



# 國科會 高效能核心科學計算中心 (NSTCCORE)

<https://nstccore.twgrid.org>

嚴漢偉 Eric Yen

中央研究院 物理所 網格中心  
2026. 1. 13

# NSTCCore @TPS2026

- **Seeking for advice, comment and discussion**
  - In order to improve the service and quality
- **Technical details could be explained at Booth #7**

<https://indico4.twgrid.org/event/70/>

**4:00 PM → 4:10 PM** 一、NSTCCore 計劃簡介 NSTCCore Project Overview

Convener: Prof. Song-Ming Wang (ASGC)

**4:10 PM → 4:30 PM** 二、資源與服務概況 Overview of Resources and Services

Convener: Prof. Eric Yen (ASGC)

**4:30 PM → 5:25 PM** 三、交流與討論 Discussion and Exchange

Convener: Mr Felix Lee (ASGC)

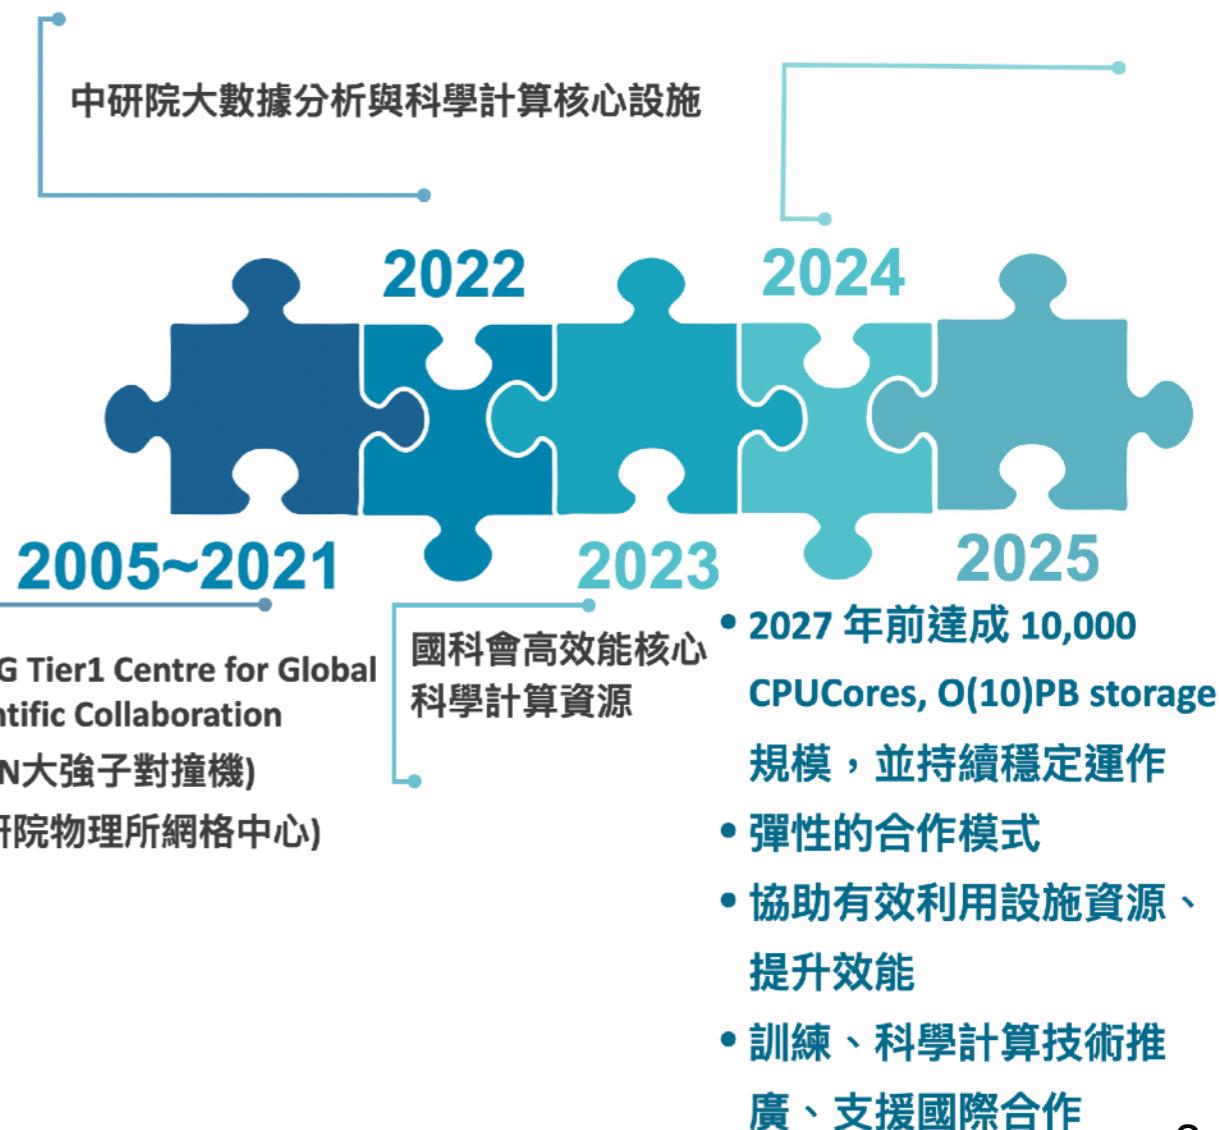
**5:25 PM → 5:30 PM** 四、問卷調查 Questionnaire

# Mid-Scale Scientific Computing Services for Accelerating Discovery & Innovation

- **Funded by NSTC, June 2023 - May 2026**
  - Jointly supported with the core facility funded by Academia Sinica since 2023
- **Reliability & Efficiency are the top priority**
- **Serving as your computational research partner, so that scientific groups could focus on sciences**
- **Flexible and customized services: resource arrangement, workflow integration, performance improvement, etc.**
- **Resource scale: providing services utilizing O(10,000) CPUCores, O(100) GPUs and O(10)PB Disk storage resources**
  - Able to support O(1,000) CPUCores parallel computing, multiple GPUs and big data analysis
  - Affordable cost and lower than operating by your own
  - Aligning with NCHC, and enhancing the advanced computing ecosystem in Taiwan

## User-Oriented Services

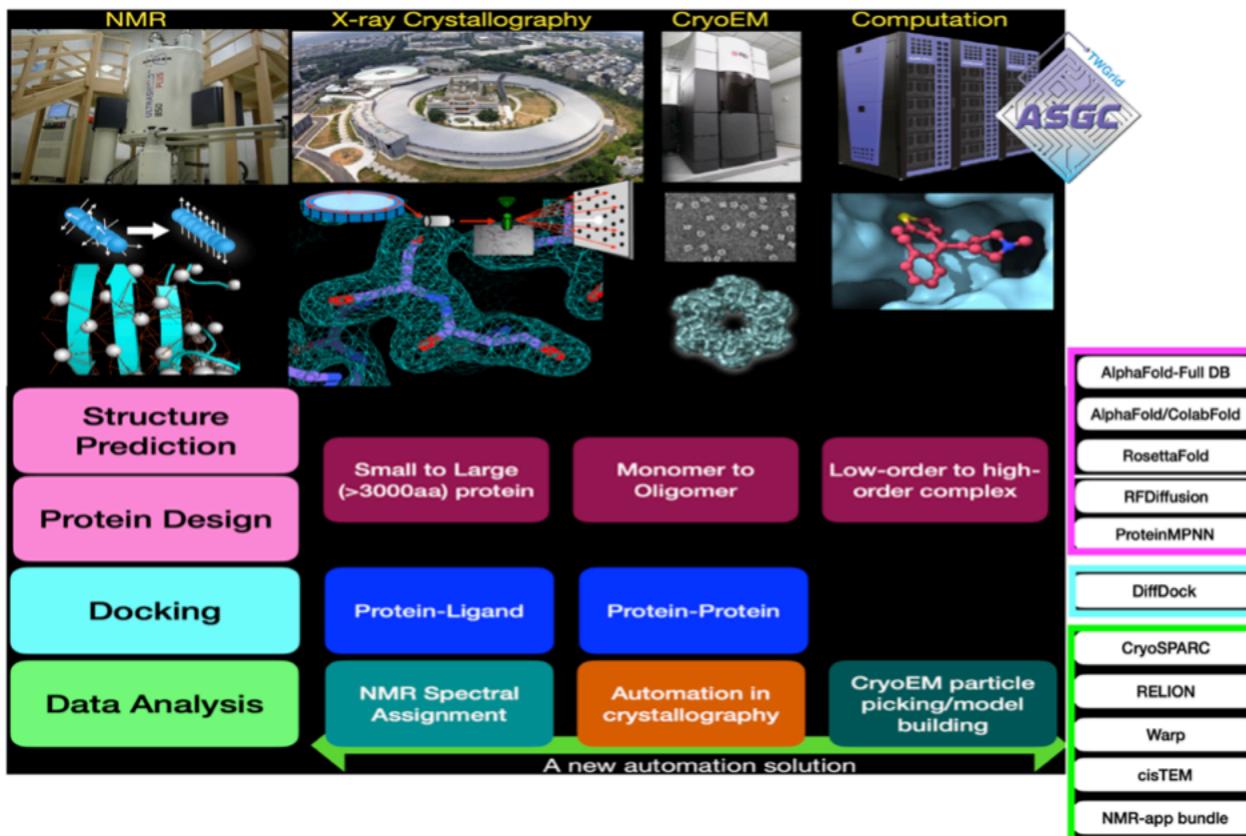
24x7x365



# Research Infrastructure Keeps Advancing With Requirements of Multiple Disciplinary Scientific Communities

From 2006, we have been supporting user communities in multiple disciplines to make better use of available scientific computing resources with efficiency.

