National Supercomputing Centre (NSCC) Singapore
1 Fusionopolis Way, Connexis South,
#17-01, Singapore 138632
E: contact@nscc.sg
Reg no: 199702110H

ASPIRE 2B General QuickStart Guide

National Supercomputing Centre (NSCC) Singapore
1 Fusionopolis Way, Connexis South,
#17-01, Singapore 138632
E: contact@nscc.sg
Reg no: 199702110H

Table of Contents

National Supercomputing Centre (NSCC) Singapore
1 Fusionopolis Way, Connexis South,
#17-01, Singapore 138632
E: contact@nscc.sg
Reg no: 199702110H

1. About This Guide

This Quick Start Guide provides the essential steps for ASPIRE 2B users to access the system, use the available storage and software environment, submit and monitor jobs through PBS Professional, and obtain support.

2. Before You Start

Before accessing ASPIRE 2B, ensure that the following prerequisites have been completed:

1. You have an active NSCC Account.
2. You have been added to an approved ASPIRE 2B project.
3. You have completed Duo enrolment.
4. You have completed Zoho enrolment.
5. You have installed the supported NSCC VPN client. (Refer to the guide here)
6. You have an SSH client available on your workstation.
7. You know your assigned ASPIRE 2B Project ID.

Access to ASPIRE 2B is provided through the NSCC VPN. After connecting to the VPN, users can access the ASPIRE 2B login nodes using SSH.

If you are unable to access ASPIRE 2B, your project directory, or other assigned resources, contact the ASPIRE 2B Service Desk. (servicedesk@nscc.sg)

3. Important Usage Rules

To ensure fair, secure and stable use of ASPIRE 2B, users must observe the following rules:

- **Use the scheduler for computational workloads** - All computational jobs must be submitted through the ASPIRE 2B workload manager/scheduler.
- **Do not run compute-intensive jobs on login nodes** - Login nodes are intended for activities such as preparing job scripts, compiling software, transferring files and performing light interactive work.
- **Use the appropriate storage location** - Use your assigned home, project and scratch storage according to their intended purpose. Project-specific datasets, results and shared applications should be stored in the assigned project directory where applicable.
- **Use only resources assigned to your project**. - Access to project directories and computing resources is controlled through approved project membership and system permissions.
- **Protect your account and credentials** - User accounts are personal and must not be shared. Users are responsible for protecting their passwords, SSH credentials and MFA access.

- **Manage project and research data responsibly**- Users are responsible for managing their data in accordance with applicable NSCC data management and retention requirements.
- **Install software only in permitted locations** - Where user-managed software installation is permitted, software should be installed under the user's home directory or the appropriate project directory.
- **Do not bypass system controls** - Users must not attempt to circumvent scheduler policies, access controls, quotas, security controls or other platform restrictions.
- **Report problems through the approved support channel** – ASPIRE 2B system issues, access problems and service requests should be submitted through the ASPIRE 2B Service Desk.

All use of ASPIRE 2B remains subject to the prevailing **NSCC Acceptable Use Policy** and **NSCC Data Management and Retention Policy**

4. Identity, MFA and Service Desk Portal Access

Before accessing the NSCC Service Desk Portal, users must complete the required Duo and Zoho enrolment. These services are used to support authentication and access to NSCC services including VPN and Service Desk Portal access.

Please refer to the dedicated User Onboarding Guide for the enrolment steps.

