SPONSORED BY THE

IDIUM

Federal Ministry
of Education
and Research

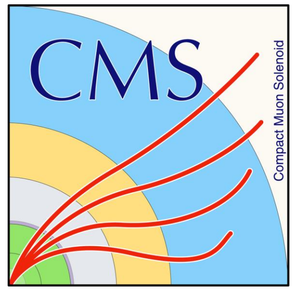

RWTH AACHEN UNIVERSITY

Aachen: Status and Plans

Alexander Jung[1], Thomas Kreß[2], Martin Lipinski[3], **Andreas Nowack[2]**, Alexander Schmidt[1], Shawn Zaleski[1]

[1] III. Physics Institue A
[2] III. Physics Institue B
[3] I. Physics Institue B

FIDIUM Collaboration Meeting
FC-AC Kickoff Meeting

2024-10-01

Aachen CMS Grid Status

- As always, very reliable service to CMS central operation and users
  - 2 weeks site performance overview:

	T2_DE_RWTH															
GGUS Tickets:																
Downtimes:																
SAM Status:	100%	100%	100%	99%	100%	100%	100%	100%	100%	100%	100%	100%	97%	100%	100%	100%
Hammer Cloud:	100%	100%	100%	99%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	100%
FTS Status:	0%	0%	0%	0%	0%	0%	100%	0%	100%	0%	0%	0%	0%	0%	0%	100%
Site Readiness:	99%	99%	96%	99%	100%	97%	90%	96%	95%	100%	97%	94%	96%	99%	95%	97%
Life Status:																
Prod Status:																
CRAB Status:																
	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
	Sep															

= ok / enabled
= warning / drain, test
= error / disabled

? = unknown / not set
WR = Waiting Room state
M = Morgue state

D = Scheduled Downtime
P = Partial Downtime
U = Ad Hoc Downtime

information as of
2024-Sep-20 12:30:06 UTC

https://cmssst.web.cern.ch/sitereadiness/report.html#T2_DE_RWTH

Aachen CMS Grid Status

- Transparent cluster of Tier-2, Tier-3
  - Almost 8000 running cores
  - $2\,PB$ disk space free
  - (`/store/user/` for German CMS users, mirrored once)

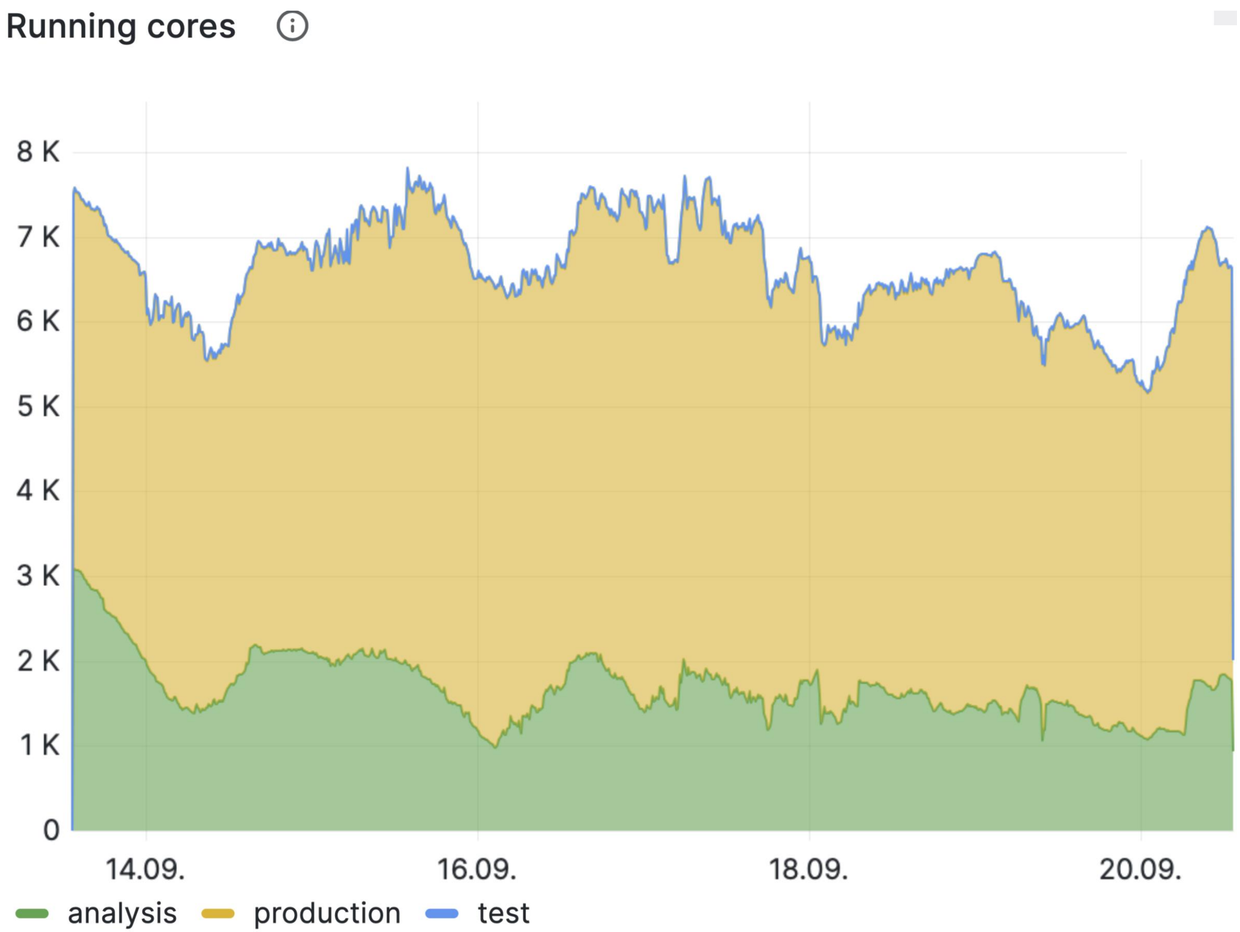

https://monit-grafana.cern.ch/d/o3dI49GMz/cms-job-monitoring-12m?from=now-24h&orgId=11&to=now&var-group_by=Type&var-Tier=All&var-CMS_WMTool=All&var-CMS_SubmissionTool=All&var-CMS_CampaignType=All&var-Site=T2_DE_RWTH&var-Type=All&var-CMS_JobType=All&var-CMSPrimaryDataTier=All&var-adhoc=data.RecordTime%7C%3E%7now-7d&var-ScheddName=All

Aachen Tier-2/3 Status

- Two remote HTCondor CEs at GridKa
  - Very stable
  - Excellent support from CMS@KIT colleagues
- Installed new Grid hardware end of 2023
  - 1568 physical cores (AMD Zen4 Bergamo with HT off)
  - Storage box with $84 \cdot 22\,TB$ disks, $\mathcal{O}(1.6\,PB)$ usable
  - Finally decommission of some $> 10y$ old hardware
- Renewed Tier-2/3 cooling system

