

Science Cloud in Taiwan


**Jing-Ya You, Felix Lee
and Eric Yen**

**Academia Sinica Grid Computing Centre (ASGC)
Taiwan**

**ISGC 2025
18 March 2025**

ASGC - Core Facility for Scientific Computing

- WLCG Tier Centre - building capacity by participating the development & deployment of the world largest and advanced distributed cloud research infrastructure
 - WLCG is moving towards O(1000)PB scale in RUN4
- Core Facility - computing arms of HEP experiment, research facility and research communities, based on WLCG core technologies
 - Also part of the research infrastructure



Neutrino/MHEP

CryoEM

Quantum Materials
Physics

Bioimaging

Physics of Active &
Living Matter

Drug Discovery

Astrophysics

Computational
Chemistry

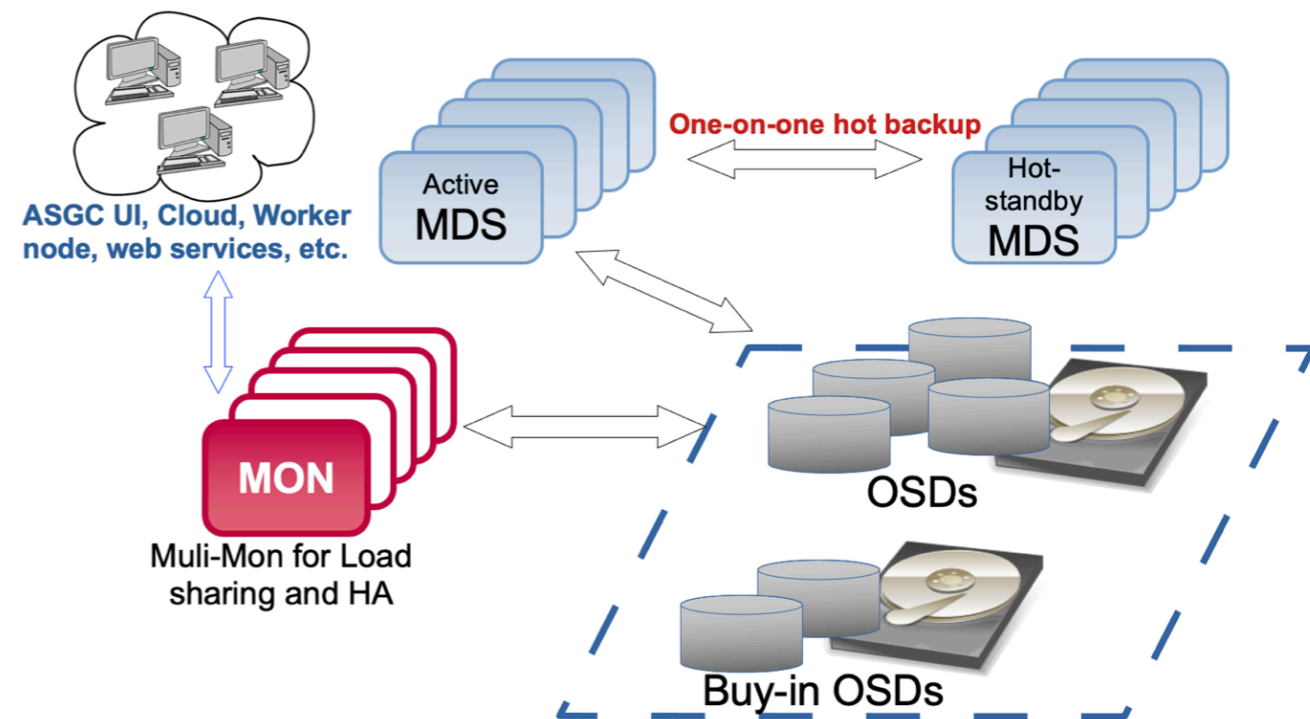
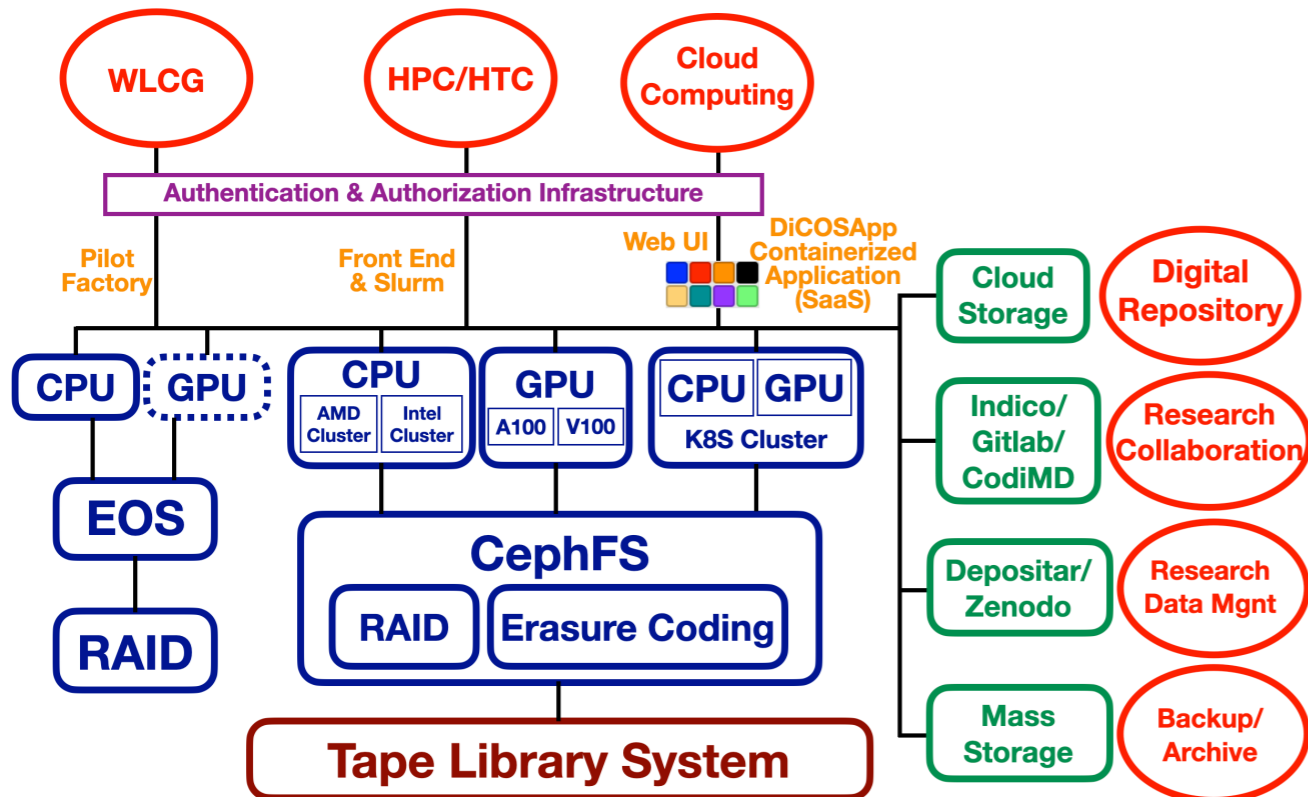
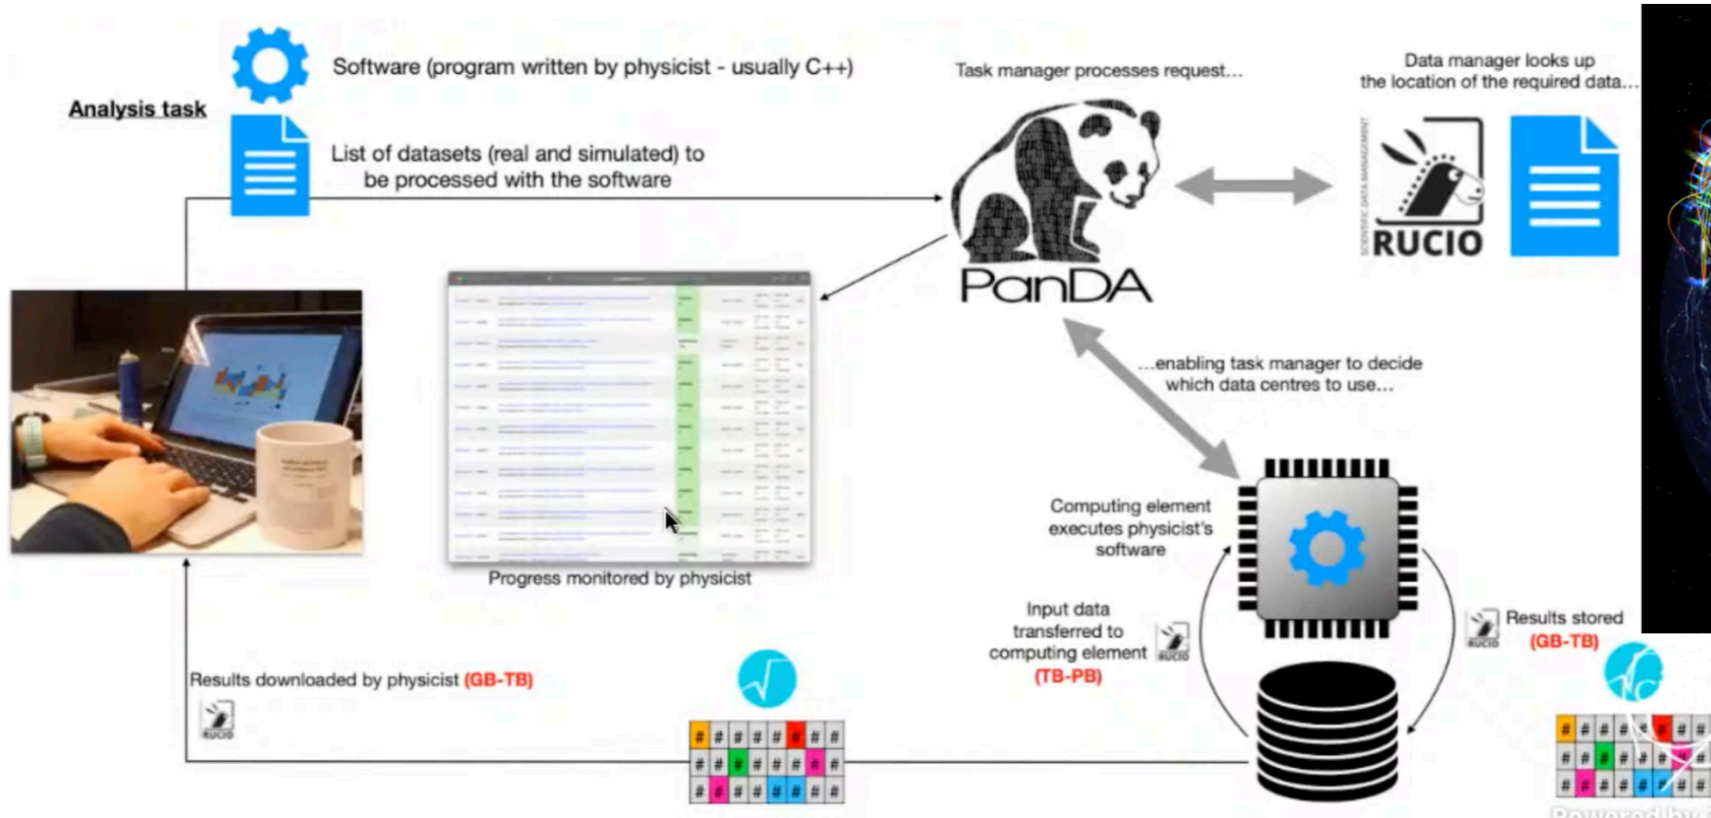
Earth Science

