



ASGC Site Report


Felix Lee,
Eric Yen

Academia Sinica Grid Computing Centre (ASGC)
Taiwan

HEPiX Spring 2021
15 March 2021

Mission


- Acting as a WLCG Tier-1 center and developing advanced distributed cloud infrastructure and technologies for O(1000)PB big data analysis globally through the collaboration with WLCG
 - Building capacity of large-scale distributed cloud for efficient big data analysis of AS
- Developing DiCOS and enhancing science-enabling capabilities
 - DiCOS 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)**


- 20,922 CPU Cores
- 868,352 GPU CUDA Cores/160 GPU Cards
 - A100 Server *3 (8 GPU Card for each) will be deployed in 2021
- 20 PB Disk Storage
- 2x10Gb links to CERN and primary NRENs worldwide

- **WLCG Tier-1 Center since 2005**

- **Supporting high-performance & high-throughput computing in Academia Sinica by distributed cloud operating system (DiCOS)**

- Usage > 1M CPUCore-Days in 2015 (ATLAS used 61%)
- Usage > 2M CPUCore-Days in 2019 (ATLAS used 68%)

- **R&D on system efficiency optimization by intelligent monitoring & control**



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

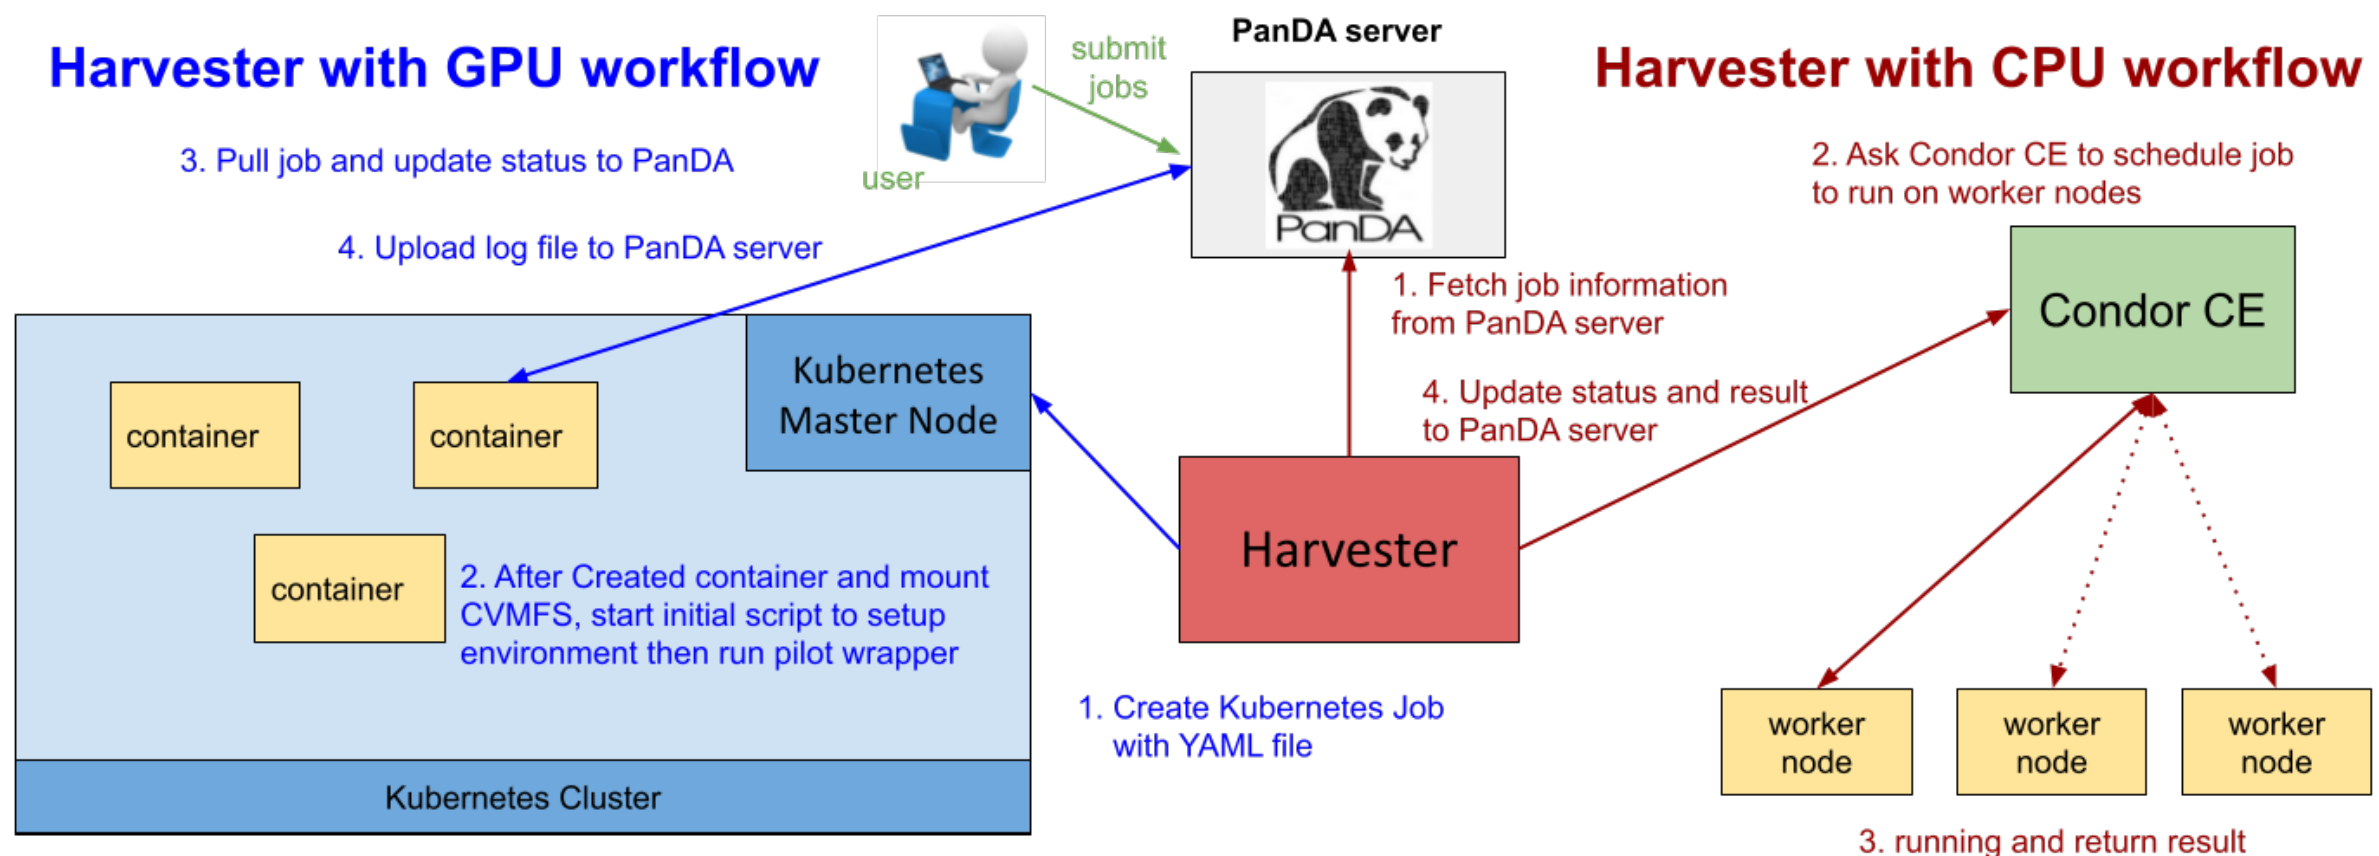
All software used are open-source codes developed by ASGC and an international collaboration led by CERN

