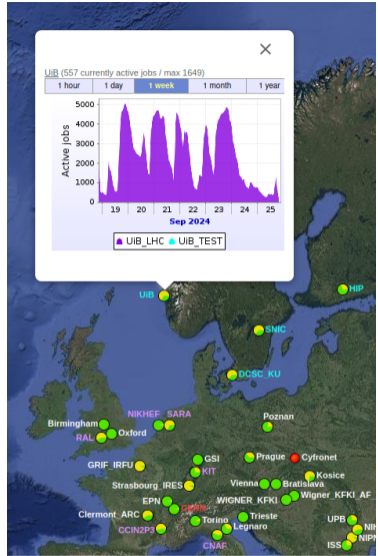


UiB site status report

Topics


- Grid site compute
- dCache Disk storage
- dCache Tape storage
- Network
- Monitoring



<http://alimonitor.cern.ch/map.jsp>

Grid site compute

Total wall time for ALICE jobs



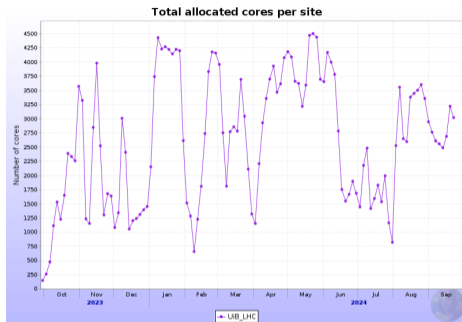
- installed: 325 virtual nodes with 16 cores
→ 5200 cores
- pledged: 55 kHS06 (target 38 kHS06)
ALICE is using ~ 60% of this (33 kHS06)
- all instances upgraded to Almalinux 9
- Site configured for JAliEn multicore jobs

ALICE is working on the following issues (progressing slowly) :

- too few jobs over long period → ALICE has to find balance between simulation jobs, pre-filtering at CERN, data analysis and other type of jobs
- large load on network, the 10 Gbit/s link to Bergen is often exhausted
- inherently inefficient jobs running the old analysis framework

⇒ We are in a transition phase

Grid site compute



- installed: 325 virtual nodes with 16 cores
→ 5200 cores
- pledged: 55 kHS06 (target 38 kHS06)
ALICE is using $\sim 60\%$ of this (33 kHS06)
- all instances upgraded to Almalinux 9
- Site configured for JAliEn multicore jobs

ALICE is working on the following issues (progressing slowly) :

- too few jobs over long period → ALICE has to find balance between simulation jobs, pre-filtering at CERN, data analysis and other type of jobs
- large load on network, the 10 Gbit/s link to Bergen is often exhausted
- inherently inefficient jobs running the old analysis framework ⇒ We are **still** in a transition phase

UiB dCache disk pool status

node001_dcachec_lhc_uib_no_Domain	251658240	42254589	
node002_dcachec_lhc_uib_no_Domain	251658240	22644398	
node003_dcachec_lhc_uib_no_Domain	251658240	89497282	
node004_dcachec_lhc_uib_no_Domain	251658240	84005385	
node005_dcachec_lhc_uib_no_Domain	251658240	90658469	
node006_dcachec_lhc_uib_no_Domain	251658240	62600156	
node007_dcachec_lhc_uib_no_Domain	251658240	57627042	
node008_dcachec_lhc_uib_no_Domain	251658240	51346153	
node009_dcachec_lhc_uib_no_Domain	251658240	51842805	
node010_dcachec_lhc_uib_no_Domain	251658240	50901494	
node011_dcachec_lhc_uib_no_Domain	251658240	56633760	
node012_dcachec_lhc_uib_no_Domain	251658240	54645746	
node013_dcachec_lhc_uib_no_Domain	251658240	53394864	
node014_dcachec_lhc_uib_no_Domain	251658240	60119762	
node015_dcachec_lhc_uib_no_Domain	251658240	65836712	
node016_dcachec_lhc_uib_no_Domain	251658240	109010373	
node017_dcachec_lhc_uib_no_Domain	251658240	108947477	
node018_dcachec_lhc_uib_no_Domain	251658240	80917369	
node019_dcachec_lhc_uib_no_Domain	251658240	53930476	
node020_dcachec_lhc_uib_no_Domain	251658240	60290205	
node021_dcachec_lhc_uib_no_Domain	251658240	190342414	
node022_dcachec_lhc_uib_no_Domain	251658240	191014886	
node023_dcachec_lhc_uib_no_Domain	251658240	190468091	