5. Accessing NSCC VPN and ASPIRE 2B System

Step 1: Connect to VPN using Checkpoint Client

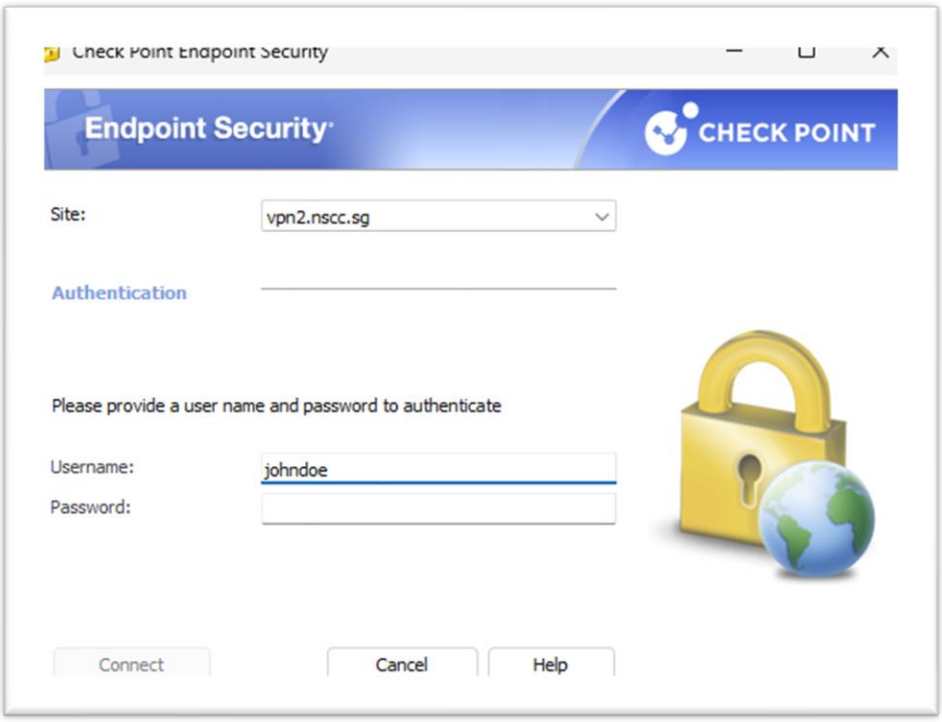

Step 2: Approve the Cisco Duo Push notification on your registered mobile device.

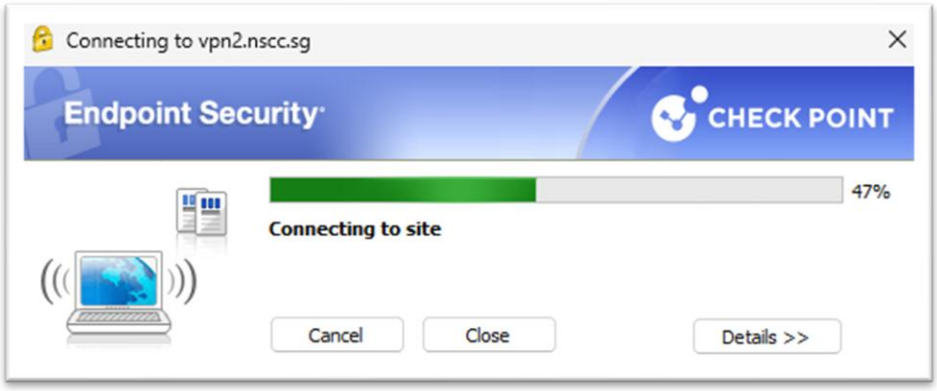

Step 3: SSH to ***aspire2b.nscc.sg*** with your preferred terminal using your HPC login and password:

FQDN	IP address
aspire2b.nscc.sg	10.105.72.23

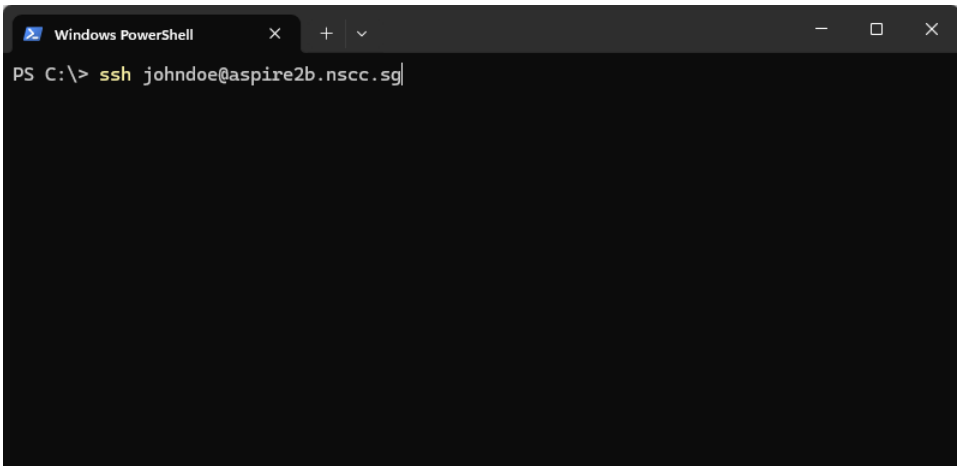

Step 4: You have successfully connected to ASPIRE 2B login node when you see this welcome message.