Biodiversity &
Ecology

WLCG-based Distributed Cloud infrastructure for Scientific computing

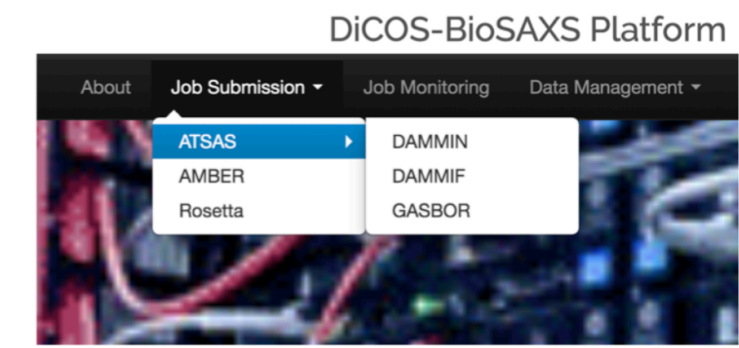
- **PanDA + RUCIO serves as the core of the common distributed infrastructure**
 - Federation of distributed institute resources
 - Federation of core facilities (including CryoEM, NSRRC, NMR, computing centre, etc.)
- **Web-based Cloud services & Slurm clusters are provided**
 - VM for core services and on-demand worker nodes managed by OpenStack
 - Containerized resources managed by Kubernetes framework - for software on-demand services and part of core services
 - Batch and interactive GUI jobs: Jupyterlab, virtual desktop
 - GPU Cloud
 - SaaS: web-based application environment with microservices
- **Resource status: 4.4M CPU jobs, and 50K GPU Jobs in 2024 (Jan. 2025) *****
 - 5,152 CPU Cores + 2,176 CPU Cores (by end 2024)
 - 24x A100 GPU Boards, 16x4090, 48xV100, 56x3090 (+ 16x L40S soon): high demanding
 - Ceph is the common scalable storage pools: > 10PB is online (+2.5PB by end 2024)
 - 6 MDS + 6 hot-standby (one-on-one backup); 7 MONs
 - 462 OSDs, 51 hosts.
 - 1x full rack 10PB LTO9 tape storage will be available in Q4 2024 - backup
- **Both core technology and application platform are evolving with user experiences**
 - 170 PI Groups, 565 users from 64 Institutes in Taiwan
- **OSG software stack has been deployed for IGWN and will be part of the infrastructure**

