Core Facility for Networking, Science, Technology and Advanced Computing (NSTCCore)

PI: Yuan-Hann Chang

co-PI: Eric Yen

On behalf of ASGC

Institute of Physics, Academia Sinica

NSTCCore Introduction

- Objectives: establish a medium-sized high-performance computing (HPC)/ high-throughput computing (HTC) facility for the research & education community in Taiwan
- Project schedule: 1 June 2023 - 31 May 2026
- Target Resource Scale
  - Procure ~1,800 CPUcores, 3PB Disk storage, 4PB tape storage every year
  - First year hardware is planned to be online by end of 2023
  - Existing ASGC resource & services (2208 + 768 CPUCores, 3TB group storage) are available before the readiness of new hardware
- Reliability and efficiency are the primary goals
- NSTCCore receives funding from the National Science and Technical Council, grant no. NSTC 112-2740-M-001-003

Event Program

9:30 AM → 9:45 AM **開場 Opening**

Convener: Eric Yen

9:45 AM → 10:15 AM **一: 計算資源和使用 Computing Resources, Usage & Collaboration Model**

- 運算資源、服務內容和計價模式 (Resource, Service & Pricing)
- 技術服務、申請流程和合作模式

Convener: Eric Yen

10:15 AM → 10:45 AM **二: 高效能科學運算服務 HPC & Scientific Computing Service**

- 科學運算服務平台(Computing Service Platform)：SLURM & DiCOSApp
- 資料服務(Data Service)：資料備份及傳輸(Data Backup & Transfer)
- 軟體服務需求

Convener: Jingya You

10:45 AM → 11:15 AM 休息 Break

11:15 AM → 11:45 AM **三: 使用者需求和經驗討論 Discussion**

Convener: Eric Yen

11:45 AM → 12:15 PM **四、實作訓練: 虛擬化計算服務(雲端模式)**

- 環境部署
- 軟體需求

Convener: Zong-Tsung Wu

12:15 PM → 1:00 PM 午餐 Lunch

1:00 PM → 1:40 PM **四 實作訓練 Hands-on: SLURM**

- SLURM執行工作操作
- SLURM參數介紹

Convener: Rudy Chen

1:40 PM → 2:00 PM **四、實作訓練: 資料操作和儲存服務 (Data Access & Storage Service)**

- 資料操作及工作儲存空間 (Job working space)
- 資料傳輸 (Data transfer)
- 資料備份 (Data backup)

Convener: Chiheng Ting

2:00 PM → 2:30 PM 休息 Break

2:30 PM → 3:00 PM **四、實作訓練: 進階計算操作 Advanced Computing Job**

- 多核心程式編程及操作 Multi-Core Jobs
- GPU程式編譯及操作 GPU programming

Convener: Chan-Hin Iong

3:00 PM → 3:30 PM **問與答 Q & A**

- Improving services from users requirements and experiences
- Plan to have 4 events a year
- For training, seminar on scientific computing and big data analysis, user experiences sharing, etc.
- Interaction between users and ASGC supporting team

ASGC Aims To Accelerate Discovery and Innovation

- ASGC joined WLCG development and deployment for the Large Hadron Collider grand challenges since 2001
  - ASGC T1 and WLCG Asian Regional Operation Centre has been operational from 2005
  - Migrating to T2s for ATLAS and CMS (effective from Oct. 2023)
- ASGC has been supporting multi-disciplinary e-Science applications of Academia Sinica from 2006, based on WLCG core technologies
  - The research infrastructure, platform and services are improved progressively along with growing scientific applications of various disciplines
- System efficiency optimization (including power, thermal, system and applications, etc.) is also a strategic goal of ASGC aided by machine learning technologies
- ASGC becomes the Core Facility for big data and scientific computing of AS from 2023

