



ASGC Site Report


Eric Yen and Felix Lee

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

**Asia Tier Centre Forum (ATCF8)
TIFR, Mumbai, India
3 Sep. 2024**

ASGC Overview

- **Founded for participating WLCG and supporting the research collaborations**
 - Being/serving WLCG Tier-1 center from Dec 2005 to Oct 2023
 - Migrating to WLCG Tier-2 center for ATLAS after Q3 2023
- **ASGC is providing big data analysis and computing services for the R&E communities in Taiwan as a core facility**
 - Funded by both Academia Sinica and National Science and Technology Council
 - Primary scientific collaborations: WLCG (ATLAS, CMS), AMS, Gravitational Wave, ICECube/ Neutrino, EIC, QCD, CryoEM, condense matter, etc.
 - Based on the core technologies of WLCG
- **System efficiency as well as AI-enabled analytics are the new focus**



TAIWAN INSTRUMENTATION AND DETECTOR CONSORTIUM (TIDC)



Neutrino/MHEP

CryoEM

Quantum Materials
Physics

Bioimaging

Physics of Active &
Living Matter

Drug Discovery

Astrophysics

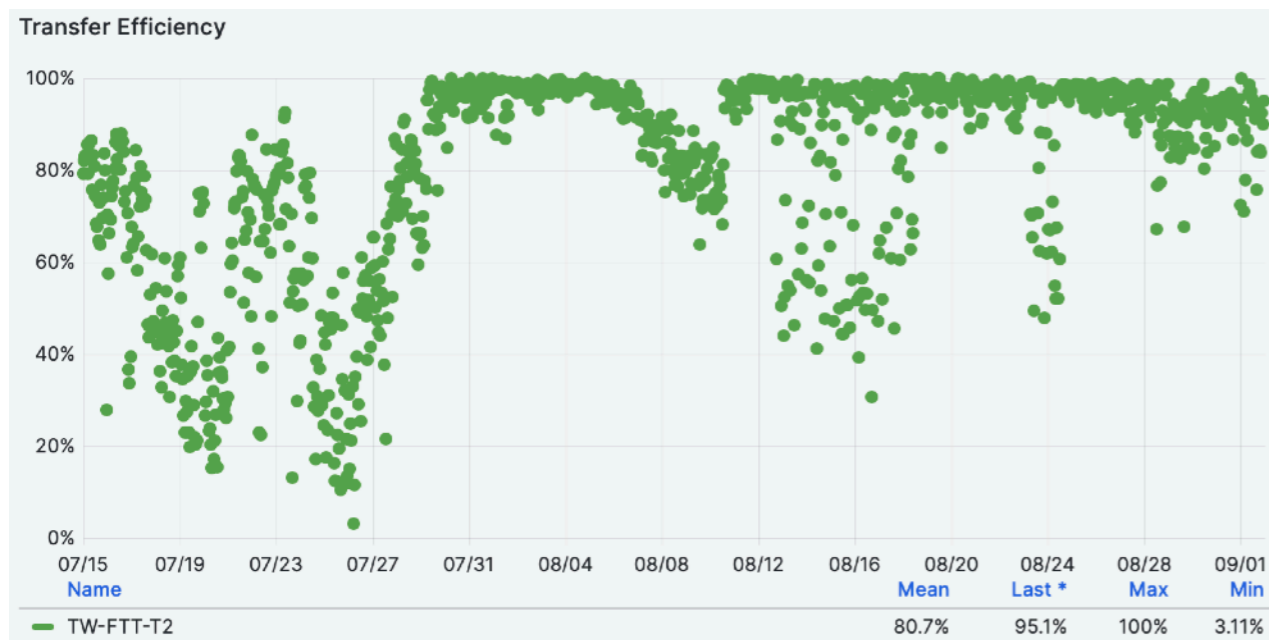
Computational
Chemistry

Earth Science

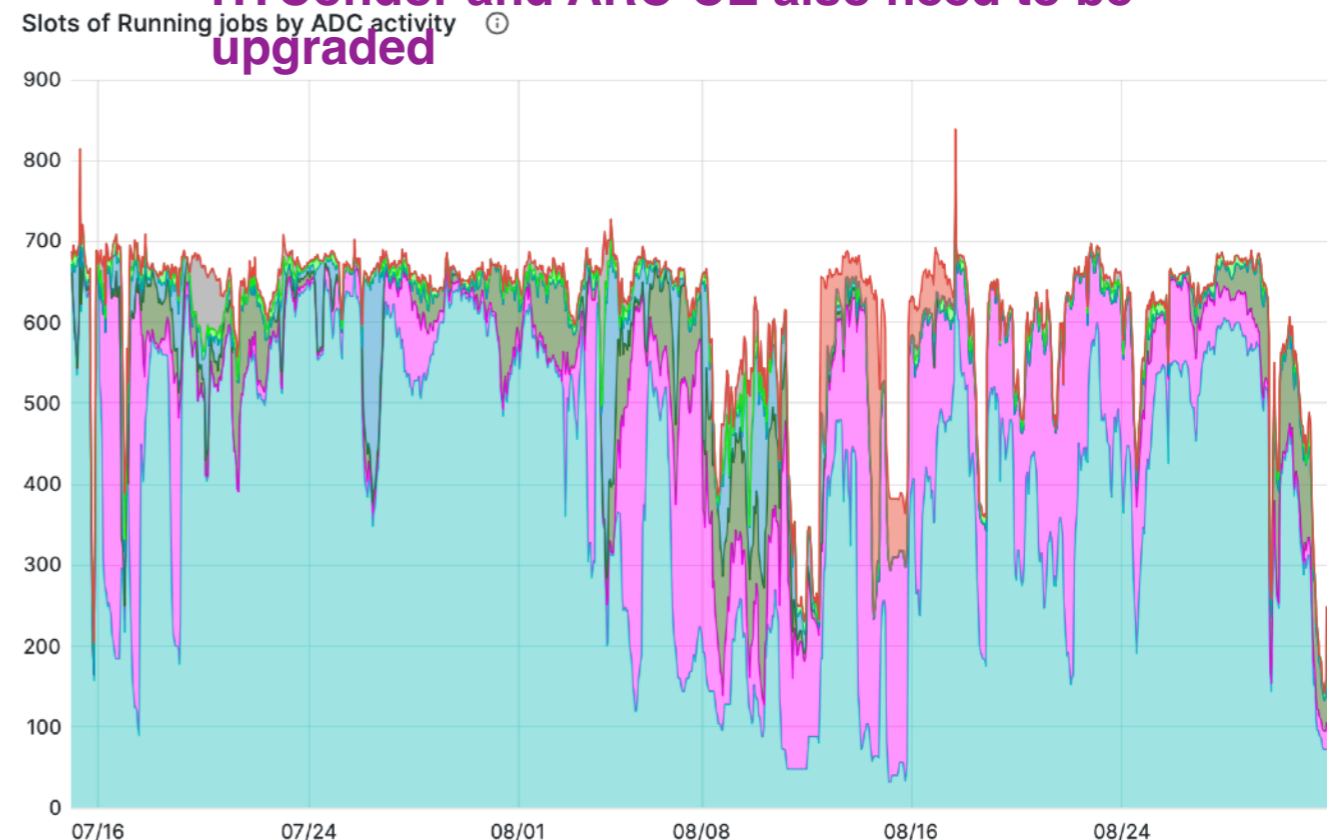
Biodiversity &
Ecology

WLCG Tier2 For ATLAS

- Site has been in production after helpful and careful investigation in July 2024
 - Connection with LHCONE was reestablished, efficiency was also affirmed
 - Local storage RSE was reconfigured and CRIC updated accordingly
 - Passed HC test jobs on 10 July, verified by ADC
- Supporting MC Simulation jobs and analysis jobs mainly afterwards
- Current Status:
 - Transfer efficiency: > 80%
 - CPU Efficiency: > 82%
 - Storage used: > 1PB
 - Issue: limited network bandwidth (< 10Gbps)
- Pledge in 2024
 - CPU: 30K HEPscore23

