ASGC Site Report

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

HEPiX Autumn 2020
12 Oct, 2020

Academia Sinica Grid Computing Centre (ASGC)

- **Mission: Research & develop advanced distributed computing infrastructure and technologies for science.**
- **Strategy: Collaboration, Openness (Open Science, Resource Federation) and Efficiency**
- **Goals**
  - **Acting as a Tier1 center of the WLCG and jointly developing advanced worldwide grid (WLCG)**
  - **Extending WLCG core technologies to facilitate big data analysis efficiency of broader scientific disciplines in AS and Taiwan by DiCOS (Distributed Cloud Operating System)**
  - **Optimizing system efficiency (incl. power, thermal, research applications & resources, operation etc.) by intelligent monitoring and control mechanism.**

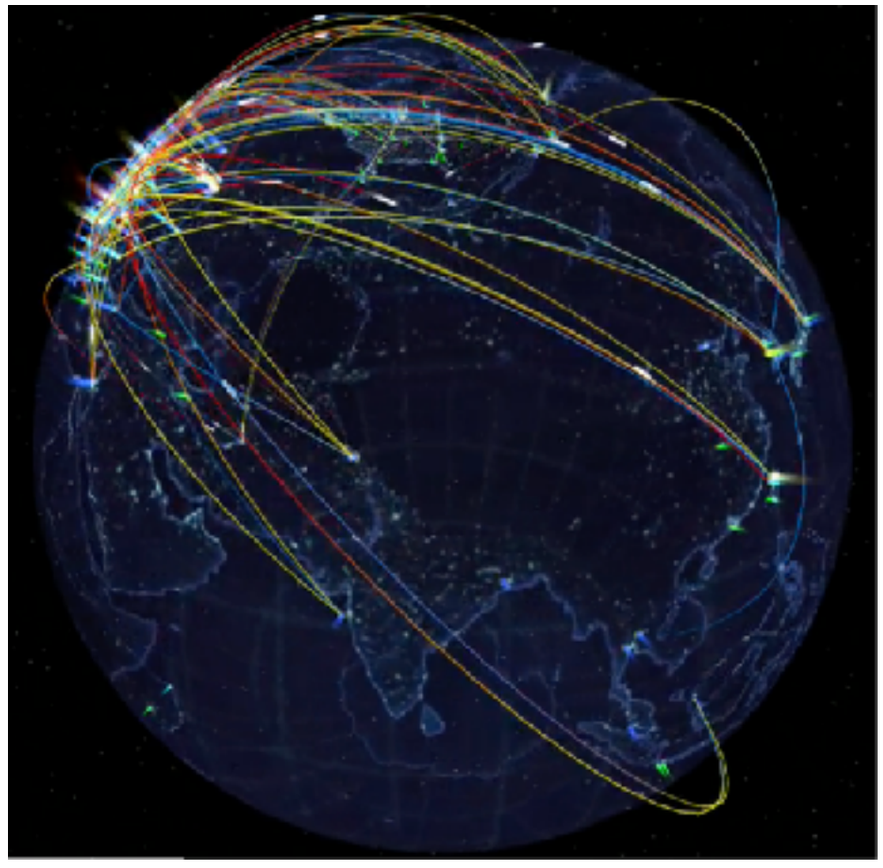

ASGC Resources

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

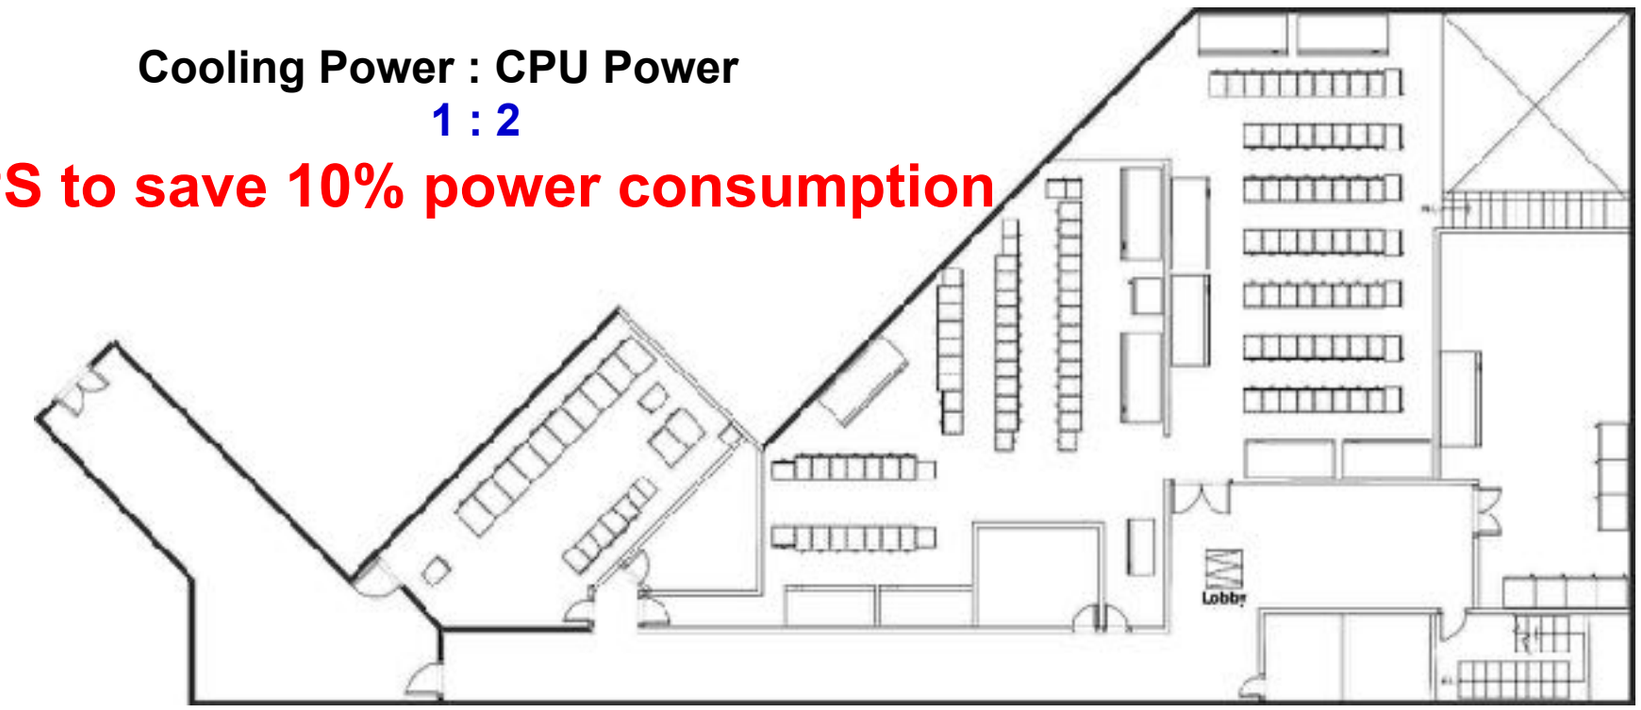

- **Total Capacity**
- 2MW, 400 tons AHUs
- 112 racks
- ~ 800 m²
- **Resources (by end 2018)**
- 20,922 CPU Cores
- 532,480 GPU Cores + 251,904
- 30,720 Tensor Cores (GPU) + 7,872
- 18.8 PB Disk Storage
- **Rack Space Usage (Racks)**
- AS: 59.2 (53.5%)
  - WLCG: 17
  - e-Science: 42.2
- ASIoP: 10.4 (9.4%)
- RCEC: 6.5 (5.9%)
- ASCC: 3.3 (3.0%)
- AI School: 2.2 (1.9%)
- Free: 28.7 (26.3%)

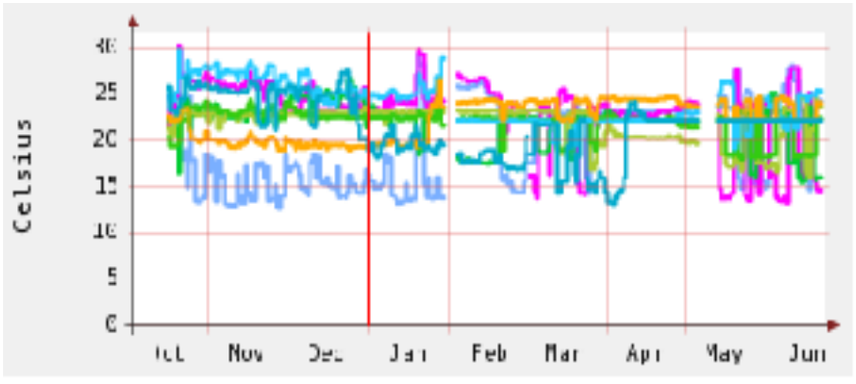

Monitoring the power consumption and temperature of every piece of equipment every 10 seconds.

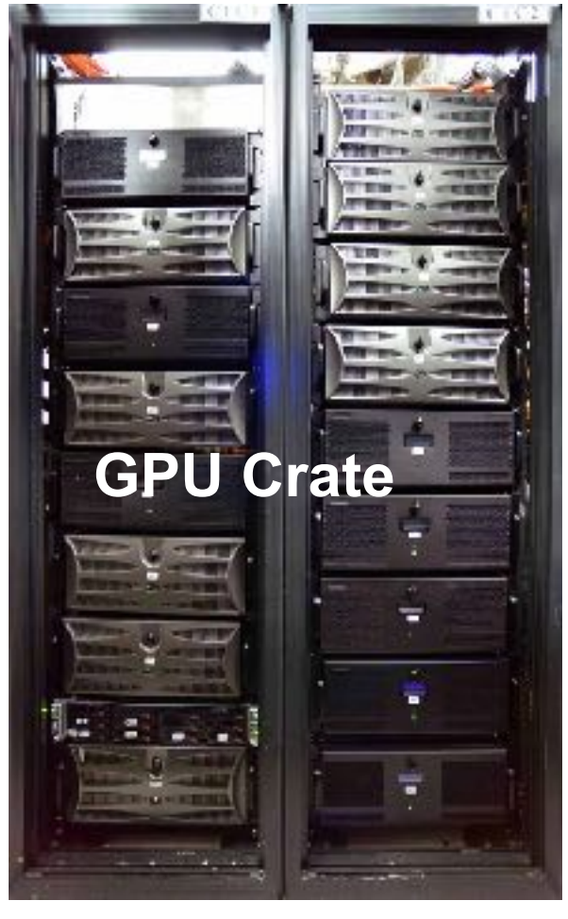

All software used are open-source codes developed by ASGC and an international collaboration led by CERN

3

ASGC Resource Usage Statistics

- CryoEM used 154K+ CPU & GPU-days computing resources in 2020.
- GPU usage is about 2x growth in 2020 comparing with 2019.

