



High Performance
Computing &
Big Data Services

-  hpc.uni.lu
-  hpc@uni.lu
-  [@ULHPC](https://twitter.com/ULHPC)

Supercomputers introduction

Supercomputing Thursdays – 2023/04/20



Outline

- Why use a supercomputer?
- What is a supercomputer?
- How to use a supercomputer?

Why use a supercomputer?

Use cases

- running an experiment takes too long on your machine
- running one experiment is fine but you need to run 1000s of experiments
- running a small experiment is fine but real size experiment is too long or not possible

Why use a supercomputer?

Running a program takes too long on your machine

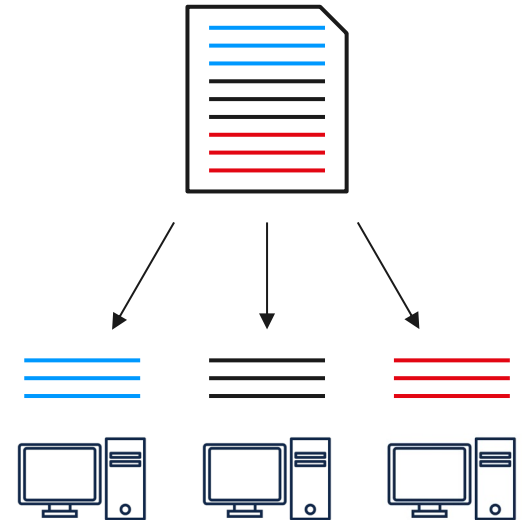
Example:

- Process one file that contains billions of lines
- Each line contain a text and we need to run an algorithm to detect its language

Solution:

- Distribute the lines onto multiple machines
- Each machine will process its own subset of lines

One program running sub-tasks on multiple computers



Split of the data for processing
on multiple computers

Why use a supercomputer?

Running one experiment is fine but you need to run 1000s of experiments


Example:

- Predict 3D shape of a protein from a chain of amino acid
- 16 000 amino acid chains to process

Solution:

- Run the protein folding program 16 000 times
- Distribute those runs on multiple machines

Many runs of the same program on a set of machines



Protein folding

Why use a supercomputer?

Running a small experiment is fine but real size experiment is too long or not possible

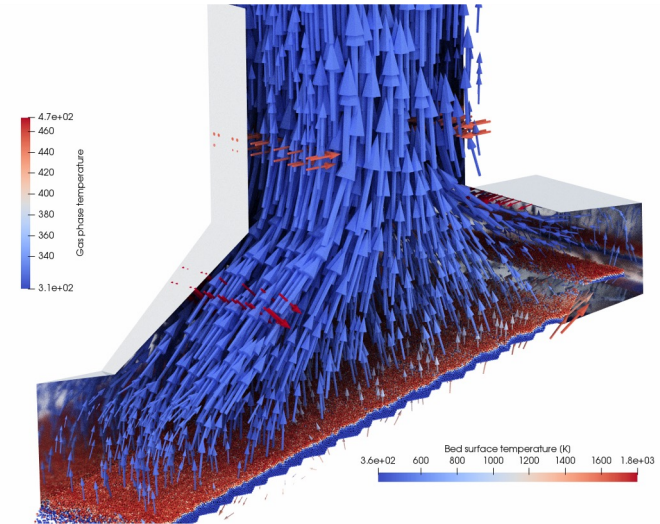
Example: simulation of Biomass Furnace

- Small number of particles can be simulated on a laptop
- Real life size number of particles requires too much memory and computation power

Solution:

- Distribute the computation over multiple computers
- Allow those computers to rapidly communicate with each other to coordinate the simulation

One program running on multiple computers acting as one



Biomass furnace simulation

What is a supercomputer?

No definitive definition, but usually:

- A lot of computational power
- Fast interconnect
- Fast and large storage(s)

What is a supercomputer?

No definitive definition, but usually:

- **A lot of computational power**
- Fast interconnect
- Fast and large storage(s)

What is a supercomputer?

One computer is not enough to solve lots of real life size problems.