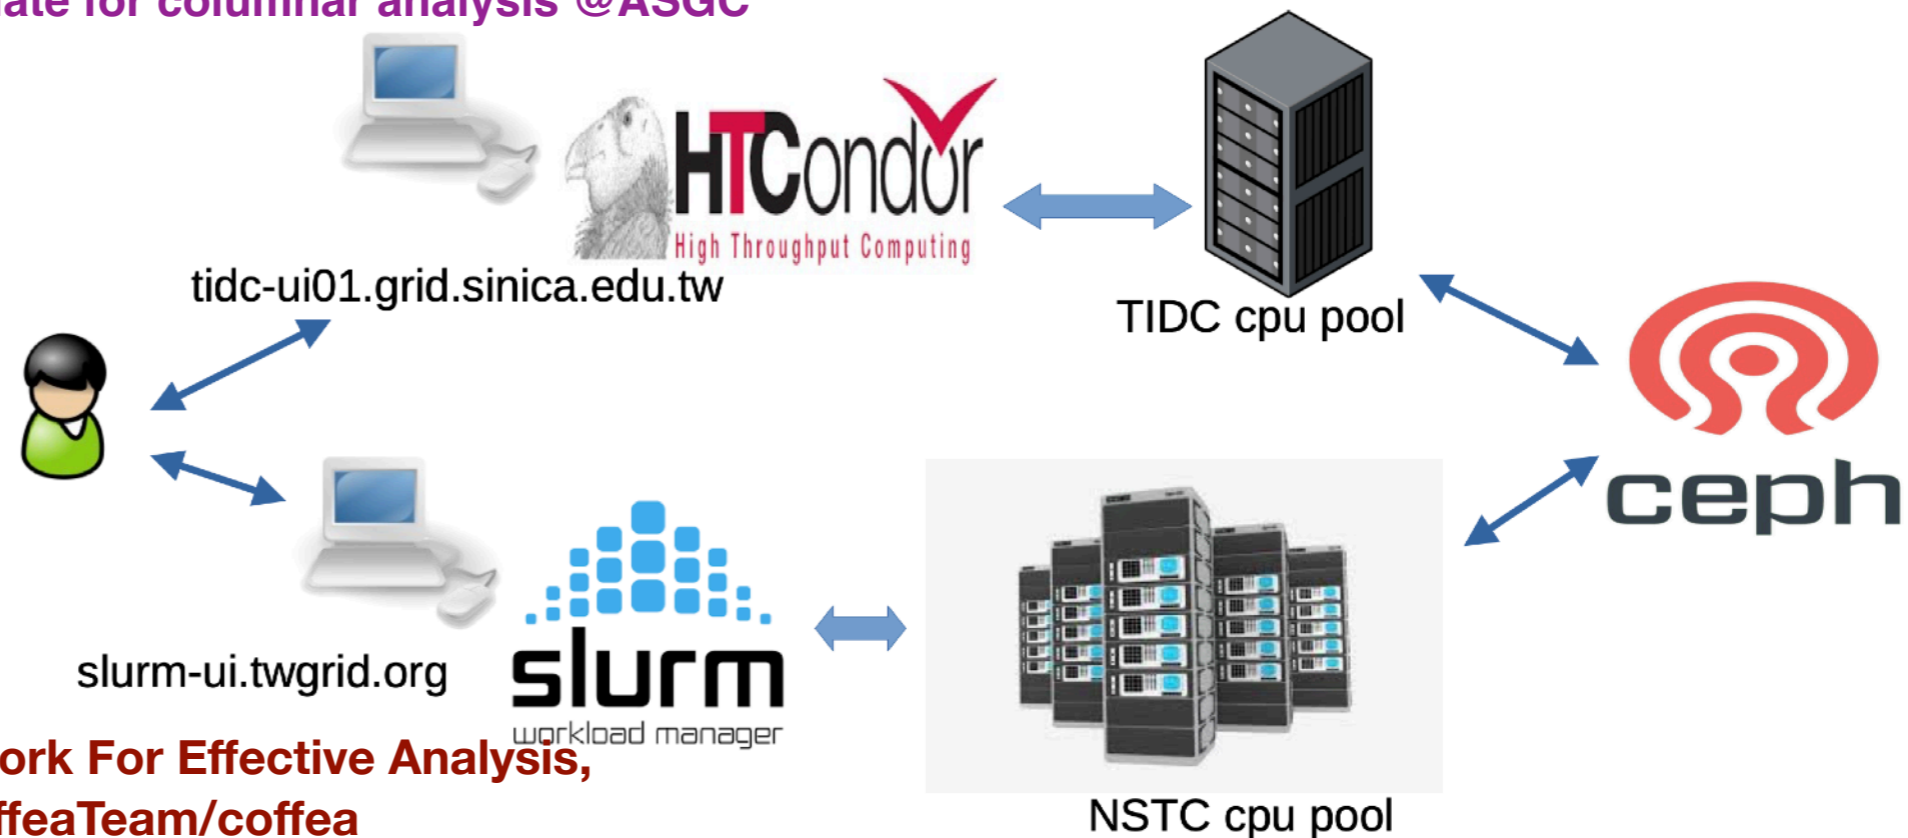


- Disk: 5 PB, managed by EOS
- Job slots - 2,208
 - ~ 839 job slots are available now because of OS migration
- Next Step:
 - Making better use of ASGC resource, e.g, pre-stage for analysis jobs
 - Will take USATLAS operation shift in Asia time zone
- Migration to AlmaLinux 9 - will be accomplished before end of 2024
 - 2-stage: ~ 1,370 job slots in AlmaLinux will be online in Sep.
 - HTCondor and ARC-CE also need to be upgraded



CMS Tier2/Tier3

- **CMS T2 Pledge in 2024 (operated by NCHC)**
 - CPU: 5K HEPscore23
 - Disk: 500TB
- **CMS T3 is managed by ASGC from 2022, in collaboration with TIDC and local CMS groups (NTU and NCU)**
- **Analysis facility**
 - **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




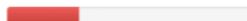

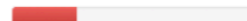