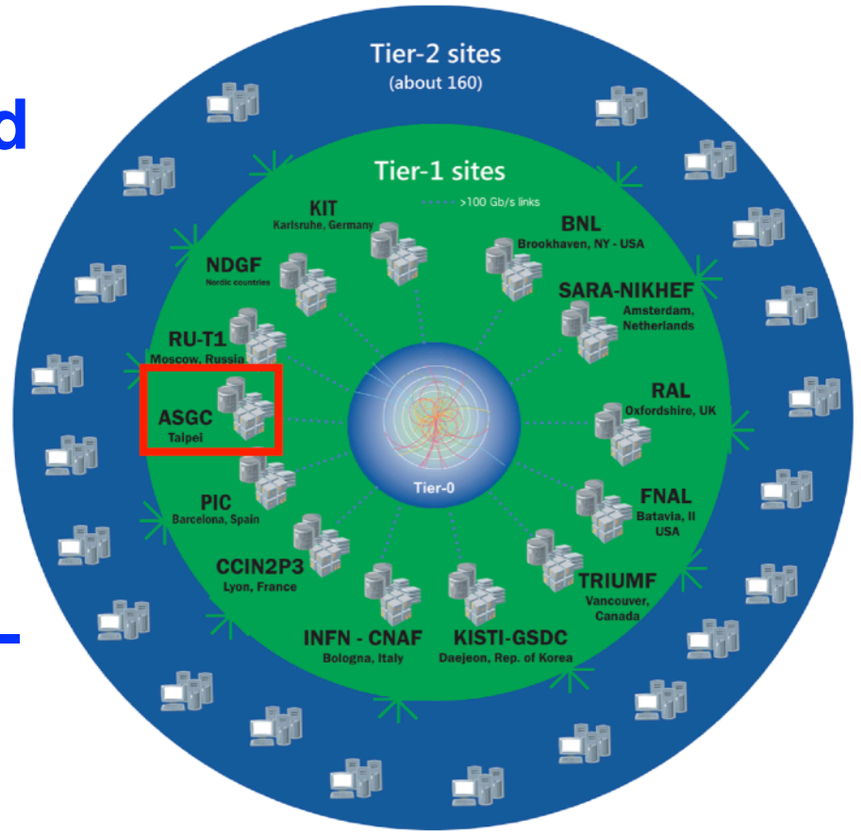

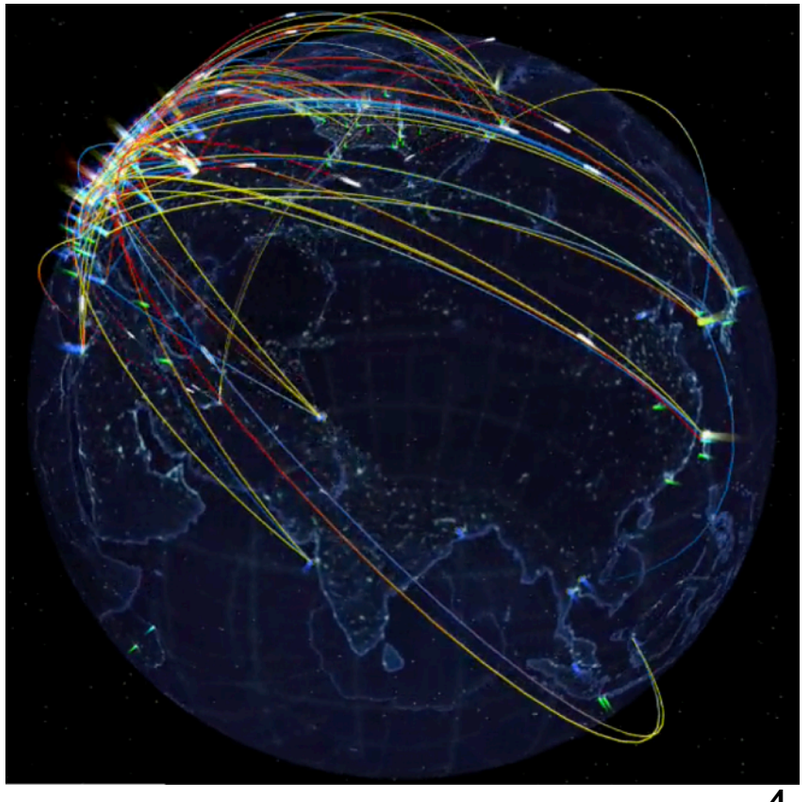

NSTCCore Service

- Accessibility /Account application
  - Group based
- Computing by Slurm and Cloud over CPU & GPU
  - Ease-of-use
  - Web-based UI and Jupyterlab
- Expandable Storage Service - typically used for frequently accessed data
  - CephFS supported - general purpose storage for computing and any type of data
  - 100GB/account; 3TB/group by default
  - User can pay for more storage space - by day or year, TB-basis
- Backup on Tape
  - Data in a specific directory will be backup onto tape automatically
  - Tape system should be available in Q4 2023
    - Benchmark on data migration will be published after the tape system is integrated
- DiCOSApp - including Jupyterlab environment
  - Web UI for packaged software environment/ tools
- Training, consulting and problem solving

