

NDGF Site Report 2019

NDGF Manager
Mattias Wadenstein
<maswan@ndgf.org>

2019-10-14
HEPiX Fall 2019
Amsterdam, Netherlands

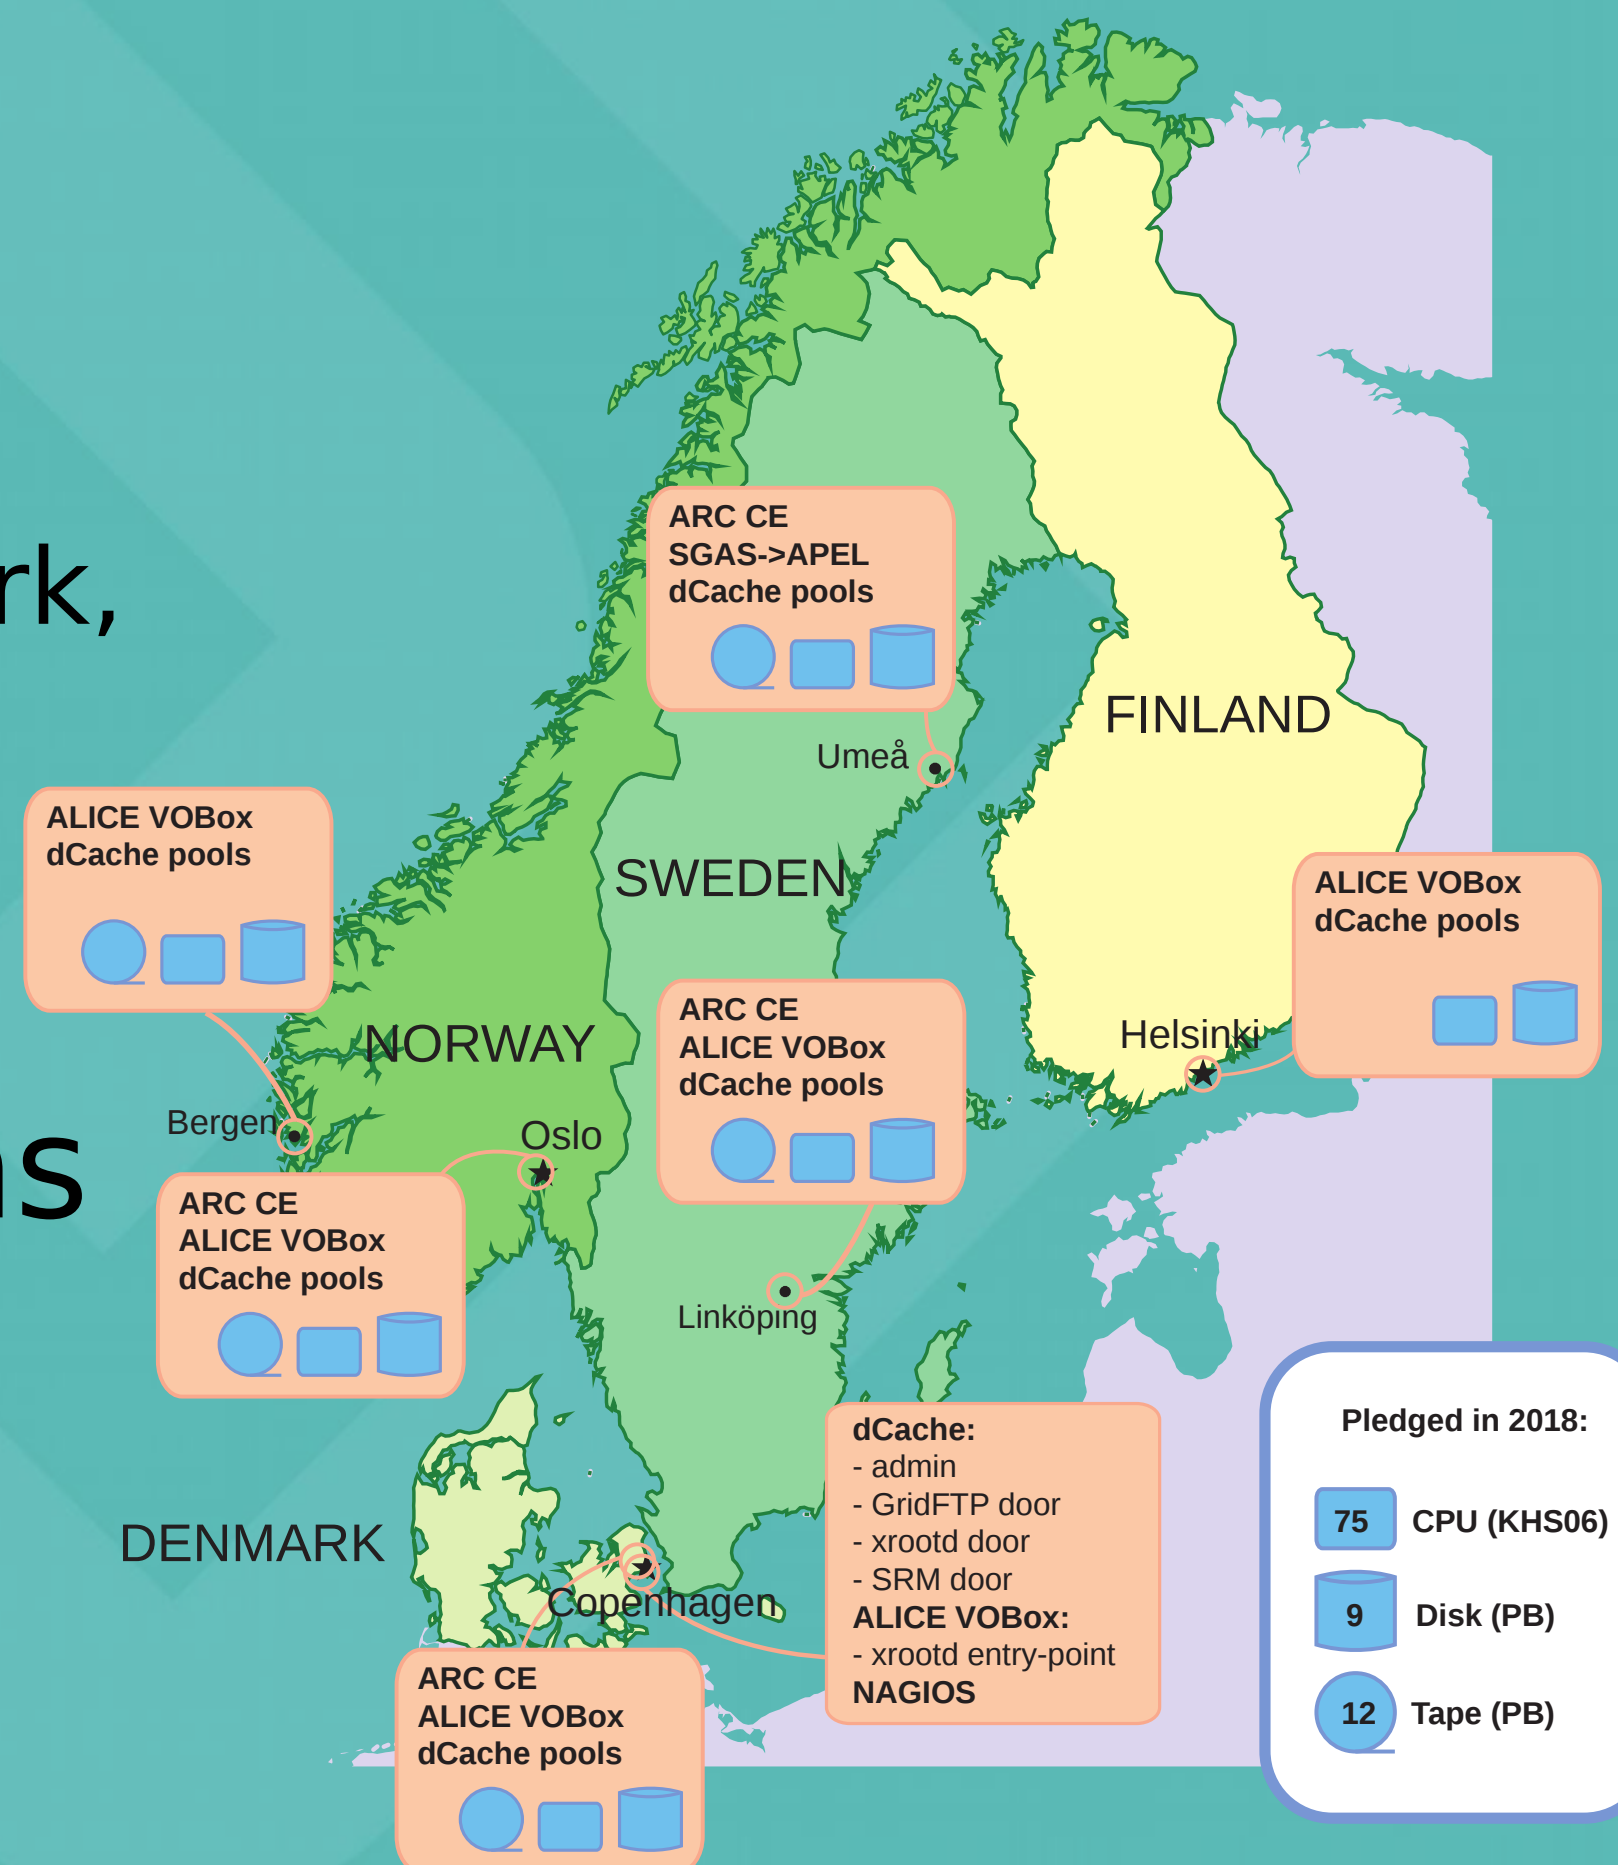
Overview

- Introduction
- Staff changes
- New central nodes
- Site news



Introduction

- NDGF is a distributed WLCG tier-1 site in the Nordic countries
 - Storage and CPU in 7 sites in Denmark, Finland, Norway, Sweden
 - Plus storage in Slovenia technically integrated
- Operated by local site admins and a central team
- ARC, friday workshop

