


Karlsruhe & GridKa CMS Report

Artur Gottmann | February 27, 2023



Topic Overview



- 1 Deployment of the new tape storage system at GridKa
- 2 Integration of data produced at opportunistic HPC resources into CMS data management system
- 3 Provisioning of opportunistic resources connected to GridKa
- 4 Integration of GPU resources at GridKa into CMS grid infrastructure
- 5 Recent developments of the meta-monitoring framework HappyFace
- 6 Maintenance, monitoring and operation tasks

Deployment of new tape storage at GridKa



- Switched to new tape library in March 2022
 - Larger tape capacity: 8 → 20 TB
 - Higher tape drive speed: 150 → 400 MB/s
 - 1 PB disk buffer as part of the system
 - In full operation for CMS, Belle 2, and LHCb
- Data from the old system fully migrated for CMS, Belle 2, and LHCb
- Migration for ATLAS started in June 2022
→ More than a half migrated until now
- Planning to finish migration for ALICE and ATLAS by the end of 2023




top: tape cartridge & drive, bottom: tape library at KIT (current total capacity: 169 PB)

Important milestone for GridKa!

Due to close collaboration between GridKa & experiment representatives (**including ATLAS & CMS**)

Schematic overview of GridKa tape system

- Write request:
 - Incoming file transfer at dCache disk pool
 - Written from dCache to HPSS disk buffer
 - Read back for checksum consistency test
 - Within HPSS, writing to tapes initiated afterwards in **file aggregates**
- Read request:
 - File read requests appear at dCache pool
 - Requests grouped by tape & aggregate
 - **Entire aggregates** read from tapes to HPSS disk buffer
 - Files read from HPSS disk to dCache pool
 - Before sent out by dCache, checksum test is performed for each file




Files < 10 GB in the same directory collected into aggregates of up to 300 GB

[more details](#)

Important fraction of in-house written interface done by **ATLAS & CMS** representatives

Testing tape system during tape challenge 2022


Reading from tapes (07 - 13.03.2022)



- Switch to operation for CMS just in time
- Good opportunity to test the system and fix remaining issues

Writing to tapes (14 - 28.03.2022)

- Deployed additional improvements
→ Load-balanced operation of dCache pools
- CMS didn't push our tape system to the limits



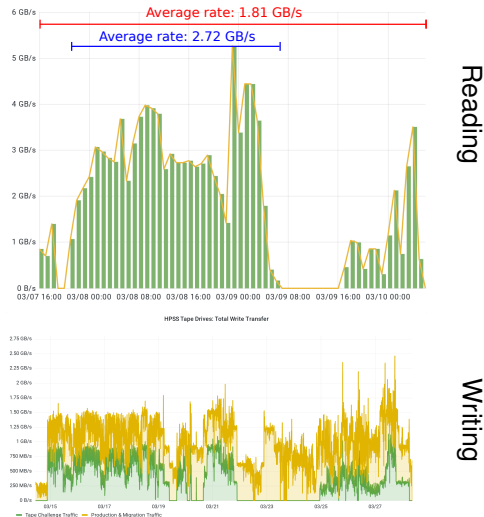
Performance for CMS at tape challenge 2022



- Achieved high total rates during normal operation
→ 2.7GB/s reading, 1.5 GB/s writing
- Reading/writing rate per drive at 300 - 400 MB/s
→ as expected

Take-away from the challenge

- New tape system performs well
- Adjustments successfully applied during the challenge
- Good opportunity to test the limits of the system



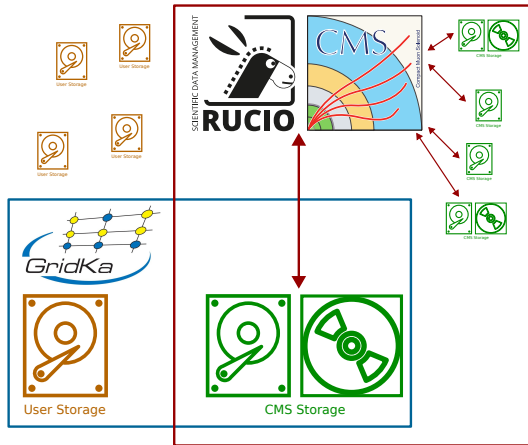
Grid storage concepts for CMS

User storage:

