Site report
National Supercomputer Centre
Linköping University, Sweden

Thomas Bellman

bellman <at> nsc.liu.se

HEPiX autumn 2022

National Supercomputer Centre

HPC centre at Linköping University (LiU), Sweden.

- Founded 1989.
- 34 employees (and hiring).
- Provides high performance compute resources to Swedish academia.
- And to Swedish, Norwegian and Finnish met-offices for weather forecasts and research.
- Also collaboration with SAAB in Linköping.
- And a little bit WLCG compute and storage.
- In total around 3600 worker nodes.

One out of six similar centres in Sweden (Lund, Gothenburg, Linköping, Stockholm, Uppsala, Umeå), but we were the first, so we could claim the name "National". :-)

Computer rooms

Two computer rooms: Hangaren ("The Hangar"), and Kärnhuset ("The Core House", as in CPU cores).
Separate buildings, located ca 20 meters apart.

- Designed for 1 MW power each.
- Currently uses 400-500 kW and $\sim$900 kW, respectively.
- District cooling.
- Different electrical and cooling feeds.
- Solar panels on roof of Kärnhuset: peak 225 kW, average 25 kW over year.

Current systems at NSC

Compute clusters

Tetralith / Sigma

- One physical cluster, split logically into two.
- For academia in Sweden (Tetralith) / at LiU (Sigma).
- 2006 worker nodes (Tetralith: 1896, Sigma: 110).
- Intel Xeon Gold 6130 CPUs (16 cores @ 3.7 GHz / socket)
- 96 Gbyte RAM / node (and some fat nodes with more memory)
- Intel OmniPath 100 Gbit/s interconnect, with 4:1 oversubscription.
- Purchased and installed in 2018.
- Ranked #74 on Top500 in June 2019 (now: #183).
- Most common usage: quantum chemistry, molecular dynamics, and computational fluid dynamics.

Compute clusters

Stratus / Nebula

- One physical cluster, split logically into two.
- Stratus part: Weather forecasting for Sweden, Norway and Finland.
- Nebula part: Norwegian met-office research.
- 656 worker nodes.
- Purchased and installed in 2018.
- Hardware identical to Tetralith / Sigma.

Cirrus

- Weather forecasting for Sweden, Norway and Finland.
- 256 worker nodes.
- Purchased and installed in 2018.
- Hardware identical to Tetralith / Sigma.
- Located at Swedish met-office (SMHI) in neighbouring city of Norrköping for redundancy.

Compute clusters

Bi

- Swedish met-office research.
- 641 worker nodes.
- 2 $\times$ Intel Xeon E5-2640v3 (8 cores @ 2.6 GHz / socket)
- 64 Gbyte RAM / node (and some fat nodes).
- Intel/Qlogic Truescale QDR Infiniband (32 Gbit/s) interconnect.
- Purchased and installed in 2014.
- No longer any hardware support
- To be replaced sometime next year (we hope); replacement will be smaller.

Compute clusters

Bluegrass

- WLCG compute cluster running ARC with SLURM.
- Part of NDGF-T1.
- 40 worker nodes; see later slide.

Storage clusters

Centre Storage

- Storage for our academic clusters (Tetralith and Sigma).
- IBM ESS system (Spectrum Scale, a.k.a. GPFS, "as an appliance", running on IBM POWER-based servers).
- 6200 TiB usable space.
- Used to be connected directly onto OmniPath fabric.
- IBM have stopped supporting OmniPath when buying new ESS building blocks. And we wanted to expand...
- Rebuilt to use 100G Ethernet on storage servers, and built routers between Ethernet and OmniPath on 10 worker nodes. Plain Linux servers, forwarding in software.

Storage clusters

Accumulus

- Storage for Swedish met-office research.
- Lustre file system, using ZFS-on-Linux as storage backend.
- Ca 18 PiB usable space (but compression is used in ZFS, so users can actually store more).