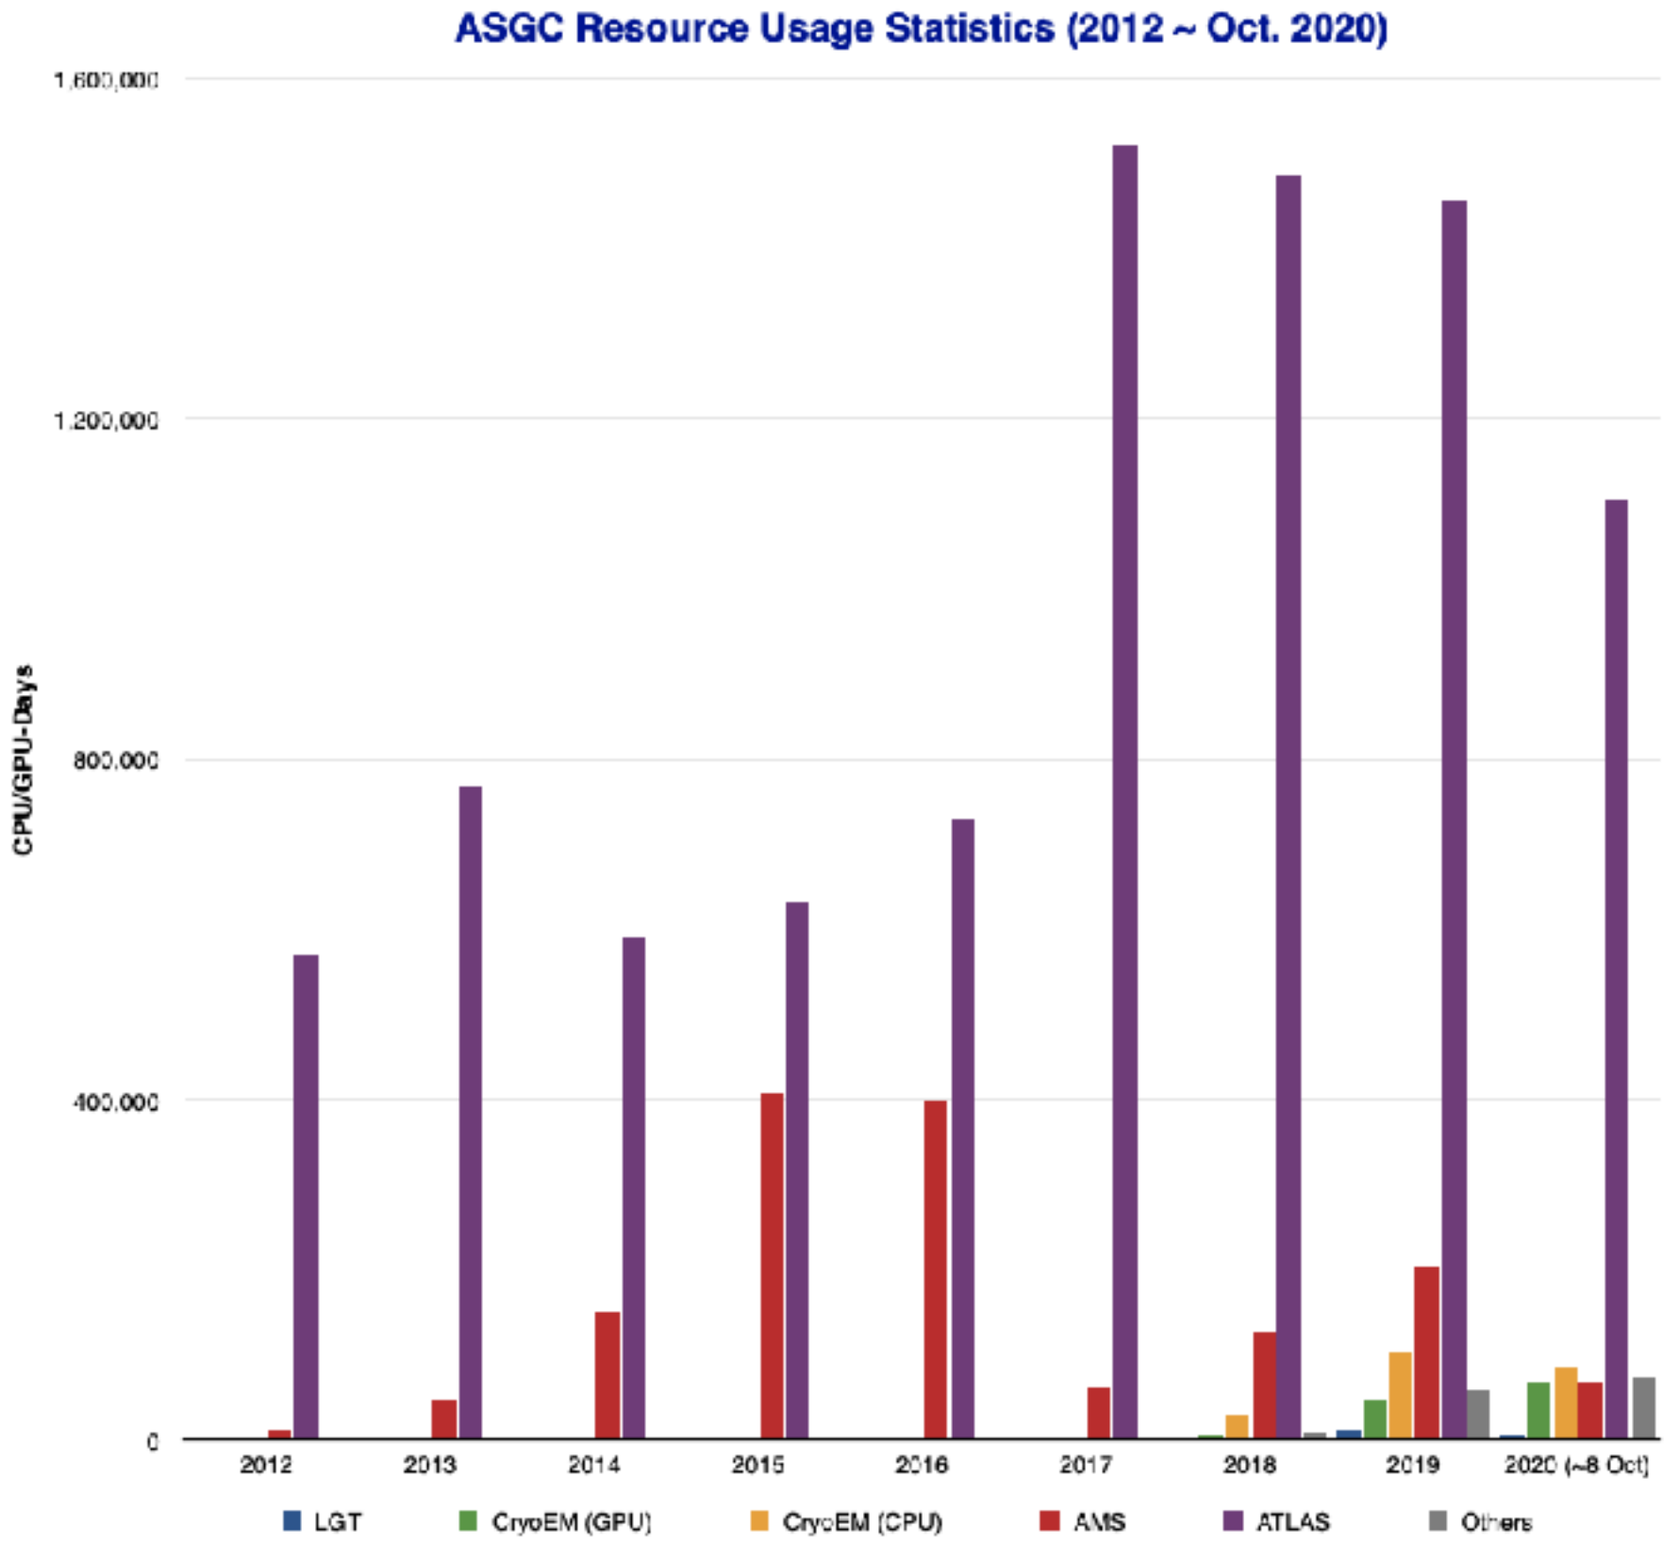

WLCG T1/T2 @ASGC

- WLCG T1/T2 resources (ATLAS): Pledged resources of 2021 will be online in Q1@2021
  - ~400 nodes, 6,000 CPU-cores,
  - TW-ASGC T1: 7.1 PB disks, 44.3K HEP-Spec06
  - TW-FTT T2: 0.9 PB Disk, 9K HEP-Spec06
  - DPM throughput: ~200TB/day (inbound + outbound) data transmission; >1GB/s throughput/storage server
  - Able to accomplish 1.5M+ jobs/month
  - Able to saturate 2x10Gb International links by ~19.8/s for inbound and outbound traffic at the same time
  - All services are IPv6 enabled by dual stack
- For RUN3
  - DPM DOME/1.14 is available;
  - DOMA/TPC (XRootD, HTTP) is ready
  - GPU for ATLAS will be provisioned from 2021
- Joint development of Harvester and Panda
  - Harvester/Condor containerization and Integration of K8S for ATLAS
  - Harvester cluster with K8S @ASGC
- Efficiency Improvement: focus on both system and operation efficiency
- Token-based AAI (IAM) has been implement and will be integrated with WLCG and e-Science applications (testbed established)