WLCG Tier-1/2 @ASGC

- **ATLAS in Taiwan**
 - Achievements: Higgs boson; Dark matter; Searches for beyond Standard Model
 - Future plan: $H \rightarrow b\bar{b}$; Di-Higgs; Dark matter
- **Computing Resource**
 - Tier-1: 4,064 CPU Cores
 - Federated Taiwan Tier2: 664 CPU Cores
 - GPU resource is under validation
- **Storage Resource: 10PB managed by DPM**
- **Activities for ATLAS**
 - Finished 891 billion events, 282 PB in 2020 (#processed data and MC events)
 - Support Folding@Home for COVID-19 studies
- **Contributions to ATLAS Software and Computing**
 - Participating development of ATLAS Harvester/Panda and RUCIO
 - Deeply involved with ATLAS data preparation activities

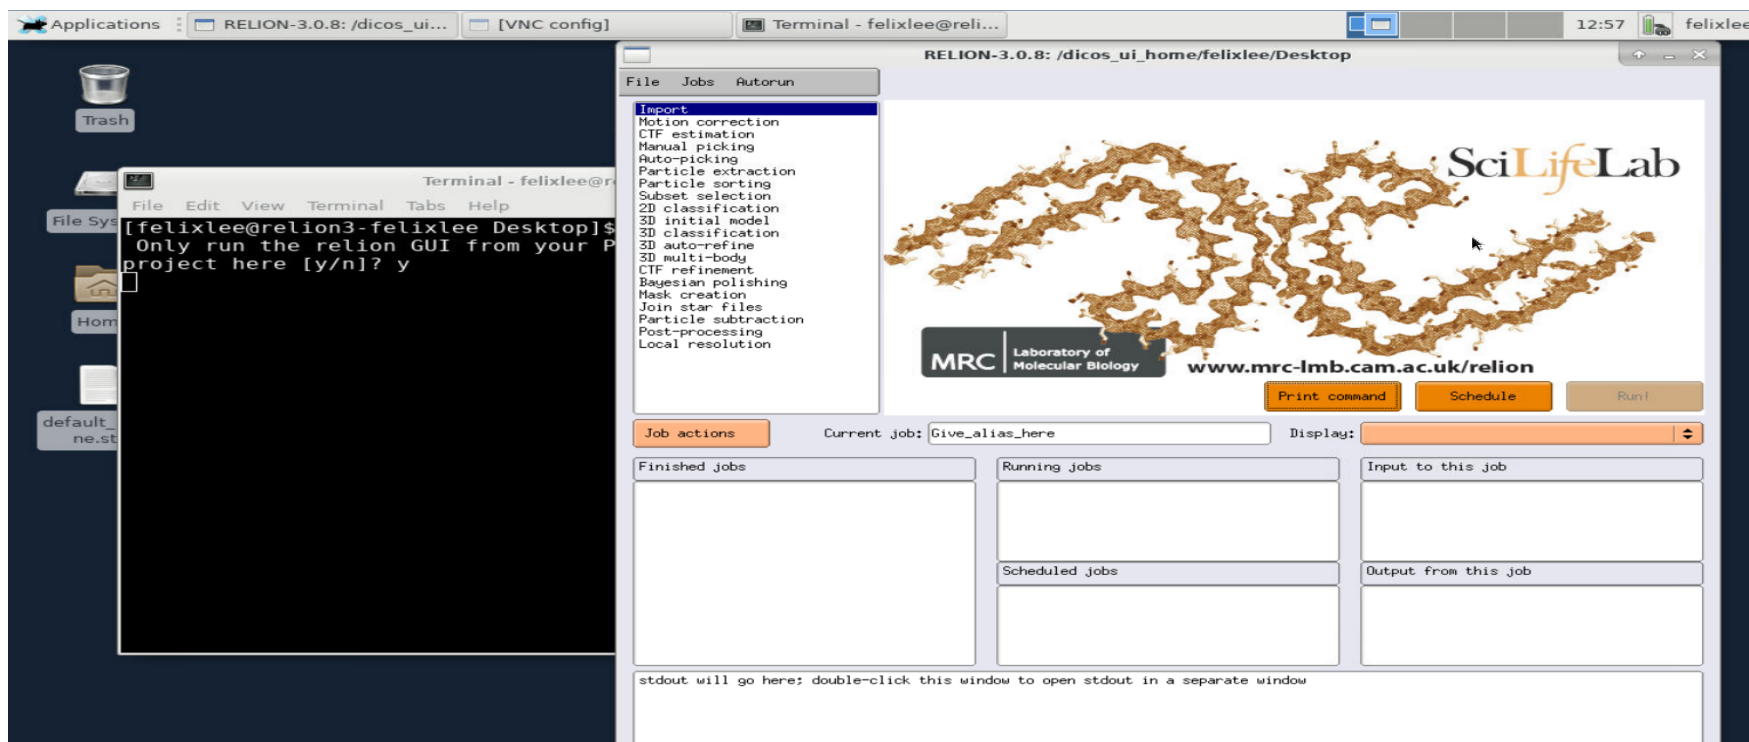
Building Distributed Cloud Infrastructure Supporting broader Scientific Applications based on WLCG

