Supporting Research Computing of IOP & AS

Eric Yen, Jingya You, Rudy Chen

Academia Sinica Grid Computing Centre (ASGC)

PS2023

29 May 2023

ASGC Is Supporting Research Computing of IOP & AS

- **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**

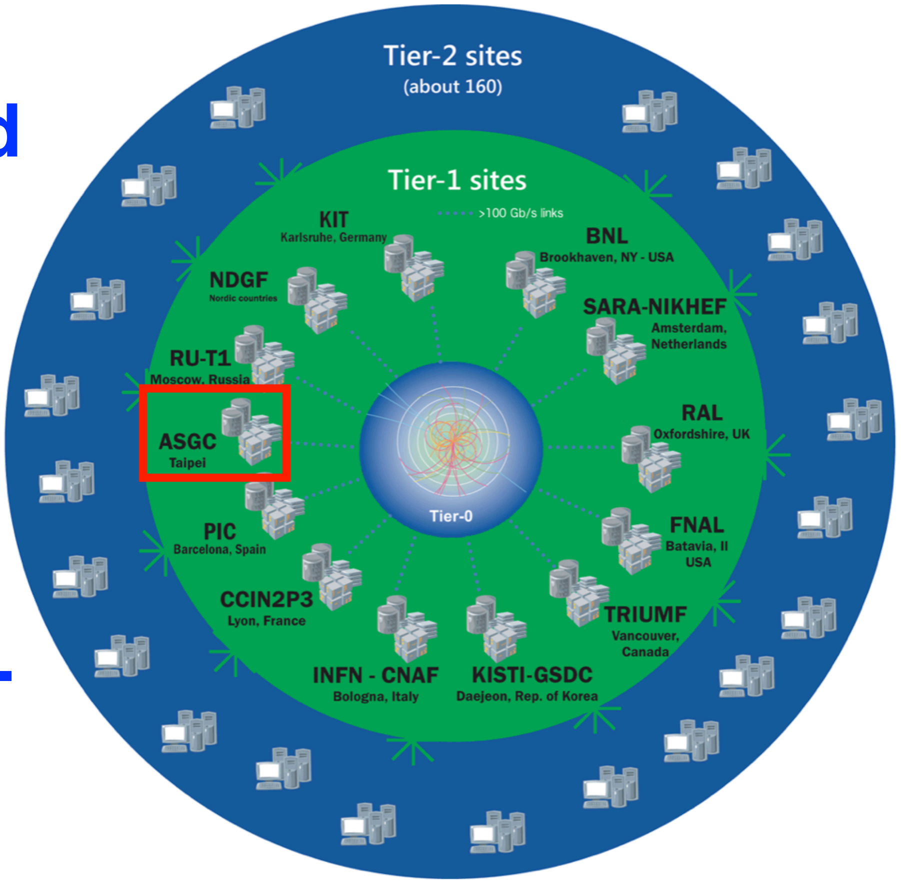

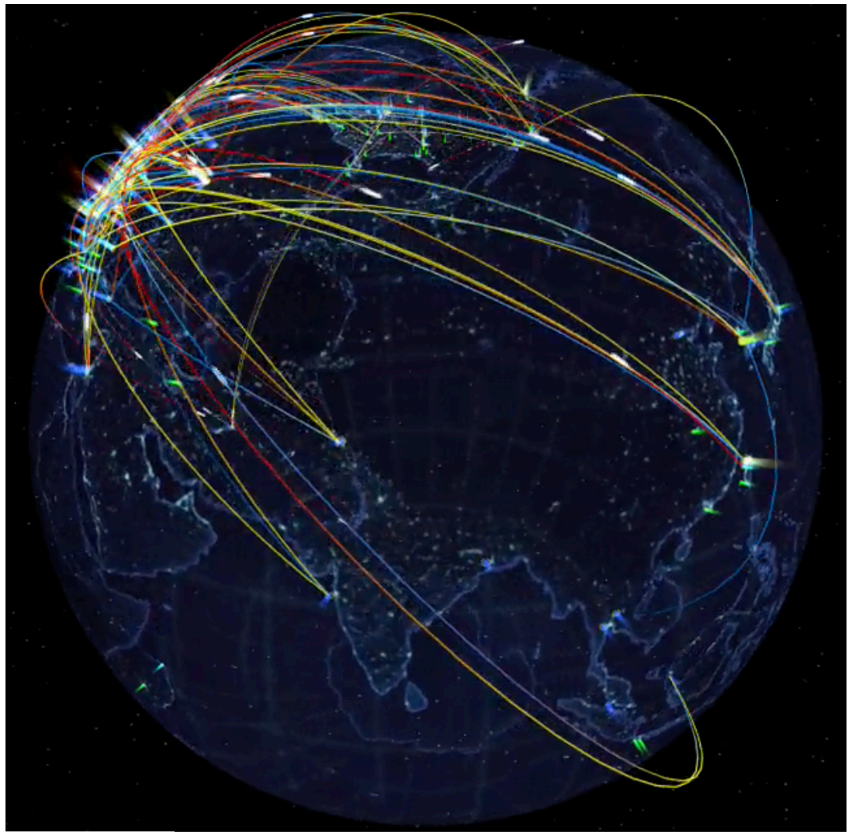

Vision - Accelerating Discovery and Innovation

- Mission: Enabling innovations by integrated research infrastructure - connecting instruments, data, minds, and computing
  - System, service, and advanced ICT R&D + collaborations
- The Fourth Paradigm: Data-Intensive Scientific Discovery (e-Science)
  - Advancing scientific discoveries by progressive computing power/capability on (big) data analytics
  - Alleviate the short-term pain of technological disruption and pave the way for long-term gain
- AI becomes a Game-changer across all disciplines/industries
  - ML-enabled algorithmic advances +
  - (significant increasing) computing power and storage +
  - huge amounts of data
- Research infrastructure & Collaboration - bridging science and computing

ATLAS Experiment: Push Frontiers of Knowledge

- **Transforming raw data from the detector into particles for analysis, with a set direction, energy and type**

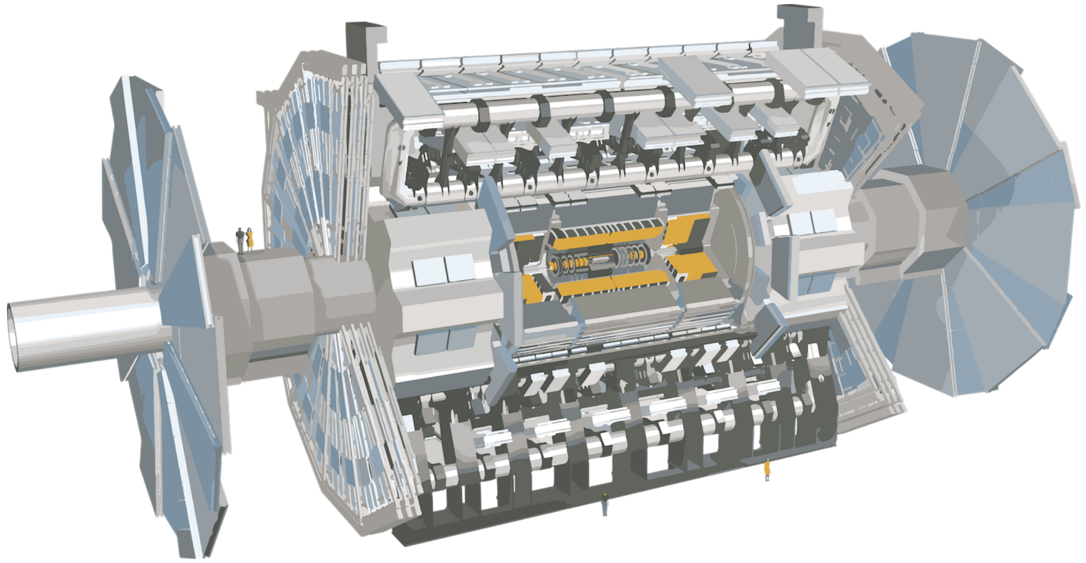

A giant 3D digital camera that can detect charged particles as well as photons

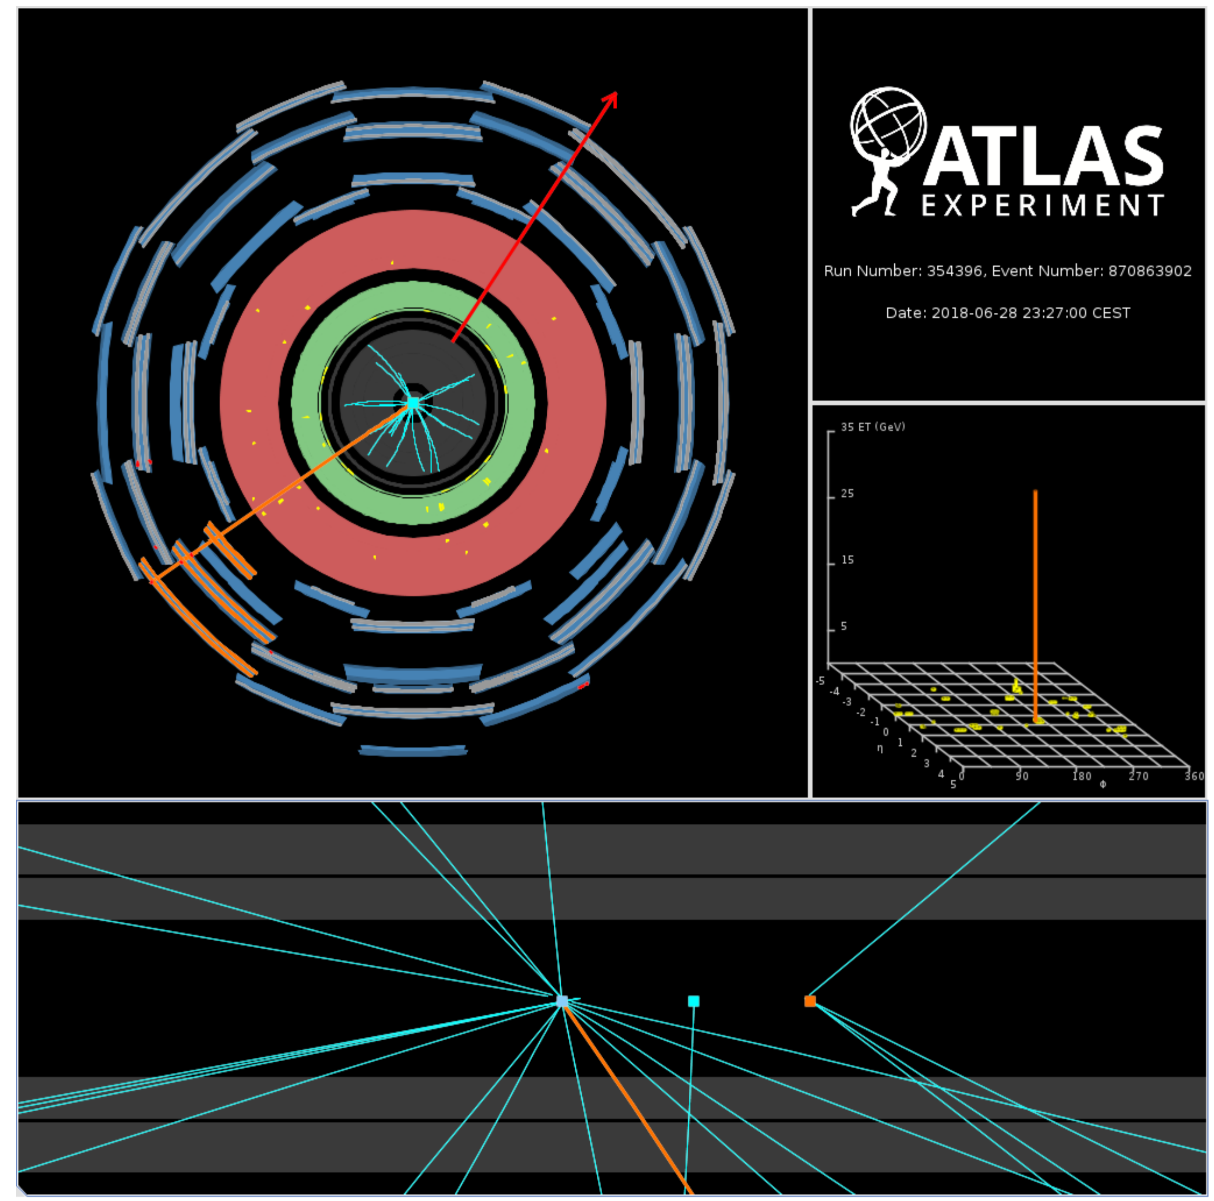

New high-precision measurements of W and Z boson properties - A major milestone towards more precise measurements of the W-boson mass (25 May 2023)

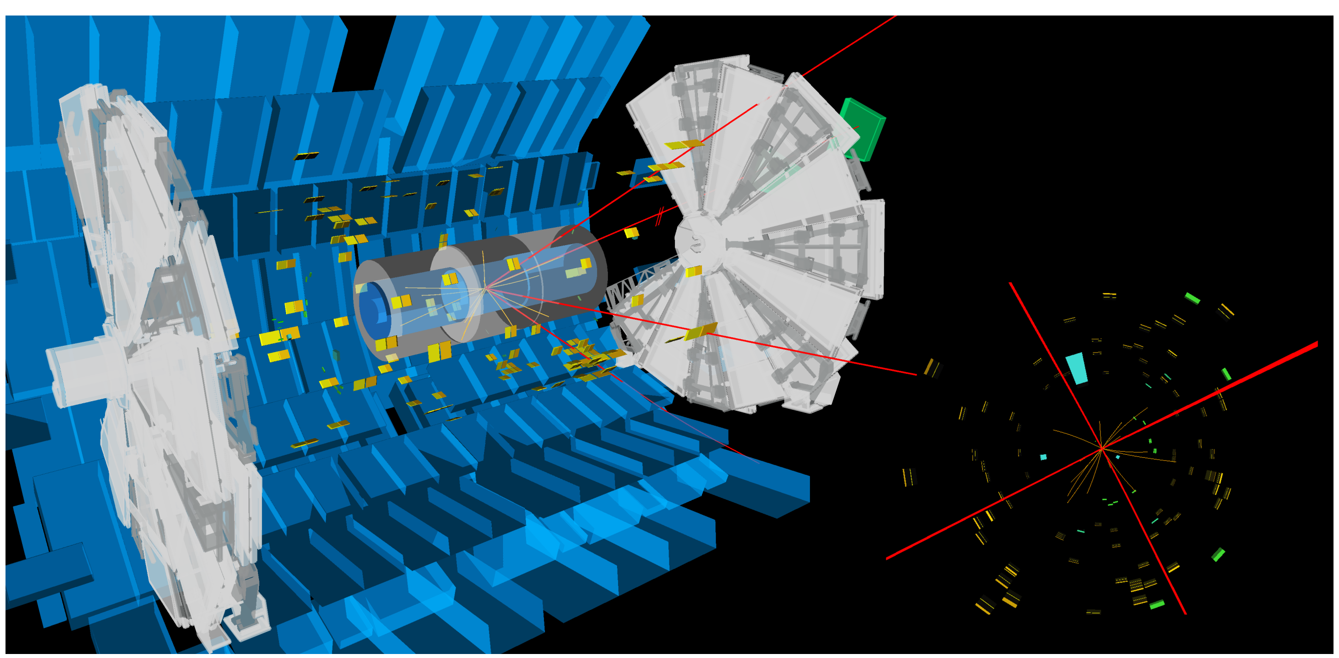

ATLAS measures the Higgs boson at 13.6 TeV (24 May 2023)

Source: ATLAS Experiment, https://atlas.cern

- **Online processing: in real time; decisions irreversible; data cannot be recovered**
- **Trigger: Event selection**
- **Data acquisition**
- **Monitoring**
- **Control**

- **Offline processing: data can be reprocessed**
- **Calibration: convert raw data to physical quantities**
- **Alignment: find out precise detector positions**
- **Event reconstruction**
- **Simulation: event generation; Detector simulation and digitization**
- **Physics analysis**

Data Analysis at the LHC

The process to transform raw data into useful physics datasets

- This is a complicated series of steps at the LHC (Run2)