- Usually disk storage, unmanaged, very limited, and **without** a backup
→ not suitable for long-term preservation
- Prominent use-case: intermediate space for analysis output

Central CMS storage:

- Both disk & tape storage available
- Centrally managed** by CMS via **Rucio**
→ Possibility to replicate to tape & other sites



Problem

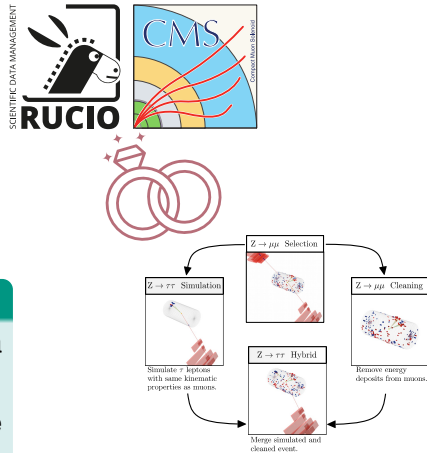
How should **validated** user data provided to the full CMS collaboration be secured & preserved?

Elevation of user data to CMS-managed storage



First large-scale show-case for the problem:

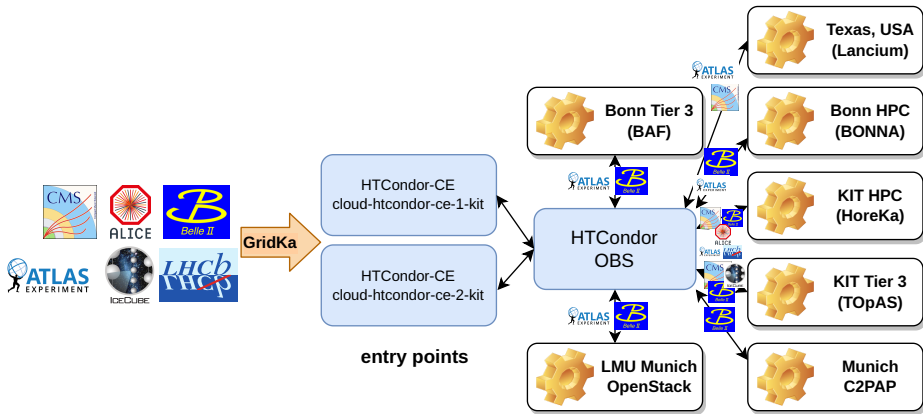
- τ -embedded datasets produced at opportunistic HPC resources at KIT (in total about 30 TB)
- τ -embedding method published in [JINST 14 \(2019\) 06, P06032](#) after careful validation
- Since then, [cited](#) in publications & used in CMS analyses → Contribution relevant to the **entire** CMS collaboration



Solution to backup & secure these datasets

- With the CMS data management team, we agreed upon a [procedure](#) to put τ -embedded datasets on **CMS storage**
- All datasets are now on disk & tape at GridKa, and can be replicated to any other site!


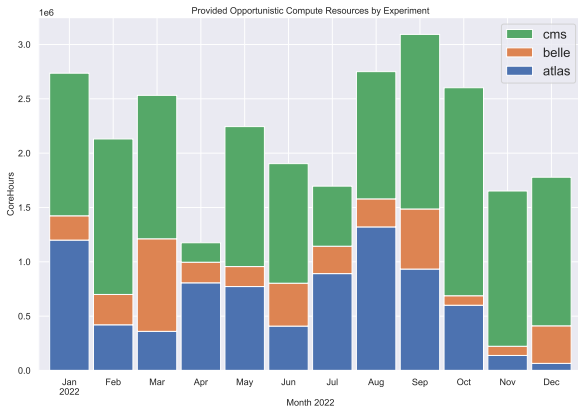
Opportunistic resources connected to GridKa



While fulfilling pledges made by experiments,

Additional opportunistic resources provided to experiments as a **natural** extension of GridKa

Usage of opportunistic resources in 2022




Amount of CPU hours in 2022 \approx **50%** of a usual contribution from a German ATLAS Tier 2 site!

