

高效能科學計算說明會暨使用者教育訓練工作坊 2023  
NSTCCore User Training Workshop 2023

# **Introduction of NSTCCore Computing & Storages Services**

**<https://nstccore.twgrid.org>**

# Introduction of NSTCCore Computing Service

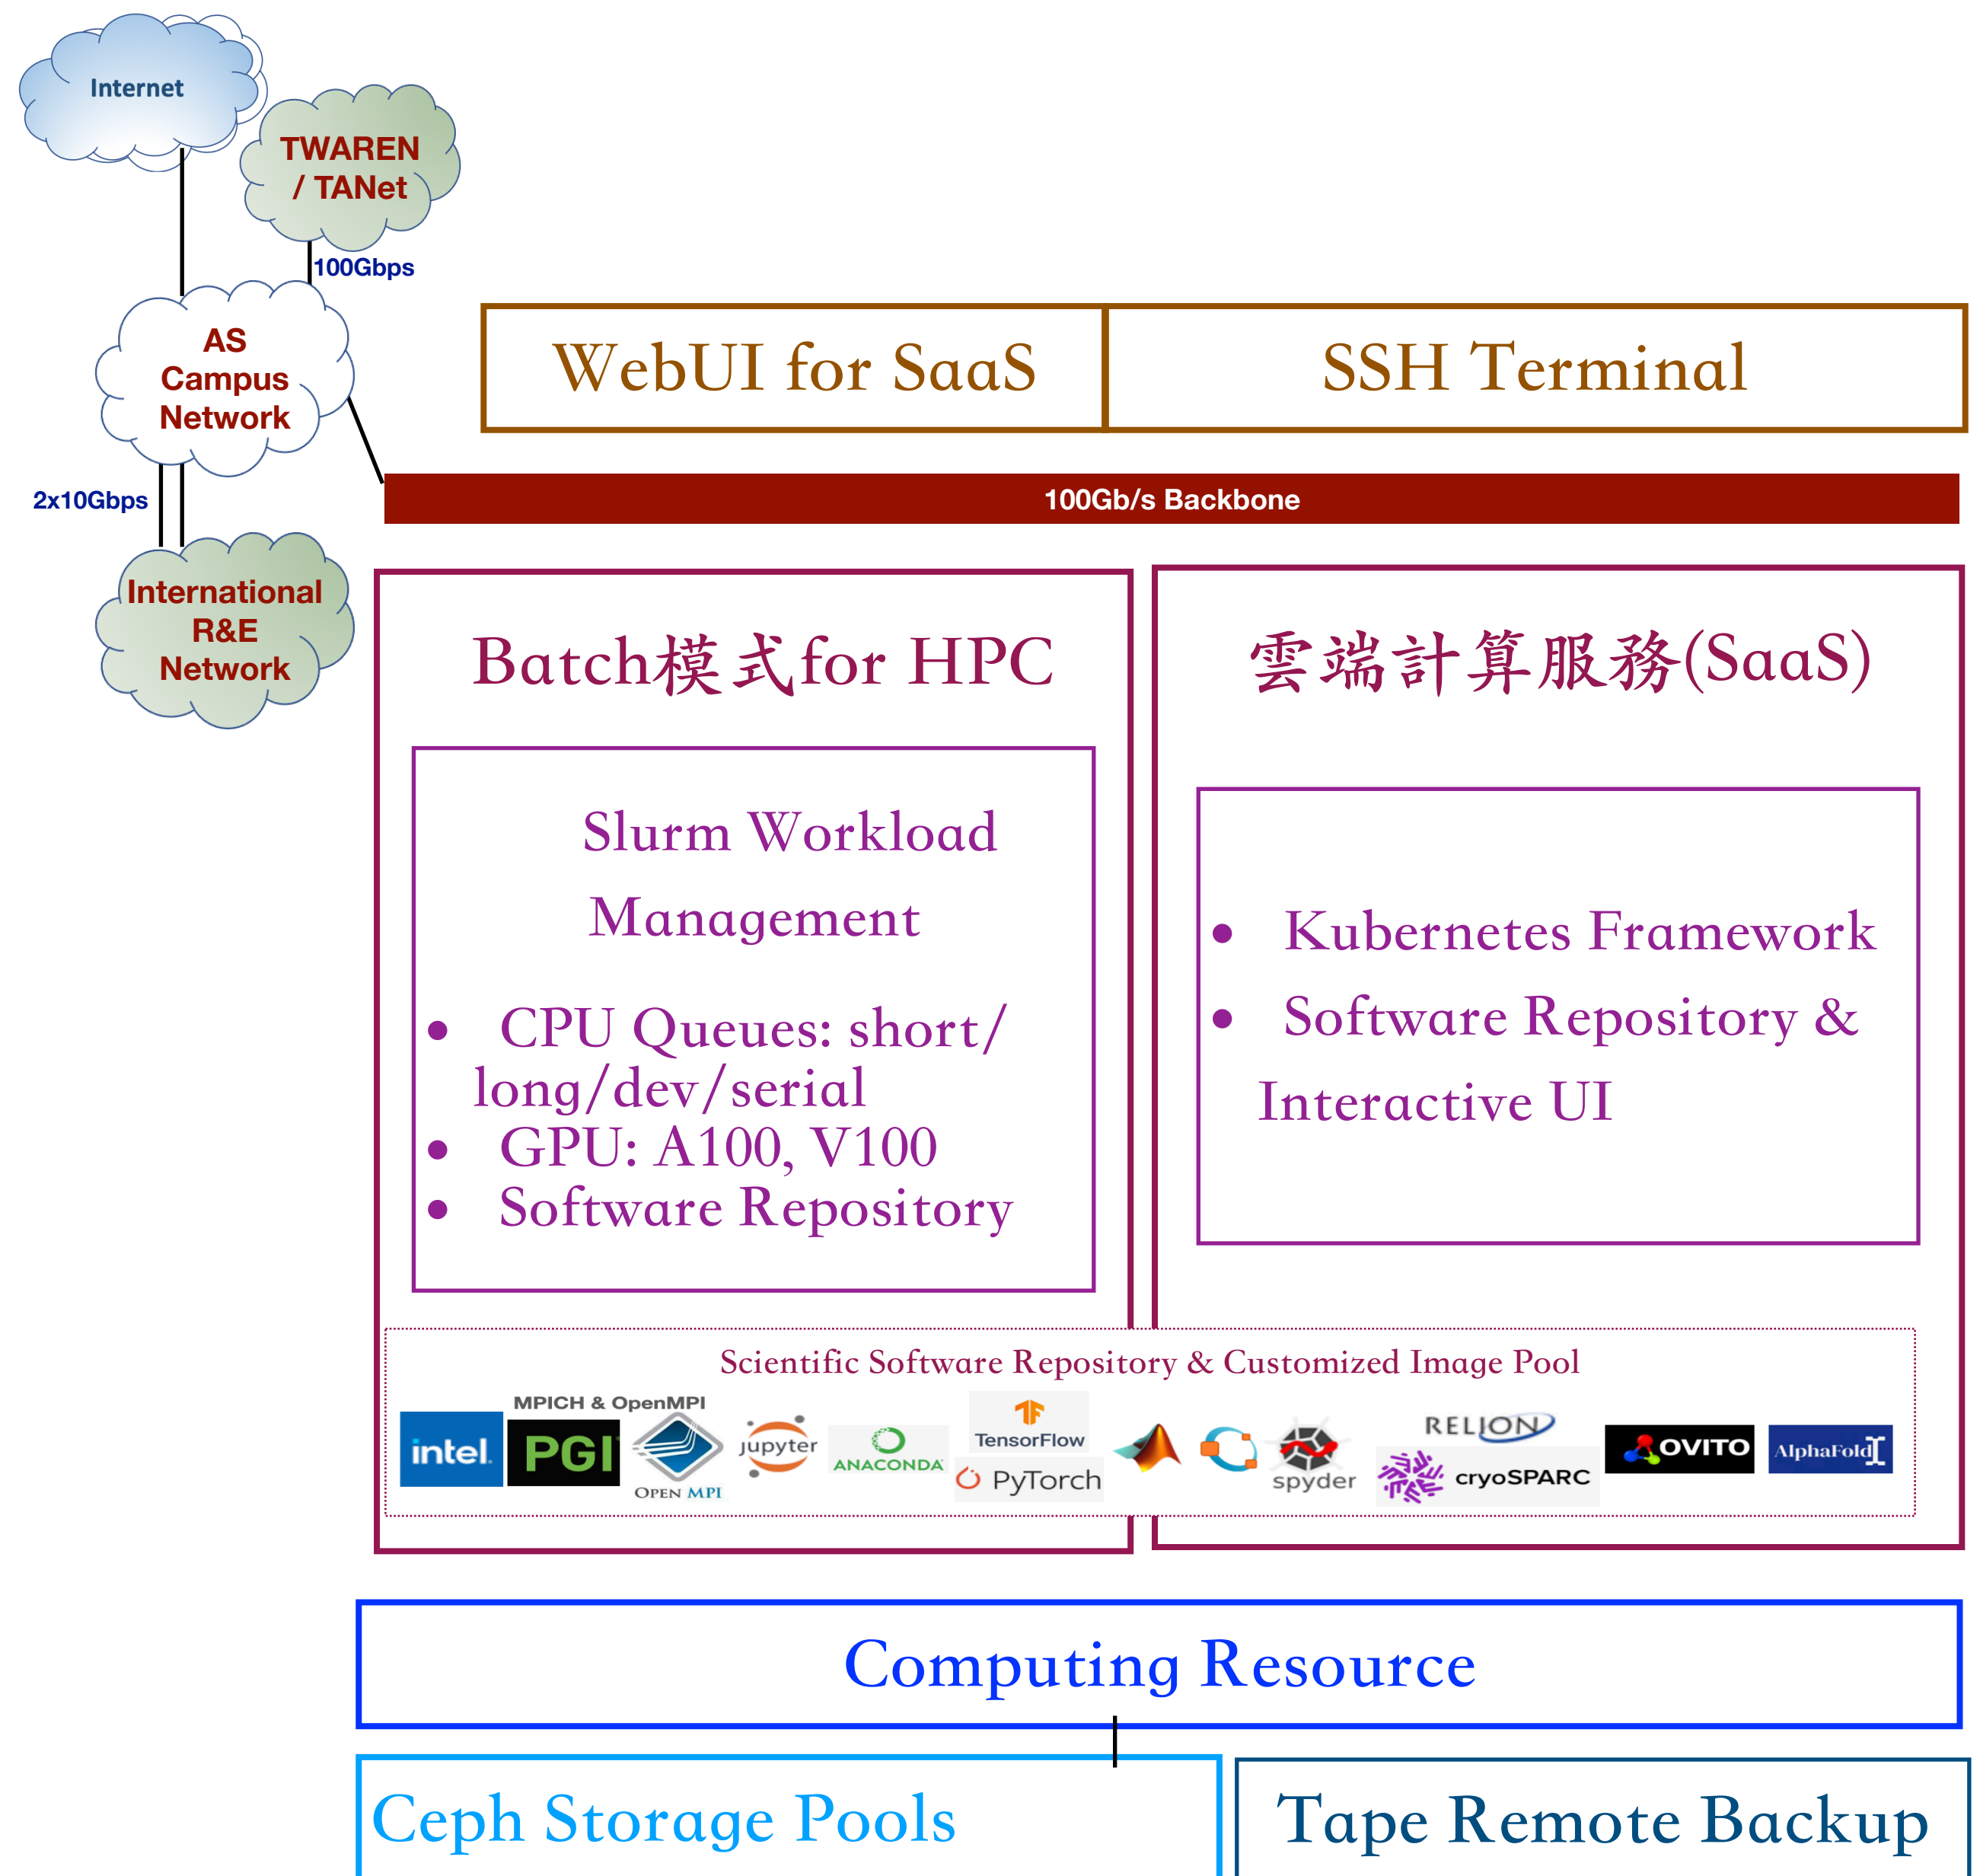
- Computing Service
- Storage Service
- Technical Support
- User Management