Services are jointly provisioned by ASGC Core Facility (SCALE, 大數據分析與科學計算核心設施)



# Computing & Storage Resources

CPU Cluster	Spec	CPU Cores	GB RAM/Core	Nodes	Performance (TFLOP/Core)	Network	Onboard Date
EDR	AMD Genoa 9645	1,920	8	10		100Gb/HDR	Feb 2026
EDR1	AMD Genoa 9654	3,840	8	20	1.1	100Gb	Mar 2024, Dec 2024
Intel-G4	Intel Xeon 6448H	512	8	4	1	100Gb	Sep 2024

GPU Cluster	Spec	GPU	Architecture	Nodes	CPU Cores	Server RAM (GB)	CPU	Network	Onboard Date
B6000	RTX 6000 96GB	16	Blackwell	2	32	1,536	Intel 6517Px2	200Gb	Jan 2026
L40S	L40S 48GB	4	Ada Lovelace	1	32	768	AMD 9374x1	10Gb	Dec 2024
4090	RTX 4090 24GB	16	Ada Lovelace	2	40	1536	Intel 4416x2	10Gb	Jul 2024
3090	RTX 3090 24GB	56	Ampere	7	32	1024	AMD 7302x2	100Gb	Dec 2020
A100	A100 NVL 80GB	24	Ampere	3	64	512	AMD 7542x2	100Gb, NVL	Jun 2020
V100	Tesla V100 32GB	48	Volta	6	24	768	Intel 6126x2	10Gb, NVL	Dec 2018
P100	Tesla P100 16GB	16	Pascal	4	24	128	Intel 2650x2	10Gb, IB-FDR	Dec 2017

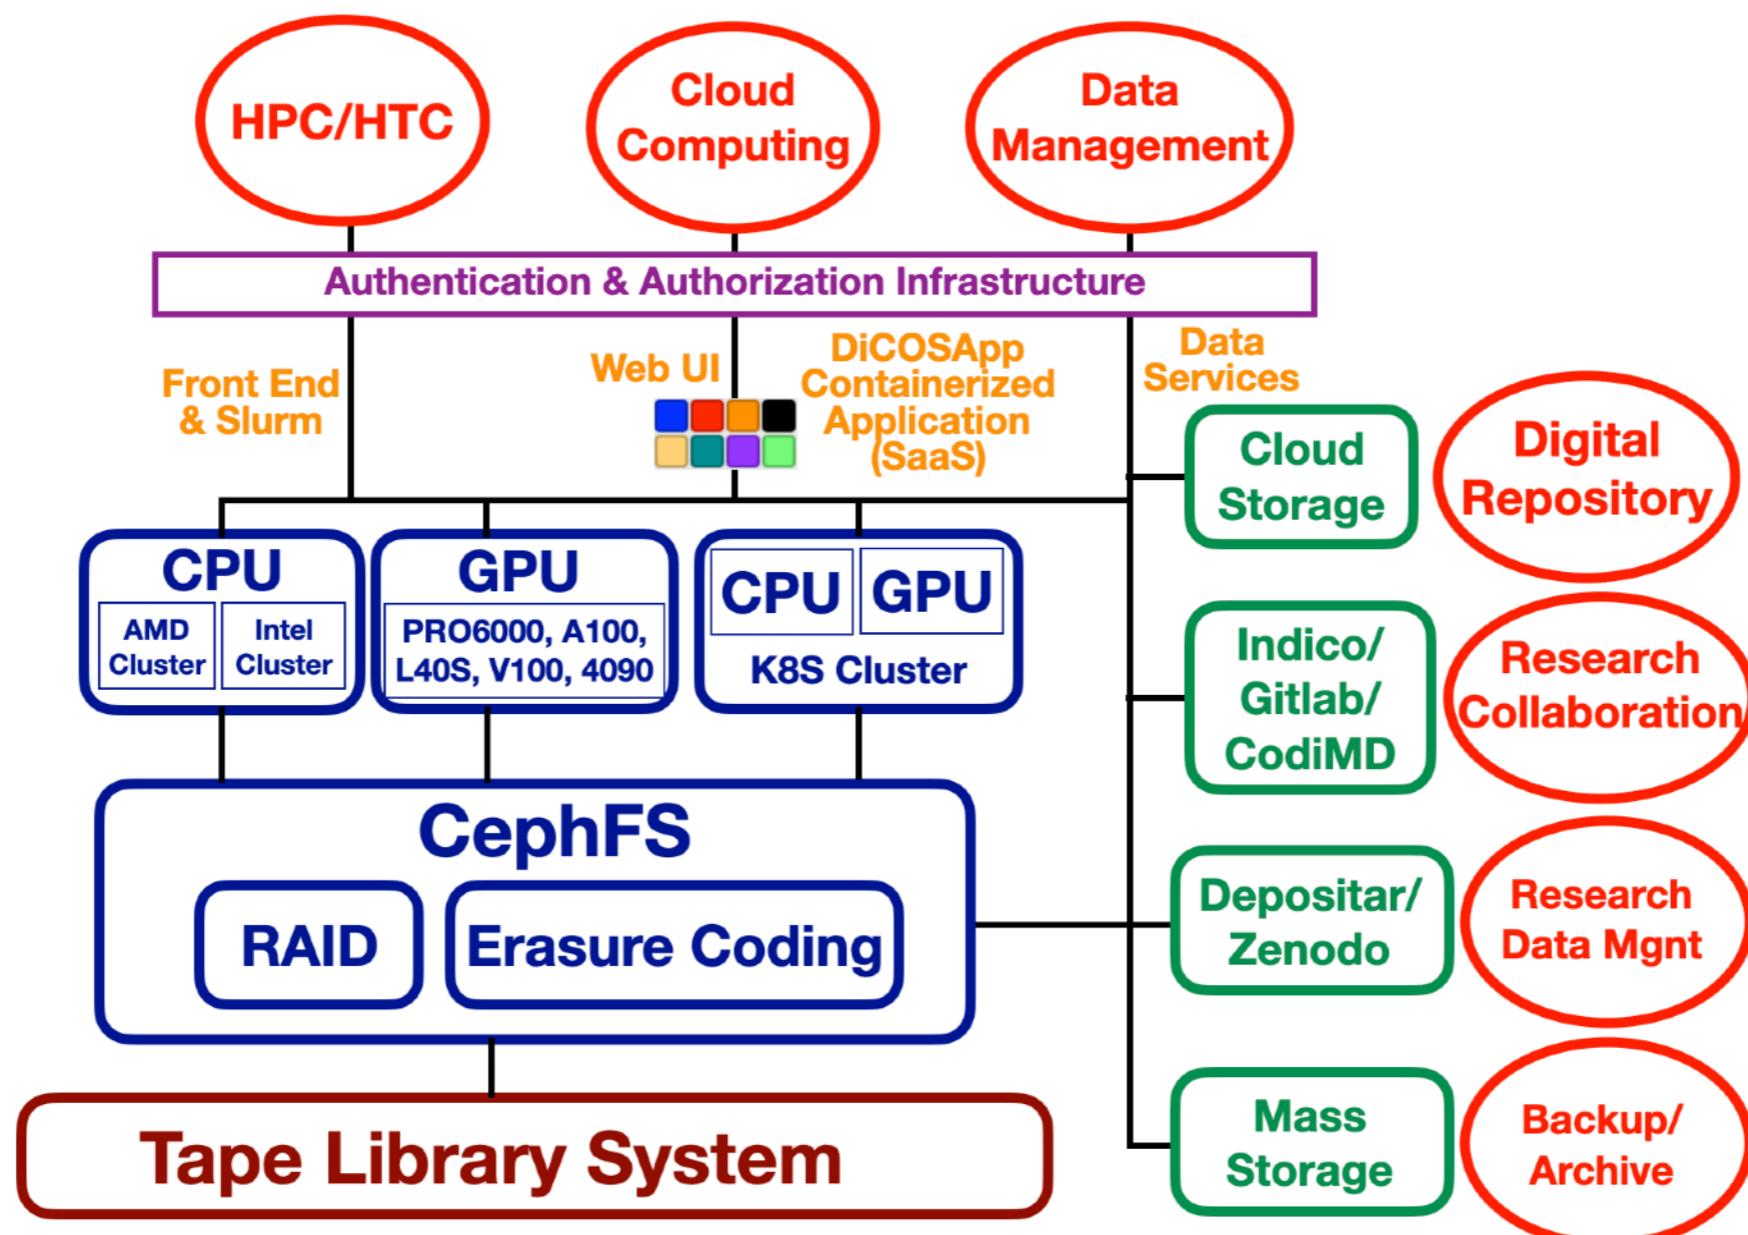
Storage	Capacity (TB)	System	Price	Remarks
Ceph File System	12,000	Online storage managed by Ceph distributed file system, with (8, 3) erasure coding	1,000 per TB-year	+3PB by end of Feb 2026
Tape Lib System	10,000	LTO9 (18TB/Tape)	300 per TB-year	Service from 2025