Can you spot the difference? → Less hours in Sep 2022 in the "By Site" plot due to data with a missing site tag ;-)

GPUs at GridKa made available for CMS



- Several GPU's deployed at GridKa TOpAS cluster and provided to entire CMS through the grid
 - 24 × Nvidia A100
 - 24 × Nvidia V100
 - 8 × Nvidia V100S
- GPU workflows sent by CMS were successfully completed
 - High Level Trigger Test Workflow
 - Release Validation Workflow
- In addition, CMS users can access our GPU's from grid with CMS analysis tools

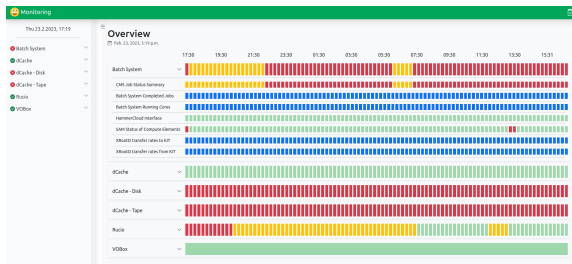


Conclusion

We are well prepared for heterogeneous computing era!

Recent developments of HappyFace

- Ideas behind monitoring with HappyFace:
 - Summarize & correlate information from different sources at **one place**
 - Provide details for **expert usage**
 - A comprehensive view to **non-experts**
 - **Fast and simple** development
- Recent HappyFace 4 started in 2018
→ Successfully operated since several years
- Continuous developments by students & post-docs
- In 2022, **several new modules** introduced



Don't panic


- **Red** \neq GridKa broken!
→ Should rather have a look at the issues
- **All red and yellow** issues known & understood :-)

Use-case at ETP institute

Important tool to monitor CMS at GridKa!

HappyFace: monitoring of CMS jobs

Overview on completed jobs at GridKa & subsites



Details on failed jobs including their properties & a protected link to log files

[Details on failed jobs](#)

Show entries Job Type: 4 items selected Error Type: 4 items selected Search:

Date	Job Type (Campaign)	Error Type (Exit Code)	Computing Element	Host	Pilot Id	Global Job Id
23.2.2023, 20:02:00	Production RunIISummer20UL18wmLHEGEN	Other 99108	htcondor-ce-1-kit.gridka.de	c01-011-115	18461539	vooms0256.cern.ch/#29693.98#1677164241
23.2.2023, 19:55:57	Production RunIISummer20UL18wmLHEGEN	Other 99108	htcondor-ce-2-kit.gridka.de	c01-014-111	12134645	vooms0256.cern.ch/#29696.92#1677164244
23.2.2023, 19:55:08	Processing UNKNOWN	Other 120	cloud-htcondor-ce-2-kit.gridka.de	hkn0829.localdomain	1346073	cmsgwms-submit3.fnal.gov/#541011.72#1677170667

[link to the module](#)

Maintenance, monitoring and operation



Maintenance, monitoring and operation - **bread & butter work** of an administrator.

At GridKa, it includes:

- Continuous monitoring of all GridKa components, using primary sources, and **HappyFace**
→ Organisation of 24/7 shifts at ETP, on-call duties at SCC, etc.
- In case of incidents:
 - Reporting at ETP, to administrators & experts at GridKa/SCC, or to the CMS computing team
→ **Very important**: make experts at CMS or at GridKa aware of problems **proactively**
 - If possible, resolving the issue by ourselves to prevent more complications
 - Example: Summer 2022 was **too hot to handle** for the CPU machines at GridKa
→ Thanks to fast reaction of GridKa experts, nothing bad happened, since machines were off in time
- Active participation in events at the computing center → (Un-)planned downtimes at GridKa, CERN tape challenges, benchmark tests, deployment of hardware & software, etc.
- Working on tickets from CERN: **27 closed CMS GGUS tickets** between Oct 2021 and Dec 2022

Made possible with a **good team** of shifters & **close communication** with experts at CMS & GridKa!

