Getting Access to LUMI

Kurt Lust LUMI User Support Team (LUST) University of Antwerp

June 2025

EuroHPC supercomputers

- 5 petascale supercomputers: Meluxina, VEGA, Karolina, Discoverer, Deucalion
- 3 pre-exascale computers: LUMI, Leonardo and MareNostrum5
- Two exascale computers under constructions or procurement: Jupiter and Alice Recoque
- 2 new mid-range systems coming up: Arrhenius and Daedalus
- Future: "AI-optimised" supercomputer
- All are a joint investment of EuroHPC with one or more countries

LUMI: Who pays the bills?

- EuroHPC machine so joint funding of:
 - EuroHPC JU (50%)
 - Consortium of 11 countries (The Netherlands joined after the start)
- The resources of LUMI are allocated proportional to the investments
- Each LUMI consortium country sets its own policies for a national access program
 - o See <u>www.lumi-supercomputer.eu/get-started</u>
- So LUST does not manage access to LUMI!



Projects and users

LUMI

• A project

- Corresponds to a coherent amount of work done by a single person or a collaboration between a group of users.
 - Typically a research project
 - Project for a course
 - Some projects for organisational issues, e.g., local support team project
- The basis for most resource allocations on LUMI
 - Compute budget: CPU core-hours for LUMI-C, GPU hours for LUMI-G and visualisation nodes
 - Storage budget: Expressed in TB hours
 - Budgets are assigned and managed by the resource allocators, not by the LUMI User Support Team
- LUMI projects: project_465XXXXXX or project_462XXXXXX (Finland only)
 - This is the number that you should mention when contacting LUMI User Support

Projects and users (2)

LUMI

• A user account

- One physical person per account
 - Do not share accounts!
- Some physical persons have more than one account
 - An unfortunate consequence of decisions made very early on in the project
- Needs a project to do anything useful on LUMI
- Many-to-many mapping between projects and user accounts
 - Projects can of course have several users who collaborate
 - Users can be a member of multiple projects (and this is more common than you think)
- Resources:
 - Mostly attached to projects
 - Bare minimum for user accounts: just a fixed size home directory

Projects management

- Different systems in different countries
 - Finland: MyCSC, completely independent management
 - Other countries and EuroHPC projects are managed via Puhuri
 - Web-based portal developed by the Nordic countries for project and resource allocation management (and not just for LUMI)
 - Some countries have their own front-end, other countries use a Puhuri front-end
 - Login to Puhuri via MyAccessID
 - MyAccessID is a GÉANT service that then interfaces with your institute identity provider and several alternatives
 - Always use the same credentials!
 - This is also the place for ssh key management for Puhuri projects
- Quick check of your resources on the system command line: lumi-workspaces

File spaces – User-specific

- Home directory: /users/<my_uid>
 - Limited in size and not extensible
 - Should be used only for very personal stuff: user-specific configuration files, etc.
 - Not meant as a way to transfer data to future projects
 - Not billed

File spaces – Project based (1)

- All billed against the storage budget
- Permanent storage in /project/project_46YXXXXX
 - For historical reasons, also appears as /projappl/project_46YXXXXX
 - Place for, e.g., software installations, permanent input data sets
 - Billed at 1TB-hour per TB per hour used
- Disk-based scratch storage in /scratch/project_46YXXXXX
 - May be erased after 90 days, but this is not active
 - Billed at 1TB-hour per TB per hour used

File spaces – Project based (2)

- Flash-based scratch storage in /flash/project_46YXXXXX
 - May be erased after 30 days, but this is not active
 - Billed at 3 TB hour per TB per hour used
- Permanent object storage (LUMI-O)
 - Billed at 0.25 TB hour per TB per hour used

File spaces - Quota

Goal	Where?	Capacity	Files	Retention
User home	/users/ <username></username>	20 GB	100k	User lifetime
Project persistent	/project/ <project></project>	50-500 GB	100k	Project lifetime
Project scratch	/scratch/ <project></project>	50-500 TB	2M	90 days (not active)
Project fast scratch	/flash/ <project></project>	2-100 TB	1M	30 days (not active)