disk pool status Sep 26 2024

- pledged 2024: 5.6 PB (free 1.8 PB)
- 23 disk pool instances Almalinux 8, each mounting 240 TiB Ceph volume
- Ceph backend: 49 Dell R740XD servers, 8.6 PB raw storage
5.6 PB Ceph storage
- Pledged 5.6 PB (target 4.1 PB), free space 1.8 PB as of Sep 25 2024
- Next extension planned to be operational Oct 2025, expecting min 7 PB raw

UiB dCache tape pool status

Ideling ... waiting for data, there have been only very few transfer requests

but they tell us there will be soon a lot


Transfer requests (add new request)								
ID	Path	Target SE	Status	Progress	Files	Total size	Started	Ended
23523.	Copy 7 runs to ALICE:NDGF:DCACHE_TAPE	ALICE:NDGF:DCACHE_TAPE	Done	<div style="width:100%; height:10px; background-color:blue;"></div>	199310	1.479 PB	09 Feb 2024 00:42	02 Apr 2024 05:16
22032.	/alice/data/2023/LHC23u/537397/raw/*_root	ALICE:NDGF:DCACHE_TAPE	Done	<div style="width:100%; height:10px; background-color:blue;"></div>	2200	19.51 TB	12 Jun 2023 12:08	12 Jun 2023 18:53
21674.	/alice/data/2022/LHC22q/529035/raw/*_o2*	ALICE:NDGF:DCACHE_TAPE	Done	<div style="width:100%; height:10px; background-color:blue;"></div>	456	763.1 GB	02 May 2023 15:01	04 May 2023 21:34
21647.	/alice/data/2022/LHC22e/519497/raw/*_o2*	ALICE:NDGF:DCACHE_TAPE	Done	<div style="width:100%; height:10px; background-color:blue;"></div>	2249	5.033 TB	02 May 2023 13:14	04 May 2023 21:11

- 2 tape write pools: one active, one standby
- 1 tape read pool, spare slot on the infrastructure to make a second one
- pledged 2024: 4 PB (target 5.71 PB), still free \sim 2 PB
- upgrade 2024: currently running vendor competition, expect to place order in week 41
 - ▶ all tape library rearrangement in the server room finished, ready to host new drives and frame
 - ▶ delivery by Nov 15, operational and integrated to dCache Dec 15
 - ▶ min 4 PB, max 8 PB depending on the price we get

Network

- 10 GB/s link is bottleneck
- JAliEn job do not implement data staging, the philosophy is that jobs are mosly running on data in the local storage → does not work for distributed T1
- Still negotiating with SIKT (the national provider) - got some horrible prices for the redundant 100 GB/s link
- There is probably a reasonable technical solution to have redundancy for the UiB site/in NREC; then running over a single link to the hubs (Oslo or Kristiansand)

Uninett LHC link Bergen



⇒ in any case, upgrade is delayed but to be finished before the end of 2024

Summary and plans

- Study Hepscore on the virtual resources
- Network upgrade to be finished ~~before summer~~ by end of 2024
- Tape upgrade started, expecting extension 4 to 8 PB by end of 2024
- Purchasing new compute and disk storage resources in 2025, planning starting spring 2025
- More detailed Grid job performance studies are necessary to understand the poor resource usage by ALICE