



The CSC Kajaani datacenter



HEPIX Fall 2013 Ulf Tigerstedt

Wings of a butterfly

- 2007: Nurses union wanted more pay, mass resignations
- Paper industry was going fine, but demand going down
- Paper union wanted a lot more pay, the factories shut down
- Enter Google in Hamina
- Then current government wanted decentralization
- CSC needed new resources, extremely expensive to build in the capital area
- A bid was sent out for a suitable site for a new datacenter

Kajaani

- ➔ Kajaani won the bid
- ➔ Big paper mill, 40000 square meters, 10 m high
- ➔ Hydroelectric power close by
- ➔ Flowing water nearby for cooling
- ➔ Fiber already present nearby



Kajaani

- Situated in the middle of the county
- Population of 38k
- $64^{\circ}13'30''\text{N}$, $27^{\circ}44'00''\text{E}$
- Average temperature $+1,1 - 2^{\circ}\text{C}$
- July average $+15^{\circ}\text{C}$
- February average -17°C
- Average 75 cm of snow in March

We are however not Google

- One big cluster and one supercomputer won't fill a really big room
- Supercomputers have special cooling needs
- Cooling a big, half empty room is hard
- More tenants planned, what about fires?

Problems rack up

- ➔ Property owner could not manage it.



Alternatives on site

- ➔ Slab of asphalt
- ➔ Old paper warehouse



Alternatives on site

- ➔ Slab of asphalt
- ➔ Old paper warehouse

Modular Data Centre

Talk to vendors

- ➔ We can sell you anything, but it won't work.
- ➔ Problems:



Talk to vendors

- ➔ We can sell you anything, but it won't work.
- ➔ Problems:
 - Cold

Talk to vendors

- ➔ We can sell you anything, but it won't work.
- ➔ Problems:
 - Cold
 - Snow

Talk to vendors

- ➔ We can sell you anything, but it won't work.
- ➔ Problems:
 - Cold
 - Snow
 - Using water for cooling

MDC inside the paper warehouse

- MDC supplied by SGI (IceCube Air)
- Different containers for the super and the cluster
- Watercooled Cray XC30 in one
- Free air cooled HP Infiniband based cluster in another
- Storage from DDN
- All parts will be expanded, this was just Phase 1

Old datacenter in Espoo


- Hosted previously 3 HP clusters, one Cray XE4/XE5 hybrid and some storage.
- PuE of 1,20 when in full production during winter, 1,40 during summer.
- PuE now 1,80 after the Cray got shut down, management is not happy.

Cluster

- HP scaleouts: SL230s, 2*8 core Xeon
- Used both as a HPC cluster and cloud resources
- Will most likely be used for Alice and CMS at some point
- Was 9 weeks late due to the InfiniBand fabric by Mellanox not being stable

Lessons learned

- ➔ Don't do a project this scale again



Lessons learned

- Don't do a project this scale again
- Everybody outsources or subcontracts. In a rural location, all express couriers might use the same man with a van for the actual delivery. If they have attitude problems, you are screwed

Lessons learned

- Don't do a project this scale again
- Everybody outsources. In a rural location, all express couriers might use the same van for the actual delivery. If they have attitude problems, you are screwed
- Remote management works for computers, not humans (especially contractors)

Lessons learned

- Don't do a project this scale again
- Everybody outsources. In a rural location, all express couriers might use the same van for the actual delivery. If they have attitude problems, you are screwed
- Remote management works for computers, not humans (especially contractors)
- Mellanox really need a competitor