- 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, and ML/DL applications.



Cloud services

- Openstack
 - For Generic services and on-demand worker nodes.
 - OS Upgrade: CentOS7 → CentOS8(oops)
 - ~~Upgrade~~ Migrate from Newton to Ussuri
- Kubernetes
 - For GPU cloud and other core services
 - Offering batch, interactive GUI jobs and services
 - Such as remote Jupyterlab and virtual desktop



Storage services

- **Two Ceph instances**


- For Openstack
 - Upgraded from Hammer to Nautilus
 - Mainly provides RBD
 - 1.4PB
- For users.
 - Luminous
 - Mainly provides Cephfs for our local users.
 - 4.37PB




System Efficiency Optimization

- Scope: Power, Thermal and Distributed Cloud System management
- Strategy: intelligent monitoring and control through ML
- Example: Thermal management, Compute/storage/network anomaly detection, Power saving of worknodes
- 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
- System Anomaly Detection
 - Classify machine status into 5 clusters daily: based on CPU-user, CPU-wio, CPU-system, CPU-idle, Network In/Out
 - 146 events in 14 types identified during March 2020 - March 2021


AHU Performance Monitoring



Worknode Monitoring: High ratio of WaitIO & System

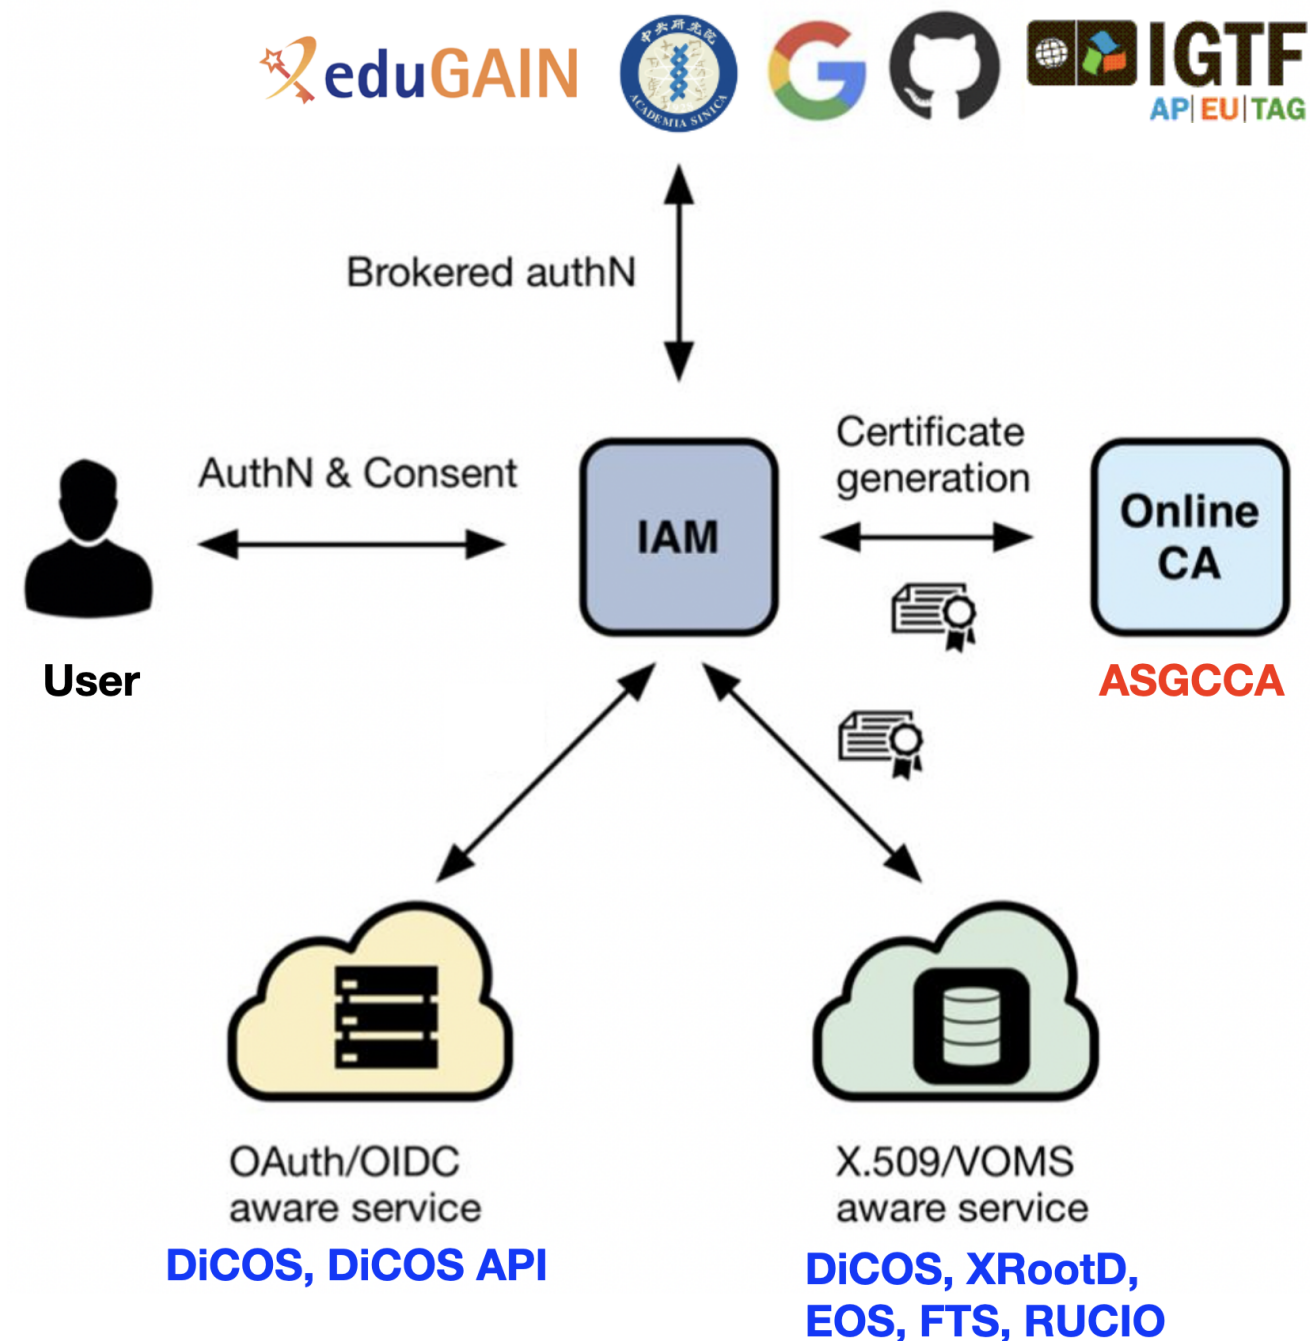


Worknode Monitoring: Misconfigured hyper-threading



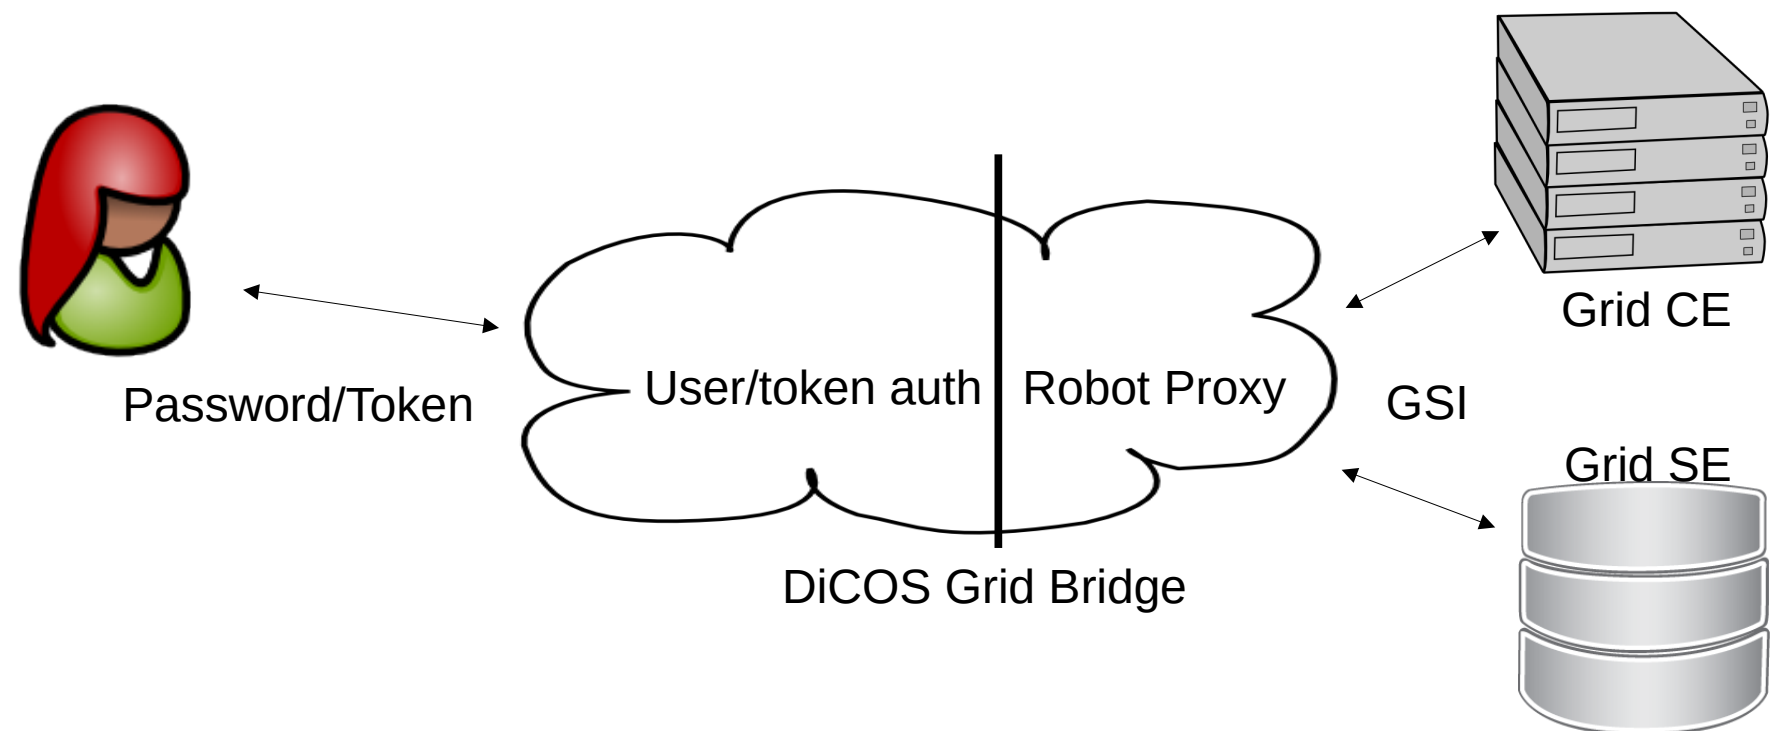
Token-based AAI

- **Basic Services according to community needs**
 - User Identity, Community Attribute Services, Access Protocol Translation, Authorisation, and End Services
- **User's home institute credential based authentication**
 - Academia Sinica SSO, SAML, OpenID Connect (Google, Microsoft, Github) and X.509 certificates
- **Supporting community AAI**
 - Either local community or international community
 - Trust building between users and service providers
- **Supporting federation of infrastructure services: identity management, access, resource, etc.**
- **Improving AAI services based on user experiences over IAM**