# Scientist Could Focus on Scientific Questions

- Utilizing high-performance computing service by affordable cost 筆電價格即可使用整年高規格計算資源
  - 100CPUCore-Yr < NTD\$51,100
  - 1TB-Yr Disk = NTD\$1,000
  - 1xA100 (80GB RAM) GPU-Yr = NTD\$ 43,800
- Free quota for new every new user: NT\$ 200
  - 新使用者可有 NT\$ 200 免付費測試使用額度
- 若有特殊或急迫需求，請逕洽本團隊
- Storage System :
  - Free: 3TB shared space/group, 100GB/ user
  - Tape Storage: NT\$300/TB-Year
- Welcomed to share your requirements and questions

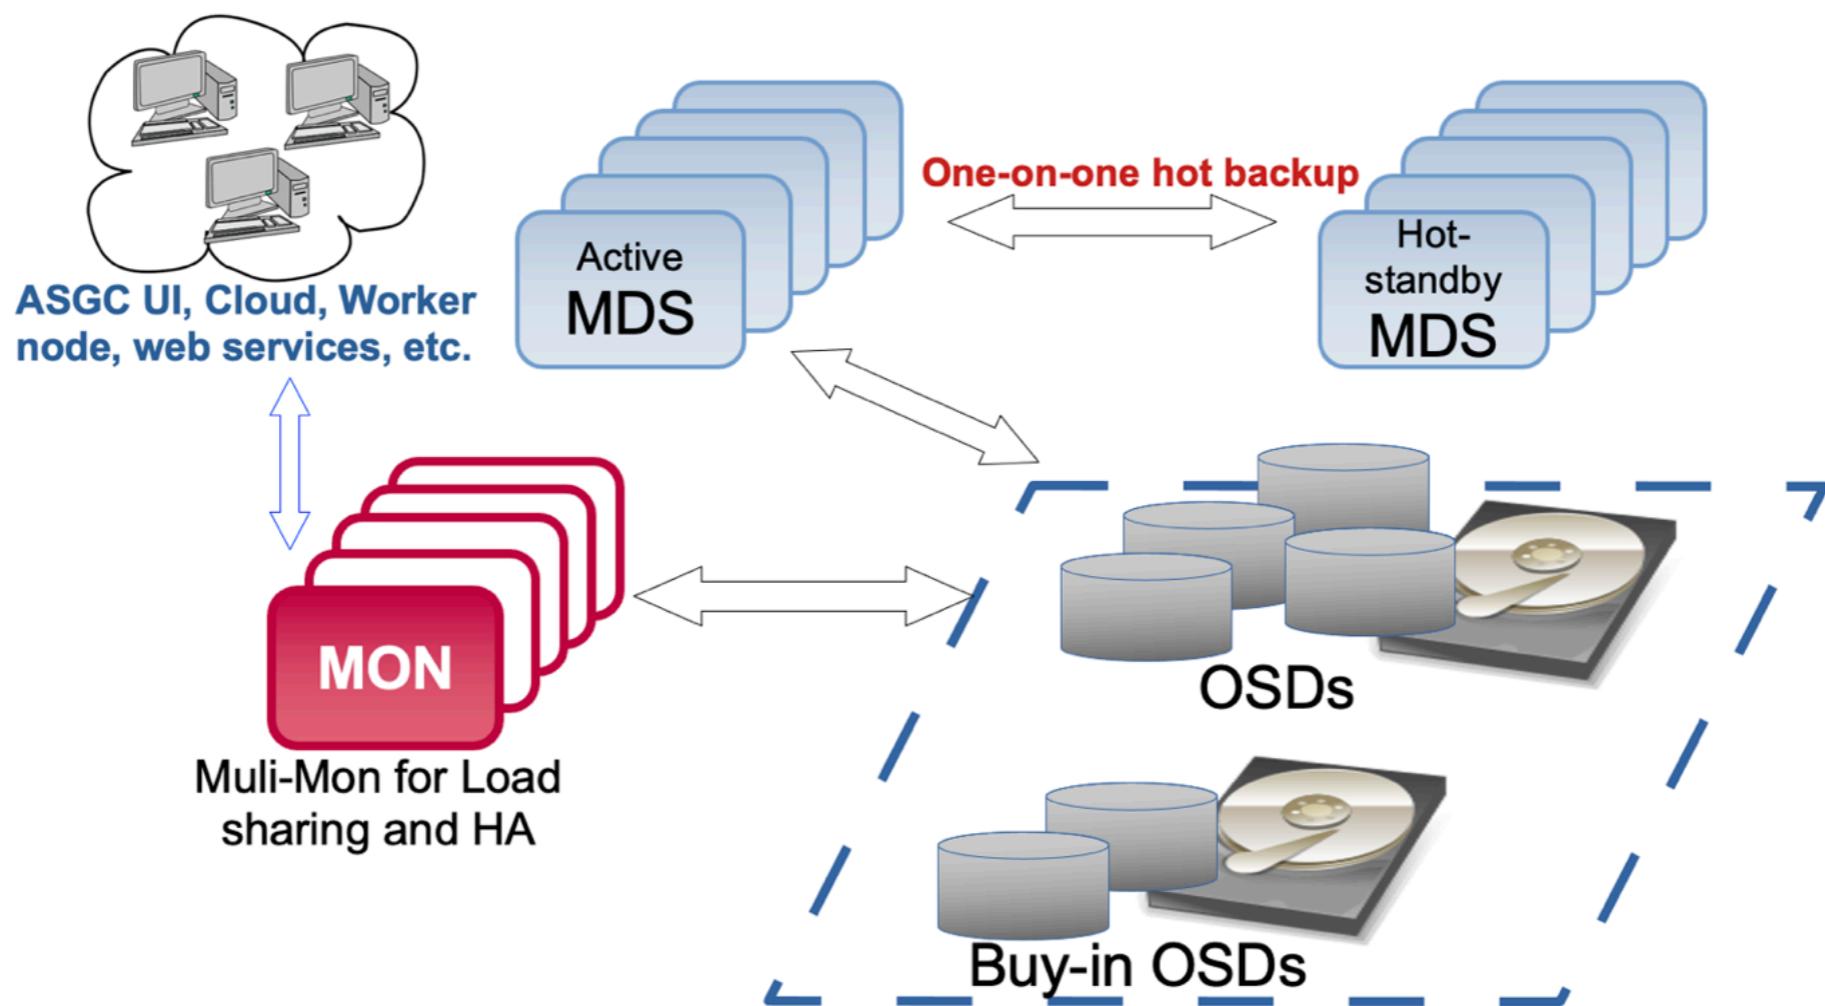
# ASGC Service Architecture - Scalable Science Cloud

- Integrating data, algorithm and computing
- Scientific workflow implementation with integrated compute, storage and network resources
- Persistent services: Analysis facility, Web portals
- New resource and new services online yearly, based on users requirements



# Data & Storage Services - CephFS

- Ceph distributed storage system -  $O(10)$ PB scalable data pool
- High performance - 4.2GB/s throughput, 1.5K op/s rd, 2K op/s wr
- Reliability - Erasure coding (8, 3), with 462 OSDs, 51 hosts and multiple MDS & MONs
- Scalability - from  $O(10)$ TB in 2012 to 17PB in 2025
- Data transmission and local backup supported
- New Services in 2026: data transfer, cloud storage and tape backup