|              | 2023.08~2023.11   | 2023.12~2024.09   | 2024.09~2025.08   | 2025.08~2026.06   |                  |
|--------------|---|---|---|---|------------------|
| CPU          | 2976 Cores<br>*Current Resources                          | 1792 Cores<br>*New<br>+<br>768 Cores<br>*AMD Rome         | 3584 Cores<br>*New<br>+<br>768 Cores<br>*AMD Rome         | 5376 Cores<br>*New<br>+<br>768 Cores<br>*AMD Rome         | *後續計算能量依計畫核定狀況決定 |
| GPU          | V100 - 32 boards<br>A100 - 8 boards<br>*Current Resources | V100 - 32 boards<br>A100 - 8 boards<br>*Current Resources | V100 - 32 boards<br>A100 - 8 boards<br>*Current Resources | V100 - 32 boards<br>A100 - 8 boards<br>*Current Resources |                  |
| Storage (PB) |   | 3<br>*Buy-in every year                                   | 6   | 9   |                  |
| Tape (PB)    |   | 4<br>*Buy-in every year                                   | 8   | 12  |                  |

表、計算資源購置規劃表 2023.7

# Scientific & HPC Computing Service

- Batch Jobs Computing Service
  - Slurm Work Management System
- Interactive Jobs Computing Service
  - Scalable & Virtualized Service-as-Service (SaaS) Service Model
  - Software on-demand Web-based UI
  - Customized Application Deployment



System Architecture of NSTCCore Computing Service

# Batch Jobs Computing Service

- Slurm System Architecture
  - Scalable Cluster Management and Job Scheduling System