ASGC Science Cloud Architecture

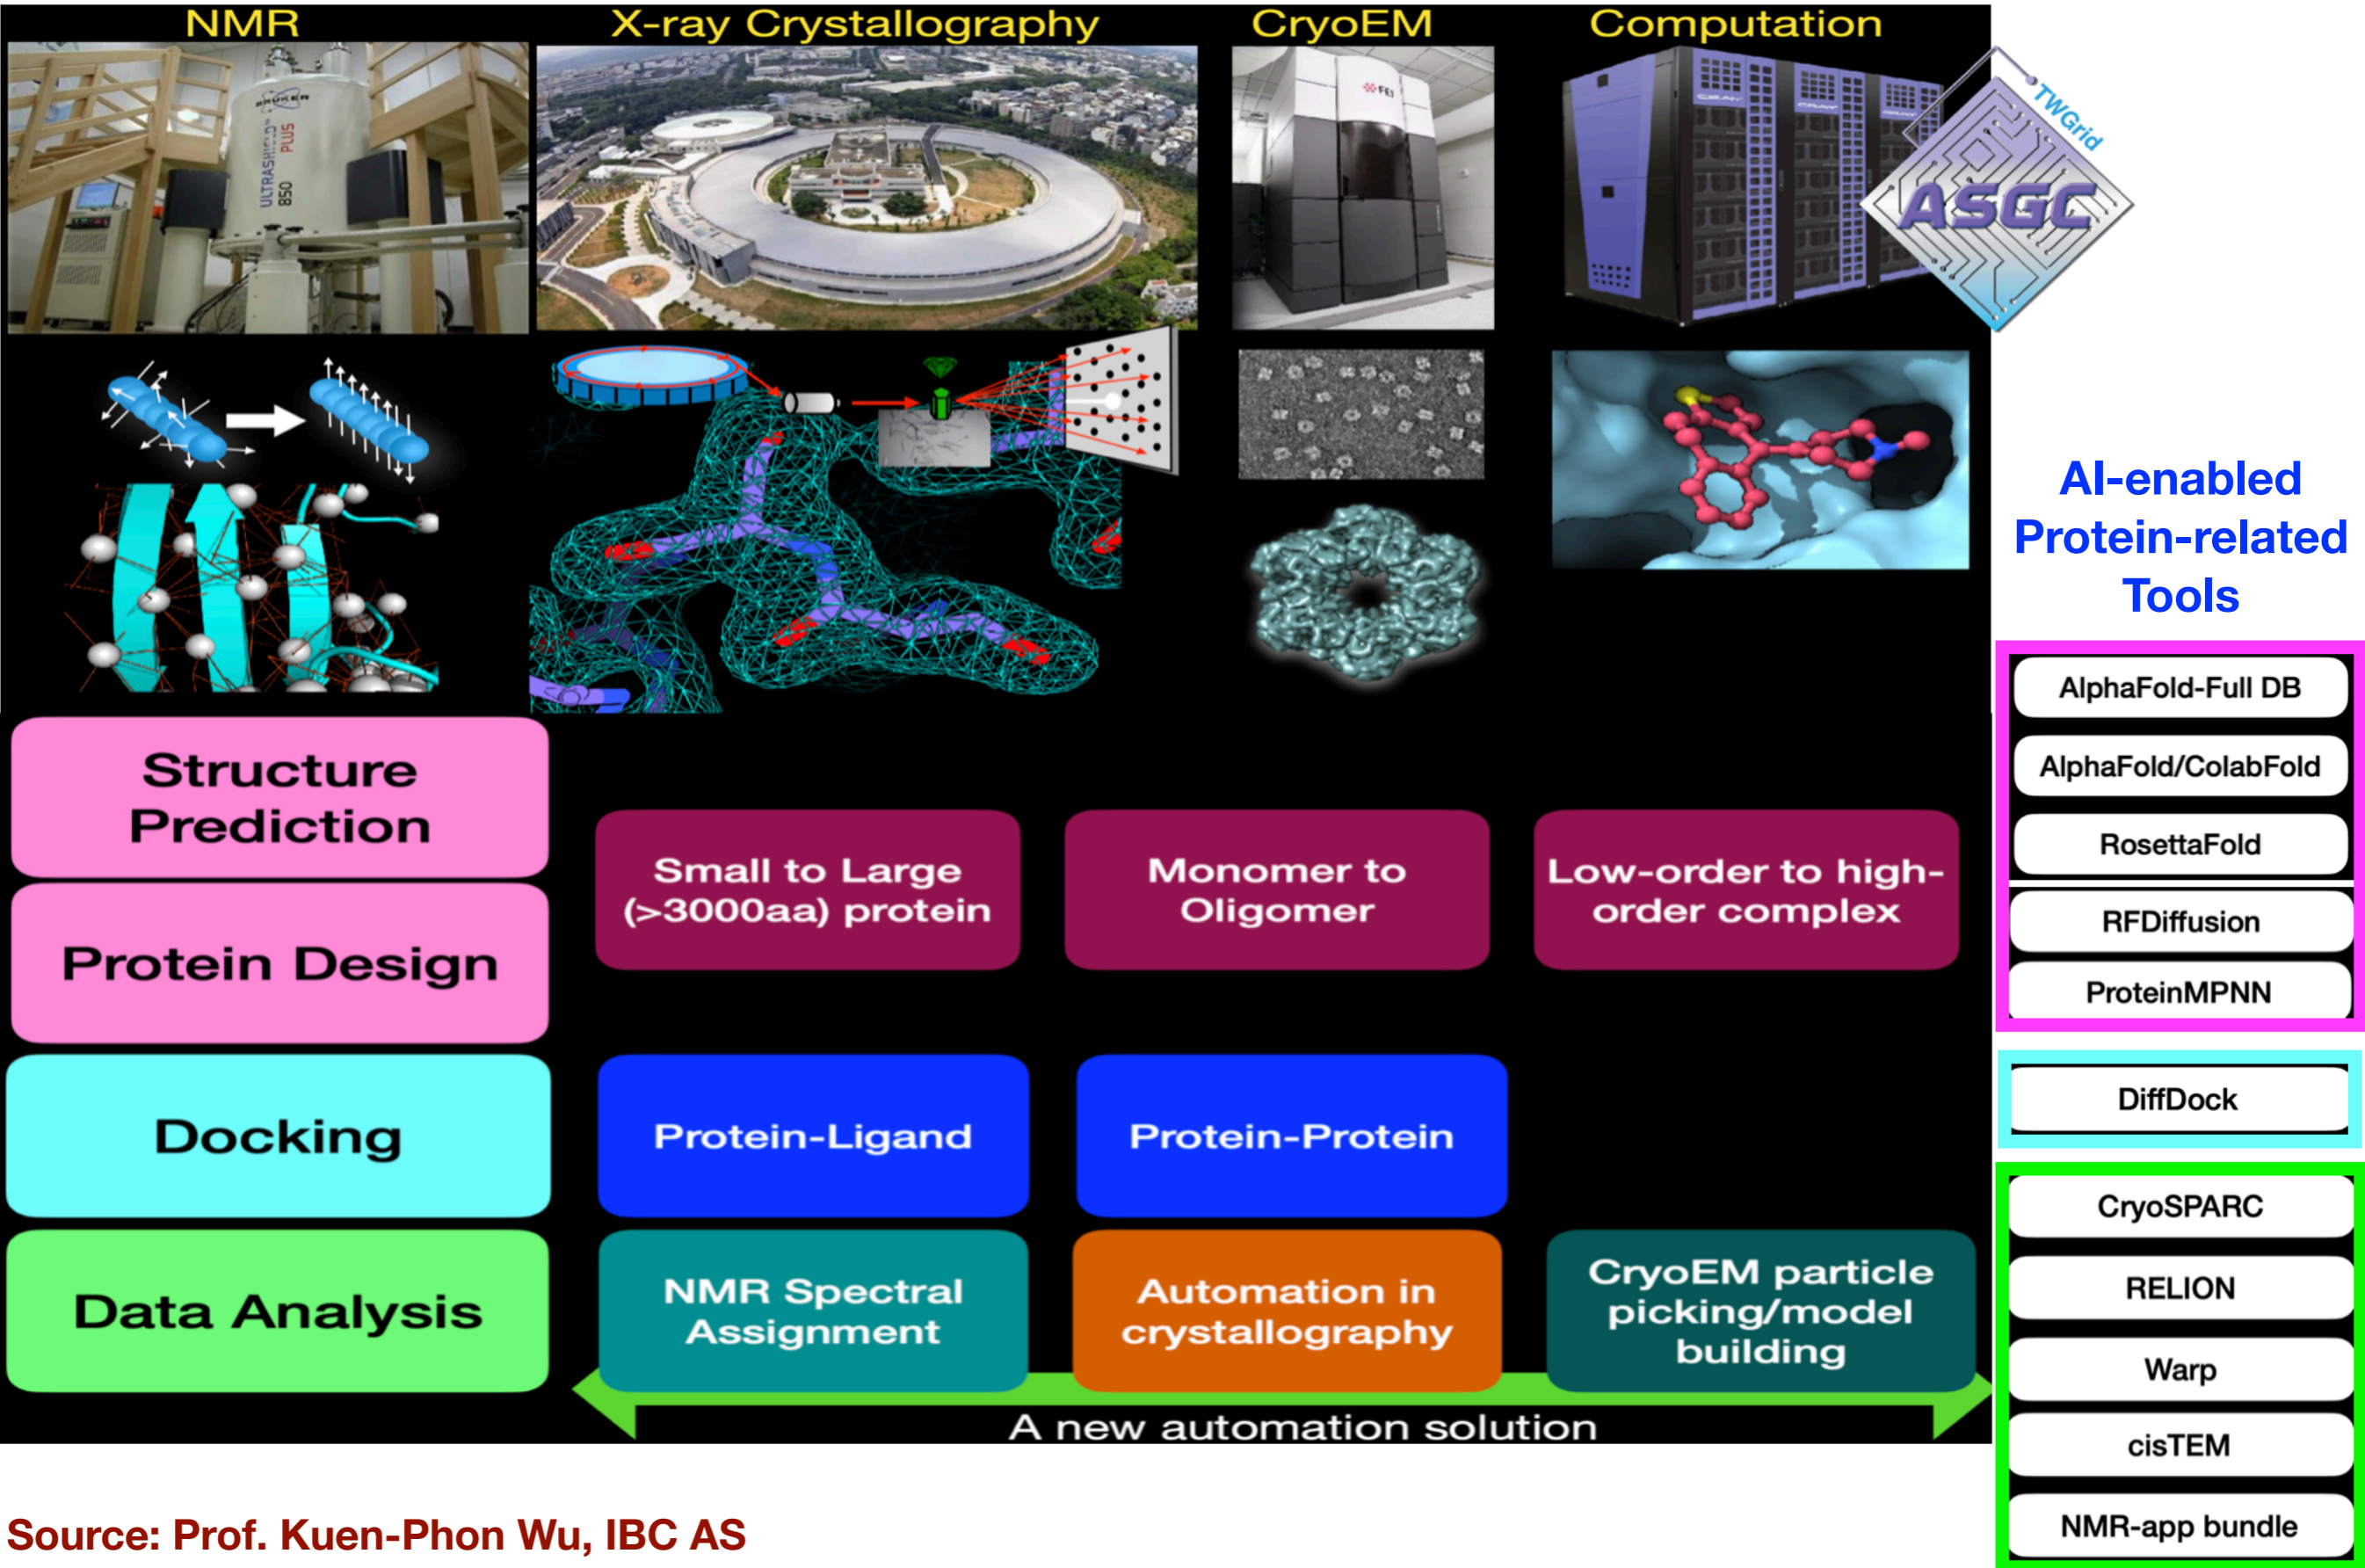


73+ Web Applications Provided (Software-as-a-Service)