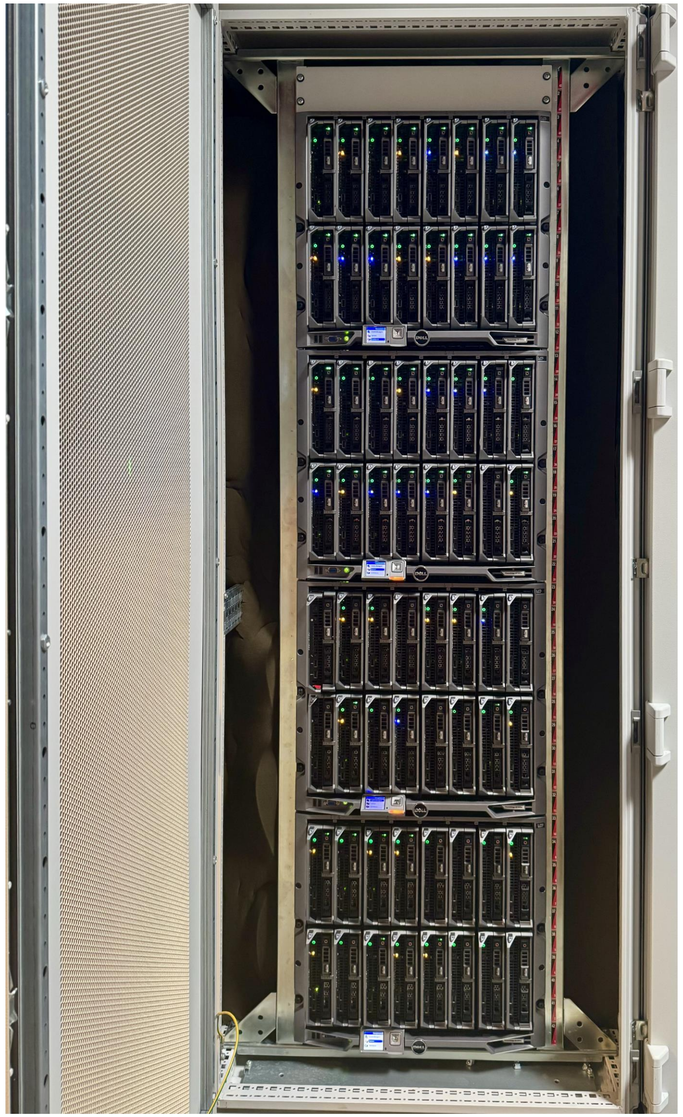

Aachen Tier-2/3 Status

- OS migration almost finished
  - All worker nodes updated to Alma Linux 9 before July 2024
  - dCache and SE nodes still on Scientific Linux 7 (not accessible from outside)
  - Other OS flavours transparently accessible by `/cvmfs` containers
  - Was transparent for CMS operation team and users
- HepScore23 benchmarks:

CPU	Model	Frequency [GHz]	#Logical Cores	Total Score [HepScore23]	Score/Logical Core [HepScore23]
AMD Opteron	6272	2.1	32	178.74	5.59
Intel Haswell	2630v3	2.4	32	315.64	9.86
Intel Broadwell	2630v4	2.2	40	323.68	8.09
AMD Epyc Rome	7452	2.35	128	1745.04	13.63
AMD Epyc Milan	7543	2.8	128	2141.91	16.73
AMD Epyc Bergamo	9734	2.3	224	5764.88	25.74 (HT off)
Intel Xeon	8468	2.1	96	2784.00	29.00 (HT off)

Overview CLAIX @ RWTH

- CLAIX-2023 (NHR-Tier-2 + NHR-Tier-3), in production since April this year
- Same number of cores expected for next upgrade in 2025
- HepScore23 benchmark: 29 (HT off, corrected by +5% due to CVMFS overhead)

Peak performance	• 2.6 + 1.4 PFlops (CPU) • 4.4 + 0.7 PFlops (GPU)
Available resources	• 346 + 185 MCoreHours (CPU) • 27 + 4 MCoreHours (GPU) [1 GPU-h = 24 CPU-h]
HPC segment	**412 + 220 nodes** • 2-socket Intel Sapphire Rapids • Xeon 8468, 2x48 cores, 2.1 GHz • 470 nodes with 256 GB • 160 with 512 GB • 2 nodes with 1024 GB
ML segment	**32 + 5 nodes** • 2-socket Intel Sapphire Rapids • Xeon 8468, 2x48 cores, 2.1 GHz, 256 GB • 4 NVIDIA H100, 96 GB HBM2e per node
Interactive segment	• Additional nodes with smaller GPUs (e.g. for JupyterHub usage)
Fabric	• Infiniband NDR network (OPA) • 2:1 blocking
Storage	• 26 PiB Lustre Storage • BEEOND on SSDs (1.4 TB per node)

Future German CMS Tier-2 Concept

- German ATLAS and CMS Tier-2 concept undergoes major adaption
- Foreseen scenario:
  - BMBF funded Tier-2 hardware at universities (age $< 5y$) still part of WLCG pledge
  - Starting 2025: gradual shift from university to federated resources
    - For storage: Data Lakes at Hamburg and Karlsruhe (Helmholtz sites)
    - For CPUs: NHR centres at Aachen and Karlsruhe (and DESY continues to provide $2/3$ of the total pledge)
  - Hybrid operation during fadeout

Future German CMS Tier-2 Concept

- Implies hybrid operation for next $5y$, then no more (except a few specialised servers) Tier-2 hardware at university sites
- In all scenarios CMS HEP/IT experts indispensable for successful WLCG participation
- Tier-3 resources (mainly extra `/store/user` disk space) still to be provided at universities and/or NAF