50+ Web Applications Provided

- Web Portal
- Application over Cloud
- Jupyterlab
- Web Terminal

PHYS

Deepmd-kit
Version: GPU with A100
Resources: 12%
Launch

Deepmd-kit
Version: GPU with V100
Resources: 80%
Launch

MAML
Version: GPU with A100
Resources: 12%
Launch

MAML
Version: GPU with V100
Resources: 80%
Launch

PVserver
Version: 5.8.0 (GPU 1080Ti)
Resources: 66%
Launch

Paraview Client
Version: 5.8.0
Resources: 97%
Launch

PyRoot
Version: GPU with 1080ti
Resources: 66%
Launch

Jupyter

Jupyter Lab
Version: CPU with Tensorflow v1
Resources: 97%
Launch

Jupyter Lab gpu 3090
Version: GPU with Tensorflow 3090
Resources: 51%
Launch

Jupyter Lab GPU V100
Version: GPU with Tensorflow V100
Resources: 80%
Launch

Jupyter Lab GPU A100
Version: GPU with Tensorflow A100
Resources: 12%
Launch

LabVIEW Run-Time Engine
Version: 2019
Launch

DiCOS-BioSAXS Platform

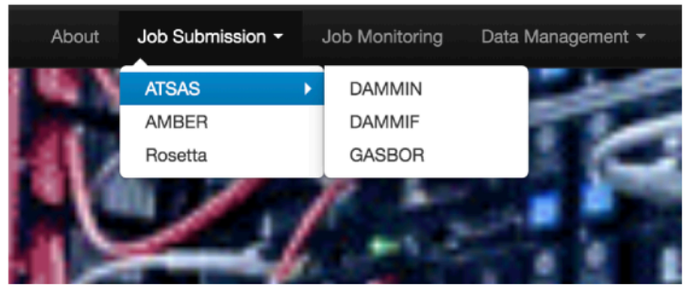

Other

spyder cpu/eman2
Version:
Resources: 97%
Launch

Octave
Version: V5.2
Resources: 66%
Launch

Transfer Data
Version:
Resources: 97%
Launch

cisTEM
Version:
Resources: 100.0%
Launch

Ovito
Version:
Resources: 97%
Launch

OpenACC
Version: GPU P100
Resources: 50%
Launch

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

AlphaFold
Version: GPU with V100
Resources: 80%
Launch

AlphaFold