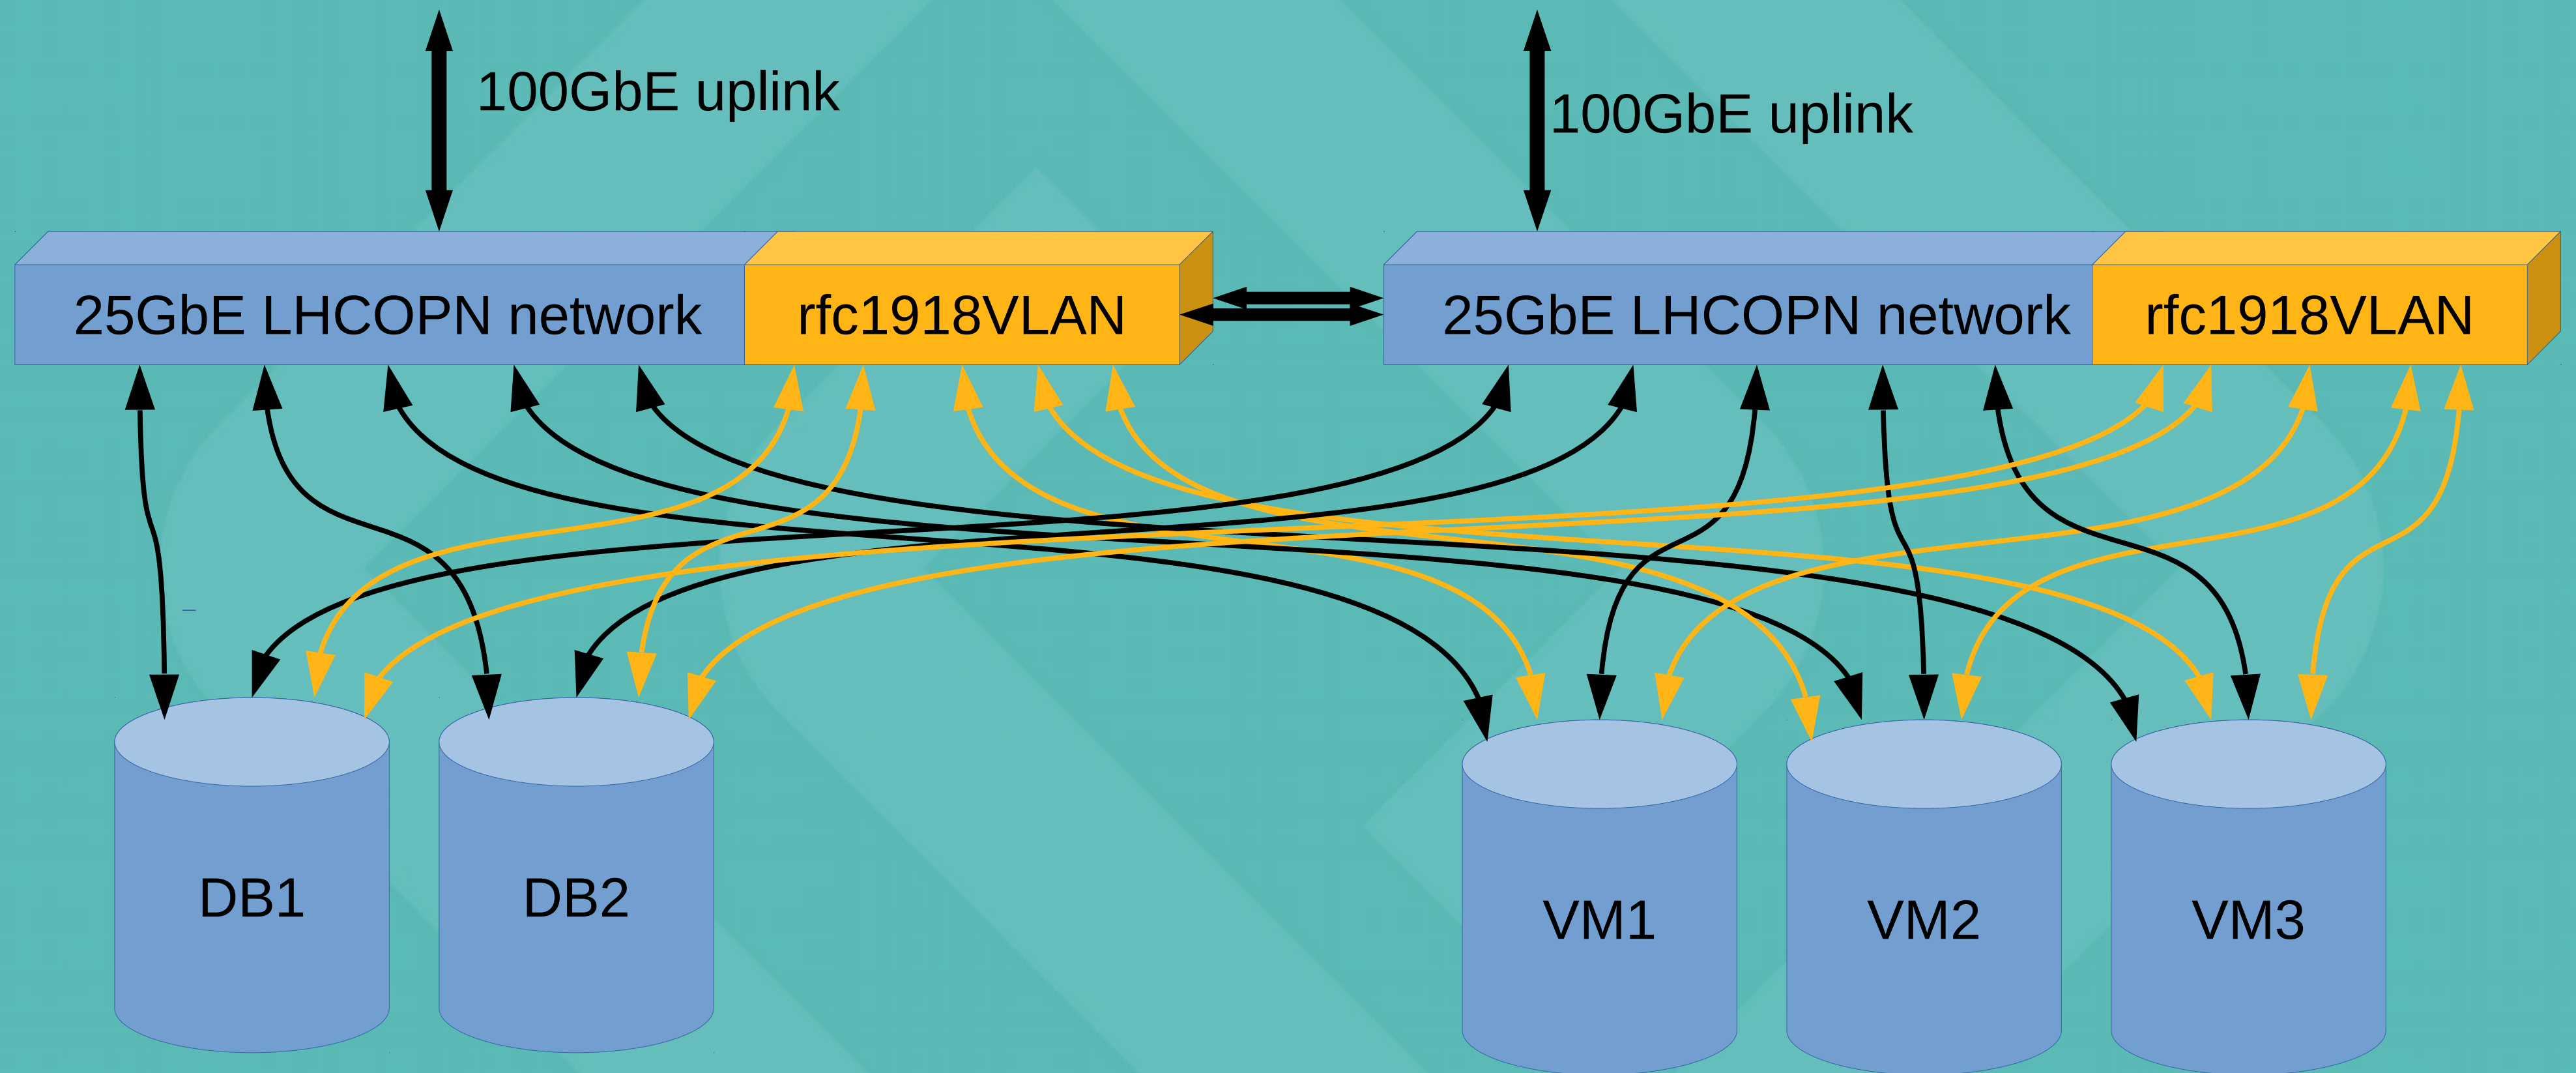


Staff Changes

- Dmytro has left for a new job
 - Replacement job ad currently out
 - <https://neic.no/news/2019/09/20/open-position-nt1/>
- University of Oslo also has another position advertised:
 - HPC- and storage expert (175885) | University of Oslo
 - <https://www.jobbnorge.no/en/available-jobs/job/175885/hpc-and-storage-expert>