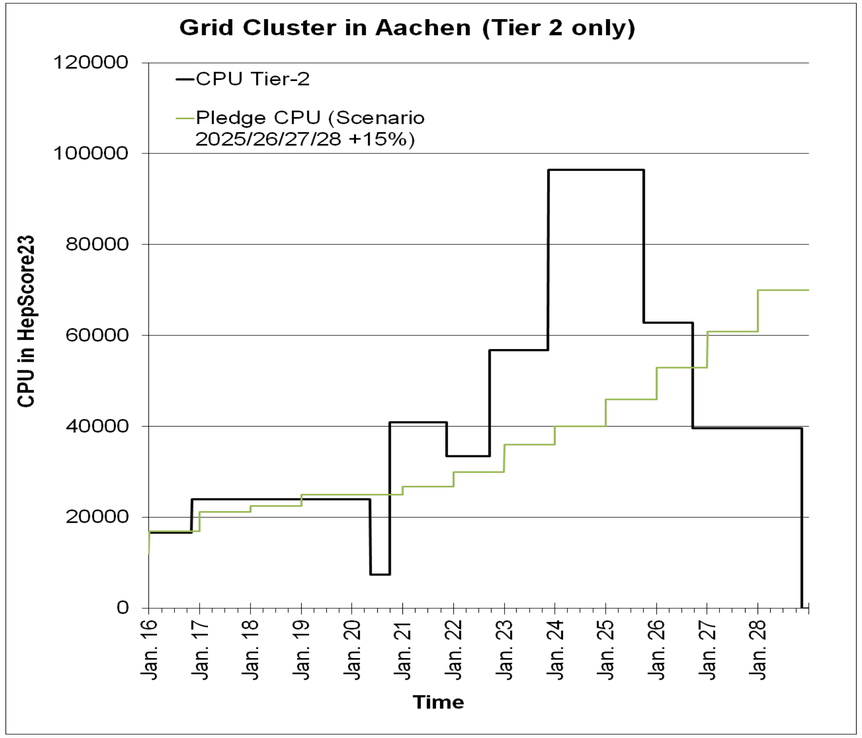

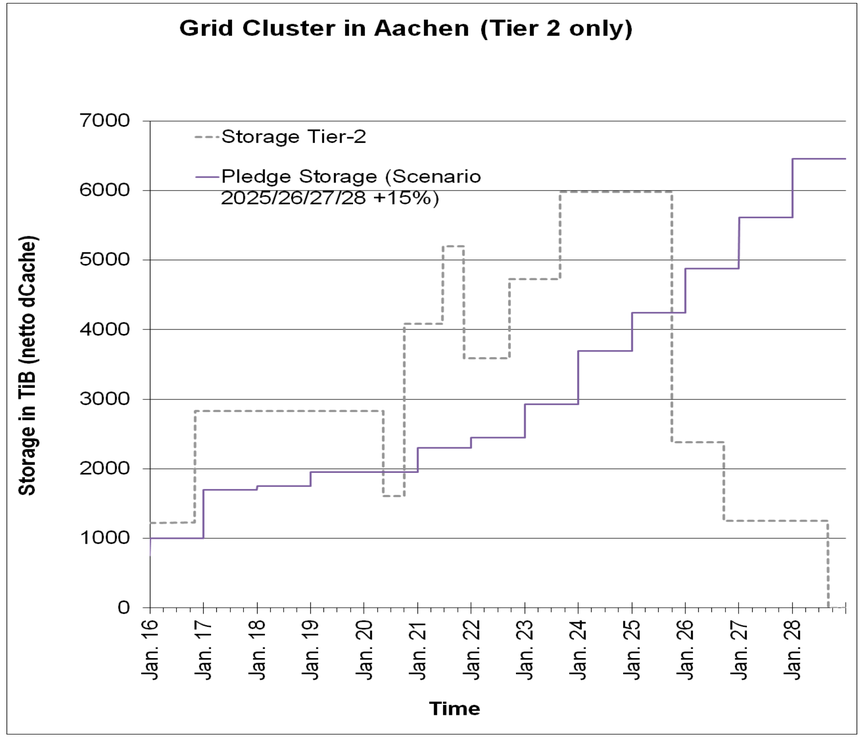

Preparing the NHR application

- Small preparation project with 0.5 million core hours
  - Starting 1st of October, already approved
- Setup proven in production environment
  - Using COBalD/TARDIS
  - Scale test performed successfully $2y$ ago
  - Next step: multi factor authentication (MFA)
    - Modification of COBalD/TARDIS, to be tested
  - Transfer knowledge to new team next month

Preparing the NHR application

- Full NHR application starting April 2025
  - Total 2025 NHR CPU pledge covered by RWTH (CLAIX: about 4 million core hours)
  - Helmholtz fraction 2025: storage pledge covered 50-50 by KIT and DESY ($1.3\,PB$)
  - Starting April 2026: 50-50 split between RWTH and KIT for CPU

Federation	Tier	VO	Country	Year	Type	Pledge
DE-DESY-RWTH-CMS-T2	2	CMS	Germany	2025	CPU	47,500 (RWTH) + 95,000 (DESY) = 142,500 (HepScore23)
DE-DESY-RWTH-CMS-T2	2	CMS	Germany	2025	Disk	3,075 TB (RWTH) + 8,750 TB (DESY) + 650 TB (DESY-Uni) = 12,475 TB
DE-KIT_t2	2	CMS	Germany	2025	Disk	650 TB

https://wlcg-cric.cern.ch/core/pledge/list/

Föderierte Digitale Infrastrukturen für die Erforschung von Universum und Materie

- FIDIUM testing interoperability between Helmholtz, NHR and Tier-2 sites
- Additional upcoming tasks
  - Probably investigating caching
    - dCache satellite setup (control nodes, doors; operated at DESY; SE storage servers locally)
    - XCache

Summary

- Very reliable CMS Grid service at RWTH Aachen
- New hardware and new OS in production use without problems
- In the future NHR RWTH HPC center will provide Tier-2 CPU pledge resources to CMS
  - Using COBalD/TARDIS service developed by KIT
  - Setup already stress tested and transparently integrated in T2_DE_RWTH

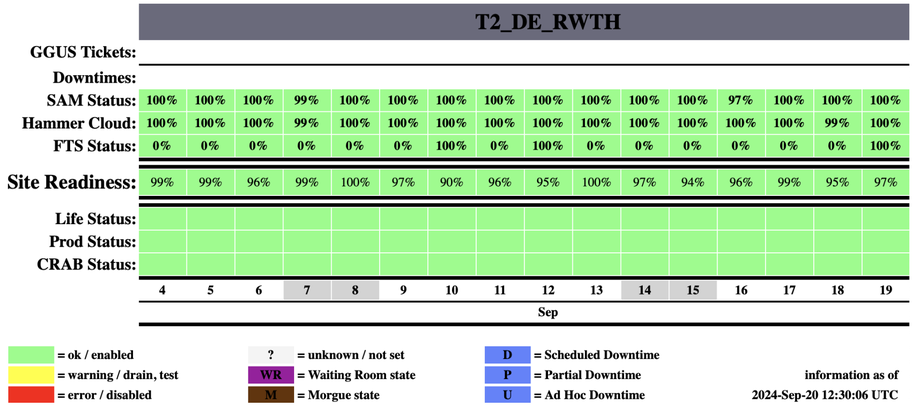

T2_DE_RWTH																
GGUS Tickets:																
Downtimes:																
SAM Status:	100%	100%	100%	99%	100%	100%	100%	100%	100%	100%	100%	100%	97%	100%	100%	100%
Hammer Cloud:	100%	100%	100%	99%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	100%
FTS Status:	0%	0%	0%	0%	0%	0%	100%	0%	100%	0%	0%	0%	0%	0%	0%	100%
Site Readiness:	99%	99%	96%	99%	100%	97%	90%	96%	95%	100%	97%	94%	96%	99%	95%	97%
Life Status:																
Prod Status:																
CRAB Status:																
	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
	Sep															

= ok / enabled
= warning / drain, test
= error / disabled
? = unknown / not set
WR = Waiting Room state
M = Morgue state
D = Scheduled Downtime
P = Partial Downtime
U = Ad Hoc Downtime
information as of 2024-Sep-20 12:30:06 UTC

