Science Cloud in Academia Sinica

Felix Lee
Academia Sinica Grid Computing Centre (ASGC), Taiwan

ISGC 2022
22 March 2022

ASGC Mission

- Founded as a WLCG Tier-1 Centre in 2005 and jointly develop & deploy advanced distributed cloud infrastructure for WLCG
- Enhancing science-enabling capabilities by advanced research infrastructure
  - Building capacity of large-scale distributed cloud for efficient big data analysis of Academia Sinica
  - Distributed cloud technologies and infrastructure are improved progressively with growing scientific applications of various disciplines
  - ML-enabled data analysis framework is also equipped
- System efficiency optimization: power, thermal, application, operation, system, etc.

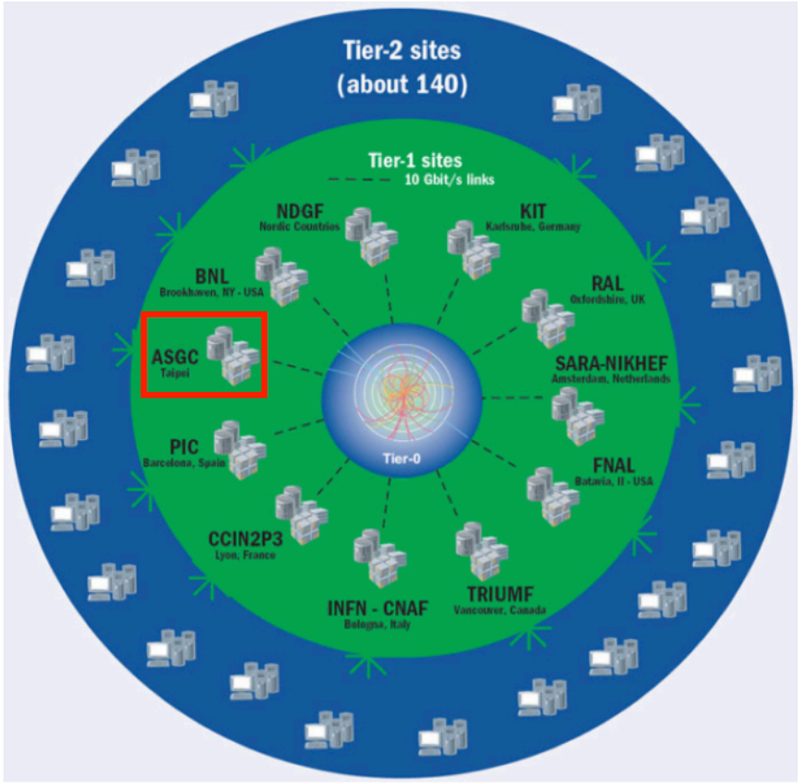

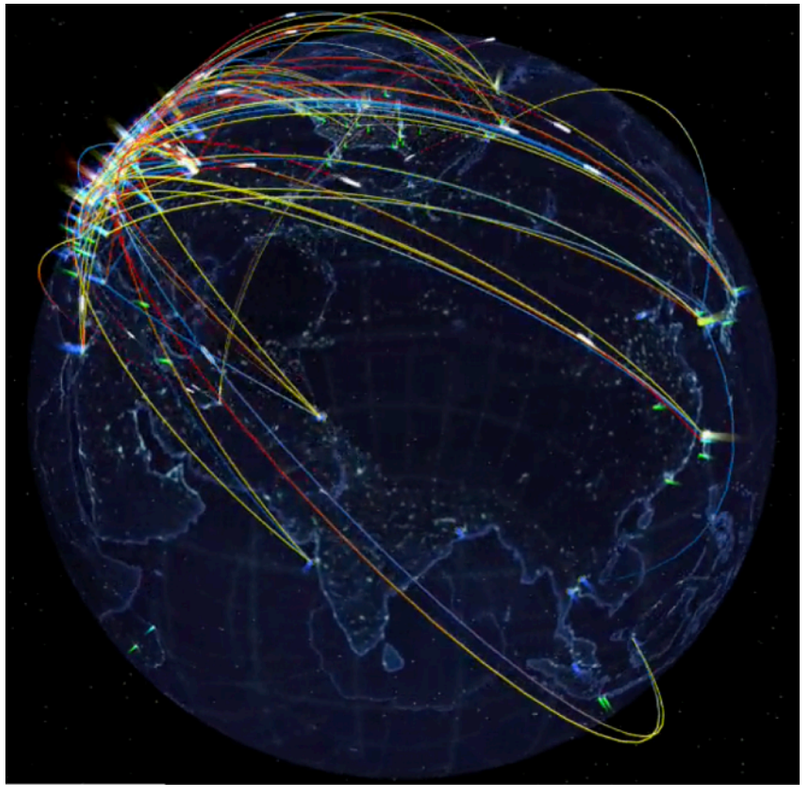

ASGC Resources

Cooling Power : CPU Power
1 : 2
No UPS to save 10% power consumption

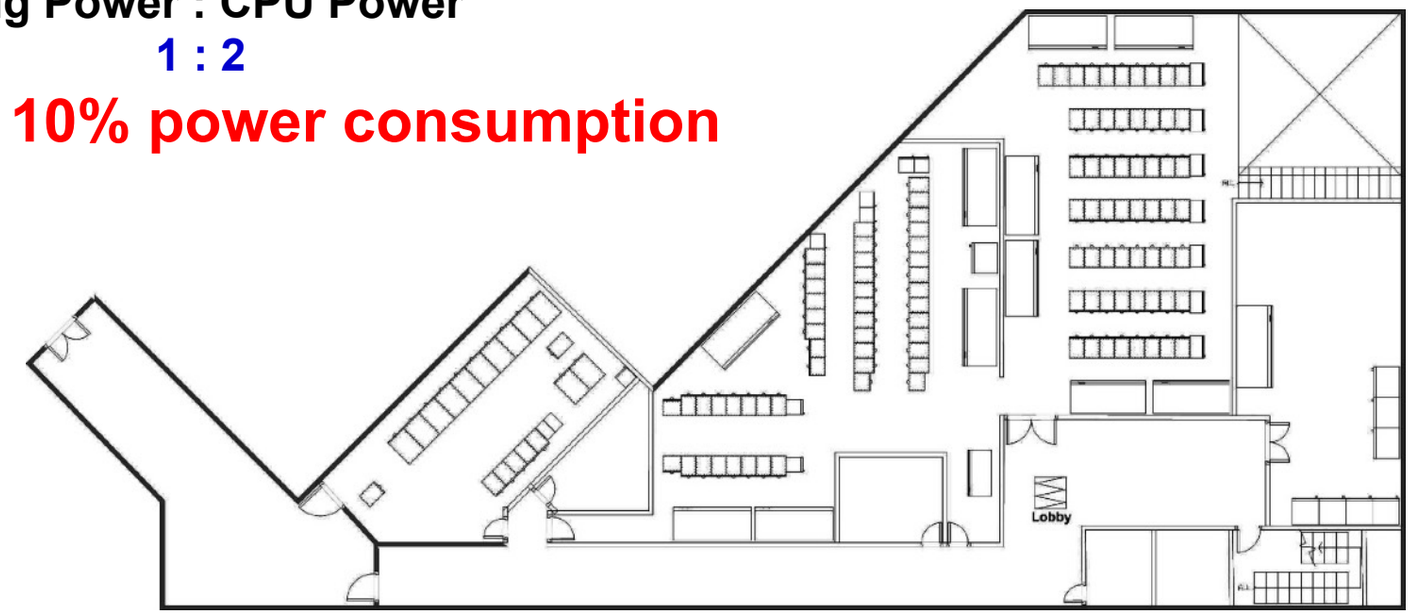

- **Total Capacity**
  - 2MW, 400 tons AHUs
  - 112 racks in ~ 800 m²
- **Resources (2021)**
  - 13,652 CPU Cores
  - 192 GPU Cards
    - A100 Server *2 (8 GPU Card for each) are available in 2021
  - 24 PB Disk Storage
  - 2x10Gb links to CERN and primary NRENs worldwide
- WLCG Tier-1 Center since 2005
- Supporting HPC & HTC in Academia Sinica by distributed cloud operating system (DiCOS)
  - Usage > 1M CPUCore-Days in 2015
  - Usage > 2M CPUCore-Days in 2019
- R&D on system efficiency optimization by intelligent monitoring & control

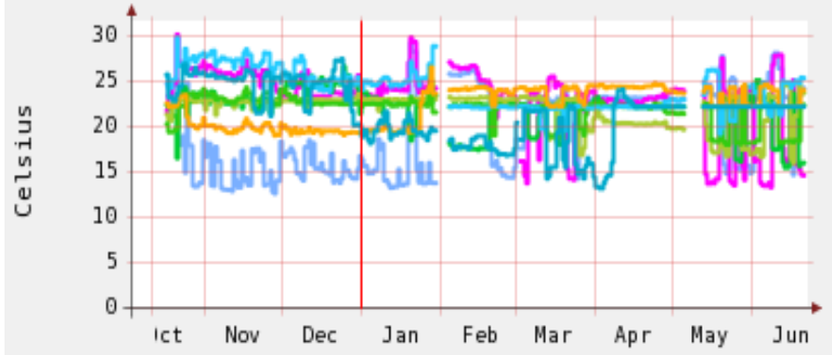

Monitoring the power consumption and temperature of every piece of equipment every 10 seconds.

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

3

Primary User Communities

- **User Communities**
  - **ATLAS/CMS/WLCG: search for the origin of matter**
  - **AMS: explore the universe with Space-borne detector**
    - **Launched to ISS in 2011; ~200BN cosmic ray events recorded**
  - **TEXONO/Neutrino: search for the origin of left-right asymmetry;**