Serenity

- dCache storage for NDGF-T1, and Swedish national storage (which includes the SE-SNIC-T2 WLCG site).
- 6400 Tbyte available space ($\sim$3200 TB each for NDGF-T1 and national storage).

Tape storage

- IBM TS4500 library with 5 frames.
- 10 LTO-8 drives, 6 LTO-7 drives.
- 5502 tape slots.
- Currently 5097 tapes (mixture of LTO versions).
- IBM Spectrum Protect (the software formerly known as "TSM", Tivoli Storage Manager).
- Usage:
  - Backups for ourselves, and some of our sister sites in Sweden.
  - WLCG via NDGF-T1.
  - Swedish met-office archival data (climate data).

News since autumn 2019

COVID-19 aftermaths

- No-one seriously sick from COVID at NSC.
- We are now allowed working 40% from home.
  - Some take advantage of that, others don't.
  - Nice to have the ability/freedom to do so.
  - But communication challenges: less face-to-face talks.
- Zoom is a boon: less travelling for meetings.
- And attending conferences you otherwise didn't have time for.

WLCG storage expansion

- Purchased summer 2020.
- 7 $\times$ HPE ProLiant DL385 Gen10 Plus
  - dCache disk pools for NDGF-T1
  - 128 Gbyte RAM
  - 16 $\times$ 14 Tbyte disk (224 TB) each
  - $\sim$1.2 PiB available space across all 7 servers.
- 2 $\times$ Dell PowerEdge R7515
  - dCache tape pools for NDGF-T1
  - 64 Gbyte RAM
  - 8 $\times$ 1.92 Tbyte (15 Tbyte) SSD
- 25 Gbit/s network.
- Used for NDGF Tier-1.

New WLCG worker nodes

- Purchased summer 2020.
- 40 $\times$ Dell PowerEdge C6525
  - 4 nodes in 2U chassis
  - 2 $\times$ AMD EPYC 7302 (Rome)
  - 16 cores @ 3 GHz / socket
  - 256 Gbyte RAM
  - 1.92 Tbyte SATA SSD
  - 25 Gbit/s Ethernet
- Replaces old worker nodes in Bluegrass cluster (used for NDGF-T1).

New AI/ML system

BerzeLIUs

- Grant from Knut & Alice Wallenberg (KAW) foundation.
- KAW pointed to NVidia: "buy whatever they tell you"
  - ... and please have it running yesterday...
- 60 NVidia DGX worker nodes:
  - 2 × AMD Epyc 7742 CPUs (Rome)
  - 64 cores @ 2.25 GHz / socket
  - 1 Tbyte RAM
  - 8 × NVIDIA A100 Tensor Core GPUs, 40 Gbyte memory
  - 15 Tbyte local NVMe SSD storage
  - 8 × Mellanox HDR InfiniBand @ 200 Gbit/s (total 1.6 Tbit/s, per node!)
- 1 Pbyte shared NVMe SSD storage
  - DDN A^3I storage
  - Dedicated HDR InfiniBand storage network

New AI/ML system

- Water-cooled rear doors
- Bright cluster manager
  - Came with the purchase.
  - Not entirely positive experience...
  - Need to run Puppet as well, to manage things Bright can't manage.
- Installed spring 2021.
- Ranked #82 on Top500 in June 2021 (now: #102).
- Largest user is human language modelling.
- Expansion of 34 nodes ordered.
  - 80 Gbyte GPU memory.
  - Air-cooled; no water-cooled rear doors.
  - Expected installation in spring 2023.

New AI/ML system

- NVidia DGX server not very high quality
  - Replacing approx. one motherboard per week (on just 60 nodes...).
  - Pro-actively replaced power distribution boards in all nodes, after one caught fire.