https://cmssst.web.cern.ch/sitereadiness/report.html#T2_DE_RWTH

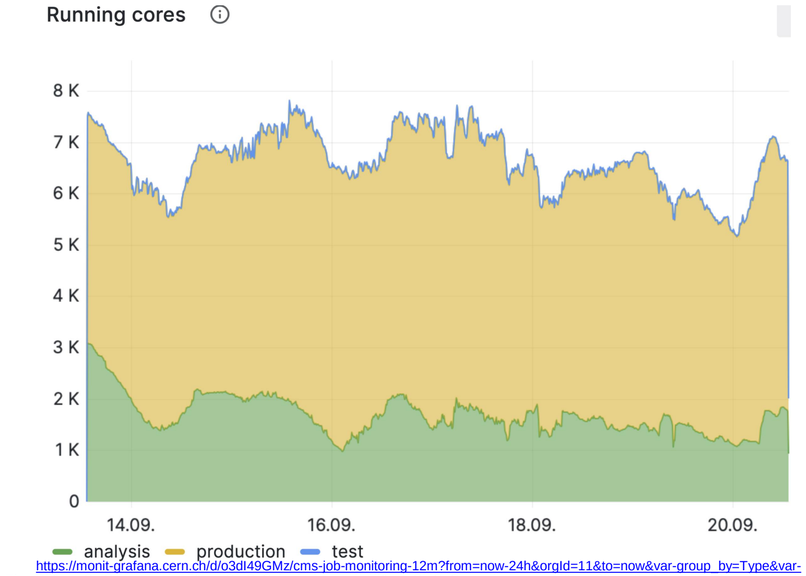

https://monit-grafana.cern.ch/d/o3dl49GMz/cms-job-monitoring-12m?from=now-24h&orgId=11&to=now&var-group_by=Type&var-Tier=All&var-CMS_WMTool=All&var-CMS_SubmissionTool=All&var-CMS_CampaignType=All&var-Site=T2_DE_RWTH&var-Type=All&var-CMS_JobType=All&var-CMSPrimaryDataTier=All&var-adhoc=data.RecordTime%7C%3E%7now-7d&var-ScheddName=All

Backup

Physis Department Linux Cluster

- $\mathcal{O}(100)$ desktops and $\mathcal{O}(20)$ home/data servers
  - HTCondor batch system
  - Also used by Aachen CMS users („low threshold“)
  - Read access from Aachen dCache at Tier-2/3 for Aachen CMS users
  - GPUs (8 ·NVIDIA A40, 6 ·NVIDIA L40 and older GTX)
    - Interactive access to a few GPU servers
    - Access mainly by HTCondor

Physis Department Linux Cluster

NVIDIA Multi-Instance GPU (A100/H100)	Our solution (A40/L40 with HTCondor)
• Expensive	• Much better FP32 performance/price
• 80 GB VRAM	• 48 GB VRAM
• Static (7+1) VRAM partitoning possible (not supported on A40/L40)	• Dynamic partitioning of VRAM by Condor ClassAds
	• Fair-share possible
	• GPUs accessible by `condor_ssh_to_job`

Physis Department Linux Cluster

- OS fully migrated to Alma Linux 9
- Fabric tools: Foreman & Puppet
- Always try to have OS in sync with our Tier-2/3
- Replaced or removed a lot of 4-core AMD Phenom II nodes
  - RHEL9 requires CPU vector instructions in later versions
  - Version 8 still necessary for some older server hardware, because of CPU support
  - Trend goes to ByoD anyway

Föderierte Digitale Infrastrukturen für die Erforschung von Universum und Materie

- FIDIUM collaboration consist of 14 university groups and research centres
- Successor of ErUM-Data IDT
- 3 research areas:
  - Development of tools for heterogeneous resource integration
  - Data lakes, distributed data, caching
  - Adaptation, testing and optimization of production and analysis environments

Motivation

- Estimated luminosity for Run 4 and beyond
- HL-LHC: challenges in the area of data acquisition, processing, simulation and analysis

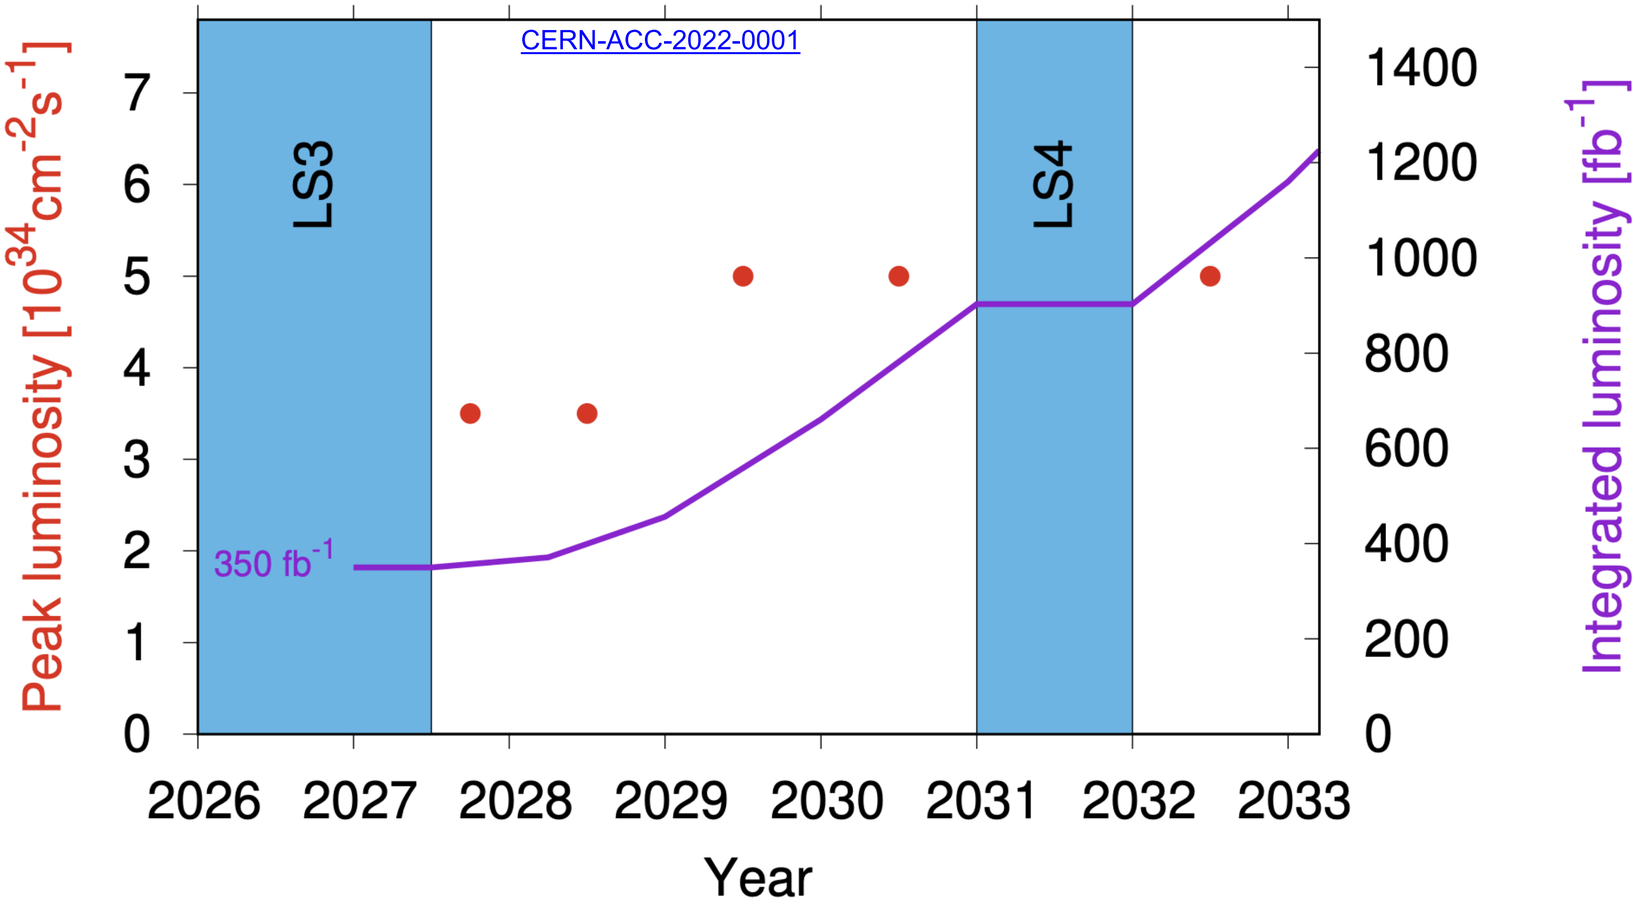

Motivation

- Expected pledged grid resources may no longer be sufficient for Run 4
- Possible solution: additional opportunistic ressources

