




NTU College of Public Health

Introduction of computational resources



IT support of NTUCPH

Po-Chang Hsiao

2025.04.10

Computational resources

HPC




GPU



VDI



Structure of IT support



High-performance computing (HPC)

• Linux



- DELL R740 * 2
 - Intel Xeon Gold 2.3 GHz (3.2 GHz)
 - 24C/48T
 - 384 GB RAM
 - Ubuntu 20.04 LTS Linux
- HPE ProLiant DL380 Gen10
 - Intel Xeon Gold 2.9 GHz (3.9 GHz)
 - 32C/64T
 - 960 GB RAM
 - Ubuntu 20.04 LTS Linux
- HPE ProLiant DL385 Gen11
 - AMD EPYC 9354 3.25 GHz (3.8 GHz)
 - 64C/128T
 - 1024 GB RAM
 - Ubuntu 22.04 LTS Linux

• Windows




- DELL R740
 - Intel Xeon Gold 2.3 GHz (3.2 GHz)
 - 24C/48T
 - 384 GB RAM
 - Windows 10 Professional

```
top - 09:53:56 up 7 days, 23:18, 1 user, load average: 1.08, 1.04, 1.00
tasks: 819 total, 2 running, 817 sleeping, 0 stopped, 0 zombie
%Cpu(s): 1.6 us, 0.0 sy, 0.0 ni, 98.4 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
%Mem0: 3.1 us, 0.0 sy, 0.0 ni, 96.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
%Mem1: 0.0 us, 0.0 sy, 0.0 ni, 99.9 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
Mem Mem : 907219.1 total, 373999.9 free, 8756.4 used, 544532.0 buff/cache
Mem Swap: 2048.0 total, 2042.4 free, 5.6 used, 952325.4 avail Mem

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  MEM%   TIME+  COMMAND
191356 kevin    20   0 13.7g 3.3g 109532 R 100.0  0.4 528:50.30 glime
221319 root      20   0 13268 7924 6812 S  0.7  0.0  0:00.02 sshd
125 root      0   0 0      0      0 S  0.0  0.0  0:00:01 migration/10
1360 systemd  20   0 24980 13472 9228 S  0.3  0.0  0:16.37 systemd-resolve
1368 avahi    20   0 9232 4616 3388 S  0.3  0.0  2:34.19 avahi-daemon
1370 message+ 20   0 16096 6788 3920 S  0.3  0.0  0:31.75 dbus-daemon
1371 root      20   0 747048 21024 17256 S  0.3  0.0 18:35.76 NetworkManager
1376 root      20   0 82300 4040 3264 S  0.3  0.0  7:22.48 irqbalance
130517 root      20   0 0      0      0 I  0.3  0.0  0:02.61 kworker/14:2-mm_percpu_wq
188557 root      20   0 28992 9436 7572 S  0.0  0.0  0:00.15 cupsd
188559 root      20   0 185400 20044 11564 S  0.3  0.0  0:00.90 cups-browsed
221128 franstie 20   0 14120 5936 4448 S  0.3  0.0  0:00.02 sshd
221302 franstie 20   0 11136 4616 3248 R  0.3  0.0  0:00.24 top
1 root      20   0 168528 11992 8524 S  0.0  0.0  0:32.80 systemd
2 root      20   0 0      0      0 S  0.0  0.0  0:00.18 kthreadd
3 root      0 -20 0      0      0 I  0.1  0.0  0:00.00 rcu_gp
4 root      0 -20 0      0      0 I  0.1  0.0  0:00.00 rcu_perf_gp
5 root      0 -20 0      0      0 I  0.1  0.0  0:00.00 netns
```


- ssh
- Visual Studio Code



- RDP (Windows)
- Windows App (MacOS)