CryoSPARC 32 P100 Version: 3.2 Resources: 83%	CryoSPARC 1080ti Version: 3.3.2 Resources: 29%	CryoSPARC RTX3090 Version: 3.3.2 Resources: 26%	spyder cpu/eman2 Version: Resources: 17%	Octave Version: V5.2 Resources: 29%	Transfer Data Version: Resources: 86%
CryoSPARC RTX3090 Version: 4.0.2 Resources: 26%	CryoSPARC RTX3090 Version: 4.1.1 Resources: 26%	CryoSPARC RTX3090 Version: 4.4.1 Resources: 26%	cisTEM Version: Resources: 86%	Ovito Version: Resources: 86%	OpenACC Version: GPU P100 Resources: 83%
AlphaFold Version: GPU with A100 Resources: 75%	AlphaFold (Full DB) Version: GPU with A100 Resources: 75%	IMOD Version: GPU with 1080ti Resources: 29%	Triton Version: 22.01-py3 (GPU P100) Resources: 83%	AlphaFold Version: GPU with RTX3090 Resources: 26%	AlphaFold Version: GPU with V100 Resources: 17%
RoseTTAFold Version: GPU with rtx3090 Resources: 26%	Dynamo Version: GPU with 1080ti Resources: 29%	MATLAB Version: R2018b on GPU V00 Resources: 17%	Jupyter Lab Version: CPU with Tensorflow v1 Resources: 86%	Jupyter Lab gpu 3090 Version: GPU with Tensorflow 3090 Resources: 26%	Jupyter Lab GPU 1080ti Version: GPU with Tensorflow v2 Resources: 29%
RFDIFFUSION Version: 2023 on GPU V00 Resources: 17%	diffdock Version: 2023 on GPU V00 Resources: 17%	EvoDiff Version: V100 Resources: 17%	Jupyter Lab GPU V100 Version: GPU with Tensorflow V100 Resources: 17%	Jupyter Lab GPU A100 Version: GPU with Tensorflow A100 Resources: 75%	Jupyter Lab Cryocare GPU Version: GPU with 1080ti Resources: 29%
QIIME2 Version: Genome Resources: %	Scipion3 Version: P100 Resources: 83%	Phenix Version: Resources: 86%	Jupyter Lab GPU A100 Version: GPU with Tensorflow v2.6 Resources: 75%	<ul style="list-style-type: none"> • Web Portal • Application over Cloud 	
MorphoGraphX Version: GPU with P100 Resources: 83%	Deepmd-kit Version: GPU with A100 Resources: 75%	Deepmd-kit Version: GPU with V100 Resources: 17%	MAML Version: GPU with A100 Resources: 75%	<ul style="list-style-type: none"> • Jupyterlab • Web Terminal 	
Warp Version: rtx4090 Resources: 700%	MAML Version: GPU with V100 Resources: 17%	PVserver Version: 5.8.0 (GPU 1080T) Resources: 29%	Paraview Client Version: 5.8.0 Resources: 86%	LabVIEW Run-Time Engine Version: 2019	
R studio Version: 10 CPU Cores Resources: 54%	PyRoot Version: GPU with 1080ti Resources: 29%	qiskit Version: Resources: 86%	NVIDIA CUDA-Quantum A100 Version: Resources: 100%		