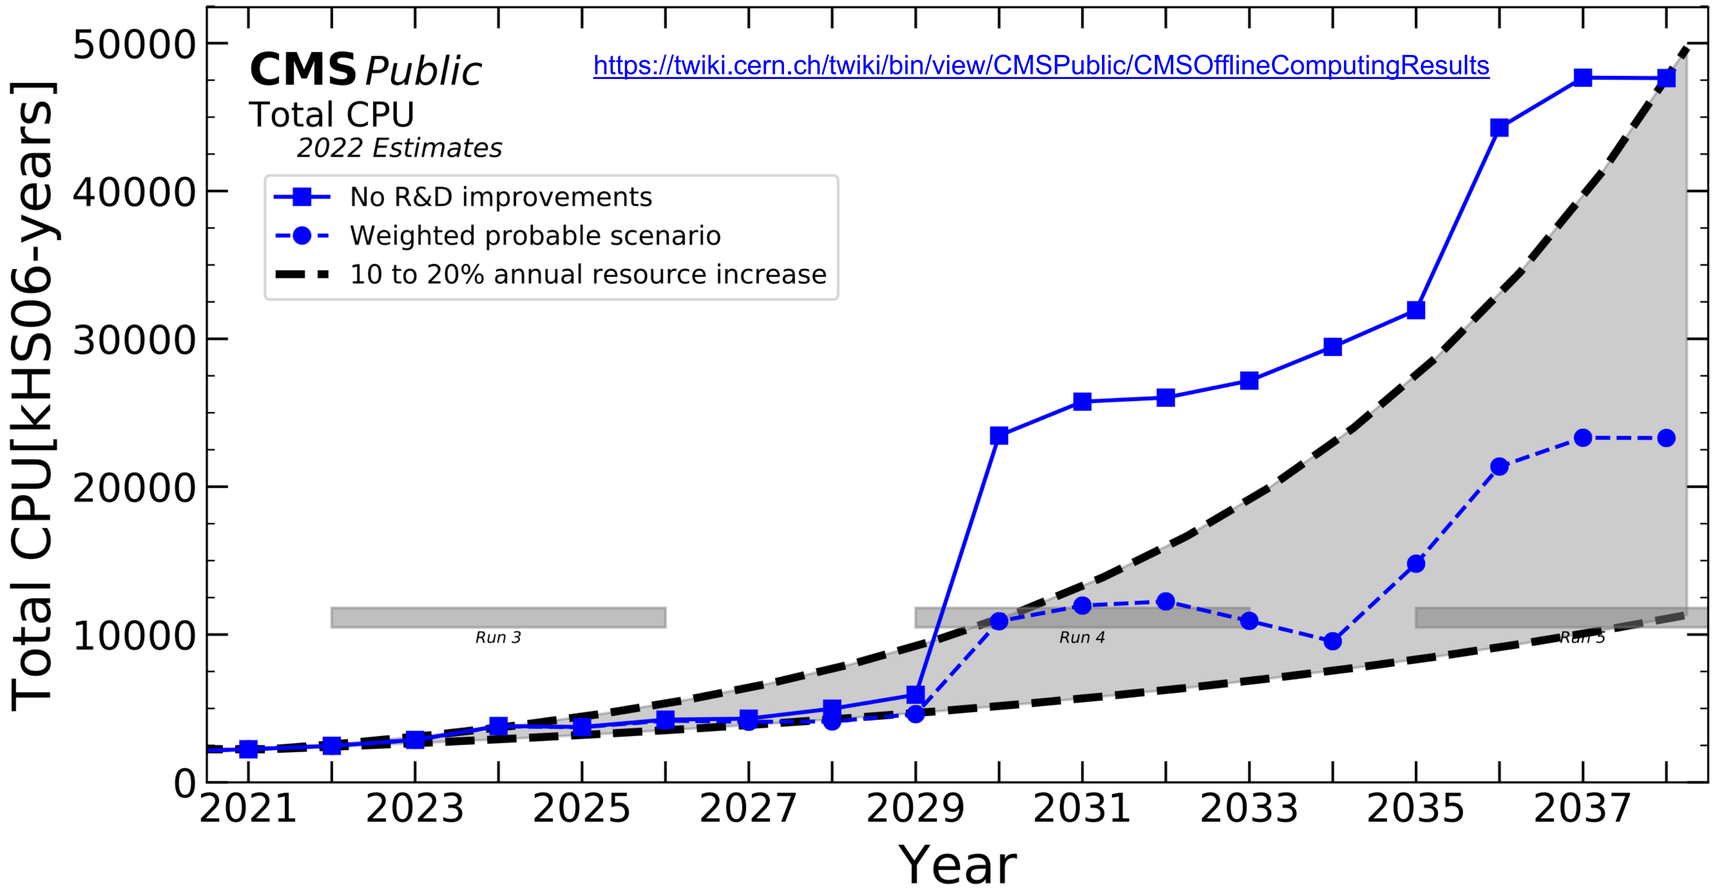

CMS's Expectations for Opportunistic Resources

- Transparent for end users (integration into global CMS collector)
- Dynamic, without knowing/predicting future demand
- Expected WLCG environment is provided

Collector:

- Maintains inventory of all pieces of a pool
- Especially state and available resources of a machine

https://videos.cern.ch/record/2640380

RWTH HPC (CLAIX)

- Running Rocky Linux 8, no problems for CMS noticeable
- At the moment only small grant of one worker node to ensure that CMS jobs and COBalD/TARDIS are running properly, preparation for NHR (see later slides)
- Operational since two years already, already successfully "Xmas" stress tested with $\mathcal{O}(10{,}000)$ cores for about two weeks
- From October 2024: 0.5 million core hours