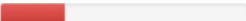

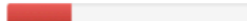

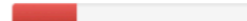











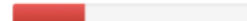



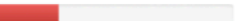

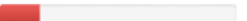

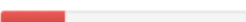
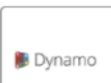
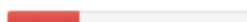

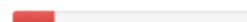



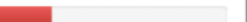

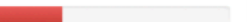




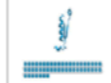
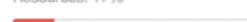







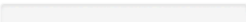











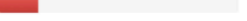




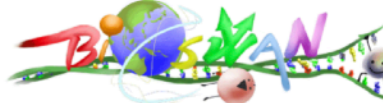






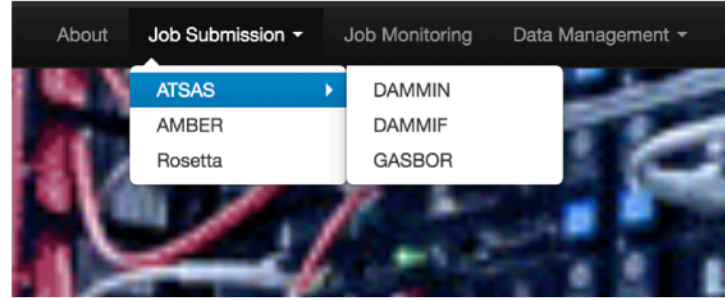






Coffea = Column Object Framework For Effective Analysis,
<https://github.com/CoffeaTeam/coffea>