- Computing Nodes

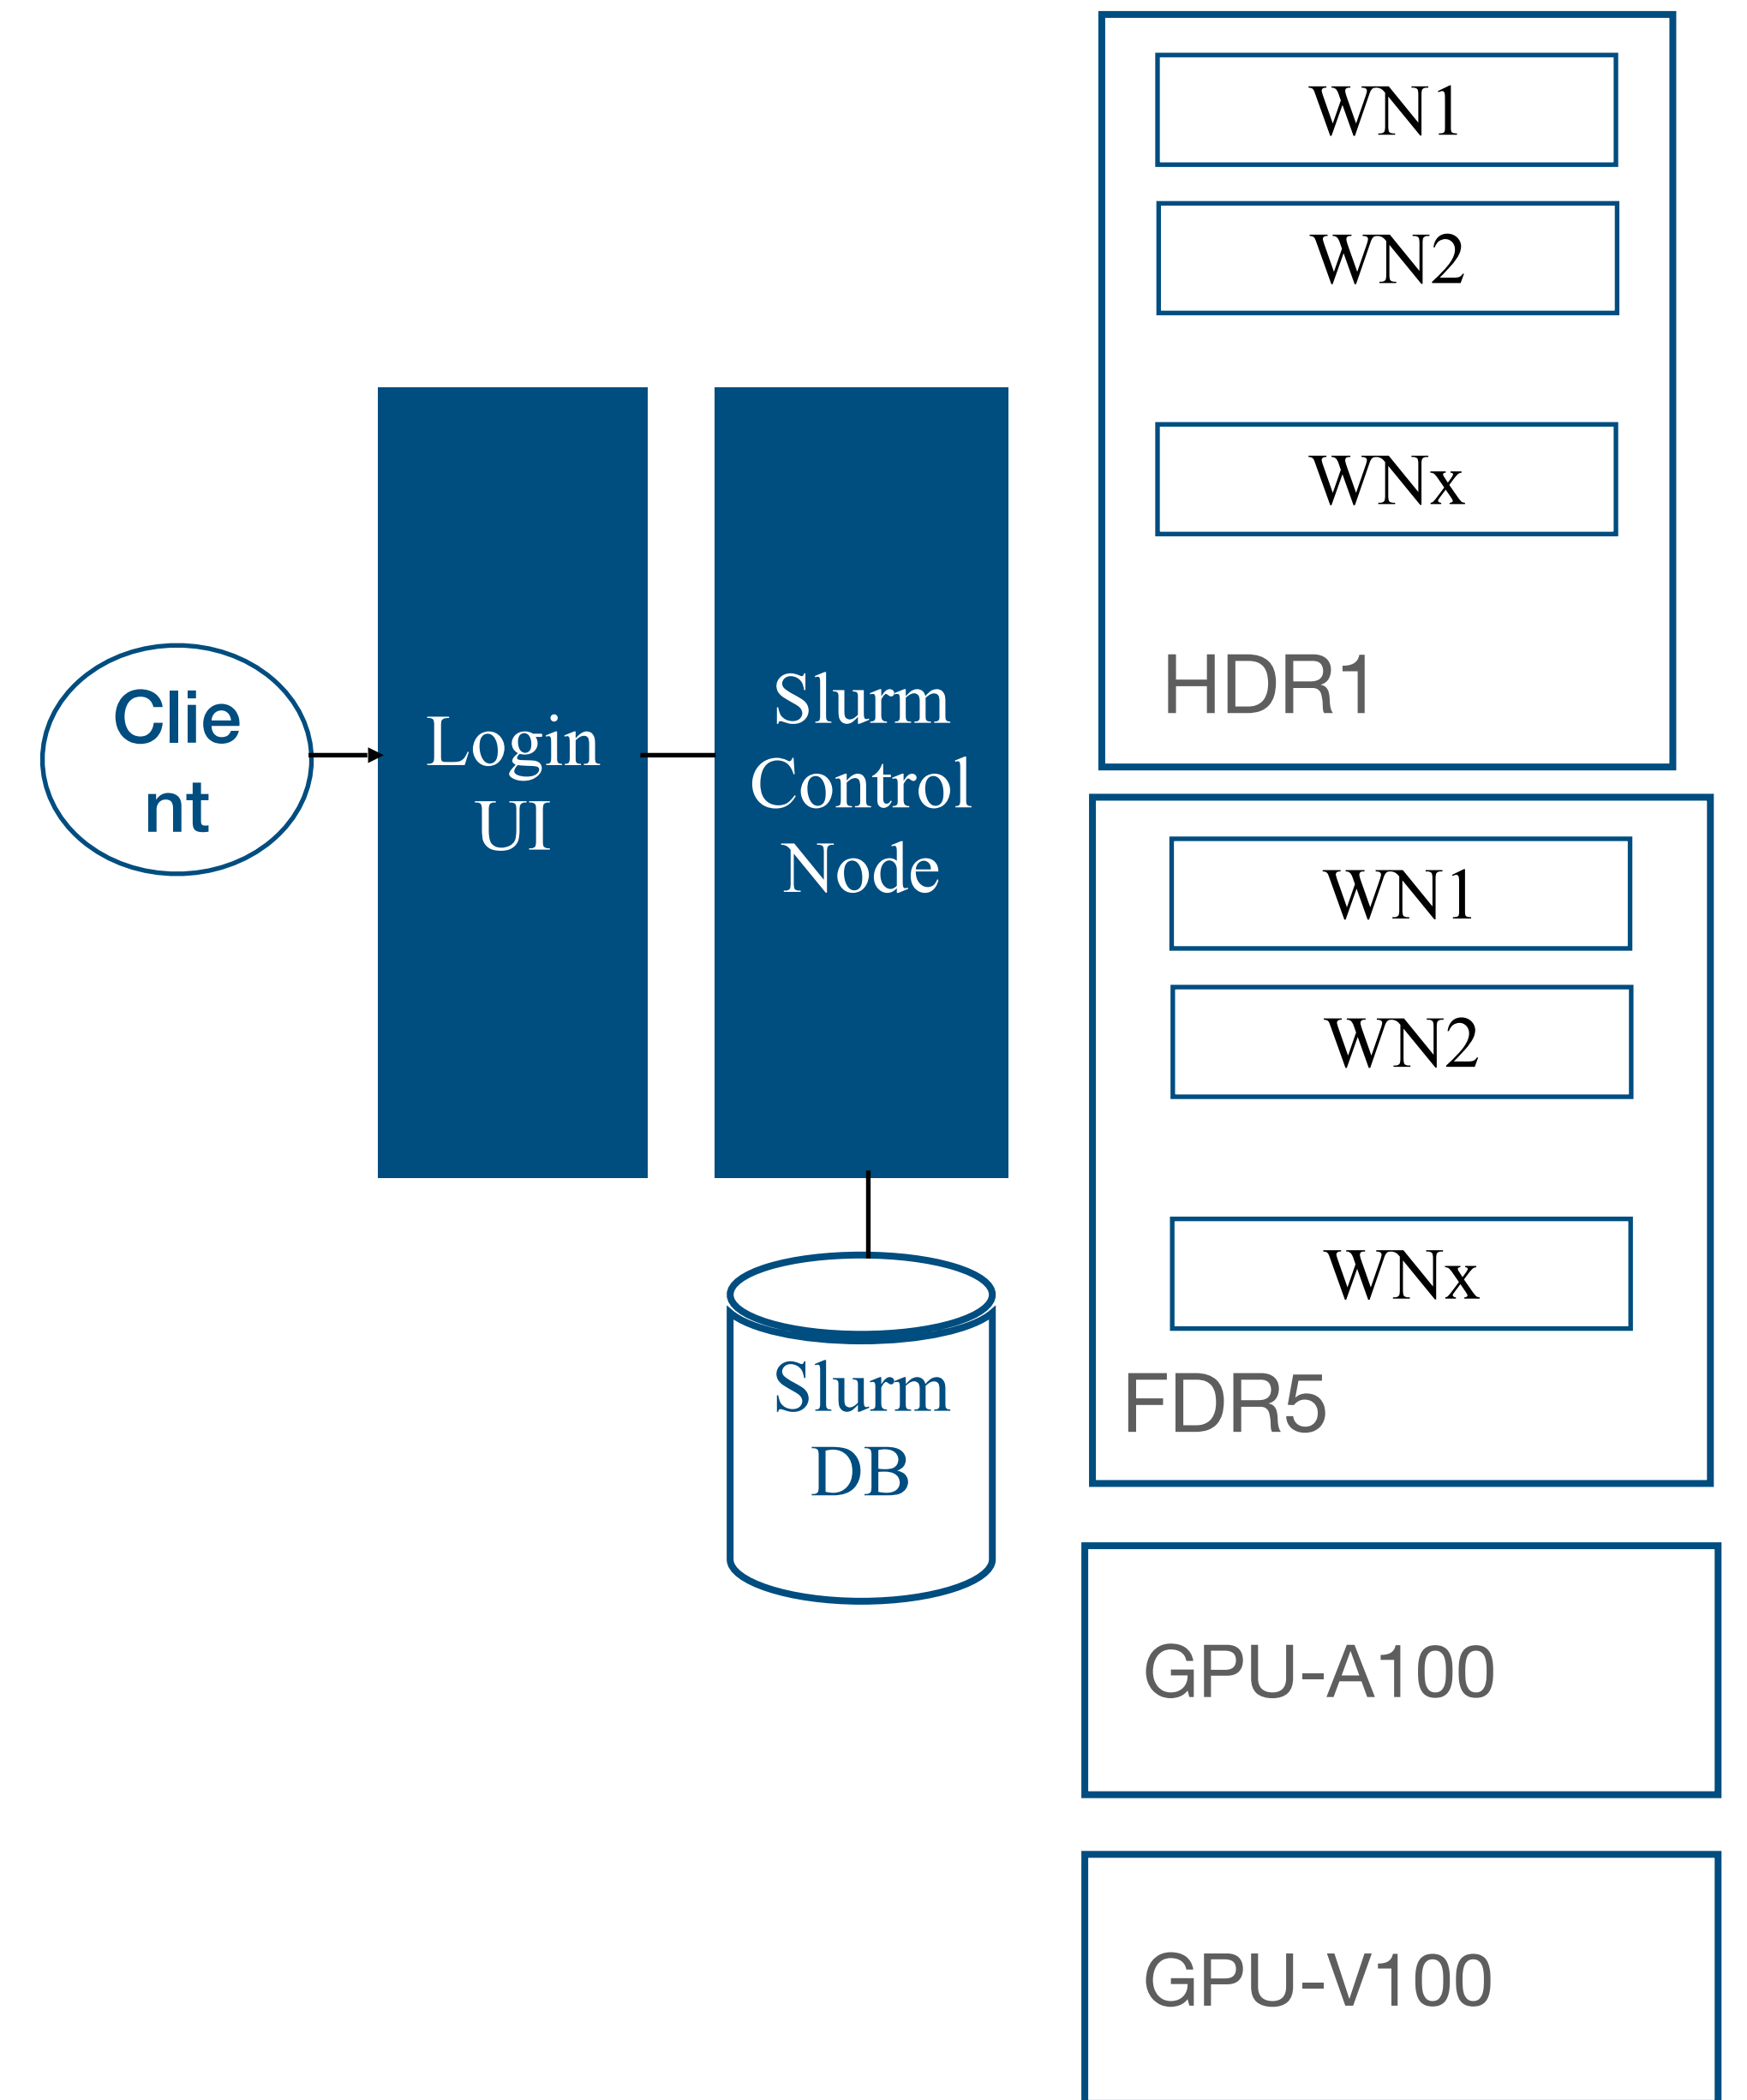
- CPU

| Cluster | CPU                                 | Nodes | Cores-Per-Node | Cores | RAM(GB) | Inter-connection    |
|---------|-------------------------------------|-------|----------------|-------|---------|---------------------|
| FDR5    | Intel® Xeon® CPU E5-2650 v4@2.20GHz | 92    | 24             | 2208  | 128     | IB:FDR, Eth:10Gbps  |
| HDR1    | AMD Rome 7662 @2.0GHz               | 6     | 128            | 768   | 1536    | IB:HDR, Eth:100Gbps |

- GPU - A100 、 V100

- Jobs Working Space:

- Ceph Cloud Filesystem
  - Local Disk [HDR1] SSD 1TB & [FDR5] NO SSD



Slurm System Architecture

# Batch Jobs Computing Service - Slurm

## System Spec

- OS : CentOS 7
- Login (SSH) : [slurm-ui.twgrid.org](http://slurm-ui.twgrid.org)
- Mount Space :
  - User Space : /dicos\_ui\_home/{user}
  - Group Space : /ceph/work/{group}
  - Backup Space : /ceph/project/{group}(in the future)
- Compiler : Intel gcc 、 AOCC, openACC and MPI & OpenMP repository
- Software Repository :
  - Scientific Packages : Root 、 MATLAB 、 R 、 Octave
  - Anaconda Python packages : TensorFlow, PyTorch, PyRoot..etc
    - GPU - CUDA v12.1