New headnodes





5 x Dell r640 with 2x5222 (2x4 core 3.8GHz) and 192GB RAM, 2TB SSD
4x25GbE for redundant (LACP) internal and external network
Network 2xS5212F-ON (12x25+3x100) with multi-chassi link aggregation



New headnodes

New nodes ->



Site news: HPC2N

- Switched to BBR as TCP algorithm for dCache pools
- EasyBuild + CVMFS + Singularity for scientific software
 - For the regular HPC side of the site
- Software is built using EasyBuild which creates modules in Lmod format using a Hierarchical Module Naming System.
 - Some 450 different softwares and libraries installed.
 - Around 2100 different versions of the above.
 - Four architecture builds of most of these.
 - In total ~1.2TB of software



Site news: HPC2N

- Software is distributed with CVMFS.
 - Stratum 0 is currently running in one corner of a small 2 core, 16GB system during initial testing.
 - Master build file system is ZFS that is NFS exported from the same server (compression factor ~2x).
 - Currently no Stratum 1 or proxies.
 - ~1000 clients over two clusters.
 - So far no performance problems seen on the test server.
 - Dedicated Stratum 0 server specified and will be ordered soon.
 - Each cluster will have it's own proxy, mainly to protect from Stratum 0 reboots.



Site news: HPC2N

- No dedicated build nodes.
- Using a minimalistic singularity container to do the builds in a consistent way.
 - Almost no OS distribution development packages.
 - Only security (openssl etc), Infiniband and Nvidia driver and required packages for EasyBuild itself are installed in the container.
 - Everything else is built by EasyBuild.
 - Bind mounts the master NFS mount point of the software tree replacing the CVMFS tree.
 - Post-build sync of NFS into CVMFS.
- For large builds batch jobs are submitted.
 - EasyBuild submits the builds with correct dependency order of build jobs.



Site report: UiO - T1 CPU

- Abel HPC is being decommissioned
 - ce01.grid.uio.no put out of service as of 15.09.2019
 - Low priority queue (ce03.grid.uio.no) is still going strong
- New Tier 1 on local OpenStack facility
 - A small scale cluster has been up and running since summer 2019 (~200 cores) AMD Epyc
 - ARC has been running in pilot mode, some nodes are backfilled with ATLAS@home boinc jobs to experiment with CPU efficiency
 - Currently setting up ARC with ARC cache and a separate ARC data staging service.
 - Openstack grid facility is expanding: AMD EPYC Rome - 3 nodes with 2 x 48 cores - 512 GB Memory - 2.35GHz - in total 1152 cores - w/ hyperthreading 2304



Site report: UiO T1 storage

- Current dCache grid storage pools will be decommissioned in March 2020
- The new dCache grid storage pools are on Ceph
 - 2PB of available space
 - Erasure Code
- New pools are being commissioned
 - Already in preproduction for local UiO physics group




Site report: NBI/KU

- Benchmarks of AMD Rome done in preparation for procurement
- I get to share some numbers here
 - PowerEdge C6525 all are dual CPU with HT of
 - 256GB RAM (16x16GB 3200 MTps)
 - GCC 8 with `-std=c++03`, only had an unprivileged account
 - HPE DL385 Gen10, 256GB 2933MTps memory, HT off (except one)
 - GCC 4.8.5, so closer to rule-abiding run
 - 256GB 2933MTps memory



Results, per core (=jobslot)



Site report: NSC

- New tape library
- Purchased in November 2018
 - IBM TS4500, 5 frames.
 - Capacity 5896 LTO tapes
 - 6 LTO-7 drives (retained from old library)
 - 8 LTO-8 drives
 - And 2 more LTO-8 drives incoming.
 - But IBM have sent them to the Customs Office in Luleå instead of to us, so still waiting. (Customs Office was confused.)
- New Spectrum Protect (aka TSM) servers, in process of being purchased.