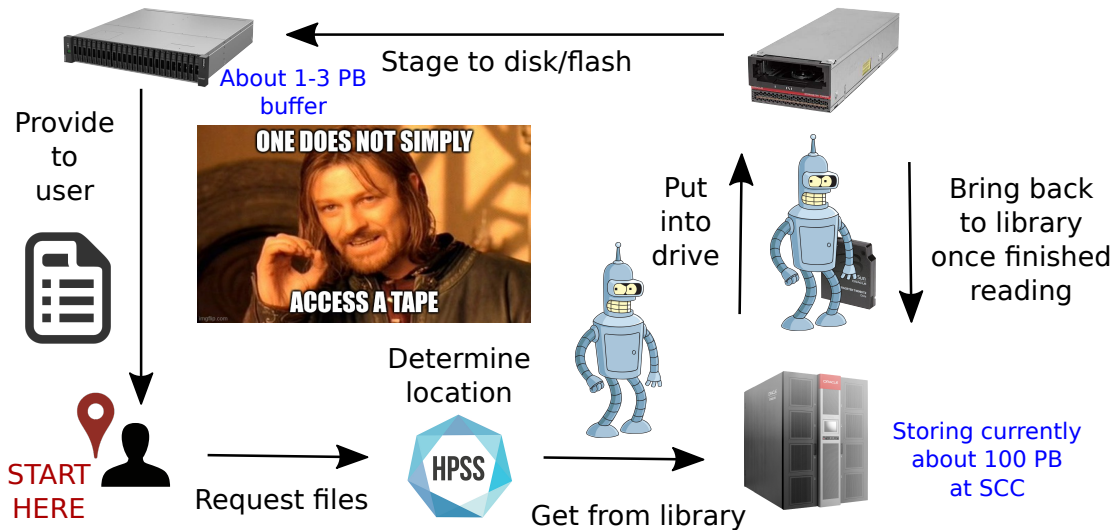
Summary



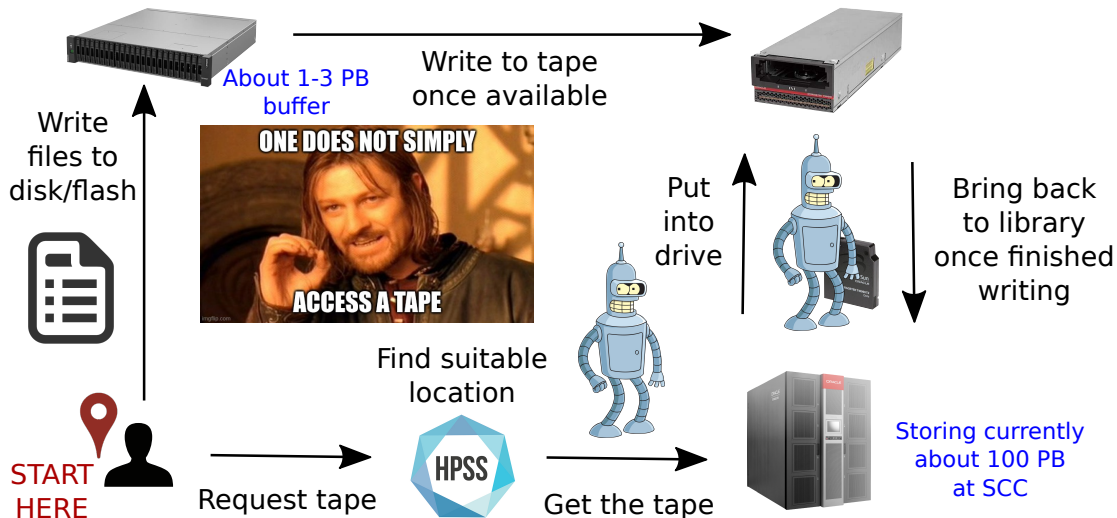
- Year 2022 was eventful for us at KIT & GridKa → a lot of milestones achieved, a lot of work done
- Important for that was close collaboration of everyone involved:
 - Students (from Bachelor to PhD) taking part in developments
 - Team of shifters (PhD students & post-docs) looking at GridKa 24/7
 - Experiment representatives at GridKa to make contact between experiment & site
 - Team of administrators & experts at GridKa, and the computing teams of the experiments
 - Everyone from higher management (professors, department leaders, etc.) setting the conditions
- Looking forward to the year 2023!

Backup

Reading from Tape



Writing to Tape



Run 2 UL τ -embedded datasets at CMS storage



Thanks to the Rucio development team, datasets are in the CMS data management system now

- Prerequisites:
 - CRIC (group:pog_tau) & Rucio (pog_tau_group) accounts, both linked with a CERN e-group (cms-embedding-pog-tau)
 - Set quota for this group at a site in Rucio (together with site admins). Our case: T1_DE_KIT_Disk
 - Setup dataset schemes (naming, blocks, USER datatier, file paths) that comply CMS conventions
- Provide a list of (remote) file paths to the files per dataset, where they are currently stored
- Let CMS O&C admins register the dataset in Rucio. Done in case of τ -embedding:
`rucio list-dids --filter 'type=CONTAINER' group.pog_tau_group:*`
- Files copied to /store/group/rucio/<rucio-group-account-name> as defined by Rucio rules:
`rucio list-rules --account pog_tau_group`
- Afterwards, datasets can be replicated to other sites or tape storage. Our case: T1_DE_KIT_Tape