WLCG-Based Common Infrastructure for Scientific Computing

- **PanDA + RUCIO** serve as the core of the common distributed infrastructure
 - Federation of distributed institute resources
 - Federation of core facilities (including CryoEM, NSRRC, MRI, 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**
- **Resource status: 4.24M CPU jobs, and 44K GPU Jobs in 2024 (till end Aug)**
 - 5,152 CPU Cores + 2,176 CPU Cores (by end 2024)
 - 24x A100 GPU Boards, 16x4090, 48xV100, 56x3090: 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 12PB LTO9 tape storage will be available in Q4 2024 - dark data, backup
- **Both core technology and application platform** are evolving with user experiences
- **OSG software stack** has been deployed for IGWN and will be part of the infrastructure

60+ Web Applications Provided

| | | | | | |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  CryoSPARC 32 P100 Version: 3.2 Resources: 83%  <input type="button" value="Launch"/> |  CryoSPARC 1080ti Version: 3.3.2 Resources: 29%  <input type="button" value="Launch"/> |  CryoSPARC RTX3090 Version: 3.3.2 Resources: 29%  <input type="button" value="Launch"/> |  spyder cpu/eman2 Version: Resources: 17%  <input type="button" value="Launch"/> |  Octave Version: V5.2 Resources: 29%  <input type="button" value="Launch"/> |  Transfer Data Version: Resources: 86%  <input type="button" value="Launch"/> |
|  CryoSPARC RTX3090 Version: 4.0.2 Resources: 26%  <input type="button" value="Launch"/> |  CryoSPARC RTX3090 Version: 4.1.1 Resources: 26%  <input type="button" value="Launch"/> |  CryoSPARC RTX3090 Version: 4.4.1 Resources: 26%  <input type="button" value="Launch"/> |  cisTEM Version: Resources: 86%  <input type="button" value="Launch"/> |  Ovito Version: Resources: 86%  <input type="button" value="Launch"/> |  OpenACC Version: GPU P100 Resources: 83%  <input type="button" value="Launch"/> |
|  AlphaFold Version: GPU with A100 Resources: 75%  <input type="button" value="Launch"/> |  AlphaFold (Full DB) Version: GPU with A100 Resources: 75%  <input type="button" value="Launch"/> |  IMOD Version: GPU with 1080ti Resources: 29%  <input type="button" value="Launch"/> |  Triton Version: 22.01-py3 (GPU P100) Resources: 83%  <input type="button" value="Launch"/> |  AlphaFold Version: GPU with RTX3090 Resources: 26%  <input type="button" value="Launch"/> |  AlphaFold Version: GPU with V100 Resources: 17%  <input type="button" value="Launch"/> |
|  RoseTTAFold Version: GPU with rtx3090 Resources: 26%  <input type="button" value="Launch"/> |  Dynamo Version: GPU with 1080ti Resources: 29%  <input type="button" value="Launch"/> |  MATLAB Version: R2018b on GPU V00 Resources: 17%  <input type="button" value="Launch"/> |  Jupyter Lab Version: CPU with Tensorflow v1 Resources: 86%  <input type="button" value="Launch"/> |  Jupyter Lab gpu 3090 Version: GPU with Tensorflow 3090 Resources: 26%  <input type="button" value="Launch"/> |  Jupyter Lab GPU 1080ti Version: GPU with Tensorflow v2 Resources: 29%  <input type="button" value="Launch"/> |
|  RFDIFFUSION Version: 2023 on GPU V00 Resources: 17%  <input type="button" value="Launch"/> |  diffdock Version: 2023 on GPU V00 Resources: 17%  <input type="button" value="Launch"/> |  EvoDiff Version: V100 Resources: 17%  <input type="button" value="Launch"/> |  Jupyter Lab GPU V100 Version: GPU with Tensorflow V100 Resources: 17%  <input type="button" value="Launch"/> |  Jupyter Lab GPU A100 Version: GPU with Tensorflow A100 Resources: 75%  <input type="button" value="Launch"/> |  Jupyter Lab Cryocare GPU Version: GPU with 1080ti Resources: 29%  <input type="button" value="Launch"/> |
|  QIIME2 Version: Genome Resources: %  <input type="button" value="Launch"/> |  Scipion3 Version: P100 Resources: 83%  <input type="button" value="Launch"/> |  Phenix Version: Resources: 86%  <input type="button" value="Launch"/> |  Jupyter Lab GPU A100 Version: GPU with Tensorflow v2.6 Resources: 75%  <input type="button" value="Launch"/> | <ul style="list-style-type: none"> • Web Portal • Application over Cloud • Jupyterlab • Web Terminal | |
|  MorphoGraphX Version: GPU with P100 Resources: 83%  <input type="button" value="Launch"/> |  Deepmd-kit Version: GPU with A100 Resources: 75%  <input type="button" value="Launch"/> |  Deepmd-kit Version: GPU with V100 Resources: 17%  <input type="button" value="Launch"/> |  MAML Version: GPU with A100 Resources: 75%  <input type="button" value="Launch"/> |  LabVIEW Run-Time Engine Version: 2019  <input type="button" value="Launch"/> |  |
|  MAML Version: GPU with V100 Resources: 17%  <input type="button" value="Launch"/> |  PVserver Version: 5.8.0 (GPU 1080Ti) Resources: 29%  <input type="button" value="Launch"/> |  Paraview Client Version: 5.8.0 Resources: 86%  <input type="button" value="Launch"/> |  <p>DiCOS-BioSAXS Platform</p> <p>Navigation: About, Job Submission, Job Monitoring, Data Management</p> <p>Application List: ATASAS, AMBER, Rosetta, DAMMIN, DAMMIF, GASBOR</p> | | |
|  PyRoot Version: GPU with 1080ti Resources: 29%  <input type="button" value="Launch"/> |  qiskit Version: Resources: 86%  <input type="button" value="Launch"/> | | | | |