Core network upgrade

- Old core (spine) routers:
  - 2 × Juniper QFX 5100-48S (48 × 10G + 6 × 40G each)
  - Total 10×10G uplinks to university border routers, which have one 100G uplink each to NREN network.
  - Purchased and installed in 2014.
- Finally got around to order new routers this summer:
  - 3 × Juniper QFX 5120-48Y (48 × 25G + 8 × 100G each)
- Increased numer of spines from 2 to 3:
  - More bandwidth
  - More redundancy
  - More bandwidth in degraded mode
- Will initially upgrade uplinks to 3×4×10G uplinks to border routers, later hopefully upgrade to 25G (requires new border routers from the NREN).
- Expected delivery time: 12-20 weeks, according to retailer.
- Actual delivery time: 3-4 days (arrived same week we ordered them...)
- Haven't had time to take into production yet.

Data-driven Life Sciences

Recently started project building storage and (non-HPC) compute system for sensitive data (genetic and medical data about living humans).

- Kubernetes-based compute.
  - Kubernetes clusters within virtual machines.
  - VMs spun up and down on-demand.
- Each project gets their own, private Kubernetes cluster.
- Different projects must not be able to communicate with each other.
- Ceph-based storage system.
- Unholy network machinations to keep projects separate, with (encrypted) VPN tunnels between VMs belonging to same project. VPN done by the VM *host*, since guests can't be trusted to do the right thing.
  - VXLAN + EVPN?
  - IPsec / Wireguard / MACsec?

Two-factor authentication

Two-factor authentication for users of our compute clusters:

- Introduced in 2021; optional in spring, mandatory after summer.
- TOTP (Time-based One-Time Passwords).
- Users get link to get a TOTP secret emailed to them.
  - Need to log in to our user/project portal to use, so not enough to steal someones email account.
- Problems: Users
  - lose or change their smart phones
  - don't complete all steps in the TOTP registration
  - use wrong account in their TOTP app
  - have incorrect time on their phone
- A small handful of user problems / TOTP resets each week.
  - About the expected amount.

Two-factor authentication

Yubikeys for NSC personnel:

- Introduced spring 2022.
- SSH keys on Yubikeys.
- Uses PIV (Personal Identity Verification) functionality.
  - Works with standard SSH servers.
  - Client tooling easily available in/for all operating systems (including EL-7, FreeBSD, MacOS, Windows, Fedora, Ubuntu, et.c).
- One person is using another kind of smartcard. Works well, and does some things better than Yubikey.
- One person also has personal grid certificate on Yubikey (for authenticating against web services). Some issues with Safari web browser, but otherwise works.

UPS woes

UPS woes, room #1

- UPS:es in computer room #1 getting old (bought when room was built in 2007).
- Batteries replaced once during lifetime, but also getting old.
- Bought new UPS system, installed spring 2021:
  - 6 $\times$ Eaton 93PM 200 kW for main load.
  - 1 $\times$ Eaton 93PM 80 kW for cooling system.
  - 4 power modules @ 50 kW in each unit; can be powered off individually to save energy when low load; automatically powered on when load increases.
  - Economy mode: run in internal bypass when input power is of "good quality".
- Room designed for 1 MW, but currently using around 400-500 kW.

UPS non-woes, room #1

Battery monitoring system!

Generex BACS system. Monitors temperatures, voltages and resistance of each individual battery. Also tries to actively ensure all batteries have the same voltage.

Highly recommended to get some way of monitoring temperatures and voltages of individual batteries.

UPS woes, room #1

- When more than 21 power modules (of total 24) active, UPSes generated 450 Hz overtone on output.
- Large enough to cause power subcentral to physically, and audibly, vibrate...
- Took Eaton long time to reproduce in their lab (mostly time for them to decide to actually test it themselves...).
- Solved with a software update.
- Now works very well.

More UPS woes, room #2

- UPS:es in second computer room had broken down; also needed to be replaced.
- Since Eaton UPS:es in room #1 now worked well, we decided to buy identical system for room #2.
- Installed winter 2021/2022.
- Seemed to work well initially.
- But...

