# NDGF Site Report

**NeIC NT1 Manager**
**Mattias Wadenstein**
**<maswan@ndgf.org>**

2018-10-08
HEPiX Fall 2018 Workshop
Barcelona, Spain

NORDIC E-INFRASTRUCTURE COLLABORATION

# Overview

- Storage and comuting
- Recent changes
- Site reports

# Distributed storage

- NDGF is a single dCache instance with common namespace etc with disk and tape distributed over many sites
  - Central administration team
  - Storage pools bought and run by local sites
- Tier-1 pledged storage plus some integrated tier-2 storage
  - NDGF pledges for ATLAS and ALICE tier-1 storage
  - ATLAS storage part of Swedish and Slovenian tier-2 pledges

# Distributed storage

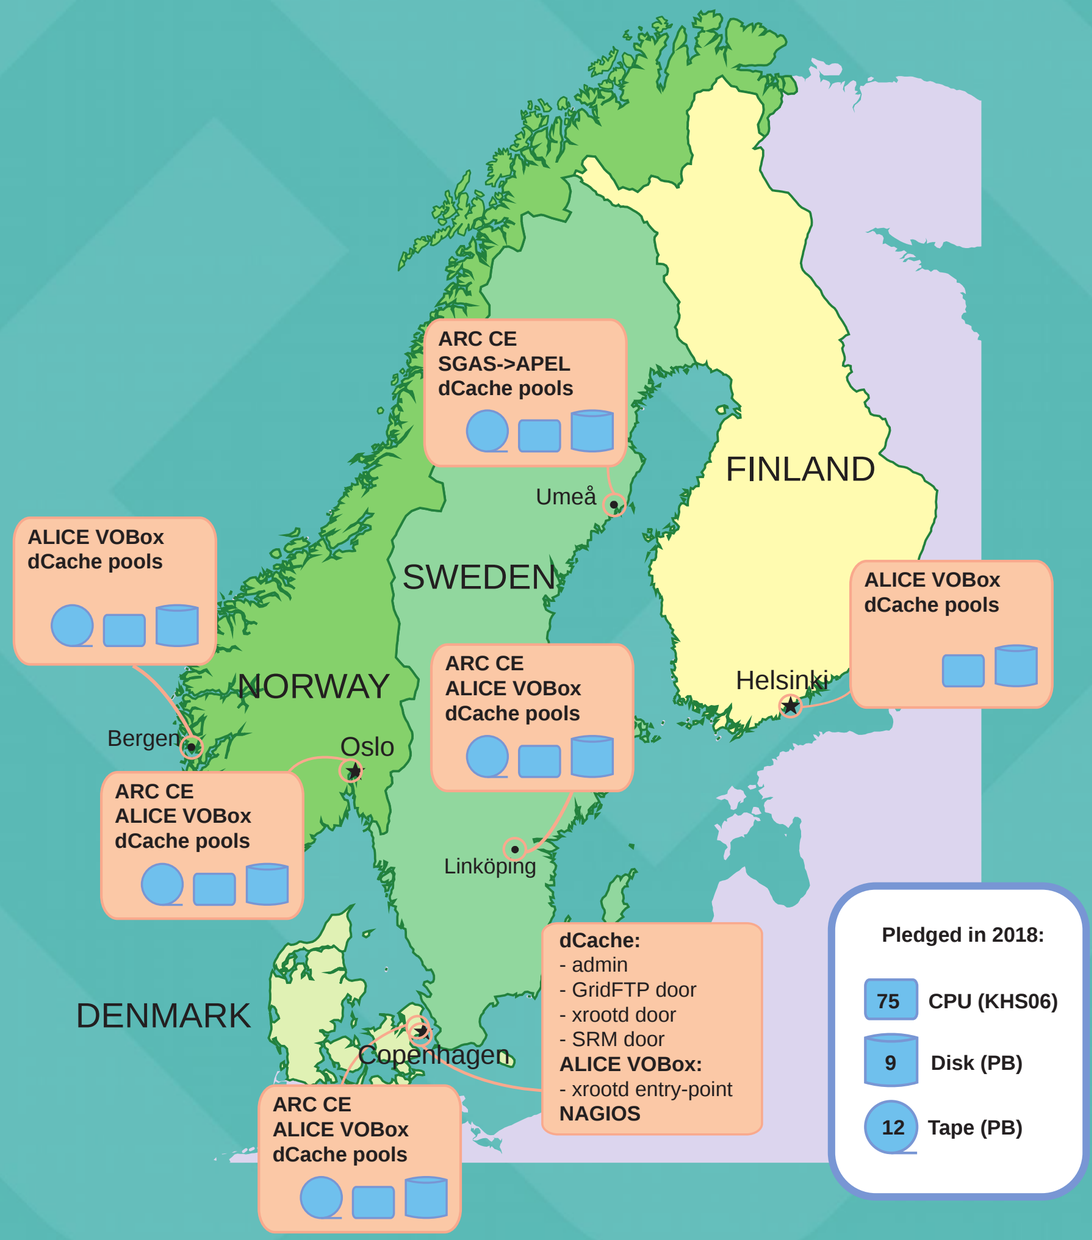

# Distributed storage

- New model of site operations for disk pools
  - Maintain hardware, networking, and OS
  - Provide a non-privileged user account with ssh login to central admin team
  - Takes very little effort, and only normal Linux sysadmin tasks
  - Only providing office hours coverage (distributed setup means loss of a site is only a “service degradation <20%”)
- For tape pools
  - Some more complications in configuring the HSM connector
  - In addition to running tape library etc

# ARC in data staging mode

- ARC is location agnostic
- No problem getting some data from other sites
- By staging data before running jobs, good CPU efficiency