Energy Saving

- **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 some 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 usage achieves around 20% reduction rate achieved per annum in 2023 and 2024 (till end Aug.)**

Regional & International collaborations

- **Asia could make significant contributions to the international collaboration, although network latency and time zone might be issues**
 - Hubs in Asia have been providing multiple 100Gbps links with other continents
 - Regional data lake and caches could cope with latency issue
 - Operation in Asia time zone could make up around the clock collaborations
- **In addition to deployment, participating to core technology development and extending to other e-Science applications could enhance both the technology evolution and also the capacity building.**
 - e.g, Improving site and WLCG efficiency intelligently
- **Active participating experiment/WLCG oriented events and share with regional partners**
 - CHEP, HEPIX, GDB, LHCONE/OPN come to Asia regularly. We also have some routine collaboration events, such as ATCF, ISGC, APAN, etc.
 - Differentiate focus of these events and make them complementary to each other
 - e.g, Summer school in Asia, WLCG Workshop in Asia when CHEP is here.
- **Pilot All-Photonics Network (APN) and IOWN in Asia has been constructed (3,000 KM). Technologies are expected to be mature by 2030**
 - 100x energy efficiency, 100x transmission capacity; 1/200 latency

Welcome To ISGC2025 in Taipei



- **Schedule: 17-21 March 2025**
- **Venue: Academia Sinica, Taipei, Taiwan**
- **Call for Abstract/ Session will be open soon**
- **Keynote speech, Asia Partner Updates, and workshop/session are major components**
 - **Security workshop, sessions of HEP, Grid & Clouds, Networking, infrastructure and AI are typical arrangements. Domain-specific workshop or sessions are also welcomed**
- **Contact: ISGC Secretariat**
 - **vic@twgrid.org**