National Supercomputing Centre

ASPIRE2A

General QuickStart Guide

Table of Contents

1. Introduction 2
2. User Policy 2
3. ASPIRE 2A Overview 3
    3.1 Compute Node Specification 3
4. Stakeholders Access to ASPIRE 2A 4
5. Non-stakeholders Access to ASPIRE 2A 5
6. Software Environment 7
    6.1 GNU Environment 7
    6.2 Intel Environment 7
7. Submitting Your First Job 8
    7.1 How to Submit Serial Jobs? 9
    7.2 How to Submit Array Jobs? 9
    7.3 How to Submit Parallel Jobs? 9
8. Need Help? 10
9. Frequently Asked Questions (FAQs) 11
    1. How do I submit a Serial Job? 11
    2. How do I submit Array Jobs? 11
    3. How do I submit Parallel Jobs? 11
    4. What are the queues available in ASPIRE2A? 11
    5. What is the maximum memory that can be used? 11
    6. What is my disk quota (/scratch & /home) on ASPIRE2A? 11
    7. How do I check my user and project quota on ASPIRE2A? 11
    8. Are there other options besides installing the Duo 2FA mobile app? 12
    9. Is it possible to use VPN clients such as OpenVPN or Tunnelblick? 12
    10. Who should I contact if the Duo 2FA does not work? 12

1. Introduction

This guide contains information on how to access and submit jobs on ASPIRE 2A.

2. User Policy

- All computational jobs must be submitted and run via the workload manager/scheduler.
- No computational jobs are allowed on the login nodes.
- Use project-directory assigned (**/home/project/<Project ID>/)** for sharing and storing data (dataset, database, results, input-output, etc.) and for installing controlled/licensed or unique-to-the-project applications.
- Users are welcome to install applications/software under their own `$HOME` or project-directory.
- Directory is ACL controlled. Please seek permission if you are unable to access the directory.
- Project members can access project directory based on project group permission (each project directory will be assigned to a group with the same name as the project ID). Project owners are given full access and will be able to adjust group permission within the project folder.
- Users are responsible for management and control of their own data. NSCC accepts no liability for any data loss or unauthorised data access.
- All terms of use are subject to the prevailing **NSCC Acceptable Use Policy** https://help.nscc.sg/aup/

3. ASPIRE 2A Overview

ASPIRE 2A is an AMD-Based Cray EX supercomputer with 10 PB of Lustre FS storage and Slingshot interconnect. The following diagram shows the summary infrastructure of ASPIRE 2A.

