

Table of Contents

- Introduction
 - Benefits for users
 - Benefits for for Faculty/Research Groups
 - Informations
- How To connect
- Tabs
 - Dashboard
 - File Manager
 - jobs
 - Active jobs
 - Job composer
 - Interactive Apps
 - HPC Usage Dashboard
 - Quick Overview of Your HPC Environment
 - Slurm Account Information
 - Group and Partition Access Overview
 - Intra-PI Consumption Comparison
 - SandBox
- Logout
- Launching applications
 - Session persistence
- Contribution

Introduction

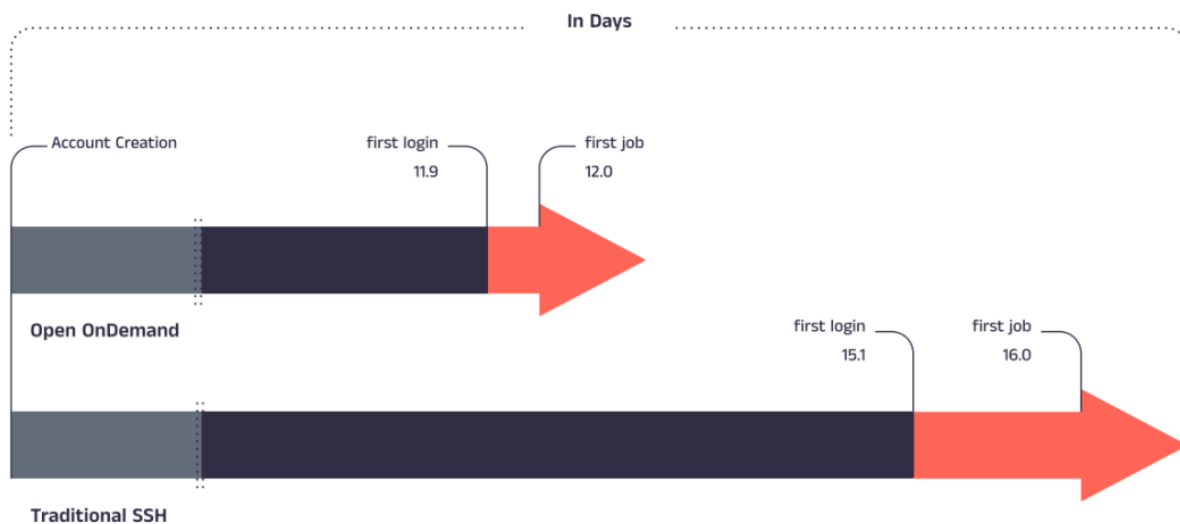
OpenOnDemand makes it easy for users to run jobs with graphics, such as JupyterLab, MATLAB, or a desktop. No need for complex setups like ssh tunneling or tools like x2go.

We are confident that this new tool will prove valuable to many individuals, especially those who are beginners in or have limited experience in computing. Its primary aim is to provide a workspace centered around your research project while striving to relieve you of as many technical requirements as possible.

Benefits for users

Instead of going to the lab or installing client software, users can access supercomputers through the web. They can compute from anywhere, on any device, regardless of skill level.

According to the OpenOnDemand team, New users get up to speed faster:



OnDemand users start work faster than traditional users, both in terms of first login and job submission.

We have tested it, and we concur that even novice users, without extensive computer skills, will find it easy to write code, submit jobs, or engage in interactive tasks (such as MatLab, JupyterLab, etc.) with ease.

Benefits for Faculty/Research Groups

OpenOnDemand offers significant advantages by streamlining job submission processes, with a key feature being the ability to create job templates. These templates simplify task submission by capturing frequently used parameters, reducing manual input and saving valuable time. They also ensure consistency and accuracy in job submissions, enhance customization options, and facilitate easy sharing among collaborators, ultimately boosting research productivity.