Version: GPU with A100
Resources: 12%
Launch

IMOD
Version:
Resources: 66%
Launch

RoseTTAFold
Version:
Resources: 51%
Launch

Dynamo
Version:
Resources: 66%
Launch

Supporting Big Data & AI in Innovations

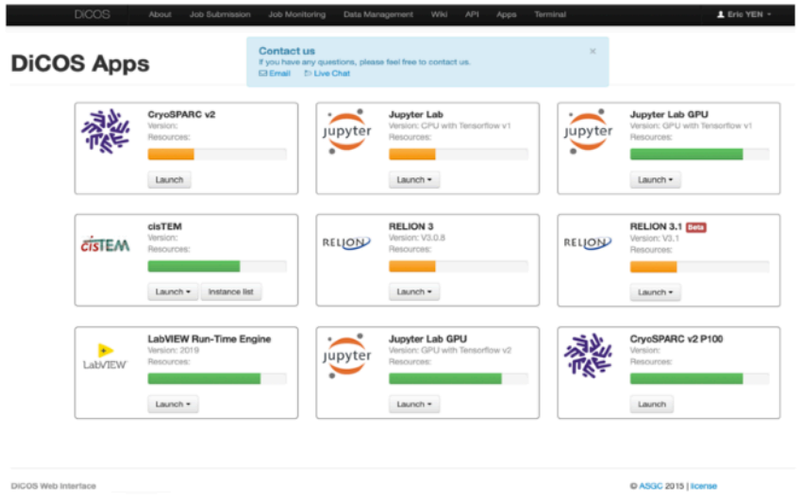

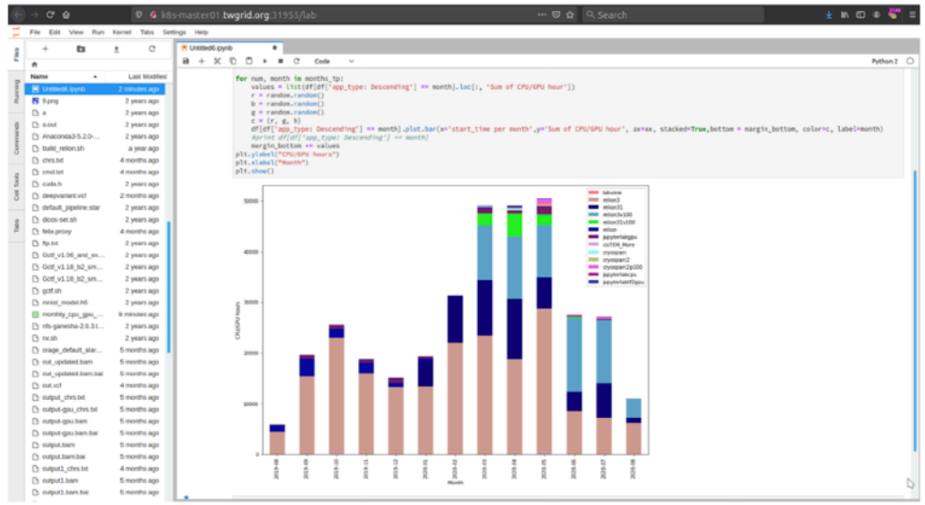

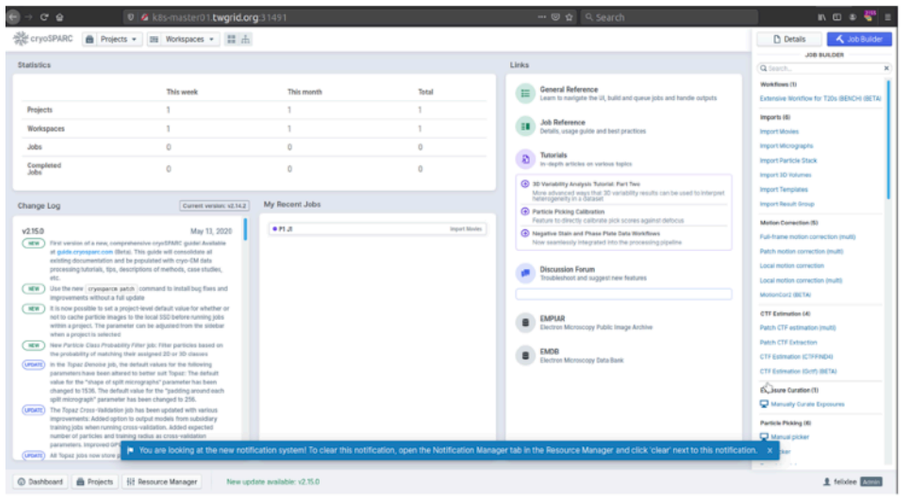

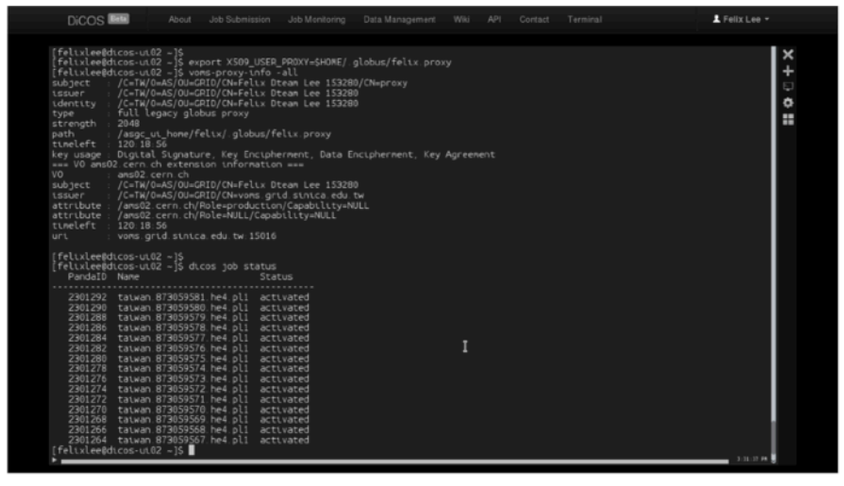

CLI | **Web Portal** | **DiCOS APP** | **Jupyter Notebook** | **Science Portal** | **Web Browser/ Terminal**

Application-specific/ Generic Learning Engines

Deep Learning Engines/Frameworks

Computing Resource (Cloud/Grid/Slurm)

Storage Resource (Ceph/EOS)

Distributed Data Management & Cloud Storage Services

Network & Data Transmission Services

Pricing Model

- **Pricing policy will be reviewed by User Committee meeting annually**
- **Payment scheme**
  - **Monthly accounting report will be delivered to PI before the 15th each month**
  - **Payment could be made by transfer, mailed check or pay in counter.**
  - **Detail account information and payment channels will be explained in the monthly report**

Research Service	Pricing
Computing Service	Based on normalized resource unit (SRU), NTD1/SRU AMD ROME(Core-Day) = 1.1 SRU; NVIDIA V-100(Board-Day) = 140 SRU; A-100(Board-Day) =346 SRU