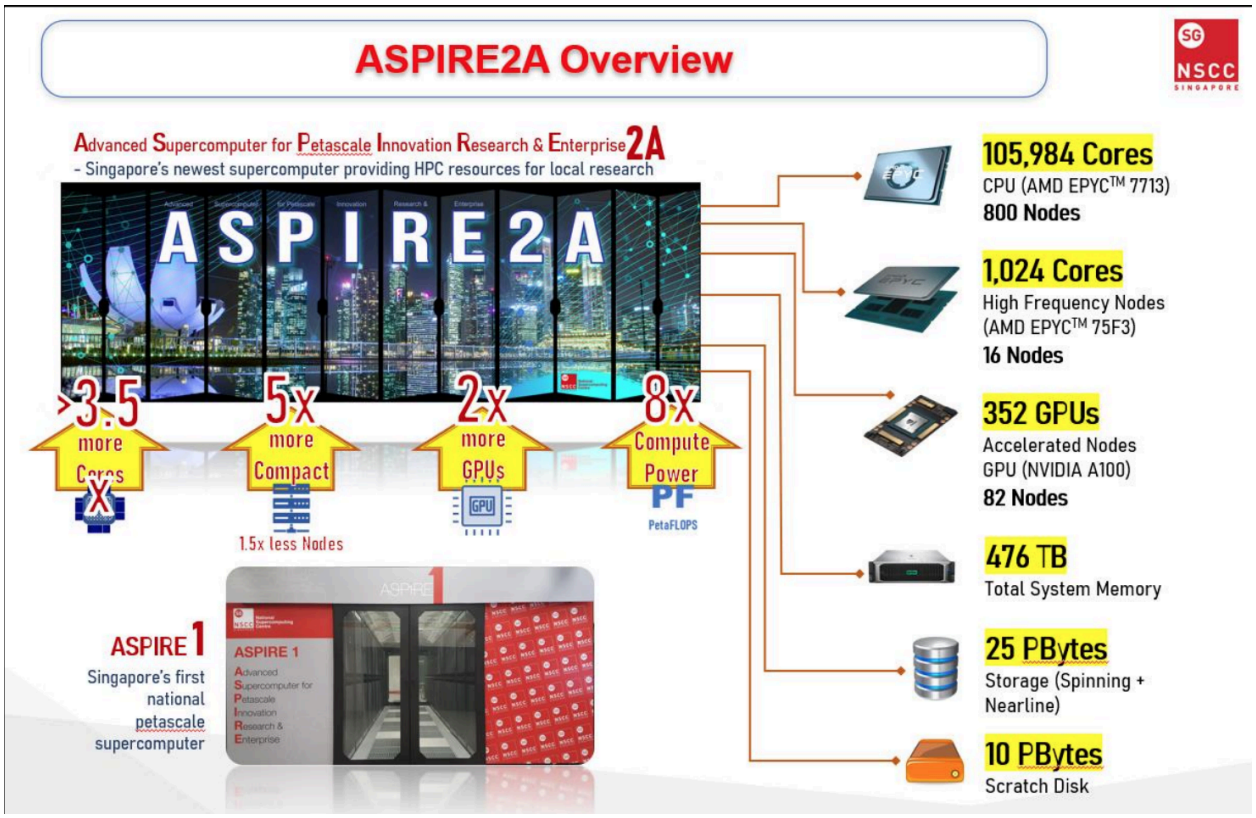

3.1 Compute Node Specification

Full list of HPC cluster node types and specification:

Server	CPU Model	Cores per socket	Socket per server	Total Physical cores per server	Available RAM (DDR4)	GPUs
Standard Compute Node (768 nodes)	Dual-CPU AMD EPYC 7713	64	2	128	512 GB	No GPU
GPU compute node (64 nodes)	Single-CPU AMD EPYC 7713	64	1	64	512 GB	**4x Nvidia A100 40GB**
GPU AI Node (12 nodes)	Single-CPU AMD EPYC 7713	64	1	64	512 GB	**4xNvidia A100 40GB** (11TB nvme)
GPU AI Node (6 nodes)	Dual-CPU AMD EPYC 7713	64	2	128	1 TB	**8xNvidia A100 40GB** (14TB nvme)

Large memory node (12 nodes)	Dual-CPU AMD EPYC 7713	64	2	128	2 TB	No GPU
Large memory node (4 nodes)	Dual-CPU AMD EPYC 7713	64	2	128	4 TB	No GPU
High frequency node (16 nodes)	Dual-CPU AMD EPYC 75F3	32	2	64	512 GB	No GPU

4. Stakeholders Access to ASPIRE 2A

All pilot users from stakeholders are recommended to connect via high speed access link to access the system. Users can access via high speed links to ASPIRE 2A as shown in Diagram 1 below (Blue colour path) using respective FQDN as listed below:-

Stakeholder	FQDN for LoginNode	If connect from outside campus network
NUS	aspire2a.nus.edu.sg	Connect to NUS VPN first
NTU (Note1)	aspire2antu.nscc.sg	Request for NTU Jumphost
A*STAR	aspire2a.a-star.edu.sg	Connect to A*STAR VPN first
SUTD	aspire2a.sutd.edu.sg	Connect to SUTD VPN first

Example, User from NUS could reach login node using ssh client via

"**`ssh <userid>@aspire2a.nus.edu.sg`**"

Note1:- NTU users will need to request NTU jumphost to access ASPIRE2A via aspire2antu.nscc.sg. Please email your jumphost access request to hpcsupport@ntu.edu.sg