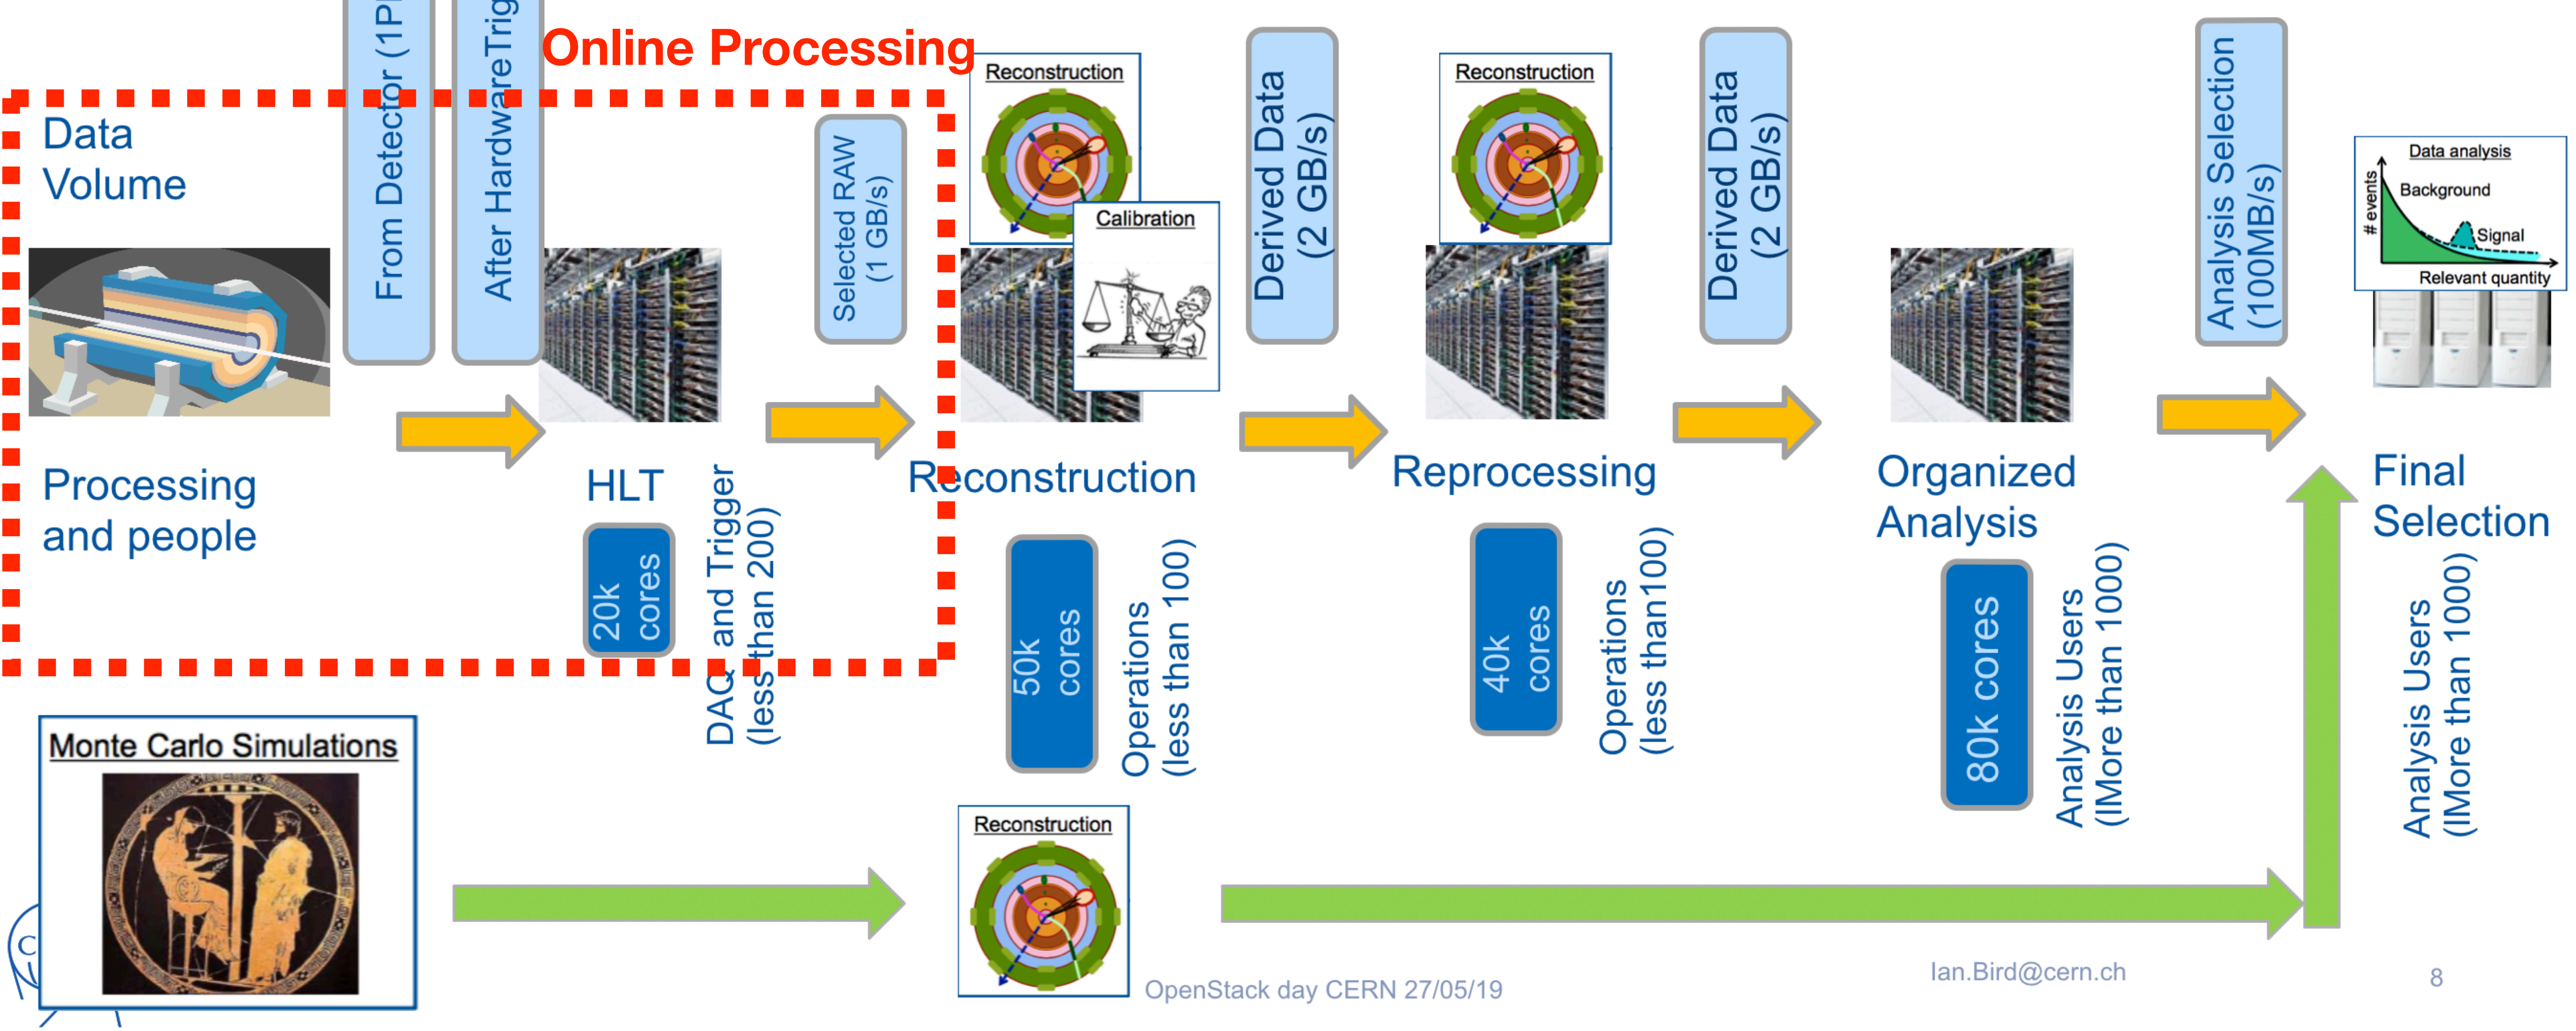

Data Analysis

- **Step1: bulk analysis - done on the grid**
  - Scan over thousands of files, millions/billions events (TB-PB)
  - Select events of interest
  - Apply calibration/corrections
  - Retain only required variables - Save to disk (GB-TB)
- **Step2: final analysis - usually done locally**
  - ML training; background studies; systematics; statistical analysis; …; final plots

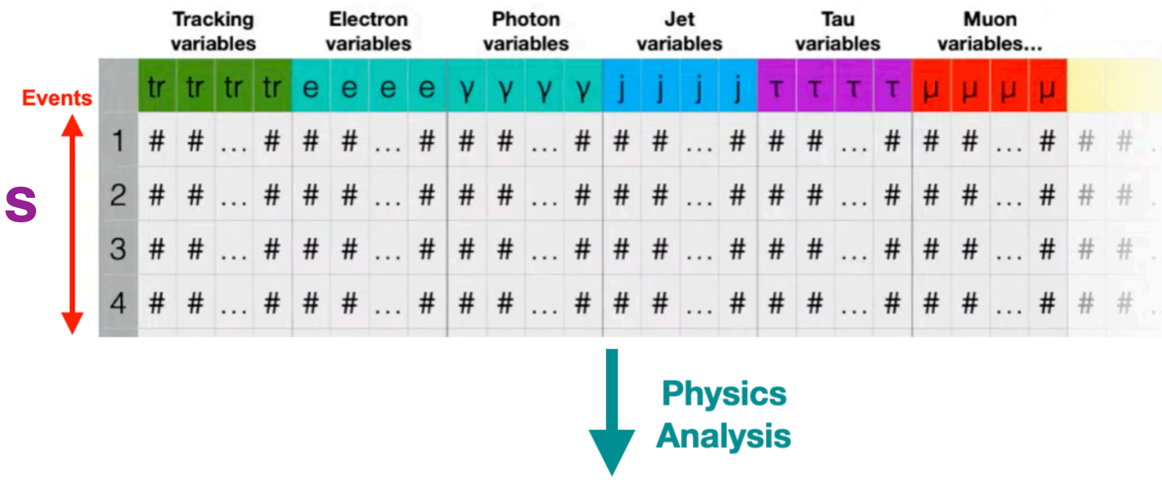

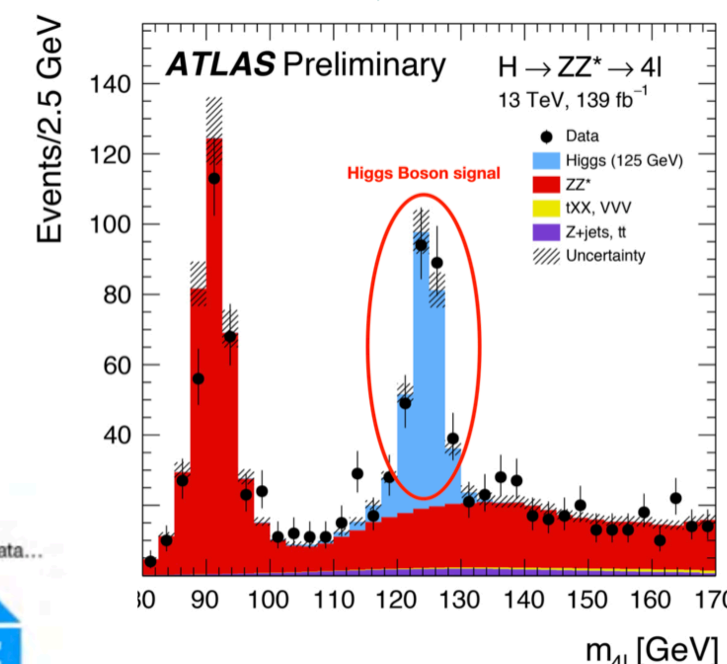

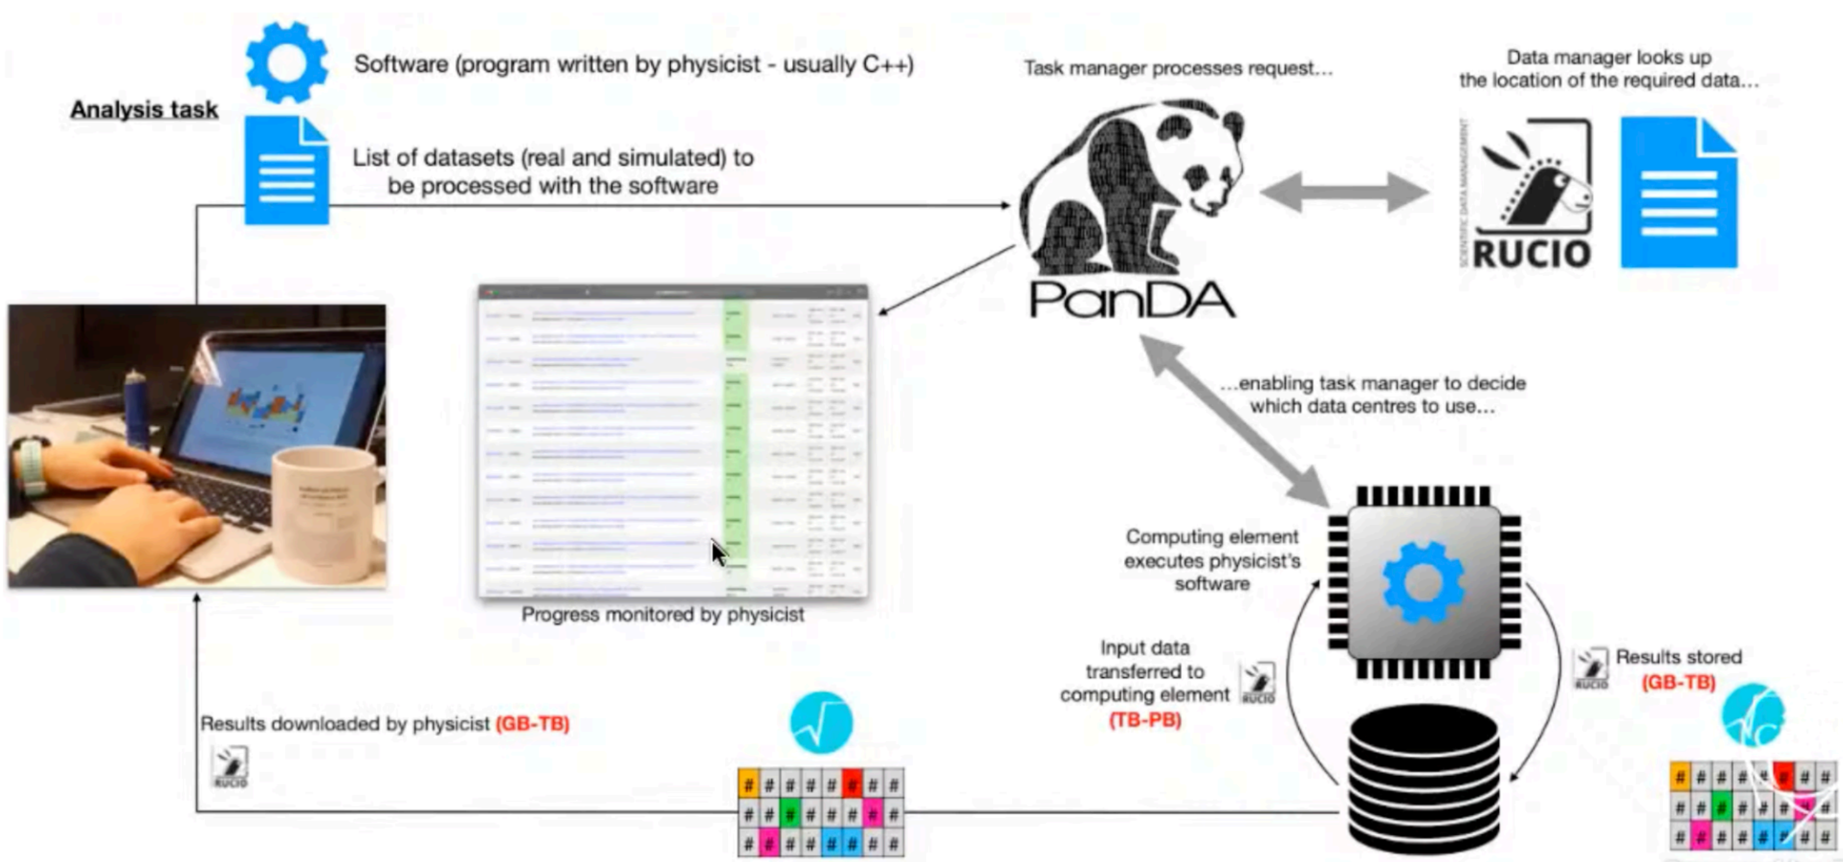

Source: James Catemore

Grid Computing for ATLAS

- Different experiments organize the way they use the grid in different ways
- Save & Spreading the raw data around
- Reconstruction to turn raw data into objects and have the reconstructed data sitting on the grid
- Reconstruct our simulated events
- Physics analysis

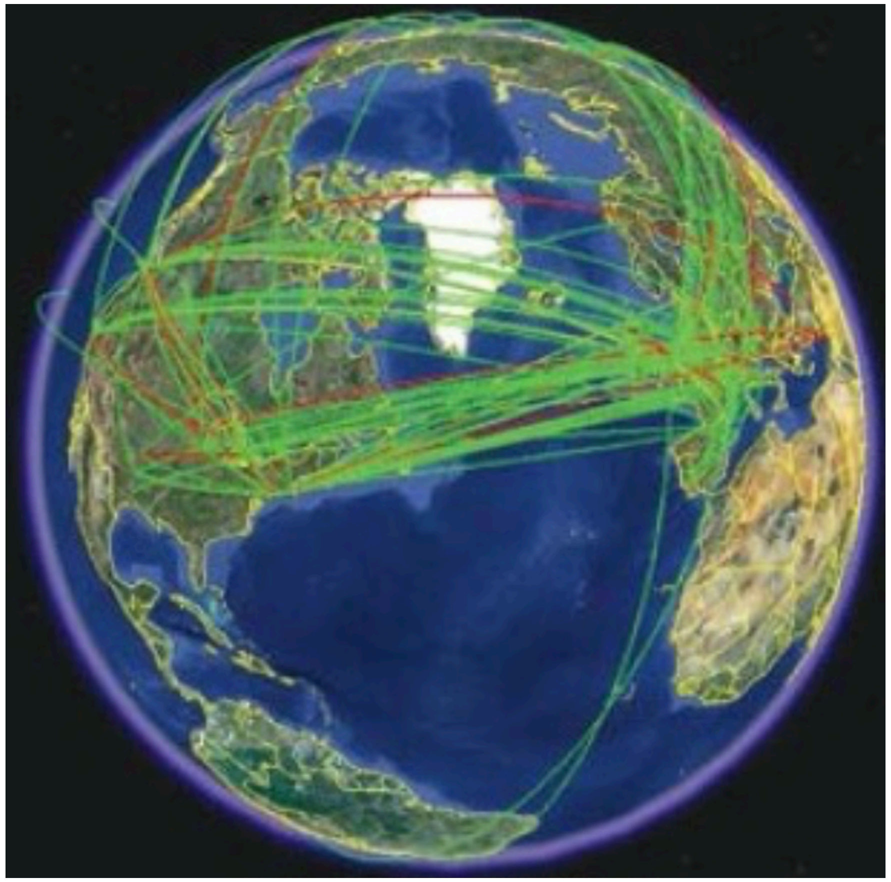

- Collaboration: 170 data centres; 42 countries; shared by the four LHC experiments
- Network: 50GB/s; 50 million files/week
- Computing: 2 million tasks/day; 1 million cores
- Data & Storage: > 1000 PB on disk and tape

