ASGC Site Report

Chien-De Li,
Felix Lee,
Eric Yen
Academia Sinica Grid Computing Centre (ASGC)
Taiwan

HEPiX Autumn 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 Academia Sinica
- 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.

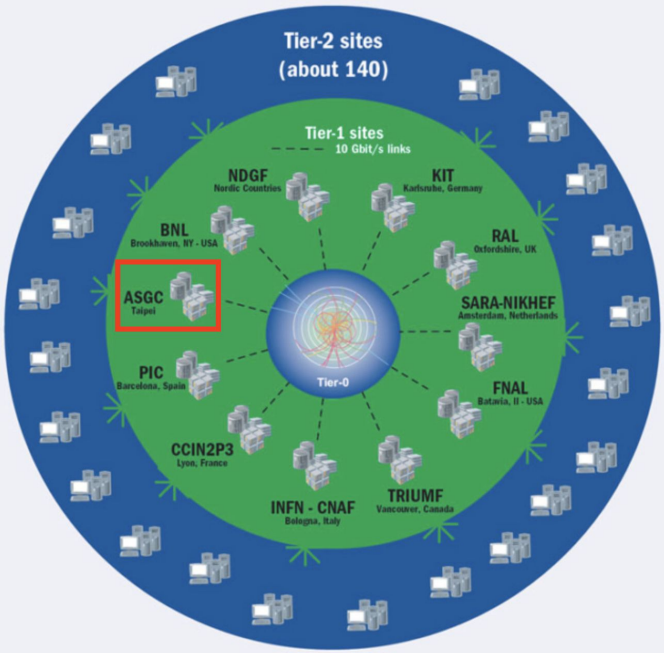

ASGC Resources

PUE : 1.6

- **Total Capacity**
  - 2MW, 400 tons AHUs
  - 112 racks in ~ 800 m^2
- **Resources (2021)**
  - 17,250 CPU Cores
  - 950,272 GPU CUDA Cores/168 GPU Cards
  - 24 PB Disk Storage
  - 2x10Gb links to CERN and primary NRENs worldwide
- **WLCG Tier-1 Center since 2005**

  84849 CPU-Days in 2016

  1132296 CPU-Days in 2021(to Oct. 20)
- **Supporting high-performance and high-throughput computing in Academia Sinica by distributed cloud operating system (DiCOS)**
- **R&D on system efficiency optimization by intelligent monitoring & control**

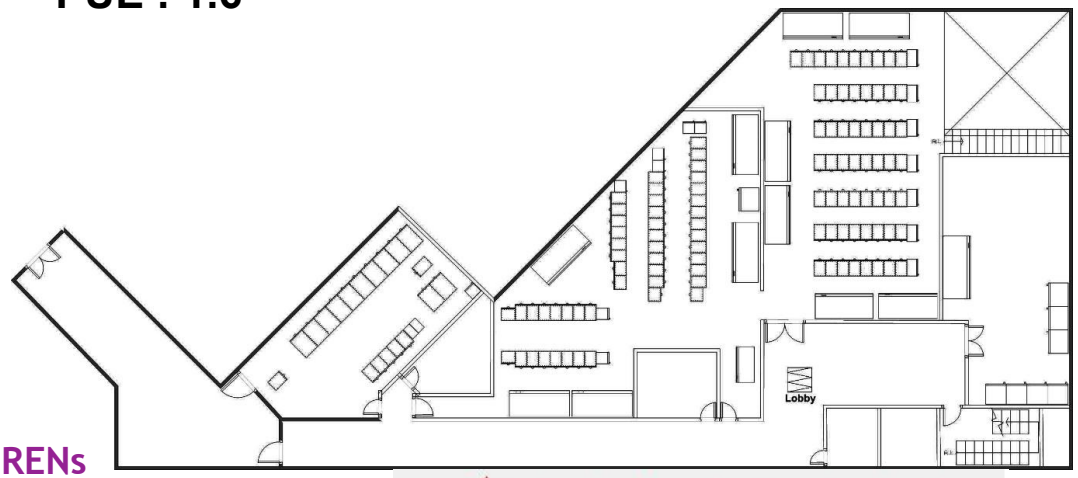

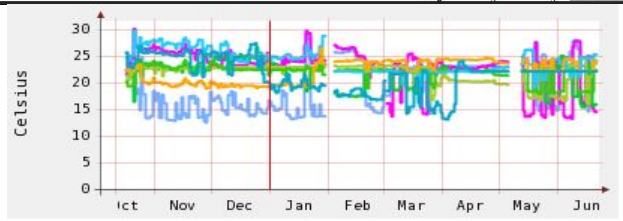

Monitoring the power consumption and temperature of every piece of equipment every 10 seconds.