Glimpse on Rucio content



```
> rucio list-dids --filter 'type=CONTAINER' group.pog_tau_group:*
```

```
+-----+
| SCOPE:NAME | [DID TYPE] |
+-----+
| group.pog_tau_group:/EmbeddingRun2016-HIPM_B_ver1/ElMuFinalState-inputDoubleMu_106X_ULegacy_miniA0D-v1/USER | DIDType.CONTAINER |
...
+-----+
```

```
> rucio list-rules group.pog_tau_group:/EmbeddingRun2018B/MuTauFinalState-inputDoubleMu_106X_ULegacy_miniA0D-v1/USER
```

ID	ACCOUNT	SCOPE:NAME	... STATE[OK/REPL/STUCK]	RSE_EXPRESSION ...	CREATED (UTC)
ca9299715c6a4eef9385de7b0c640a37	pog_tau_group	group.pog_tau_group:/EmbeddingRun2018B/...	OK[117/0/0]	T1_DE_KIT_Disk ...	2022-12-16 16:16:16
36e419b3daf043c9ac4f22137fac5b6d	pog_tau_group	group.pog_tau_group:/EmbeddingRun2018B/...	OK[117/0/0]	T1_DE_KIT_Tape ...	2023-01-12 16:52:42

GPU HLT Test Workflow



Request: CMSSW_12_3_0__PdmVHLTestGPU_MBTEST

Days since request approval 301.3
Production status done
McM status undefined ([request details](#))
Number of events requested 7,618,804
Number of events produced 0 (0.0%)
Primary output
Unified [Log](#)
JIRA [Find tickets](#)

Outputs

	Dataset name	Status	Events ▾	Size	DBS First	DBS Last
1	/MinimumBias0/CMSSW_12_3_0_KIT-123X_dataRun3_HLT_reval_v3_PdmVHLTestGPU_MBTEST-v1/FEVTDEBUGHLT	VALID	7,618,804	2,886,567,050,179	2022.05.13 11:22	2022.05.13 11:22
2	/MinimumBias9/CMSSW_12_3_0-123X_dataRun3_HLT_reval_v3_PdmVHLTestGPU_MBTEST-v1/FEVTDEBUGHLT		0	0		

Workflows

	Workflow name	Age (days) ▲	Current State (days)	Prepld	Type	ReqManager Status	Unified Status	nEvents	Priority
1	kskovpen_CMSSW_12_3_0HLTTestGPU_MBTEST_220426_134810_6554 (ReqMgr , Transferor , Request Details , WorkQueue , ErrorReport , summary)	301.3	284.4	CMSSW_12_3_0__PdmVHLTestGPU_MBTEST	ReReco	normal-archived		7,618,804	120,000
2	kskovpen_CMSSW_12_3_0HLTTestGPU_MBTEST_220415_152621_9803 (ReqMgr , Transferor , Request Details , WorkQueue , ErrorReport , summary)	312.2	301.2	CMSSW_12_3_0__PdmVHLTestGPU_MBTEST	ReReco	aborted-archived		0	120,000

GPU ReVal Workflow



Request: CMSSW_12_6_0_pre4__Run3MC_KITGPU-TTbar_14TeV-00003

Days since request approval 104.3
Production status done
McM status undefined ([request details](#))
Number of events requested 1,000
Number of events produced 1,000 (100.0%)
Primary output [/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-PU_125X_mcRun4_realistic_v2_GPUv3-v1/MINIAODSIM](#)
Unified [Log](#)
JIRA [Find tickets](#)