Storage Service	Disk Storage System：1,500NTD/TB-Year Tape Storage System：300NTD/TB-Year
Data Transmission	Free of charge at this moment
Advanced Service	Charged by working hours (NTD 120,000 per 168 man-hr)

ASGC 資源使用統計月報表

ASGC Resource Usage Monthly Report

Reporting Period: 01 – 31 July, 2023 **Date of Report: 10 August, 2023**

PI Group: **(Group ID: ASGC)**

單位(Institute): 中央研究院 物理研究所

使用費總計(TOTAL COST ESTIMATE): NT $15,006

Notes

1. 總使用費 = (計算資源使用費 + 儲存資源使用費) * 折扣
   Total Cost = (Cost of Computing + Cost of Storage) * (1-Discount (%))
   $15,006 = ($9,950 + $5,056) * (1-0%)
2. 折扣率依據合作模式而定。
   Discount rate is defined by collaboration model with ASGC.

- 計算資源使用統計(COMPUTING USAGE OF GROUP) – 使用費(INITIAL COST): 9,950

使用者(username)	姓名(name)	CPU (SRU)	GPU (SRU)	費用 (Initial Cost Estimation)
c	(	105	4,864	9,938
r		3	0	6
t		1	2	6
總計	單價SRU=NT$2	109	4,866	(109+4,866)*2=9,950

- 儲存資源使用統計(STORAGE USAGE OF GROUP) – 使用費(INITIAL COST): 5,056

帳號 (ID)	名稱(name)	儲存空間用量 (storage(TB))	費用(Initial Cost Estimation)	說明 (note)
ASGC	/ceph/work/ASGC/	10.0	756	day 方案
ASGC	/ceph/project/ASGC/	10.0	300	day 方案
ASGC	/ceph/work/ASGC/	4.0	4,000	year 方案
總計	單價NT$/TB-mo=3 NT$/TB-yr=1000	24.0	(30 day(s) * 6T * 3) + (18 day(s) * 4T * 3) + (2 day(s) * 5T * 3) + (18 day(s) * 1T * 3) + (18 day(s) * 4T * 3) + (4 T * 1000) = 5056	

使用明細(USAGE DETAILED)

