




Site Report HPC2N

NDGF All Hands Bern
2026-04-08



HPC2N LHC OPN switch(router)

- Purchased together with the new/replacement dCache disk pools
- Dell S5448F-ON 48x100GbE, 8x400G QSFP56-DD
- Used as a 25G/100G switch to begin with
 - Aiming for 400G in a pilot project later on
- Quirks so far
 - PSU has C15 (high-temperature) connectors, shipped with C15-Shuko cables...
 - Shipped with SONIC, we purchased it with OS10
 - Squealing fans, don't know yet if it's broken speed control or...
 - Needs front-to-back airflow to handle many 400G modules
- SUNET still wants us to do BGP, waiting for feedback to do testing



Tape/backup

- Same as last time
- IBM TS4500 library (2550 slot capacity)
- 6x TS1155 tape drives (JD tapes, 15T, 360 MB/s)
- 6x TS1170 tape drives (JF tapes, 50T, 400 MB/s)
- Dell R750
 - 2x100G Ethernet
 - 4x32G FC
 - Approx 30T NVMe for DB and incoming stgpool
 - A few TB of SAS SSD for log mirrors etc
 - Approx 250T spinning disk for on-disk backup storage
 - 256 G RAM, 2xIntel Gold 6334 (total 16 cores @ 3.6 GHz)

New dCache disk pools

- 7x Dell R760xd2
 - 28x24T HDD, PERC H965i, 128 G RAM, 2x 12core Xeon Silver 4510
- Slow and confused delivery, but finally got all Dell stuff
 - And after even more delays the PDUs and rack arrived
- Final RAID setup was two RAID6:s instead of one RAID60
 - Much better bandwidth
- 2x25G Ethernet ports in LACP
 - Silly since the switch is 100G but there are no such NICs
- Quirk: Needs to slide out to access mid-bay, so rails and cable arms needed!

WLCG compute – BUGZ!

- The arc-arex service started OOM-ing middle of January
 - Without any local changes arc-arex began restarting 10-12 times per hour while using up all memory
 - Limit the number of restarts per hour
 - Give up instead of infinitely restarting
 - Nordugrid bug 4284
 - The problem looked to be related to datastaging (
 - Reducing core-count (1024) made it only slightly more stable, would still suddenly decide to OOM a lot
 - Increasing VM size from 10GB to 32GB improved stability (could now handle 1050 cores. The extra memory avoids immediate crashes during OOM-mode, allowing sufficient time for datastaging to cleanup and eventually stabilize.
 - With proposed fix to 4284, it looks like it is at least able to run ~15 nodes without eating up all memory.
 - Max memory usage does climb up above 12GB (normal is around 2GB)
- Datastaging still stalling with ARC 7.1
 - Nordugrid bug 4191
 - Reduced core count since load seems to exacerbate the problem
 - Maybe the ARC disk cache is too small?



WLCG compute – Aigert (aka g-ce01)

- Same as last time (except when marked)
 - 31 PowerEdge R6625, 256 cores/node, 3GB/core, 3.2TB/node
 - 18 nodes actually used
 - OS – Ubuntu 22.04 Jammy
 - CE running Ubuntu 24.04 Noble
 - Upgraded to try to fix OOM issue. Did not!!
 - New python version, more ARC bugs reported upstream
 - Only one minor service host (grid-home & user mapping) still Focal
 - ARC 7 nightly(!)
 - Fixing(?) OOM issue



Buying Storage

- HPC2N needs more storage
- World AI is eating **all** storage
- Currently seems impossible to procure storage at sane prices
- Our Lustre storage was full
 - Old NDGF dCache nodes being used as Lustre servers as an emergency measure
- Planning/planned to buy cold/lukewarm storage this year
 - Pending whether vendors can deliver anything, and at which cost
- Kludging together a local dCache/tape thing as offload storage