


NDGF All Hands 2021-1

HPC2N Site Report




New WLCG/NDGF-T1 dCache disk pools

- Dell R740xd2
 - 8 servers each with:
 - 25 GigE networking
 - 26 of 16TB SAS 7kRPM HDDs
 - PERC H730P RAID-controller
 - BOSS controller for OS storage (ie mirrored M.2 sticks)
 - 2 of Intel Xeon Silver 4210R 2.4G, 10core
 - Needed to populate 2 sockets due to PCIe slot layout (motherboard is based on R540)
 - 96G RAM
- Delivered on time
- One machine was stuck on low cpu clock, fixed with motherboard replacement
- One HDD replaced (pre-fail, lots of media errors)

New WLCG/NDGF-T1 dCache disk pools (2)

- Dell S5248F-ON switch (25G/100G)
 - Connects stuff to LHCOPN
 - 25G-capable disk pools, tape pools and LHCOPN interface of TSM server
 - Currently 4x10G to LHCOPN switchrouter
 - More on LHCOPN uplink later

ARC 6

- Upgrade to ARC6 relatively easy
- Local admin not understanding
 - how mapping and authentication works
 - authgroup, userlist, voms, lsc — how, when, why
 - sometimes explicit DNs needed in config
 - what services should actually be running
 - acix - must be great we had it before
- Most difficult thing was not ARC but puppet! 
- For instance: Problems getting the cache/session structure to interested parties
- Seems a lot more stable than ARC5
 - Still missing the auto restart script for a-rex we had

Abisko vs Kebnekaise

- Abisko is no more, hello Kebnekaise
- Kebnekaise cores are a lot faster than Abisko cores 😄
 - 2.6GHz Skylake (vs 2.6GHz Interlagos bulldozer)
 - 24.15 HepSpec (vs 9.89)
- Less cores available on Kebnekaise 😞
 - 31 nodes with 28 cores (vs 50+ nodes with 48 cores)
 - Hard limit of 868 vs more than 2500 cores
- Local SSD, fewer cores/node ⇒ use for local scratch.
 - Worked fine until it jobs seemed to get bigger
 - Now back to using NFS sessionsdirs during job runs 😞
- The grid-nodes about to be upgraded to Ubuntu focal
 - The rest of the cluster is already upgraded

Future plans


- New ARC-cache
- 100G LHCOPN on HPC2N segment

New ARC-cache

- We have a separate session on requirements
 - Hint: Lowering size means endurance comes into play when using SSDs!
- Kebnekaise has four free 40G ports in the switch for the grid nodes
 - So 4 ARC cache hosts each with 40G connectivity to grid compute
- Tentatively $8 \times 1.92\text{T}$ NVME SSD:s in each
- 25G connection for each server to the LHCOPN switch
- Guesstimating that a 16-core CPU with 128G RAM will be enough
- Quirk: We use these servers for the session directory storage as well

100G LHCOPN

- Yup, we're sometimes maxing out our 20G (2x10G) uplink.
- Current hardware has 10G ports
- 100G timeline is "during next year".
- More 10G links meanwhile.
 - Aiming for 4x10G.
 - Campus network people discovered shortage of free SM fiber pairs on campus, working on solving that.
 - Lots of FDDI grade MM though...



The End