Customized Cloud Services for Proteomics/Life Science Communities



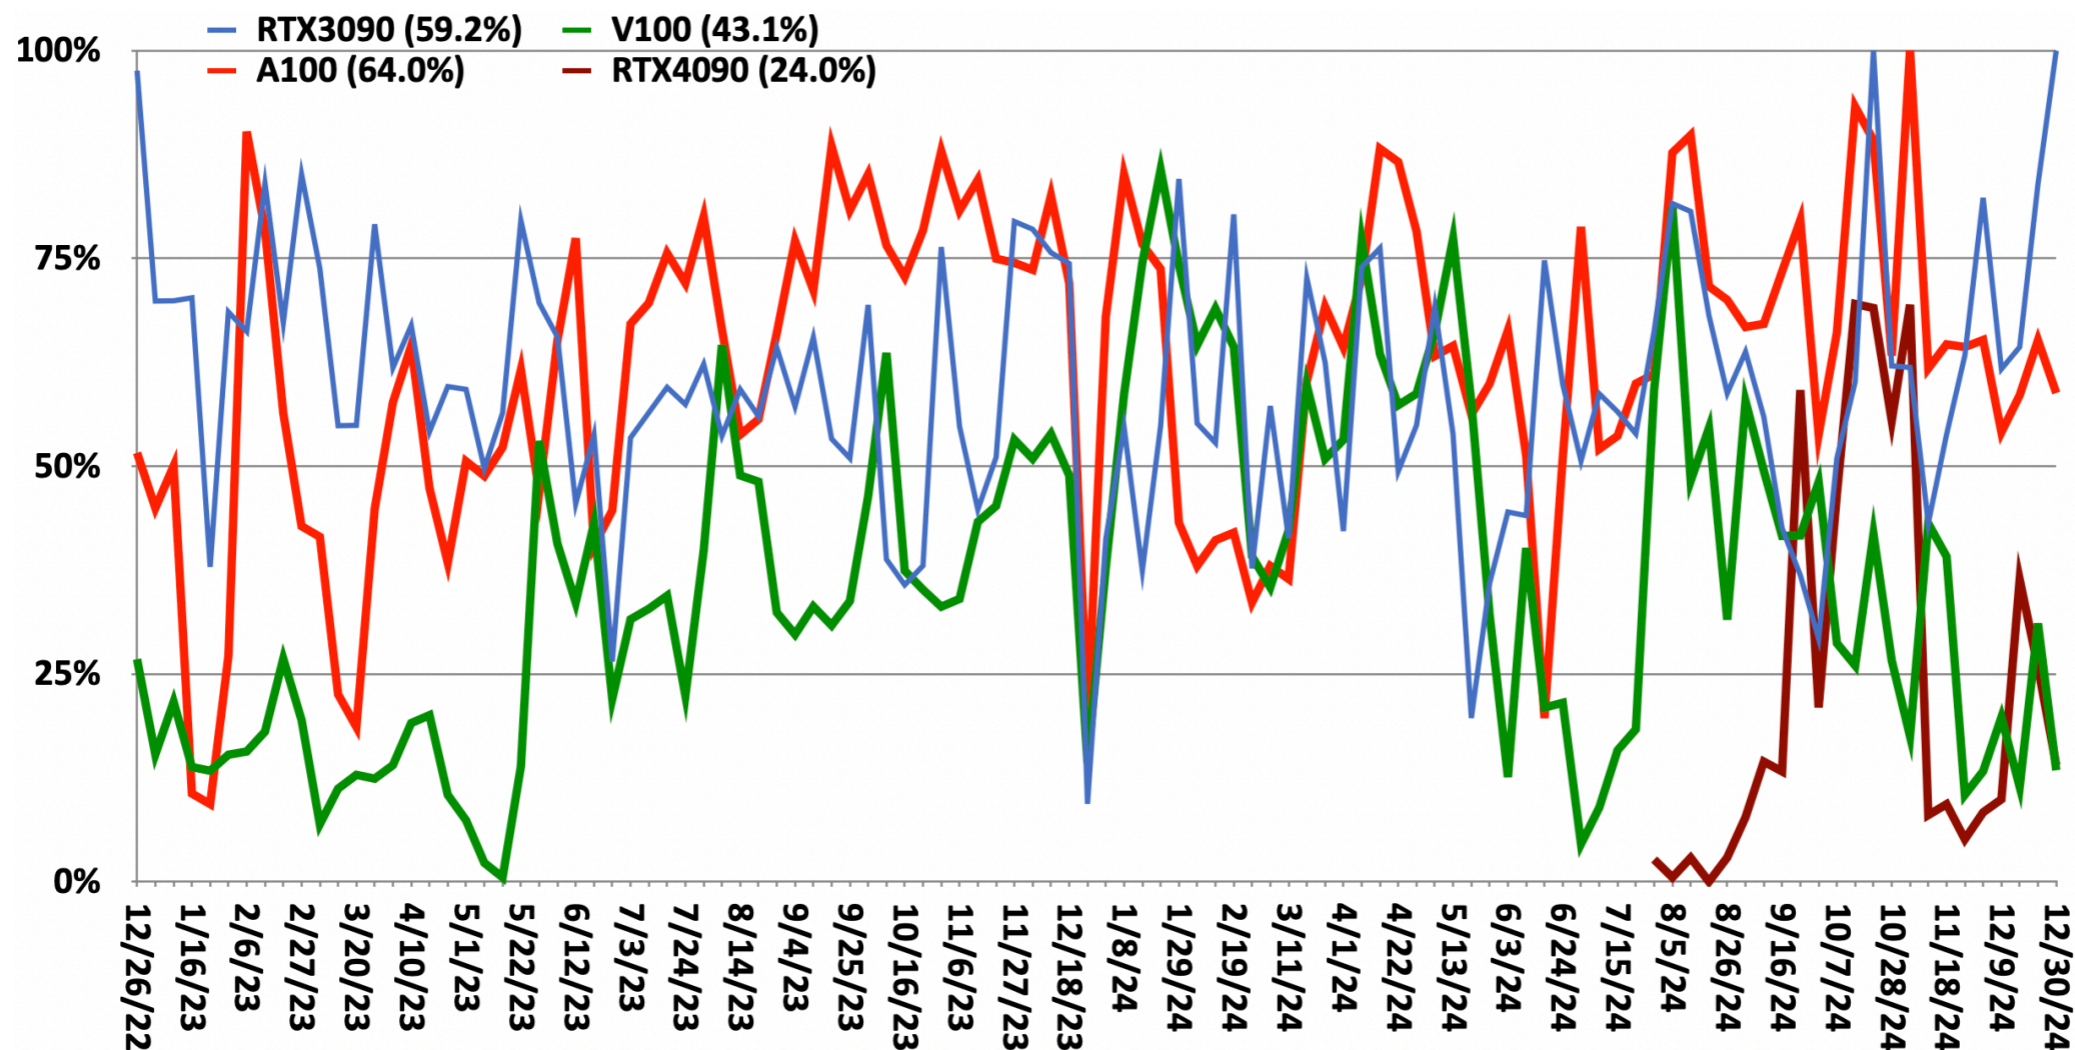
ASGC Science Cloud Resource Usage

- **Sustained growth of users, utilization, reliability and satisfaction in 2024**
 - 174 PI Groups, 595 users from 25 Institutes in Taiwan
 - Finished 4.94M CPU jobs, and 52.8K GPU Jobs, with 99%+ reliability
 - 26+ supported research publications
 - Demands of advanced GPU keeps growing in 2024 - 20% utilization growth in high usage situation
 - CPU usage has increase of 333% in 2024
- **Gaining trust of users/ building collaboration with user communities**
 - Customized services: just like tailored-made services in the cloud/ serving as the computing arm
- **Onboarding new hardware and retiring legacy ones for power saving and performance improvements: 5,760 CPU Cores, 210 GPU Boards, 12PB Disk, 410PB Tape**
 - New procured HW (1,920 + 256 CPU Cores, 2.5PB Disk Storage) had been available on 20 Dec 2024




GPU Resource Usage in 2024

- GPU demand keeps growing in 2024: 20% utilization growth comparing to 2023 under high usage situations
 - Diversified user communities
 - Not just AI model training and development, scientific computing also requires large amount of high-end GPU resource
- Users can hardly access to the most advanced GPU like H100 - Urgent action is needed before being GPU disadvantaged
 - Unaffordable usage cost, not mentioning the acquisition of powerful enough high-end GPUs
 - Long waiting queue for competing the rare resources for the academic
- A100 servers @ASGC with large embedded GPU memory (80GB) and interconnection bandwidth (NVLink) becomes the alternative
 - Congested jobs/requests for limited A100 servers (24 boards in 3 servers) has been since 2023
 - A100, 3090 are almost saturated - average utilization in 2023&2024 are around 60%. Users were forced to choose less advanced V100. Usage of V100 in 2024 has grown by 1.4 times (average utilization is about 45%). Even the primitive GPUs (P100 & 1080ti) are also highly utilized.
 - Although newly procured 4090 (8 boards only) served as the substitute in time
- More investment in advanced GPU is needed, for both AI development as well as scientific computing