How to access jump host:- Using-NTU-JumpHost-to-NSCC-ASPIRE-2A

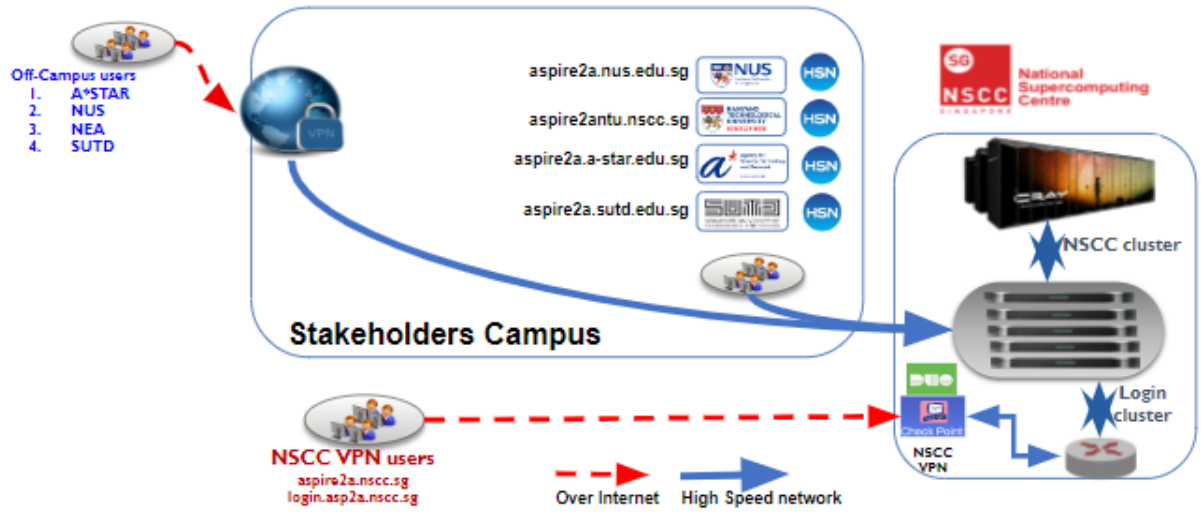

Diagram 1 : Users accessing ASPIRE 2A via High speed link and NSCC VPN

5. Non-stakeholders Access to ASPIRE 2A

All pilot users from non-stakeholders are required to connect via VPN to access the system. For more information on how to connect to the VPN, please follow the respective Windows, Mac or Unix VPN connection guides.

Non-stakeholders are referring to those user not belong to any stakeholders such as NTU, NUS, SUTD, A*STAR and NEA.

Step 1: Users are required to enrol in **DUO 2FA** before connecting to the ASPIRE 2A VPN. Users will receive an instruction email from DUO to install the DUO mobile app on their mobile phones. Once installed, users may proceed to connect to the VPN.

Please refer to the following guide for the DUO 2FA enrollment steps.

DUO 2FA Enrollment Guide

Note: User MUST complete Step 1 before proceeding to Step 2 to successfully connect to the VPN.

Step 2: Access via NSCC VPN and logging in to ASPIRE 2A
Please refer to the following guides for the setup and connection of the VPN to ASPIRE2A for different Operating Systems:

ASPIRE 2A VPN for Windows

ASPIRE 2A VPN for MAC

ASPIRE 2A VPN for Linux

Note: Please check your 2FA mobile device pop-up request for the 2FA approval while logging into the VPN using your login credentials. **VPN connection will fail without the DUO 2FA approval.**

Step 3: SSH

Login node FQDN: aspire2a.nscc.sg
Alternate FQDN: login.asp2a.nscc.sg

```
ssh -Y <userid>@aspire2a.nscc.sg
OR
ssh -Y <userid>@login.asp2a.nscc.sg
```

Should you have any query, please contact help@nscc.sg for further help.

Upon successful login you will be presented with a screen similar to Figure 1 displaying the current status of the system.

Figure 1: Successful login to ASPIRE2A login node

6. Software Environment

Figure 2 represents the software environment of ASPIRE 2A.

ASPIRE 2A uses the Red Hat Enterprise Linux 8 (RHEL) as its operating system (OS). The Cray development environment will be loaded by default.

Type `module list` to see currently loaded modules.

Type `module av` to see the available modules on ASPIRE2A.

6.1 GNU Environment

To load the GNU environment and GNU compilers please type:

```
module swap PrgEnv-cray PrgEnv-gnu/8.3.3
```

6.2 Intel Environment

To load the Intel environment and Intel compilers please type:

```
module swap PrgEnv-cray PrgEnv-intel/8.3.3
```

