

HPC on the Cloud – A Norwegian Perspective

Nico Reissmann: nico.reissmann@ntnu.no

Jan Christian Meyer: jan.christian.meyer@ntnu.no

* Einar Næss Jensen: einar.nass.jensen@ntnu.no

<http://www.hpc.ntnu.no/>

Who we are

- NTNU IT Department
- Providing HPC services for
 - NTNU researchers and students on local clusters
 - Sintef
 - National infrastructure and users as part of Sigma2



NTNU HPC

What is HPC? (High Performance Computing)

- It depends.....
 - Lots of Memory
 - Blazing fast CPU
 - Millions of iops
 - Extreme Network
- All of the above x1000
- Basically:
 - When your laptop is not enough.
(my laptop have been running Fluent for 4 days and I cannot get any work done)
 - Or maybe you don't even have laptop.



So you need a new Computer

- Blazing fast monster machine!!!
 - 2 x 28 cores CPU
 - 1.5TB memory
 - 2x nvme drives
 - 25Gbit ethernet
- basically a new laptop
- with same problems



We need HPC



Lets pool our resources together

- and build a really, really large computer (Abel, Vilje, Stallo, Hexagon on top100 2012)

	Year	Cores	GPU	Interconnect
Idun (NTNU only)	2016	1700	62	FDR/EDR
Abel	2012	10000	X	FDR
Stallo	2012	10000	X	FDR
Hexagon	2011	22000	X	Gray Gemini
Vilje	2012	22000	X	FDR
Fram	2017	30000	X	EDR
Saga	2019	9000	32	FDR
B1	2020	???	???	???

B1 (the next big thing)

- Procured by Sigma2
- Process started early 2018
- Set to replace Vilje and Hexagon
- Goal was to get at least 50k - 60k cores with high speed interconnect and at least 2.5 PB storage.
- Budget: ca 90-100million NKR
- Public announcement soon(ish)

Cloud evaluation

- All purchases related to IT infrastructure must do evaluation on whether a cloud based solution is a viable alternative.
- Use Fram supercomputer as a reference and get public pricing from cloud vendors

FRAM

- 1 Gbit ethernet for admin
- EDR Infiniband 100Gbps interconnect, RDMA capable
- 2.4 PB parallell filesystem storage
- 150TB transfer of data within cluster each month
- at least 100TB outbound each month

Class	#nodes	Cores/nodes	Memory (GB)	Disk (GB)
Compute A	940	32	64	120
Compute B	8	32	512	2x960
Compute C	2	28	6000	14400
Frontend	10	32	128	2x800

Feature comparison

- Many cloud vendors. Investigated “the usual suspects”, Google (Cloud), Azure, Amazon, IBM (Softlayer). Only public available pricing considered.

	Compute A	Compute B	Compute C	Frontend	Storage	Network	RDMA
Fram	yes	yes	yes	yes	yes	yes	yes
Cloud	yes	yes	no	yes	no	no	no
AWZ	yes	unclear	no	unclear	unclear	unclear	no
Azure	unclear	unclear	no	yes	no	unclear	yes
Softlayer	yes	unclear	no	yes	no	no	no

Monthly cost comparison in Mnok

- The cost of Fram includes total HW cost, electricity, cooling, housing
- assuming four years life span
- When cloud alternative missing, assuming closes match
- The only alternative with RDMA alternative is 7 times more expensive

Vendor	Fram	Cloud	Aws	Azure	Softlayer
Cost	1,687	6,8311	7,6003	12,3383	11,3959

Should we do HPC in the cloud?


- Probably. But only for specific workload
 - “HPC”
 - long running, single core (or single node)
 - small input/output datasets
 - Not requiring expensive hardware
 - where cloud competition is good.
 - remember that HPC is moving target.