Informations

Deployment has just begun, and we are currently in the testing and user observation phase. Every piece of feedback will help us to identify user needs, which in turn will enable us to evolve the platform to better meet those needs. Your input is greatly appreciated, as we work towards continuous improvement.

How To connect

First each cluster will have his own OpenOnDemand instance. To allow all users (unige, extern and outsider users) to connect to OpenOnDemand, you must to authenticate through switch.

You MUST use the same email bound to your HPC account. It means:

- Unige email (@unige.ch) for unige members (Student/collaborator/external)

- Email used for registration for Outsider members

If not, you will not be able to connect through OpenOnDemand web Interface

Access to Baobab OpenOnDemand web instance: <https://openondemand.baobab.hpc.unige.ch>

Access to Bamboo OpenOnDemand web instance: <https://openondemand.bamboo.hpc.unige.ch>

Access to Yggdrasil OpenOnDemand web instance: <https://openondemand.yggdrasil.hpc.unige.ch>

Tabs

Dashboard

The **Dashboard** tab serves as a central hub and the homepage within the OnDemand web app, offering users quick access to essential tools and services through pinned apps. Additionally, it provides crucial cluster information, including status updates on any ongoing issues, their resolutions, and other informative notices. Users can also find the 'Message of the Day' section, featuring the latest announcements and important updates to keep them informed about the platform and cluster developments.

File Manager

The **File Manager** tab serves as your centralized workspace for all your files within the OnDemand application. Here, you can view all the files and create, copy, move, or delete.

It provides a user-friendly interface for efficiently organizing and managing your data. Additionally, you have the option to open a Linux terminal directly from this section, enabling you to execute command-line operations for advanced tasks.

For quick access, shortcuts to the **Home directory** and **Scratch directory** are also available, simplifying file navigation and work with your data.

jobs

Active jobs

The **Job Activity** tab aggregates all job activities of the day, presenting the same information as obtained from the `sacct` command. You can expand the entries to access more details about a particular job, such as queued jobs that have not yet started, potentially due to constraint violations.



While you have permission to view job information of other users, this privilege is not meant

for surveillance or reporting on users you suspect of wrongdoing. If you have any concerns, please consider posting them on the HPC Community forum with courtesy and respect. The HPC team will respond, offering guidance on best practices and will contact the user if necessary.

Job composer

The Jobs Composer is a convenient tool designed to assist new users in getting started with HPC. However, it sometimes operates like a black box, leaving users unsure about its inner workings.

For instance, attempting to include the `##SBATCH --array=` option directly in the `sbatch` command won't work, as this option is available within **Job Options** subtab.'

Interactive Apps

The **Interactive Apps** tab serves as a centralized hub for all available applications, For now, JupyterLab, a desktop environment, and MatLAB are available. These tools are provided to enhance user productivity and convenience. Explore a range of interactive applications to support your research tasks, all within easy reach from this tab.

HPC Usage Dashboard


The HPC team has released a new dashboard integrated directly into OpenOnDemand. This dashboard provides a clear graphical overview of your resource usage and helps you better understand how you consume HPC resources.

The dashboard is designed to help you with a quick and simplified view:

- Track and understand your HPC usage
- Identify your accounts and partitions
- Compare consumption patterns within your PI
- Navigate the HPC infrastructure more confidently

You can access it via:

* **Accounting** → **my_hpc_usage** (Available on [Baobab](#), [Yggdrasil](#), and [Bamboo](#) OpenOnDemand instances.)




Quick Overview of Your HPC Environment

The dashboard offers a graphical summary of your usage over a predefined period (from November to now). This gives you immediate insight into your consumption patterns.

Slurm Account Information

The dashboard displays:

- A list of all Slurm accounts associated with your user profile
- Clear identification of your **default Slurm account**, automatically used when no account is specified during job submission



Group and Partition Access Overview

A summary table shows:

- The HPC groups you belong to
- Associated Slurm servers
- Slurm partitions you can access

This helps you understand which parts of the HPC infrastructure are available to you.