Potential reasons:

- Too much time to compute
- Not enough memory to compute

Solution: **multiple computers** acting as one

Supercomputer terminology

- 1 computer = 1 computing node
- 1 Blade
 - 1 “shelf row”
 - 1 or more computing nodes



3 computers in 1 blade

Supercomputer terminology

- 1 computer = 1 computing node
- 1 Blade
 - 1 “shelf row”
 - 1 or more computing nodes
- 1 Rack = 1 “shelf” containing multiple blades
- Supercomputer = multiple racks



4 racks with 20 blades each (240 computers)

What is a supercomputer?

No definitive definition, but usually:

- A lot of computational power
- **Fast interconnect** (multiple computers acting as one)
- Fast and large storage(s)


What is a supercomputer?

Why do we need fast network?

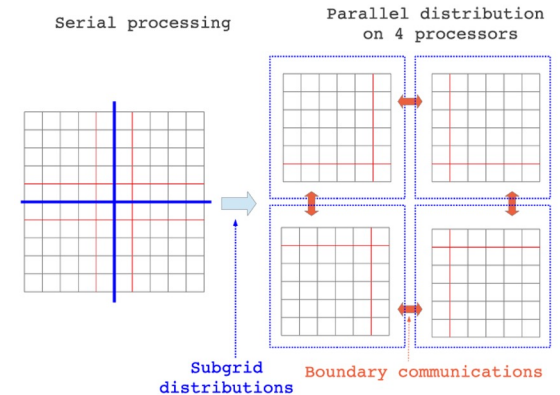
- Multiple orders of magnitude slower than a computer's memory
- Bottleneck to computer collaboration (scaling)
- Collaboration is necessary when a problem is decomposed in non-independent sub-problems

Characteristics of a fast interconnect:

- High bandwidth
- Low latency



Simulations of debris within a dam break flow



Sub-problems and necessary communications

What is a supercomputer?

No definitive definition, but usually:

- A lot of computational power
- Fast interconnect
- **Fast and large storage(s)**

What is a supercomputer?

- Usually in Petabytes (1000 TB, 1 000 000 GB), e.g. LUMI has 118 PB of storage.
- Technology that supports lots of concurrent users and offers, e.g. Lustre, GPFS
- Very dense appliances with lots of hard-drives



Storage appliance with 60 hard drives

What is a supercomputer?

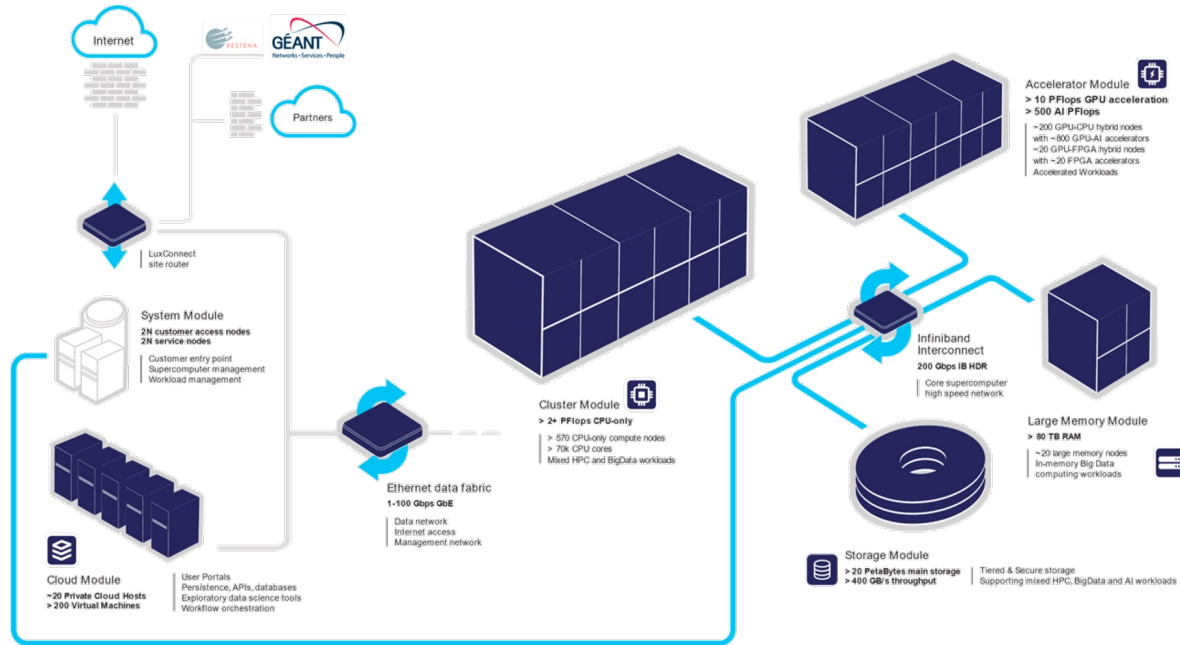
Historically:

- Lots of CPUs - homogeneous design

Nowadays:

- More specialisation - HPC centers have modules
- Accelerators:
 - GPU (very frequent)
 - FPGA (less frequent)
 - Quantum (rare but increasing interest)

Modular design - MeluXina



MeluXina modular design

How does a supercomputer look like?



MeluXina - Luxembourg



LUMI - Finland

I want to use a supercomputer, how can I do that?



Slightly different than using a personal computer

- Traditional way to use HPC is not interactive but batch jobs
- Current trend is a mix of interactive and batch
- What you cannot do: host a permanent web-service (Cloud services)

Two main ways to do this:

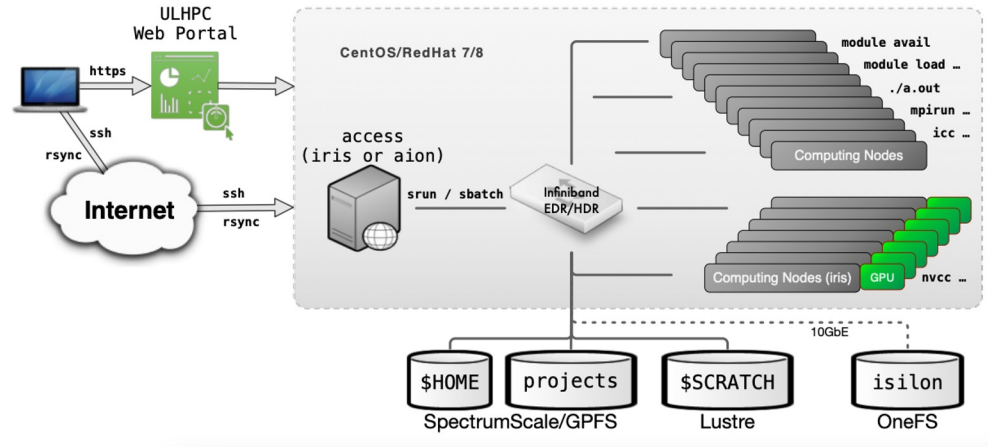
- Traditional SSH connection
- Web portal

I want to use a supercomputer, how can I do that?

General workflow:

- Connect to a login / access node to request some resources
- When the resources are granted you can use them

Example: you want to run a program A on the HPC



Architecture overview of a supercomputer

I want to use a supercomputer, how can I do that?

Traditional SSH connection

1. Transfer your program and / or your input files
2. Connect to a login / access node
3. Submit a job with your program
4. Monitor the job execution
5. Check your results on the HPC
6. Transfer your results on your personal computer

```
> rsync -avP my-program aion-cluster
building file list ...
1 file to consider
my-program
      0 100%   0.00kB/s   0:00:00 (xfer#1, to-check=0/1)

sent 127 bytes  received 42 bytes  338.00 bytes/sec
total size is 0  speedup is 0.00
```

Transfer files via command line

I want to use a supercomputer, how can I do that?