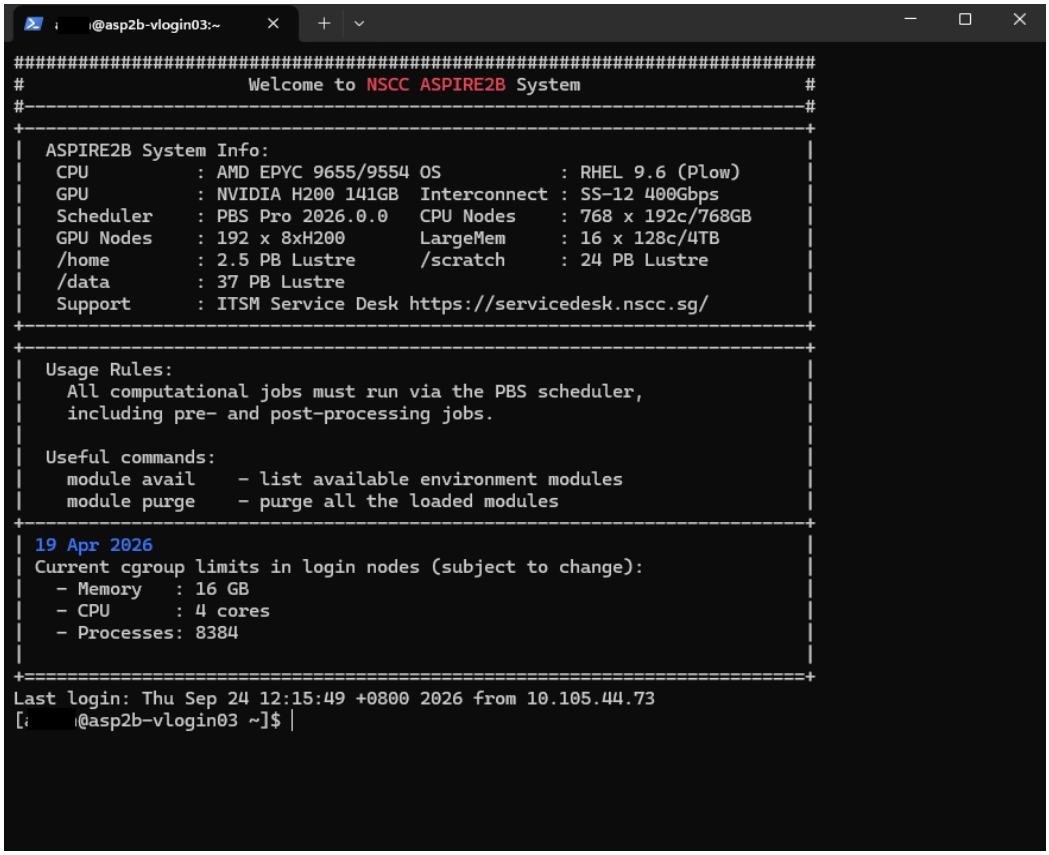

6. ASPIRE 2B System Overview

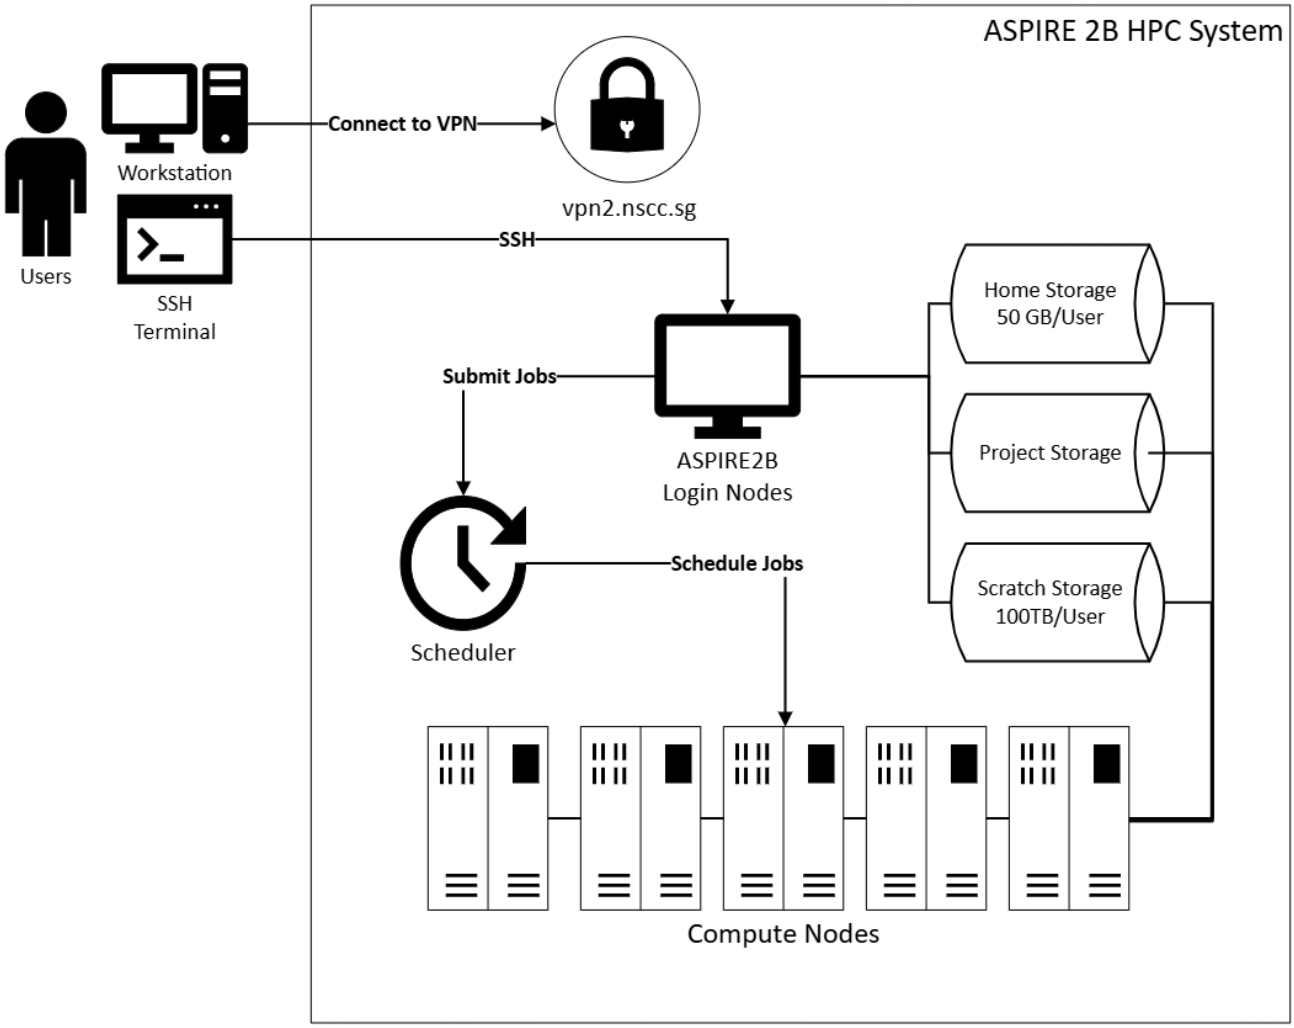

ASPIRE 2B is NSCC's high-performance computing platform providing CPU, GPU and large-memory computing resources for research workloads:

Resource	Number of Nodes	Specification per node
CPU Compute	768	HPE Cray EX4252 Dual Socket CPU per Compute Node **Processor:** 2x AMD EPYC 9655 96-Core **Installed Memory:** 768GB DDR5

GPU Compute	192	ProLiant Compute XD685 Dual Socket 8-Way GPU Server **Processor:** 2x AMD EPYC 9655 96-Core **Installed Memory**: 2.2 TB DDR5 **GPU:** 8 x Nvidia H200 **GPU Memory:** 141 GB HBM3e
Large Memory Compute	16	ProLiant DL385 Gen11 2U dual-socket rack server **Processor:** 2x AMD EPYC 9554 64-Core **Installed Memory:** 4 TB DDR5
Scheduler	-	PBS Professional
Parallel Filesystem	-	HPE ClusterStor - Lustre Filesystem • Home – 2.5 PB • Project – 34 PB • Scratch – 22 PB
Operating System	-	Red Hat Enterprise Linux 9
Interconnect	-	HPE Slingshot 400Gbps

7. Understand Your Storage

ASPIRE 2B provides different storage areas for different purposes. Users should select the appropriate location based on how the data will be used.

Storage Area	Typical Purpose	Default Quota
Home Directory	Personal files, scripts, configuration files and small user-specific data	50GB per user
Project Directory	Shared project datasets, results, input/output files and project-specific applications.	Based on approved allocation.
Scratch Directory	Temporary, high-performance storage for active computational workloads and intermediate job data.	100 TB per user

Project directories are access-controlled and are available only to authorised members of the corresponding project.