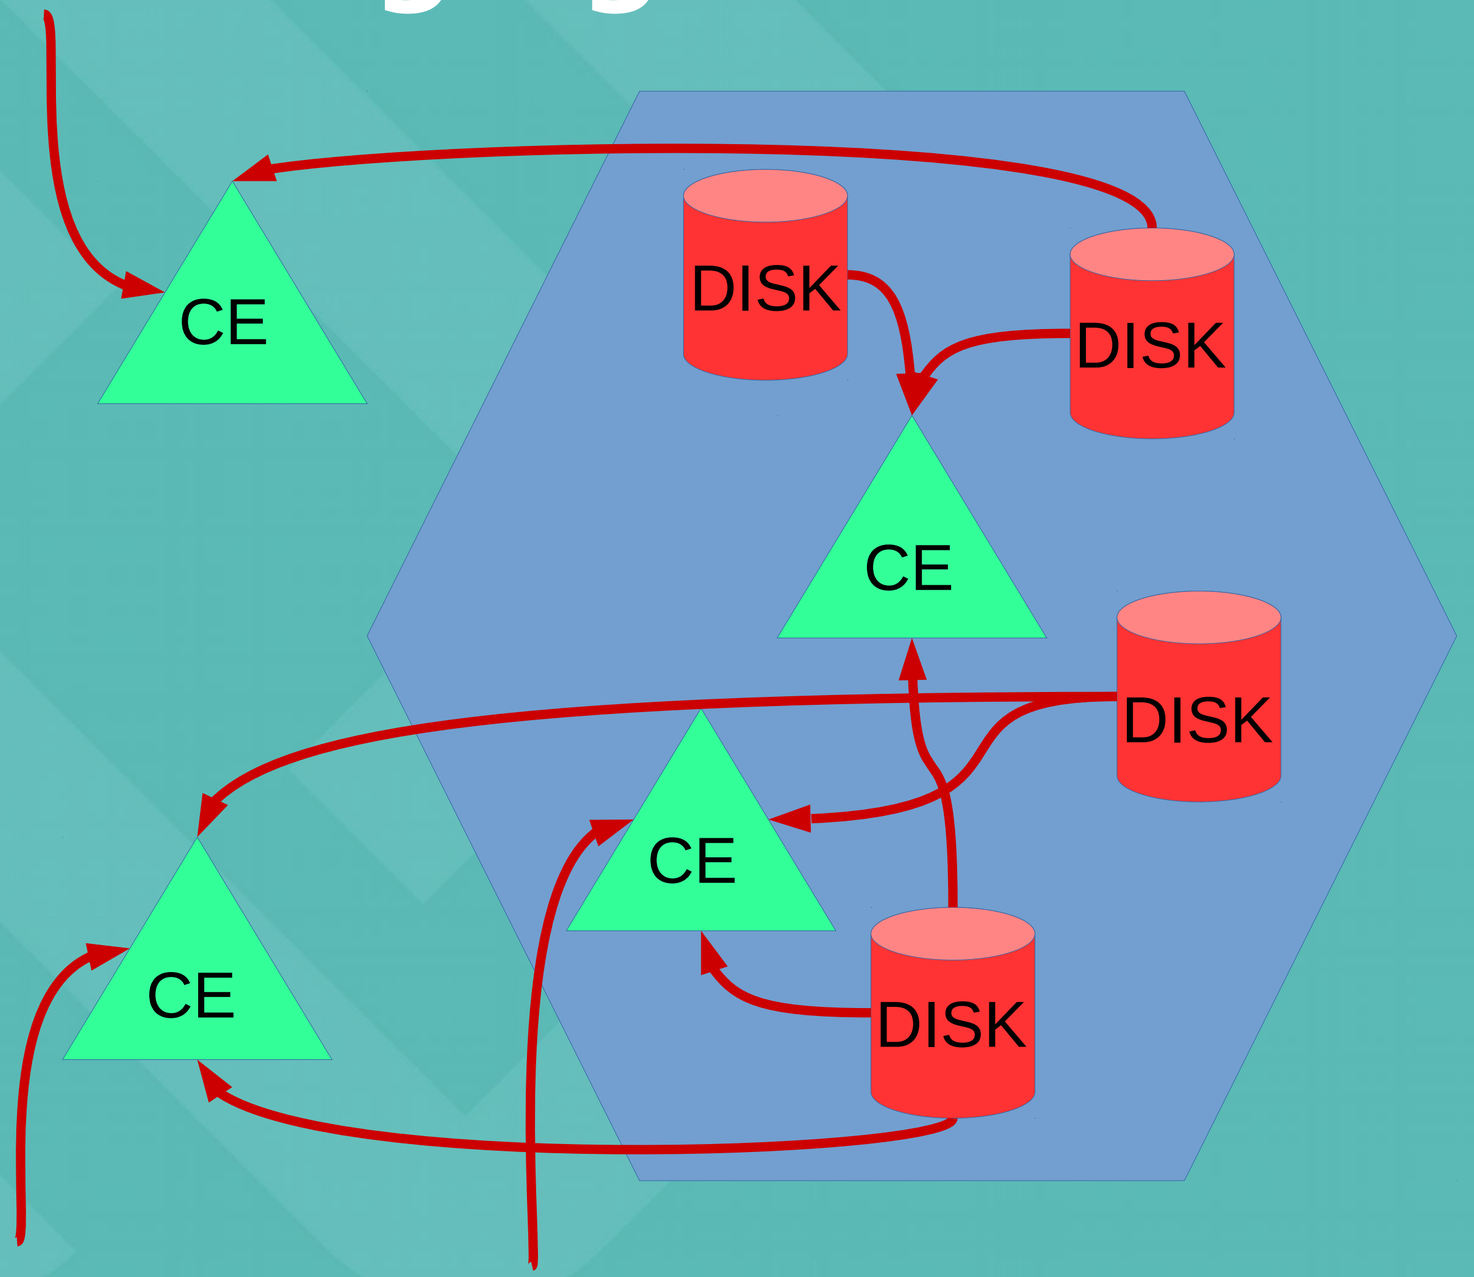

# ARC for caching

- ARC with 100TB of cache can support full ATLAS data flow on 5-10k cores
  - A size where IOPS-friendly storage is not horribly expensive
  - Reliability only needs to be OK, impact is crashed jobs not lost data
  - Makes it possible to buy the bulk storage optimized for size, aggregate bandwidth, and reliability
- Good cache reuse (>80% hit rates)
  - Size of cache impacts cache hit rate over time
- Protocol agnostic
  - Easy to add new protocols to inputs, no application support needed

# ARC for caching

- Green is downloads from SE to cache
- Blue is serving data from cache to WN (NFS)