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—>bb; Di-Higgs; Dark matter
- Computing Resource
  - Tier-1: 3,268 CPUCores
  - Federated Taiwan Tier2: 672 CPUCores
- 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, RoseTTAFold and ML/DL applications.

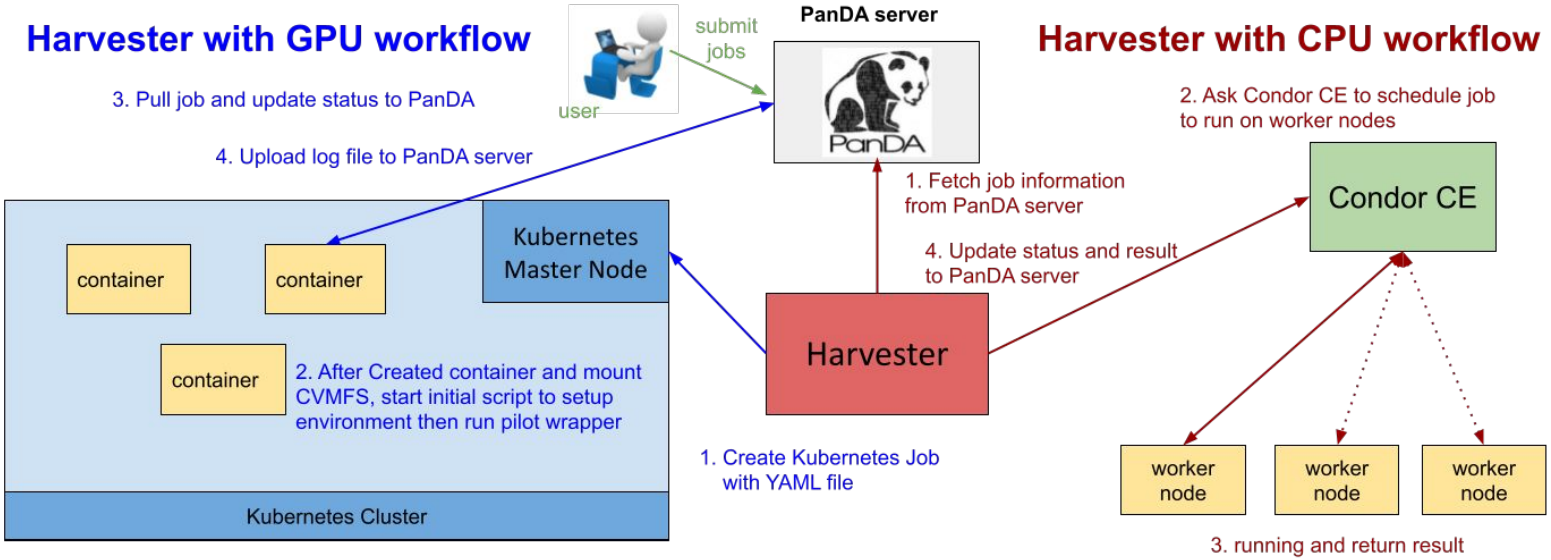

Building Distributed Cloud Infrastructure Supporting broader Scientific Applications based on WLCG

- **GPU update**
  - 2 servers
    - HPE XL675d Gen10 Plus
    - CPU： AMD EPYC 7302 (Rome) 16-Core Processor *2
    - 1TB memory and 1.2TB Nvme SSD
    - GPU：A100-SXM-80GB *8
    - Network：100Gb/s (EDR)*4,
      will enable NVLink in future

  - 6 servers
  - Supermicro AS-4124GS-TNR
    - CPU：AMD EPYC 7302 (Rome) 16-Core Processor *2
    - 1TB memory and 1TB Nvme SSD
    - GPU：GeForce RTX 3090 *8
    - Network：100Gb/s (EDR)