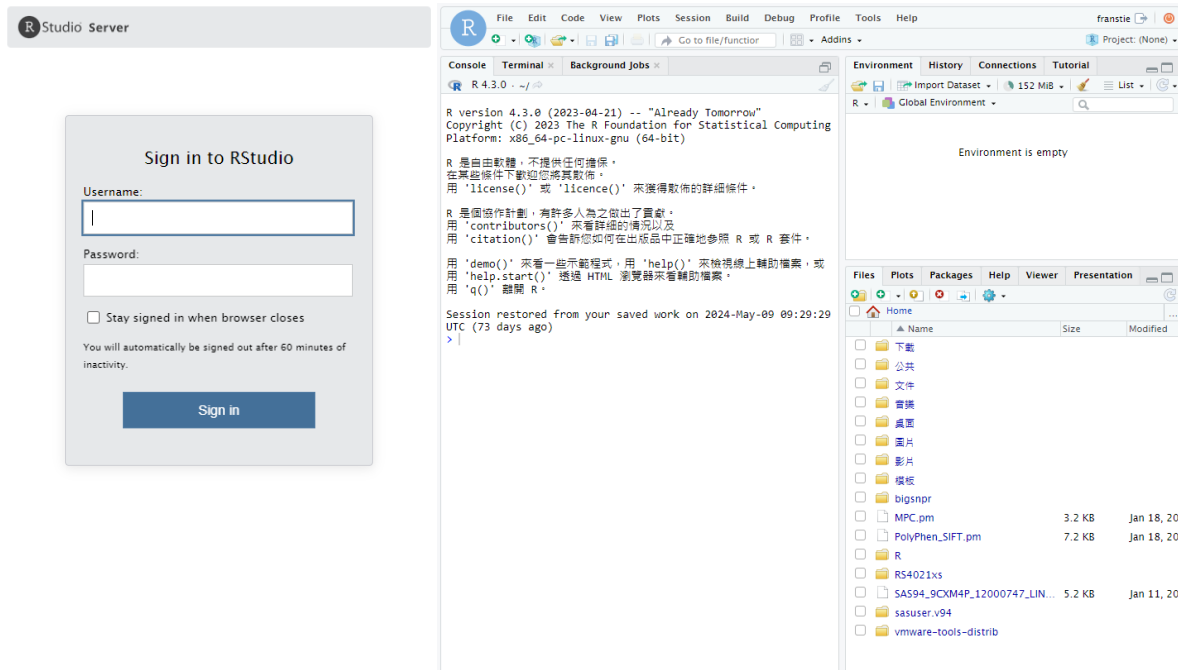
Softwares & Applications

```
----- /usr/share/modules/modulefiles -----
biology/BCFtools-1.2          genetics/GCTA/1.94.1
biology/Bedtools2-2.31.0     genetics/GLIMPSE2-2.0.1
biology/BWA-0.7.17           genetics/KING-2.7.7
biology/BWA-MEM2-2.2.1       genetics/MAGMA-1.10
biology/CrossMap-0.6.4       genetics/PLINK/1.07
biology/FastQC-0.1.2         genetics/PLINK/1.90
biology/GATK-4.4.0.0         genetics/PLINK/2.00
biology/HISAT-2.2.1          genetics/PLINK/multivariate
biology/HTSlib-1.17          genetics/PRScs_1.1.0
biology/Picard-3.1.0         genetics/PRScsx
biology/Samtools-1.17        genetics/PRSite_2.3.5
biology/Subread-2.0.5        genetics/REGENIE-3.3
biology/Trimmomatic-0.39     genetics/SHAPEIT5-5.1.1
biology/VCFtools-0.1.13     genetics/SnpTracker-1.0
conda_envs/bigsnp-1.11.8     genetics/VEP-110
conda_envs/GATK-4            program/Anaconda3/2024.10
conda_envs/LDSC-1.0.1        program/gcc/9.4.0
conda_envs/PICRUSt-2.5.1     program/Java/JDK-11.0.18
conda_envs/Q2-PICRUSt2-2023.2 program/Java/JDK-17.0.7
conda_envs/QIIME2/2023.2     program/Python/2.7.18
conda_envs/SAIGE-1.1.8       program/Python/3.9.13
conda_envs/TB-Profiler-4.4.0 statistics/Metal-2020-05-05
conda_envs/TB-Profiler-6.2.0 statistics/R/4.2.3
genetics/BOLT-LMM-2.4.1      statistics/R/4.4.0
genetics/GCTA/1.25.3         statistics/SAS-9.4
```




Web-based applications

R-Studio



JupyterHub



GPU computing

• Linux



- HPE ProLiant DL385 Gen11
 - AMD EPYC 9354 3.25 GHz (3.8 GHz)
 - 64C/128T
 - 1024 GB RAM
 - Ubuntu 22.04 LTS Linux

- Nvidia Tesla L40s * 2
 - RAM: 48 GB DDR6
 - CUDA core: 18,176
 - Tensor core: 568
 - RT core: 142
- Version
 - Python: 3.12.7
 - NVIDIA: 535.183.01
 - CUDA: 12.2.91
 - cudnn: 8.9.7
 - tensorflow-gpu: 2.16.1
 - Pytorch: 2.51+cu121



• Windows



- DELL R740
 - Intel Xeon Gold 2.3 GHz (3.2 GHz)
 - 4C/8T
 - 256 GB RAM
 - Windows 10 Professional

- Nvidia Tesla T4
 - RAM: 16 GB DDR6
 - CUDA core: 2,560
 - Tensor core: 320
- Version
 - Python: 3.7
 - NVIDIA: 412.31
 - CUDA: 10.0
 - cudnn: 7.6.0
 - tensorflow-gpu: 2.0
 - Pytorch: 1.4