  - **Gravitational wave: KAGRA and LIGO**
  - **CryoEM: Assist analyzing protein structure and cellular ultrastructure for further solutions of bio-technology issues and serious disease problems.**
  - **BioSAXS: advanced bio-structural researches by biological small-angle X-ray scattering beamlines**
  - **NGS: precision medicine;**
  - **Astronomy & Astrophysics**
  - **ML-enabled analytics;**
  - **Others: Drug discovery; Proton therapy; Lattice gauge theory; Condense Matter; Earth science; Environmental changes; Hazard risk analysis; Biodiversity and ecology monitoring; etc.**
- **Types of services: HPC, HTC, resource federation, data backup and archive, data transmission, analysis pipeline integration, web portal development and customization, as well as ML-enabled application platform, etc.**

Regional Collaborations

- **Thematic Data Space**
  - **WLCG: research infrastructure and user support**
  - **Gravitational Wave**
  - **Disaster Mitigation: capacity building & case study**
    - **Sentinel Asia: making better use of earth observation data**
    - **Hazard risk analysis based on deeper understanding approaches and numerical simulations**
    - **Master class, training, workshop**
    - **Hazard types: tsunami, storm surge, extreme weather, flood, forest fire/smoke monitoring,**
    - **Collaborators: 11 Asian countries, ASEAN, UNESCO**
  - **Agriculture: Deeper understanding approach and case studies**
- **Collaboration Framework: APAN, WLCG, ISGC, regional infrastructure**
  - **WGs @APAN: Disaster Mitigation, Agriculture, Open & Sharing Data, Cloud, Security, Network, Advanced Research Platform, Identity & Access Management**
  - **EGI, EGI-ACE/EOSC: distributed infrastructure, application market place, technology/platform/service**

ASGC 2012 - Oct. 2021 Trends of Resource Utilization

- CPU Computing: linear growth rate ~15%
- GPU Computing: Exponential growth rate ~220%