	CA	CERN	DE	ES	FR	IT	ND	NL	RU	TW	UK	US
CA	99%	99%	92%	77%	90%	97%	98%	98%	92%	100%	96%	98%
CERN	100%	100%	95%	100%	96%	99%	100%	98%	84%	100%	94%	99%
DE	56%	98%	95%	99%	75%	62%	98%	99%	97%	97%	96%	96%
ES	96%	86%	100%	94%	65%	40%	98%	98%	99%	100%	97%	97%
FR	92%	51%	81%	82%	57%	30%	70%	81%	67%	43%	82%	78%
IT	70%	91%	87%	82%	99%	93%	75%	90%	76%	87%	74%	77%
ND	100%	100%	96%	91%	99%	100%	70%	96%	94%	100%	100%	92%
NL	100%	100%	97%	100%	97%	97%	98%	97%	99%	100%	93%	90%
RU	84%	80%	84%	90%	80%	90%	82%	94%	97%	87%	82%	78%
TW	100%	100%	92%	100%	99%	94%	98%	100%	100%	-	98%	99%
UK	98%	98%	98%	99%	78%	55%	97%	83%	92%	98%	66%	92%
US	98%	95%	92%	95%	83%	58%	93%	85%	96%	94%	95%	94%

DiCOS

- Enable high-performance scientific computing by distributed cloud, driven by research applications of AS
  - Based on ATLAS software (Harvester, PanDA, Rucio, EOS, FTS, XRootD, CVMFS, ARC-CE, etc.), as well as CERNBox, HTCondor, Docker and JupyterLab
  - Ceph serves as the generic storage pools (~5PB currently)
  - AAI: IAM implemented and is integrating with Application workflow
- Integration Model: application-specific web portal integrates computing model over distributed cloud efficiently
  - Interactive applications: Web portal, K8S, container, Ceph
  - Batch: Web portal, Harvester/Panda, RUCIO, EOS/Ceph
  - Web terminal and CLI are also supported by default
- Collaborations: CryoEM, BioSAXS, NGS, Drug Discovery (Docking), Earth Science (Tsunami, Weather, Storm Surge, Earthquake), Ecology/Biodiversity, Gravitational Wave (KAGRA/IGWN), Lattice Gauge Theory, Computational Quantum Mechanics (VASP), Brain Image Analysis, Condense Matter, AMS, TEXONO
- Lessons Learned
  - MPI jobs: configuration for Harvester and flexibility should be improved;
  - JEDI might be not needed for non-WLCG applications for now
  - Ever experienced slow, faulty, and multiple data movements on big input file by RUCIO (> 100GB)
  - Web portal is the best model integrating workflow, UI, user applications, various resources & platform (e.g., VM, container) and computing models
- Future plan
  - PRMON deployment is underway, and will be integrated with Prometheus, Kibana, etc.
  - System efficiency optimization
  - Interoperability with WLCG and EGI/EOSC
  - Acts as a reference platform for Asia regional collaborations
  - Archive and Community-based features for: reproduction, repurposing, search & access, knowledge base
  - Cloud-based ML platform & data analytics tools
  - Educational materials: better utilization, integration, case study, benchmark reference, etc.

