

Mimir-2

High performance computer cluster for bioinformatic work

Multiple CPUs, large memory, high data input/output

Part of the IHPC resource cluster

Hosted at UTS UI

wiki: [irhpc.github.io](https://github.com/irhpc)

Snævar Sigurðsson, snaevar@hi.is

Elvar Örn Jónsson, elvarorn@hi.is,



HÍ

UPPLÝSINGATÆKNISVIÐ

Elja



(diligence, perseverance)

- Resources at IHPC/IREI:

Available to Icelandic universities and government research institutes.

Scientists in independent institutes funded by public grants.

Funded by Rannís/Innviðasjóður + Vegvísir

- Elja :

login-node “elja.hi.is”

controller-node (runs the show); deploys OS to compute nodes

HPC-Elja - 63 cpu nodes / 8 gpu nodes (+16 cpu nodes later this year)

HTC-Mimir - 9 cpu nodes / 1 high mem node

GPU8-Elja (coming this fall)

Elja



(diligence, perseverance)

Hardware:

- 4736 cores / 13 GPUs

28 x standard compute nodes:

- Partition : 48cpu_192mem
- 2 x Intel(R) Xeon(R) Gold 6248R CPU, 2 x 24 @ 3.00 GHz
- 192 (12 x 16) GB DDR4 @ 2667 MHz

44 x standard compute nodes:

- Partition : 64cpu_256mem, mimir
- 2 x Intel(R) Xeon(R) Platinum 8358 CPU, 2 x 32 @ 2.60 GHz
- 256 (16 x 16) GB DDR4 @ 2667 MHz

1 x high-memory compute node:

- Partition : mimir-himem
- 2 x Intel(R) Xeon(R) Platinum 8358 CPU, 2 x 32 @ 2.60 GHz
- 2048 (32 x 64) GB DDR4 @ 2667 MHz

3 x single GPU card nodes:

- Partition : gpu-1xA100
- 2 x Intel(R) Xeon(R) Gold 6248R CPU, 2 x 24 @ 3.00GHz
- 1 x NVIDIA A100 Tensor Core GPU
- 192 (12 x 16) GB DDR4, 2667 MHz

5 x dual GPU card nodes:

- Partition : gpu-2xA100
- 2 x Intel(R) Xeon(R) Gold 6248R CPU, 2 x 24 @ 3.00GHz
- 2 x NVIDIA A100 Tensor Core GPU
- 192 (12 x 16) GB DDR4, 2667 MHz

-
- InfiniBand (4 x EDR) Mellanox, 100 Gb/s
 - 3.5 Tb /scratch/ disk

HPC-Elja


Optimized software:

EasyBuild package manager:

- libraries+dep built against OS libs
- MPI w. full IB support

Test system: OpenMPI, OpenBLAS, ScaLAPACK, FFTW
GPAW Copper-filament benchmark

LUMI-C			ELJA-21			ELJA-22		
Sec	Cores	Nodes	Sec	Cores	Nodes	Sec	Cores	Nodes
121	48	3/8	204	48	1	74	48	1
109	64	1/2	110	96	2	43	96	2
53	128	1	89	144	3	36	144	3
29	256	2	41	288	6	19	288	6
19	384	3	32	384	8	13	384	8



LUMI:

- pre-exascale supercomputer for processing big data
- hosted in Finland

Would be nice to have reference for Mimir-I vs. Mimir-II?

File systems:

- **NFS server:**

User account home directory (/users/home/uname), SSH credentials
Good for personal compiled software / post processing / programming
~300-500 GB of data per active user

Not for large data! Slow I/O

- **NetApp:**

parallel-NFS server

Directly connected to Mellanox switches (~100Gb/s) : **fast I/O!**

HTC-Mimir users get access to : /hpcdata/Mimir/uname/ & /hpcdata/Mimir/shared/
hosted on a 100TB 'hot' disk

For large data!

Data left untouched for **7** days is moved to 'cold' disk ~1PB (~10Gb/s, not tested)

Data left untouched for **TBD** is deleted.