Traditional SSH connection

1. Transfer you program and / or your input files
2. **Connect to a login / access node**
3. Submit a job with your program
4. Monitor the job execution
5. Check your results on the HPC
6. Transfer your results on your personal computer

```
> ssh aion-cluster
=====
Welcome to access1.aion-cluster.uni.lu
=====

  Aion Cluster

=====
Atos BullSequana XH2000 Direct Liquid Cooling (DLC) supercomputer
https://hpc-docs.uni.lu/systems/aion/
=== Computing Nodes === #RAM/n === #Cores ==
aion-[0001-0318] 318 Atos X2410 AMD compute blade      256GB      40704
                (2 AMD Epyc ROME 7H12 @ 2.6 GHz [64c/280W])
=====
```

SSH connection to a login node

I want to use a supercomputer, how can I do that?

Traditional SSH connection

1. Transfer your program and / or your input files
2. Connect to a login / access node
3. **Submit a job with your program**
4. Monitor the job execution
5. Check your results on the HPC
6. Transfer your results on your personal computer

```
#!/bin/bash -l
#SBATCH -N 10
#SBATCH --ntasks-per-node=128
#SBATCH --time=01:00:00
#SBATCH -p batch
#SBATCH -J my-program-name

srun my-program
```

Simple job and its requested resources

```
0 [jschleich@access2 ~]$ sbatch job.sh
Submitted batch job 692495
```

Submission of the job

I want to use a supercomputer, how can I do that?

Traditional SSH connection

1. Transfer you program and / or your input files
2. Connect to a login / access node
3. Submit a job with your program
4. **Monitor the job execution**
5. Check your results on the HPC
6. Transfer your results on your personal computer

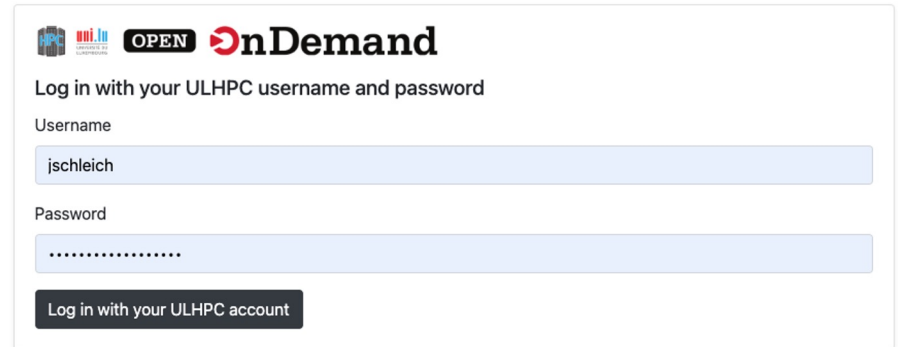
```
0 [jschleich@access2 ~]$ sbatch job.sh
Submitted batch job 692495
0 [jschleich@access2 ~]$ sq
# squeue -u jschleich
JOBID PARTIT  QOS          NAME          USER NODE  CPUS ST    TIME  TIME_LEFT  PRIORITY NODELIST(REASON)
692495  batch    normal    my-program-name  jschleich  10  1280  R    0:01    59:59    12112  aion-[0061,0064-0067,0082-0084,0319-0320]
```

Monitoring of the job

I want to use a supercomputer, how can I do that?

Access HPC through portal:

- Same idea but with a **web-interface**
- Often used to propose apps with a GUI



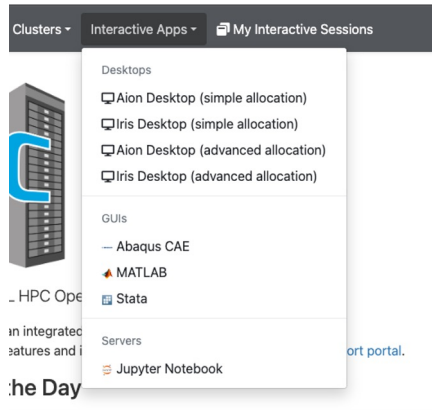
The screenshot shows the login interface for the OPEN OnDemand portal. At the top left is the uni.lu logo. To its right is the text "OPEN OnDemand" with "OPEN" in a black box and "OnDemand" in a red font. Below this is the instruction "Log in with your ULHPC username and password". There are two input fields: "Username" with the value "jschleich" and "Password" with masked characters ".....". At the bottom is a black button with the text "Log in with your ULHPC account".