Refer to ASPIRE2A Software List: ASPIRE2A Software-List

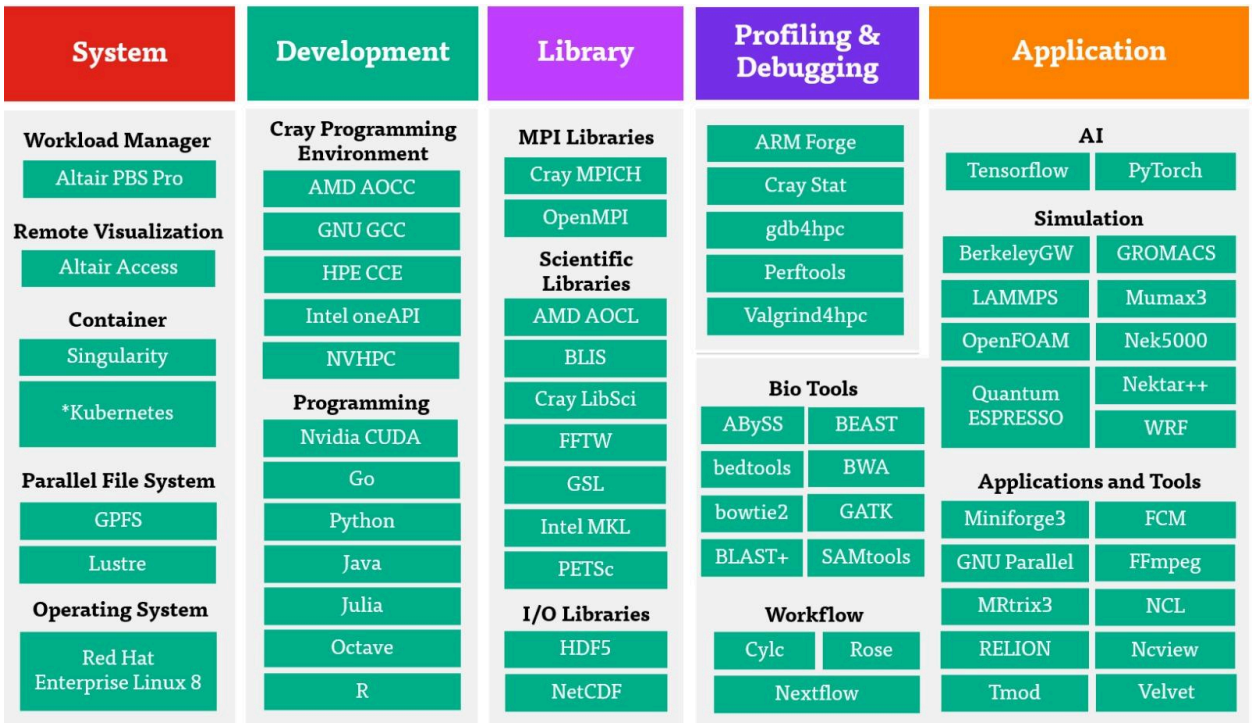

Figure 2: ASPIRE2A Software Environment

7. Submitting Your First Job

On the login nodes, you will be able to submit a PBS job script to the scheduler to run your HPC workload. Below is an example PBS job script for your reference.

```
#!/bin/bash
#PBS -N UM2KM
#PBS -l select=8:ncpus=128:mem=440G:mpiprocs=128:ompthreads=1
#PBS -l walltime=3:00:00
#PBS -j oe
#PBS -o out-run.txt
#PBS -P <Project ID>
#PBS -q normal

module load miniforge3
```

#PBS -P <Project ID>

Users are advised to always specify the **Project ID** accessible to users while submitting a job.

#PBS -q normal

Currently only **normal** and **ai** queues are accessible to users while submitting jobs, while all other queues are execution queues, which are not directly accessible to users. However, user jobs will be routed to respective queues for execution automatically based on job configuration.

- **Normal** queue is for **CPU** and **GPU node jobs.**
- **AI** queues are for **AI workloads** where each node has been configured with high speed NVme local storage for AI/ML use cases.
  - The NVme local storage is mounted as `/raid`.
- The AI queue is not accessible to users by default. they will need to request AI queue access during project submission.

NOTE: GPU Job will have fixed CPU and Memory ratio enforced i.e. for every 1 GPU requested, 16 CPU and 110 GB of memory will be allocated.

The PBS directives request for 8 nodes each with 256 Hyper-Threading threads (128 physical cores) and 440 GB memory. The `ncpus=128` specified corresponds to 256 Hyper-Threading threads, which is equivalent to 128 physical cores. Each node has 512 GB of main memory, but the `mem=440G` is used considering that some of the memory will be used for the system OS.