Future - Meeting the HL-LHC Challenge

- Make existing software run more efficiently -> optimization
- Use new computational technologies (e.g., GPUs) - migrating existing workflows & developing ML workflows
- Make more use of ML, including for simulation (e.g., generative adversarial techniques) and beyond
- Write smaller data formats (using less data)
- Make more use of tape
- Data-on-demand for analysis

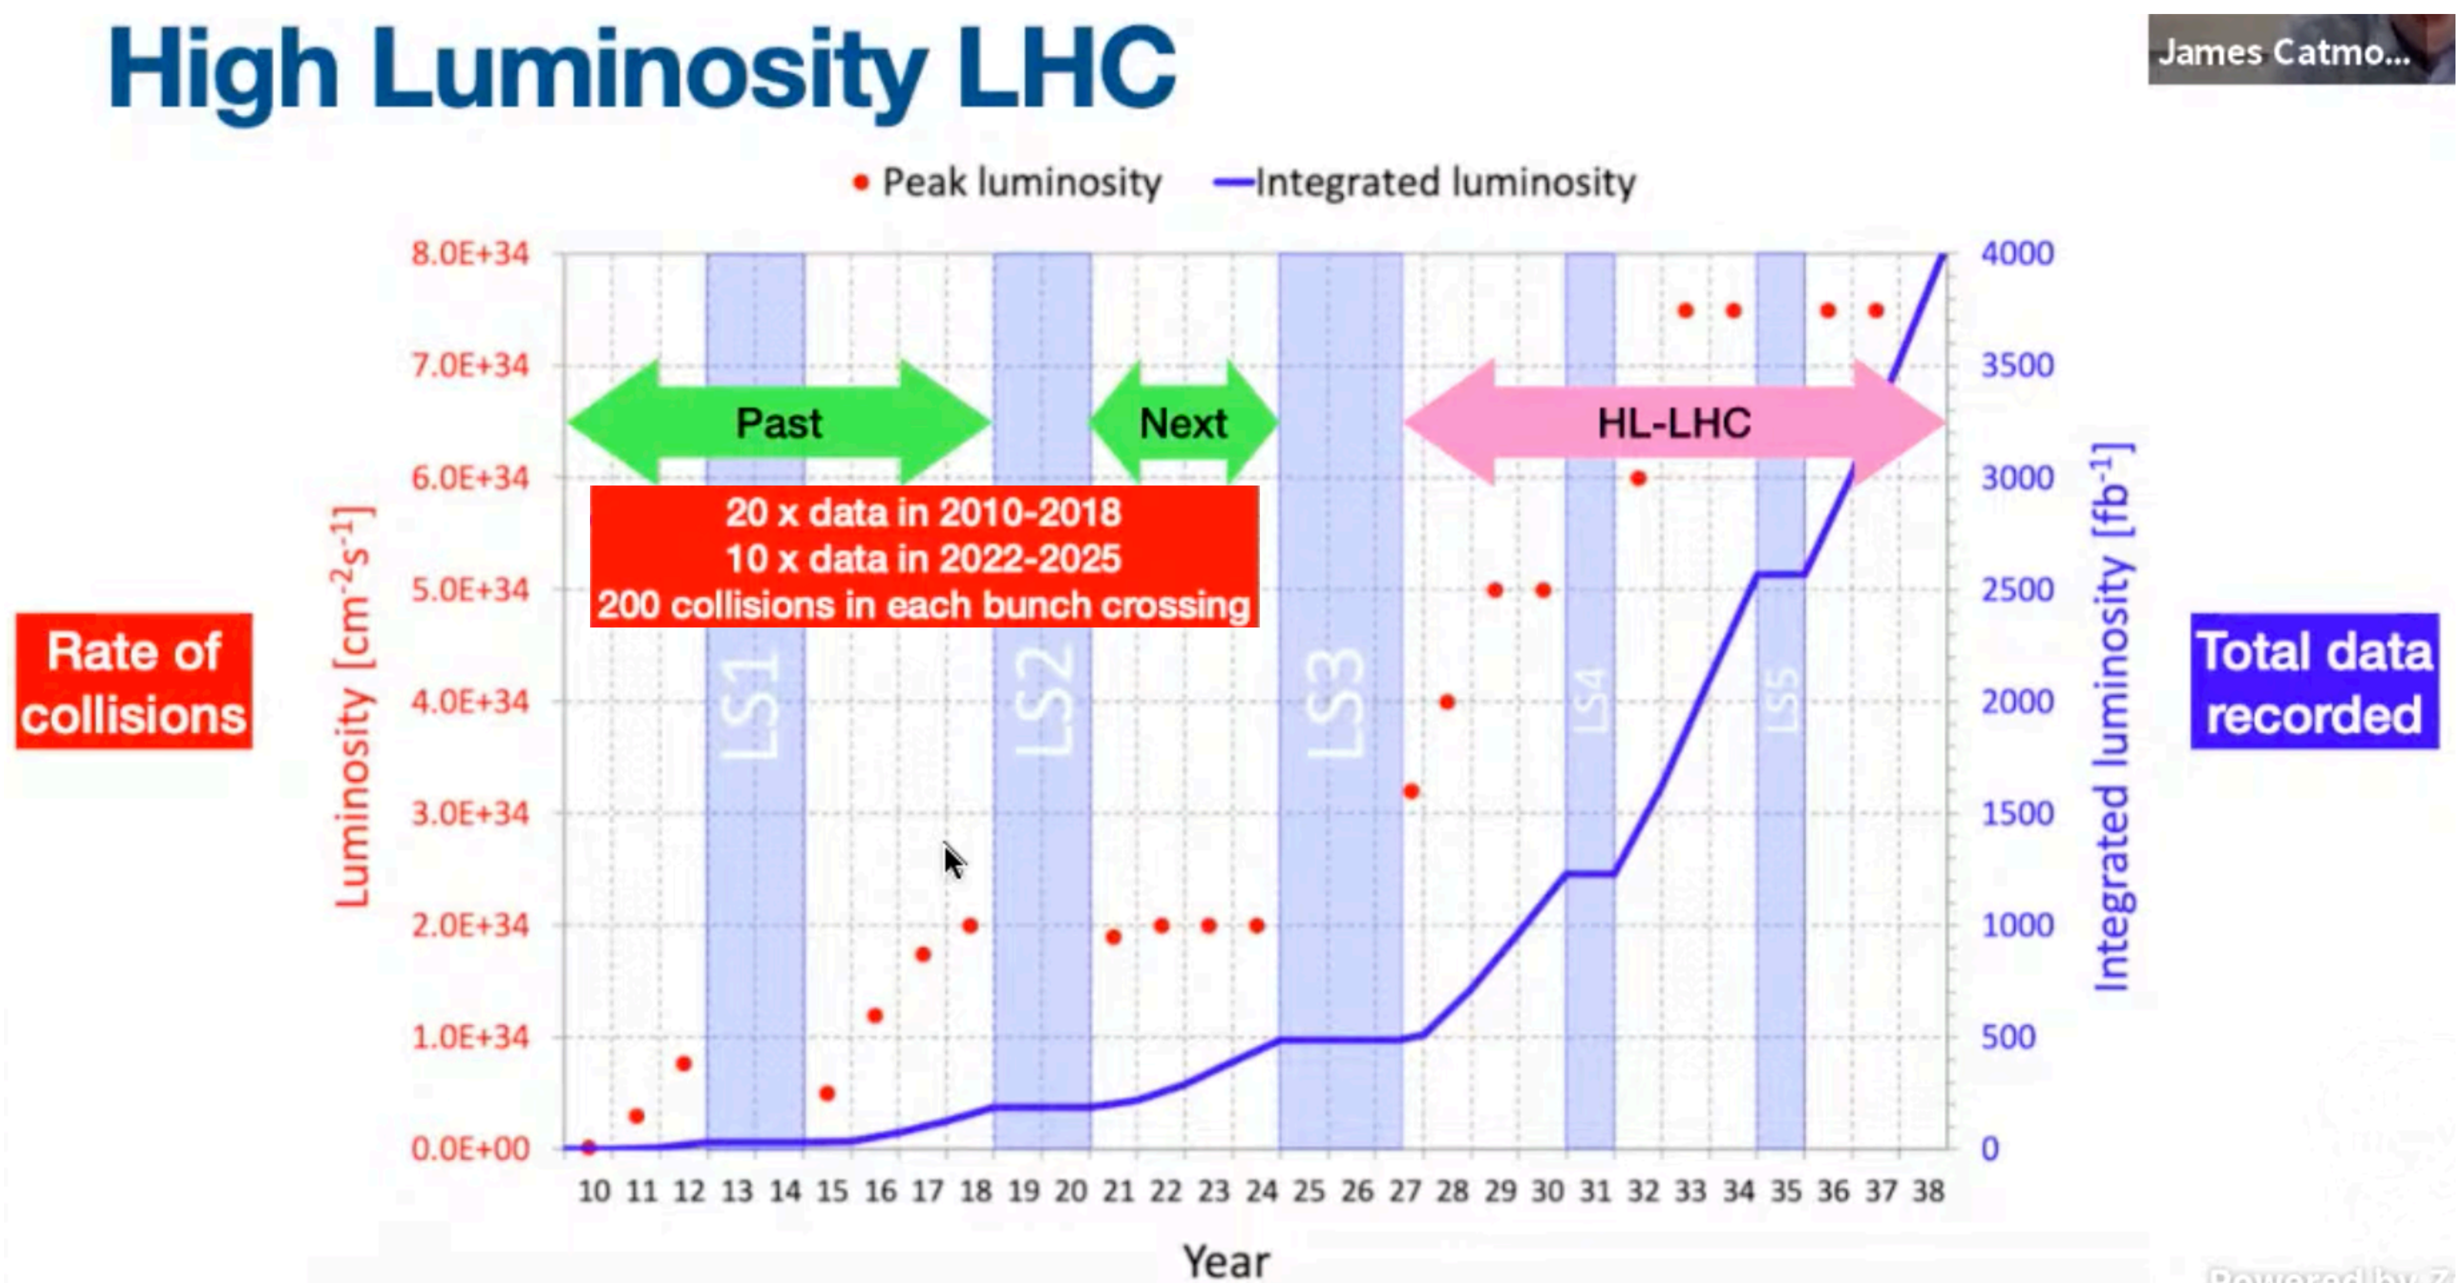

WLCG Tier-1/2 @ASGC

- **ATLAS in Taiwan**
  - Achievements: Higgs boson; Dark matter; Searches for beyond Standard Model
  - Future plan: H—>bb; Di-Higgs; Dark matter
- **Computing Resource - retirement of legacy hardware for energy saving**
  - ASGC Tier-1 (2023): 58,760 HEPSpec06 (3,200 CPUCores)
  - Federated Taiwan Tier-2 (2023): 10,896 HEPSpec06 (1,536 CPUCores)
  - GPU would be available after validation of new computing models (ATLAS, CMS)
- **Storage Resource (2023) of ASGC T1 and FTT T2: 9.6PB + 1.1 PB**
  - Migration from DPM to EOS is under validation by ATLAS
- **Activities for ATLAS**
  - Finished 1,300 billion events, 350 PB in 2010 - 2022 (#processed data and MC events)
  - Development of High Granularity Timing Detector (HGTD) DB and backup support
  - Support Folding@Home for COVID-19 studies
- **Data Networking**
  - 30+ PB data (Inbound + Outbound) transferred in 2022
  - Able to fully utilize the 2x10Gbps links between TW and CERN - reached 19.8bps at peak
- **Contributions to ATLAS Software and Computing**
  - Participating development of ATLAS Harvester/Panda and RUCIO
  - Deeply involved with ATLAS data preparation activities

ASGC Is the Core Facility for Big Data Analysis and Scientific Computing

- Particle physics: ATLAS, CMS, AMS, KAGRA/LIGO/IGWN, ICECube, EIC, proton therapy
- Physics: Surface, Material, Quantum Field, Biophysics
- Astronomy and Astrophysics
- Structural Biology, Drug Discovery, NGS, Bioinformatics, CryEM
- Ecology and Biodiversity informatics
- Computational Chemistry, Biophysical Chemistry, Chemoinfomatics
- Seismology and earth science
- Environmental changes and hazard risk analysis
- ML-enabled data analysis
- Research infrastructure and e-Science: including open data, research data management

GW/KAGRA Computing

- **GPU-Accelerated Parameter Estimation for Gravitational Wave**
- **Numerical Relativity (NR) Challenges**
  - Definition of initial states
  - Step-by-step evolution with General Relativity
  - Huge hierarchical scale
- **Bayesian inference of ~15 parameters (CBC, parameterization of NR and external conditions)**
  - MCMC, Nested sampling, …
- **Advantages from GPU**
  - Detector noise still dominates accuracy
    - Single precision
  - Evaluation of likelihood function
    - Repetition over ~100K frequency bins
- **Source: S. Haino (IoP, AS)**