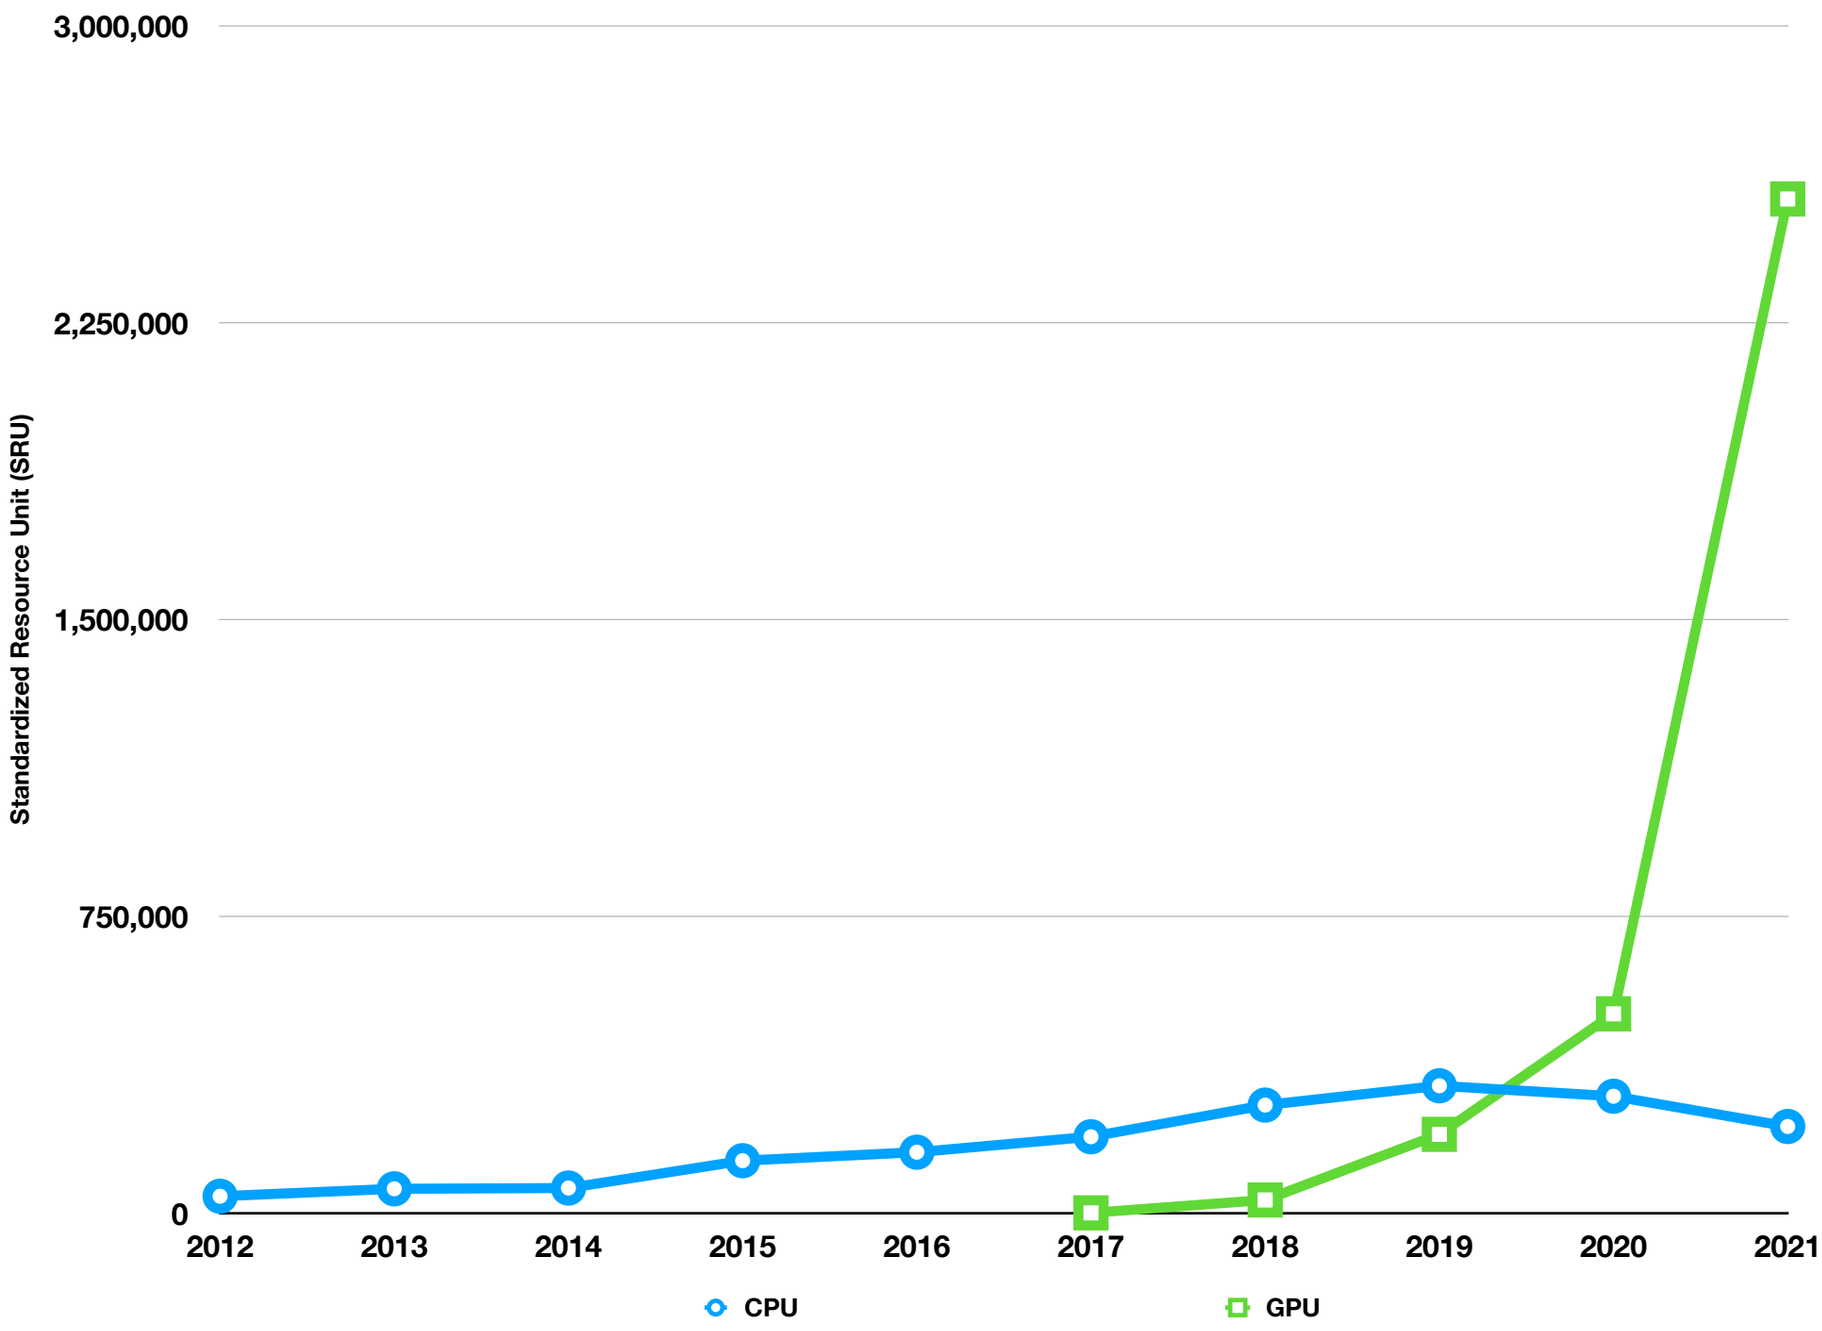

Cloud Infrastructure

- OpenStack Cloud: for core services and on-demand worker nodes
  - Separate regions for scalability
  - Multiple cells/Region for various configurations and capabilities
    - e.g. GPU, Netron Compute, Nova Compute, …
  - Single hypervisor type: KVM
    - #hypervisors: 100+
    - #VMs: 500+, dynamic provisioning
  - Networking: flat and segmented
- Containerization by Kubernetes framework
  - User cluster:
    - batch, interactive GUI jobs: remote Jupyterlab, virtual desktop
    - GPU Cloud
  - Core Services: distributed cloud cores; Elasticsearch + Kibana
- Storage system: Ceph-based
  - OpenStack services (RBD): 1.5PB
  - Users: CephFS (7+PB)
- UI: Web UI/Terminal; JupyterLab
- Operation and management: automate everything !
  - Source control: Gitlab
  - Puppet-based deployment of components
  - HELM

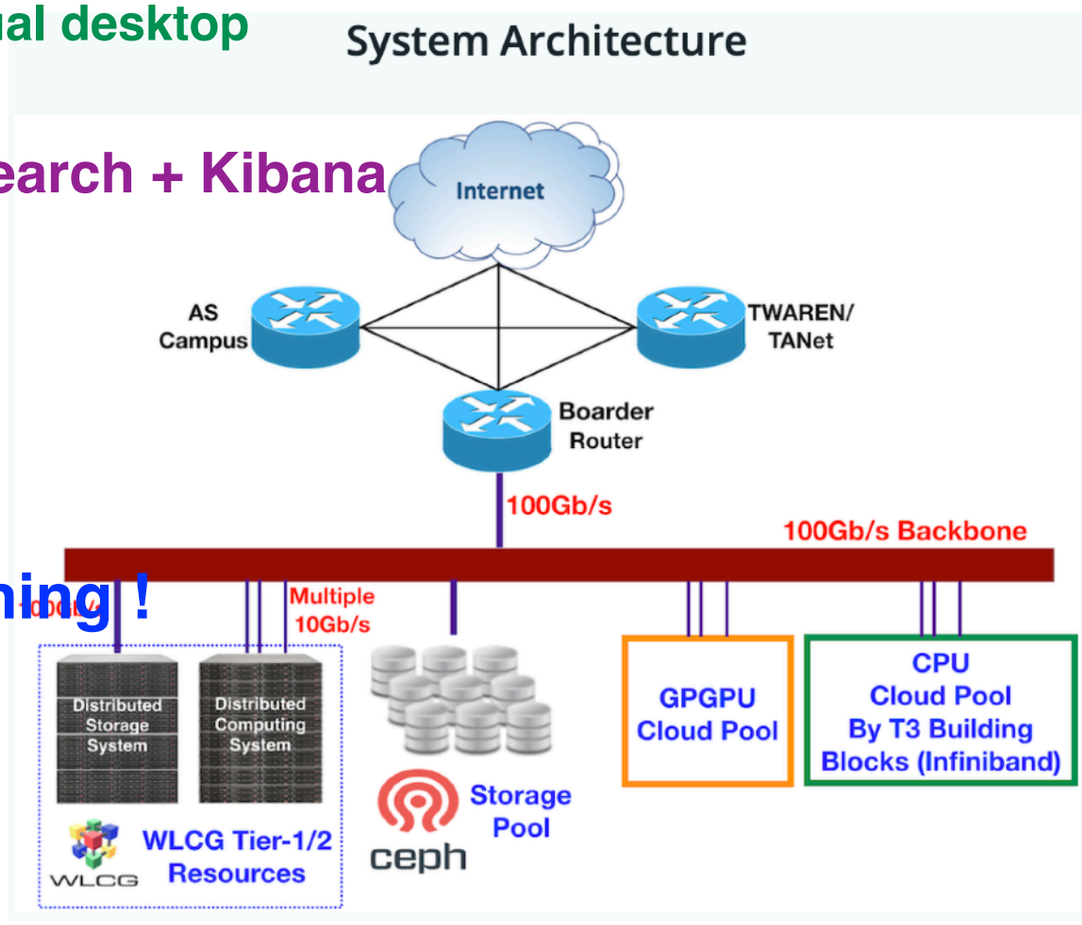

Resource Federation (Distributed Cloud)

- Building Distributed Cloud Infrastructure Supporting broader Scientific Applications based on WLCG
- Integrate the whole data analysis pipeline, develop web portal/science gateway, and optimize system efficiency
- Facilitate GPU computing for big data analytics through DiCOS: >100K GPUCard-Day used in 2020
- Computing model, system architecture and services, solution and technology are continuously improved by user experiences and advanced ICT
- Flexible virtual cluster over distributed heterogeneous resources
  - GPU, CPU with/without infiniband
  - Shared filesystem/storage by Ceph
  - Job scheduler through Slurm, HTCondor and Kubernetes (with containers)
  - Containerization of DiCOS core components: analysis pipeline robustness; portability; maintainability;
- Disciplines: AMS, TEXONO, Gravitational Wave(KAGRA, LIGO, IGWN), NGS, CryoEM, BioSAXS, Drug Discovery, Earth Science, Environmental Changes, Biodiversity and Ecological Monitoring, Lattice Gauge Theory, Condense Matter, proton therapy, RoseTTAFold and ML/DL applications.