The figure is a dark-themed dashboard titled "Private Resources" with a help icon. It features a table titled "Associated Partitions" with the following data:

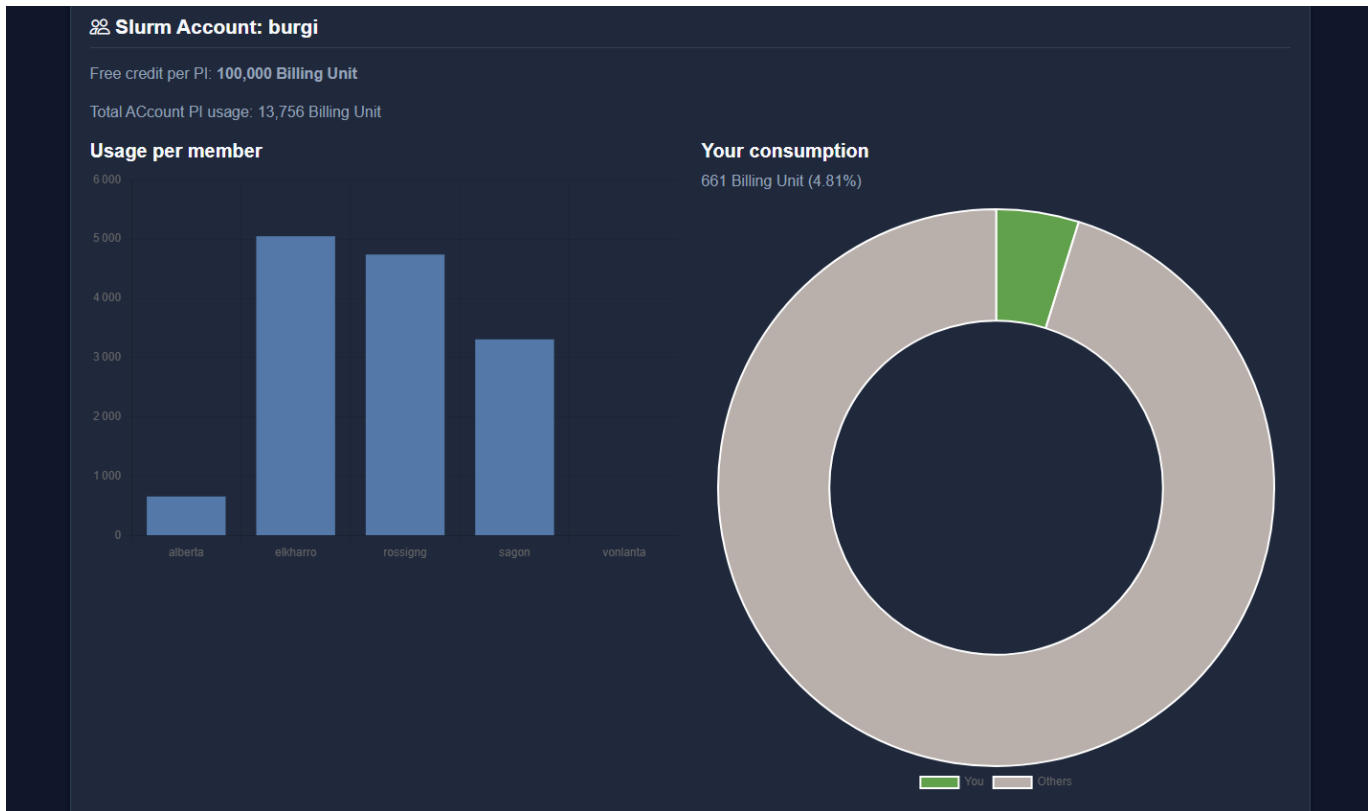
Group	Partition	NodeSet
hepia	private-hepia-cpu	cpu[265-272]
hepia	private-hepia-gpu	gpu[006,011]

Intra-PI Consumption Comparison

A comparative graph displays resource consumption (in **Billing Units**) for:

- You
- Other users belonging to the same PI

This allows you to see how computing resources are used within your research group.