Energy Efficiency Improvement

- Reliability enhanced by intelligent monitoring and control is the key approach
- Retirement of legacy hardware
- Improvement of AHU efficiency, including the replacement by top-flow cold air
 - Anomaly detection
 - Well-prepared backup plan
- Energy-sensitive operation:
 - Plan for power efficient hardware: e.g., non-X86 CPUs
 - Power saving - shutdown idle WNs when the waiting queue is quite short
 - 20% power usage reduction in 2023 - Effective on 3 CPU clusters (> 3,000 CPU Cores) from May 2023
- Overall, DC Power consumption reduced 20% for two straight years from 2023



Capacity Building

- **User community engagement with flexible collaboration model**
 - QCD, DFT, First Principle, Astronomy/Astrophysics, Neutrino, GW, HEP, Proteomics, Computational Chemistry, Drug discovery, NGS, Earth Science, Social Science, etc.
 - Engagement with other core facilities: ASCEM, NMR, NSRRC
 - Understanding user requirements and supporting workflow integration
 - Collaboration model: pay-as-you-go, buy-in, facility/service collocation, reservation
- **Supporting user access to ASGC services easier and enhance application efficiency**
- **Training and workshop/conference: 250+ participants in 2024**
 - 4x training a year on ASGC services with hands-on
 - Thematic computing/application training together with user communities: 3x events in 2024
 - Also covers AI applications, HEP data analysis and more activities.
- **Core technology and international collaboration: ISGC (annual event since 2002)**



Action Plans in 2025

- **Broader user community engagement**
 - Condense Matter, Earth science, life science, engineering, AI & AI-enabled researches, research data management etc.
- **Enhanced resilience & efficiency of storage system**
 - Moving cold data to Tape: KAGRA, Soundscape, Earthquake, Satellite images, etc.
 - Kick off tape backup services
- **AI-enabled big data analytics: extending from development platform and software/program integration**
- **Gaining certification of ISO27001 + ISO27017 (for Cloud Service)**
- **HW procurement and retirement plan**
 - Another 2,176+ CPU Cores (higher than AMD Genoa) and GPUs (L40S or higher-end) online
 - 3PB+ disk and 4PB+ tape on board
 - 1080ti retirement
- **Capacity Building: another 4x training events, ISGC2025**
 - GPU Computing, AI application, ASGC Core Technology (distributed Cloud, BDA, AAI, data repository), etc.
- **Effective communication**
 - Contact point
 - Responsive
 - Information service
 - Weekly user meeting
- **Call for collaboration - flexibility and R&D oriented**
 - AI-enabled BDA for research applications and new services
 - Efficiency
 - Security - intelligent monitoring and responses

Summary and Future Works (Update Needed)



- **Supporting high-performance data analysis needs of multiple-disciplinary applications by making better use of ASGC resources, with reliability and efficiency**
- **Based on WLCG core technology, a scalable, consistent progress distributed computing system and services are available for TIDC**
 - PanDA+RUCIO, AAI, CVMFS, monitoring & control, system efficiency and security
 - Supporting both batch (HTC/HPC) and cloud services (SaaS) at the same time
- **Both the capacity of users and ASGC are growing according to collaborations and ICT evolution**
 - Extended from local ATLAS/CMS groups, to broader scientific communities
 - Providing ASGC service training and thematic training/workshop
 - Collaboration platform: ISGC annual event since 2002
- **Upcoming Services**
 - Data efficiency and reliability: enrich backup and archive services by integration of disk & tape storages
 - Migration to CTA will be kicked off in 2025
 - Two Factor Authentication service has been effective since Jan 2025
- **Prototyping Hybrid quantum cloud service in the near future**

Thank you

Backup Slides

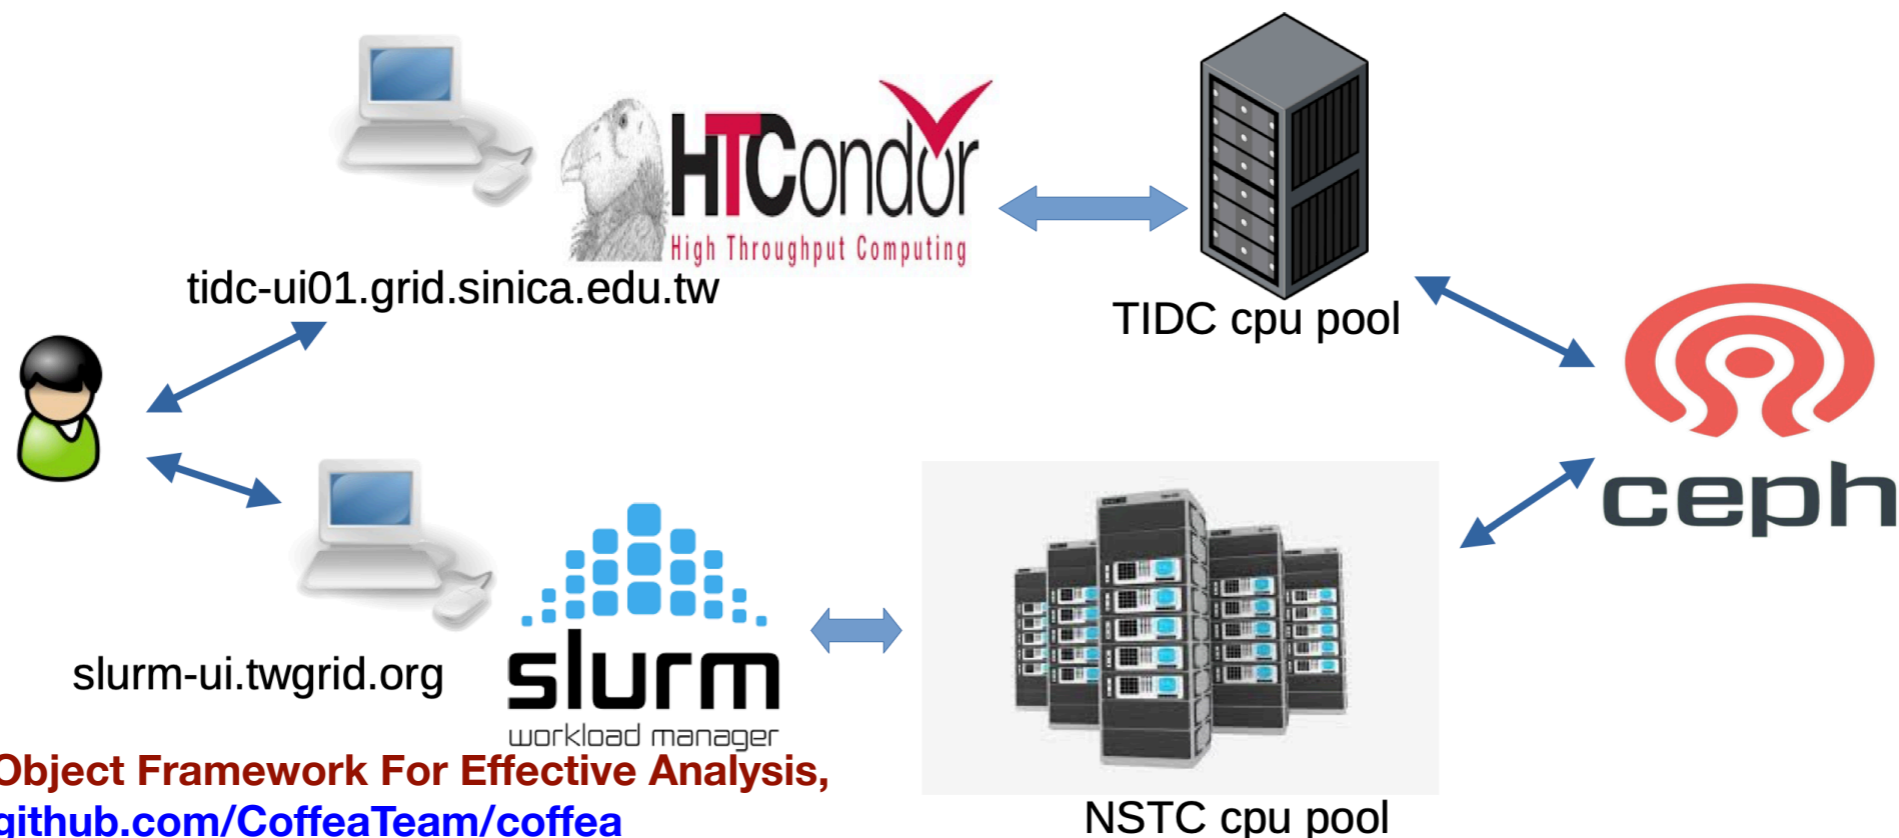
WLCG Tier2 for ATLAS is running smoothly

