European Supercomputing (HPC) Resources

Ramon Brasser

Origins Research Institute

CSFK

My cluster

I have a small workstation cluster in the server room. 74 CPU cores, 9 consumer-grade GPUs.

CPU type and clock speed (Turbo) (GPU)

8C Intel i7-6900K @ 3.2 (3.7); RTX 2070 S

12C AMD R9-3900X @ 3.8 GHz (4.1); GTX 1080Ti (ID6), RTX 2070 (ID5)

6C Intel i5-8600K @ 3.6 (4.1); RTX 2070 S

6C Intel i5-9600K @ 3.7 (4.3)

12C AMD R9-3900X @ 3.8 (4.1); RTX 2070 S (ID5), RTX 2080 (ID8)

6C AMD R5-2600X @ 3.6 (4.0); RTX 3070

10C Intel i9-10850K @ 3.6 (4.6); RTX 3070 Desktop in my office

12C AMD R9-3900X @ 3.8 (4.1); GTX 2070 S (ID5), RTX 3070 (ID6)

12C AMD R9-3900X @ 3.8 (4.1); GTX 1650

Idle power usage: ~700 W.
Max load power usage: ~2.3 kW.

Electricity price: 357 Ft/kWh (insert sarcastic comment here)

Idle yearly electricity cost: 2.2 M HUF

Full load yearly electricity cost: 7.2 M HUF

Expensive! So what can I do to reduce this?

National HPC resources

Helyszín	Budapest	Budapest2	Szeged	Debrecen	Debrecen2-GPU (Leo)	Debrecen3-Phi (Apollo)	Pécs	Miskolc
Típus	HP CP4000SL	HP SL250s	HP CP4000BL	SGI ICE8400EX	HP SL250s	HP Apollo 8000	SGI UV 1000	SGI UV 2000
CPU-k / node	2	2	4	2	2	2	192	44
Core-ok / CPU	12	10	12	6	8	12	6	8
Memória / node	66 GB	63 GB	132 GB	47 GB	125 GB	125 GB	6 TB	1.4 TB
Memória / core	2.6 GB	3 GB	2.6 GB	2.6 GB	7 GB	5 GB	5 GB	3.75 GB
CPU	AMD Opteron 6174 @ 2.2GHz	Intel Xeon E5-2680 v2 @ 2.80GHz	AMD Opteron 6174 @ 2.2GHz	Intel Xeon X5680 @ 3.33 GHz	Intel Xeon E5-2650 v2 @ 2.60GHz	Intel Xeon E5-2670 v3 @ 2.30GHz	Intel Xeon X7542 @ 2.66 GHz	Intel Xeon E5-4627 v2 @ 3.33 GHz
GPU	-	-	2 * 6 Nvidia M2070	-	68 * 3 Nvidia K20x + 16 * 3 Nvidia K40x	-	-	-
Intel Xeon Phi (KNC)	-	14 * 2 * Intel(R) Xeon Phi(TM) MIC SE10/7120 Intel Xeon Phi	-	-	-	45 * 2 * Intel(R) Xeon Phi(TM) MIC SE10/7120	-	-
Linpack teljesítmény (Rmax)	5 Tflops	27 Tflops	20 Tflops	18 Tlops	254 Tflops	~106 Tflops	10 Tflops	8 Tflops
Compute node-ok száma	32	14	50	128	84	45	1	1
Dedikált storage	50 TB	500 TB	250 TB	500 TB	585 TB (Phi-vel közös)	585 TB (GPU-val közös)	500 TB	240 TB
Interconnect	IB QDR	IB NB FDR	IB QDR	IB QDR	IB NB FDR	IB NB FDR	Numalink 5	Numalink 6
Scheduler	SLURM	SLURM	SLURM	SLURM	SLURM	SLURM	SLURM	SLURM
MPI	OpenMPI (ompi)	IntelMPI (impi)	OpenMPI (ompi)	SGI MPT (mpt)	OpenMPI (ompi)	OpenMPI (ompi)	SGI MPT (mpt)	SGI MPT (mpt)

These HPC resources are very outdated. Debrecen (CPU), Debrecen-2 (GPU) and Miskolc (CPU) I would consider, but there are far better options elsewhere.