Site report: NSC

- SE-SNIC-T2 dCache
- New headnodes purchased.
 - Four Dell R340
 - 64 Gbyte RAM
 - 6-core Intel Xeon E-2186G (Coffee Lake @ 3.80GHz)



Site report: NSC

- Update on 2018 HPC clusters
- Academic compute cluster
 - Purchase reported in Barcelona.
 - Second half installed late autumn last year.
 - Now 2016 worker nodes.
- Numerical Weather Prediction clusters
 - Purchase reported in Barcelona.
 - Installed during autumn last year.
 - In full production since late spring



Site report: NSC

- Cluster vendor bankrupt
- Both academic and NWP clusters bought from ClusterVision BV.
- They went bankrupt this spring...
- Luckily, they sold the clusters to us in consortium with Hammer PLC to sell the clusters to us. Now Hammer are responsible for HW warranty and service.
- Service under Hammer works well. But Hammer is a UK company, so we'll see what happens after Brexit...



Site report: NSC

- OmniPath - One year perspective
- Both academic and NWP clusters use Intel OmniPath for their fast interconnect.
- OPA (OmniPath Architecture) is Intel's answer to InfiniBand. Heavily based on Qlogic's implementation of InfiniBand (TrueScale). Runs at 100 Gbit/s.
- 48 QSFP28 ports in 1u switches, and also large chassis-based switches.
- Works reasonably well. Roughly on-par with Mellanox InfiniBand; some things better, some things worse.



Site report: NSC

- Management tools are nicer than the InfiniBand tools.
- The "verbs" transport is not as well-supported as under Mellanox InfiniBand.
- IP-over-Fabric somewhat fickle
 - Occasional problems with ARP, where nodes temporarily fail to get MAC address of other node. Most normal IP communication retry enough so not a practical problem.
- Drivers from IBM on POWER buggier than RedHat's on x86_64.
 - And GPFS is one of those that IP-over-Fabric users that do **not** retry enough...
- We **seem** to have had more problems with IP over OPA than we have had with IP over IB, but many differences between our IB and OPA systems (CentOS 6 vs CentOS 7, number of nodes, cores per node, ...) so not sure if problems are because of OPA.



Site report: NSC

- Intel announced this summer that they are not going to continue developing OmniPath, so no OPA 2 in the future...
- Leaves us with Mellanox (now NVidia) as single vendor for high-speed low-latency interconnect. (Oh, and Cray as well.)



Site report: NSC

- Thunderstorm #1 (2019-06-12)
- Caused breakers for bypass power to *and* from one UPS (in a group of 5) to trip. Estimated 3-4 kA to trip on momentaneous current. That UPS disconnected from group. Took some time to convince the group to become healthy again.
- Almost same current must have passed through the other four UPS:es, as they are connected in parallel, so at least 15 kA in total.
- All the IBM POWER servers in our GPFS filesystem for academic HPC cluster logged over-voltage, and shut themselves down.
- No other equipment cared. But without filesystem...



Site report: NSC

- Thunderstorm #2 (2019-08-06)
- District cooling stopped, for both computer rooms (despite being fed from different cooling stations).
- Common occurrence when there are power glitches.
- Personnel from district cooling company need to go out to each station and restart, so it can take some time before cooling is restored.
- Usually restored quick enough for us, but this time one of the stations took several hours.
- Had to power off about half of worker nodes in that room in order to not overheat.



Site report: NSC

- Thunderstorm #2 (continued)
- In the other computer room, three UPS:es, in group of six, didn't like the voltage spike, went to battery power, and then immediately disconnected their outputs. One of them also tripped its battery breaker.
- Luckily, that room is currently not used at capacity, so remaining three UPS:es were enough to power everything.
- Not so luckily, a 7th UPS, powering ventilation system for cooling, did the same thing, including tripping its battery breaker...
- But luckily, room is not used at capacity, so temperature rose slow enough for us to notice, drive back to work, and put cooling UPS on external bypass.





Questions?