

Infrastructure of the ALICE Grid computing setup in Norway

- Tier-1 hardware is hosted within the Norwegian Research and Education Cloud (NREC)
- All nodes are provisioned and managed as Infrastructure as Code (IaC)



Current resources:

- 44 compute servers, hosting compute workers and dCache pool instances
- 325 worker instances with 16 CPUs providing 45 MCPUh/a, ~ 60 kHS06
- 50 Dell disk servers, 7.8 PiB raw, 5.6 PB usable space
- 10 PB tape storage



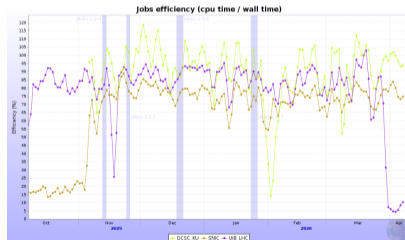
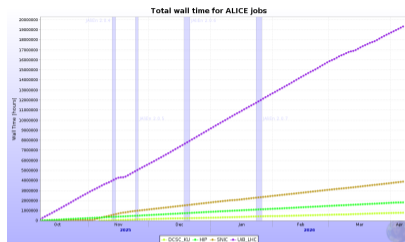
BGO

Status

- Fulfilling pledges for 2026 despite unplanned increase to cover loss of Russian resources
- Operating opportunistic resources - hardware from 2019 with end-of-life in 2024 - to meet computing pledges, in fact providing more on request by ALICE
- Disk storage resources 90% full which brings challenges in the operation of the Ceph backend, instabilities are more likely now, improvement is expected with the installation of new resources and the possibility to separate Grid Ceph from the general NREC Ceph
- Purchase of new hardware January 2026, got roughly half of the planned resources due to price development
- 10 PB Tape pool has been nearly filled: balancing within NT1 (thanks Mattias) and recent migration from ALICE
- Lots of dialog to establish long-term funding for Grid computing in Norway

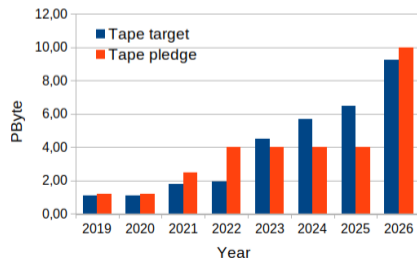
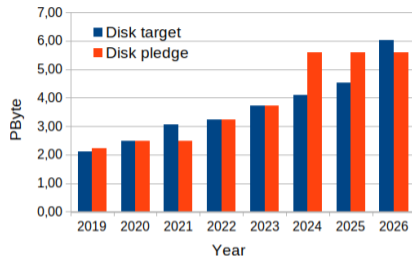
Compute resources

- operating 325 virtual nodes with 16 cores
→ 5200 cores
- $45 \cdot 10^6$ CPUh per year
- pledged: 55 kHS23 (target 51.48 kHS23 in 2026)
- operating opportunistic resources
- Small investment with the rest of NorLHC-II infrastructure project in 2027; applying for funding for investments in 2028 and 2030 (NorLHC III)
- Using JALiEn middleware



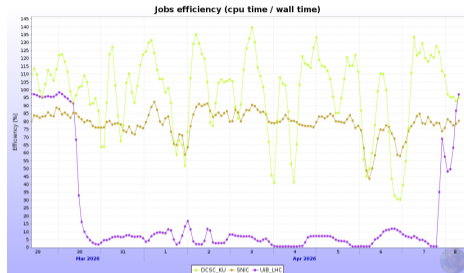
Storage resources

- 5.6 PB Disk storage as Ceph cluster
- expected disk upgrade in 2026: ~~~ 2.5 PB~~ - only got ~ 1.2 PB due to price explosion
- Ceph cluster is nearly full - situation is expected to relax once new resources are installed
- tape upgrade 2025 finished: new drives, new server, extended frame, same robot
- 4 PB (2019) plus 6 PB (2025) tape storage now ~ 90% filled
- Performance problem on tape pool instances will be investigated after finishing upgrade to Alma 9