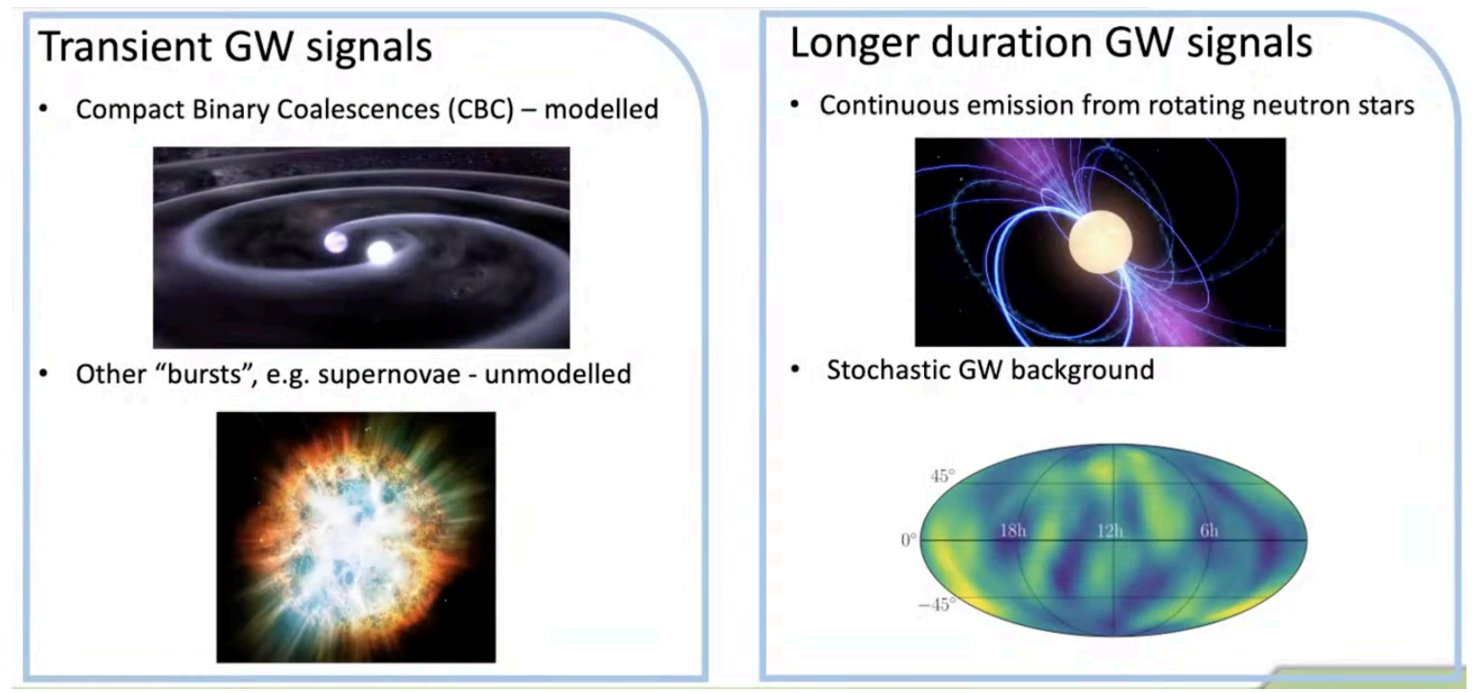

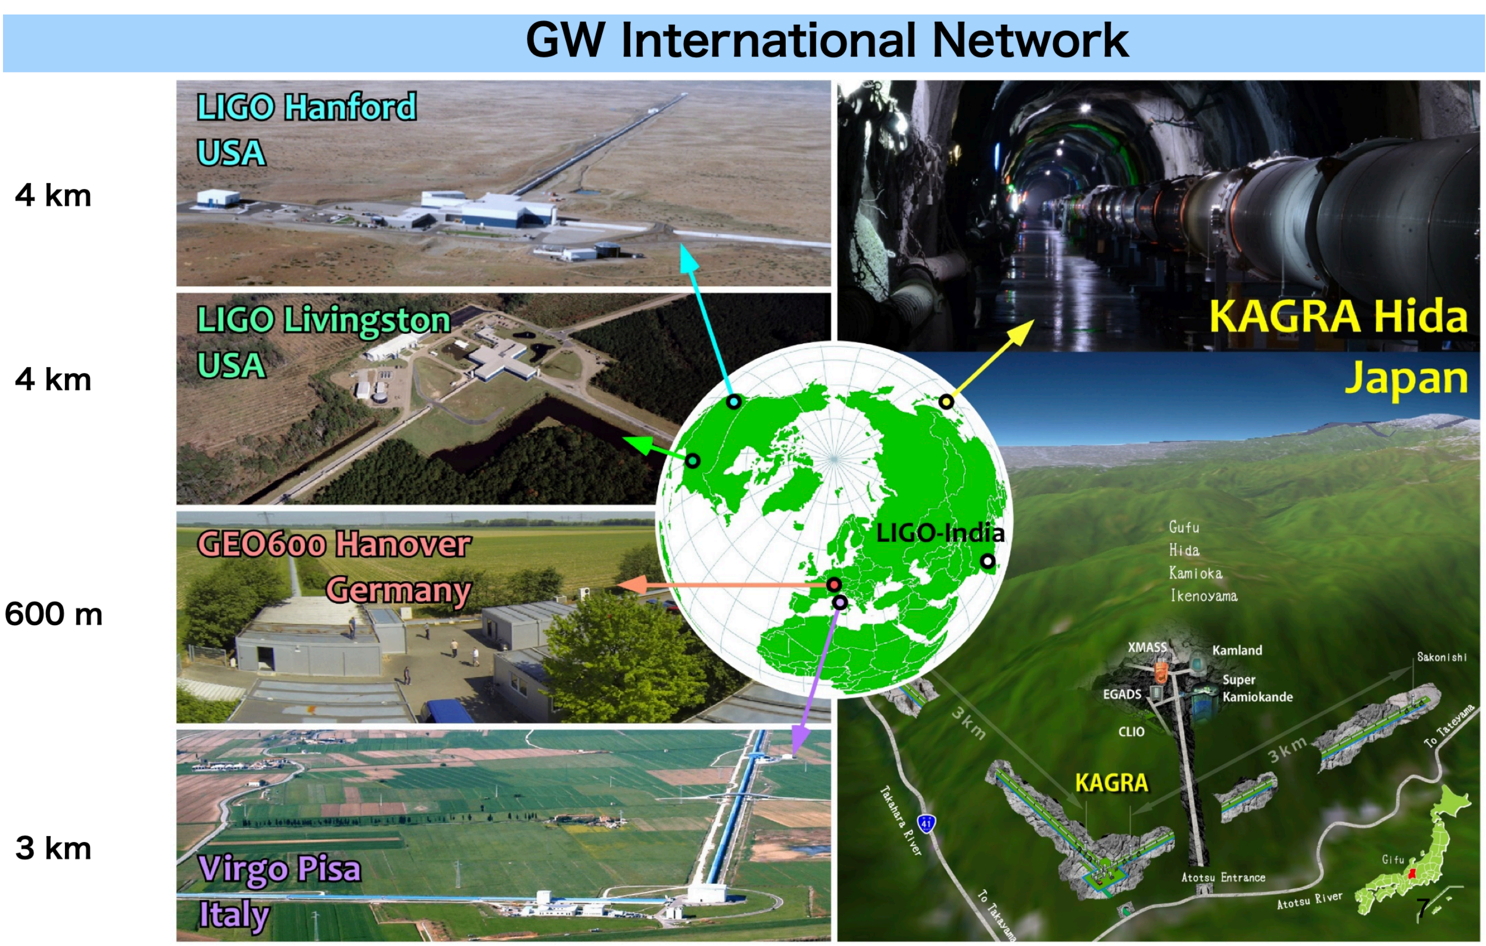

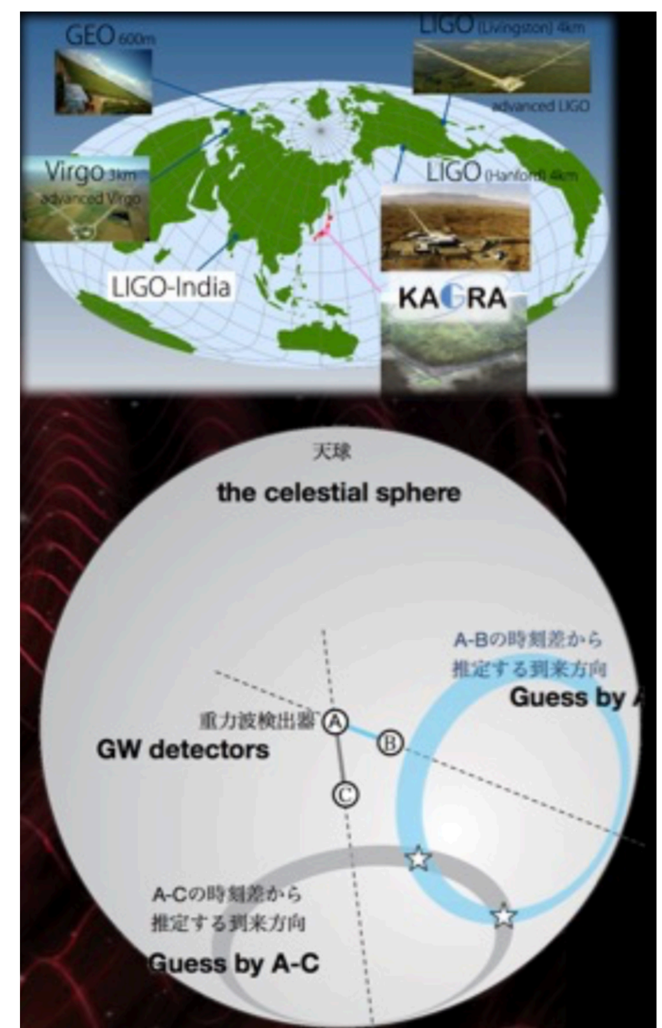

International GW Network

LIGO | VIRGO | KAGRA

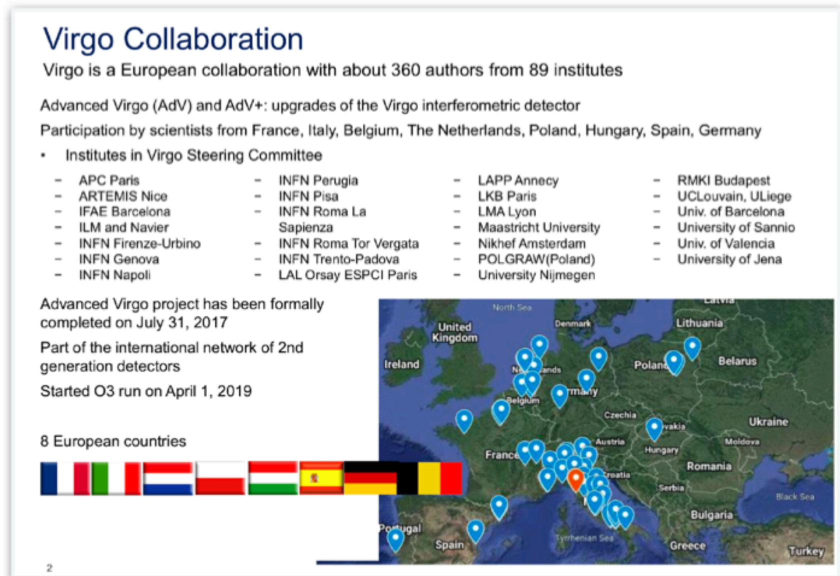

1330 members	**465 members**	**410 members**
860 authors	**360 authors**	**240 authors**
101 groups	**96 groups**	**115 groups**
20 countries	**8 countries**	**14 regions**

- **ASGC has been a member of IGWN collaboration**

Customization Services

- **Principle: Common requirements first, for user community/institute**
  - **Cloud Service：better scalability, efficiency, performance**
  - **Integration & Optimization: data pipeline and workflow (computing model)**
    - **Reduced latency between experiment facility (data source) and data analysis facility**
    - **Integrating required software, application framework, storage and analysis workflow**
    - **Reduced latency between data and training model**
  - **Web Service: developing Web App, Web Portal or Science Gateway**
  - **Continuous Evolution: research pull + technology push**
  - **Generalization and new service creation**
- **Example：Structural Biology: NSRRC (BioSAXS) and Taiwan Protein Project (CryoEM Core Facility)**

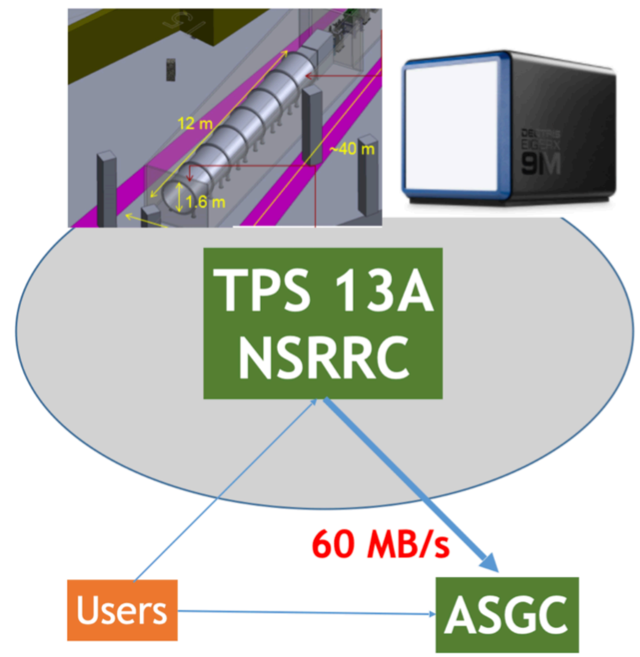

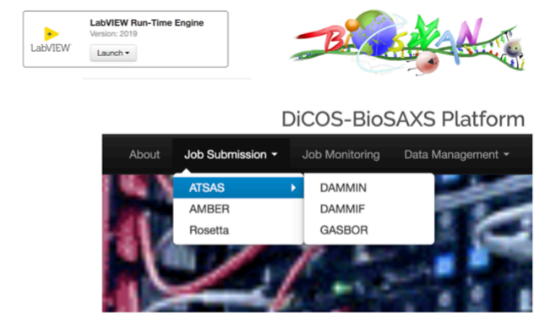

- **Web Portal**
- **Application over Cloud**
- **Jupyterlab**
- **Web Terminal**

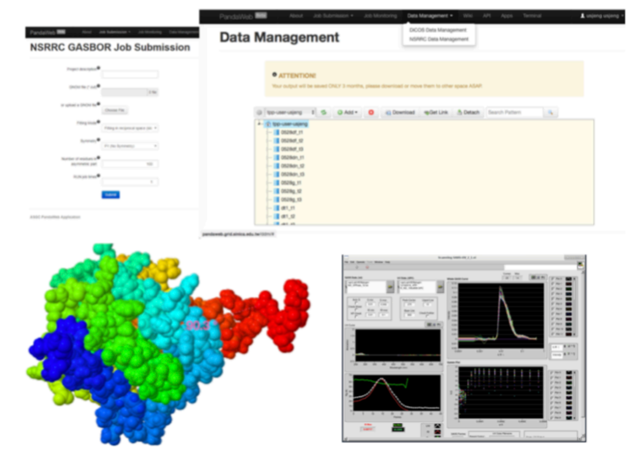

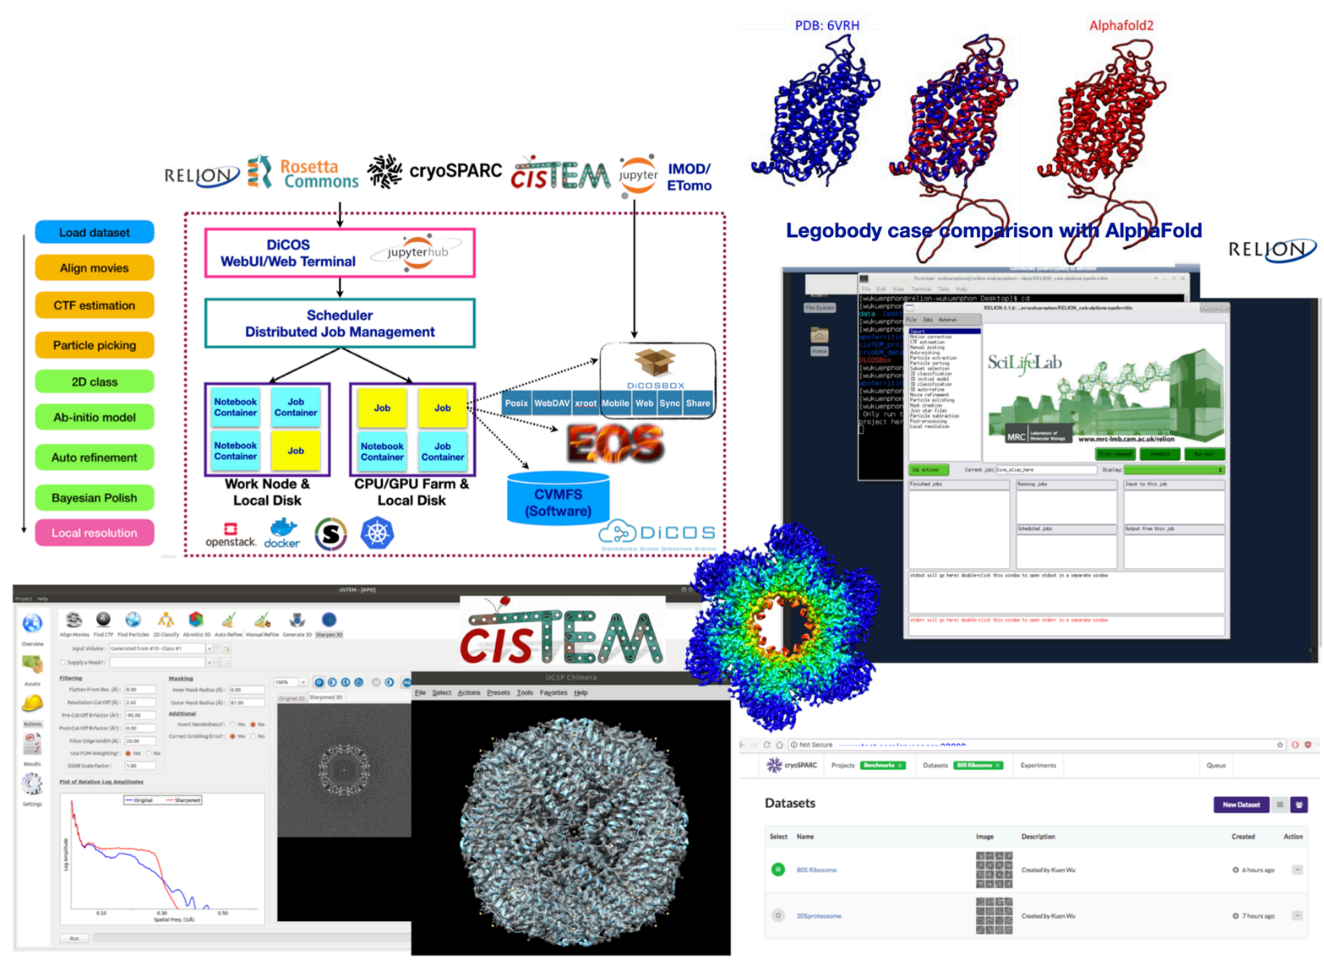

Protein Design

- **Application platform + Computing Infrastructure + Workflow Integration + Efficiency Optimization**
- **AlphaFold, RosettaFold, RosettaFold Diffusion, DiffDock**