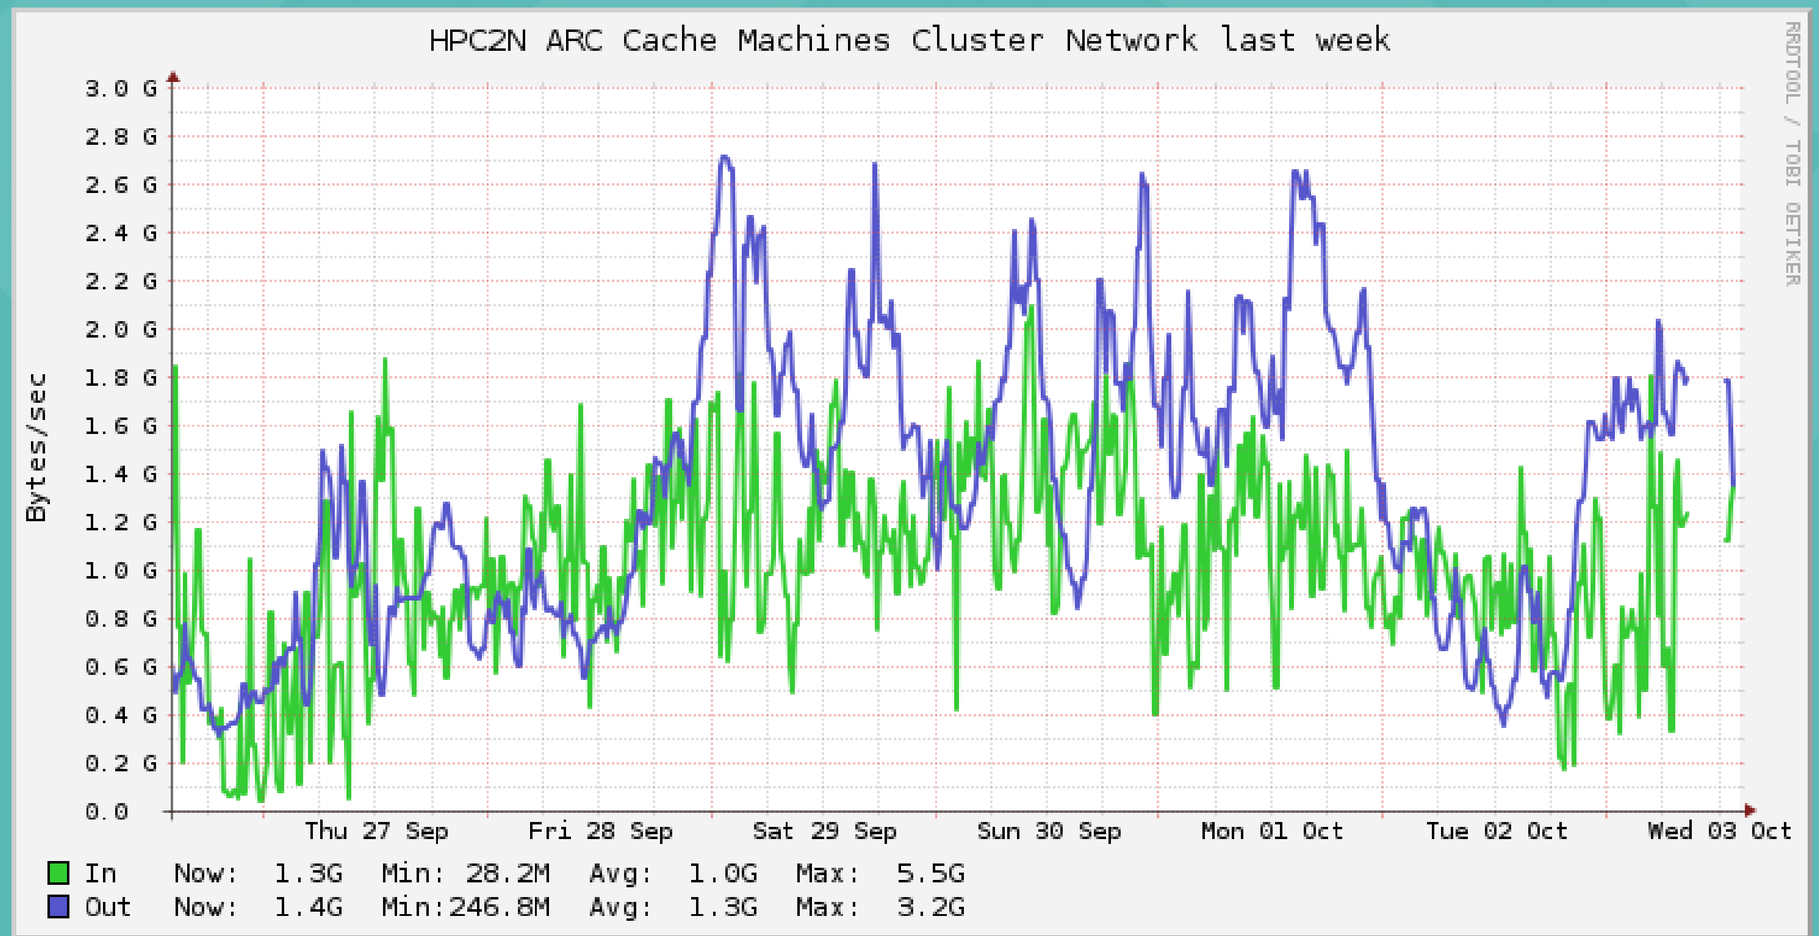

# Nordic Data Glaciers and Fjords

Photo: Erlend Bjørtvedt (CC-BY-SA)