HTC-Mimir

44 x standard compute nodes:

- Partition : 64cpu_256mem, mimir
- 2 x Intel(R) Xeon(R) Platinum 8358 CPU, 2 x 32 @ 2.60 GHz
- 256 (16 x 16) GB DDR4 @ 2667 MHz

1 x high-memory compute node:

- Partition : mimir-himem
- 2 x Intel(R) Xeon(R) Platinum 8358 CPU, 2 x 32 @ 2.60 GHz
- 2048 (32 x 64) GB DDR4 @ 2667 MHz

Partitions:

[..]\$ sinfo

```
mimir*   up 14-00:00:0   1   mix  compute-71
mimir*   up 14-00:00:0   8   idle compute-[64-70,72]
himem-mimir up 14-00:00:0   1   idle mem-compute-1
```

Software stack:

- Easybuild 4.6.0

```
----- /hpcapps/lib-mimir/modules/all -----
Anaconda3/2021.11          HTSLib/1.14-GCC-11.2.0      file/5.41-GCCcore-11.2.0
BCFtools/1.14-GCC-11.2.0  Java/11.0.2                 (11)  kallisto/0.48.0
BEDTools/2.30.0-GCC-11.2.0 Java/15.0.1                 (D:15) matplotlib/3.4.3-foss-2021b
BamTools/2.5.2-GCC-11.2.0 MACS2/2.2.7.1-foss-2021b    megalodon/2.3.0
Biopython/1.79-foss-2021b R/4.1.2-foss-2021b          minimap2/2.24-GCCcore-11.2.0
Bowtie2/2.4.4-GCC-11.2.0 SAMtools/1.15-GCC-11.2.0  modbam2bed/0.6.2
CellRanger/7.0.0          TrimGalore/0.6.7           nanopolish/0.14.0
EasyBuild/4.6.0           (D)  bcl2fastq2/2.20.0-GCC-11.2.0 ont-guppy/6.2.1
FastQC/0.11.9-Java-11     cutadapt/4.1               picard/2.26.10-Java-15
HISAT2/2.2.1-gompi-2021b  deepTools/3.5.1-foss-2021b remora/1.1.1
```

[..]\$ module avail

[..]\$ module load Biopython

[..]\$ python version

Python 3.9.6

For Conda, Python & R see [wiki](#)

Your user account directory is : **/users/home/uname**

Your work/data directory is : **/hpcdata/Mimir/uname**

Sharing resources:

- Resources are allocated on a no. cpu cores / memory per cpu core
- Only allocate the number of cpu-cores you need :
[..]\$ srun ... --cpus-per-task **10** --mem-per-cpu **3900**

! 52 cpus are still available and 3900*52 Mb RAM !

Sharing data:

Option 1:

- move file(s) 'data' to /hpcdata/Mimir/shared/ : \$ mv data /hpcdata/Mimir/shared/.
- change ownership/perm :
\$ chown -R username:HTC-Mimir /hpcdata/Mimir/shared/data/
\$ chmod -R 771 /hpcdata/Mimir/shared/data/

Option 2:

- If you have sensitive data contact admins
 - special groups for select users can be made
 - or select users can be grouped together

How to apply for access?

HTC-Mimir:

- Send email to "help@hi.is" with the subject line "Access to the bioinformatic server Mimir 2"
- Include: Name, affiliation, your PI (professor / supervisor)
- Please put snaevar@hi.is in CC.

HPC-Elja:

- A formal application is sent by the **PI** to "help@hi.is" with subject line "Access to Elja HPC"
- Requires project description, approximate start / end date, intended use of the resources (software and cpu/gpu both?).
- Application is reviewed by the Steering Committee of the IHPC.
- See: <https://irhpc.github.io/docs/policies>

GPU8-Elja: TBD

User Agreement:

- [Here](#)

Disk policy:

- **Elja is not a data storage device!** After processing your data, please remove it from the NFS disk as a courtesy to your fellow coworkers / scientists.

NetApp policy:

- Need to decide on **policies** for /hpcdata/Mimir/