PRACE

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

(Partnership for Advanced Computing in Europe)

- https://prace-ri.eu
- Seeks to facilitate access to research infrastructure.
- 25 member countries, and 5 hosting countries. Member countries provide national resources (KIFÜ in Hungary), and some pan-national resources.
- EC funded.
- Training (coding) and user support
- Infrastructure support (for admins)
- Events (PRACEDays, summits, programming
- guides etc)
- Check it out

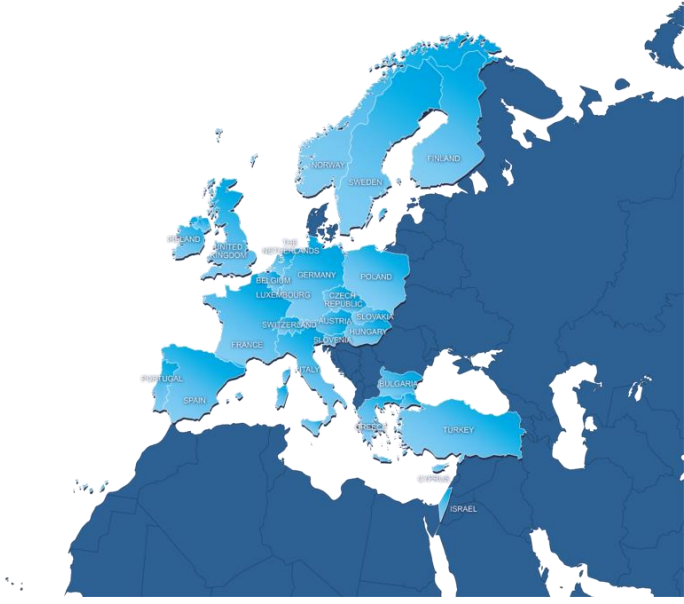

PRACE Resources

- HPC systems in France, Germany, Italy, Spain and Switzerland (blue = GPUs)
- *Free* access for academics at an EU-based institution
- Project Access applications 2/yr (deadlines: March & September).
- Online proposal submission through https://pracecalls.eu
- **Minimum** request 15-35M CPU core hours.
- Focused on large, massively parallel simulations.
- Preparatory Access available (code testing and debugging). Always open, quarterly cut off dates – 1st working day of the quarter (March; June; September; December).
- Duration: 1 yr.
- Hardware: Intel Skylake 2x24 @ 2.7 GHz, AMD Epyc 2x64 @ 2.5 GHz. Nvidia P100 or V100.

EuroHPC Resources

- https://prace-ri.eu/hpc-access/eurohpc-access/
- **Very new** systems in Bulgaria, Czechia, Finland*, Luxumbourg, Slovenia. Pending: Italy, Spain, Portugal.
- Online proposal submission through https://pracecalls.eu
- Benchmark and Development access (call deadlines every 1st of the month)
- Regular Access (deadlines March, July, November)
- RA: **Minimum** 10M CPU/1M GPU hours.
- RA: Proposal is MS Word, 10 pages, much technical info required.
- **Strongly suggested** to get Benchmark access first to get familiar with the system, run scaling tests, and only then write Regular Access proposal.
- Exascale Access (new!)
- EA: Minimum 20M (Italy; new), 80M (Finland).
- EA: Same proposal format as RA.
- Duration: 1 yr
- Hardware: AMD Epyc 2x64 @ 2.6 GHz or 2.45 GHz. Nvidia A100 or AMD MI250 (LUMI)

*LUMI in Finland uses AMD GPUs rather than Nvidia. Largest supercomputer in the EU.

Collaborative Calls

(Interactive Computing E-Infrastructure)

- https://prace-ri.eu/hpc-access/collaborative-calls/
- Provided by the Fenix institute, part of the Human Brain Project.
- Systems in France, Germany, Italy, Spain, Switzerland.
- Call deadlines: February, April, July, October (approximately).
- Various services (scalable, interactive, storage).
- **Minimum** request for HPC: 5M CPU hours.
- Proposal: MS Word, 4-8 pages. Not too technical. Submitted through e-mail.
- Duration: 1 yr.
- Hardware: AMD Epyc 2x64 @ 2.25 GHz or Intel Xeon 2x28 @ 2.2 GHz. Nvidia V100.
- Their administrative support is excellent!
- Fun fact: the scientific supervisor at the German JUSUF cluster is Susanne Pfalzner, who works in astrophysics and planetary science.