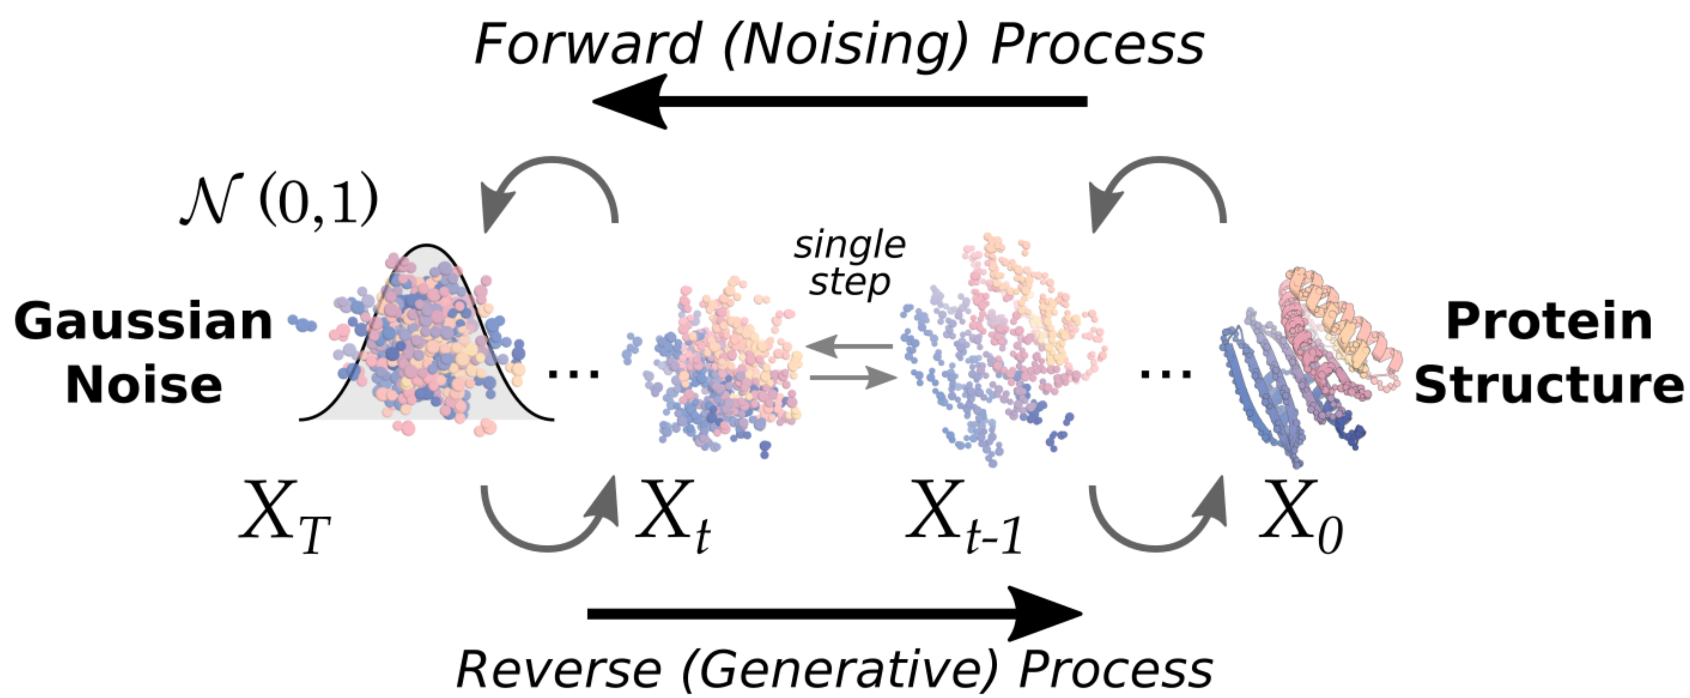

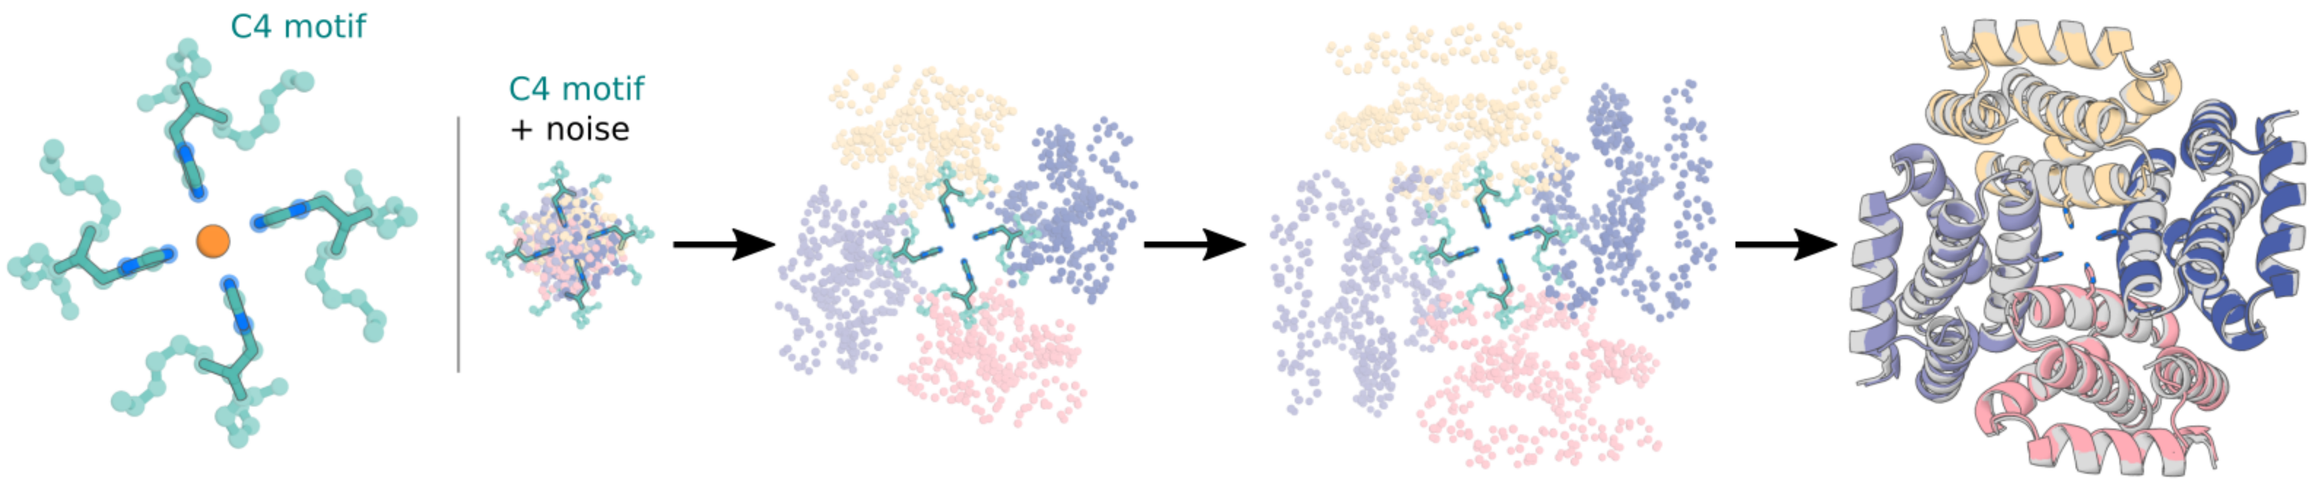

Source: https://github.com/RosettaCommons/RFdiffusion

Supporting BioSAXS Data Analysis for Protein Structure together with NSRRC and ASIBC

DiCOS-BioSAXS Platform

DiCOS-BioSAXS Platform provides TPS 13A BioSAXS beamline users a friendly interface to access their experimental data, analyze data, and submit SAXS simulation jobs.

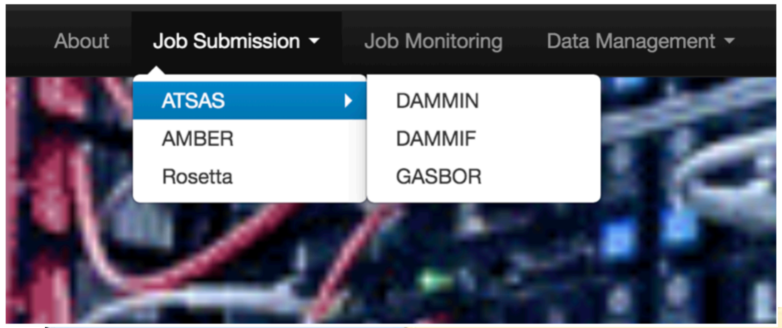

BioSAXS: Biological Small Angle X-ray Scattering

- Goal: Biological structures and structural kinetics in atomic-to-micrometer length scales and in μs – min time resolution
- Operational since April 2019 (~100 users)

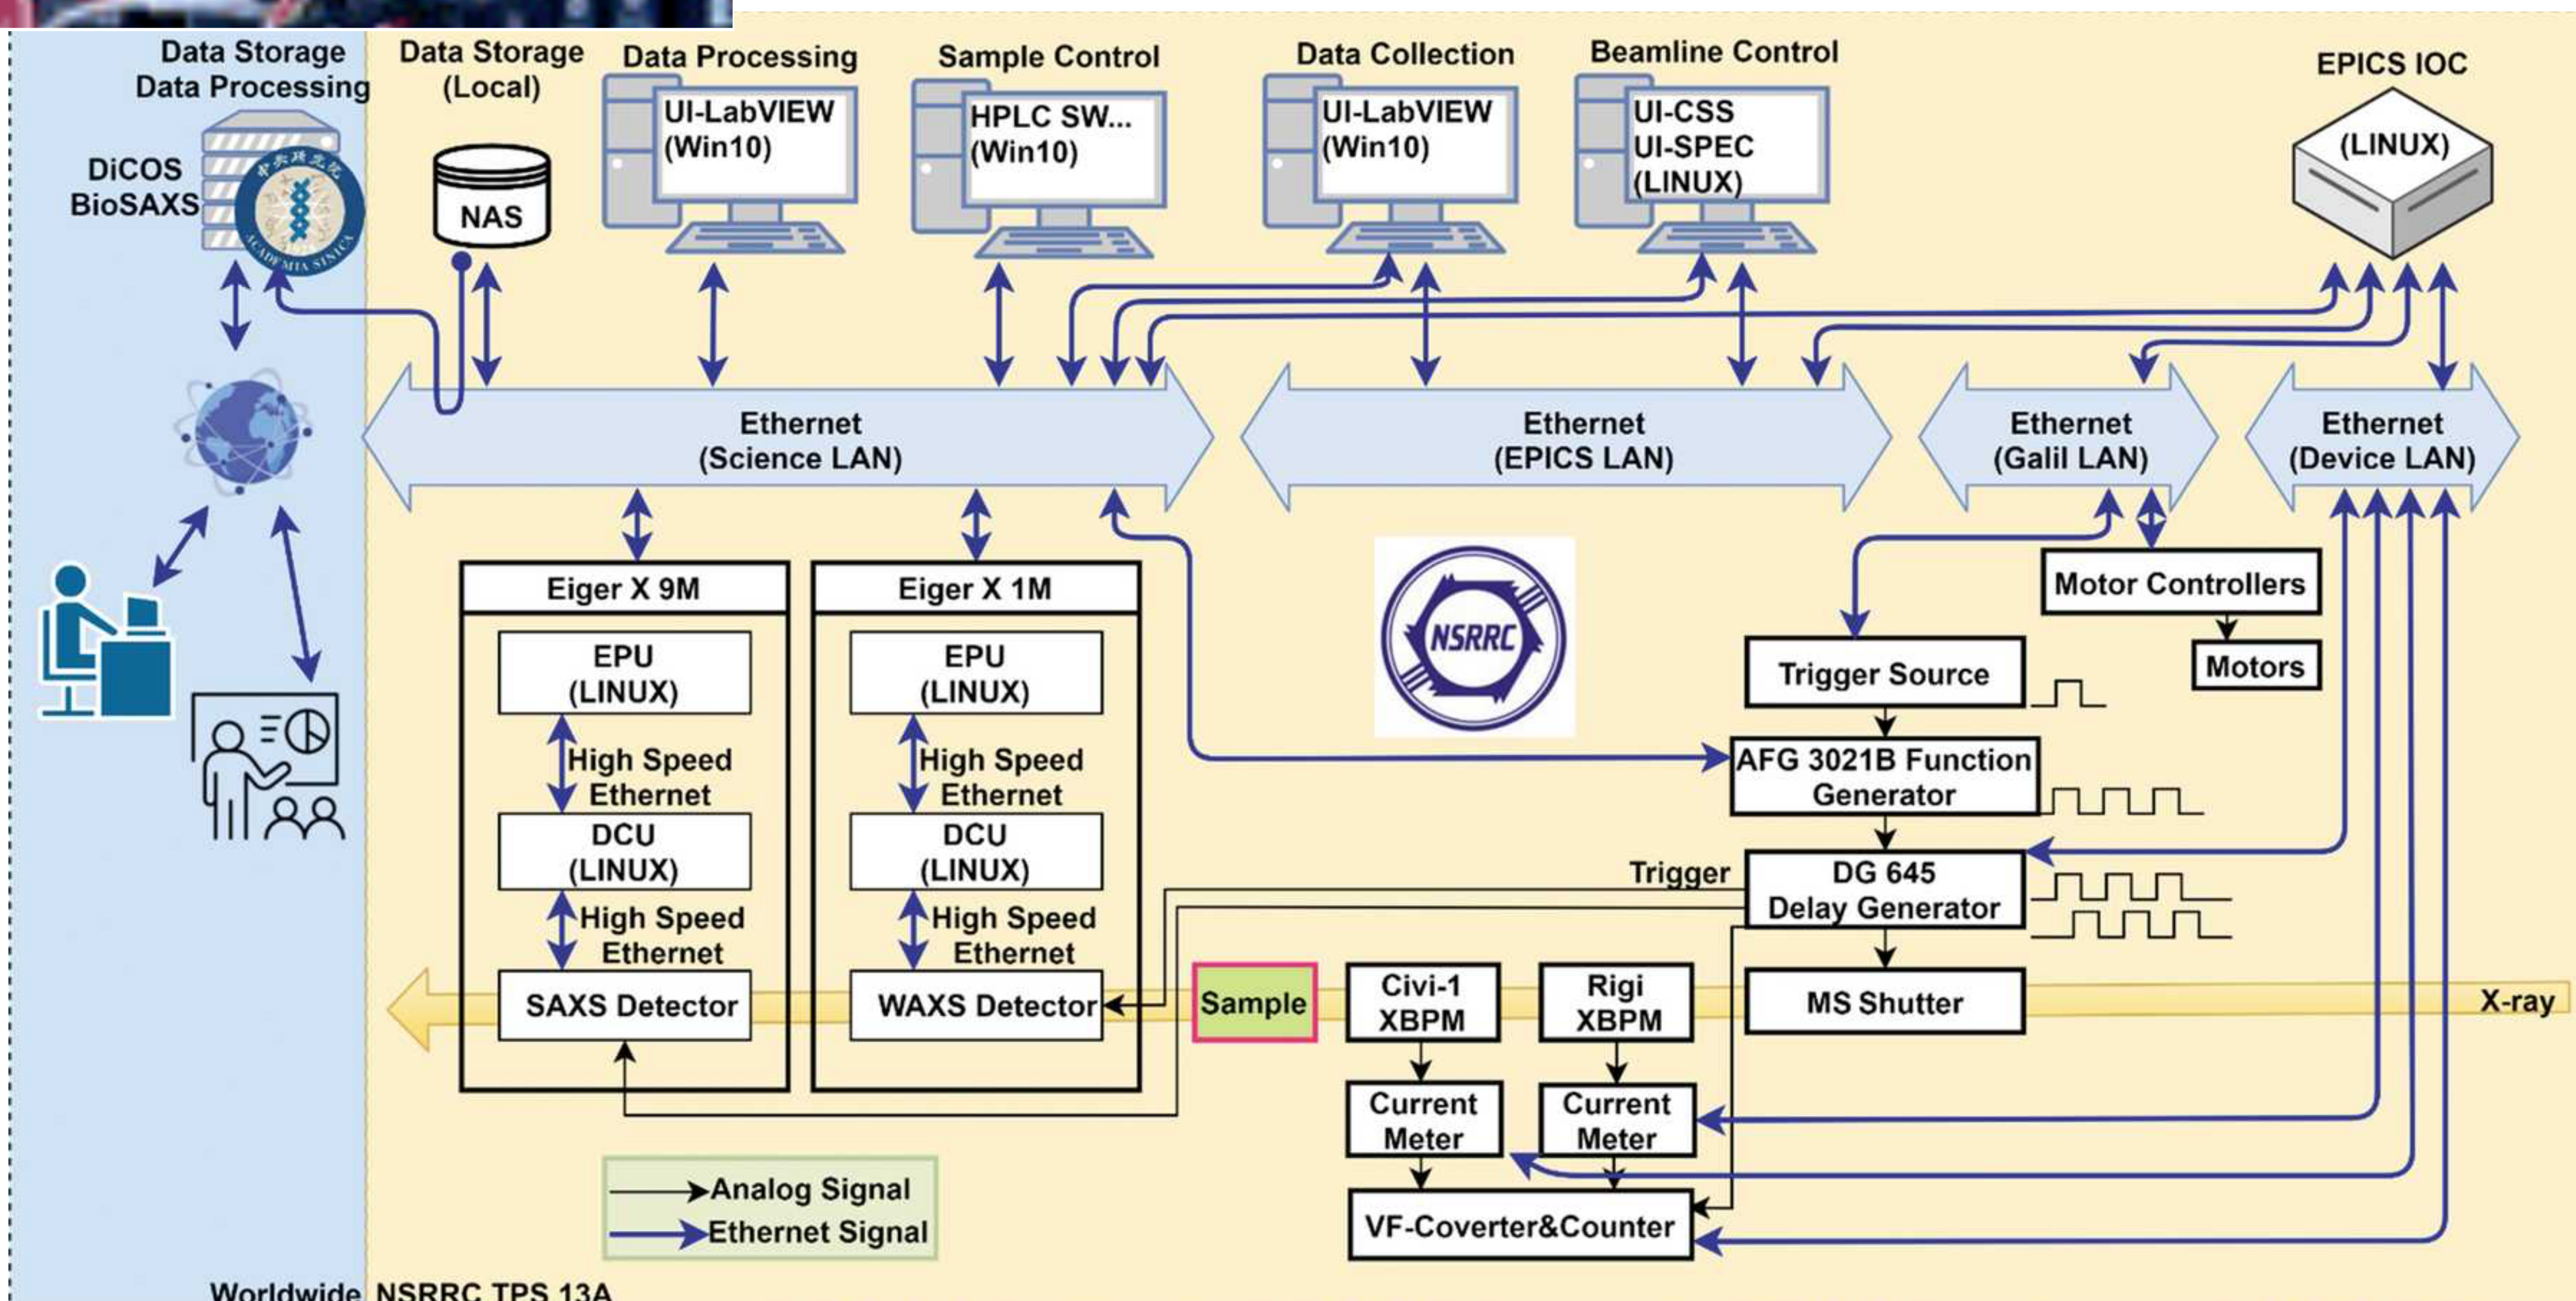

Earthquake Applications

- **P-Alert onsite warning system: Strong Motion Network**
  - P wave brings information and S wave brings energy —> gaining warning time in the blind zone
  - 700+ sites installed mostly in elementary schools
- **Applications**
  - Real-time observed shakeup
  - PGA attenuation relationship
  - Scientific and educational applications