Scratch storage is intended for temporary working data and should not be used as the sole location for important or long-term data.

Users are responsible for managing their own data and should ensure that important data is stored in the appropriate location in accordance with the prevailing NSCC Data Management and Retention Policy.

To check your storage usage and quota

```
$ myquota
```

```
[      @asp2b-vlogin01 ~]$ myquota
+-------------------------------------------------+-------------------------+
|                                 Block Limit     |       Inode Limit       |
| Filesystem Type                 Usage     Quota |       Usage       Quota |
+-------------------------------------------------+-------------------------+
| home       LUSTRE                100K    50.00G |          23           0 |
| scratch    LUSTRE               4.00K      100T |           1      200000 |
+=================================================+=========================+
Quota ID 12611 - covers your files under:
home: /home/users/industry/ptcsys/alden
scratch: /scratch/users/industry/ptcsys/alden

+-------------------------------------------------+-------------------------+
|                                 Block Limit     |       Inode Limit       |
| Project    Type                 Usage     Quota |       Usage       Quota |
+-------------------------------------------------+-------------------------+
| 90000003   LUSTRE               6.71T      100T |      285117           0 |
+=================================================+=========================+
Your /data projects - covers files under:
90000003: /data/projects/90000003
[      @asp2b-vlogin01 ~]$
```