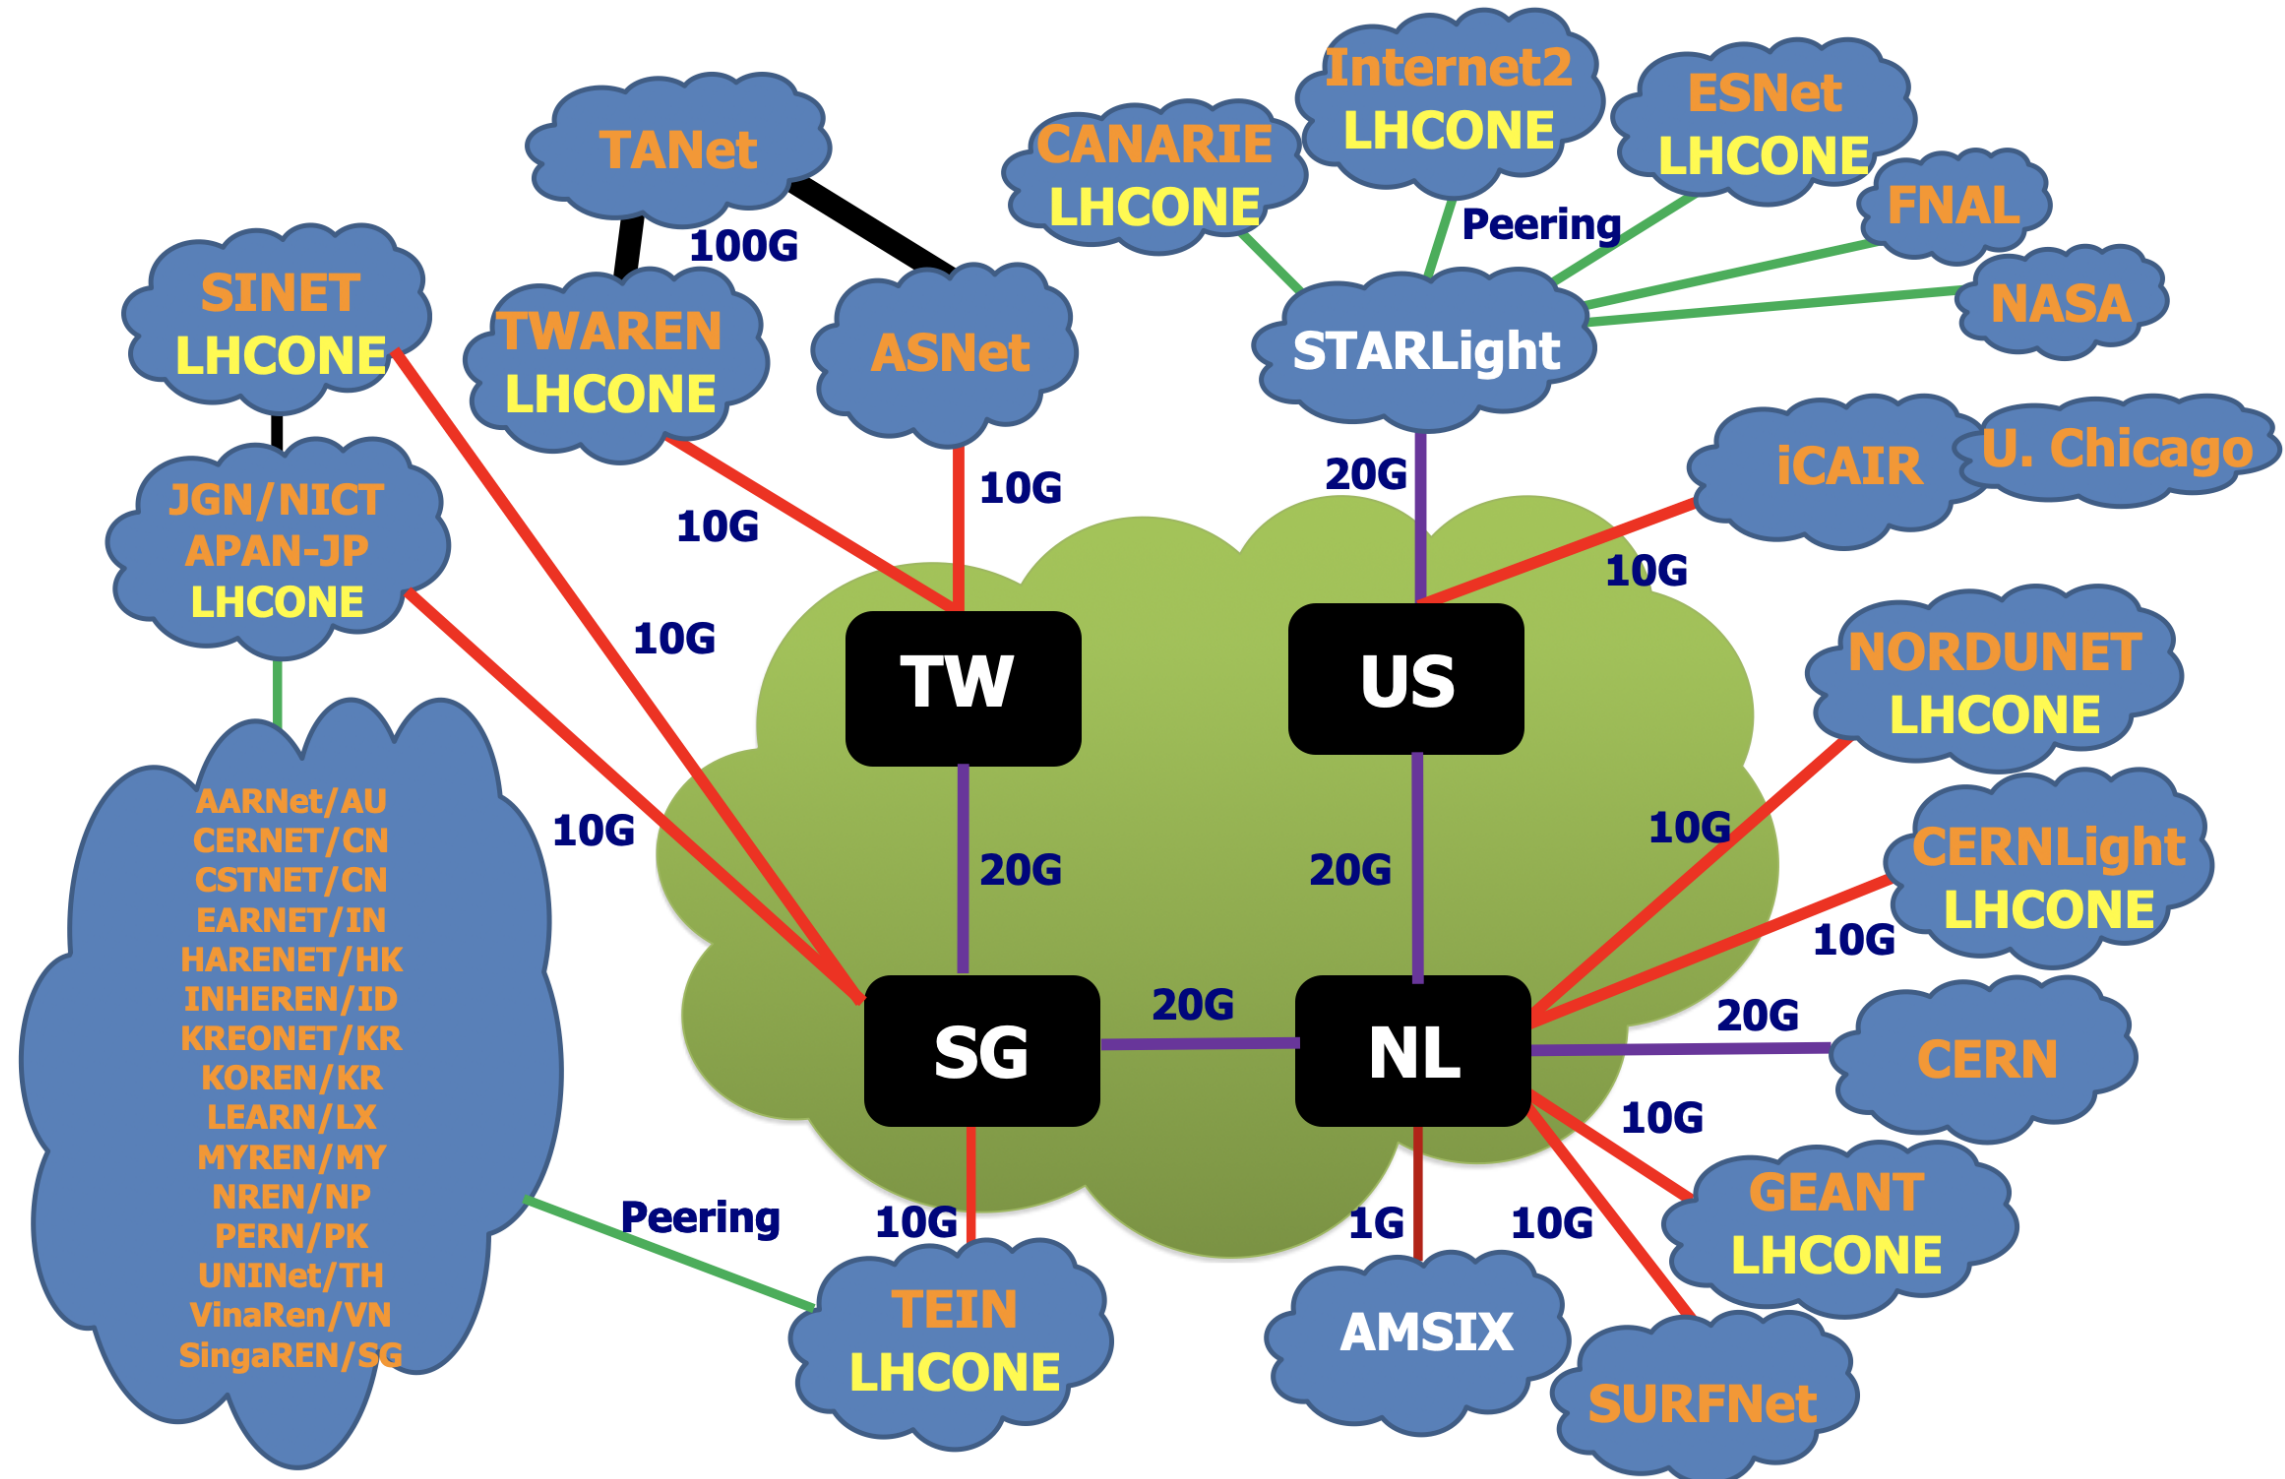
GSI -> token transition

- DiCOS Grid bridge acts as middleware to do GSI authentication for users.
- User doesn't need to manage Grid proxy anymore.