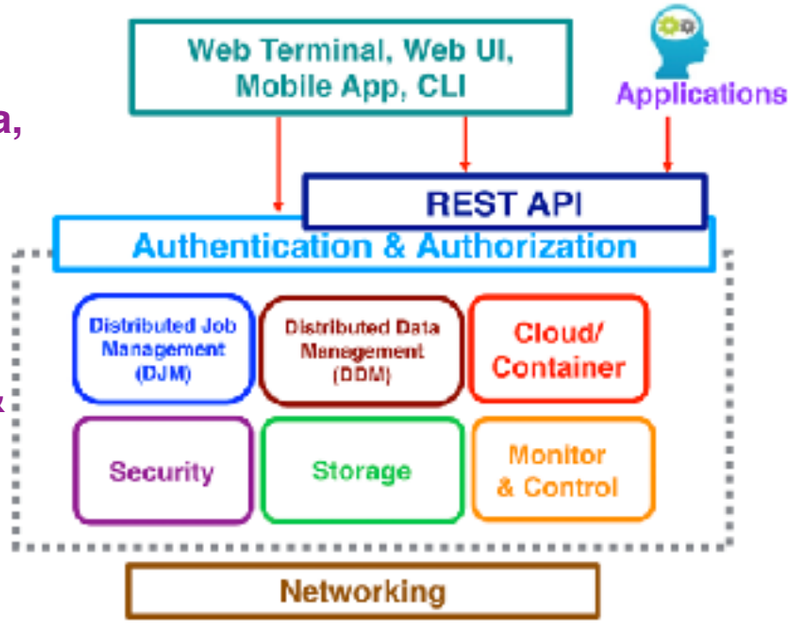

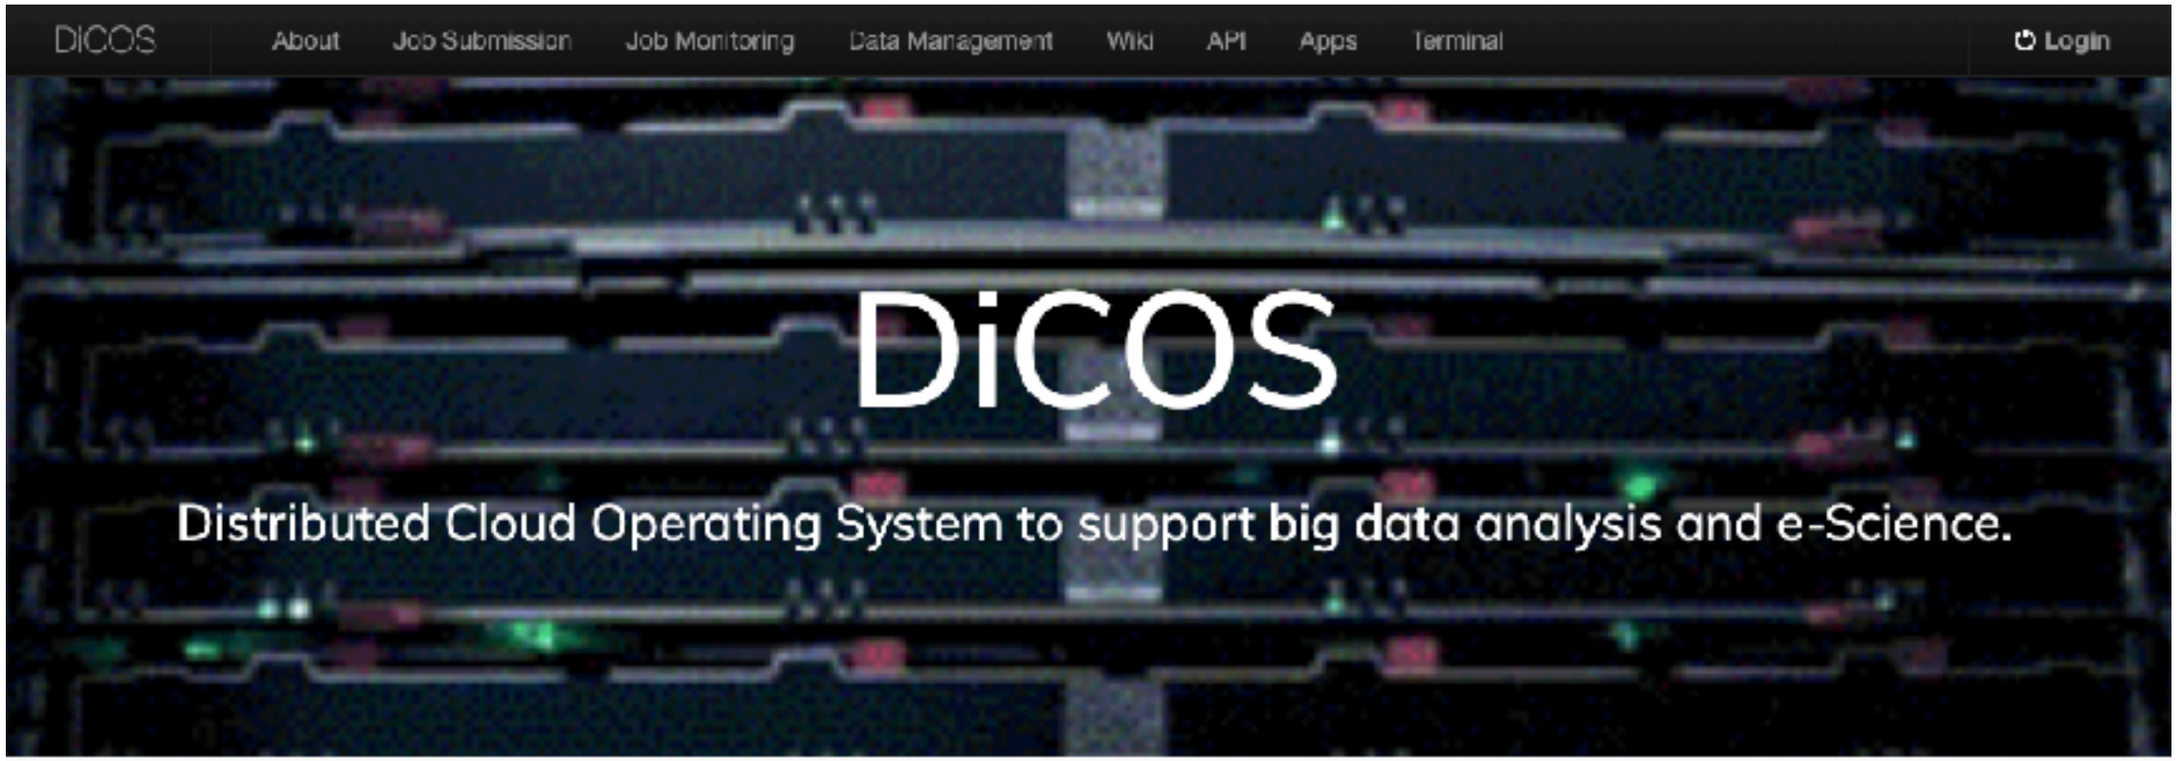

DiCOS

DISTRIBUTED CLOUD OPERATING SYSTEM

https://dicos.twgrid.org

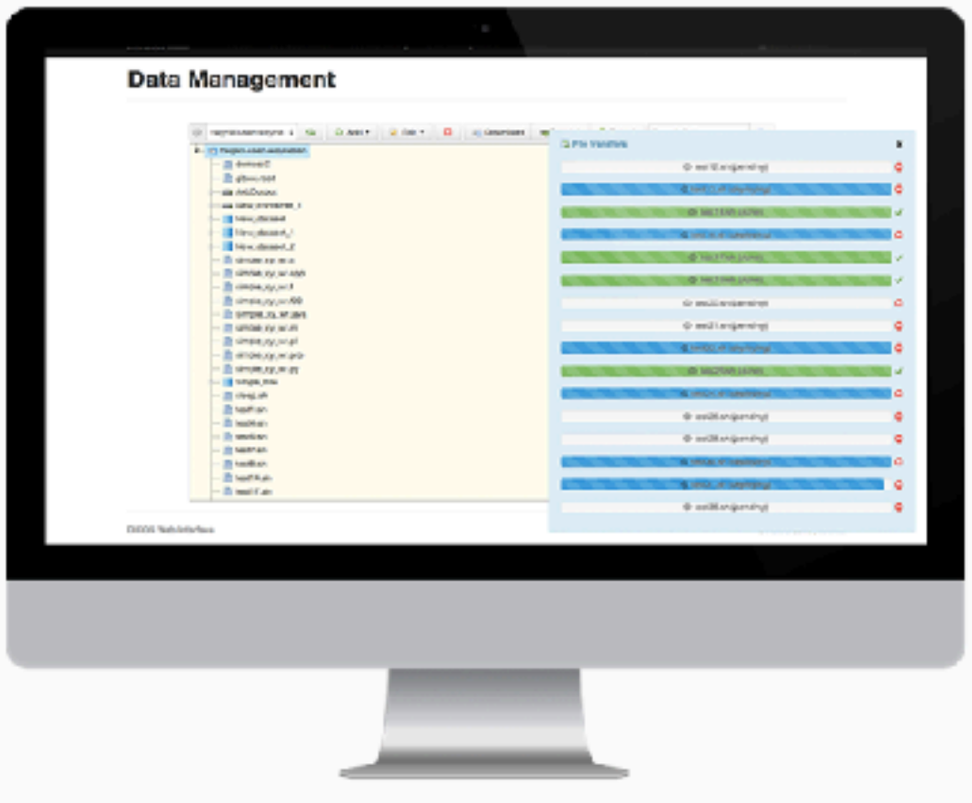

About DiCOS

DiCOS is the operating system for the distributed Cloud environment, to provide scalable infrastructure, flexible platform as well as intelligent data and job tools for high performance scientific and generic computing purposes.

In addition, together with the fanless rack system, DiCOS also serve as the building block of various levels of resource centers for the distributed Cloud.

DiCOS-UserManual (English) DiCOS-UserManual (Chinese)

Supporting Cryo-EM Applications by DiCOS

- Primary GPU (single precision) and CPU (300-600 threads) users, O(TB) input/job (>2,000 jobs in 2020)
- Customization and development
  - Web UI
  - SW package as container
  - Data flow and performance optimization
  - Mass production over DiCOS