8. Using the Software Environment

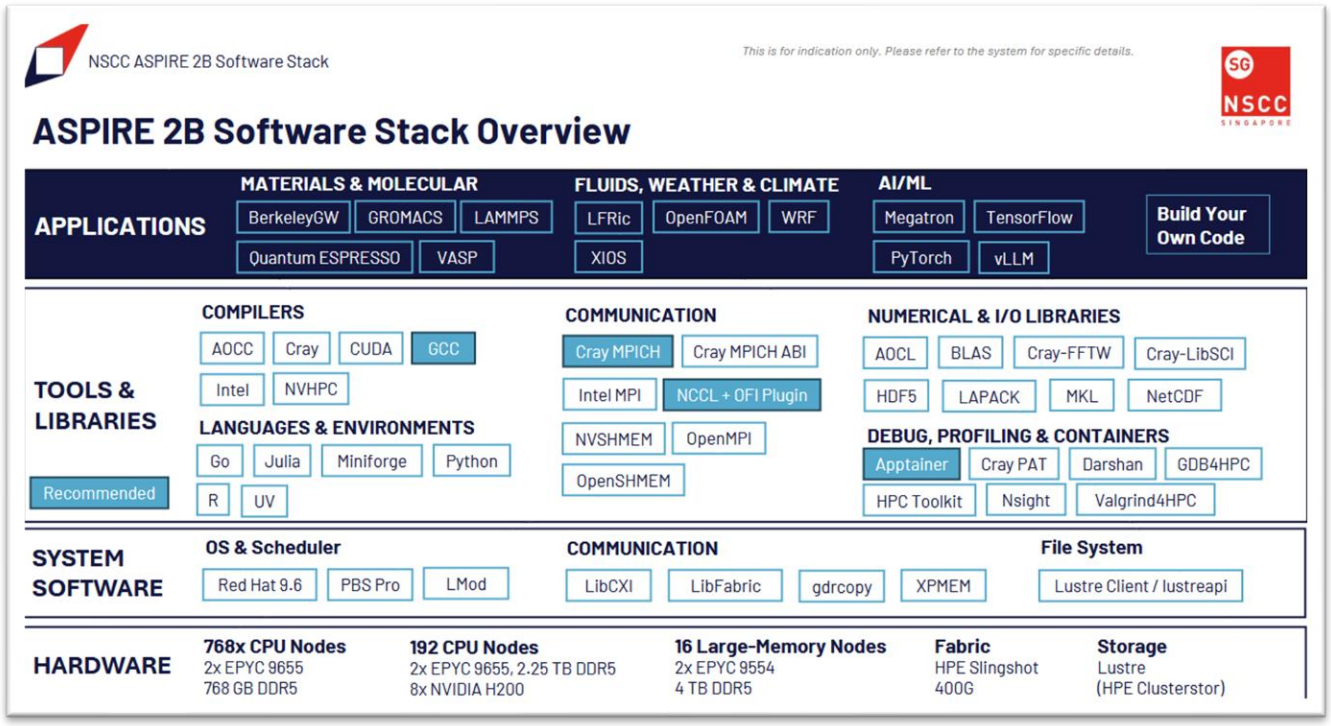

ASPIRE 2B provides centrally managed software through environment modules. Users can use the module system to view, load and manage the software available for their jobs.

#1 To displays the software modules currently loaded in your environment.

```
$ module list
```

```
[      @asp2b-vlogin02 ~]$ module list

Currently Loaded Modules:
  1) craype-x86-turin     4) perftools-base/25.09.0          7) craype/2.7.35       10) cray-pals/1.7.1
  2) libfabric/2.3.1      5) xpmem/1.0.1-1.4_gf67ddfb318d4   8) cray-mpich/9.0.1    11) PrgEnv-gnu/8.6.0
  3) craype-network-ofi   6) gcc-native/14                   9) cray-libpals/1.7.1

[      @asp2b-vlogin02 ~]$
```