Note that **if the job queue is not specified** in the job script the **normal queue** is automatically selected.

7.1 How to Submit Serial Jobs?

To submit a serial job, simply request a single chunk and single ncpus resource.

Example: `qsub -l select=1:ncpus=1:mem=2GB`

Wrapping multiple serial commands in a single job script.
Example:
A single PBS job script can contain script1,script2 as two serial jobs / commands. The drawback with this example is that any output from job1 or job2 will get mixed up in the batch jobs output file.

The script would look like:

```
#!/bin/bash
#PBS -l walltime=1:00:00
#PBS -l select=1:ncpus=2
cd $PBS_O_WORKDIR
# Use '&' to start the first job in the background
./script1 &
./script2
# Use 'wait' as a barrier to collect both executables when they are
done. If not, the PBS job will finish when the script2 program
finishes and kill script1 if it is still running.
wait
```

7.2 How to Submit Array Jobs?

Job arrays are a great way to organise the execution of multiple short jobs with similar properties or if they are using similar data with different algorithms or if they are using serial input file numbering system eg., file01, file02, file03

Example: Submit 10 jobs with consecutive index numbers.

```
#!/bin/sh Examples Using a Job Script
#PBS -N Simn1010Jobs
#PBS -J 1-10
echo "Main script: index " $PBS_ARRAY_INDEX
/opt/AppA -input /home/user01/runcase1/scriptlet_$PBS_ARRAY_INDEX
```

7.3 How to Submit Parallel Jobs?

For an MPI job, select= and ncpus= can be one or more, with np= >1.

```
#!/bin/bash
#PBS -N testrun01
#PBS -l select=6:ncpus=32:mem=128gb:mpiprocs=32:ompthreads=1
#PBS -o /home/user01/my_o_files
#PBS -e /home/user01/my_e_files
#PBS -m be
mpirun -np $MPI_NPROCS -machinefile $PBS_NODEFILE ./my_scripts
```

8. Need Help?

If you encounter any issues, please contact us at help@nscc.sg during this transition period and until further notice. For enquiries regarding **ASPIRE 2A**, please indicate the following in your email subject heading - **“ASPIRE 2A: [issue]”**

9. Frequently Asked Questions (FAQs)

1. **How do I submit a Serial Job?**

   To submit a serial job, ensure you request one CPU resource only, (i.e., `-l select=1:ncpus=1`). Refer to **Section 6.1** for example.

2. **How do I submit Array Jobs?**

   Refer to **Section 6.2** for example.

3. **How do I submit Parallel Jobs?**

   To submit a parallel job, you can request 1 or more ncpus resources per chunk. Refer to **Section 6.3** for example.

4. **What are the queues available in ASPIRE2A?**

   The queues available for job submission are:
   - normal
   - ai

5. **What is the maximum memory that can be used?**

   There are 3 types of nodes in the cluster based on their memory, which are 502GB, 2TB and 4TB. A job can request multiple chunks of memory with a maximum of 4TB(4000GB) per chunk.
   (For example: `-l select=2:ncpus=128:mem=4000g` is a job that is requesting 8TB(8000GB) memory in total and will run on 2 nodes). However, **"ai"** queues will have fixed CPU and Memory ratio enforced i.e. for every 1 GPU requested, 16 CPU and 110GB of memory will be allocated.

6. **What is my disk quota (`/scratch` & `/home`) on ASPIRE2A?**

   Scratch and Home quota remains the same as below.
   `/home` Home Directory - 50 GB (Total inodes - 400000000)
   `/scratch` Scratch Directory - 100TB (Total inodes - 200000000)

7. **How do I check my user and project quota on ASPIRE2A?**

   Users may use **`myusage`** to check resources usage, **`myprojects`** to check project balance and **`myquota`** to check storage usage.

8. **Are there other options besides installing the Duo 2FA mobile app?**

   Only the mobile app is supported currently. NSCC will continue to explore other alternatives at a later stage.

9. **Is it possible to use VPN clients such as OpenVPN or Tunnelblick?**

   Only Checkpoint VPN clients are supported.

10. **Who should I contact if the Duo 2FA does not work?**

    Please contact us at help@nscc.sg during this transition period and until further notice. For enquiries regarding **ASPIRE 2A**, please indicate the following in your email subject heading - **“ASPIRE 2A: [issue]”**