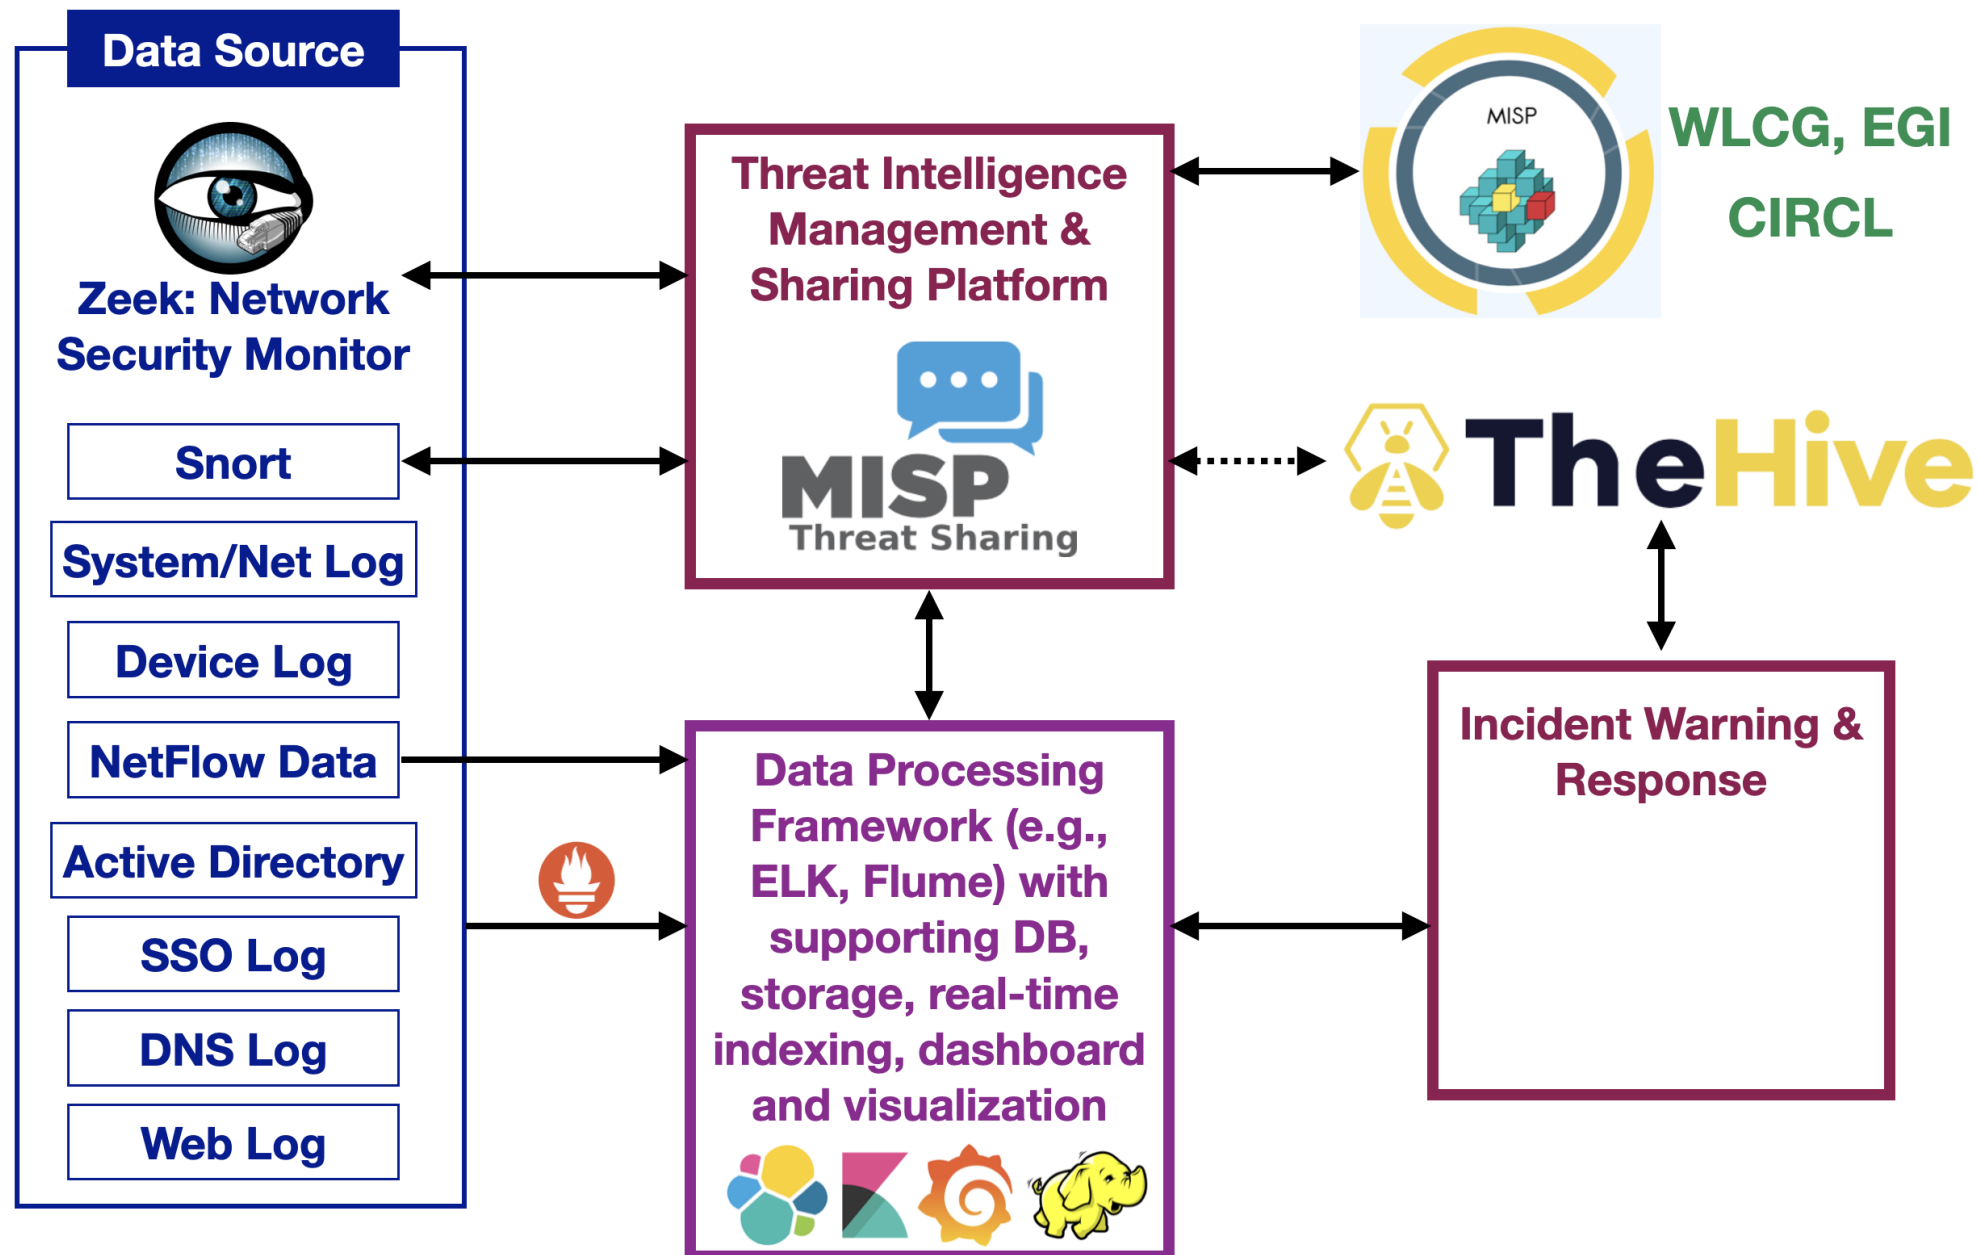
ASGCNet is responsible for providing high-throughput research network from Taiwan to Europe and Asia

- **Achieved 19.5Gb/s and 19.6Gb/s In/Out performance respectively at the same time over 2x10Gb/s international backbone between TW-SG-NL-CERN**
- **Automatic backup with JGN and TEIN for LHCONE/LHCOPN traffic**



WLCG SOC

- All components are in place and operating
- Improving intelligence and efficiency of the data collection, ingestion, analytics and storage





- **Thanks a lot!**