#2 To displays the software modules available on the system.

```
$ module avail
```

```
[      @asp2b-vlogin02 ~]$ module avail

------------------------------------------ /opt/cray/pe/lmod/modulefiles/mpi/gnu/10.0/ofi/1.0/cray-mpich/8.0 ------------------------------------------
   cray-hdf5-parallel/1.14.3.7    cray-mpixlate/1.0.9    cray-parallel-netcdf/1.12.3.19

---------------------------------------- /opt/cray/pe/lmod/modulefiles/comnet/gnu/10.0/ofi/1.0 ----------------------------------------
   cray-mpich-abi/8.1.33    cray-mpich-abi/9.0.1 (D)    cray-mpich/8.1.33    cray-mpich/9.0.1 (L,D)

-------------------------------------------------- /opt/cray/pe/lmod/modulefiles/compiler/gnu/10.0 --------------------------------------------------
   cray-hdf5/1.14.3.7    cray-libsci/25.09.0

---------------------------------------------------- /opt/cray/pe/lmod/modulefiles/perftools/25.09.0 ----------------------------------------------------
   perftools-lite-events    perftools-lite-gpu    perftools-lite-hbm    perftools-lite-loops    perftools-lite    perftools-preload    perftools

------------------------------------------------------ /opt/cray/pe/lmod/modulefiles/net/ofi/1.0 -------------------------------------------------------
   cray-openshmemx/11.8.0

------------------------------------------------------ /opt/cray/pe/lmod/modulefiles/cpu/x86-turin/1.0 -------------------------------------------------
   cray-fftw/3.3.10.11

---------------------------------------------------------- /opt/cray/pe/lmod/modulefiles/core ----------------------------------------------------------
   PrgEnv-cray/8.6.0          cce/20.0.0         cray-dsmml/0.3.1          cray-stat/4.12.6          cuda/13.0      (D)    papi/7.2.0.2
   PrgEnv-gnu/8.6.0    (L)    cpe/25.09          cray-dyninst/12.3.6       cray-zmqnet/1.3.2         gcc-native/14  (L)    perftools-base/25.09.0 (L)
   PrgEnv-intel/8.6.0         cray-R/4.4.0       cray-libsci_acc/25.09.2   craype/2.7.35      (L)    gdb4hpc/4.16.5        sanitizers4hpc/1.1.6
   PrgEnv-nvidia/8.6.0        cray-ccdb/5.0.7    cray-mrnet/5.1.6          craypkg-gen/1.3.36        intel/2025.3          valgrind4hpc/2.13.6
   atp/3.15.7                 cray-cti/2.20.0    cray-pmi/6.1.16           cuda/12.9                 nvidia/25.11

------------------------------------------------- /opt/cray/pe/lmod/modulefiles/craype-targets/default --------------------------------------------------
   craype-accel-host    craype-accel-nvidia90    craype-network-ofi (L)    craype-x86-genoa    craype-x86-turin (L)

------------------------------------------------------------ /opt/cray/pals/modulefiles ------------------------------------------------------------
   cray-libpals/1.7.1 (L)    cray-pals/1.7.1 (L)

-------------------------------------------------------------- /data/apps/modulefiles --------------------------------------------------------------
   apptainer/1.5.1                      intel/mkl/2026.0              openmpi/5.0.10-gcc15.3-cpu
   aws-ofi-nccl/1.19.2                  intel/mpi/2021.18             openmpi/5.0.10-gcc15.3-cuda13.0          (D)
   berkeleygw/4.0-cpu                   miniforge3/26.3.2             pelec/1.0.0-cpu
   berkeleygw/4.0-gpu            (D)    mkl/2026.0                    pelec/1.0.0-gpu                          (D)
   cmake/4.3.4                          nccl/2.30.7                   python/3.11
   darshan/3.5.0-cpe25-gnu-cmpi         nekrs/26.0                    python/3.14t
   elpa/2026.02.001-gcc14-libsci        nvshmem/3.4.5-gcc11.5         python/3.14                              (D)
   elpa/2026.02.001-nv25.11-mkl  (D)    nvshmem/3.7.2-cuda12    (D)   qe/7.0-cpu
   gcc/15.3.0-znver5                    nvtop/3.3.2                   qe/7.5-cpu
   go/1.26.4                            openblas/0.3.33-gcc14.2       qe/7.5-gpu                               (D)
   gromacs/2026.2-cpu                   openblas/0.3.33-icc25.3 (D)   scalapack/2.2.2-gcc14.2-cmpi9.0-openblas
   gromacs/2026.2-gpu            (D)    openfoam/9                    scalapack/2.2.2-icc25.3-cmpi9.0-openblas (D)
   hpctoolkit/2026.0.1-cpe25-gnu        openfoam/13                   svn/1.14.5-cpe25-gnu
   htop/3.5.1                           openfoam/2312                 tmux/3.6a
   intel/compiler/2026.0                openfoam/2512           (D)   uv/0.11.21

------------------------------------------------------------- /opt/cray/modulefiles --------------------------------------------------------------
   libfabric/2.3.1 (L)    nvhpc-nompi/25.11    xpmem/1.0.1-1.4_gf67ddfb318d4 (L)
```

#3 To loads a software module and configures your environment so that the software can be used.

```
$ module load <name>/<version>
```

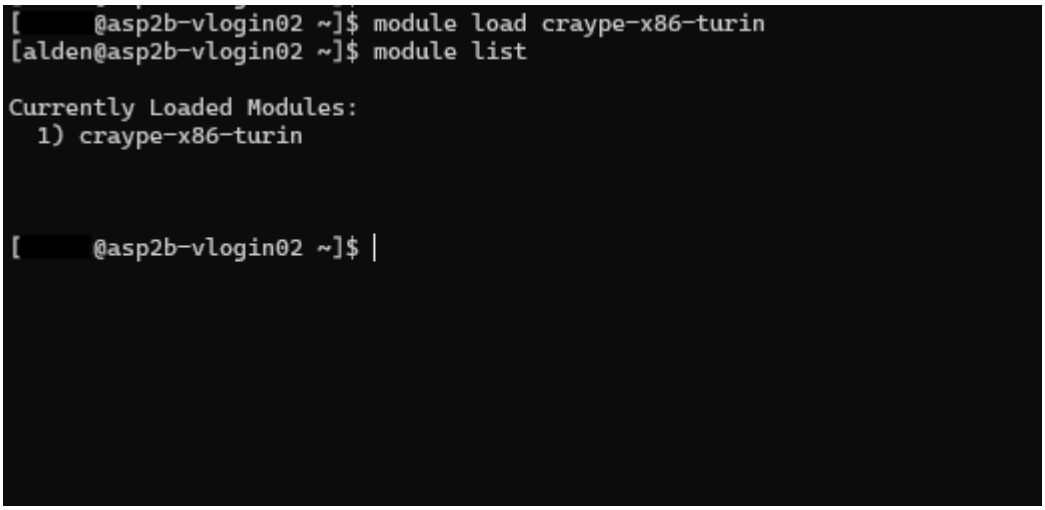

#4 Clear all loaded modules

```
$ module purge
```

Removes all modules currently loaded in your environment. This is useful when you want to start with a clean module environment before loading the software required for your workload.