- Site has been in production since July 2024 with helpful and careful investigation of USATLAS
 - Connection with LHCONE was reestablished, efficiency was also affirmed
- Supporting MC Simulation jobs and analysis jobs mainly afterwards
- Current Status: > 111K CPUCore-Days (2024.8 - 2025.2.26)
 - Transfer efficiency: > 92%
 - CPU Efficiency: > 89%
 - Storage used: > 1.2PB
 - Data Transmission: > 2PB (in+out)
 - Issue: limited network bandwidth (3Gbps) connecting to LHCONE
- Pledge in 2025
 - CPU: 30K HEPscore23
 - Disk: 5 PB, managed by EOS
- Job slots - 2,208
 - ~ 839 job slots are available now because of OS migration
- Next Step:
 - Making better use of ASGC resource, e.g, pre-stage for analysis jobs
 - Will take USATLAS operation shift in Asia time zone in 2025
- Migration to AlmaLinux 9 - will be accomplished before Q1 of 2025
 - 2-stage: ~ 1,370 job slots in AlmaLinux will be online in Sep.



CMS Tier3

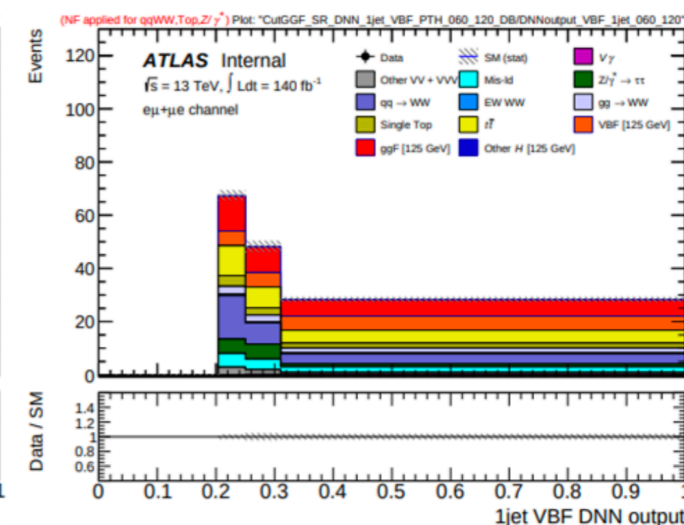




- CMS T3 is managed by ASGC from 2022, in collaboration with TIDC and local CMS groups (NTU and NCU)
- Analysis facility: 6,496 CPUCore-Days, 1,727 GPUBoard-Day, 463,422 events processed in 2024
 - Both Condor/UI and CRAB/ARC-CE are available
 - CephFS shared filesystem: 3TB/group by default
 - EOS by xrootd and fuse: 1PB
 - Condor cluster
 - 768 cores(AMD EPYC 7713) + 768 cores (Intel CPU E5- 2650 v4)
 - Supporting user's access to CMS data or CRAB jobs submission
 - JupyterLab is also available for local user analysis
- User training and support are also provisioned
 - e.g., using Coffea with template for columnar analysis @ASGC



Coffea = Column Object Framework For Effective Analysis,
<https://github.com/CoffeaTeam/coffea>

Taiwan Contributions in ATLAS

- AS served as Online Data Quality coordinator for 7 yrs (2017, 2018, 2023, 2025)
- Higgs measurements in $H \rightarrow WW$ decay (NTHU + AS)
 - Differential kinematic measurement in $HH \rightarrow WW^* \rightarrow l\nu l\nu$
 - An overall 20% improvements across different kinematic regions
 - Quantum entanglement
- Search for double Higgs boson in $HH \rightarrow bb\tau\tau$ decay
 - World best expected sensitivity for single channel
- Search CP-odd Higgs boson in $A \rightarrow ZH \rightarrow llbb$
- High Granularity Timing Detector (HGTD) for ATLAS in HL-LHC
 - Taiwan-ATLAS responsible for producing electrical and optical fiber cables, and implementing a production database



HH trigger acceptance improvement

GGF and VBF DNN in the 1 jet category