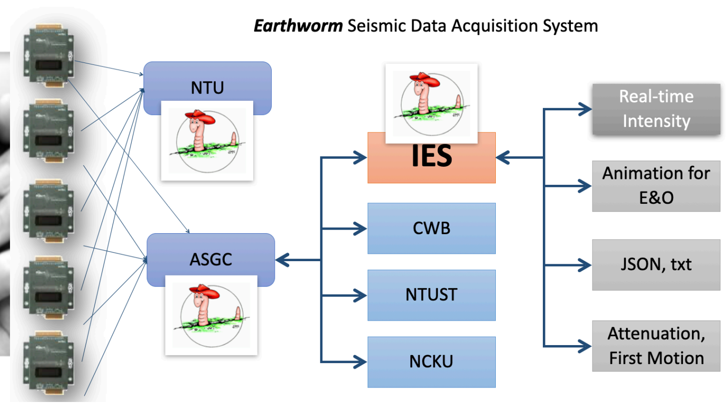

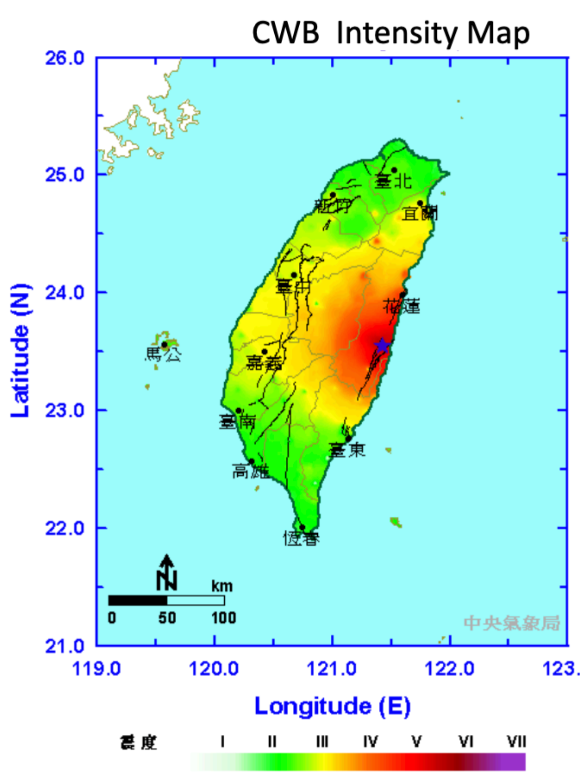

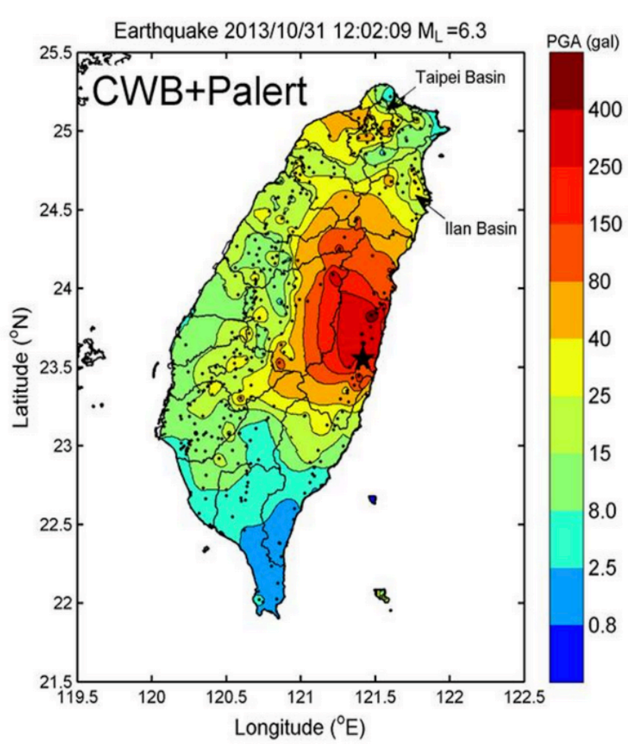

- **MiDAS(光纖訊號搜集地震連續資料)**
  - Making use of intelligent distributed acoustic sensing through fibre optic cables across the Milun fault in Taiwan.
- **Applications**
  - Long-term continuous and intensive wave form collection
  - Improve accuracy of seismic tomography
  - High-resolution rupture characteristics analysis

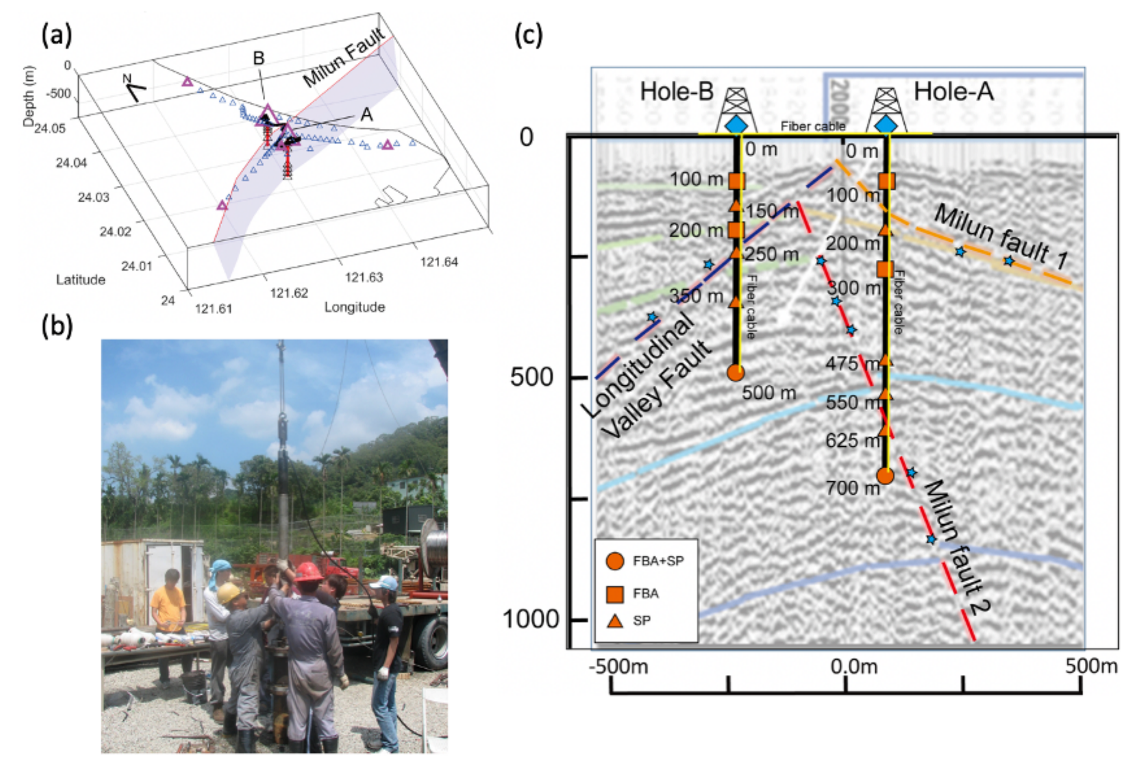

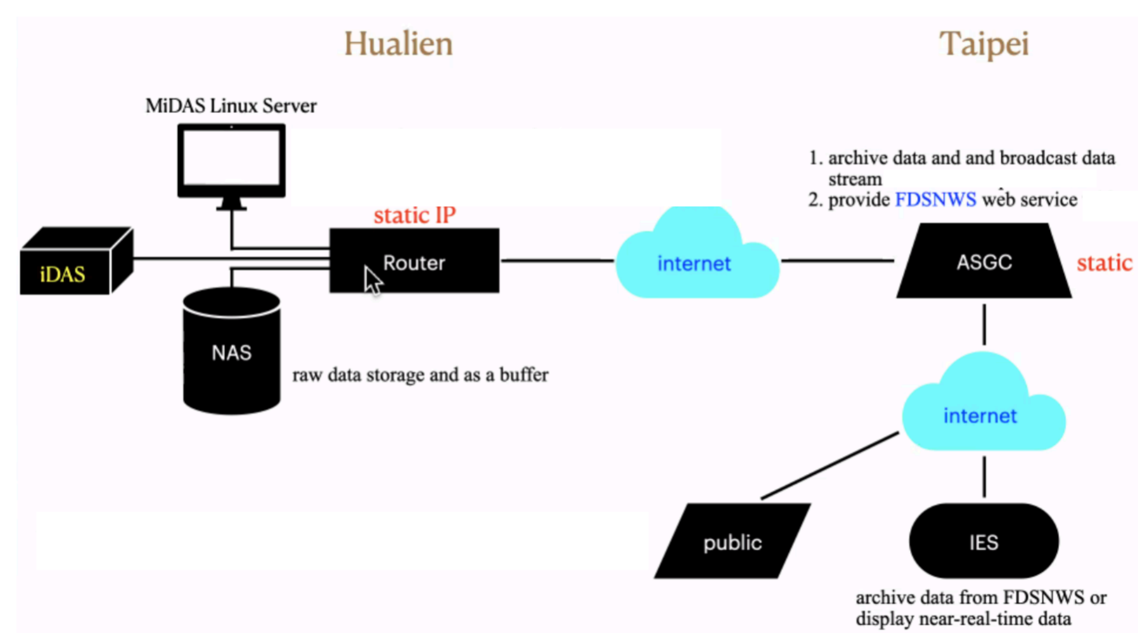

2013 Super Typhoon Haiyan (Yolanda)

Typhoon Life Cycle: November 3rd – 11th

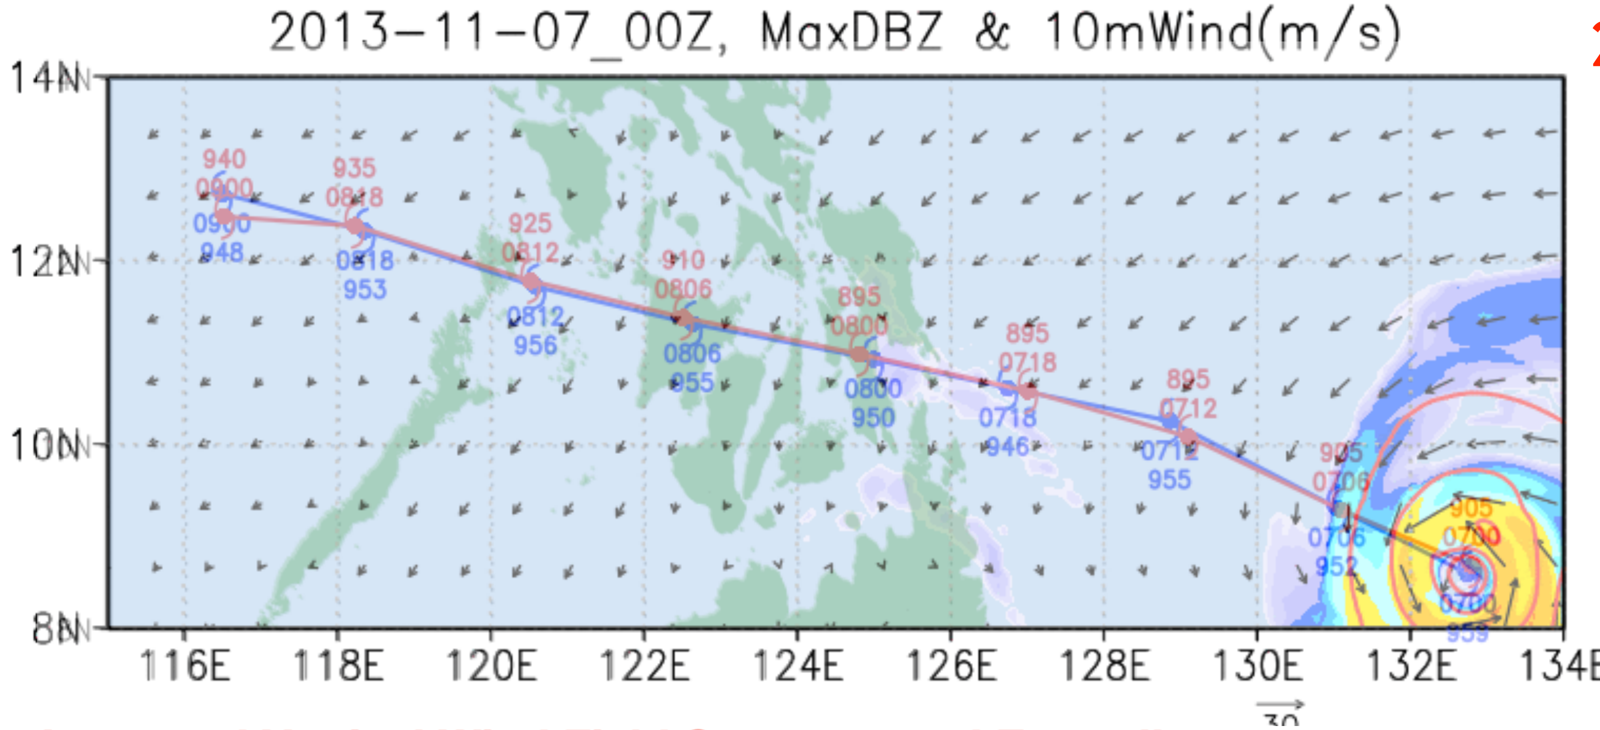

Simulation of impacts from Storm surge caused by super typhoon by combining atmospheric model and oceanic model

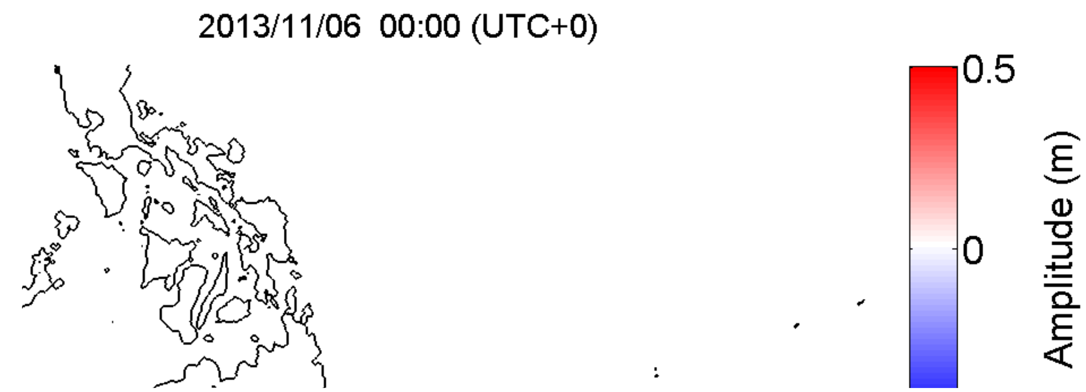

Improved Vertical Wind Field Structure and Eyewall Contraction for Typhoon Haiyan