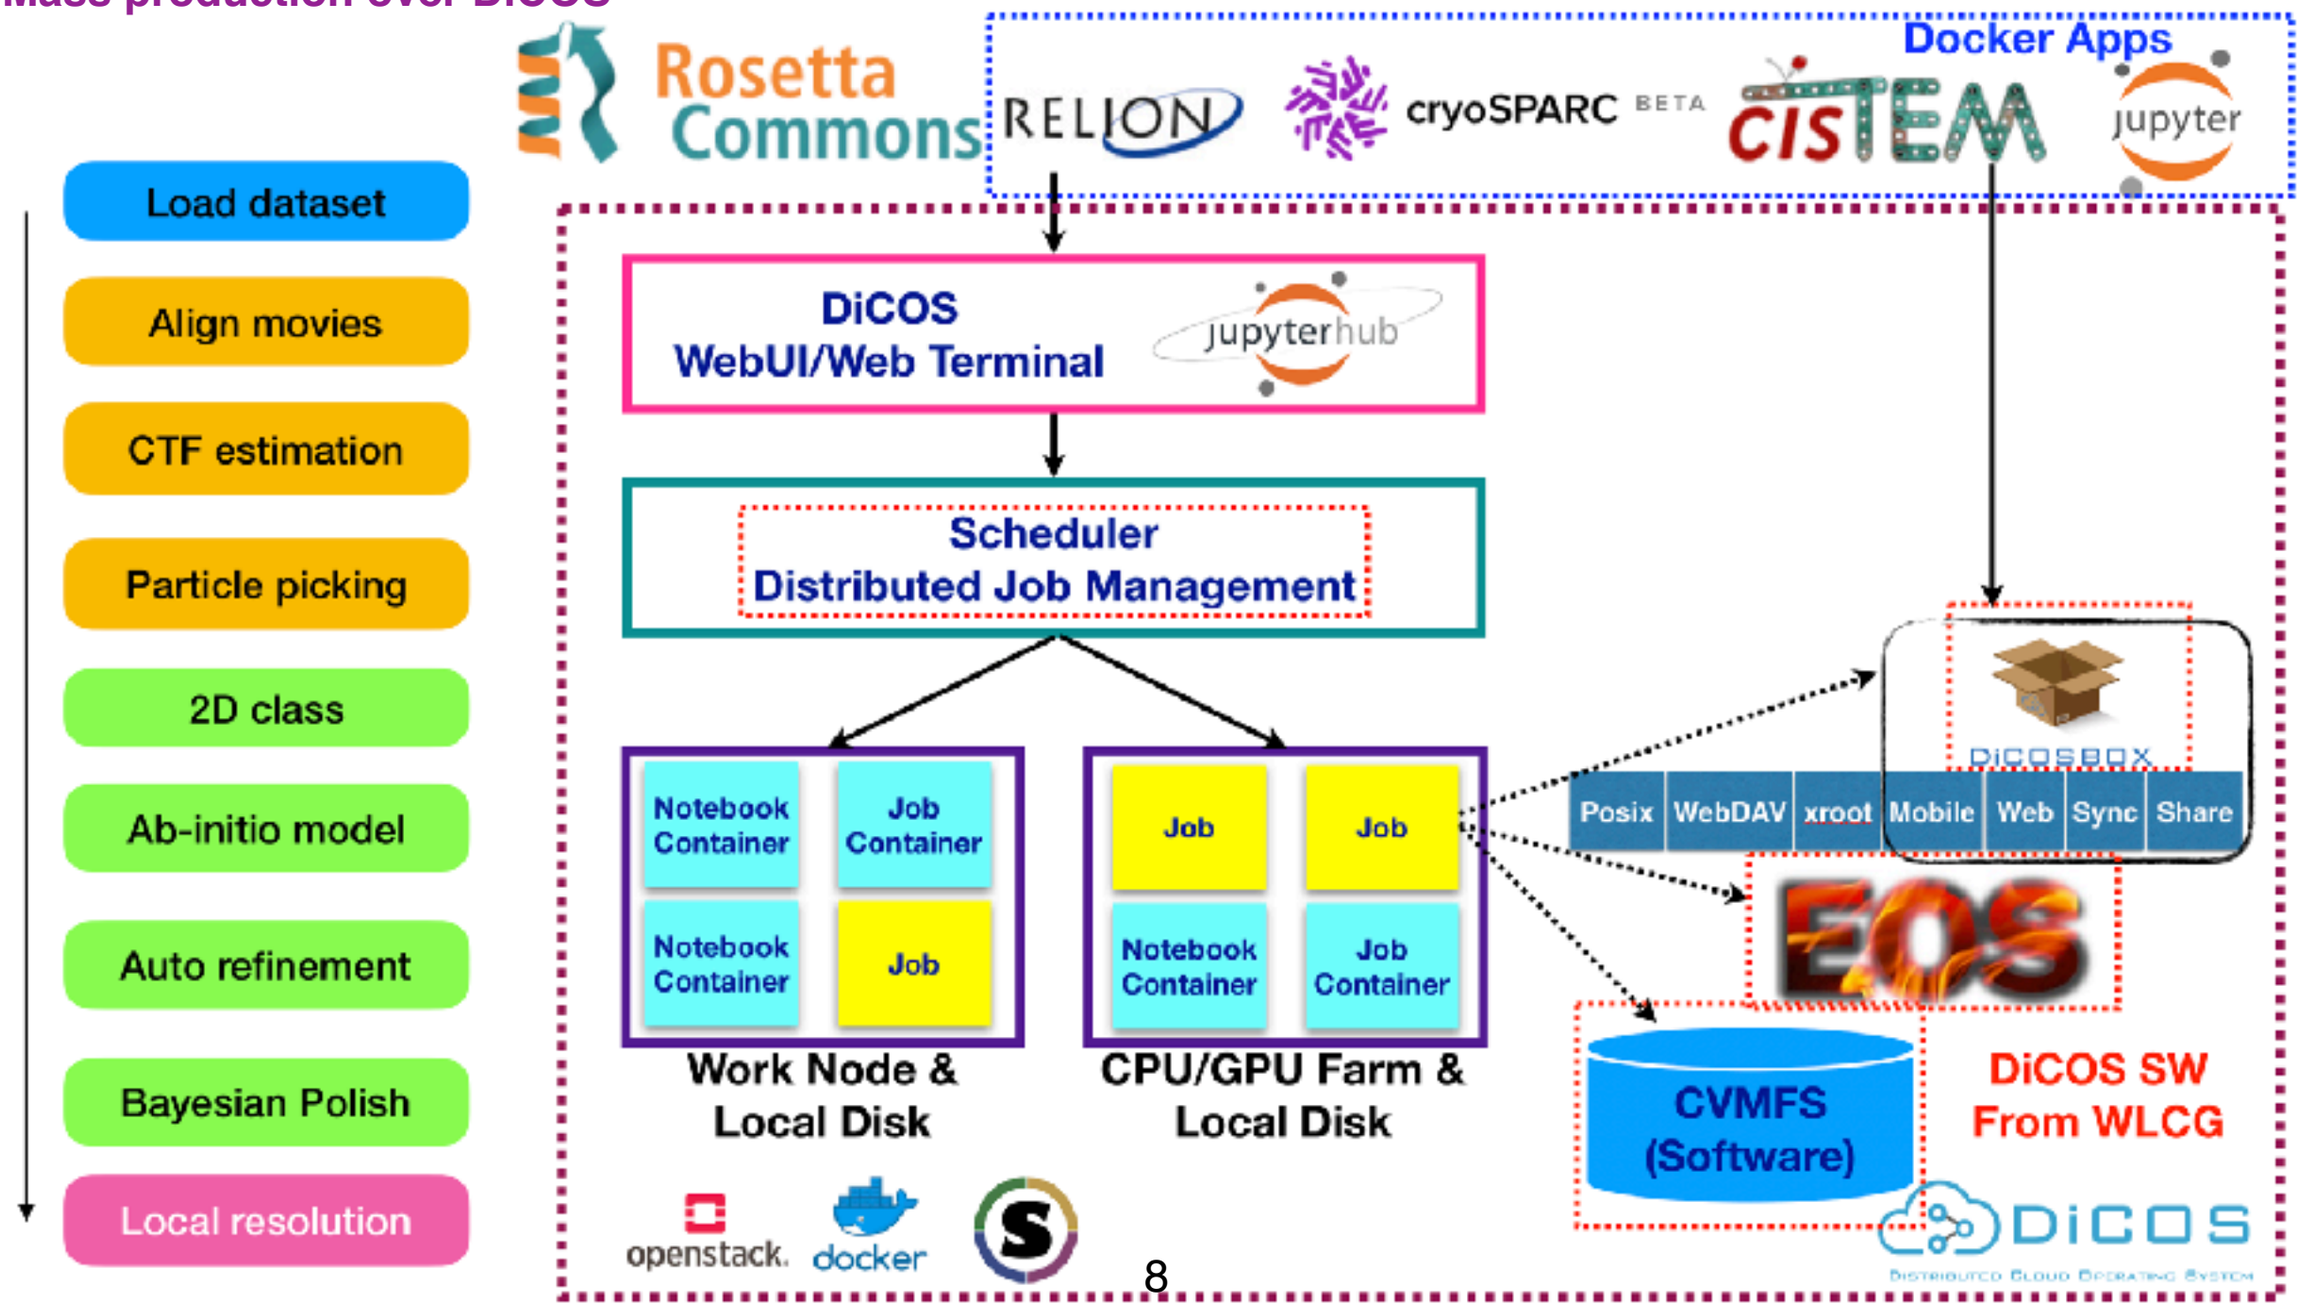

8

CryoEM Applications on DiCOS

- Web portal for CryoSPARC, cisTEM, and RELION
- JupyterLab: interactive environment for coding, analysis, visualization, and output enrichment, etc.