Moving to IPv6-only

- As of monday March 30 no IPv4 over the LHCOPN in Nordunet anymore
- This impacted full Grid production at UiB but not the other NT1 sites
- Typical situation with many involved “stakeholders”, difficult to locate, bad timing
- Long story short: ALICE is not ready for this, see talk at HEPiX IPv6 Working Group meeting January 2026
- Traced down to IPv4 connection to ALICE central services, they should be IPv6
- Probably some asymmetric flow, will contact Sikt to check this



ALICE on Collider Equipment



IPv6-only ?

Storage nodes: not any time soon, too many WNs on many sites are not dual stacked

- any WN must be able to reach any storage at any time
- the storage itself is close to being fully dual stacked

WNs: could be IPv6-only, as soon as the above is solved

- ok to test with one but in production we will increase the job error rate unless all data is available

LHCOPN: should be ok to drop IPv4

- T1 storages are dual stacked, raw data replication as 3rd party XRrootID transfers
- provided IPv4 still works as a fallback (CNAF WNs reading from CERN EOS for example)

LHCONE: should also be ok

- T2s exchange more data between them (job output -> external SEs to the originating site)
- As per the previous slide, IPv4 connectivity is a must
- Site and storage node-level firewalls are prepared for traffic to come from LHCONE
- but it won't be efficient

6

Request by ALICE for better integration of dCache

- ALICE team is asking for support to query information on NT1 storage
- Basically ALICE central services do not have access to information on available and used capacity in NT1
- two queries they would like to be supported:

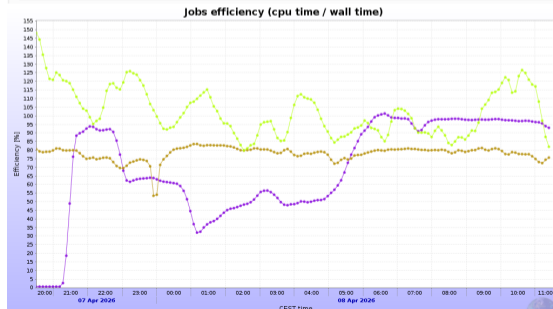
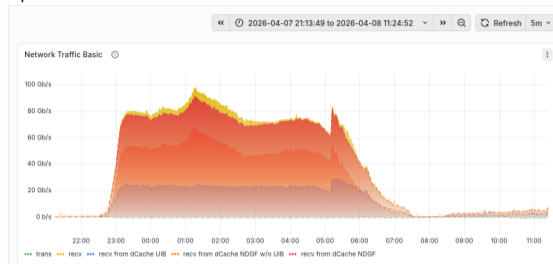
```
xrdfs srm.ndgf.org:1094 query space /alice/disk
[ERROR] Server responded with an error: [3013] Request 3001 not supported
xrdfs srm.ndgf.org:1094 spaceinfo /alice/disk
[ERROR] Server responded with an error: [3013] Request 3001 not supported
```

- Is this NT1 or general dCache?
- ALICE would like to have a contact point in NT1

ALICE Grid jobs on NT1 dCache

- Data locality is not given within NT1 because of distributed storage
- An ALICE Grid job is processing multiple data files, job splitting done by JAlIEn based on the SE
- All files in NT1 storage appear to be in the same SE
- Currently investigating if information from NT1 dCache can be used to optimize traffic
- Impact to job efficiency

Input traffic to worker instances



Status and Plans

- Integration of new hardware soon to happen
- Tape pools going to be upgraded to AlmaLinux 9, last three hypervisors going to be upgraded
- We can fulfill pledges until 2027 with another small investments in 2027
- Infrastructure application NorLHC III including Norwegian Grid resources has been submitted in Nov 2025
- More work on the Ceph cluster to improve stability and increase capacity
- Taking up ALICE's request for better integration of NT1 dCache