

# To all users

- [Brief introduction to the cluster](#)
- [Access the cluster](#)
- [Builtin software](#)
- [Finding help](#)

# Brief introduction to the cluster

The cluster consists of many components to provide a good experience for various tasks for various users. Below are some highlights:

1. Variety in compute nodes
2. Parallel file system storage
3. Fast Infiniband and Ethernet network
4. SSH servers cluster
5. Web portal
6. CLI client
7. Software by modules and containers

This user guide will not cover the detailed hardware spec but instead focus on the experience. If you are interested in those technical details, please get in touch with us.

## Compute nodes

We want to provide a heterogenous cluster with a wide variety of hardware and software for users to experience different combinations. Compute nodes may have very different models, architecture, and performances. We carefully build and fine-tune the software on the cluster to ensure they fully leverage the computing power. We provide tools on our web portal to help you choose what suits you.

Besides the OneAsia resources, bringing in hardware is also welcome. Our billing system is smart enough to charge jobs by individual nodes. That means one can submit a giant job to allocate computing power owned by multiple providers. To align the terminology, we group them into three pools:

1. OneAsia
  - Hardware owned by OneAsia Network Limited
2. Bring-in shared
  - Hardware brought by external but willing to share with others
  - You could control priority with quota, priority, preemption, and fair share
3. Bring-in dedicated
  - Hardware brought by external but not willing to share

# Storage

The cluster has a parallel file system which provides fast and reliable access to your data. We are charging monthly by the maximum allowed quota. You can quickly check your quota through our web portal or CLI client. You may request a larger quota anytime by submitting a ticket to us.

One should see at least three file sets, they are:

1. User home directory
  - To store your persistent data. It is mounted at /pfss/home/\$USER and has a default quota of 10GB
  - You may access the path with an environment variable: **\$HOME**
2. User scratch directory
  - To be used for I/O during your job. It is mounted at /pfss/scratch01/\$USER and has a default quota of 100GB
  - We may purge inactive files every 30 days
  - You may access the path with an environment variable: **\$SCRATCH**
3. Group scratch directory
  - To share your file with your group mate. It is mounted at /pass/scratch02/\$GROUP and has a default quota of 1TB.
  - You may access the path with an environment variable: **\$SCRATCH\_<GROUP NAME>**

There are many ways you can access your files.

1. From the web portal file browser
2. SSH / SFTP
3. Mount to your local computer using our CLI client

When running jobs, no matter whether you are using our modules or containers. All file sets you have access to will be available.

# Networking

Traffic between compute nodes or between compute nodes and the parallel file systems is going through our Infiniband network. Both our modules or containers are compiled with the latest MPI toolchain to utilize the bandwidth fully.

# Login nodes farm

Whether you access the cluster through our web portal or your SSH client, you will be connecting to our SSH servers cluster. Our load balancer will connect you to the server with the least connections. We only grant connections by a private key. No password authentication is allowed.

You may connect through the web portal or the CLI client if you don't want to keep the private key.

You will have 8 shared CPU cores, and 8GB of memory. Free of charge for you to prepare your software, job script, and data. You will have access to all your file sets, all modules, containers, SLURM commands, and our CLI client.

Please leverage compute nodes for heavy workloads. If you need more resources on the login node, please submit a ticket and let us help.

## Web portal

Our web portal provides many features to make the journey easier. Our goal is to enable users from different backgrounds, to consume HPC resources quickly and efficiently. We also leverage the web portal internally for research and management. We will cover the details in later chapters. Below are some highlights:

1. Web Terminal
2. File browser
3. Software browser
4. Quick jobs launcher
5. Job efficiency viewer and alert
6. Quota control
7. Team management
8. Ticket system
9. Cost allocation

## CLI client

To further accelerate the workflow, we created our command line client. We will cover the details later, but below are some example use cases:

1. Connect to the login nodes farm without the private key
2. Mount a file sets to your local computer
3. Allocate ports from compute node for GUI workloads
4. Check quota and usage
5. Check cluster healthiness

## Software

The cluster currently provides free software in two ways: Lmod and Containers. Our team is working hard to provide state-of-the-art software which fine-tuned for the cluster's compute nodes. You may log in to our web portal to browse the available software.

Besides software, we also provide pre-trained models and popular data sets. We will cover the details later.

# Access the cluster

There are three ways to access the cluster: web portal, SSH, and CLI client. This article will cover how they authenticate users.

## Web portal

You should be able to log in to the web portal <https://oasishpc.hk> using the provided username and password.

If this is your first-time login, the system will ask you to set up your second factor for authentication. Please install the Google Authenticator app, and follow the instruction to set that up.