SandBox

The sandbox allows you to develop your own apps in your own environment. To enable the sandbox tab, run the following command:

```
mkdir -p $HOME/ondemand/dev
```

Restart your environment **“Help > Restart web server”** (Don't be scared about this function, it restarts only YOUR environment) and the `</>develop`` tab should appear next to the ``Help`` tab.



For the most enthusiastic and adventurous users, you can develop your own app by following the procedure outlined in the

<https://osc.github.io/ood-documentation/latest/how-tos/app-development.html>

OpenOnDemand and its partners already share their apps, which you can explore here:

<https://osc.github.io/ood-documentation/latest/install-ihpc-apps.html>

Please note that apps must be configured to function properly on the Unige cluster. If the chosen app (for fake example, Dynamito) is not installed or available through a `module``, it will not work. Don't hesitate to share it with us so that we can add it to the system of interactive apps available to everyone.

To install an already developed app, navigate to the new tab ``**develop > My Sandbox app**``. This section lists all your installed apps, and you can click "**New**" to clone a git repository.



Some adjustment are required to make your own app work.

Logout

To Logout you need to close you browser and not only the current tab.

Launching applications

This documentation has been taken from <https://docs.hpc.qmul.ac.uk/ondemand/>

After choosing from the available applications, you can requested the necessary resource for the session, and then click Launch.

Number of cores

Select 1 unless you know your code can use multiple cores

Memory per core (GB)

Maximum running time

1hr job has a much shorter queueing time. 240hrs is the maximum allowed

I would like to receive an email when the session starts

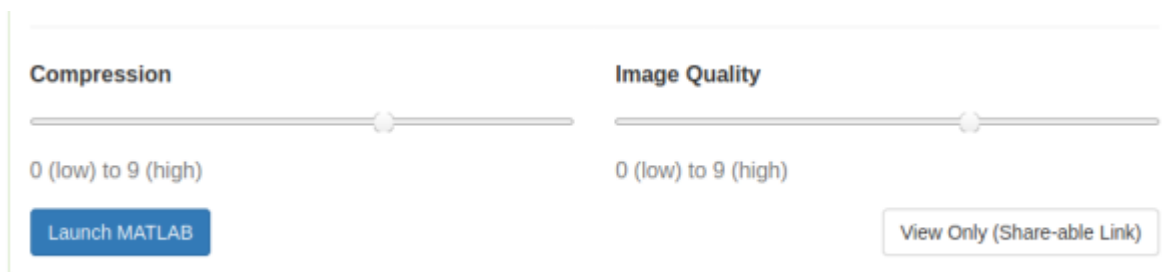
Launch

Typically, when requesting resources, you will be given the option to choose number of cores, maximum running time and memory per core. If multiple versions of the application are available, this may be offered too.


For many use-cases, including non-computational visualisation, and many Jupyter/R Studio sessions, a maximum of 1 core will be used by the application.

There is also the option to request an email to notify you when the job starts, in the case of sessions which require longer queueing time.

For graphical applications, once the session is established, you have the option to alter the image quality and compression level. For most cases, using the defaults should be fine, but for precise graphical work, or taking screenshots, you may want to increase the image quality.



The screenshot shows two sliders: 'Compression' and 'Image Quality', both ranging from 0 (low) to 9 (high). Below the sliders are two buttons: a blue 'Launch MATLAB' button and a grey 'View Only (Share-able Link)' button.