  - Some Customized Requirements needs to deploy by Singularity, like AlphaFold. (\* Build fee)

# Interactive Jobs - SaaS Computing Service

- Kubernetes and Openstack
  - High extensible and reliable virtual environment
- Customized Application Deployment
- Images Repository
  - JupyterLab and various scientific applications
  - Built by user's requirements
- Software-on-demand Web UI
  - No installation and easy to adopt
- Working Space
  - Ceph Filesystem

## System Spec

- Node Spec : NVIDIA GPU V100 、 A100
- Service Web Portal :
  - [dicos.grid.sinica.edu.tw](https://dicos.grid.sinica.edu.tw)
- Mount Space :
  - User Space : /dicos\_ui\_home/{user}
  - Group Space : /ceph/work/{group}
  - Backup Space : /ceph/project/{group}

# SaaS for Virtualized Computing Service

## Scientific Software Repository

- Interactive : Ovito(Molecular Dynamics) 、 cisTEM 、 RELION(Medical Image Reconstruction)
- BioMedical : Cryosparc (\* License required from users )
- Anaconda Python packages for ML : JupyterLab 、 TensorFlow 、 PyTorch 、 PyRoot 、 DeepMD(Molecular Dynamics)···etc.




## When your job needs

- Interactive UI
- Specific OS or Application required
- Dedicate node for rapid development for multi-core or GPU to develop and testing your task