GPU applications

- AI
- Machine Learning
- GPU-accelerated computing

```
[5]: import os
os.environ['TF_ENABLE_ONEDNN_OPTS'] = '0'

import tensorflow as tf

# 確保 TensorFlow 2.16 被正確加載
print("TensorFlow Version:", tf.__version__)

# 檢查可用的 GPU
gpus = tf.config.list_physical_devices('GPU')
print("Available GPUs:", gpus)

if len(gpus) < 2:
    raise RuntimeError("至少需要兩張 GPU 才能執行此測試")

# 設置 TensorFlow 使用所有可用 GPU
try:
    for gpu in gpus:
        tf.config.experimental.set_memory_growth(gpu, True)
except RuntimeError as e:
    print(e)


# 創建策略來自動分配計算到多張 GPU
distributed_strategy = tf.distribute.MirroredStrategy()

with distributed_strategy.scope():
    # 建立簡單的 TensorFlow 計算圖
    a = tf.random.normal([10000, 10000])
    b = tf.random.normal([10000, 10000])
    c = tf.matmul(a, b) # 簡單的矩陣相乘操作

    print("Computation completed on multiple GPUs.")

TensorFlow Version: 2.18.0
Available GPUs: [PhysicalDevice(name='/physical_device:GPU:0', device_type='GPU'), PhysicalDevice(name='/physical_device:GPU:1', device_type='GPU')]
INFO:tensorflow:Using MirroredStrategy with devices ('/job:localhost/replica:0/task:0/device:GPU:0', '/job:localhost/replica:0/task:0/device:GPU:1')
Computation completed on multiple GPUs.
```

- 3D graphics
- Geographic Information System (GIS)




Virtual Desktop Infrastructure (VDI)

- DELL R730 * 2
 - Intel Xeon Gold 2.1 GHz (3.0 GHz)
 - 24C/48T
 - 384 GB RAM
- DELL R740 * 4
 - Intel Xeon Gold 2.3 GHz (3.2 GHz)
 - 24C/48T
 - 384 GB RAM




Purpose

- Server
 - DB Server (MariaDB, SQL Server...)
 - Web Server (Apache, IIS, Tomcat, Nginx...)
 - Data host
 - APP
 - IoT
- Customized computing
 - Hardware & software customized
 - Non-Ubuntu Linux dist.
 - Windows
 - Authorized software
 - Non-Chinese language version




Network Attached Storage (NAS)

- Extra storage for computing
- Total 650 TB
- Important Notes
 - **Not personal cloud storage**
 - **Without backup**



Requirement of computational resources



Case illustration:

1. 24 threads, 256 GB memory for data analysis
2. Additional 20TB hard disk space for data analysis
3. 10 VDIs with personal licensed software for a workshop
4. A Linux server for sensor data receiving and building an alarming system

FAQ 1. Who may want to apply?

➤ Computing

- pc, laptop user
 - CPU clock is low
 - CPU cores are not enough
 - Memory is not enough
 - Need High-end GPU
 - Hard disk space is not enough
- software platform issue
 - e.g. SAS

➤ Teaching

- Practice
- Workshop

➤ Server

FAQ 2. How to select?

➤ Windows

- Graphic User Interface (GUI): GIS 、 Java GUI...
- Windows software: SAS 、 SPSS 、 STATA 、 STELLA...

➤ Linux

- Command Line Interface (CLI)
- Linux software
- Web-Based application: R-Studio 、 JupyterHub...

▣ Ask for advice at R551 if had no idea

FAQ 3. About application

➤ Windows

- Period of use: 4 weeks (including holidays)
 - Can extend for 4 weeks time by time without other applications get in line
- Only teachers can apply a long-term whole machine (with other rules)

➤ Linux

- Apply once until graduation or leave the college

➤ VDI

- Customized computing
 - Same as Windows
- Server
 - Only teachers can apply
 - Ask for advice at R551 in advance

FAQ 4. Other rules

- Utilization
 - Do not use VDI as personal cloud storage space
 - Files will be deleted and won't be backed up after the expiration of the (extended) application
 - Illegal download, bitcoin mining, network attack, computer games, etc. are strictly forbidden

How to apply

Who can apply:

- HPC, GPU, computing VDI
 - NTUCPH teachers, post-doc fellows, students, research assistants
- Server VDI
 - NTUCPH teachers


Contact & IT advice

- **Mr. Po-Chang Hsiao**
- R551
- 3366-8043
- hsiaopc@ntu.edu.tw



Options

Type	(GPU) HPC			VDI	
	Linux	Windows (4 weeks)	GPU Windows (4 weeks)	Server Computing	Windows (4 weeks)
OS	Linux Ubuntu 20	Windows 10	Windows 10	- Linux - Windows Server - Windows 10/11	Windows 10/11
Hardware	Machine spec	- CPU: 24C/48T - RAM: 384GB	- CPU: 4C/8T - RAM: 384GB - VRAM: 16GB	Customized	Customized
Identity	- Teacher - Post-doc - student - Research assistant	- Teacher - Post-doc - student - Research assistant	- Teacher - Post-doc - student - Research assistant	Teacher	- Post-doc - student - Research assistant
Connection	- SSH - Web-based	RDP	RDP	- SSH - RDP	- RDP ² - Horizon Client ³
Location	Anywhere	Anywhere	Anywhere	Anywhere ¹	Anywhere ^{2, 3}
Memo				1: prepare NTU global ip : 140.112.117.X	2: prepare NTUCPH lan ip 3: NTUCPH lan



NewsNews

News Releases >

Featured Stories >

Mentorship
System

Handbook for
New Faculty

Information Technology Support



High Performance Computing (HPC)

Suitable for those who needs lots of CPU cores or memories for long time computing or complex analysis. Linux and Windows operating system are all available.

Who can apply:

- NTUCPH faculty members, post-doc fellows, students, research



Thanks for your attention.