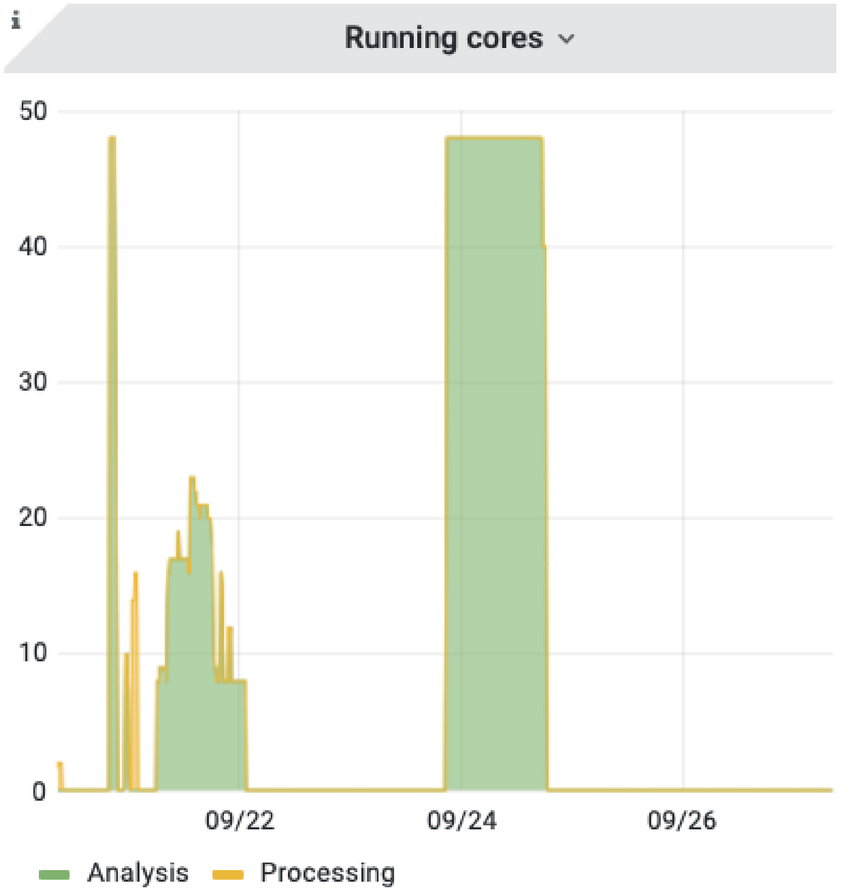

COBalD/TARDIS (developed by KIT)

- Already in use to extend German Tier-1 and Tier-2 to local HPC centres, for example
  - KIT's Tier-1 extended to HoreKa
  - RWTH's Tier-2 extended to CLAIX

} ~ 23 Million core hours provided to CMS in 2022

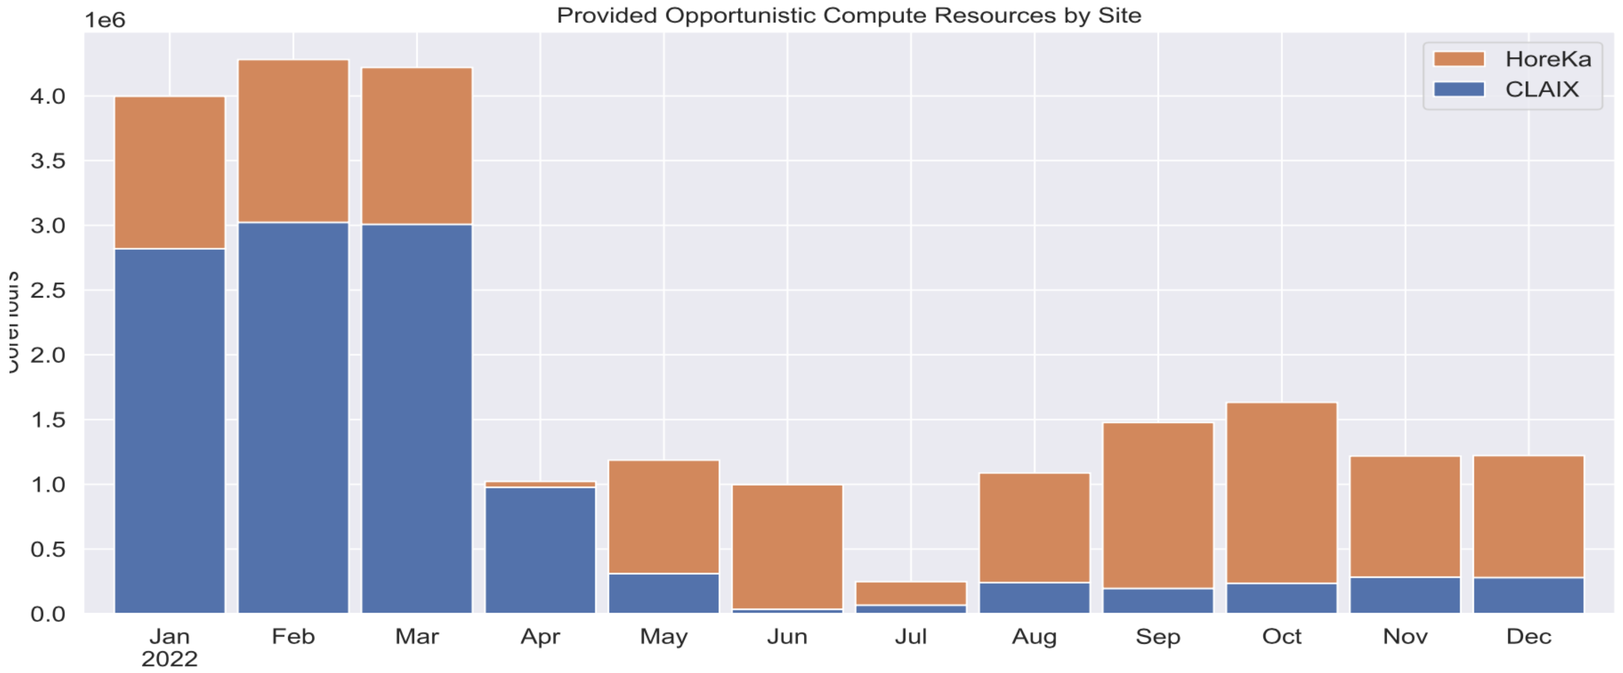

COBalD/TARDIS (developed by KIT)

- COBalD = Opportunistic Balancing Daemon
  - Monitors allocation and utilisation of resources (e.g. RAM, CPUs...)
  - Forms abstract metrics
  - Requests more glideins or lets existing ones expire, based on metrics
- TARDIS = Transparent Adaptive Resource Dynamic Integration System
  - Interface between jobs, external resources,

  batch systems and COBalD

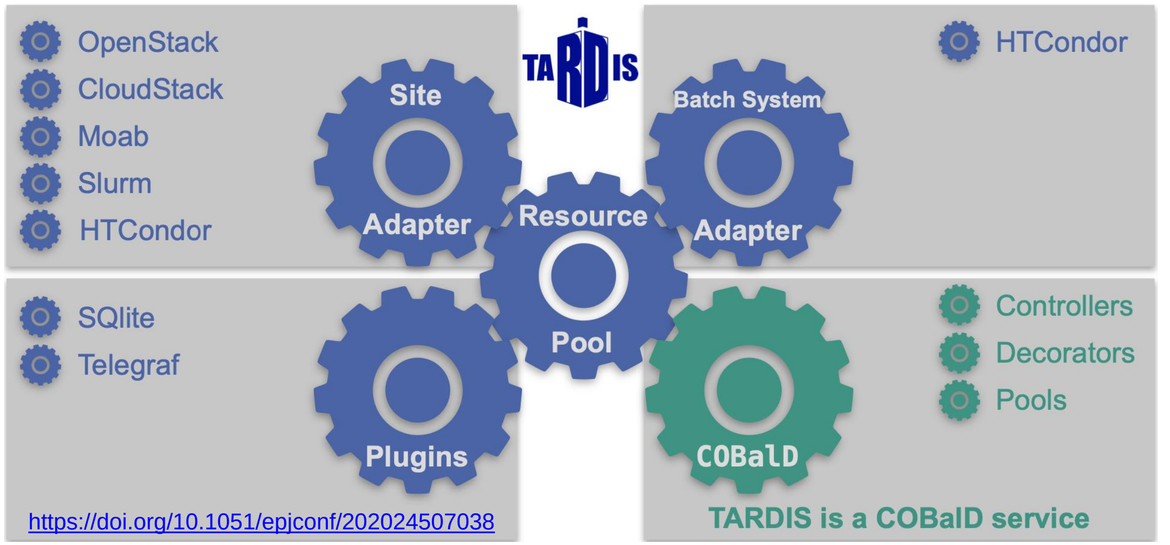

https://doi.org/10.1051/epjconf/202024507038

COBalD/TARDIS (developed by KIT)