# <https://dicos.grid.sinica.edu.tw/dockerapps/>




**CryoSPARC 1080ti**  
Version: 3.3.2  
Resources: 52%

Launch




**CryoSPARC RTX3090**  
Version: 4.0.2  
Resources: 44%

Launch




**PyRoot**  
Version: GPU with 1080ti  
Resources: 52%

Launch ▾




**RELION 4 beta**  
Version: V4  
Resources: 52%

Launch ▾



**Triton**  
Version: 22.01-py3 (GPU P100)  
Resources: 50%

Launch ▾



**Jupyter Lab GPU A100**  
Version: GPU with Tensorflow A100  
Resources: 62%

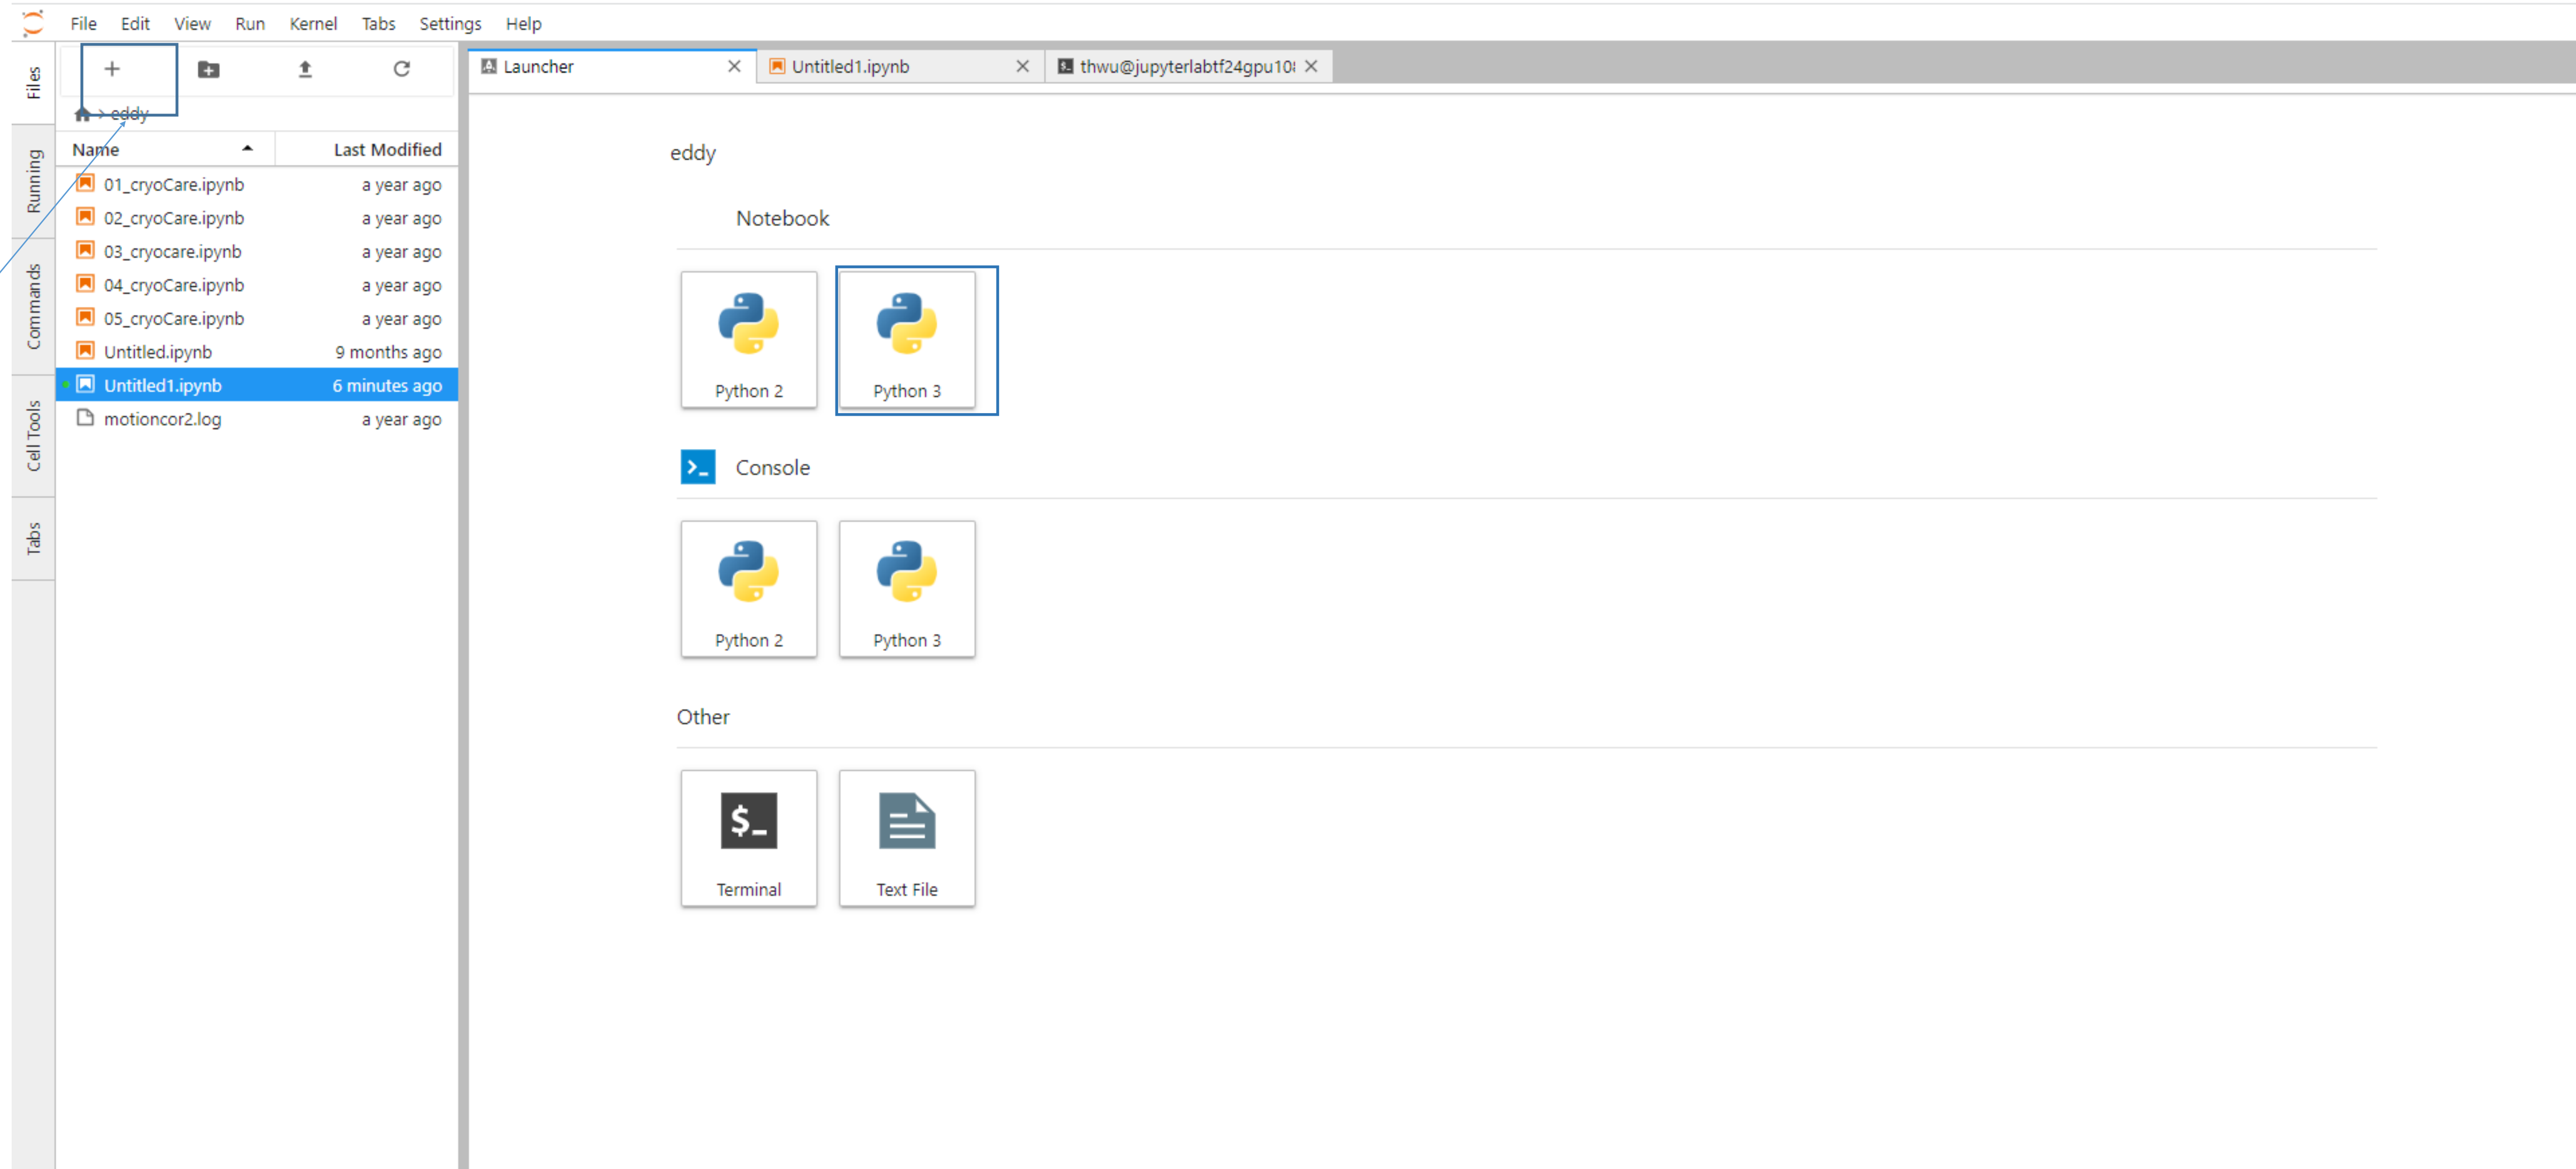
Launch ▾

- 3 days
- 7 days
- 10 days**

# The path of Disk Space

- All user:  
  /dicos\_ui\_home/{user} (UI)
- cryoEM group:  
  /activeEM/data/{group}/{user}
- NSTCCore group:  
  /ceph/work/{group}/{user}

# Example 1.1 (Open a Jupyter RTX 3090)



The screenshot displays the JupyterLab interface. The top menu bar includes File, Edit, View, Run, Kernel, Tabs, Settings, and Help. The left sidebar shows the 'Files' view for the 'eddy' environment, listing several IPython notebooks and a log file. The main area is the 'Launcher' view, which offers options to create a new Notebook, Console, Terminal, or Text File. The 'Python 3' notebook option is highlighted with a blue border.

Files

Running

| Name              | Last Modified |
|-------------------|---------------|
| 01_cryoCare.ipynb | a year ago    |
| 02_cryoCare.ipynb | a year ago    |
| 03_cryocare.ipynb | a year ago    |
| 04_cryoCare.ipynb | a year ago    |
| 05_cryoCare.ipynb | a year ago    |
| Untitled.ipynb    | 9 months ago  |
| Untitled1.ipynb   | 6 minutes ago |
| motioncor2.log    | a year ago    |

Commands

Cell Tools

Tabs

Launcher

eddy

Notebook

Python 2 Python 3

Console

Python 2 Python 3

Other

Terminal Text File

# Example 1.2 (Open a Jupyter RTX 3090)

The screenshot displays a JupyterLab environment. On the left, a sidebar shows a file browser with a table of files:

| Name              | Last Modified |
|-------------------|---------------|
| 01_cryoCare.ipynb | a year ago    |
| 02_cryoCare.ipynb | a year ago    |
| 03_cryocare.ipynb | a year ago    |
| 04_cryoCare.ipynb | a year ago    |
| 05_cryoCare.ipynb | a year ago    |
| Untitled.ipynb    | 9 months ago  |
| Untitled1.ipynb   | 4 minutes ago |
| motioncor2.log    | a year ago    |

The main area shows a code cell for 'Untitled1.ipynb' with the following Python code:

```
In [2]: # -*- coding: utf-8 -*-
import numpy as np
import math

# Create random input and output data
x = np.linspace(-math.pi, math.pi, 2000)
y = np.sin(x)

# Randomly initialize weights
a = np.random.randn()
