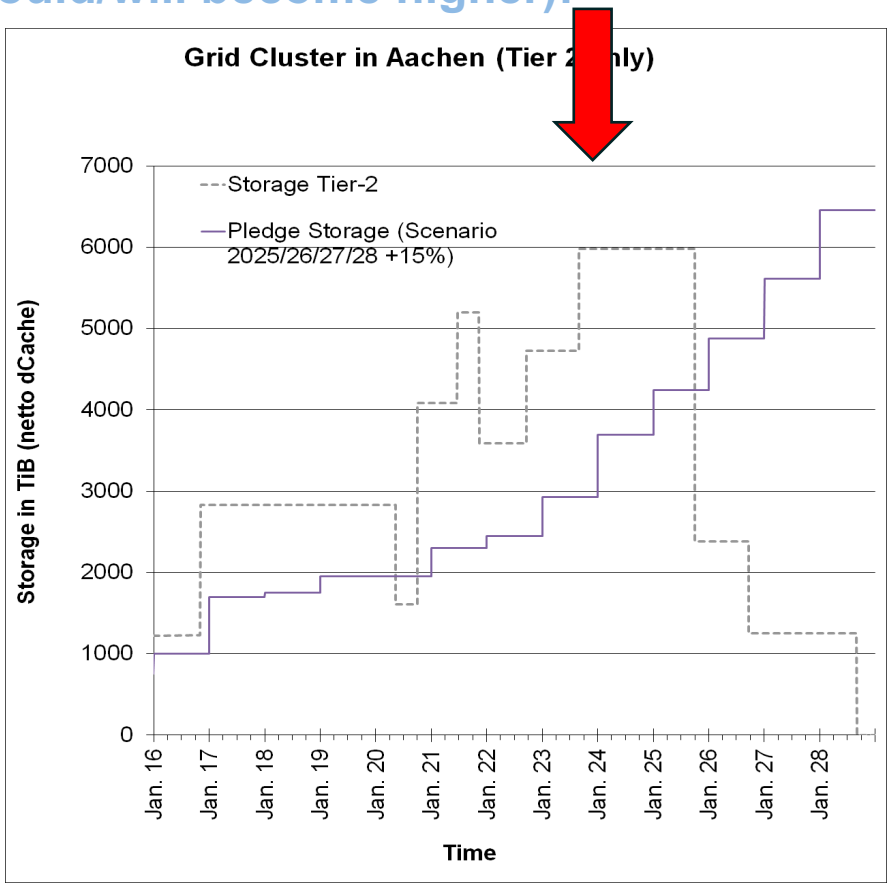

Integration of RWTH HPC CL AIX(-La-Chapelle) Cluster

Using COBalD/Tardis:

- setup 2020 by F. v.Cube and KIT team
- maintained by A. Jung & KIT
- T2+T3+CLAIX transparent for CMS
- dedicated job monitoring possible

Successful integration:

- stress test in Xmas'20 period
- more than 6k cores for 2 weeks
- visible in CMS opportunistic HPC resources

- presently "easy grant" for only one node to constantly test usability
  - CentOS 7 -> Rocky 8
  - MFA for ssh login & commands (still needs final setup)
  - Claix2016/18 -> Claix2018/23+ …

- preparation for 10% NHR pilot application to test full future D-CMS KIT/RWTH/DESY setup

M.Giffels

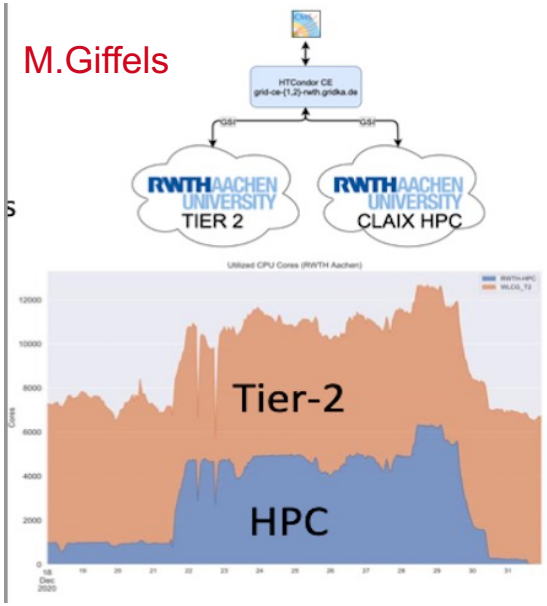

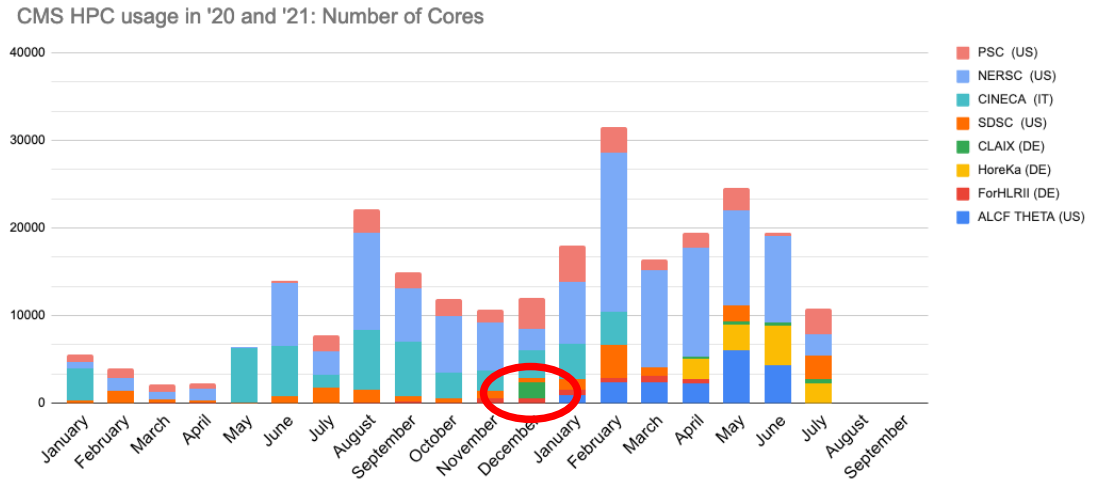

Forthcoming Actions

Pedges 2024:

April 1st official deployment date for T2s, CPUs already used opportunistically by CMS

Software upgrades/migrations:

SL 7 to Alma 9 Linux OS soon

dCache new release soon

of our two remote HTCondor CEs at KIT (probably in summer/fall)

- **change access from X.509 certificates to tokens mandatory when moving to new OS/HTCondor version**
- **implies that users can only use official CMS tools but direct job submission no longer possible (a few users affected with taken-over scripts), since AIM write tokens for USERS not yet on the horizon for CMS**

Other possible changes:

CMS operations pushes to operate 8-core pilots with 8+2 payload cores to increase efficiency

- **concerns by some multi-VO sites (e.g. KIT) to steal resources from other VOs**
- **CMS claims ”on average no problem” … or protectable by cgroups (at least with OS v8/9)**
- **our new AMD Zen4 with many cores and a lot of memory channels do not have much RAM contingency**
- **since our T2/3 is CMS VO only, “stealing from other VOs” not an issue**
- **might be even better for efficiency to move to whole node scheduling (used by all CMS US Tier-2s)**
- **evaluate possible implications for NHR setup**