# Recent Changes

- Staff changes
  - Vincent Garonne storage developer since last fall
  - Ulf Tigerstedt has moved on to a new job
  - Ville Salmela is a new sysadmin in the central team
- New ENDIT version
  - Parallelism, logging, and stuff improvements.
- New site disk pool operations
  - Running dCache up to central operations team, via ssh and ansible
- Central dCache news
  - Proper pre-production setup, almost identical to production
  - Test before deployment, with real user load

# Site reports, HPC2N

- Newest HPC cluster Kebnekaise expanded
  - 52 new Skylake nodes (2x14 cores, 192GB ram)
  - 10 new Volta nodes, Skylake nodes with 2xNVidia V100
- New UPS for storage&services computer room

# NSC: Academic HPC cluster

- Contract awarded to ClusterVision BV
  - 2016 worker nodes
  - Intel HNS2600BPB nodes in Intel H2204XXLRE chassis (4 nodes in 2u)
  - Dual Intel Xeon Gold 6130 (Skylake, 16 cores, 2.1 GHz)
  - 96 Gbyte RAM (64 nodes with 384 Gbyte)
  - Gigabit ethernet
  - Intel OmniPath 100 Gbit/s interconnect.  Folded Clos ("spine and leaf") architecture with 3.8:1 oversubscribe
  - Central OmniPath director switch with __ ports (can be built out to 768 ports)
- First half installed this summer, second half to be installed in October/November.

# NSC: Numerical Weather Prediction clusters

- Two clusters, for Swedish, Norwegian and Finnish met offices.
- One located at NSC in Linköping, the other at Swedish met office in neighbouring city of Norrköping (for redundancy).
- Contract awarded to ClusterVision BV
  - 880 worker nodes in total (524 at NSC, 256 at SMHI)
  - Identical hardware to academic cluster
  - But 2:1 oversubscribe on OmniPath fabric
- In process of being installed

# NSC: GSS + Infiniband -> ESS + OmniPath

- Have had IBM/Lenovo GSS (x86-based servers with GPFS) connected with Infiniband to old academic cluster.
- Rebuilt to an IBM ESS (POWER-based servers) with OmniPath this summer.

# NSC: OmniPath - is it any good?

- Some problems:
- A linecard in central director switch hung, with strange symptoms.  Fabric manager probably got confused by the strange state.  Reboot of switch solved the problem.
- Multicast group sharing for IPv6 neighbour discovery doesn't work (at least using the fabric manager bundled in EL-7). Didn't have time to solve problem before putting cluster into production; IPv6 thus turned off over OmniPath (but still enabled on ethernet).

# OmniPath vs IBM ESS

- Note: ESS with OmniPath is not yet officially supported by IBM.  NSC is the second customer to run this (Barcelona Supercomputing Centre the first).
- Severe stability problems
  - Several crashes.
  - Problems with IP-over-OmniPath driver in the ESS (this is IPv4 only, no relation to IPv6 problems)
  - RDMA also give problems.  At least one problem found when links are not perfect.
  - Working with IBM to figure out
  - No solutions yet, but some workarounds deployed (use IP over ethernet, while still running RDMA over OPA; detect hangs and power off servers automatically)

# NSC: Network outage

- October 20th, 2017, 08:35 Swedish time (15:35 Japan time) Friday of HEPiX in Tsukuba. :-)

- All network connectivity to Linköping University lost.

- But we do have two redundant fibers, with at least 10m separation.

# NSC: Network outage

# NSC: Network restored

- One link restored at 11:28 Swedish time. Other link restored 16:20.
- When the NREN purchased the fiber paths, they got maps showing two different fiber paths with sufficient separaion, but this got repatched at a later time without telling the customer.
- Municipality fiber company has now **labeled** our fibers, with notes to not put them in the same duct. (Why are they not doing that by default?)

# NSC: UPS battery meltdown

- April 13th, 2018.  Friday evening (17 o'clock): One of our personnel went in to UPS battery room in one of our computer rooms.  Strong smell of battery acid.
- One battery bank (out of 8) was ~~warm~~hot. 50°C on the front of the batteries, almost 80°C on the top. Felt like standing in front of an oven...


**SPEAKER | Mattias Wadenstein <maswan@ndgf.org>**

# NSC: UPS battery meltdown

- Top of batteries no longer flat.