- Flexibility in block quota (within limits) but less flexibility in file quota
 - See day 2 session: Big parallel file systems don't like small files
 - Singularity containers should be used for software installations with lots of small files
 - Quota extensions currently done by the LUMI User Support Team

File spaces – Further information

- 4 disk based file systems for /users, /project and /scratch
 - Your user home directory may be on a different file system as your /project and /scratch directory
 - And no, the LUMI User Support Team cannot change that
- /flash is also a parallel file system...
- LUMI is not a data archiving or data publishing system
 - "Permanent" = for the duration of the project
 - Data that is not needed anymore should be moved to your home institute or an archiving service
 - No backup. Repeat: NO BACKUP.

Access

- 4 login nodes accessible via key-based ssh
 - Generic name: lumi.csc.fi
 - Specific login nodes: lumi-uano1.csc.fi, lumi-uano2.csc.fi, lumi-uano3.csc.fi, lumiuano4.csc.fi
 - May be needed for tools for remote editing etc.
 - Key management:
 - Most users: Via MyAccessID: <u>mms.myaccessid.org</u>
 - Users who entered first via CSC: my.csc.fi
- Web interface via Open OnDemand: <u>www.lumi.csc.fi</u>
 - Own set of login nodes
 - Simple GUI environment via the "Desktop" app, based on VNC
- Little support for GUI applications on LUMI through other technologies
 - X11 over ssh is unbearibly slow for most users
 - Additional primitive VNC support outside of OOD (lumi-vnc module)

Open OnDemand (1)

LUM



Open OnDemand (2)



Open OnDemand (3)

Pinned Apps					
Home Directory	Compute node shell	Login node shell	Desktop	Cloud storage configuration	
Disk quotas	Project view	C Active Jobs	Jupyter Jupyter	Jupyter for courses	
Julia-Jupyter	MATLAB	MLflow	TensorBoard	Visual Studio Code	

Open OnDemand (4) – Login node shell



LUMI uses warm-water cooling, which enables its waste heat to be utilized in the district heating network of the city of Kajaani, and thus replaces heat produced by fossil fuels.

The waste heat from LUMI that can be used in Kajaani's district heating network is equivalent to up to 20 per cent of the energy that Kajaani needs to use in the area's district heating. This reuse of waste heat will reduce the annual carbon dioxide footprint of Kajaani by 13,500 tons - an amount that equals the output from 4000 passenger cars.

There are storage guota and storage billing unit warnings: WARNING: project 465001361 is out of storage hours WARNING: project 465001363 is out of storage hours

login nodes

[lumi][kulust@uan09-1000 ~]\$

Open OnDemand (5)

Pinned Ap	ps			
Home Director	ry Compute node shell	Login node shell	Desktop	Cloud storage configuration
Disk quotas	Project view	C Active Jobs	Jupyter Jupyter	Jupyter for courses
Julia-Jupyter	MATLAB	MLflow	TensorBoard	Visual Studio Code

Open OnDemand (6) – Home Directory

O Dashboard - LUMI.csc.fi ×	+			~
\leftrightarrow \rightarrow C \bigcirc lumi.csc.fi/pun/sys/dasht	board/files/fs/users/ku	lust		☆ ⊉
LUMI Files - Jobs -	Apps - Tools	- j		0- ≛ ↔
	>_ Ope	n in Terminal 🕨 📿 Refresh 🕇 New File	🗈 New Directory 🚨 Upload 🛓 Download	Copy/Move Tolete
A Home Directory		re / kulust / 🕞 channe diasa		(Converte
/projappl/project_465001603	↑ User	Change directory		Copy path
🖿 /projappl/project_465000095			□ Show Owner/Mode □ Show Dotfiles	Filter:
/projappl/project_465000844	•	This app does no	ot run in the conte	xt of a job
🖿 /projappl/project_465001361				
/projappl/project_465001362	•	Can also be used	a for uploading and	a downloading data,
/projappl/project_465001363		but will fail for b	oig transfers	
/scratch/project_465001603		60 01 0 C		IT E EDED ISIDE: I I
/scratch/project_465000095		Desktop	-	11-4-2024 12:24:45
/scratch/project_465000844		Documents		11.4.2024 12:24:51
/scratch/project_465001361		Documents		11-7-2024 12.24.31
/scratch/project_465001362		Downloads	; • -	26-5-2023 18:20:10
/scratch/project_465001363		EasyBuild	i .	7-10-2024 11:08:13
/flash/project_465001603			<u> </u>	