使用者(username): c

	#Instance	A100		RTX3090		V100	
		CPU (SRU)	GPU (SRU)	CPU (SRU)	GPU (SRU)	CPU (SRU)	GPU (SRU)
relion311rtx3090	2			98	4,108		
jupyterlabgpu26a100	1	1	519				
jupyterlabtf24gpu3090	1			5	237		
matlab	1					0	0

使用者(username): r

	#Instance	FDR5	
		CPU (SRU)	GPU (SRU)
FDR5 slurm	63	3	0

使用者(username): t

	#Instance	1080TI		QDR6		RTX3090	
		CPU (SRU)	GPU (SRU)	CPU (SRU)	GPU (SRU)	CPU (SRU)	GPU (SRU)
relion3	4	0	0				
relion4011080ti	4	1	1				
jupyterlabtf24gpu3090	2					0	0
spyder	2			0	0		
cistem	1			0	0		
relion31	1	0	0				
transferdata	1			0	0		

資源使用說明(Resource Usage Detail): https://dicos.grid.sinica.edu.tw/resources

聯絡專線(Contact): DiCOS-Support@twgrid.org

Collaboration Model

- Pay-as-you-go: using existing resource and services
- Buy-In: joint procurement with ASGC by PI's fund and the procured hardware is shared to all NSTCCore users managed by ASGC
  - Contributor would have priority to the resource and deducted usage fee
- Service Collocation: providing shared services developed by user group(s)
- Advanced or Customized services:
  - charged service by request
  - For shared common services mainly

User Support and Communication

- User Committee will be setup in next few month
  - Represented by user groups or communities
    - < 15 members or 10% of #PI Groups
  - Help review & enhance the services, quality and resource policy, etc.
- Online contact channel
  - email: DiCOS-Support@twgrid.org
  - Rocket.Chat
- Weekly user meeting: 1:20pm every Wed
  - Action items and service status will be reviewed
- Website: news, overview and details of services
  - https://nstccore.twgrid.org

Both Supported Researches and Services Will Keep Growing

- **Next project event: planned in Q4 2023**
  - **Could be thematic seminar, user experiences/case study sharing or more hands-on sessions**
- **Sustainability: reliability and efficiency**
  - **Evolving services provided by enhanced information & communication technologies, based on users' requirements**
  - **Collaborations: with user communities and ICT service providers**
    - **Flexible Collaboration models - turning research needs into services**
  - **Good level of utilization**
- **Acknowledgement**
  - **Please include an Acknowledgment in any publication when research outcomes that relied on resources, services or expertise from NSTCCore project are presented.**
  - **For example, "This research is partially supported by the NSTCCore project (NSTC 112-2740-M-001-003) running by Academia Sinica Grid Computing Centre (ASGC)."**

Support and Service of ML-Enabled Data Analytics by ASGC

- **ML/AI application platform service is available NOW - SW library, HW, integration and application**
  - **Build up customized ML platforms for user specified projects - Deploy ML packages ready envirnonment in order to help ML development smoothly and provide on-demand computing power**
  - **Upkeep of the application framework**
  - **Workflow and data pipeline integration**
  - **Efficiency Improvement**
- **Potential use cases**
  - **Users who bring existing source code - ASGC could help to setup a virtual environment and confirm source code running normally**
- **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.**

Available Hardware, Software & Use Cases

- GPU Servers (with local SSD enhanced)
  - A100 (8xboards/server, 80GB RAM/board) * 3
  - V100 (8xboards/server, 24GB RAM/board) * 6
  - 3090 (8xboards/server, 11GB RAM/board) * 4
- ML related framework and tools
  - TensorFlow, PyTorch, Keras, NVIDIA Triton, Scikit Learn
- Large-scale storage /file system
  - 8 Petabyte+ disk-based storage system managed by CephFS
  - Tape-based backup storage will be available by end of 2023
- Use Cases
  - CryoEM - ML-enabled bioimage processing
  - Deployment of ML-enabled protein simulation tools - AlphaFold, RosettaFold & Diffusion, DiffDock
  - Deployment of ML-enabled packages (by IOP PABS group): DeepMD-kit (with interface with LAMMPS)
  - AMS & KAGRA - programs developed by local groups
  - Data Center intelligent monitoring & control (ASGC projects): Air Handler, power saving, etc.