Outputs


	Dataset name	Status	Events	Size	DBS First	DBS Last
1	/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-125X_mcRun4_realistic_v2_GPUv3-v1/GEN-SIM	VALID	1,000	2,145,299,389	2022.11.09 15:56	2022.11.09 15:56
2	/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-PU_125X_mcRun4_realistic_v2_GPUv3-v1/GEN-SIM-DIGI-RAW	VALID	1,000	75,571,621,467	2022.11.09 15:58	2022.11.09 15:58
3	/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-PU_125X_mcRun4_realistic_v2_GPUv3-v1/GEN-SIM-RECO	VALID	1,000	102,683,096,467	2022.11.09 16:00	2022.11.09 16:00
4	/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-PU_125X_mcRun4_realistic_v2_GPUv3-v1/MINIAODSIM	VALID	1,000	281,722,818	2022.11.09 16:00	2022.11.09 16:00
5	/RelValTTbar_14TeV/CMSSW_12_6_0_pre4-PU_125X_mcRun4_realistic_v2_GPUv3-v1/DQMIO	VALID	0	64,296,688	2022.11.09 16:00	2022.11.09 16:00

Workflows

	Workflow name	Age (days)	Current State (days)	Prepid	Type	ReqManager Status	Unified Status	nEvents	Priority
1	pdmvserv_RVCMSSW_12_6_0_pre4TTbar_14TeV_GPUv3_221109_110455_5970 (ReqMgr , Transferor , Request Details , WorkQueue , ErrorReport , summary)	104.3	104.1	CMSSW_12_6_0_pre4__Run3MC_KITGPU-TTbar_14TeV-00003	TaskChain	normal-archived	done	1,000	500,000

HappyFace: monitoring CMS data requests

Overview on queued data volume to GridKa tape system



Details on individual requests their properties & a protected link to request page

[show current rucio rules](#)

Show entries Activity: 3 items selected State: 2 items selected Only Problematic Search:

Rule ID	Dataset	Activity	State	Account	Created	Last Updated
g3f89ad6- c7f78-4a21- af3a- 12d156bf63d7	/LambdaBToPsi12sLambda_Psi12sToJPsiPiPi1_Psi1ToMuMu_TuneCP5_13TeV-pythia8-evtgen/RunIISummer20UL17REC0-186X_mc2017_realistic_v0-v1/AODSIM	Production Output	Stuck	wmcore_output	14.1.2021, 19:58:00	23.2.2023, 20:24:15
389ad09c- 3faf-4713- 8cc1- 41f382014805	/NoBPTX/Commissioning2023-v1/RAI	T0 Tape	Replicating	tier0_prod	8.2.2023, 13:41:23	23.2.2023, 20:21:18

[link to the module](#)

HappyFace: monitoring CMS transfer errors



Overview on transfer error types (incoming to GridKa tape) grouped by rating & remote sites

Summary of the last hours Search: T2

Error Type	Remote RSE	Critical	Warning	OK
DESTINATION [17]	T2_CH_CSCS	5	0	0
DESTINATION [17]	T2_US_MIT	3	0	0
DESTINATION [17]	T2_BR_SPRACE	3	0	0
TRANSFER [5]	T2_US_Caltech	0	1	0
TRANSFER [110]	T2_ES_IFCA	0	0	2
SOURCE [13]	T2_PT_NCG_Lisbon	0	0	2

[Details on incidents](#)

Details on individual errors including their properties, failure reason & a protected link to log files

[Details on incidents](#)

Show entries Activity: 2 items selected Remote RSE: 8 items selected Search:


Time	Activity	Remote RSE	Failure Reason	File Name with Transfer Link
23.2.2023, 14:04:16	Production Output	T2_US_Caltech	TRANSFER [5] SOURCE CHECKSUM MISMATCH User defined checksum and source checksum do not match 57727ae9 != 00000000	/store/mc/RunIISummer20UL17MiniAODv2/NMSSM_XToYHTo2G2B_MX-4000_MY-1800_TuneCP5_13TeV-madgraph-pythia8/MINIAODSIM/106X_mc2017_realistic_v9-v2/80000/56437CB6-1575-A74F-B9FC-39FBDDF47AD1.root

Showing 1 to 1 of 1 entries (filtered from 22 total entries) Previous **1** Next

[link to the module](#)

HappyFace: monitoring of CMS jobs

Overview on completed jobs at GridKa & subsites



Details on failed jobs including their properties & a protected link to log files

[Details on failed jobs](#)

Show 10 entries Job Type: 4 items selected Error Type: 4 items selected Search:

Date	Job Type (Campaign)	Error Type (Exit Code)	Computing Element	Host	Pilot Id	Global Job Id
23.2.2023, 20:02:00	Production RunIISummer20UL18wmLHEGEN	Other 99108	htcondor-ce-1-kit.gridka.de	c01-011-115	18461539	vooms0256.cern.ch/#29693.98#1677164241
23.2.2023, 19:55:57	Production RunIISummer20UL18wmLHEGEN	Other 99108	htcondor-ce-2-kit.gridka.de	c01-014-111	12134645	vooms0256.cern.ch/#29696.92#1677164244
23.2.2023, 19:55:08	Processing UNKNOWN	Other 120	cloud-htcondor-ce-2-kit.gridka.de	hkn0829.localdomain	1346073	cmsgwms-submit3.fnal.gov/#541011.72#1677170667

[link to the module](#)

GGUS Tickets



27 of 27 Tickets

Ticket-ID	Type	VO	Site	Priority	Resp. Unit	Status	Last Update	Subject	Scope
159911	cms	FZK-LCG2	urgent	NGI_DE ▶ involved	closed	2022-12-29	SAM WN-xrootd-access failing at ...	WLCCG	
159853	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2023-02-14	CMS data world accessible at T1_DE_KIT	WLCCG	
159482	cms	FZK-LCG2	top priority	NGI_DE	closed	2022-12-02	Aptainer issues at T1_DE_KIT	WLCCG	
159157	cms	FZK-LCG2	very urgent	NGI_DE	closed	2022-10-21	subsite tag in site-local-config.xml at ...	WLCCG	
159151	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned	closed	2022-10-20	Massive production failures at T1_DE_KIT	WLCCG	
158491	cms	FZK-LCG2	less urgent	NGI_DE	closed	2022-09-01	Testing CMS tickets	WLCCG	
158488	cms	FZK-LCG2	urgent	NGI_DE	closed	2022-09-07	Testing CMS feature	WLCCG	
158380	cms	FZK-LCG2	urgent	NGI_DE	closed	2022-08-22	Testing CMS tickets	WLCCG	
157773	cms	FZK-LCG2	very urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-07-04	SAM tests for SRM and WEBDAV failing ...	WLCCG	
157367	cms	FZK-LCG2	top priority	NGI_DE ▶ assigned	closed	2022-06-21	Deletion of Tape Challenge Spring 2022 ...	WLCCG	
157202	cms	FZK-LCG2	urgent	VOSupport	verified	2022-05-17	Removing 2017 and 2018 RAW data from ...	WLCCG	
157163	cms	FZK-LCG2	top priority	NGI_DE	closed	2022-05-24	Migration of JINR_Tape files to KIT_Tape	WLCCG	
156488	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned	closed	2022-04-07	Destination Overwrite error at ...	WLCCG	
156380	cms	FZK-LCG2	urgent	NGI_DE	verified	2022-03-16	Enabling token authorization on CEs for ...	WLCCG	
156274	cms	FZK-LCG2	top priority	NGI_DE	closed	2022-04-21	Tape Challenge Spring 2022 KIT_Tape	WLCCG	
156216	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-03-16	MAKE_PARENT Error on T1_DE_KIT_Tape	WLCCG	
155871	cms	FZK-LCG2	urgent	NGI_DE	closed	2022-02-18	File exists and overwrite not enabled - ...	WLCCG	
155678	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-02-21	File exists and overwrite not enabled - ...	WLCCG	
155491	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-01-31	"Incomplete configuration" ...	WLCCG	
155362	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-02-03	Transfers failing to T1_DE_KIT	WLCCG	
155256	cms	FZK-LCG2	urgent	NGI_DE ▶ involved	closed	2022-01-05	File exists and overwrite not enabled - ...	WLCCG	
155188	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2021-12-21	SAM dataset at T1_DE_KIT_Tape	WLCCG	
155003	cms	FZK-LCG2	urgent	NGI_DE ▶ assigned ▶ involved	closed	2022-04-13	Testing Tape access via srm+https KIT	WLCCG	
154823	cms	FZK-LCG2	urgent	NGI_DE	verified	2021-11-05	Pilots at T1_DE_KIT	WLCCG	
154663	cms	FZK-LCG2	less urgent	NGI_DE	closed	2021-11-10	TEST TICKET	WLCCG	
154513	cms	FZK-LCG2	urgent	NGI_DE	verified	2021-10-22	SITECONF storage.json out-dated at ...	WLCCG	
154246	cms	FZK-LCG2	urgent	NGI_DE	closed	2021-11-15	Tape writing test at KIT_Tape	WLCCG	