- Binds resources
- Sets up expected WLCG environment by using containers
- Integrates resources into WLCG's scheduler
- Dynamically requests resources or lets them expire, depending on the current load
- Completely transparent for CMS users
- Everything happens automatically, no intervention necessary

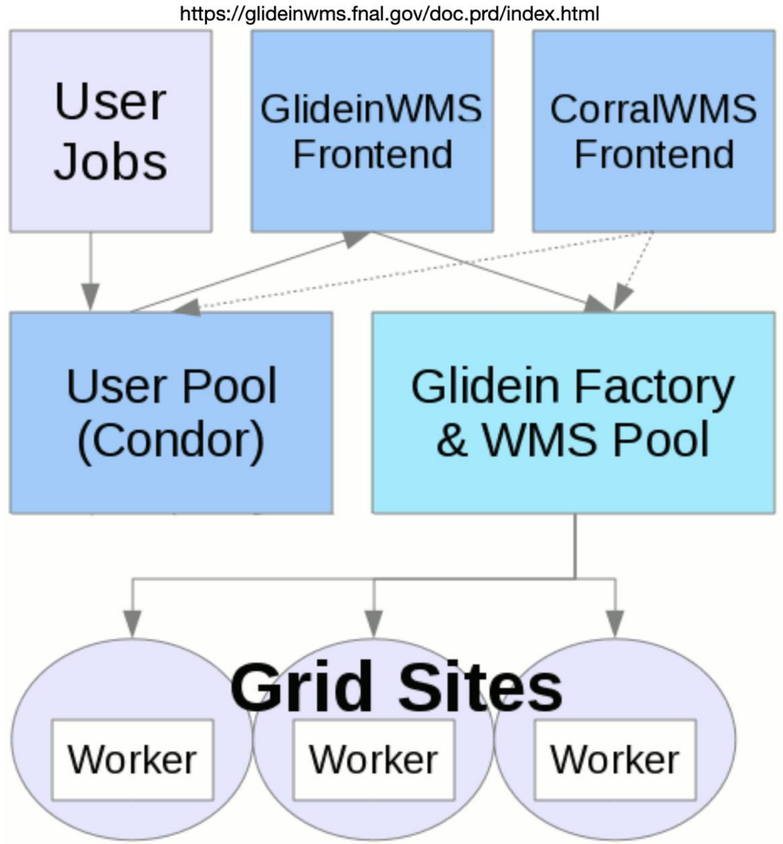

COBaID/TARDIS (developed by KIT)