# NSC: Battery meltdown

- Over several weeks, battery current (to keep batteries at loat charge) had slowly risen from ~1A to ~5A.
- UPS regularly pushes more than that through batteries, as it does charging cycles, so can't just set simple Nagios alarm on "high" battery current.  Need to check for high **sustained** current.
- Or better, monitor temperature of individual batteries. (Bonus: also monitor voltage and resistance) Highly recommended!

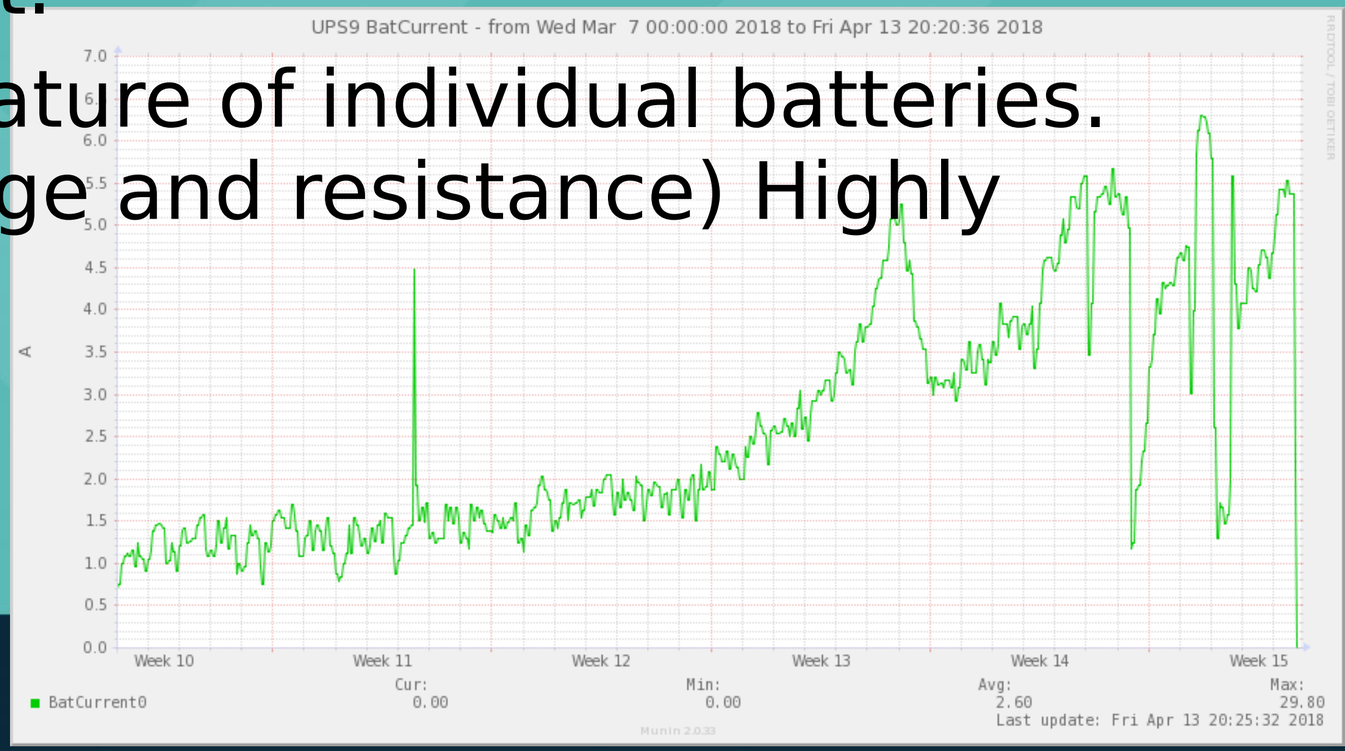


**SPEAKER | Mattias Wadenstein <maswan@ndgf.org>**

# Disturbance in the electric force

- When powering on ~100 nodes in new cluster at the same time, several other nodes and some network switches powered themselves off.
- Measured voltage in outlets: perfect balance phase-to-phase, but phase-to-neutral fluctuating between 210 V and 250 V (should be ca 230 V).
- The box connecting the power rail to one of the sub-centrals in the room was warm.  ~45°C on the outside, in a room that has excellent cooling with 20°C air.
- Loose bolt connecting neutral.  Arcing inside box.
- Had to power down entire computer room to repair.

# Questions?

NORDIC E-INFRASTRUCTURE COLLABORATION