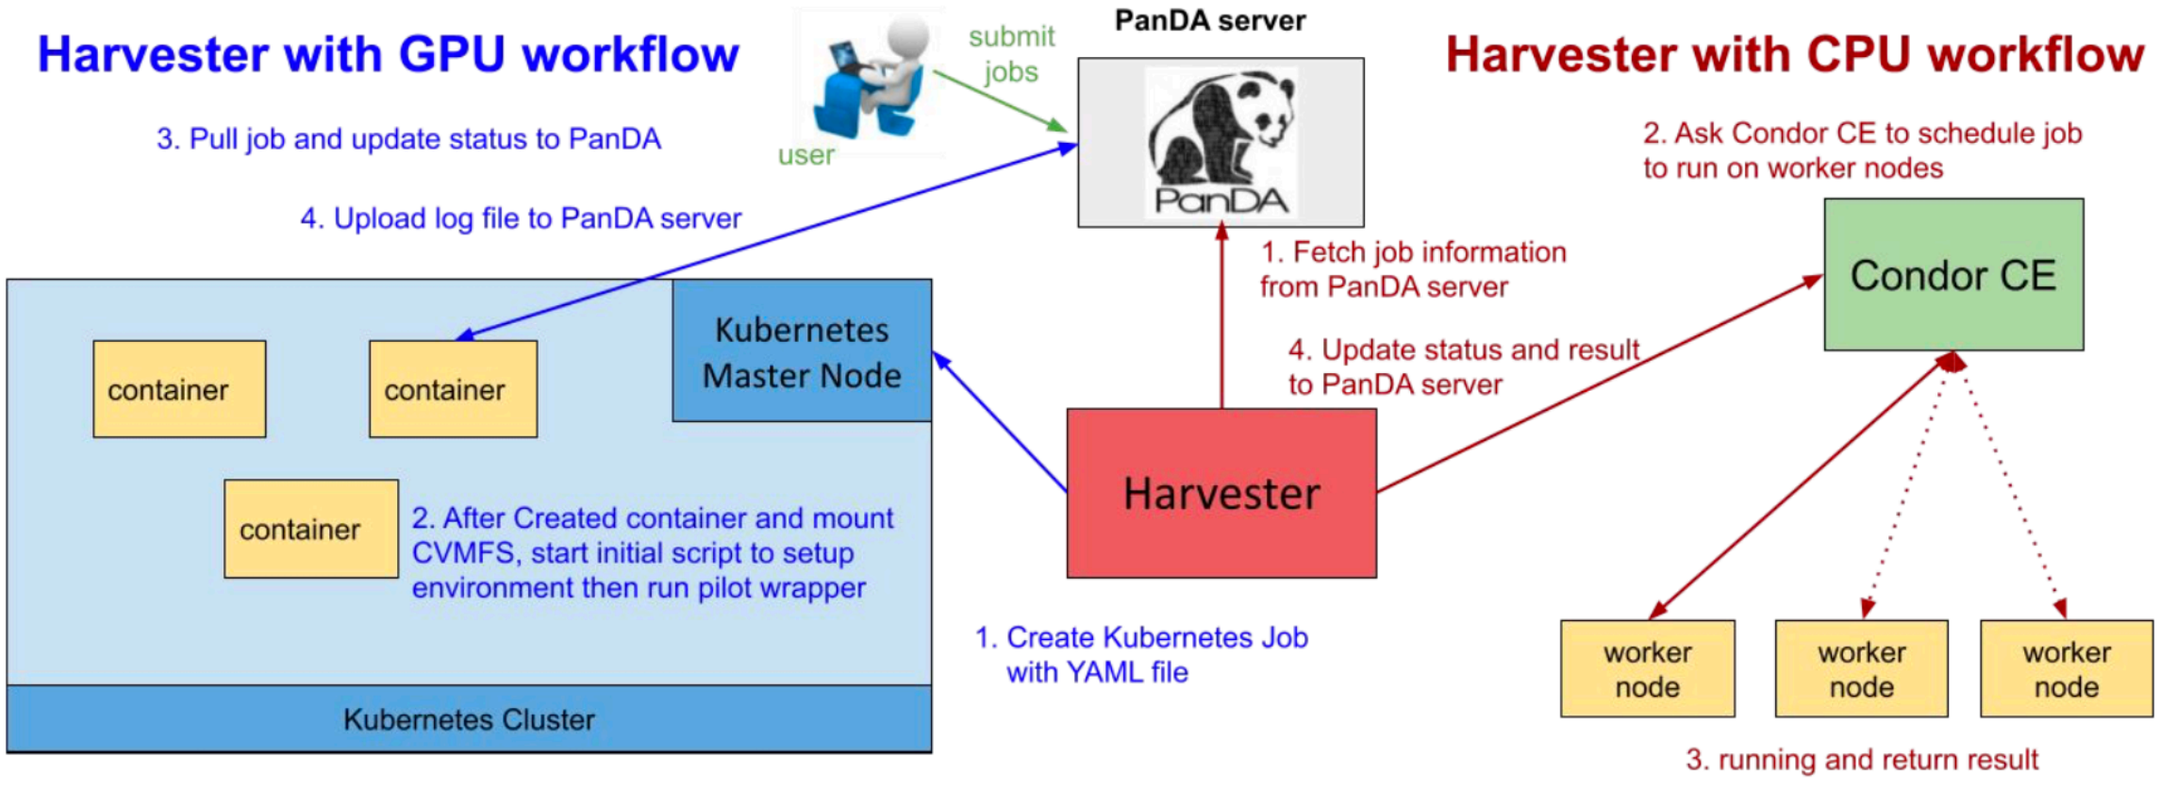

Supporting Cryo-EM Applications by DiCOS

- Primary GPU (single precision) and CPU (300-600 threads) users, ~5-10TB input/job
- Contributions of ASGC for now
  - Web UI development
  - SW package as container
  - Data flow and performance optimization
  - Mass production over DiCOS