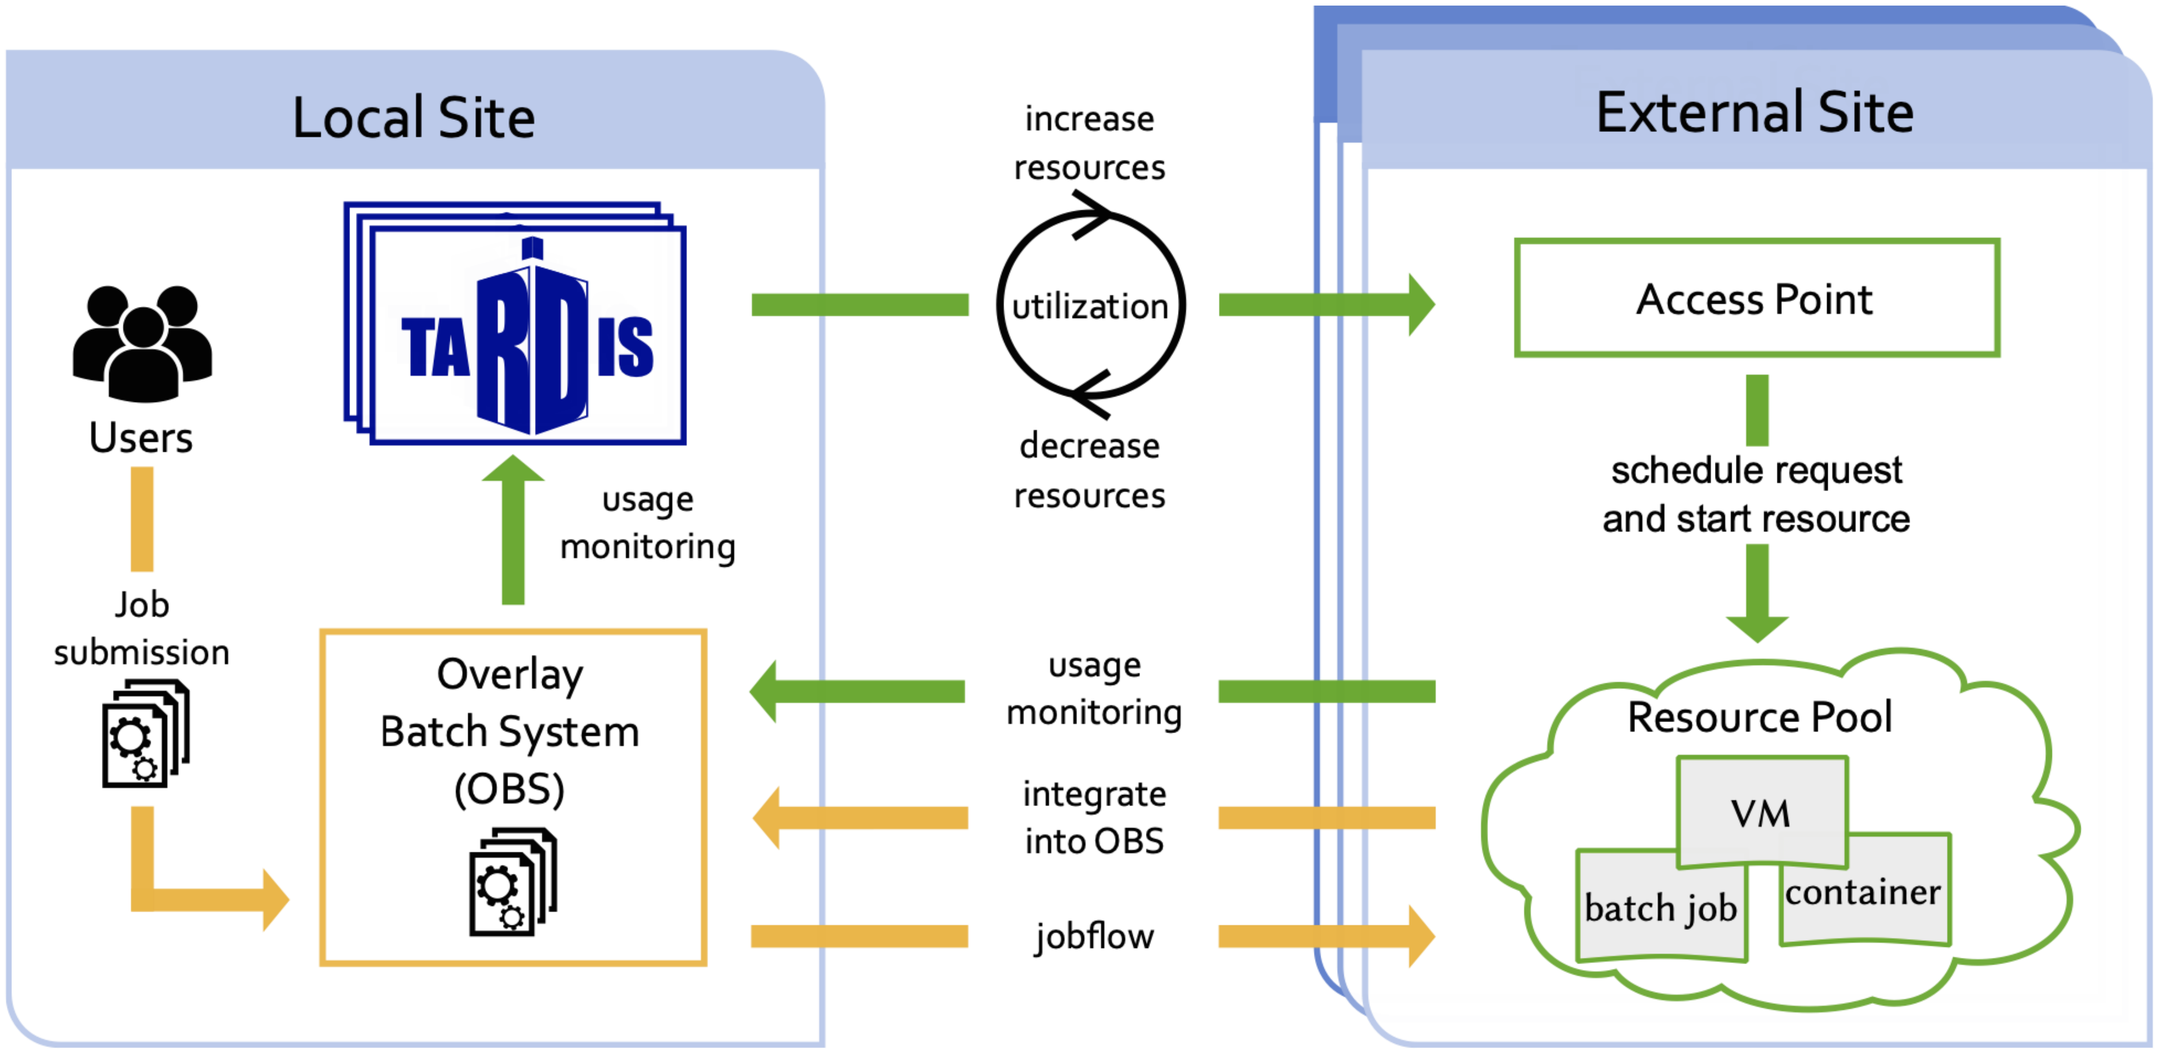

https://doi.org/10.1051/epjconf/202024507038

COBalD's metrics

- Set $\mathcal{I}$ contains resource types (CPU, RAM, …)
- Allocation: measure how full a resource is/whether jobs still fit on it:

$$allocation = \max_{i \in \mathcal{I}} \left(\frac{used(i)}{requested(i)} \right)$$

- Utilization: represents the suitability of the resource for the job (high utilization → good fit):

$$utilization = \min_{i \in \mathcal{I}} \left(\frac{used(i)}{requested(i)} \right)$$

Glidein/Drone/Pilot concept

- Placeholder jobs (=glideins) bind resources
- Environment provided by Apptainer image
- Integration into overlaying batch system (OBS)
- OBS: central management of payload jobs

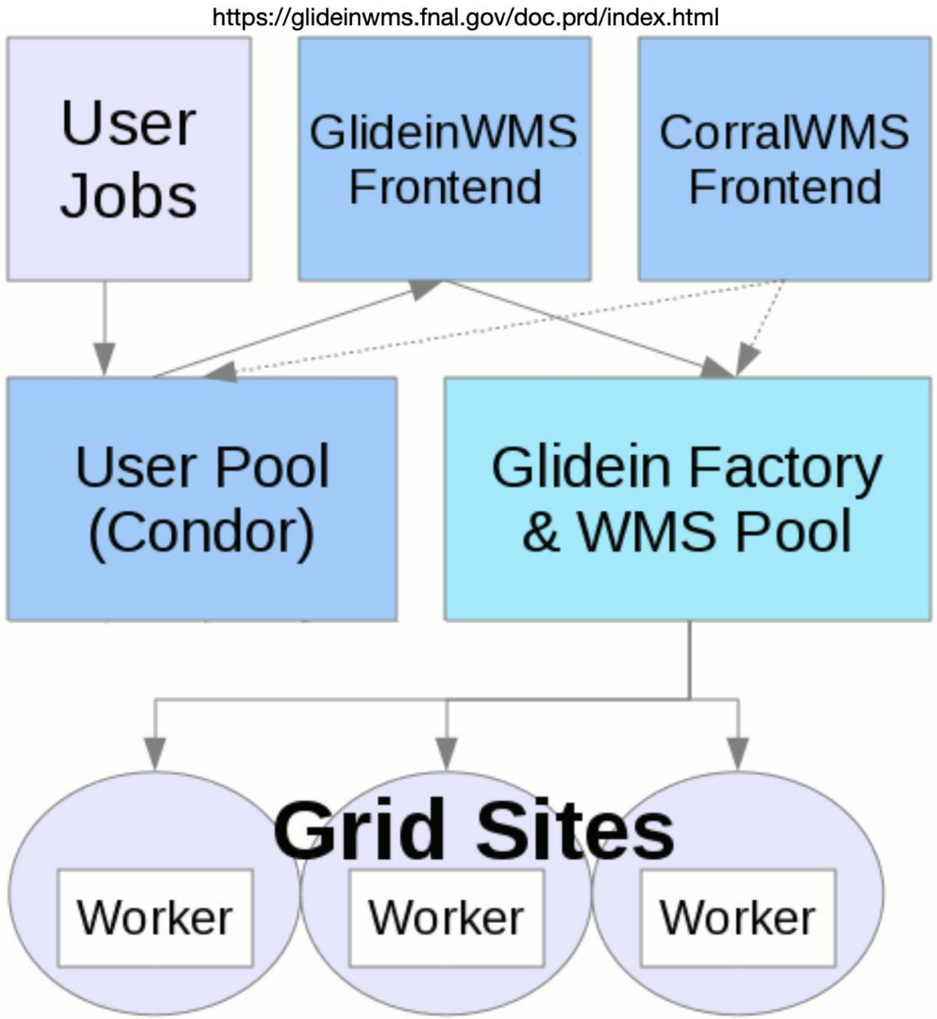

Expected setup by grid jobs

- Potentially access to input data
- Potentially access to conditions database
- CVMFS present to grant access to

experiment's software (CMSSW)

- Connection to external site to store output

CVMFS:

- Cern Virtual Machine

File System

- Read only
- Distribution of software

using http protocol