Currently, we don't have a forgot password mechanism. Please get in touch with your administrator for support if you have forgotten your login password.

The cluster is integrated with The Hong Kong Access Federation (HKAF). So if you have an HKAF account, click "login through HKAF" to log in. Extra two-factor authentication is not required.

## SSH

You may SSH directly into our cluster at **ssh.oasishpc.hk:22** through the SSH servers farm. Our load balancer will connect you to the server with the least connections. We only grant connections by a private key. You may download your private key from the web portal. Please keep your private key safe and don't share it with others.

You will have 8 shared CPU cores and 8GB of memory to prepare your software, job script, and data. If you need more resources on the login node, please submit a ticket and let us help.

If your login name is hpcuser123, you can log in by using this command line:

```
# download your SSH key from the web portal home page
# assume it is named your-key.pem

# protect your key to not be read or write by others
chmod 400 your-key.pem

# ssh with your login name and the key
```

```
ssh -i your-key.pem hpcuser123@ssh.oasishpc.hk
```

When you download a new key from the web portal, the previous one you downloaded will be deactivated.

You may also access the file system with your favorite SFTP client.

## CLI client

If you find it tedious to keep the private key, you may leverage the CLI client for keyless login. First, follow the web portal's instructions to install and configure the CLI client on your local computer. Then execute the following command to SSH without a key.

```
hc connect
```

The client supports file sets mounting like with the below command. We will cover the details in a later chapter.

```
hc filesystem-mount -t home -m /mnt/hpc-home
```

Currently, both sub-commands support Linux and Mac OSX only. Windows is not supported.

# Builtin software

The OAsis HPC cluster has some standard software built-in. We provide them via Lmod and Containers. Users may choose the way they are comfortable with.

## Lmod

All software and module files are put in the parallel file system, and accessible in any compute nodes and login nodes. To compile your code, you may load a specific MPI toolchain in a login node. And use the same MPI version to run your program on compute nodes.

We provide various versions for each software. To prevent loading incompatible sets of modules, our Lmod uses [a software hierarchy](#). For example, to load FFTW 3.3.10 using the OpenMPI 4.1.4 plus GCC 11.3 toolchain, you may execute the following statement:

```
module load GCC/11.3.0 OpenMPI/4.1.4 FFTW.MPI/3.3.10
```

Later, when you want to load, e.g., the BLAST module. You don't have to worry about the incompatible toolchain because of the help from the module system.

## Browse modules on the web portal

Log in to the web portal, click Supports, then Software, and you will see a graphical module browser. You may search for any keyword, check available versions, and copy the loading statement.

## Browse in console

To browse the catalog in log in node, use the standard Lmod spider command:

```
# see the full catalog for all available module
```

```
module spider
```

```
# search with keyword
```

```
module spider openmpi
```

```
# search for the document of a specific version
```

```
module spider openmpi/4.1.4
```



To browse modules supported by a toolchain:

```
module load GCC/11.3.0  
module avail
```

If you are interested in the details, please check out the [Lmod documentation](#).

## Containers

Another way to run software on the cluster is to use containers, which is another good way to avoid library incompatibility. Since all libraries are built into a portable container image, we don't need to load any module in advance.

Most of our examples use containers. Also, software in containers is often more up-to-date.

The cluster-provided containers are located at `/pfss/containers`. This folder is shared by all login nodes and compute nodes. In addition, we provide several GPU and MPI-ready containers for you to kick-start your workload.

Similar to Lmod, we have a polished browser on our web portal. Head to supports, software, then containers. Besides the provided containers, you should also see your containers there. The browser looks for any containers placed in the containers folder in any of your file sets. For example, `/pfss/home/loki/containers` and `/pfss/scratch02/oneasia` for the user loki of group oneasia.

# Finding help

We recommend you reach us by creating a support ticket through the web portal if you need help.

You may contact us by email or phone if you can't access the portal.

## Create a ticket in the portal

Please follow the below steps to file a ticket:

1. Log in to the web portal.
2. Locate the top menu bar, and click Tickets under support.
3. Click new ticket at the top-left corner.
4. Fill in your situation and click submit.

You will receive a notification whenever there is an update about your ticket.

## 7x24 Enquiry Support

For log in issues like forgetting the login password or losing the two-factor authentication device, you may reach us with the following information.

ECC HK

OneAsia Network Limited

Dir: (852) 3979 3961

Email: [ecc-hk@oneas1a.com](mailto:ecc-hk@oneas1a.com)

Website: [www.oneas1a.com](http://www.oneas1a.com)