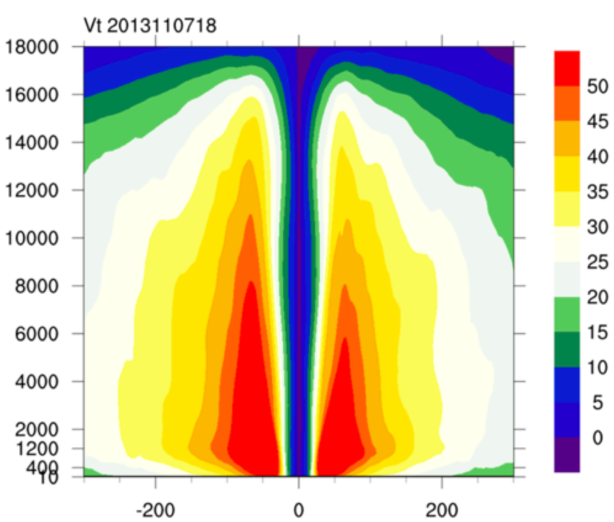

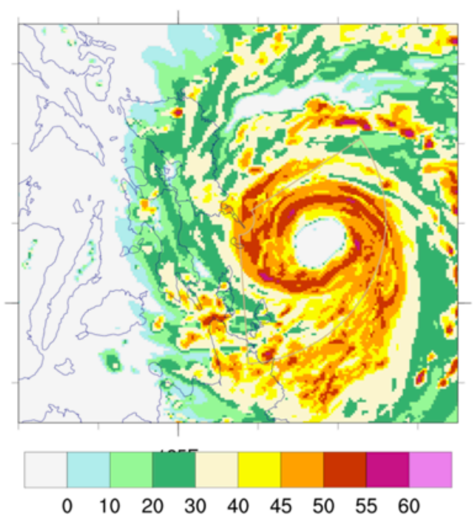

Vegetation change between 2013 and 2020 on Haiyan impacted areas from EODC

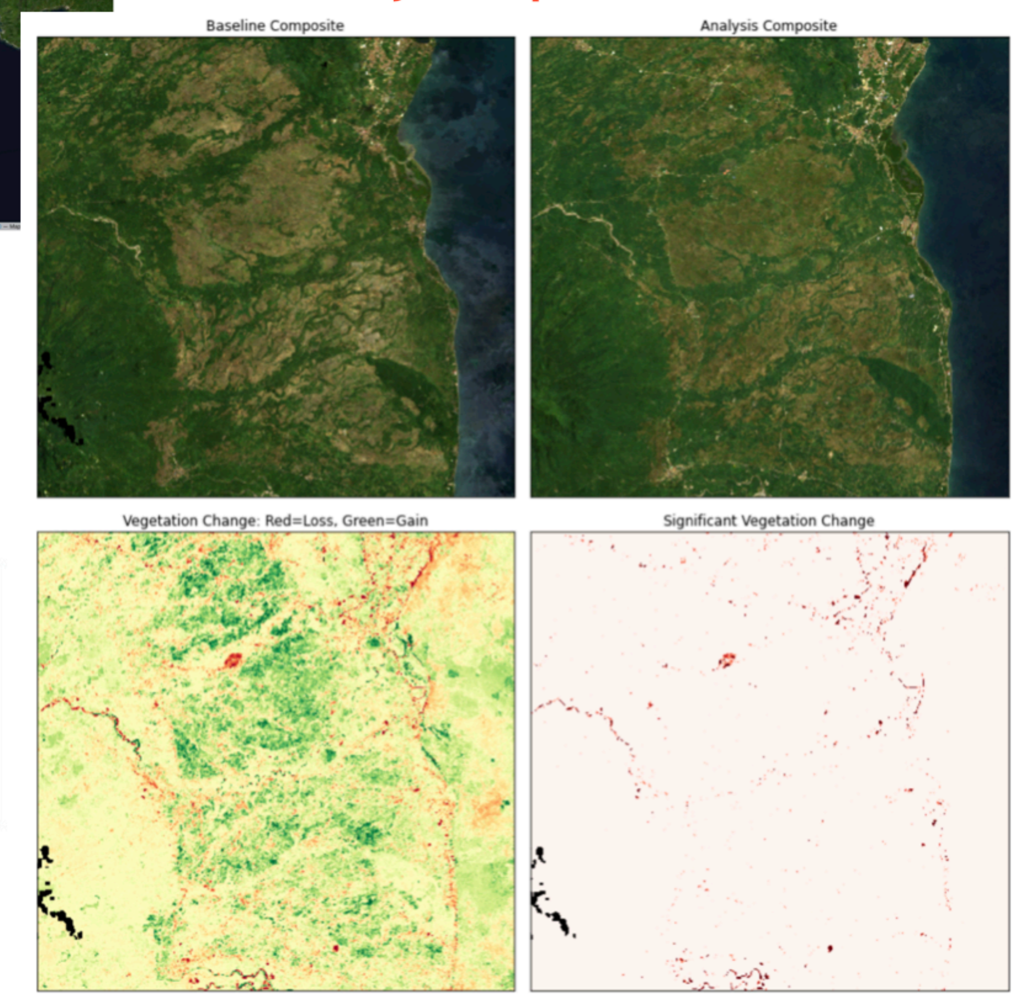

Maximum Simulated Storm Tides at Leyte Gulf

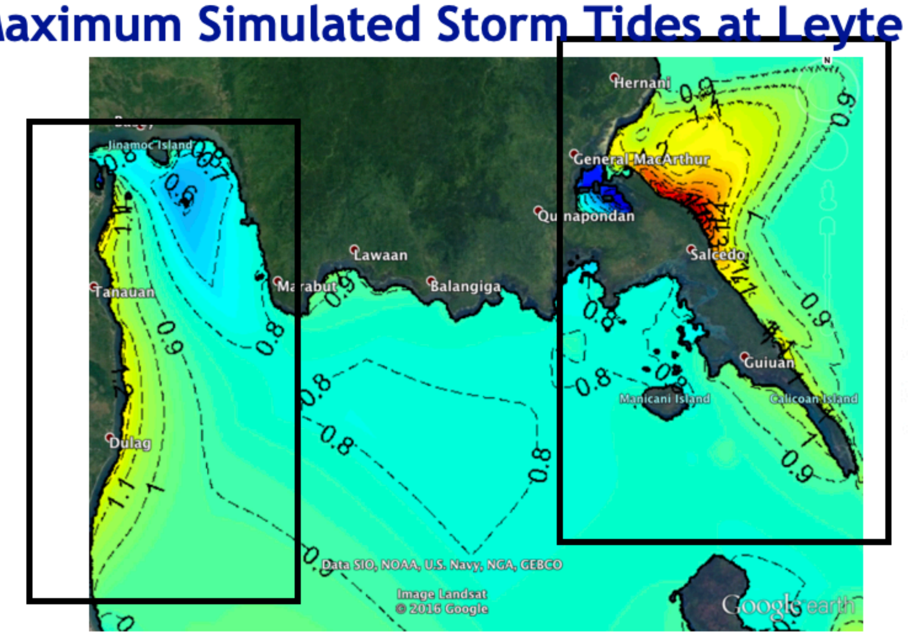

Box-and-Whisker Plot of NDVI from EODC

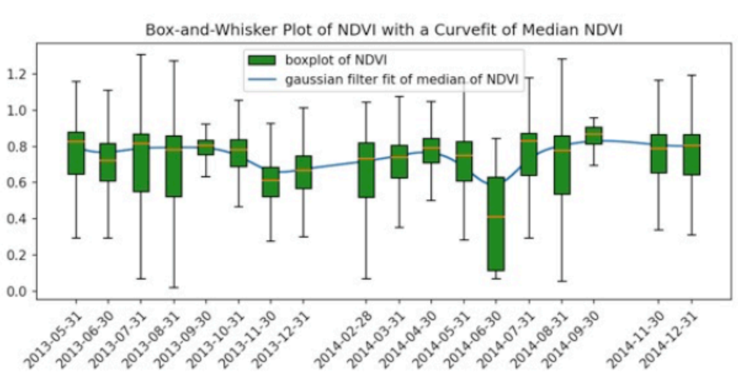

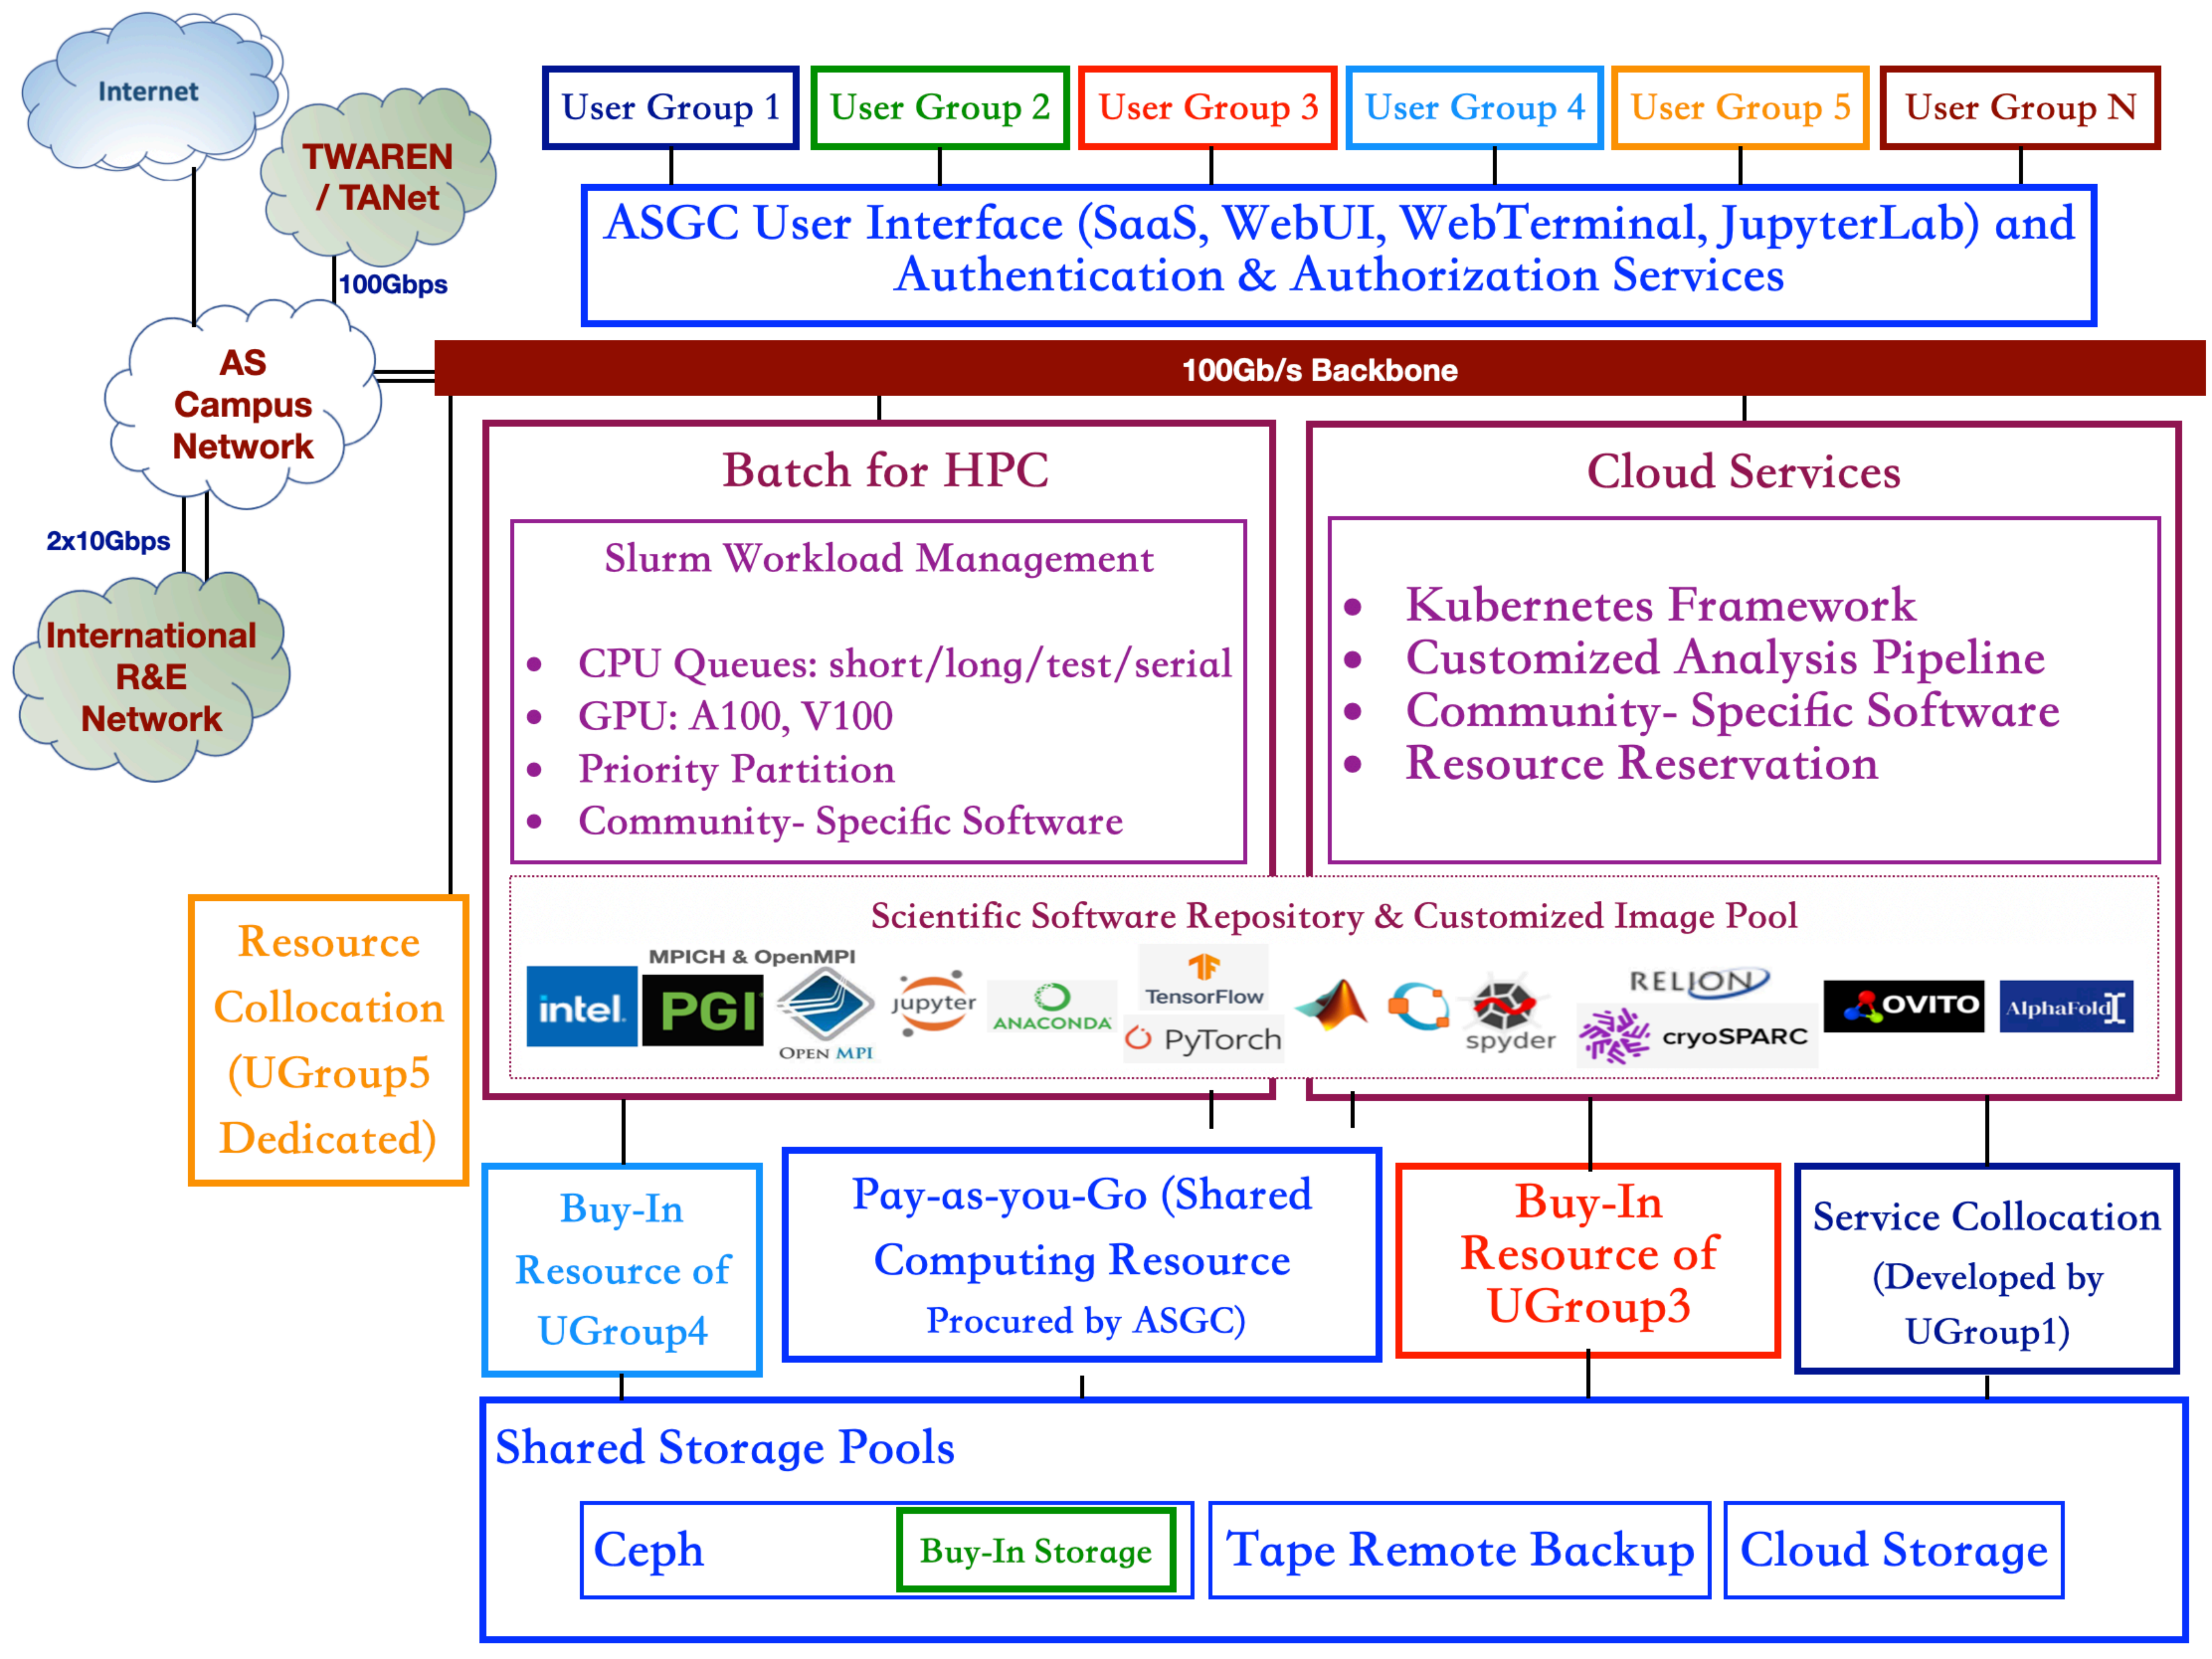
Internet
TWAREN / TANet
100Gbps
AS Campus Network
2x10Gbps
International R&E Network
User Group 1
User Group 2
User Group 3
User Group 4
User Group 5
User Group N
ASGC User Interface (SaaS, WebUI, WebTerminal, JupyterLab) and Authentication & Authorization Services
100Gb/s Backbone
Batch for HPC
Slurm Workload Management
• CPU Queues: short/long/test/serial
• GPU: A100, V100
• Priority Partition
• Community- Specific Software
Cloud Services
• Kubernetes Framework
• Customized Analysis Pipeline
• Community- Specific Software
• Resource Reservation
Scientific Software Repository & Customized Image Pool
intel
MPICH & OpenMPI
PGI
Open MPI
jupyter
ANACONDA
TensorFlow
PyTorch
spyder
RELION
cryoSPARC
OVITO
AlphaFold
Resource Collocation (UGroup5 Dedicated)
Buy-In Resource of UGroup4
Pay-as-you-Go (Shared Computing Resource Procured by ASGC)
Buy-In Resource of UGroup3
Service Collocation (Developed by UGroup1)
Shared Storage Pools
Ceph
Buy-In Storage
Tape Remote Backup
Cloud Storage