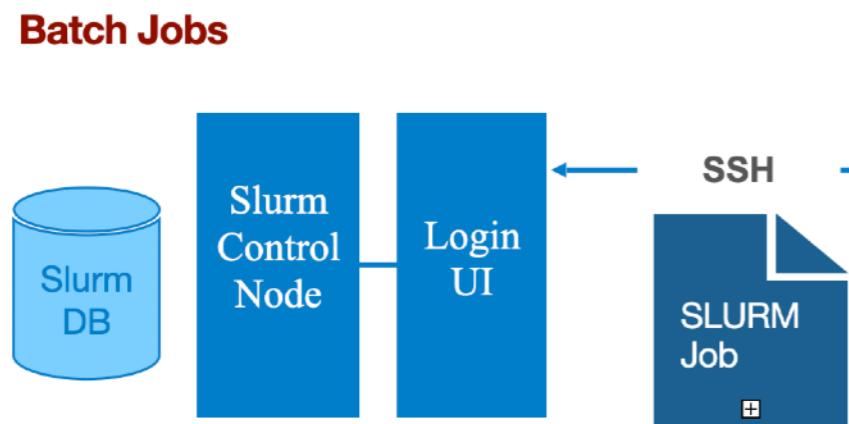


# Computing Services: Optimization of Application and System Efficiency

- **HTC/HPC & Cloud Computing Platform**
- **Integration of computing model and software**
- **Research data workflow and management**
- **Information security**
- **Application efficiency**
- **Customized IaaS, Paas, and SaaS**

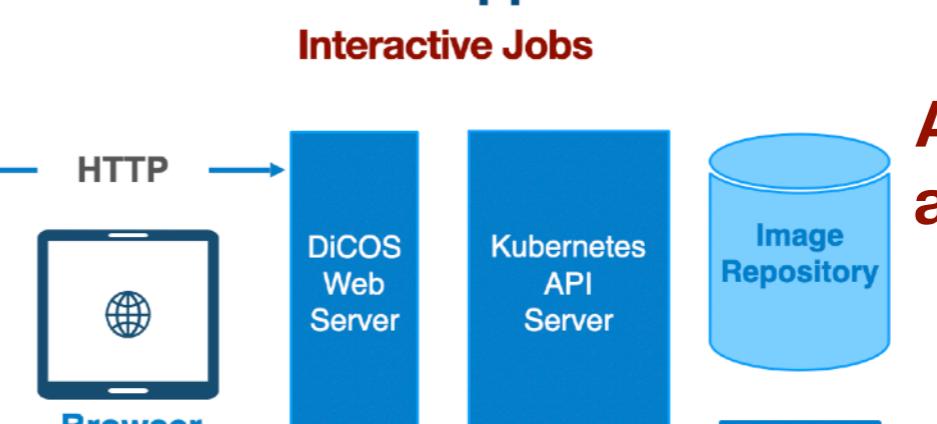
## HPC platform Slurm System Architecture

### Batch Jobs



## Software-as-a-Service Web-based Cloud Platform DiCOSApp SaaS Architecture

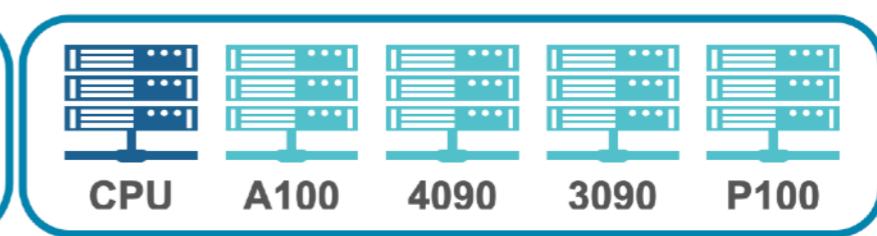
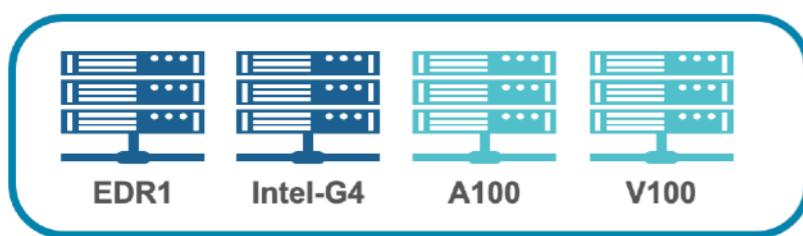
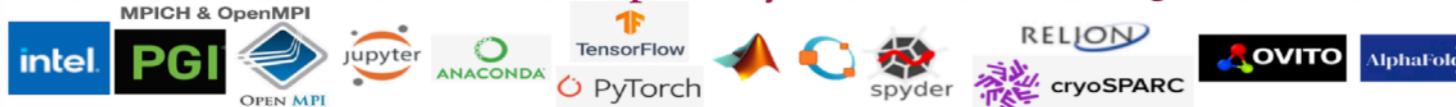
### Interactive Jobs



70+

**Applications  
are available**

### Scientific Software Repository & Customized Image Pool



# 73+ Web Applications Provided (Software-As-a-Service)

 <b>CryoSPARC 32 P100</b> Version: 3.2 Resources: 83%  <a href="#">Launch</a>	 <b>CryoSPARC 1080ti</b> Version: 3.3.2 Resources: 29%  <a href="#">Launch</a>	 <b>CryoSPARC RTX3090</b> Version: 3.3.2 Resources: 26%  <a href="#">Launch</a>	 <b>spyder cpu/eman2</b> Version: Resources: 17%  <a href="#">Launch</a>	 <b>Octave</b> Version: V5.2 Resources: 29%  <a href="#">Launch</a>	 <b>Transfer Data</b> Version: Resources: 86%  <a href="#">Launch</a>
 <b>CryoSPARC RTX3090</b> Version: 4.0.2 Resources: 26%  <a href="#">Launch</a>	 <b>CryoSPARC RTX3090</b> Version: 4.1.1 Resources: 26%  <a href="#">Launch</a>	 <b>CryoSPARC RTX3090</b> Version: 4.4.1 Resources: 26%  <a href="#">Launch</a>	 <b>cisTEM</b> Version: Resources: 86%  <a href="#">Launch</a>	 <b>OVITO</b> Version: Resources: 86%  <a href="#">Launch</a>	 <b>OpenACC</b> Version: GPU P100 Resources: 83%  <a href="#">Launch</a>
 <b>AlphaFold</b> Version: GPU with A100 Resources: 75%  <a href="#">Launch</a>	 <b>AlphaFold (Full DB)</b> Version: GPU with A100 Resources: 75%  <a href="#">Launch</a>	 <b>IMOD</b> Version: GPU with 1080ti Resources: 29%  <a href="#">Launch</a>	 <b>Triton</b> Version: 22.01-py3 (GPU P100) Resources: 83%  <a href="#">Launch</a>	 <b>AlphaFold</b> Version: GPU with RTX3090 Resources: 26%  <a href="#">Launch</a>	 <b>AlphaFold</b> Version: GPU with V100 Resources: 17%  <a href="#">Launch</a>
 <b>RoseTTAFold</b> Version: GPU with rx3090 Resources: 26%  <a href="#">Launch</a>	 <b>Dynamo</b> Version: GPU with 1080ti Resources: 29%  <a href="#">Launch</a>	 <b>MATLAB</b> Version:R2018b on GPU V00 Resources: 17%  <a href="#">Launch</a>	 <b>Jupyter Lab</b> Version: CPU with Tensorflow v1 Resources: 86%  <a href="#">Launch</a>	 <b>Jupyter Lab gpu 3090</b> Version: GPU with Tensorflow 3090 Resources: 26%  <a href="#">Launch</a>	 <b>Jupyter Lab GPU 1080ti</b> Version: GPU with Tensorflow v2 Resources: 29%  <a href="#">Launch</a>
 <b>RFDIFFUSION</b> Version: 2023 on GPU V00 Resources: 17%  <a href="#">Launch</a>	 <b>diffdock</b> Version: 2023 on GPU V00 Resources: 17%  <a href="#">Launch</a>	 <b>EvoDiff</b> Version: V100 Resources: 17%  <a href="#">Launch</a>	 <b>Jupyter Lab GPU V100</b> Version: GPU with Tensorflow V100 Resources: 17%  <a href="#">Launch</a>	 <b>Jupyter Lab GPU A100</b> Version: GPU with Tensorflow A100 Resources: 75%  <a href="#">Launch</a>	 <b>Jupyter Lab Cryocare GPU</b> Version: GPU with 1080ti Resources: 29%  <a href="#">Launch</a>
 <b>QIIME2</b> Version: Genome Resources: %  <a href="#">Launch</a>	 <b>Scipion3</b> Version: P100 Resources: 83%  <a href="#">Launch</a>	 <b>Phenix</b> Version: Resources: 86%  <a href="#">Launch</a>	 <b>Jupyter Lab GPU A100</b> Version: GPU with Tensorflow v2.6 Resources: 75%  <a href="#">Launch</a>	 <b>LabVIEW Run-Time Engine</b> Version: 2019  <a href="#">Launch</a>	 <b>DiCOS-BioSAXS Platform</b>
 <b>MorphoGraphX</b> Version: GPU with P100 Resources: 83%  <a href="#">Launch</a>	 <b>Deepmd-kit</b> Version: GPU with A100 Resources: 75%  <a href="#">Launch</a>	 <b>Deepmd-kit</b> Version: GPU with V100 Resources: 17%  <a href="#">Launch</a>	 <b>MAML</b> Version: GPU with A100 Resources: 75%  <a href="#">Launch</a>	 <b>Paraview Client</b> Version: 5.8.0 Resources: 86%  <a href="#">Launch</a>	 <b>ATSAS</b> Version:   <b>MAMM</b> Version:   <b>Qiskit</b> Version:   <b>NVIDIA CUDA-Quantum A100</b> Version: Resources: 100%  <a href="#">Launch</a>
 <b>Warp</b> Version: rtx4090 Resources: 700%  <a href="#">Launch</a>	 <b>MAML</b> Version: GPU with V100 Resources: 17%  <a href="#">Launch</a>	 <b>PVserver</b> Version: 5.8.0 (GPU 1080Ti) Resources: 29%  <a href="#">Launch</a>	 <b>Paraview Client</b> Version: 5.8.0 Resources: 86%  <a href="#">Launch</a>	 <b>LabVIEW Run-Time Engine</b> Version: 2019  <a href="#">Launch</a>	 <b>ATSAS</b> Version:   <b>AMBER</b> Version:   <b>DAMMIF</b> Version:   <b>Rosetta</b> Version:   <b>GASBOR</b> Version:   <b>NVIDIA CUDA-Quantum A100</b> Version: Resources: 100%  <a href="#">Launch</a>
 <b>R studio</b> Version: 10 CPU Cores Resources: 54%  <a href="#">Launch</a>	 <b>PyRoot</b> Version: GPU with 1080ti Resources: 29%  <a href="#">Launch</a>	 <b>Qiskit</b> Version: Resources: 86%  <a href="#">Launch</a>	 <b>NVIDIA CUDA-Quantum A100</b> Version: Resources: 100%  <a href="#">Launch</a>	 <b>Job Submission</b> Version:   <b>Job Monitoring</b> Version:   <b>Data Management</b> Version:   <b>About</b> Version:   <b>ATSAS</b> Version:   <b>AMBER</b> Version:   <b>DAMMIF</b> Version:   <b>Rosetta</b> Version:   <b>GASBOR</b> Version:   <b>NVIDIA CUDA-Quantum A100</b> Version: Resources: 100%  <a href="#">Launch</a>	