The module purge does not permanently change your default environment. When you log in again, the system's default module environment may be loaded automatically.

Users should load only the software modules required for their workload and, where applicable, specify the required modules in their job submission script so that the software environment can be reproduced consistently when the job runs.

Where centrally provided software does not meet the project requirement, users may install permitted software within their own home directory or the relevant project directory, subject to applicable NSCC policies and licensing requirements.

9. Submit Your First Job

All computational workloads on ASPIRE 2B must be submitted through PBS Professional. Users should not run computational workloads directly on the login nodes.

A PBS job is normally submitted using a job script that defines the required resources and the commands to be executed.

Syntax:

```
$ qsub -q normal -P <project_id> -l select=<chunks>:ncpus=<cores> -l walltime=<HH:MM:SS>
job.sh
```

Examples:

```
$ qsub -q normal -P 11501234 -l select=1:ncpus=48 -l walltime=04:00:00 job.sh
$ qsub -q normal -P 11501234 -l select=2:ncpus=192 -l walltime=48:00:00 job.sh
$ qsub -q normal -P 11501234 -l select=1:ngpus=8 -l walltime=12:00:00 gpu_job.sh
$ qsub -q normal -P 11501234 -l select=4:ngpus=8 -l walltime=24:00:00 multi_gpu.sh
```

In practice, all these parameters can be defined inside the job script using #PBS directives, so you do not need to type them on the command line each time:

```
#!/bin/bash
#PBS -q normal
#PBS -P 11200001
#PBS -l select=1:ncpus=48
#PBS -l walltime=04:00:00

# your commands here
```

Then simply submit with:

```
$ qsub job.sh
```

For more info on the queue structure, you may refer this guide here.

The guide is accessible after connecting to the NSCC VPN.

10. Jobs Monitoring

10.1 To show what jobs have been submitted.

```
$ qstat -answ
```

```
[root@x8006c1s0b0n1 ~]# qstat -ans

pbs201:
                                                                  Req'd  Req'd   Elap
Job ID               Username Queue    Jobname    SessID NDS TSK Memory Time  S Time
-------------------- -------- -------- ---------- ------ --- --- ------ ----- - -----
120980.pbs201        josephpo pgpu     qwen_sweep 38652*   4 768  8812g 08:00 R 04:42
   asp2b-gpu189/0*192+asp2b-gpu190/0*192+asp2b-gpu191/0*192+asp2b-gpu192/0*192
   Job run at Tue Aug 25 at 08:55 on (asp2b-gpu189:ncpus=192:ngpus=8:mem=...
120988.pbs201        zhangn3  pgpu     slm_w10    25467*  16 30*   35tb 12:00 R 02:45
   asp2b-gpu004/0*192+asp2b-gpu005/0*192+asp2b-gpu006/0*192+asp2b-gpu007/0*192+
   asp2b-gpu008/0*192+asp2b-gpu009/0*192+asp2b-gpu010/0*192+asp2b-gpu011/0*192+
```

10.2 To show full details of a job if it is currently in queue or running.

```
$ qstat -f <jobID>
```

```
[root@x8006c1s0b0n1 ~]# qstat -f 120980.pbs201
Job Id: 120980.pbs201
    Job_Name = qwen_sweep
    Job_Owner = josephpo@asp2b-vlogin12.head.cm.asp2b.nscc.sg
    resources_used.cpupercent = 1886
    resources_used.cput = 06:41:22
    resources_used.gpu_usage = asp2b-gpu189(cuda.4.boardLimitViolationTime=0.00
        0s, cuda.4.eccDoubleBit=0, cuda.4.eccSingleBit=0,
         cuda.4.energyConsumed=30644J, cuda.4.lowUtilizationTime=0.000s,
         cuda.4.maxGpuMemoryUsed=406552.000MB, cuda.4.memoryClock_min=3201MHz,
         cuda.4.memoryClock_max=3201MHz, cuda.4.memoryClock_avg=3201MHz,
         cuda.4.memoryUtilization_min=0%, cuda.4.memoryUtilization_max=79%,
         cuda.4.memoryUtilization_avg=8%, cuda.4.numOtherComputePids=0,
```

10.3 To show all the nodes in the cluster, available/total resources and jobs running on the nodes.

```
$ pbsnodes -aSj
```

```
[root@asp2b-vlogin09 ~]# pbsnodes -aSj
                                                    mem       ncpus   nmics   ngpus
vnode           state           njobs   run   susp     f/t        f/t     f/t     f/t   jobs
--------------- --------------- ------ ----- ------ ------------ ------- ------- ------- --------
x8000c0s0b0n0   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s0b0n1   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s0b0n2   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s0b0n3   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s1b0n0   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s1b0n1   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s1b0n2   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s1b0n3   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s2b0n0   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
x8000c0s2b0n1   free                 0     0      0  692gb/692gb 192/192     0/0     0/0 --
```

11. Data Transfer from Other NSCC HPC System

Data transfer between ASPIRE 2B and other NSCC supported HPC systems should be performed using the approved NSCC data-transfer method.