Hidden costs

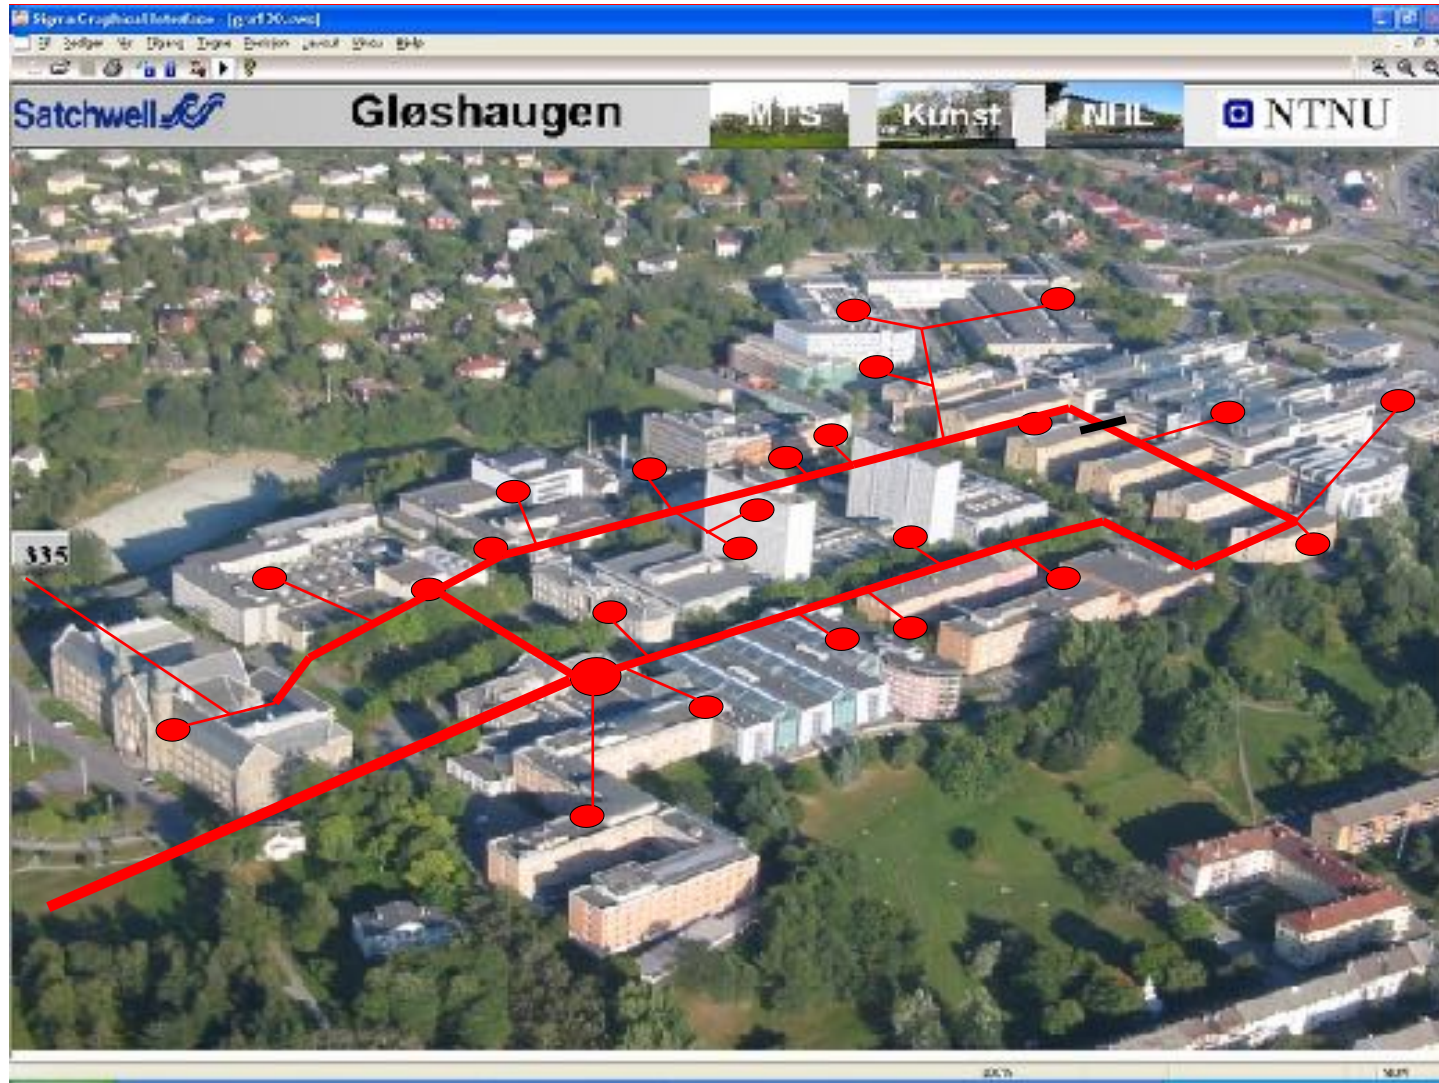
- Difficult to get exact pricing
- Asking for a deal will give huge discounts.
- Risk of “enslavement” to specific vendor
- GDPR and sensitive data?

and our hidden costs

- Takes more than two years for full process of procurement, installation and production
- Lots of resources required
- In April-May 2019, on average -250K cores
- Scientists have to wait



HPC bonus (NTNU) Heat Capture



Campus for the future

- 5.0-5.2Gwh each year (2014)
- Estimated price per kWh: 0.68nkr/0.1\$ (2014)
- We can do the same with Electricity/Power
- Heat is environmental waste, and needs to be recycled
- Heat, cooling, power, HPC consolidated in one building

More information

- Sigma2: <http://www.sigma2.no/>
- NTNU HPC: <http://www.hpc.ntnu.no/>

- einar.nass.jensen@ntnu.no