

# FGCI ARC tutorial



May 20th, 2016

Lecturers:

Luis Alves

Johan Guldmyr

Tomasz Malkiewicz

Kimmo Mattila

Ulf Tigerstedt (remotely)



# Program



09:30 – 10:00	Morning coffee & registration
10:00 – 10:15	Introduction: grid computing vs. local clusters, batch queue systems and cloud
10:15 - 11:00	Introduction to grids
11:00 - 11:30	Obtaining personal certificates hands-on
11:30 - 12:00	NorduGrid ARC middleware hands-on
12:00-12:45	Lunch (on own expense)
12:45 - 13:30	Writing job description files hands-on
13:30 - 14:00	Submitting and monitoring jobs hands-on
14:00-14:15	Coffee
14:15 - 14:45	Arcrunner grid job manager and MPI jobs on the grid hands-on
14:45 - 15:00	Wrap-up
15:00 - 15:30	Troubleshooter: Interactive session to deal with open questions and specific problems

# Practicalities



- Keep the name tag visible
- Toilets are in the lobby
- Network:
  - WIFI: eduroam, HAKA authentication
  - Ethernet cables on the tables
  - CSC-Guest accounts upon request
- Bus stops
  - Other side of the street (102,103) → Kamppi/Center (note, underpass)
  - Same side, towards the bridge (194,195,503-6) → Center/Pasila
  - Bus stops to arrive at CSC at the same positions, just on opposite sides
- *If you came by car: parking is being monitored - ask for a temporary parking permit from the reception (tell which workshop you're participating)*
- Visiting outside: doors by the reception desks are open
- Room locked during lunch
  - lobby open, use lockers
- Username and password for *workstations*: given on-site

# CSC?

- Non-profit company owned by Ministry of education and culture
- Services mainly free for researchers
- Applications, computational capacity, user support, FUNET, information management services, data services
- Participating in 15 EU projects



# Round robin



- ➡ Clément Fiere
- ➡ Matthew Hudson
- ➡ Khalid Latif
- ➡ Sofia Khan
- ➡ Pär Håkansson
- ➡ Trung Hieu Nguyen
- ➡ Markus Rauhalahti
- ➡ Eelis Solala






## Grid computing vs. local clusters, batch queue systems and cloud

# Learning target

- Know how to choose right server (resource)



# Clusters

- A cluster is a connection of separate units (nodes) via a fast network
- All larger CSC platforms (Sisu, Taito, FGCI) are clusters in a general sense





# Finnish Grid Infrastructure - FGCI

- Distributed computing capacity
- 12 universities/research institutes + CSC
- Requires a certificate
- Preinstalled software
- ARC –middleware
- From your own computer or Taito



```
arcproxy
```

```
arcsub jobscript.xrsl
```

```
arcget gsiftp://usva.fgi.csc.fi:2811/jobs/12465133890987654
```

- [FGCI guide](#)

# Server use profiles



- FGCI (Dell/HP)
- Serial and parallel (12/16/24 cores)

- cPouta (HP) Cloud
- Serial and parallel (16 cores)

- Taito-shell (HP)
- Interactive jobs
- Very long jobs
- Auto queue, shared resources

- Sisu (Cray XE40)
- Parallel from 72 up to thousands of cores
- Scaling tests 1008+

- Taito (HP)
- Serial and parallel upto 448/672 cores
- Huge memory jobs
- Lots of preinstalled software

# Main Computing capacity: Sisu,Taito, FGCI

	Sisu (Phase 2)	Taito (Phase 2)	FGCI prerelease
<b>Availability</b>	2014-	2015-	2016-
<b>CPU</b>	Intel Haswell and Sandy Bridge, 2 x 12 and 2 x 8 cores, 2.6 GHz, Xeon E5-2690v3 and E5-2670		Intel Xeon, 2 x 6 cores, 2.7 GHZ, X5650 and 4x12 Intel Xeon CPU E7-4830v3 @2.1GHz
<b>Interconnect</b>	Aries	FDR IB	QDR IB
<b>Cores</b>	40512	9768+9216	7308+3600
<b>RAM/node</b>	64 GB	64/128/256/ 1536 GB	128/256/512 GB
<b>Tflops</b>	1688	515	218
<b>GPU nodes</b>	-	50	?
<b>Disc space</b>	4 PB	4 PB	1+ PB

# Cloud computing: three service models

- **Software as a Service (SaaS)**
  - e.g. Chipster
- **Platform as a Service (PaaS)**
- **Infrastructure as a Service (IaaS)**
  - e.g. Pouta

# cPouta – computing in the Cloud

- cPouta
  - Virtual machines on demand
  - Taito hardware
  - Dedicated resources (HPC focus)
- More freedom and flexibility
- More responsibility



# cPouta's use cases




- Enhanced security – isolated virtual machines
- Advanced users – able to manage servers
- Difficult workflows – can't run on Taito
- Complex software stacks
- Ready made virtual machine images
- Deploying tools with web interfaces
- "We need root access"



*If you can run on Taito – run on Taito*  
*If not – Pouta might be for you*

- Pouta user guide: <https://research.csc.fi/pouta-user-guide>

- [illegible]




# Batch cont'd



Individual batch jobs


Batch job scheduler places jobs on compute nodes



# Compute nodes are used via queuing system

```
$ sbatch job_script.sh
```

```
$ ./my_prog &
```



# Using grid



- ➡ The jobs are submitted using the ARC middleware  
(<http://www.nordugrid.org/arc/>)
  - Using ARC resembles submitting batch jobs in Taito or Sisu
- ➡ ARC is installed Taito, but in many cases you can install it to your local machine too.



# Sample ARC job description file



```
&
(executable=runbwa.sh)
(jobname=bwa_1)
(stdout=std.out)
(stderr=std.err)
(gmlog=gridlog_1)
(walltime=24h)
(memory=8000)
(disk=4000)
(runtimeenvironment>="APPS/BIO/BWA_0.6.1")
(inputfiles=
( "query.fastq" "query.fastq" )
( "genome.fa" "genome.fa" )
)
(outputfiles=
( "output.sam" "output.sam" )
)
```

# Example serial batch job script on Taito

```
#!/bin/bash -l
#SBATCH -J myjob
#SBATCH -e myjob_err_%j
#SBATCH -o myjob_output_%j
#SBATCH --mail-type=END
#SBATCH --mail-user=a.user@foo.net
#SBATCH --mem-per-cpu=4000
#SBATCH -t 02:00:00
#SBATCH -n 1
#SBATCH -p serial
#SBATCH --constraint=snb
```