Login of the web portal

I want to use a supercomputer, how can I do that?

Access HPC through portal:

- Same idea but with a web-interface
- Often used to propose apps with a GUI



HPC portal list of services

MATLAB version: 8f683da

This app will launch a **MATLAB** GUI on the **Iris/Aion** cluster. You will be able to interact with the MATLAB GUI through a VNC session.

Cluster

MATLAB version

This defines the version of MATLAB you want to load.

Account

Can be left blank and the default (PI) account will be used.

Number of hours

Number of cores

Number of cores on node type.

Number of GPUs

Allocate #GPUs (only valid on the *gpu* partition of Iris)

Example: using MATLAB on a HPC

I want to use a supercomputer, how can I do that?

Access HPC through portal:

- Same idea but with a web-interface
- Often used to propose apps with a GUI



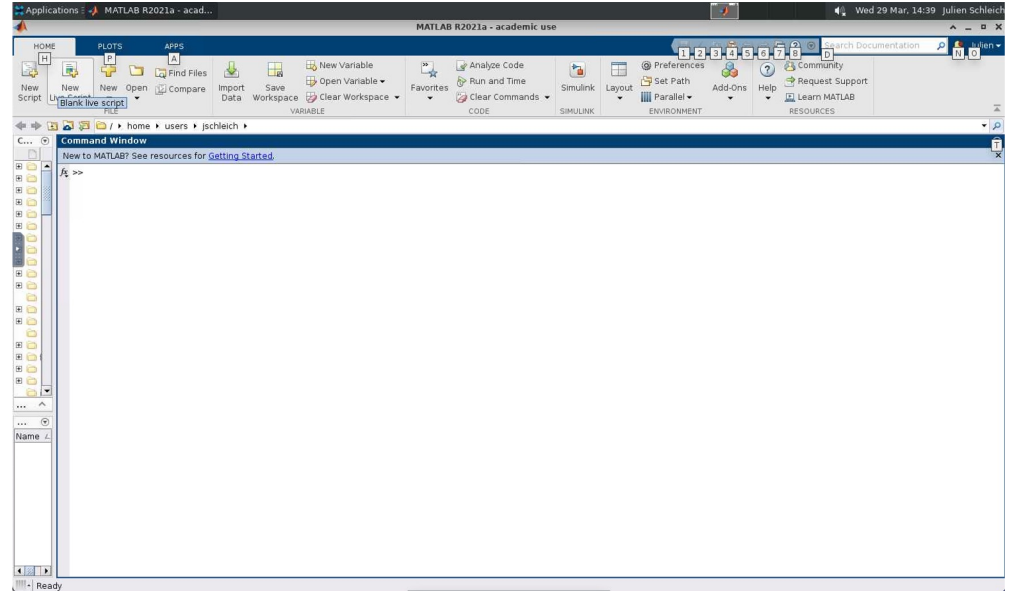
The screenshot shows a web interface for submitting a MATLAB job. At the top, it displays 'MATLAB (692864)' in a green header, followed by '1 node | 128 cores | Running'. Below this, the 'Host' is 'u_nion-0085' with a dropdown arrow, and a 'Delete' button is visible. The 'Created at' timestamp is '2023-03-29 14:38:38 CEST' and the 'Time Remaining' is '59 minutes'. The 'Session ID' is '8075ede6-569f-4430-8d2c-6589fd56a3f9'. There are two sliders: 'Compression' (0 (low) to 9 (high)) and 'Image Quality' (0 (low) to 9 (high)). At the bottom, there is a 'Launch MATLAB' button and a 'View Only (Share-able Link)' button.

Resource has been granted

I want to use a supercomputer, how can I do that?

Access HPC through portal:

- Same idea but with a web-interface
- Often used to propose apps with a GUI



MATLAB can now be used remotely and all computations run on the HPC

Thank you for attending the presentation