b = np.random.randn()
c = np.random.randn()
d = np.random.randn()

learning_rate = 1e-6
for t in range(2000):
    # Forward pass: compute predicted y
    # y = a + b x + c x^2 + d x^3
    y_pred = a + b * x + c * x ** 2 + d * x ** 3

    # Compute and print loss
    loss = np.square(y_pred - y).sum()
    if t % 100 == 99:
        print(t, loss)

    # Backprop to compute gradients of a, b, c, d with respect to loss
    grad_y_pred = 2.0 * (y_pred - y)
    grad_a = grad_y_pred.sum()
    grad_b = (grad_y_pred * x).sum()
    grad_c = (grad_y_pred * x ** 2).sum()
    grad_d = (grad_y_pred * x ** 3).sum()

    # Update weights
    a -= learning_rate * grad_a
    b -= learning_rate * grad_b
    c -= learning_rate * grad_c
    d -= learning_rate * grad_d

print(f'Result: y = {a} + {b} x + {c} x^2 + {d} x^3')
```

The output of the code cell shows the loss values at intervals of 100 iterations:

```
99 712.447821888372
199 478.3262177418421
299 322.276702381664
399 218.21288593989044
499 148.7807326213537
599 102.43000968080545
699 71.47013935582402
799 50.77823131954953
```

# Example 1.3 (Open a Jupyter RTX 3090 with terminal)



The screenshot shows a JupyterLab interface with a file browser on the left and a terminal window on the right. The file browser shows a directory named 'eddy' containing several IPynb files and a log file. The terminal window shows the output of the 'nvidia-smi' command, which displays GPU information for a GeForce GTX 1080 Ti.

```
[thwu@jupyterlabtf24gpu1080ti-thwu ~]$ nvidia-smi
Mon Nov 28 07:32:53 2022
+-----+
| NVIDIA-SMI 418.56      Driver Version: 418.56      CUDA Version: 10.1     |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|    0   GeForce GTX 108...    On          | 00000000:0F:00.0 Off  |          N/A         |
| 29%   34C   P8      8W / 250W | 1MiB / 11178MiB |    0%    Default     |
+-----+-----+
+-----+
| Processes:                                     GPU Memory |
|  GPU       PID    Type   Process name                               Usage      |
+-----+-----+
| No running processes found                       |
+-----+
[thwu@jupyterlabtf24gpu1080ti-thwu ~]$
```

# Storage Service

- Ceph Filesystem
  - An open source distributed filesystem
  - High-Throughput

## User Home Space

- /dicos\_ui\_home/{user\_account}
- 100GB Free space

## Working Space

- /ceph/work/{group\_account}
- Every Group has 3TB free space, PI has full permissions for data in this space. Buy more space according to your computing needs, 1TB/days as a purchase unit.