More UPS woes, room #2

- Ten days before planned acceptance test, we had a power outage affecting room #2.
- Noticed that nothing in that room was reachable.
  - Including Matrix chat, RT support system, our large academic cluster, and more.
  - Except cooling PLC:s, which for reasons are on our landlords network, not our own.
  - Email server, Nagios, and weather forecast cluster were reachable, but they are in computer room #1.
- Saturday late evening, February, snowy night, and my bicycle was broken...
- When I arrived, everything was silent, but lights (not on UPS) and cooling (separate UPS) were on.
- Battery breakers on all six UPSes in main group (non-cooling) had been tripped...

More UPS woes, room #2

- System designed to run everything for 4 minutes, then turn off power to worker nodes, and then allow storage, infrastructure servers and network to run for one hour.
  (Disconnecting worker nodes is done by external PLC:s, not the UPS:es themselves.)
- According to logs, the UPS:es had been shut down by "external command" 10 seconds after input power was lost.

More UPS woes, room #2

- Turned out to be feature of the network cards...
- Network cards can turn off UPS outputs for low-prio consumers when batteries are "low", for those that don't have external PLC:s doing that.
- High current to support worker nodes for a few minutes causes battery voltages to drop noticably, and network cards decided batteries were low, and turned off the UPS:es...
- UPS:es in room #1 had been configured to ignore such commands from network cards (this setting only available to Eaton service personnel, not customers).

More UPS woes, room #2

More UPS problems in room #2:

- UPSes didn't want to go into economy mode. Took months for Eaton to figure out that the UPSes were configured differently than the ones in room #1. Despite extracting the configs *twice* from them...
- After service personnel configured UPSes correctly, they now go into economy mode.

More UPS woes, room #2

- UPSes now go into economy mode, but voltage drops to zero on one phase and 80V on another phase for 1/4 period (5 ms) during transition...
- Still unsolved...

A Computer Security Ghost Story

Security ghost story

User support ticket: threatening text appeared in their terminal while logged in to our big academic cluster.

- String of ∼50 words; not entirely coherent, but not complete gibberish either.
- Kind of looked like ramblings from a madman, or what you would expect from a haunted fax machine in a horror movie.

Contained phrases like *"nobody can come to my home"*, and *"if you need help going to kill you okay okay okay yeah"*...

Security ghost story

Investigation started:

- Who were logged in at the same time?
- Who were logged in earlier using the same tty?
- What processes had been run around that time (logged using audit system)?
- Called user on phone to get more details. Language barrier: user not entirely fluent in English, but after while managed to learn:
  - Text appeared as if typed, not all at once.
  - User was using a Macintosh.

Security ghost story

After a couple of hours, suggestion from another colleague:

> *"This kind of looks like when I turn on automatic subtitling in Youtube. Could this be some kind of voice typing?"*

Hmm...

Security ghost story

After a couple of hours, suggestion from another colleague:

> *"This kind of looks like when I turn on automatic subtitling in Youtube. Could this be some kind of voice typing?"*

Hmm...

Off to yet another colleague, someone equipped with a Macintosh.

Learned from local Macintosh user:

- Some work needed to enable voice typing in MacOS.
- But when enabled, it can be activated by pressing `Ctrl` button twice.
- Icon shows up when voice typing active, but that could be overlooked when focusing on other task (and part of screen).

Security ghost story

In the meantime, got an actual screenshot from the user, containing $\sim$40 more words. This time containing the phrase *"can you take a print screen of this to my home direction"*.

Also learned (via user's PI) that user was sitting in a shared office with several people in small room.

And later got confirmation from user that voice typing was very likely the culprit.

Security ghost story

Conclusions:

- Our biggest system is *not* cracked.
- Neither is the user's laptop / workstation.
- We *don't* need to take the cluster offline for several weeks.
- We *can* go home at normal times today.
- Voice recognition is not quite "there" yet; especially when faced with language challenges in a noisy environment.

Whew!

Security ghost story

Questions:

- How would *you* react if this ghost typing happened to you?
- How long would it take *you* to figure this out if you received this support ticket?
- Is `Ctrl-Ctrl` a good user interface for activating and deactivating voice typing?

The End

for now...