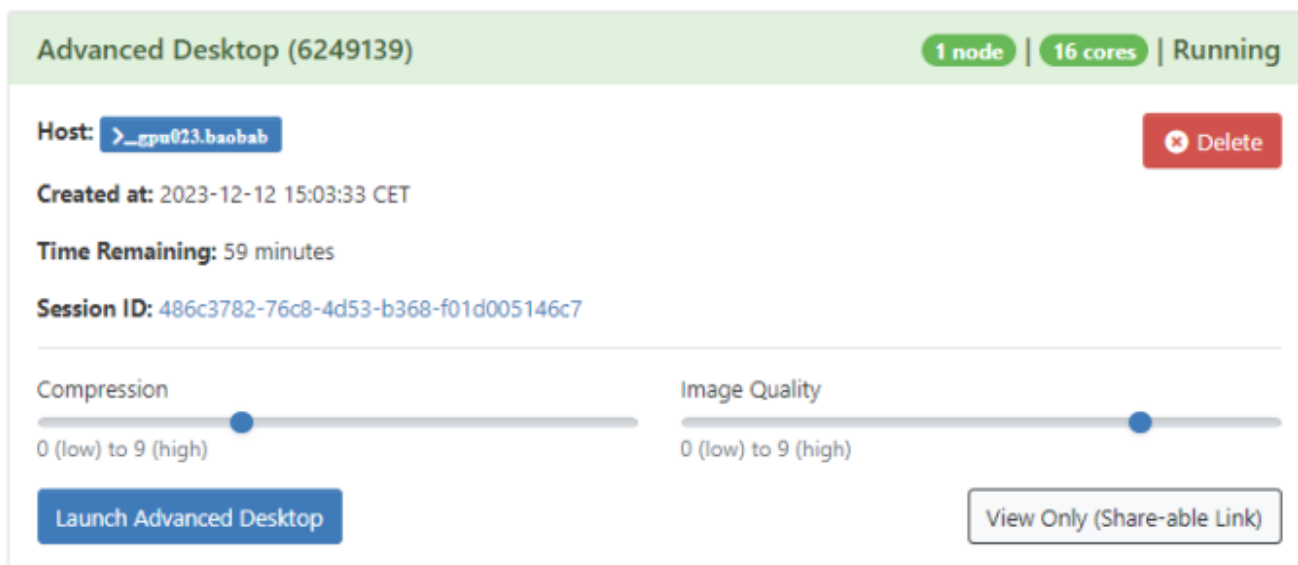
 Note also the **View only** link. OnDemand sessions are secure and confined only to you. However, you have the option to share a **read-only view** of your session for purposes of monitoring progress of a job with a collaborator or to demonstrate work to a colleague.

Session persistence

A session remains active for the time period you have requested it for. If you have finished your work, then we recommend that you close the application cleanly, followed by clicking Delete next to your session on My Interactive Sessions page. This will release resources back to the cluster.

If you are running a graphical application which uses a desktop environment, any computation will continue even after the browser tab is closed. Clicking Launch again from the My Interactive Sessions screen will reconnect to the session.

Closing a browser tab while running a web service session such as Jupyter will interrupt computation, in the same way as it would if you did the same while running Jupyter locally on your own PC. Jupyter and both offer ways to save and checkpoint your work if you want to resume later.



The screenshot shows a session titled 'Advanced Desktop (6249139)' with status '1 node | 16 cores | Running'. It includes a 'Host' field with a dropdown menu, a 'Delete' button, 'Created at' (2023-12-12 15:03:33 CET), 'Time Remaining' (59 minutes), and 'Session ID' (486c3782-76c8-4d53-b368-f01d005146c7). Below this are sliders for 'Compression' and 'Image Quality', and buttons for 'Launch Advanced Desktop' and 'View Only (Share-able Link)'.

Contribution

We have deploy a GitLab repository to centralize the Interactive apps sources and allow users to contribute, you will find more information here:

<https://hpc-community.unige.ch/t/openondemand-the-unige-openondemand-community-git-now-available/3881/1>

GitLab Repository: <https://gitlab.unige.ch/hpc/unige-openondemand>

From:

<https://doc.eresearch.unige.ch/> - **eResearch Doc**

Permanent link:

https://doc.eresearch.unige.ch/hpc/how_to_use_openondemand

Last update: **2026/02/09 15:22**