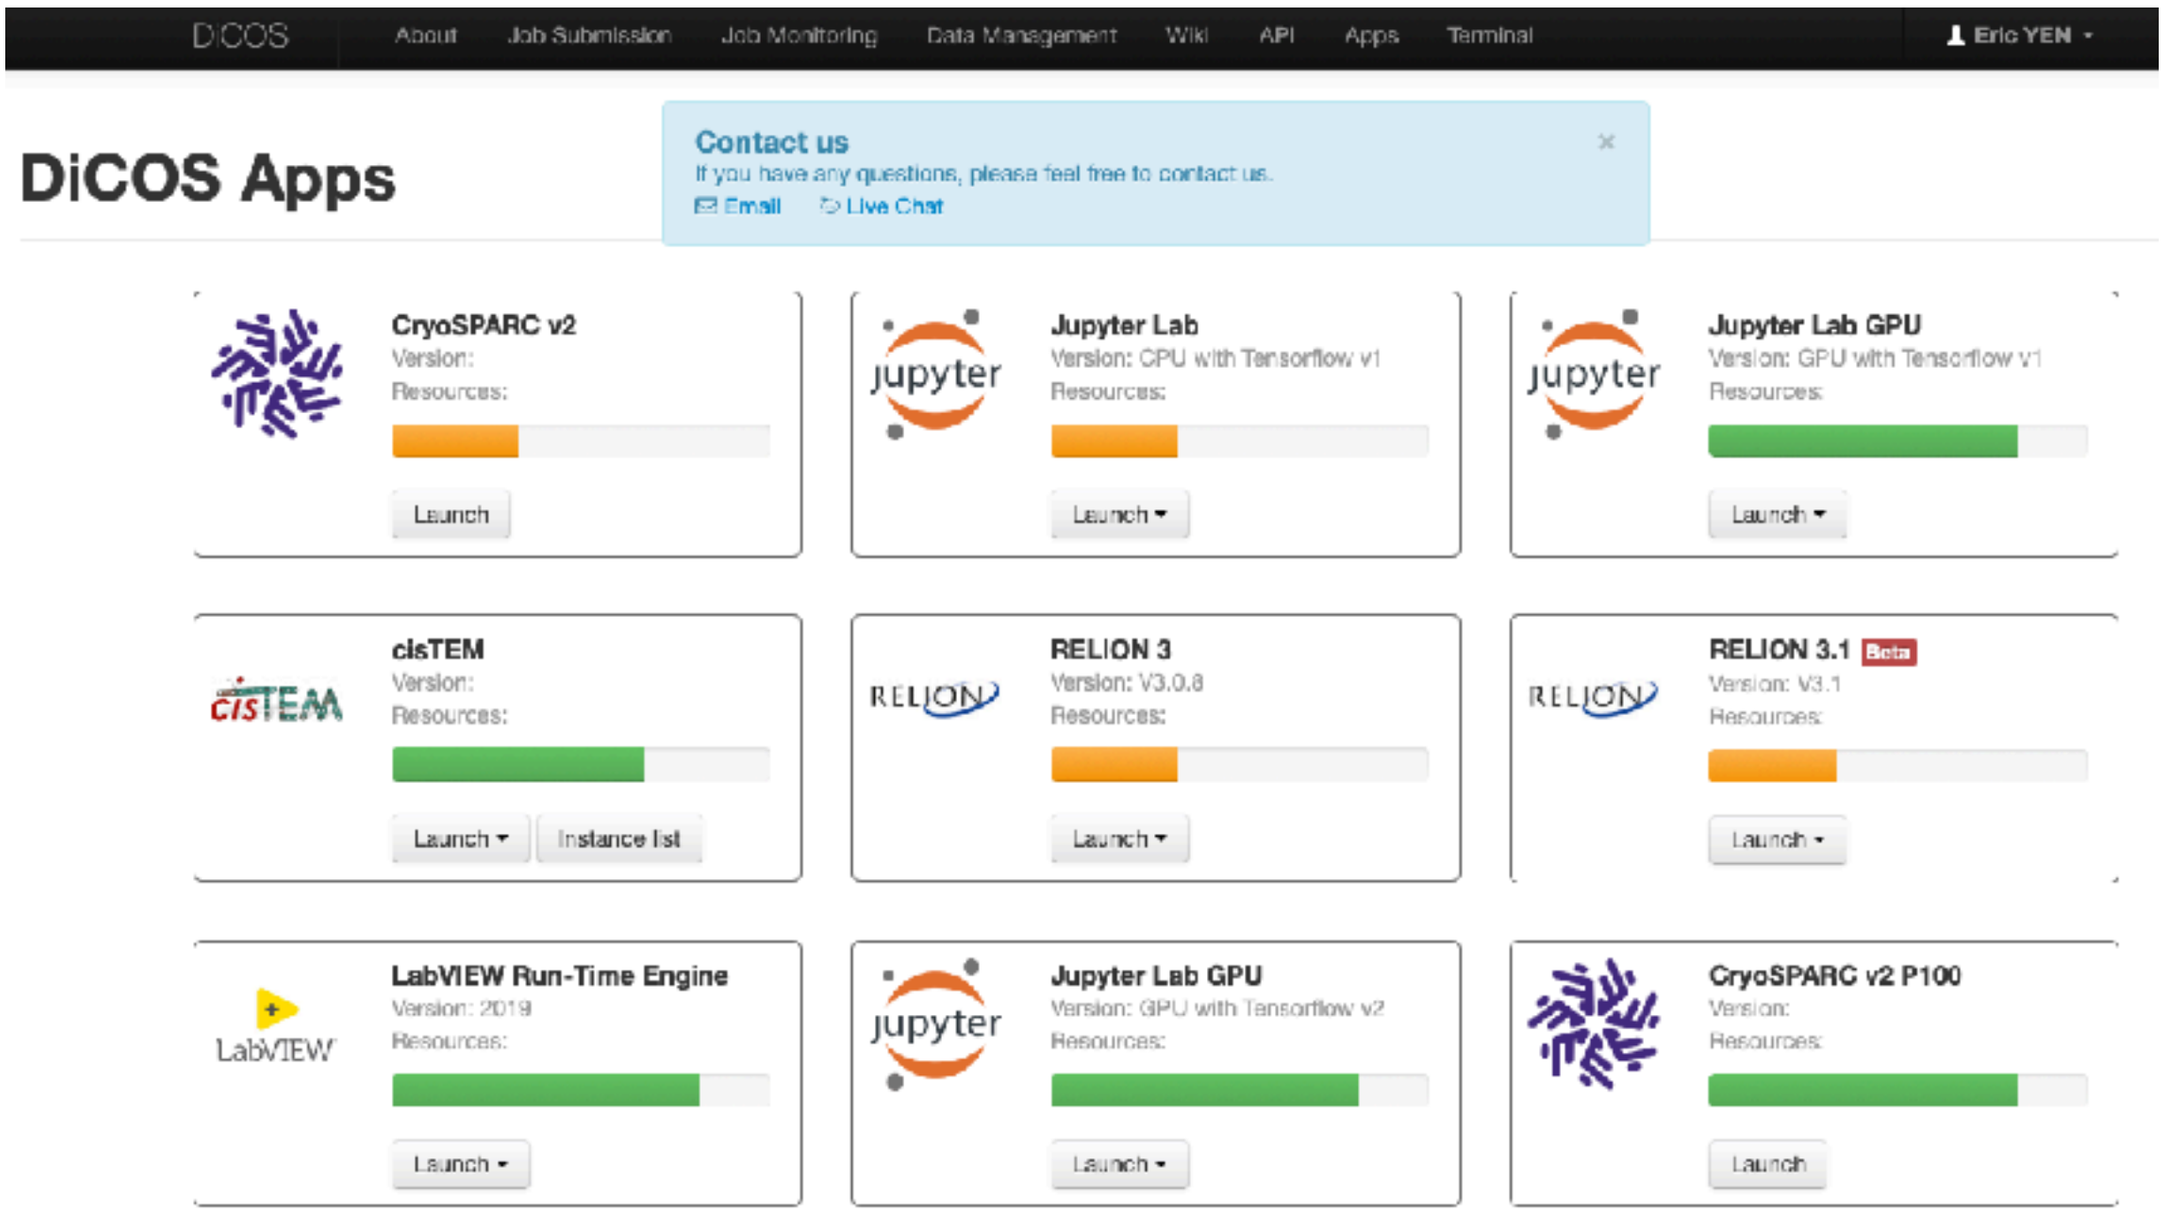

Web Terminal for Command Line Interface

DiCOS Beta | About | Job Submission | Job Monitoring | Data Management | Wiki | API | Contact | Terminal | Felix Lee

```
[felixlee@dicos-ui02 ~]$
[felixlee@dicos-ui02 ~]$ export X509_USER_PROXY=$HOME/.globus/felix.proxy
[felixlee@dicos-ui02 ~]$ voms-proxy-info -all
subject   : /C=TW/O=AS/OU=GRID/CN=Felix Dteam Lee 153280/CN=proxy
issuer    : /C=TW/O=AS/OU=GRID/CN=Felix Dteam Lee 153280
identity  : /C=TW/O=AS/OU=GRID/CN=Felix Dteam Lee 153280
type      : full legacy globus proxy
strength  : 2048
path      : /asgc_ui_home/felix/.globus/felix.proxy
timeleft  : 120:18:56
key usage : Digital Signature, Key Encipherment, Data Encipherment, Key Agreement
=== VO ams02.cern.ch extension information ===
VO        : ams02.cern.ch
subject   : /C=TW/O=AS/OU=GRID/CN=Felix Dteam Lee 153280
issuer    : /C=TW/O=AS/OU=GRID/CN=voms.grid.sinica.edu.tw
attribute : /ams02.cern.ch/Role=production/Capability=NULL
attribute : /ams02.cern.ch/Role=NULL/Capability=NULL
timeleft  : 120:18:56
uri       : voms.grid.sinica.edu.tw:15016

[felixlee@dicos-ui02 ~]$
[felixlee@dicos-ui02 ~]$ dicos job status
   PandaID  Name                        Status
-----------------------------------------------
   2301292  taiwan.873059581.he4.pl1  activated
   2301290  taiwan.873059580.he4.pl1  activated
   2301288  taiwan.873059579.he4.pl1  activated
   2301286  taiwan.873059578.he4.pl1  activated
   2301284  taiwan.873059577.he4.pl1  activated
   2301282  taiwan.873059576.he4.pl1  activated
   2301280  taiwan.873059575.he4.pl1  activated
   2301278  taiwan.873059574.he4.pl1  activated
   2301276  taiwan.873059573.he4.pl1  activated
   2301274  taiwan.873059572.he4.pl1  activated
   2301272  taiwan.873059571.he4.pl1  activated
   2301270  taiwan.873059570.he4.pl1  activated
   2301268  taiwan.873059569.he4.pl1  activated
   2301266  taiwan.873059568.he4.pl1  activated
   2301264  taiwan.873059567.he4.pl1  activated
[felixlee@dicos-ui02 ~]$
```

3:31:37 PM

DiCOSBox for All AS Users

- 1TB per user by default
- User friendly interface for file management and sharing