Refer to the ASPIRE 2B Data Transfer Guide for the supported transfer endpoints, commands and usage instructions.

The guide is accessible after connecting to the NSCC VPN.

12. Get Help

Support channels for ASPIRE2A/ASPIRE2A+ and ASPIRE 2B are different. Users should submit incidents and service requests through the support channel associated with the affected system.

System	Support Channel	Usage
ASPIRE 2A / ASPIRE 2A+	ASPIRE 2A and 2A+ Service Desk: - Web: https://keris.service-now.com/fapl_csm - Email: help@nscc.sg	Support for ASPIRE 2A and ASPIRE 2A+ services
ASPIRE 2B	ASPIRE 2B Service Desk: - Web: https://servicedesk.nscc.sg - Email: servicedesk@nscc.sg	Support for ASPIRE 2B incidents, access issues and service requests

For ASPIRE 2B System, users should use the ASPIRE 2B Service Desk as the primary support channel for:

- system access or login issues;
- project or account-related requests;
- job submission or scheduler issues;
- software and application support;
- storage and filesystem issues;
- service degradation or system incidents; and
- other ASPIRE 2B enquiries.

When reporting an issue, provide sufficient information to assist troubleshooting, such as:

- user ID;
- project ID, where applicable;
- job ID, where applicable;
- approximate date and time of the issue;
- error message or observed behaviour; and
- relevant screenshots or logs where available.

Important: For ASPIRE 2B support, always use the ASPIRE 2B Service Desk. ASPIRE 2A/ASPIRE 2A+ and ASPIRE 2B are supported through separate service desks.

National Supercomputing Centre (NSCC) Singapore
1 Fusionopolis Way, Connexis South,
#17-01, Singapore 138632
E: contact@nscc.sg
Reg no: 199702110H

13. Frequently Asked Questions (FAQs)

a. Which Service Desk should I use for ASPIRE 2B?
Ans: Use the ASPIRE 2B Service Desk at https://servicedesk.nscc.sg. Do not use the ASPIRE2A/ASPIRE2A+ Service Desk for ASPIRE 2B requests.

b. Why can I access the VPN but cannot log in to ASPIRE 2B?
Ans: Confirm that your ASPIRE 2B account is active and that you have been added to an approved ASPIRE 2B project. If the issue persists, contact the ASPIRE 2B Service Desk.

c. What is the ASPIRE 2B login address?
Ans: Connect using SSH to **`aspire2b.nscc.sg`** after establishing the NSCC VPN connection.

d. Can I run applications directly on the login node?
Ans: No. Compute-intensive workloads must be submitted through PBS Professional; login nodes are intended for job preparation, compilation, file transfer and light interactive work.

e. What is the maximum memory that can be used?
Ans:

- CPU node: Up to 192 physical CPU cores and approximately 691 GB usable memory per node.
- GPU node: Up to 8 NVIDIA H200 GPUs and approximately 2200 GB usable system memory per node.
- Large Memory Node: Up to 128 physical CPU cores and approximately 4000 GB usable memory per node.

f. What is my disk quota (/home & /scratch) on ASPIRE 2B?
Ans:

- Home Directory (/home) - 50 GB
- Scratch Directory (/scratch) - 100TB

g. How do I check my home and scratch storage quota on ASPIRE 2B?
Ans: Run `myquota` command in the login node

h. Are there other options besides installing the Duo 2FA mobile app?
Ans: Only the mobile app is supported currently.

i. Is it possible to use VPN clients such as OpenVPN or Tunnelblick?
Ans: Only Checkpoint VPN clients are supported.

j. Who should I contact if the Duo 2FA does not work?
Ans: Submit an incident through the ASPIRE 2B Service Desk Portal or email servicedesk@nscc.sg.

k. Why is my job still queued?

Ans: A job may remain queued depending on the requested resources, queue availability and current system usage. Check the job details with **`qstat -f <jobID>`**; if further assistance is required, provide the Job ID to the ASPIRE 2B Service Desk.

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

Ans: Submit an incident through the ASPIRE 2B Service Desk Portal or email servicdedesk@nscc.sg.

14. References

1) List of Documentation Guides
2) VPN Setup Guides for Window Users
3) VPN Setup Guides for MAC Users
4) User Onboarding Guide