

Infrastructure of the NT1 setup @ UiB

NT1 setup @ UiB - Fully operated as Infrastructure-as-a-service



- Tier-1 hardware is part of the Norwegian science cloud (NREC)
- Computing resources are a cluster of virtual machines
- Storage resources organized in a CEPH cluster, mounted to virtual machines and integrated into NDGF dCache
- Tape resources at UiB LHC, tape clients in NREC interface to NDGF dCache

Current resources:


- 24 compute hosts providing 27 kHS06
→ 160 worker nodes with 16 CPUs
- 24 storage hosts with 4 PB disk on distributed file system (CEPH)
- 4 PB tape storage

- current infrastructure installed early 2019, operational since June 2019
- tape server operational all the time, but integration into NDGF dCache has never been fulfilled at the UiB site until work started in early summer 2021

Computing

UiB LHC and Scandinavian (and other) data centers


Total wall time for ALICE jobs




- Deployment of new virtual setup early 2021
- Stable performance since April 2021
- Openstack/Terraform infrastructure
- Slurm cluster managed by ElastiCluster
- ALICE JAlEn middleware on the headnode
- Sporadic glitches, mostly related to the running grid tasks, all handled automatically now

Tape

- Finally, commissioning of the tape resources started Jun 2021
- Hardware problems in the NREC host machines, backlog of commissioning tasks
- Endit deployment, infrastructure upgrade, new local disks in the client instances
- Target: CERN tape data challenge Oct 2021




NDGF tape pool monitoring on Oct 15 2021



Tape continued

- ALICE is requesting 300 MB/s tape write rate to NDGF for normal data taking and all other operations, peak rate will probably be higher (as we have seen in the tape challenge)
- UiB tape performance is matching the requirement but must be improved
- Tape client instances do not allow for more than 500 MB/s parallel input and output network traffic \Rightarrow only one of the tape drives can be operated in full streaming


tape002 network rates during data challenge



\Rightarrow the instance can not keep up with the rate needed for operating two drives if there is input traffic

Verified so far:

- disk performance in the virtual instance
- disk and network performance of the host
- Localized in the network of the VM
 \Rightarrow NREC admin investigating



Rate modeling without disk writes, 800 MB/s target output, variable input

Upgrade 2022

- Purchase first half of 2022
- Prices and delivery times not yet clear, negotiation with vendors started in Sep
- We can increase either computing or storage resources by roughly 50% or distribute among those