## Backup Space \*purchase by the end of 2023

## Tape as Backup and Preservation Service in the future

- /ceph/project/{group\_account}
- Backup and long-term preserved space. Buy as needed. 1TB/years as a purchase unit.

## Data Transfer

- Transfer by SFTP via [dicos-sftp.twgrid.org](https://dicos-sftp.twgrid.org)

# Technical Support

- Help Desk & Service Notification



- Rocketchat online chat - <https://rocketchat.twgrid.org/channel/general>
- Email - [dicos-support@twgrid.org](mailto:dicos-support@twgrid.org)
- Portal - <https://nstccore.twgrid.org>
  - Release up-to-date services status, group usage, pricing and technical relevant information
- Service Portal - <https://dicos.grid.sinica.edu.tw>
  - SaaS Computing Service
  - PI & User Management
- Training & Workshop
  - Regular workshops every 3 months
  - Technical support & consulting services

# User Management - User Account



- Apply your account
  - Group Account: <http://canew.twgrid.org/ApplyAccount/groupcreate.php>
  - User Account: <http://canew.twgrid.org/ApplyAccount/ApplyAccount.php>
  - PI approval for Member's application
  - Password & Account Expiration (ISO security)
    - 1 year validation, password & account expiry notification will be sent on 7, 15 and 30 days to expiration.
    - Account Deletion: Your account & user space(UI home directory, work directory and DiCOSBox) will be removed after 6 months of expiration.



# User Management - Group

- Members management
  - Members list
  - Abnormal member usage report
  - Member's usage review
- Resource usage Management & Budget Control
  - Monthly Bill will be delivered
  - Resource Usage
  - Payment Management

| 使用者 (username) | 姓名(name)      | Email                   | Expired Date            | Active | Joined Date             | Last Login              | Storage Usage (UI Home)  |
|----------------|---------------|-------------------------|-------------------------|--------|-------------------------|-------------------------|--|
| chiong         | CHAN-HIN IONG | chiong@me.com           | Jun. 13, 2024, 00:00 AM | True   | Aug. 22, 2018, 00:00 AM | Jul. 31, 2023, 03:41 AM | <b>33.4G/100G</b><br><small>Latest Update: 2023-07-31 00:50:03</small> |
| dickie         | Dickie Chang  | dickie.chang@twgrid.org | Oct. 12, 2023, 00:00 AM | True   | Dec. 15, 2022, 07:16 AM |                         |  |
| eric           | 嚴漢偉 嚴漢偉       | Eric.Yen@twgrid.org     | Oct. 18, 2023, 00:00 AM | True   | Mar. 11, 2019, 00:12 AM | Jul. 06, 2023, 04:27 AM | <b>0.0G/100G</b><br><small>Latest Update: 2023-07-31 00:50:03</small>  |
| ericyen        | Eric YEN      | Eric.Yen@twgrid.org     |                         |        |                         |                         |  |
| etomo          | etomo etomo   | etomo@twgrid.org        |                         |        |                         |                         |  |
| felixlee       | Felix Lee     | felix@twgrid.org        |                         |        |                         |                         |  |

使用統計摘要(TOTAL SUMMARY TABLE) 使用明細(DETAILED USAGE ACCOUNTING)

• 計算資源使用統計(COMPUTING USAGE OF GROUP) – 預估使用費(INITIAL COST): NT \$15,878

| 使用者(username) | 姓名(name)      | CPU (SRU) | GPU (SRU) | 預估費用 (Initial Cost Estimation) |
|---------------|---------------|-----------|-----------|--------------------------------|
| chiong        | CHAN-HIN IONG | 175       | 7,601     | 15,552                         |
| jyou          | Jingya You    | 1         | 54        | 110                            |
| rudy          | 陳侑廷           | 3         | 0         | 6                              |
| thwu          | Tsung-Hsun Wu | 4         | 101       | 210                            |

使用統計摘要(TOTAL SUMMARY TABLE) 使用明細(DETAILED USAGE ACCOUNTING)

使用明細(DETAILED USAGE ACCOUNTING)  
使用者(username): chiong (CHAN-HIN IONG) 用量異常回報(submit issue)

| Alert                    | #Instance              | A100      |           | FDR5      |           | P100      |           | RTX3090   |           | V100      |           |
|--------------------------|------------------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
|                          |                        | CPU (SRU) | GPU (SRU) | CPU (SRU) | GPU (SRU) | CPU (SRU) | GPU (SRU) | CPU (SRU) | GPU (SRU) | CPU (SRU) | GPU (SRU) |
| <input type="checkbox"/> | relion311rtx3090       | 3         |           |           |           |           |           | 151       | 6,320     |           |           |
| <input type="checkbox"/> | jupyterlabtf24gpu3090  | 2         |           |           |           |           |           | 11        | 474       |           |           |
| <input type="checkbox"/> | matlab                 | 2         |           |           |           |           |           |           |           | 0         | 6         |
| <input type="checkbox"/> | jupyterlabgpu26a100    | 1         | 1         | 519       |           |           |           |           |           |           |           |
| <input type="checkbox"/> | openaccp100            | 1         |           |           |           | 11        | 282       |           |           |           |           |
| <input type="checkbox"/> | FDR5 slurm             | 2         |           |           | 0         | 0         |           |           |           |           |           |
| <input type="checkbox"/> | STORAGE USAGE (0.0 TB) |           |           |           |           |           |           |           |           |           |           |

說明(note)

- Group使用空間
- User使用空間
- User使用空間
- User使用空間
- User使用空間