ASGC Data Center & Resources

Cooling Power : CPU Power

1 : 2

No UPS to save 10% power consumption

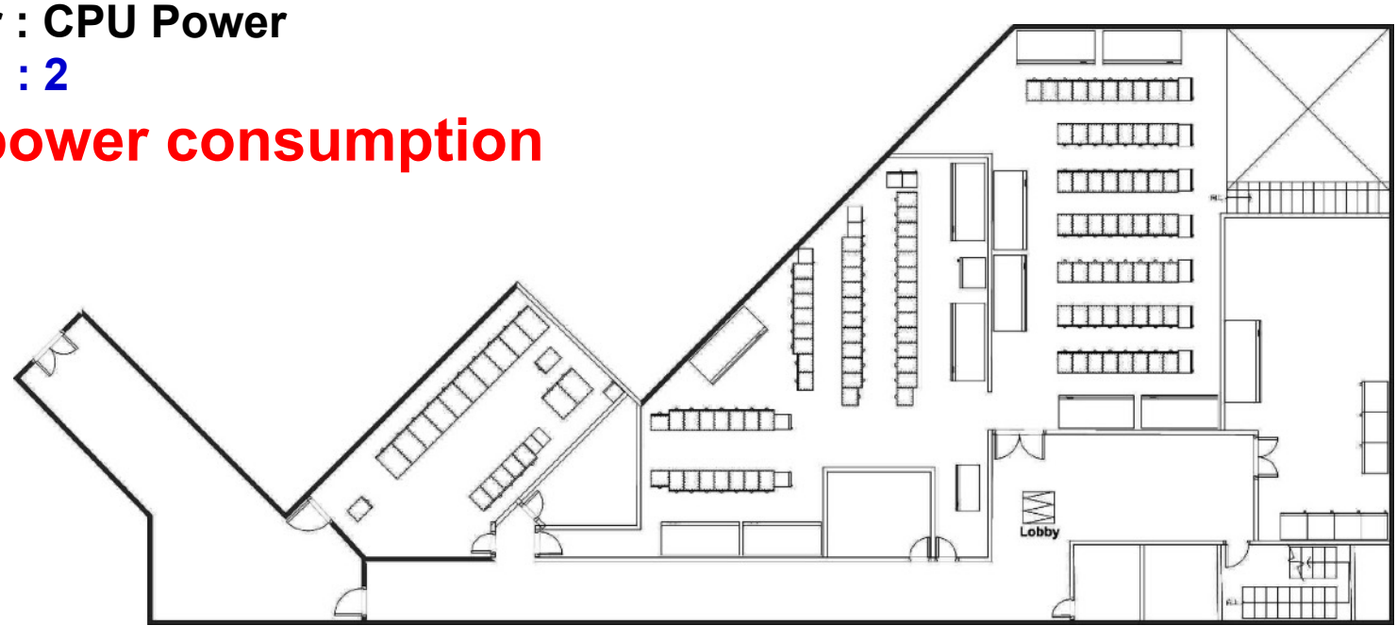

- **Total Capacity**
  - **2MW, 400 tons AHUs**
  - **112 racks in ~ 800 m²**
- **Resources (Apr. 2022)**
  - **20,090 CPU Cores**
  - **236 GPU Cards**
  - **30 PB Disk Storage**
  - **2x10Gb links to CERN and primary NRENs worldwide**
- **WLCG Tier-1 Center since 2005**
- **Supporting HPC & HTC in Academia Sinica by distributed cloud operating system (DiCOS)**
  - **Usage > 1M CPUCore-Days in 2015**
  - **Usage > 2M CPUCore-Days in 2019**
  - **GPU usage is growing exponentially from 2017**
- **Reliability: > 99.9% yearly average**
- **R&D on system efficiency optimization by intelligent monitoring & control**

All software used are open-source codes developed by ASGC and an international collaboration led by CERN

NSTCCore Services

- Project Website: https://nstccore.twgrid.org
- Access to ASGC Resources
  - https://dicos.grid.sinica.edu.tw/
- Contact point: DiCOS-Support@twgrid.org

Backup Slides