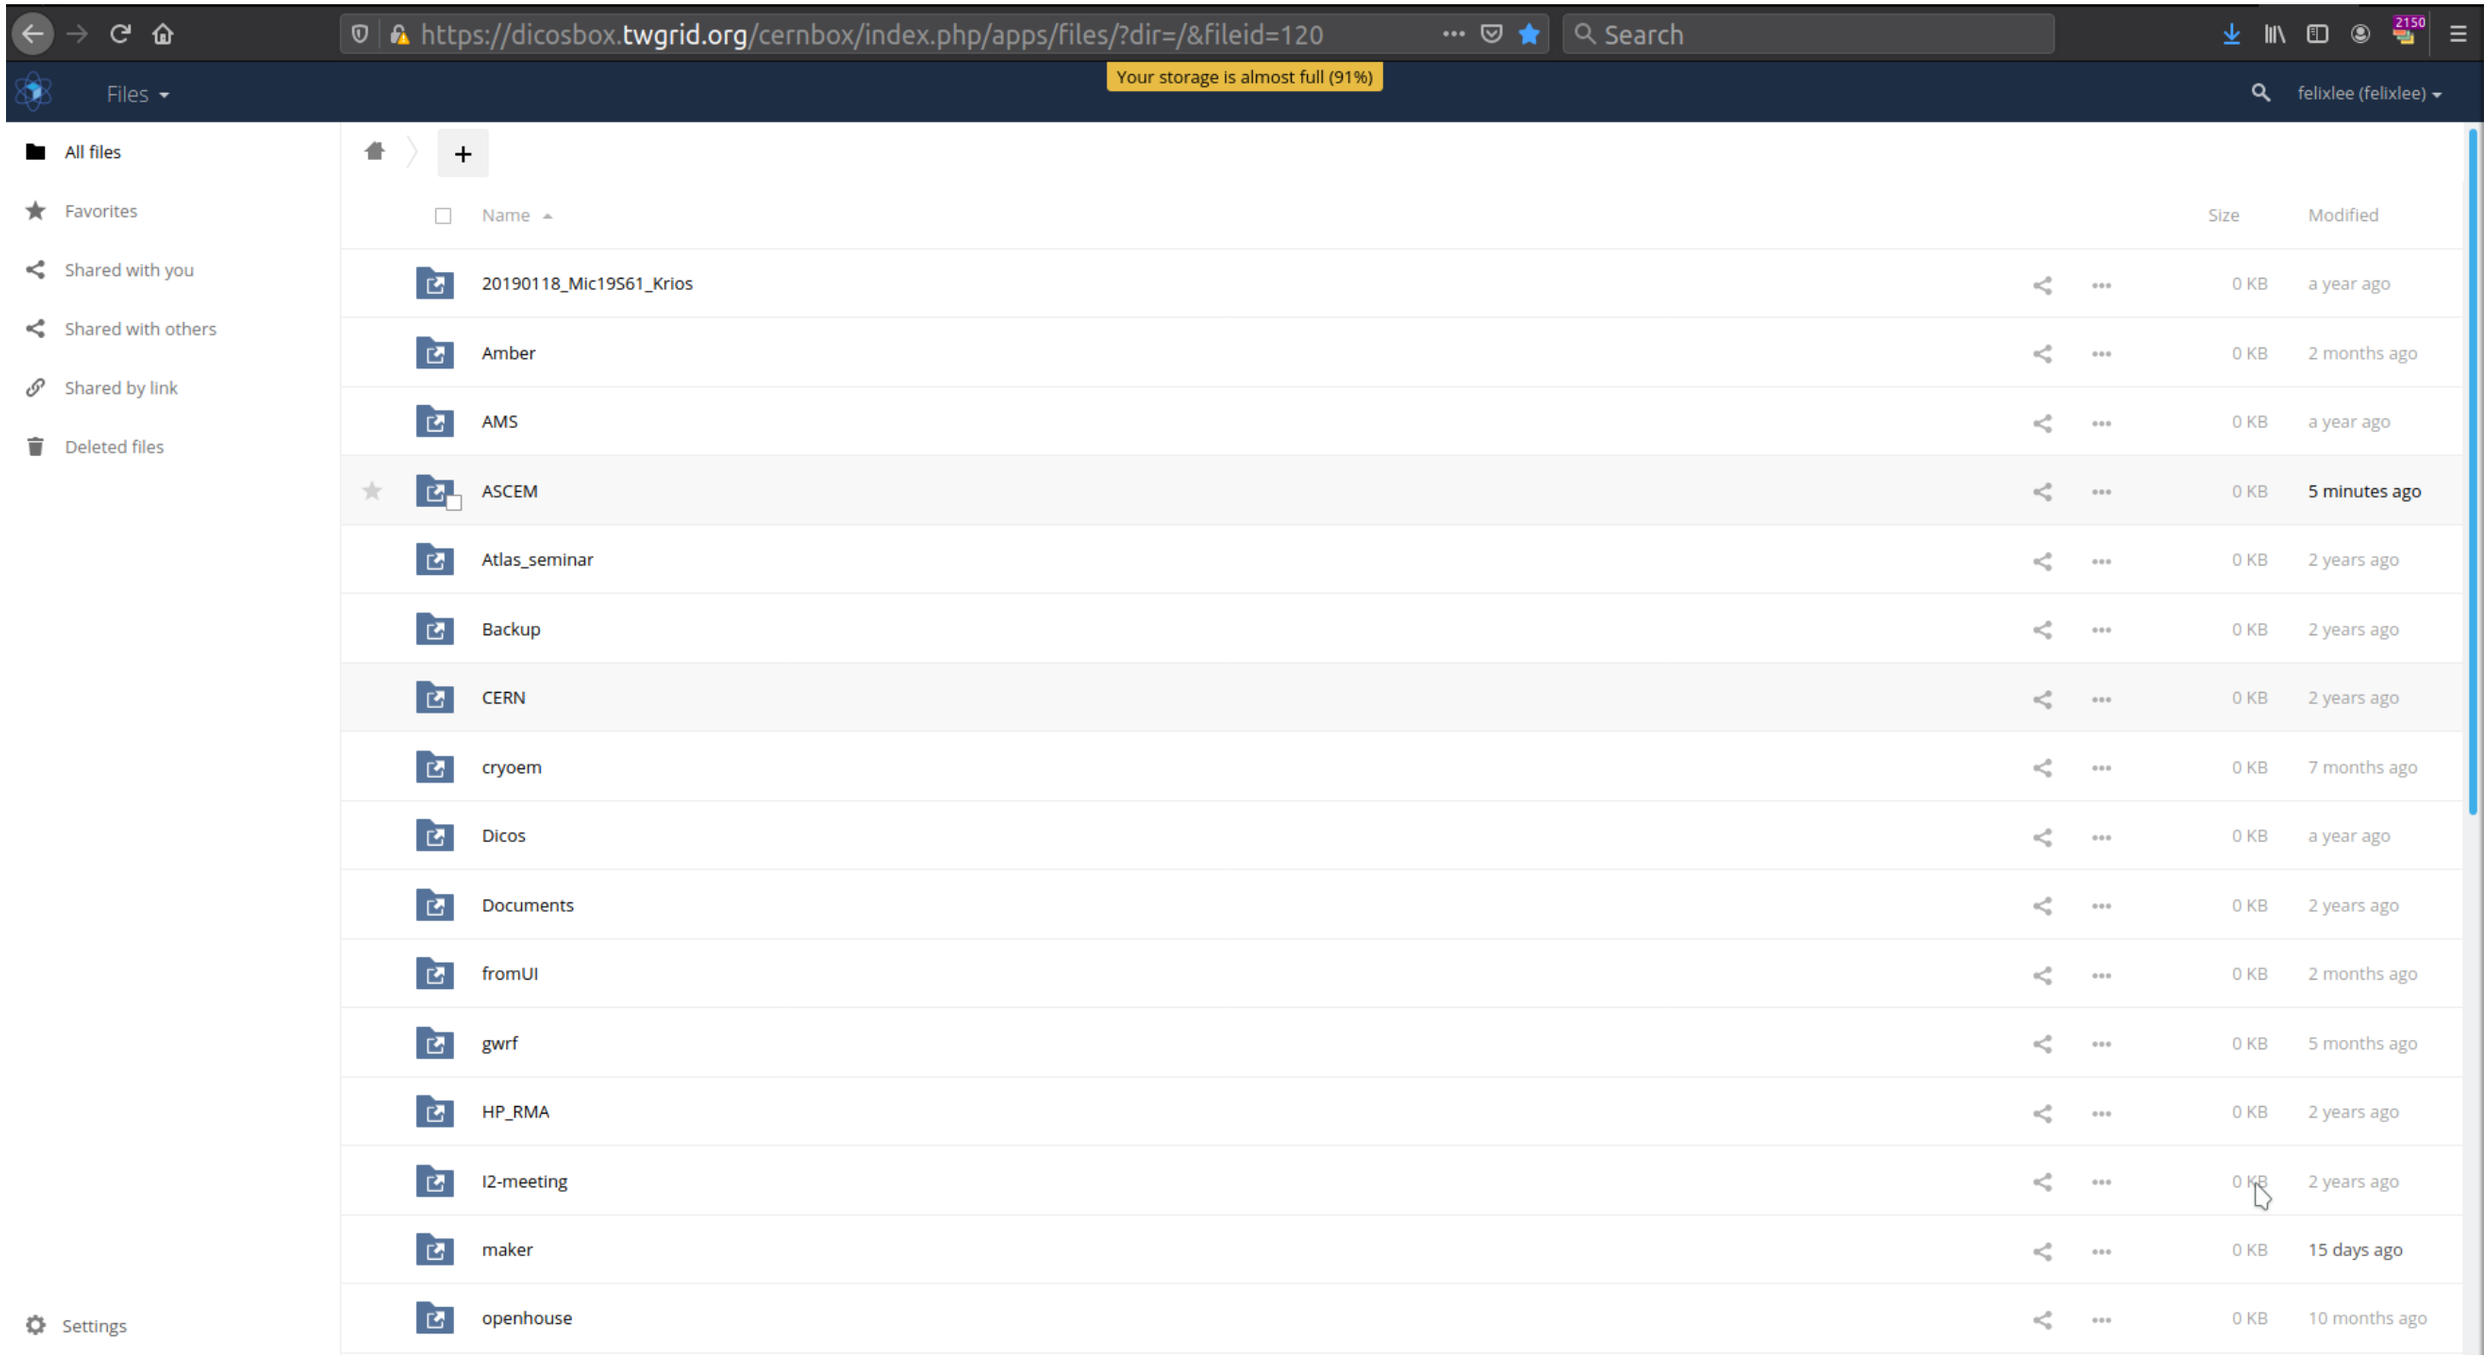

Efficiency Optimization

- Scope: Power, Thermal and System (incl. Operation)
  - e.g., resilient and energy-saving; reduce human errors, manual work and cost of service delivery, etc; ensure configuration consistency; security enhancement; early warnings, etc.
- Strategy: intelligent monitoring and control
- Examples
  - Power saving for servers
  - Thermal management: Intelligent AHU control according to server loadings at rack level
  - Anomaly detection from computing, storage and network activities
- A long-term big data analysis work