Supporting Big Data & AI in Innovations

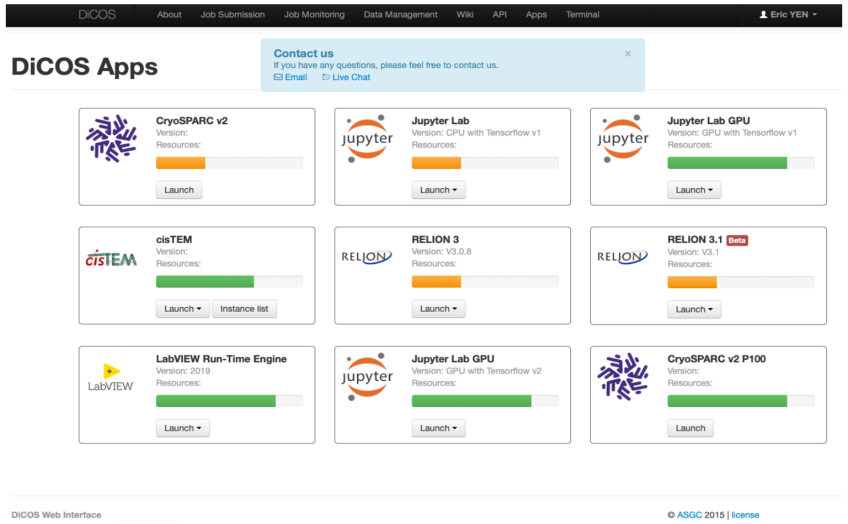

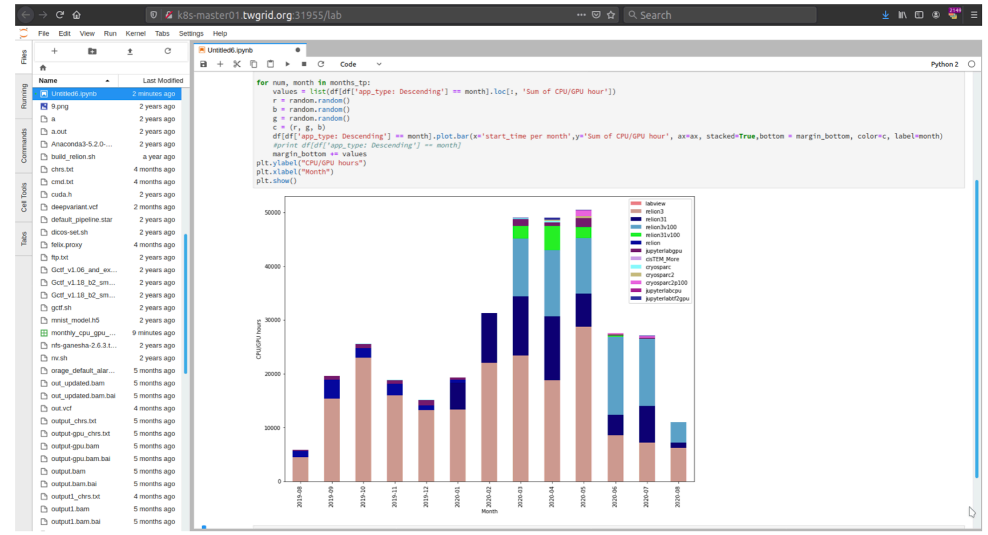

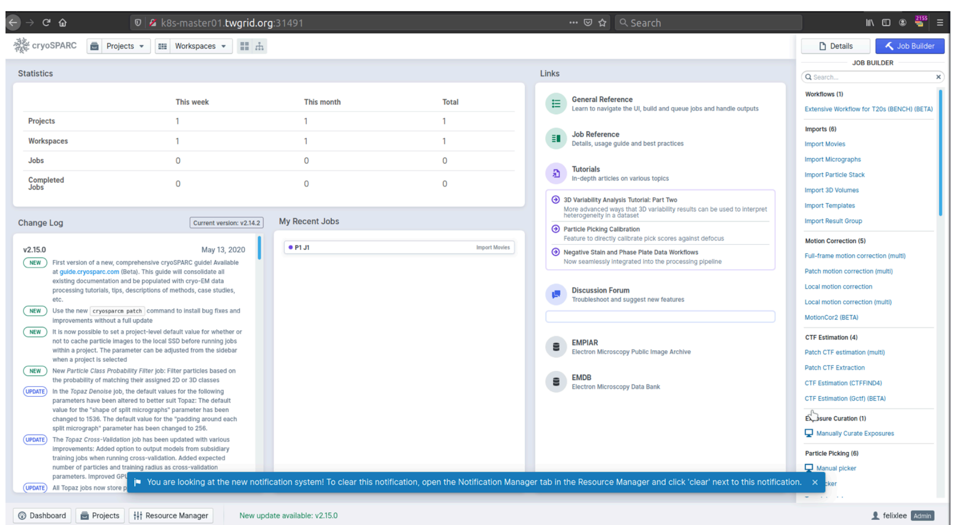

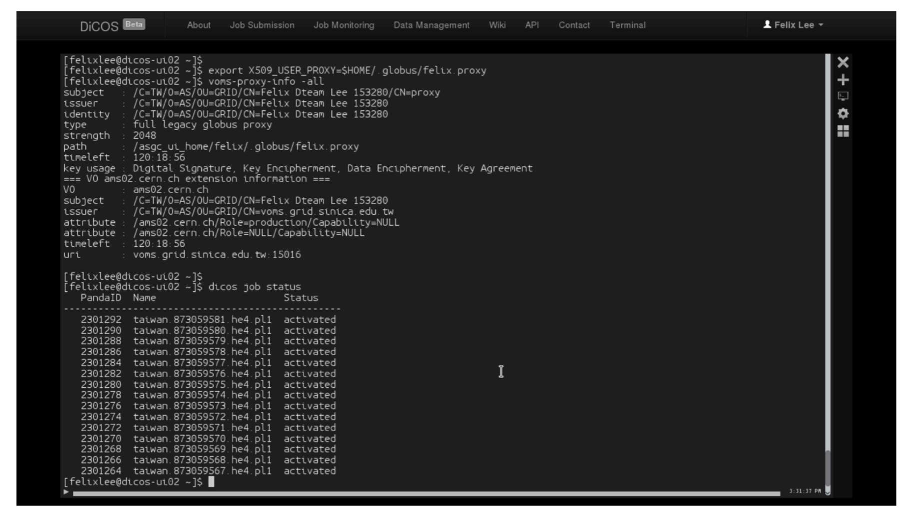

CLI	Web Portal	DiCOS APP	Jupyter Notebook	Science Portal	Web Browser/ Terminal

Application-specific/ Generic Learning Engines — scikit learn, ROOT Data Analysis Framework

Deep Learning Engines/Frameworks — NVIDIA TRITON INFERENCE SERVER, PyTorch, TensorFlow, Keras

Computing Resource (Cloud/Grid/Slurm)	Storage Resource (Ceph/EOS)	Distributed Data Management & Cloud Storage Services	Network & Data Transmission Services

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

LabVIEW Run-Time Engine
Version: 2019
Launch

DiCOS-BioSAXS Platform

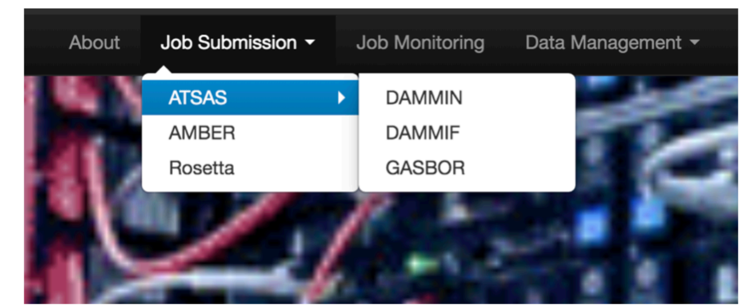

Other

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

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

Contributing to ML-Enabled Analytics

- Research groups could focus on scientific problems solving
- Starting from ML/AI application platform service - SW library, HW, integration and application
  - Build up customized ML platforms for user specified projects - Deploy customized 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
  - Users who don't have source code for the moment but come up with concrete idea - ASGC can try as much as possible to setup various types of ML packages ready environment. It should be helpful for users to test and choose an appropriate ML method.
  - Users who don't have defined idea - ASGC is not able to contribute to the prototyping at this moment
- 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.

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**
- **System Anomaly Detection**

AHU Performance Monitoring

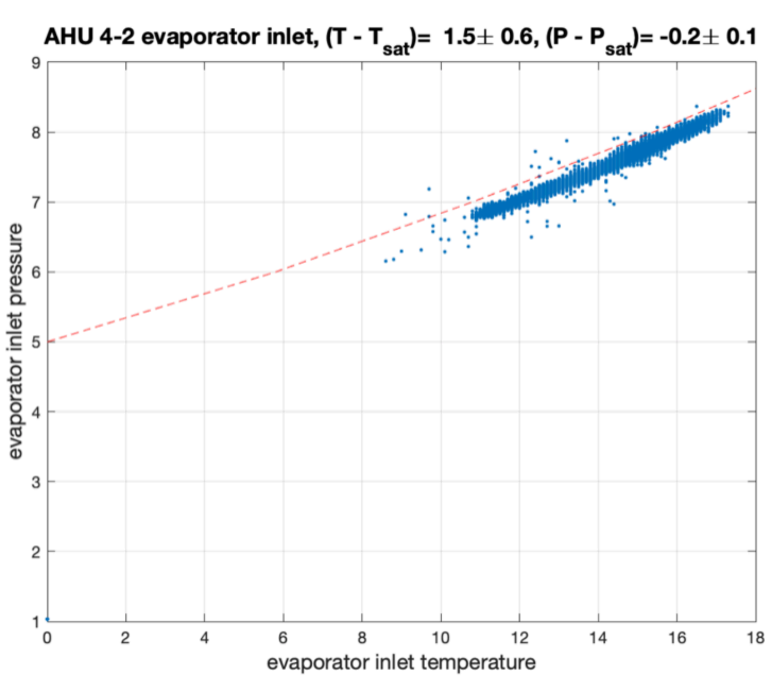

Worknode Monitoring: High ratio of WaitIO & System

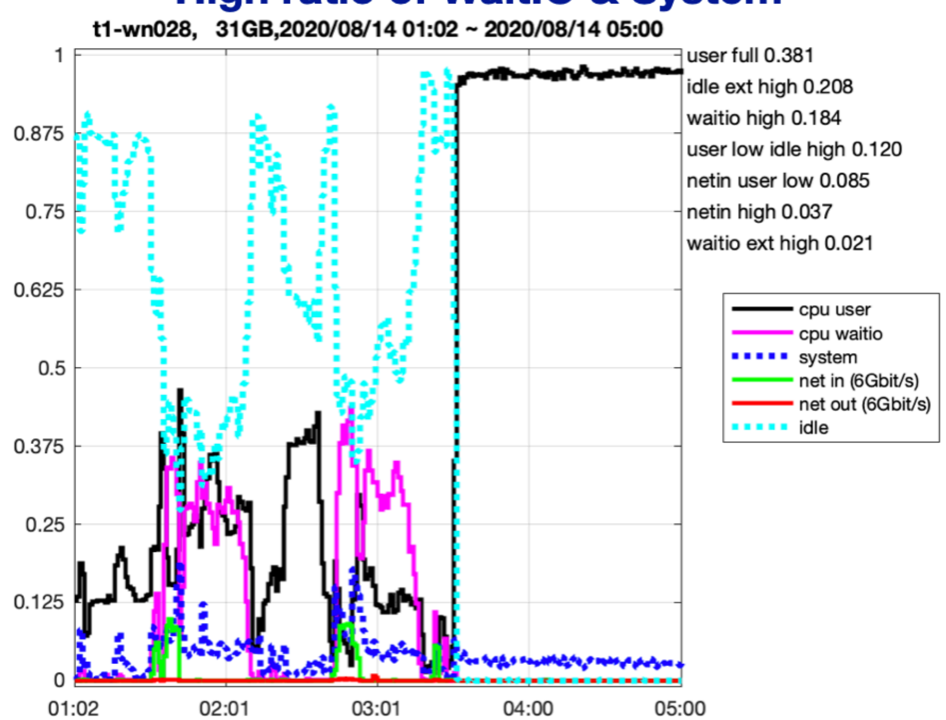

Worknode Monitoring: Misconfigured hyper-threading

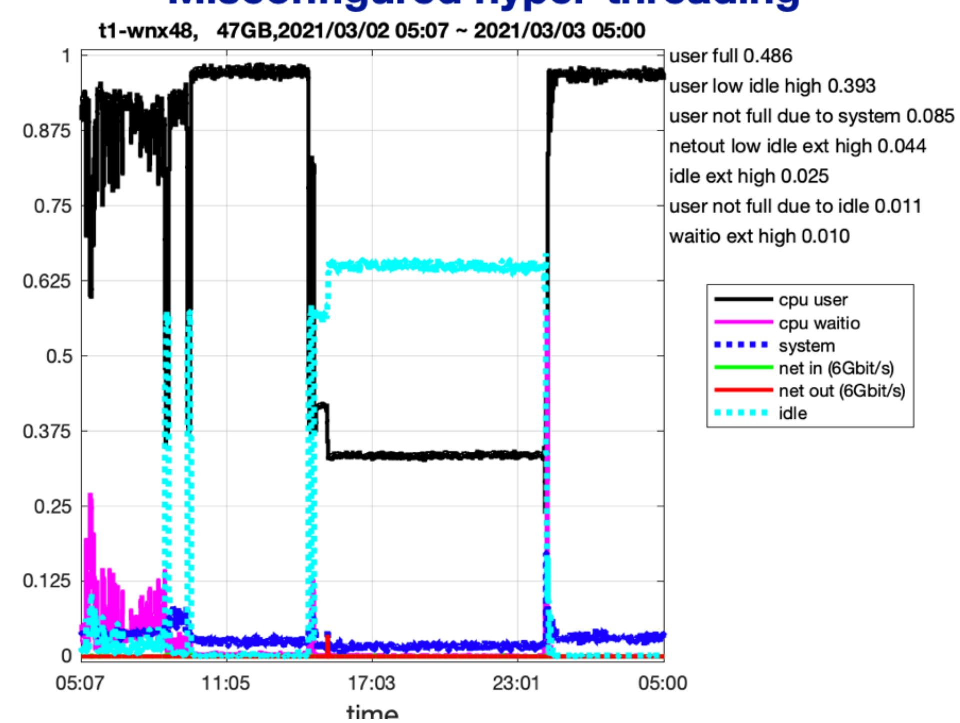

Summary

- Technological evolution in computing empowers science
  - especially in data-intensive domains, not just High Energy Physics but also life sciences and broader disciplines
  - Computing is strategical to do research efficiently on a large scale
- Reliability and efficiency are the keys for a production research infrastructure
- Collaboration and evolutionary approach are essential to the science advancement
  - about the next generation of data processing and analysis workflows that will maximize the science output
- Based on WLCG core technologies, ASGC is supporting big data analysis and AI in innovations for broader disciplines

ASGC Data Center & Resources

Cooling Power : CPU Power
1 : 2

No UPS to save 10% power consumption

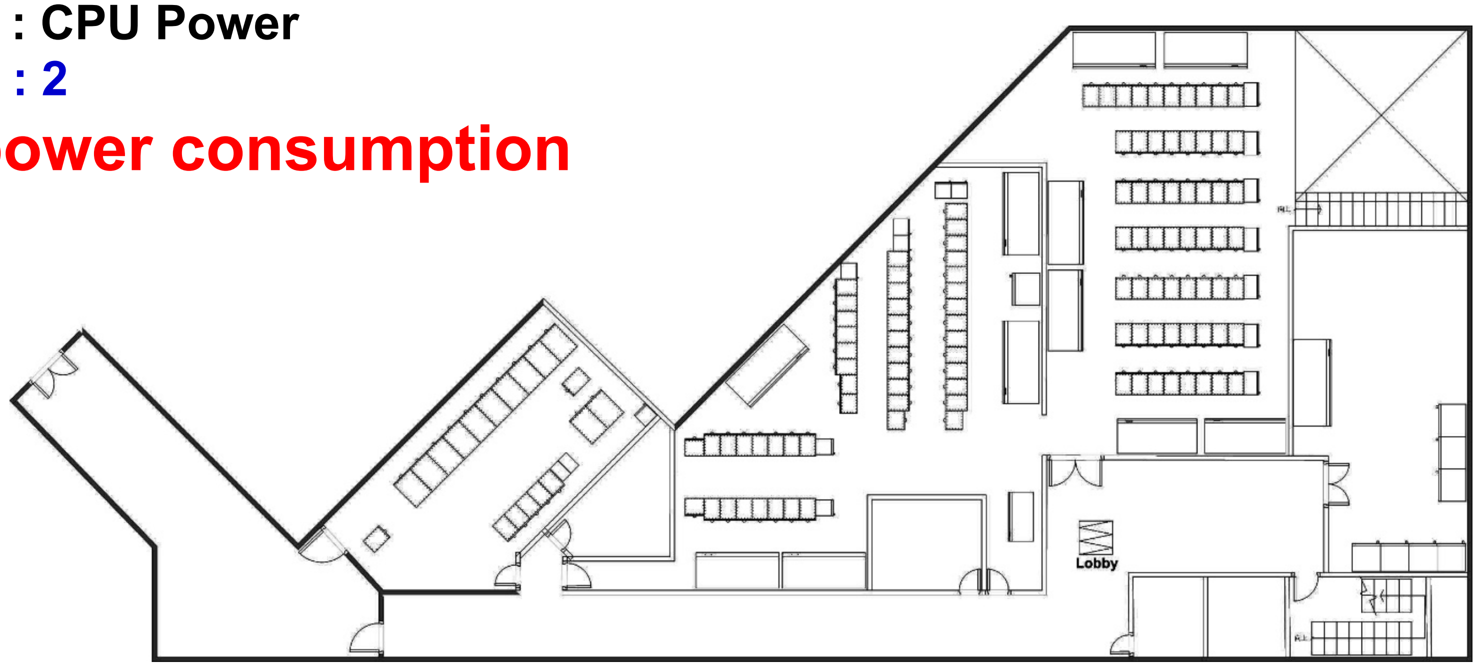

- **Total Capacity**
  - 2MW, 400 tons AHUs
  - 112 racks in ~ 800 m^2
- **Resources (Apr. 2022)**
  - 12,000 CPU Cores
  - 220 GPU Cards
  - 30 PB Disk Storage
  - 2x10Gb links to CERN and primary NRENs worldwide
- **WLCG Tier-1 Center since 2005**
- **Supporting HPC & HTC in Academia Sinica by distributed 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

ASGC Data Center

- Efficiency optimization and intelligent monitoring & control
  - Continuous evolution according to lessons learned from daily operation
- Continuous evolution of power and thermal efficiency
  - Separation of cold isle and hot isle (2010-)
  - Replacement of DC-wide UPS by small-scale UPS for storage
  - PoC on conduction-based heat dissipation for the Single Rack Data Centre (2014-2016)
  - Deployment of evaporative condensing air conditioning system (2013-)
  - Improvement of DC air interchanging system (2015-)
  - Enhancement of condensers (2022-)
- Intelligent monitoring and Control
  - Anomaly detection and warning for air handlers (2022-)
  - Intelligent power saving system for servers (2018-)
  - Scientific computing efficiency enhancement (2016-)

ASGC Services

- ASGC Web Site: https://www.twgrid.org
- Access to ASGC Resources
  - https://dicos.grid.sinica.edu.tw/
- Contact point: DiCOS-Support@twgrid.org