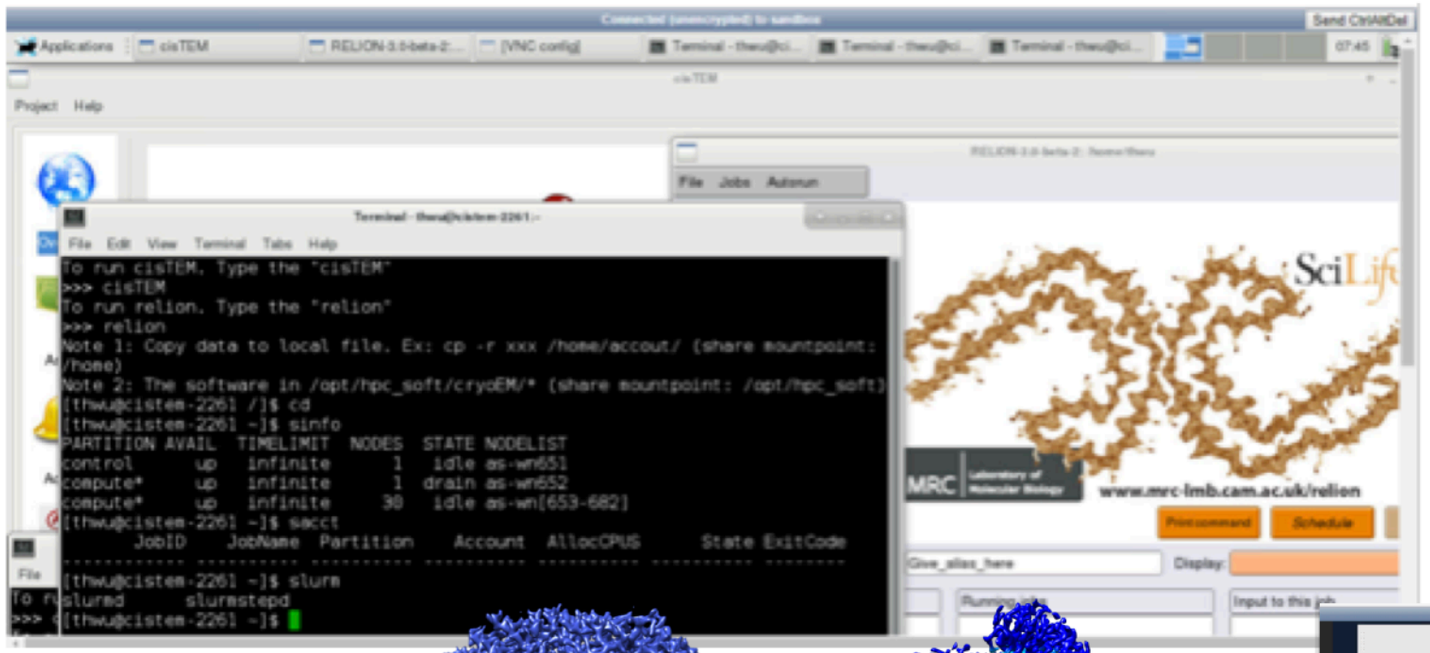

cisTEM

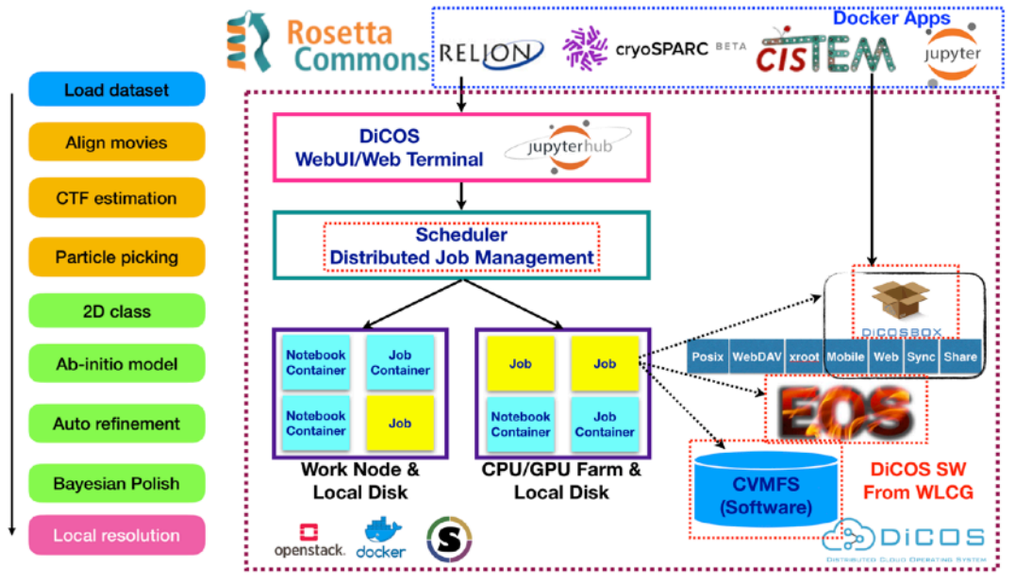

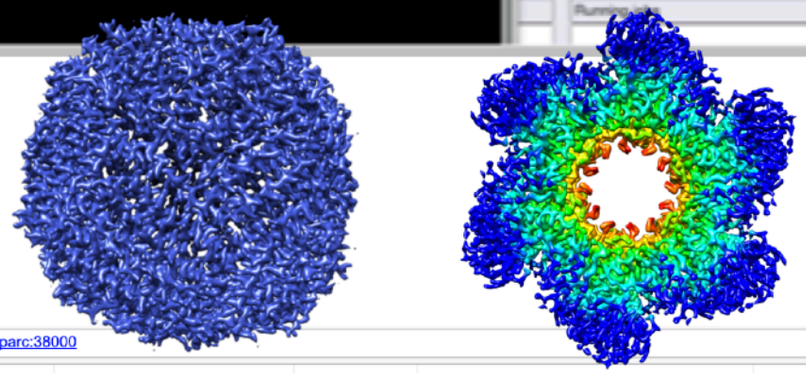

cryoSPARC

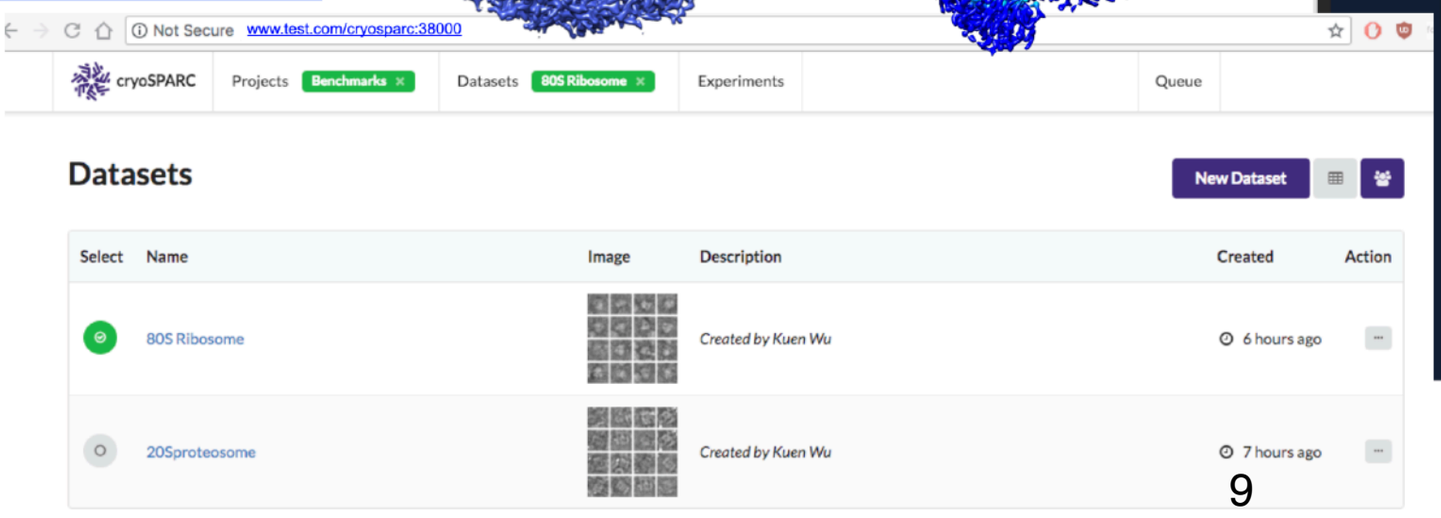

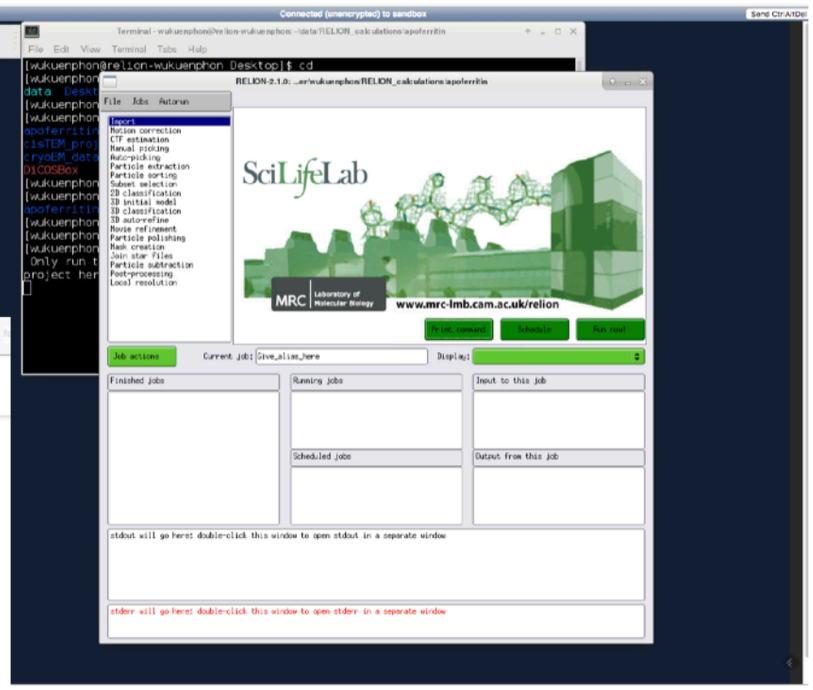

Relion

DiCOS

Distributed Cloud Operating System

System Efficiency Optimization

- Goals: maximize application performance by available resources dynamically, in terms of power, thermal and system (Comp, Storage, Network, application) efficiency
- Scope: Power, Thermal and Distributed Cloud System management
- Strategy: intelligent monitory and control assisted by ML
- Example: Thermal management, Compute/storage/network anomaly detection, Power saving of work nodes
- AHU monitoring and control
  - Detection of refrigerant operating issues and abnormal components; Efficiency optimization
  - 13 sensors; 18K data points/day;
  - Realtime monitoring, adjustment and diagnostics: refrigerant operating issue; abnormal components detection; efficiency tuning; ML-based automatic detection of critical problems;
- System Anomaly Detection
  - Classify machine status into 5 clusters daily: based on CPU-user, CPU-wio, CPU-system, CPU-idle, Network In/Out
  - >30M records/day from all systems of ASGC are covered
  - 146 events in 14 types identified during March 2020 - March 2021