LUM

Open OnDemand (7)

Pinned Apps					
Home Directory	Compute node shell	Login node shell	Desktop	Cloud storage configuration	
Disk quotas	Project view	C Active Jobs	Jupyter	Jupyter Jupyter for courses	
Julia-Jupyter	MATLAB	MLflow	TensorBoard	Visual Studio Code	

Open OnDemand (8) – Desktop app

 millosom/punjoya/daanboard/batch_ct	nineedofofood moldendulfaeoonin_contextolinem	A G I (* :
	Partition	
A MATLAB	lumid	L
🎓 TensorBoard	Runs as a jo	D
🗧 Jupyter	Resources	
Tools	Number of CPU cores	
>_Compute node	2	
shell	SMT is enabled for the selected partition. 2 threads per	
Course environments	core will be allocated.	
📁 Jupyter for	Memory (GB)	
courses	2	
	Number of GPUs (A40)	
	1	
	Time	
	4:00:00	
	d-hh:mm:ss, or hh:mm:ss	
	Launch	
	Deset to default settings	

Open OnDemand (9) – Desktop app

```
My Interactive Sessions - LUN ×
                                                                                                                                  ☆ ▷ | ① :
\leftarrow \rightarrow
       C
            25 lumi.csc.fi/pun/sys/dashboard/batch_connect/sessions
                                                                                           Inode | 4 cores | Running
                                                   Desktop (8630905)
                          Apps
                          Editors
                                                   Host: >_nid000023
                                                                                                              8 Cancel
                          🗙 Visual Studio
                          Code
                                                   Created at: 2024-12-05 10:13:40 UTC
                          Graphical applications
                                                   Time Remaining: 3 hours and 59 minutes
                         Desktop
                                                                                      Runs as a job
                                                   Session ID: c42cb9ed-alc6-4c15-
                                                                                ۲
                          Servers
                                                                                      VNC-based with a choice of
                                                                                •
                                                   If you run into issues, please inclu
                         🔉 MLflow
                                                   output.log
                                                                                      connection options, including just the
                          🔒 Julia-Jupyter
                                                   Project: project_465001603
                                                                                      web browser
                                                   Partition: lumid
                          📣 MATLAB
                                                   Cores: 4
                                                   Memory: 2048M
                          TensorBoard
                                                   GPUs (A40): 1
                          🝯 Jupyter
                                                     noVNC Connection
                                                                         Native Instructions
                          Tools
                                                   Compression
                                                                                       Image Quality
                         >_Compute node
                                                   0 (low) to 9 (high)
                                                                                       0 (low) to 9 (high)
                          shell
                                                    Launch Desktop
                                                                                               View Only (Share-able Link)
                          Course envi
                          Jupyter for
                          courses
```

Open OnDemand (10) – Desktop app

LUMI



) Menu 📃

Data transfer

- sftp to the login nodes
 - Authentication with your ssh key
 - Can be slow on high latency connections
 - Slow connections are not the fault of LUMI but of the whole path to the machine
- Data transfer via the object storage system LUMI-O
 - Transfer to LUMI-O and then to other LUMI file systems
 - Or from the file systems of LUMI to LUMI-O and then to your home institute
 - Support for various tools including rclone and S3 commands
 - Multi-stream transfers are a way to deal with high latency
 - See the <u>storage section of the LUMI documentation</u> at <u>docs.lumi-supercomputer.eu</u> and the presentation on day 2 of this course
- Unfortunately no support yet for Globus or other forms of gridFTP



Questions?