- [Web Portal](#)
- [Application over Cloud](#)
- [Jupyterlab](#)
- [Web Terminal](#)



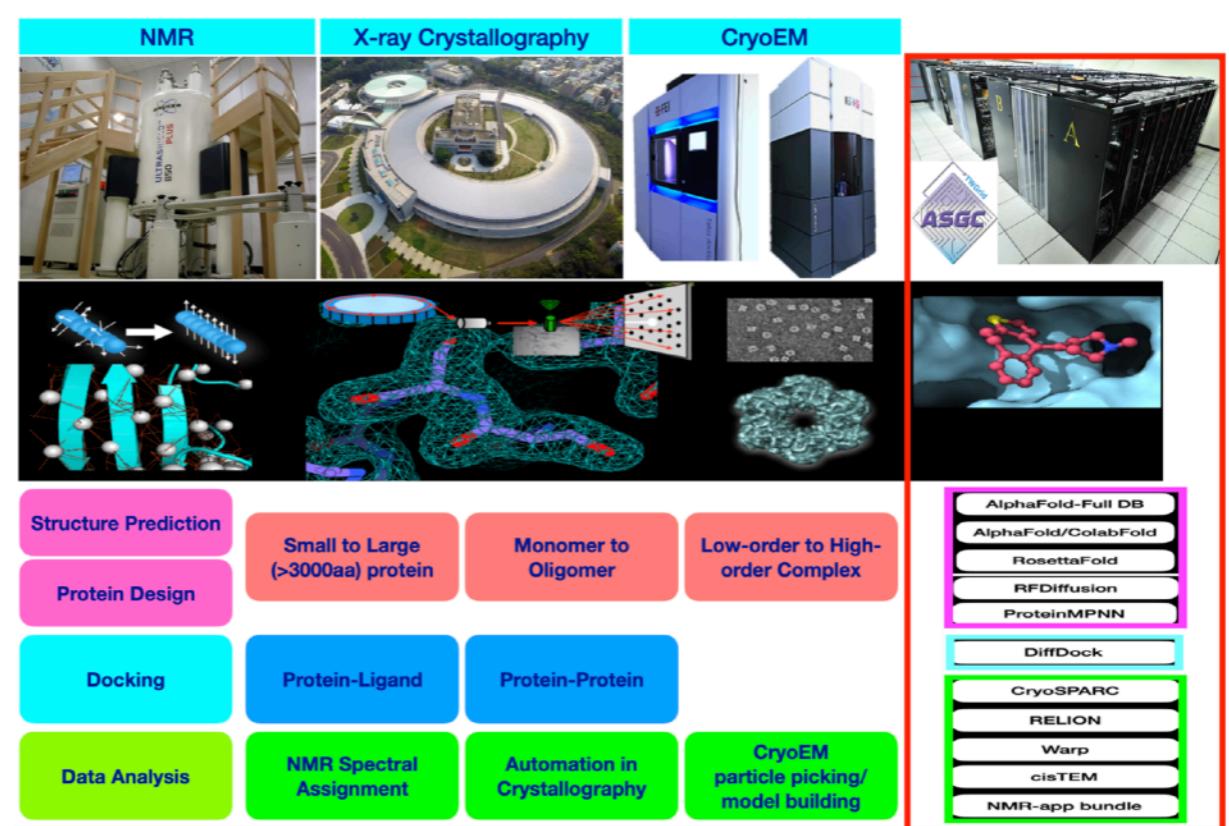
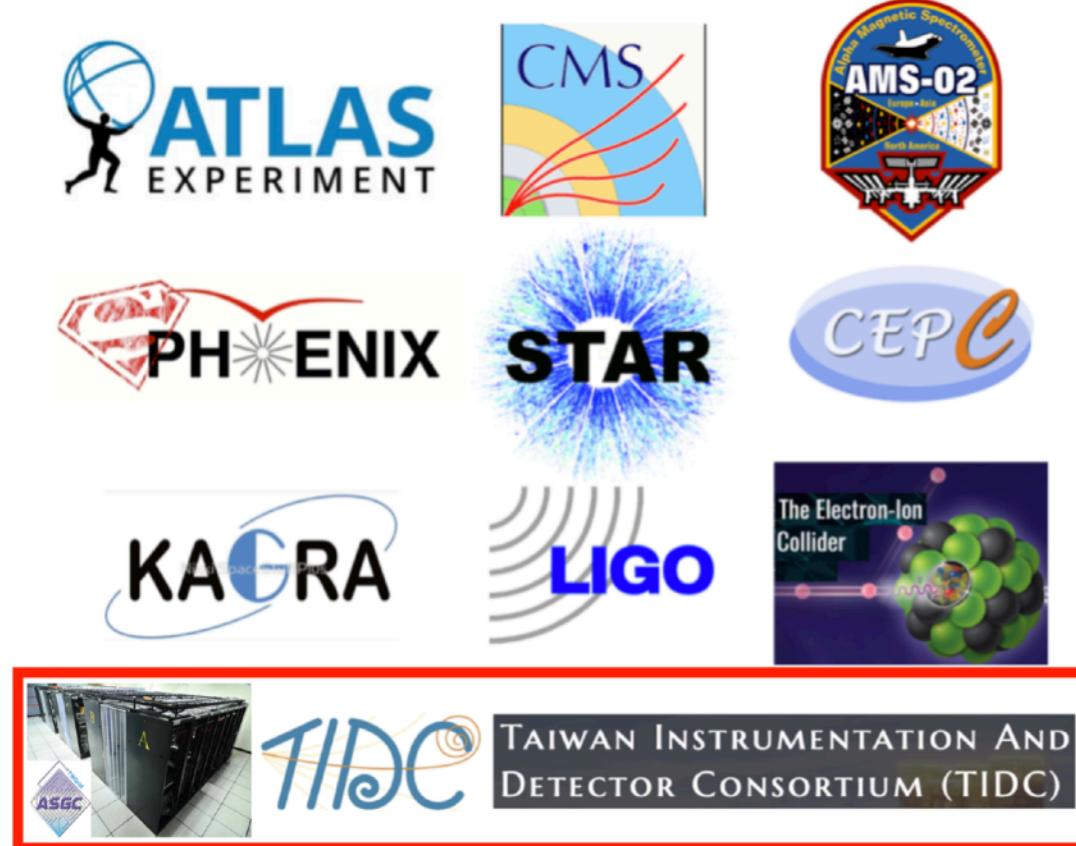
DiCOS-BioSAXS Platform