AHU Performance Monitoring

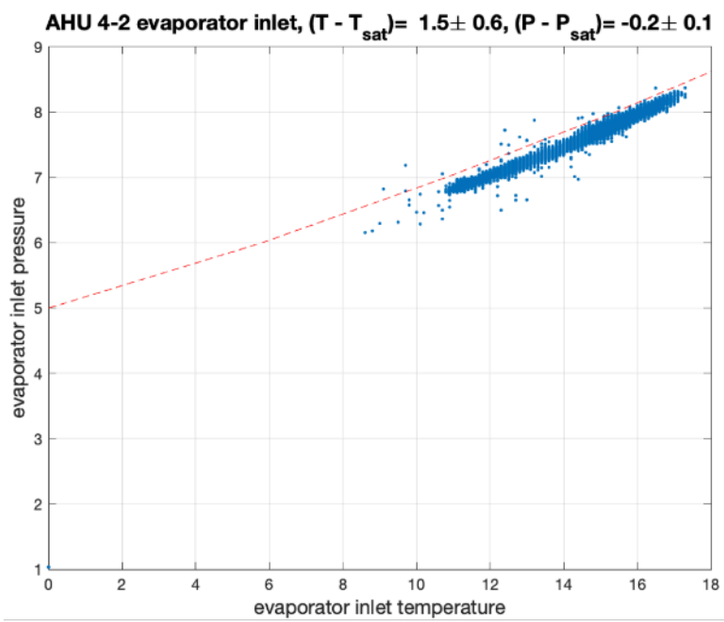

Worknode Monitoring: High ratio of WaitIO & System

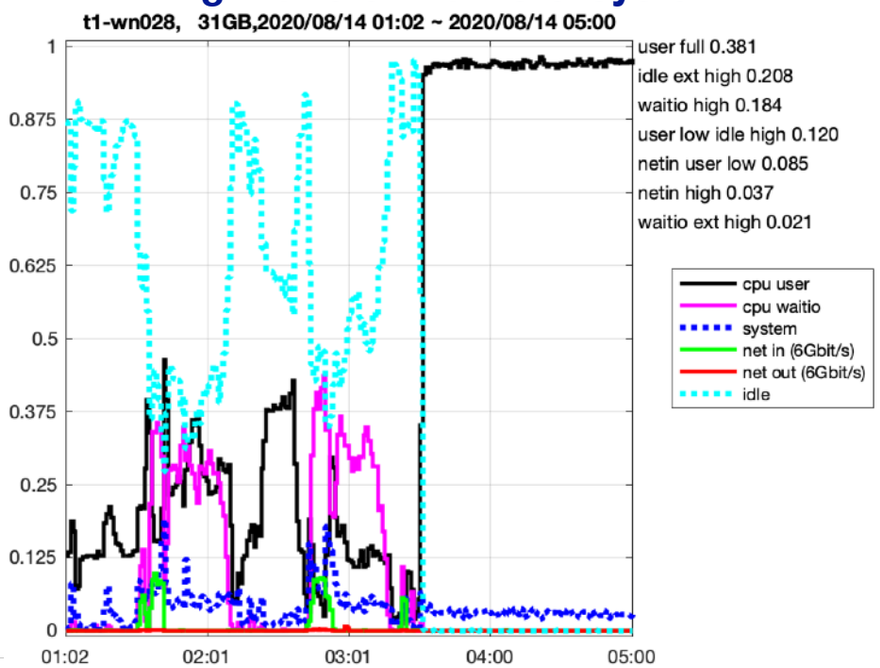

Worknode Monitoring: Misconfigured hyper-threading

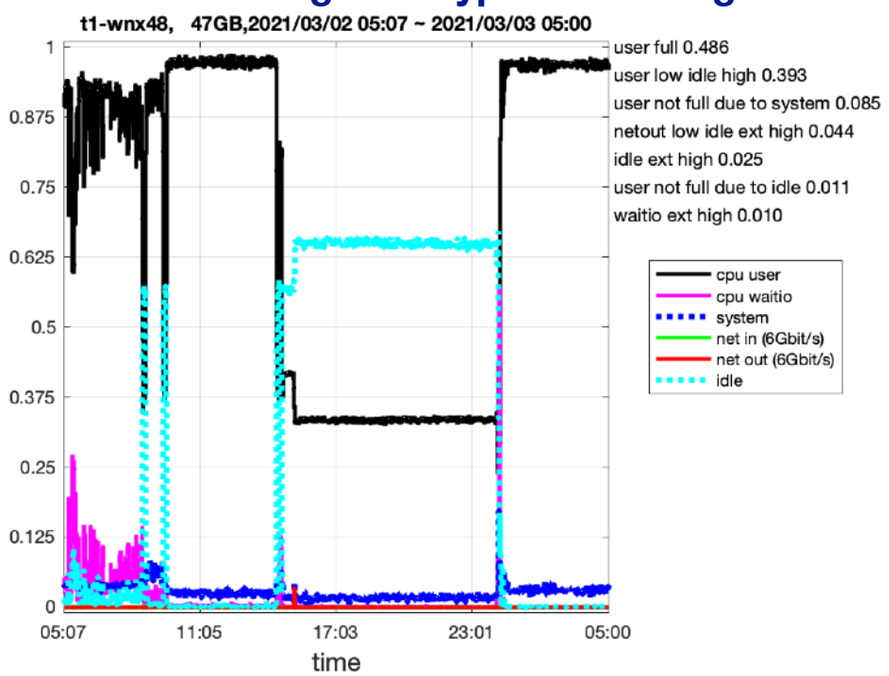

Summary

- Data-oriented, science-driven, ML-enabled collaboration framework over distributed cloud
- Services and supported technologies are driven by research requirements
- Reliability and efficiency are the key of science cloud
- Moving towards open science cloud for regional collaborations, supported by the collaboration with EGI-ACE
  - FAIR principle;
  - Reproducible, open source, open collaboration
- Challenges
  - Automation everything
  - Configuration management
  - Long-term data management
  - Monitoring and intelligent control
  - ML-enabled anomaly detection and efficiency optimization