My experience

- Awarded 6.6M (2022) and 8.4M (2023) CPU hours on JUSUF in Germany through PRACE-ICEI. Have only used GPUs so far; CPU usage on JUSUF is coming.
- Awarded 1.6M (2022) GPU hours on VEGA in Slovenia through EuroHPC. Second proposal for 2.2M (2023) GPU hours pending. No CPU usage.
- Awarded Benchmark access to VEGA in Slovenia and KAROLINA in Czechia in November 2021.
- KAROLINA access is complicated and tedious and requires an e-mail certificate issued by KIFÜ (IT Team can confirm how difficult this was). VEGA and JUSUF access is straightforward, only requires SSH keys.
- VEGA and JUSUF use the SLURM scheduler; KAROLINA uses PBS.
- VEGA technical support is fast and excellent! JUSUF is similar, but I have less experience. In contrast, my experience with KAROLINA was not good. I **cannot** recommend using this system.
- Have not tried other systems or written a PRACE Project proposal because they do not meet my needs.
- Never been in the queue for GPU usage longer than 12 hours on both systems.
- All awarded projects require a written report at the end no later than 3 months after the project has ended.
- All awarded projects should be acknowledged in resulting publications. Details on the websites.
- We should not all compete with each other on the same systems!

Public HPC limitations

- Limited running time (typically 2-7 days) → Requires frequent starting and stopping of the simulations, which can accumulate errors. Also need to find a way to have the submission script resubmit itself.
- Much lower CPU clockspeed; ~3 GHz on HPC clusters vs. ~4 GHz on desktop PCs.
- Project duration 1 yr, so sims that take longer need different resources.
- Often oversubscribed, so sometimes have to wait in the queue for >24 h.
- Proposal and report writing 

PRACE-ICEI storage options

Archival data repositories

Component	Site (Country)	Minimum request	Total Resources for Call 11	Unit
Archival	CEA (FR)	1	1 125	TByte
Archival Data Repository	CSCS (CH)1	1	380	TByte
Archival Data Repository	CINECA (IT)	1	300	TByte
Active Archive 2	BSC (ES)	90	440	TByte

Active data repositories

Component	Site (Country)	Minimum request	Total Resources for Call 11	Unit
HPST@JUELICH	JSC (DE)	10	300	TByte
Lustre Flash	CEA (FR)	10	106	TByte
Data Warp	CSCS (CH)1		5.5	TByte
HPC Storage @ CINECA	CINECA (IT)	1	1 520	TByte
HPC Storage @ BSC	BSC (ES)	2.5	2.5	TByte

The PRACE-ICEI programme also offers archival storage options.

I have no experience with this, but I wanted to mention it here for those of you who are interested.