Job Submission

- ATSAS
- AMBER
- DAMMIF
- GASBOR

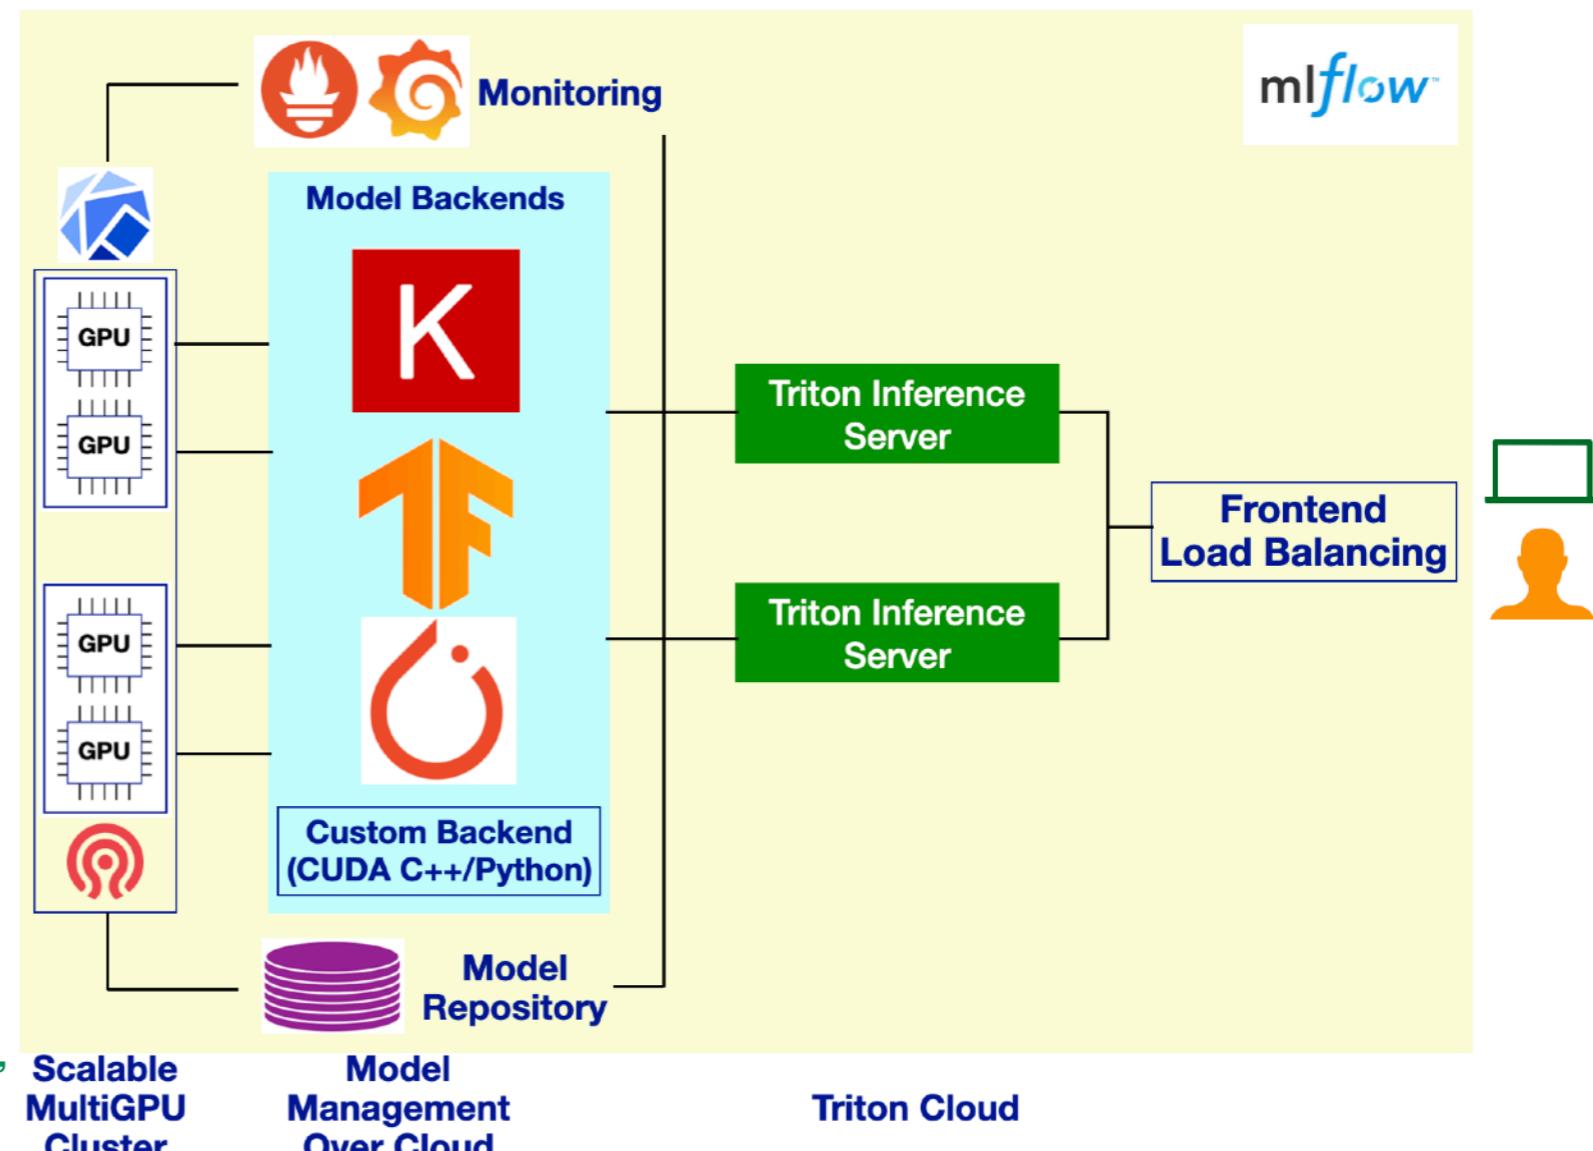
# Flexible Collaboration Model: User-Driven, Service Oriented

- Pay-as-you-go
- Resource Integration: Collocation & Buy-In
- Service Hosting
- Reservation
- Experiment & Collaboration support
- Advanced Service: Web Application and portal, efficiency tuning, workflow optimization, customization, hands-on training
- User Engagement : weekly User Meeting, User Committee, dissemination & outreach