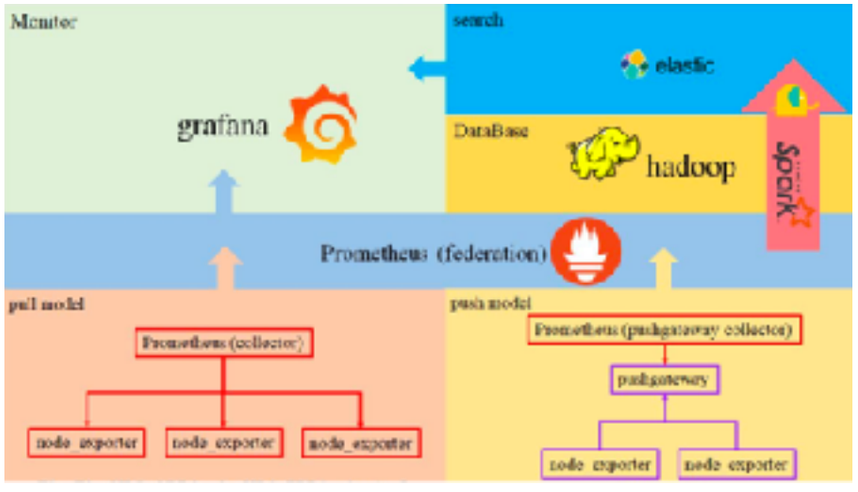

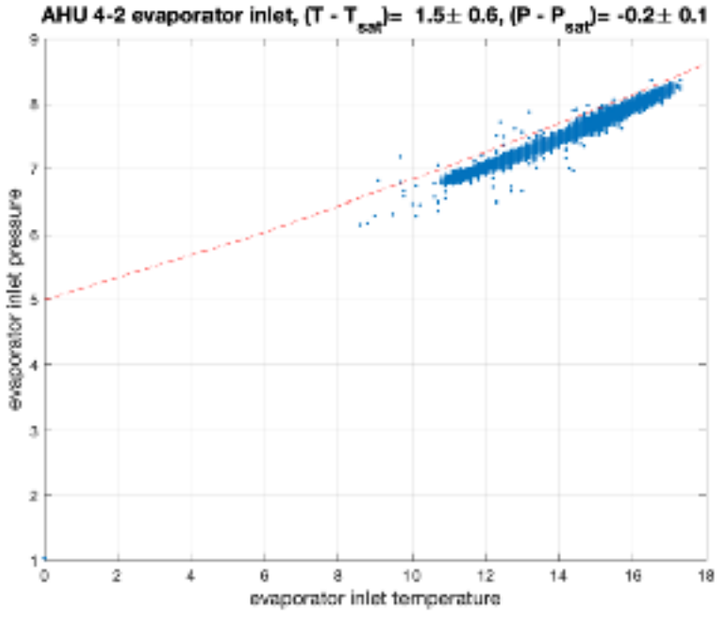

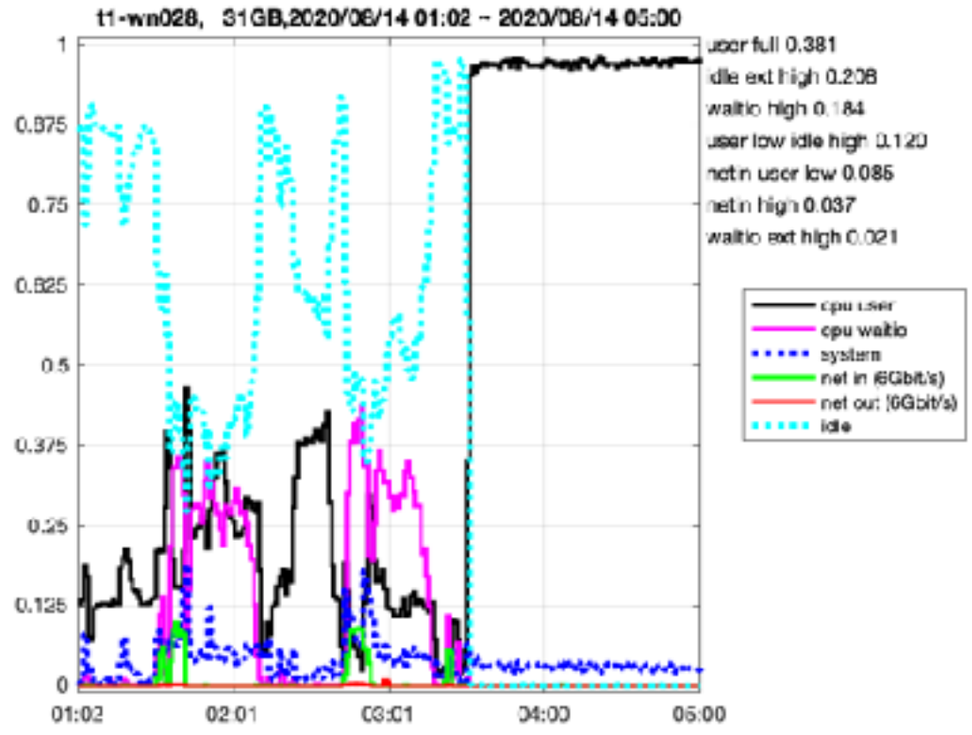

Security Operation Framework

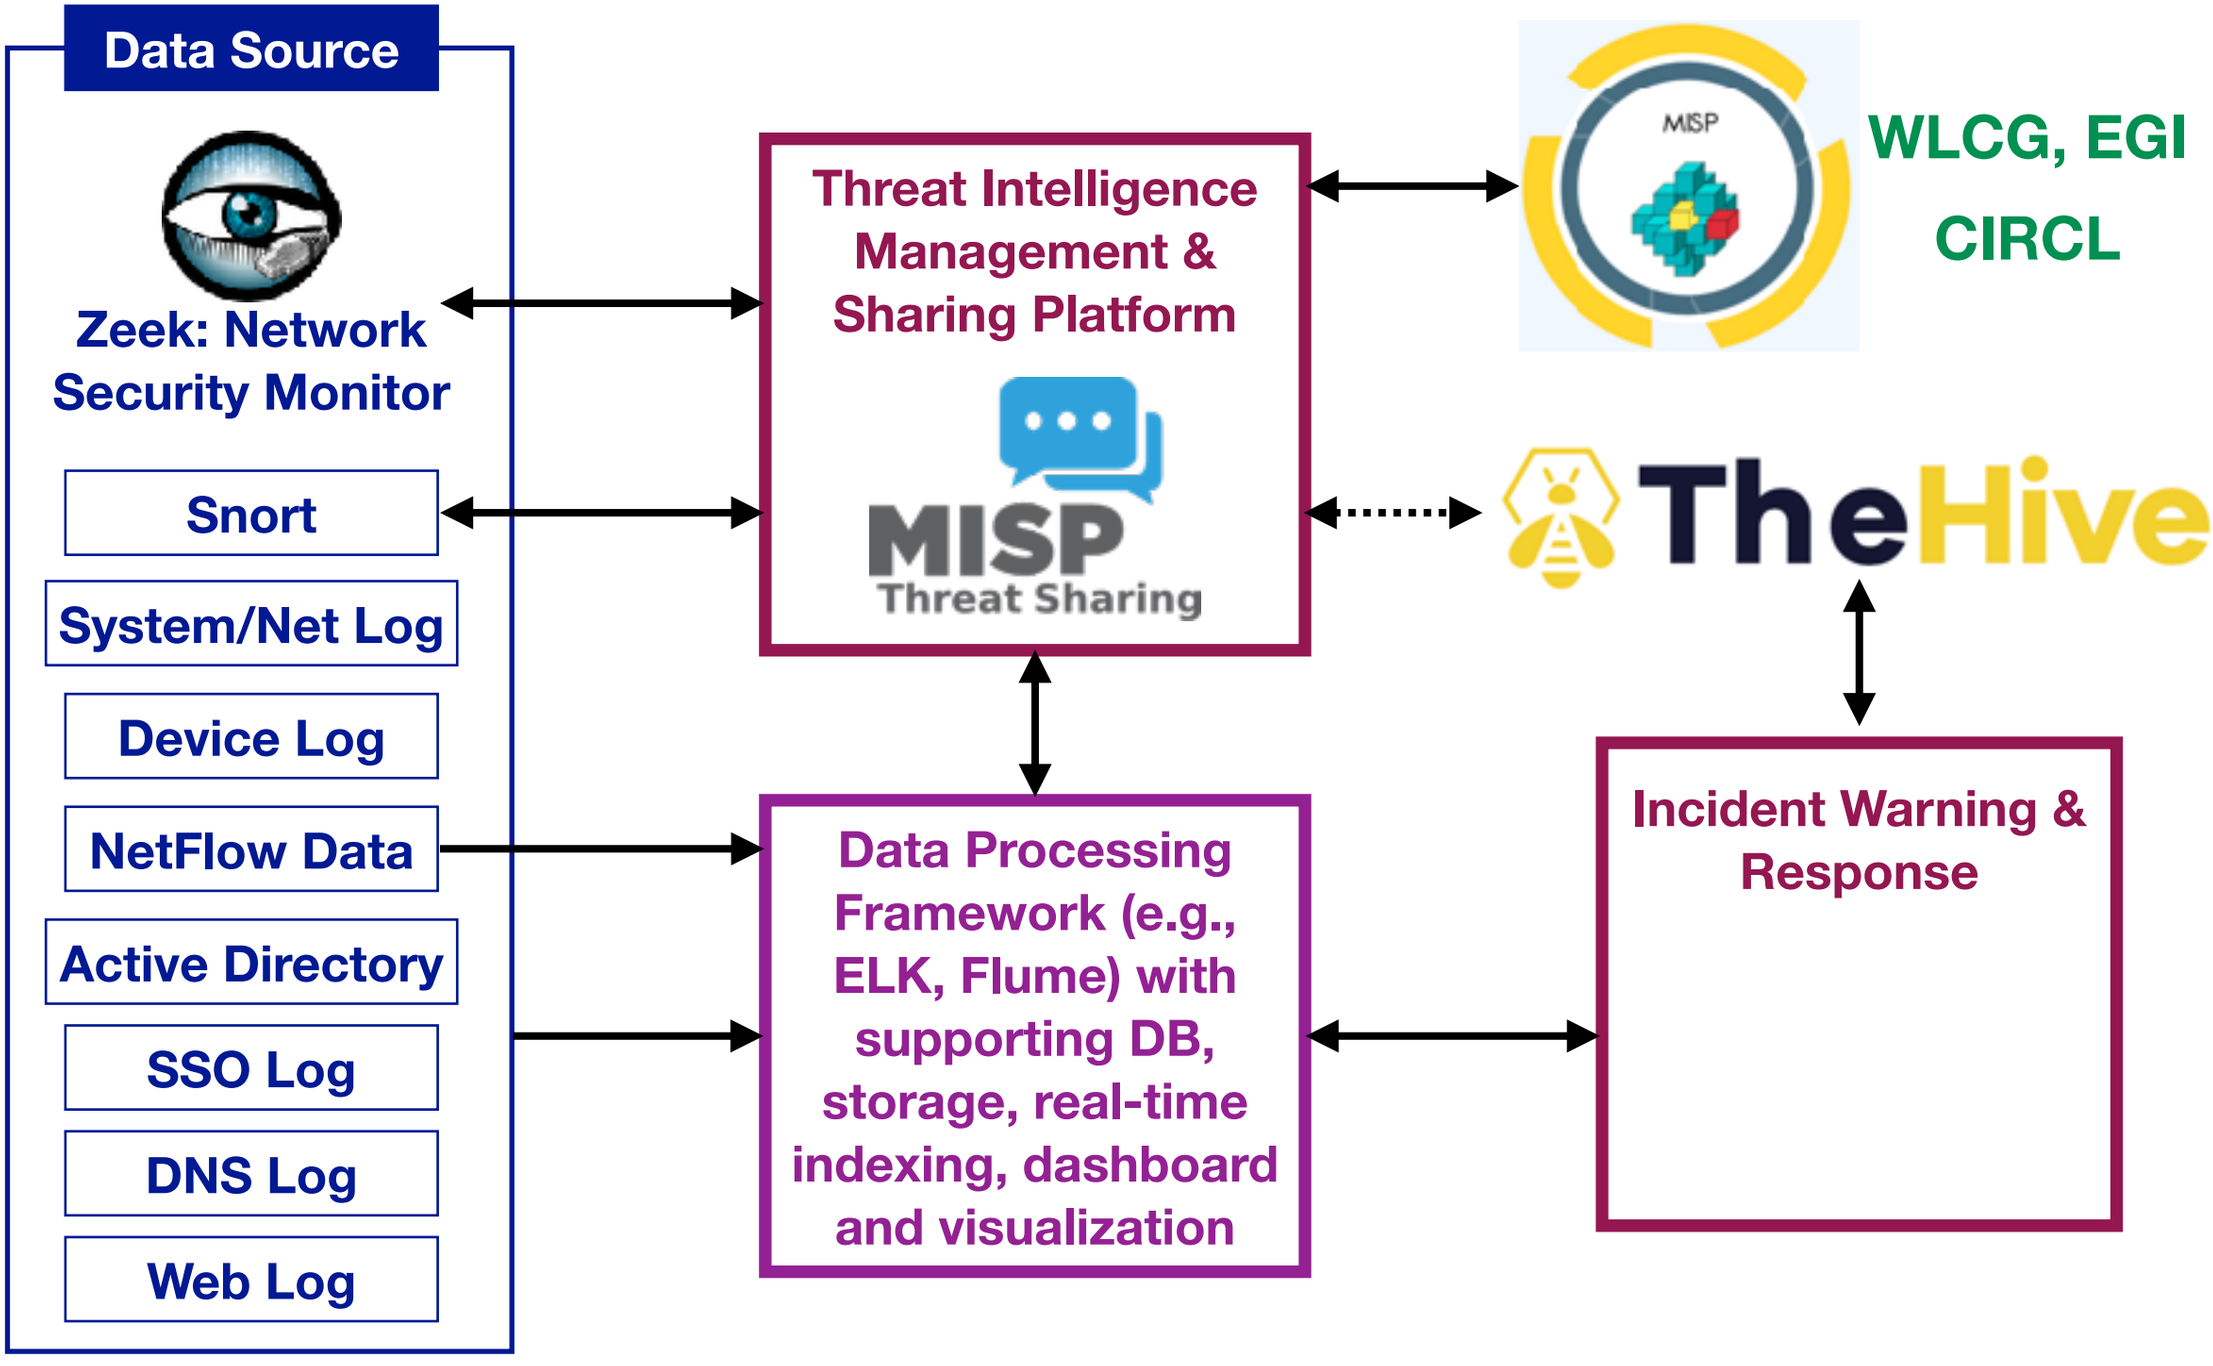

- Interoperation with WLCG SOC through MISP sharing and collaboration framework
- Improving intelligence and efficiency of the Data collection, ingestion, analytics and storage

Summary

- ASGC is extending WLCG technologies in supporting multiple O(PB) scale research applications of various disciplines by the common distributed cloud infrastructure
  - Customized workflow and optimized efficiency of whole data analysis pipeline over distribute cloud are typical results
  - Keep on evolving by WLCG collaboration and the interactions between application requirements and advanced cloud technologies
  - Supporting regional collaborations with non-HEP communities working together with EGI/EOSC
  - Experiences and outcomes are open and sharable for sure
- GPU, containerization, Ceph, system efficiency are the focus for improving application performance in a local site
- System efficiency based on intelligent monitoring and control is also our strategic focus.
- Open source, data sharing, reproducibility and open application are included in our roadmap