Some examples of online storage

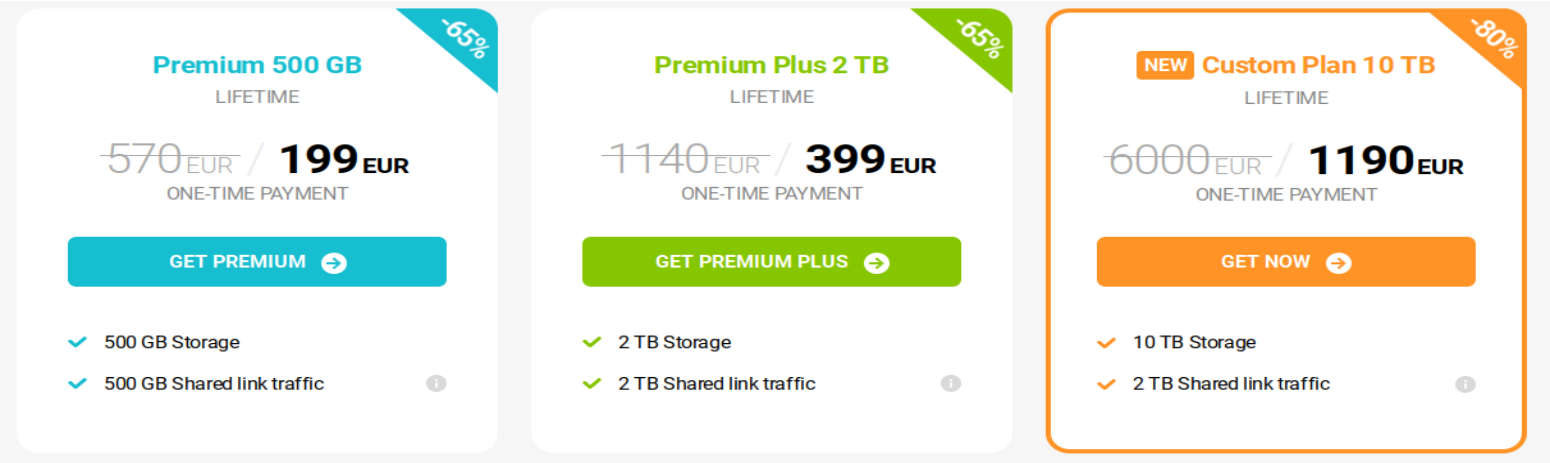

pCloud (Switzerland)

Icedrive (UK)

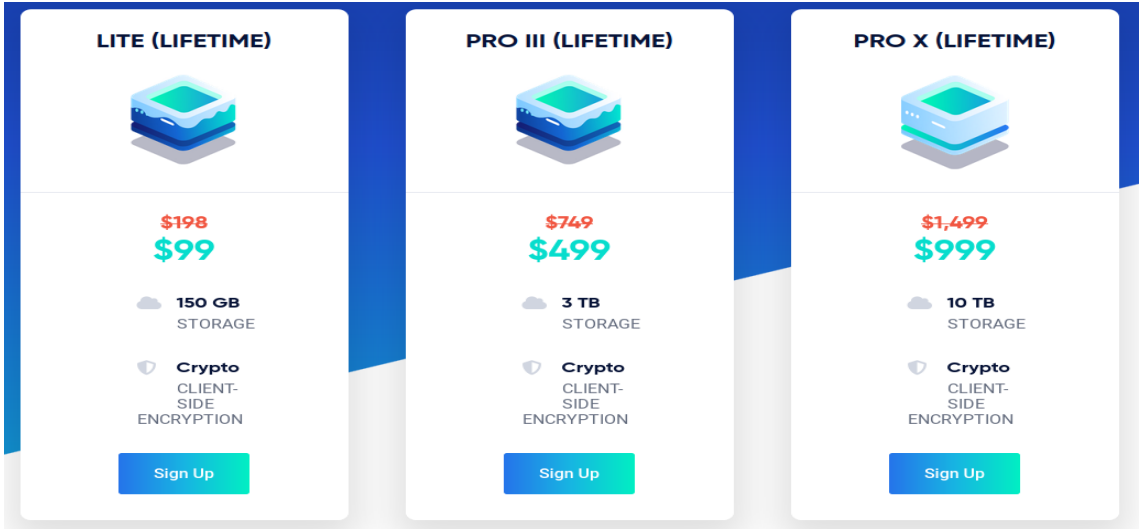

These are just some examples.

Due to privacy and surveillance, it's best to choose EU-based providers (avoid USA, China, Russia, possibly UK)

Commercial options

(rough estimate)

- Amazon AWS EC2
- 1 VM with 32 CPUs continuous usage
- USD 6.3k-7.2/yr (HUF 2.3M+).

- Google cloud
- 1 VM with 32 CPUs continuous usage
- EUR 5.3k/yr (HUF 2.1M).

- Microsoft Azure
- 1 VM with 32 CPUs continuous usage
- EUR 15k/yr (HUF 6M).

(Hey, it's blooming M$ after all...)

All of these are very expensive for continuous usage such as mine, but for spot/spike usage (need lots of CPUs now for a short time) these options are viable. Pricing usually USD 1-2/hr for 32 CPUs.

For example, it would cost USD 90 for 192 sims lasting about 13 hours, assuming they can all run at the same time (which they may not).