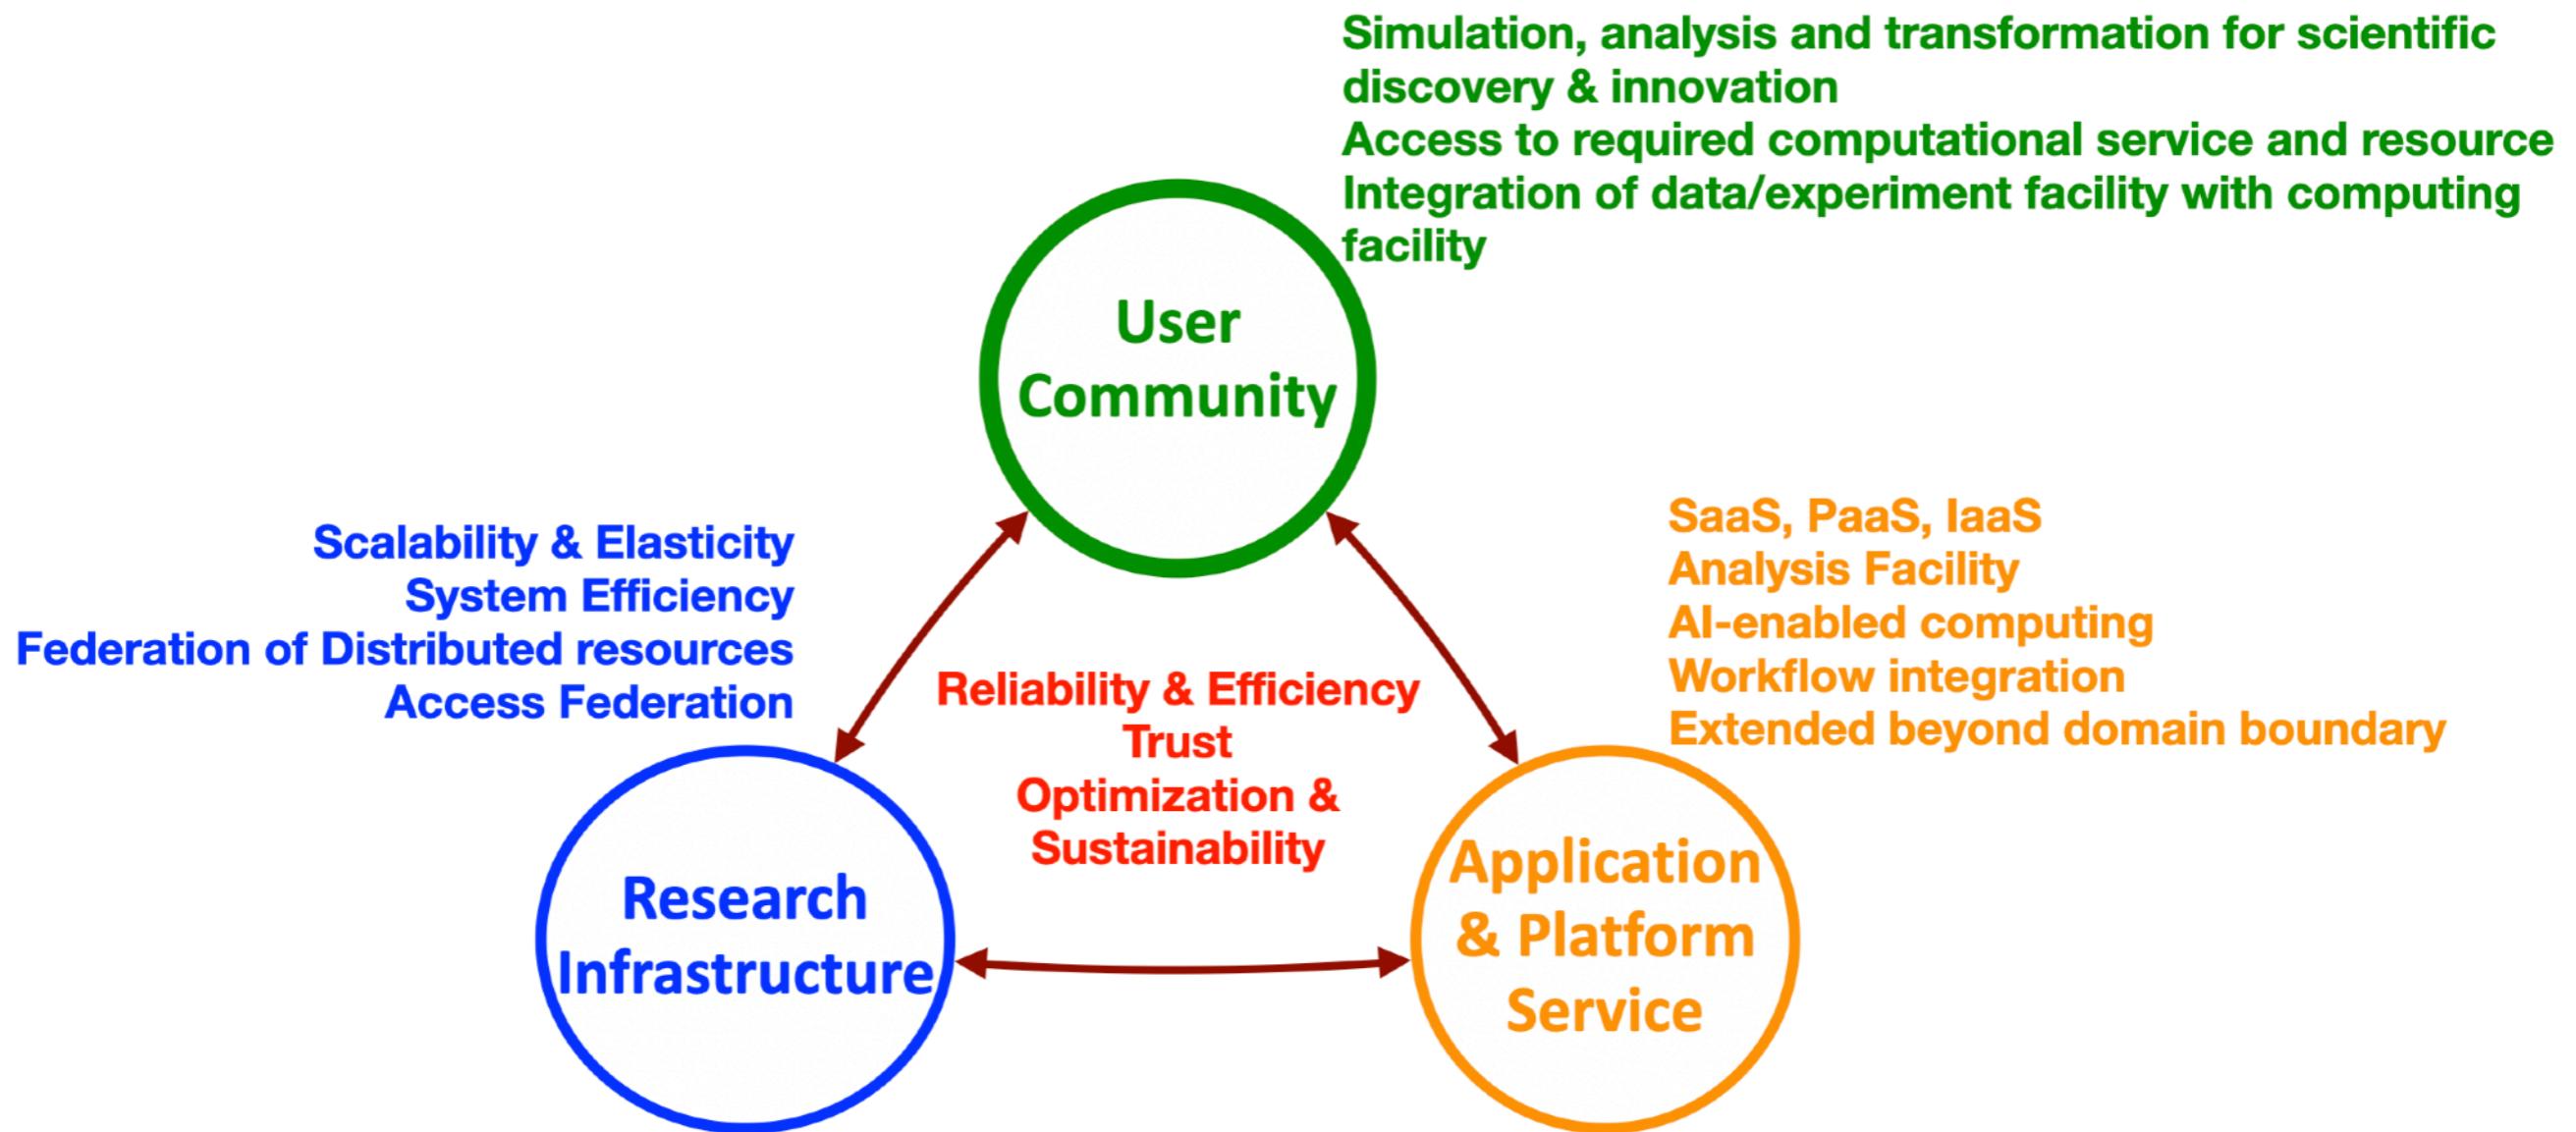


# Services of ML-Enabled Data Analytics

- ML/AI application platform service - SW library, HW, integration and application development management (e.g, MLflow, Triton)
  - Build up customized ML platforms for user specified projects - Deploy ML packages ready environment in order to help ML development smoothly and provide on-demand computing power
  - Upkeep of the application framework
  - Workflow and data pipeline integration: tracking, model serving, packaging, model registry, evaluation & tracing
  - Efficiency Improvement
  - Inference-as-a-Service (incl. Model repository): e.g., Triton-based services
- Potential use cases
  - Users who bring existing source code - ASGC could help to setup a virtual environment and confirm source code running normally
  - Share and reuse of trained/verified models
  - MLOps over the cloud
- Approaches
  - Supporting Kubernetes/Jupyter lab for development purpose
    - Create Kubernetes/Jupyter lab environment with user specified ML packages ready.
    - Support on-demand scalable CPU/GPU computing power.
  - Supporting containerized environment (e.g, Docker image) for deployment purpose
    - Create takeout images in Docker format as an option for user who wants to train/predict model
    - Docker images could be downloaded from ASGC server and deployed on users' Docker Desktop on Windows/Linux.



# Capacity Building



# Enhancing User Community Engagement

- Discipline-oriented or user community-oriented collaboration is essential
- Training, dissemination and outreach
- Improving communication and response
- Supporting user access and application efficiency
- Understanding user requirements and experiences
- Facilitating capacity building of users and NSTCCore
- Training and workshop/conference: 250+ participants in 2025
- Weekly user meeting: 2:30pm, Wednesday

## Basic

- 30min – full day
- Get access, details of using compute, data and SaaS services, w/o hands-on, User Forum

## Thematic

- 1 hrs – full day
- Discipline-based Data analysis
- ICT topics: GPU Computing, AI, CPU Trends

## Domain-Specific Workshop

- 30min – 3hrs
- Supporting hands-on, w/o ASGC service introduction

## Conference

- e-Science, community-specific workshop, core technology
- ISGC (International Symposium on Grids & Clouds, since 2002)
- HEPiX, CHEP (Computing in High Energy Physics)
- APAN (Asia Pacific Advanced Network)

## Training, Dissemination, and Outreach



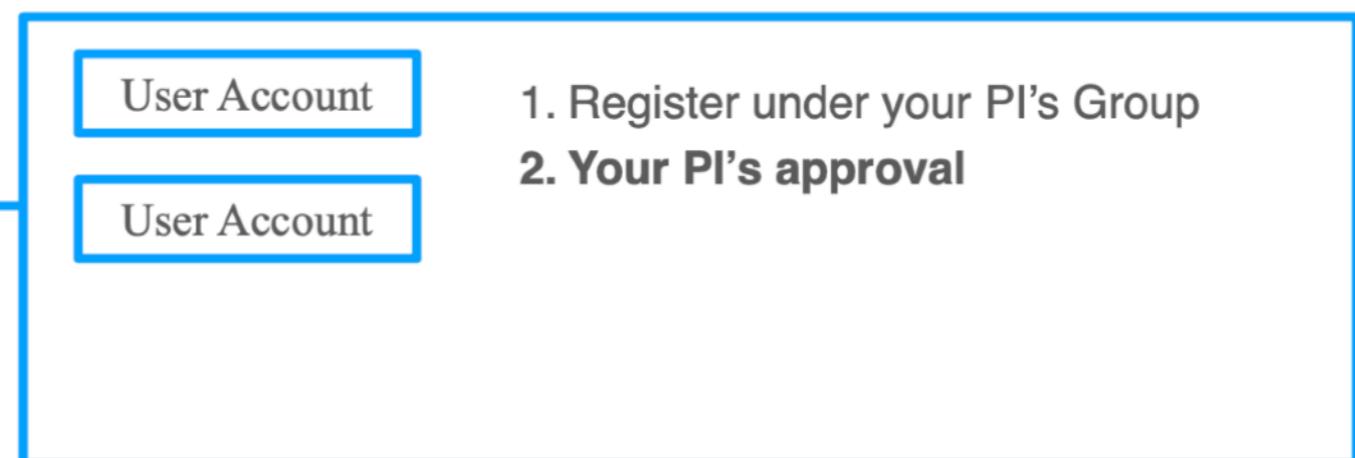
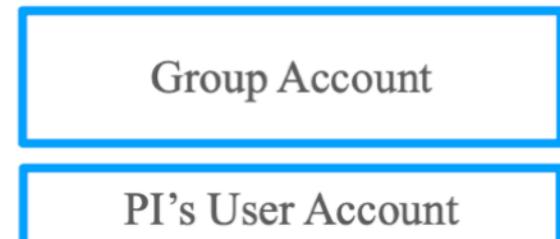
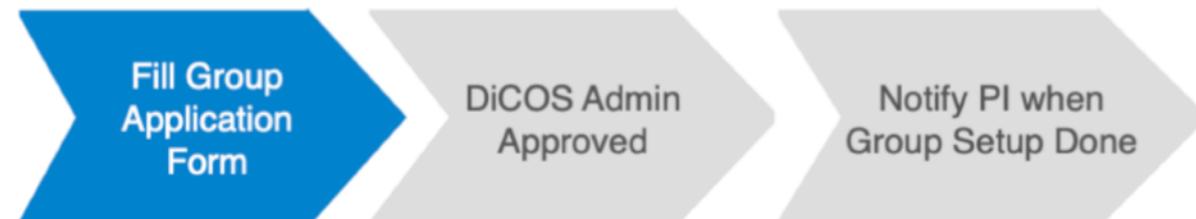
International Symposium on Grids & Clouds (ISGC, annually since 2002)



# Access to NSTCCore Services

- Two-Factor Authentication enabled from January 2025
- Free \$200 Credits for 1-month Trial
- Free 3TB shared disk storage per group, and 100GB per user

- Group Resource Usage and Payment Management
- Member accounts management support
  - Members list
  - Abnormal member usage report
  - Member's usage review



- Group Accounts
  - <https://canew.twgrid.org/ApplyAccount/groupcreate.php>

- User Accounts
  - <https://canew.twgrid.org/ApplyAccount/ApplyAccount.php>

# Group Management

- Members management
  - Members list
  - Abnormal member usage report
  - Member's usage review

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	33.4G/100G Latest Update: 2023-07-31 00:50:03
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, 09:13 AM	Jul. 06, 2023, 04:37 AM	0.0G/100G Latest Update: 2023-07-31 00:50:03
ericyen	Eric YEN	Eric.Yen@twgrid.org	Dec. 08, 2023, 00:00 AM	True	Apr. 23, 2013, 00:00 AM	Jul. 24, 2023, 03:22 AM	0.0G/100G Latest Update: 2023-07-31 00:50:03
etomo	etomo etomo	etomo@twgrid.org	Aug. 09, 2025, 00:00 AM	True	Dec. 23, 2020, 09:19 AM	Jan. 22, 2021, 10:16 AM	0.0G/100G Latest Update: 2023-07-31 00:50:03
felixlee	Felix Lee	felix@twgrid.org	Oct. 24, 2023, 00:00 AM	True	Mar. 05, 2013, 00:00 AM	Jul. 27, 2023, 07:29 AM	436.3G/100G Latest Update: 2023-07-31 00:50:03

使用統計摘要(TOTAL SUMMARY TABLE)		使用明細(DETAILED USAGE ACCOUNTING)									
使用明細(DETAILED USAGE ACCOUNTING)		用量異常回報(submit issue)									
使用者(username): chiong (CHAN-HIN IONG)		Alert									
Alert		#Instance	A100	FDR5		P100		RTX3090		V100	
			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				11	282		
<input type="checkbox"/>	openaccp100	1									
<input type="checkbox"/>	FDR5 slurm	2			0	0					
<input type="checkbox"/>	STORAGE USAGE (0.0 TB)										

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					• 使用統計摘要(TOTAL SUMMARY TABLE)				
• 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)	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
總計	單價SRU=NT\$2	183	7,756	(183+7,756)*2=15,878					
• 儲存資源使用統計(STORAGE USAGE OF GROUP) – 預估使用費(INITIAL COST): NT \$417									
帳號(ID)	名稱(name)	儲存空間用量(storage(TB))	預估費用(Initial Cost Estimation)	說明(note)	帳號(ID)	名稱(name)	儲存空間用量(storage(TB))	預估費用(Initial Cost Estimation)	說明(note)
ASGC	Group	0.0	0.0	Group使用空間	chiong	CHAN-HIN IONG	0.0	0.0	User使用空間
jyou	Jingya You	0.0	0.0	User使用空間	rudy	陳侑廷	0.0	0.0	User使用空間
thwu	Tsung-Hsun Wu	4.9	408.3	User使用空間					

# NCHC Collaboration

- **Collaboration Model**
  - **Workload distribution Principle: large-scale vs high-throughput**
  - **Quarterly meeting for milestones**
- **Starting from**
  - **Federated IdM - easy access to both sites using the same account (better to be user's home institute SSO account)**
  - **Home directory and data consistency**
  - **Faster access to required services and finish processing efficiently**
- **Challenge**
  - **Security governance - requirements and policy**
  - **Unique UID for each user**
  - **Data & storage consistency**
  - **Compatible Software and environment**
  - **Services**

# Addressing Challenges for Sustainability

- **Target: growing Users x Service x Resource, with reliability and efficiency**
- **Maximize the effectiveness of budget from AS and NSTC**
- **Short of advanced GPU resource: budget, scientific computing support are key issues**
- **Data Center reliability**
  - Power system upgrade
  - Energy efficiency
  - Upgrade for higher power consumption & liquid-cooling computing hardware in the near future
  - Deployment & R/D on Intelligent monitoring & control
- **High and increasing electricity cost**
- **Improving system efficiency**
  - Improving users' experience, utilization, and reducing the waiting time
- **HR & Capacity building: enhancing professional skills according to user needs and ICT progress**
  - Have to catch up/support evolving needs of scientific computing and AI applications
  - Hard to have backup HR for key technologies
  - User support needs more investment
  - Uncompetitive ecosystem
- **Pricing strategy**
  - **Baseline: covering power consumption (including data center operation)**
  - **Will rely mainly on data services and advanced services**
  - **Expect NSTC to approve 使用費留用**

# NSTCCore Services 聯絡資訊

- **Core Facility Services**
  - <https://nstccore.twgrid.org>
  - <https://scale.grid.sinica.edu.tw/index.php>
- **ASGC Web Site:** <https://www.twgrid.org>
- **Access to ASGC Resources**
  - <https://dicos.grid.sinica.edu.tw/>
- **Contact point:** [DiCOS-Support@twgrid.org](mailto:DiCOS-Support@twgrid.org)
- **Gentle Reminder**
  - Please include ASGC (Academia Sinica Grid-computing Center) in the acknowledgement when research outcomes that relied on ASGC resources, services or expertise are presented in your research.

# Unit Cost of Electricity Is 15.8% Higher in 2025

- Electricity unit cost in 2025 (monthly average) is 4.01/kWh, which is 48.2% higher comparing to 2022
- Both facility and cooling power usage decrease by > 20% in 2023 and 2024
- The power usage in 2025 keeps having 2.7% drop
- PUE (2025) is 1.58, which maintains the same level as in 2024 (1.56)