Cloud services

- Openstack
  - For Generic services and on-demand worker nodes
- Two kubernetes cluster
  - For users
    - Offering batch, interactive GUI jobs
      - Such as remote Jupyterlab and virtual desktop
    - GPU cloud
  - For core services
    - Harvester
    - Rucio
    - Elasticsearch + kibana

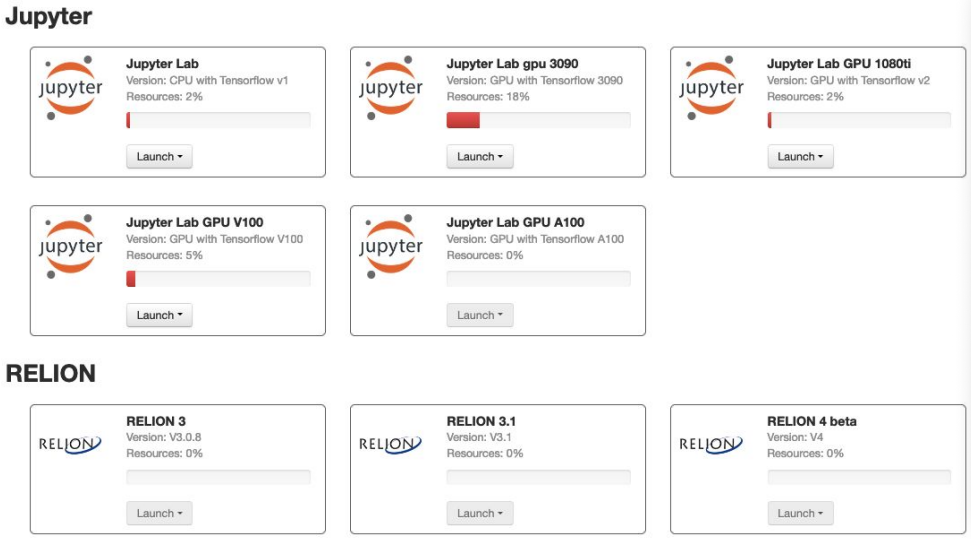

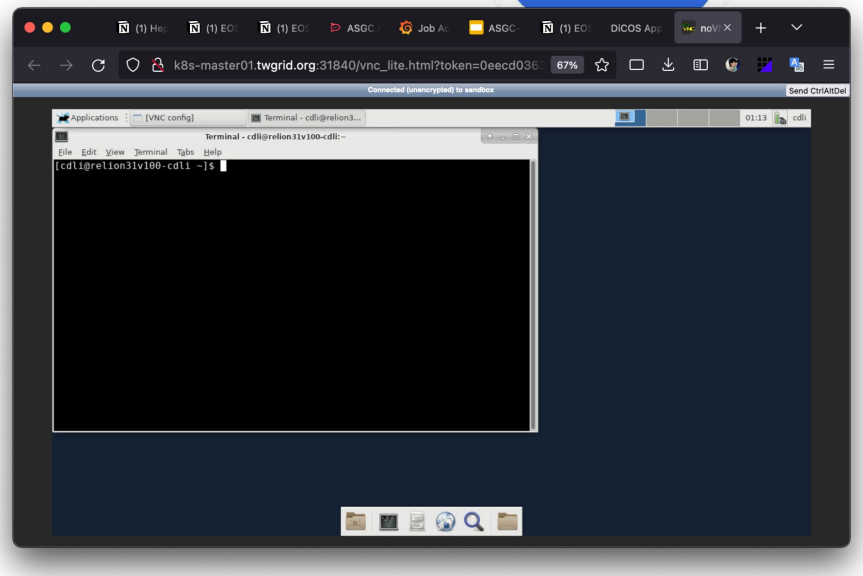

Storage services for local user and core services

- Two ceph instance
  - For OpenStack
    - Mainly provides RBD
    - 1.4PB
  - For user
    - Luminous => Nautilus
    - Mainly provides Cephfs for our local users
    - 5.5 PB
    - mds replay issue
    - expand mds from 3 to 14(7 online, 7 hot-standby)

Storage services for WLCG

- DPM
  - Mainly provides for ATLAS
  - 8.7 PB capacity
  - DOME 1.14.2
  - plan to migrate in early 2022
- EOS
  - 3 x (MGM+QuarkDB) setup
    - 256GB RAM
    - 1TB system disk
    - 2 TB SSD (for /var/eos QuarkDB)
  - Will deploy testbed at the end of 2021

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

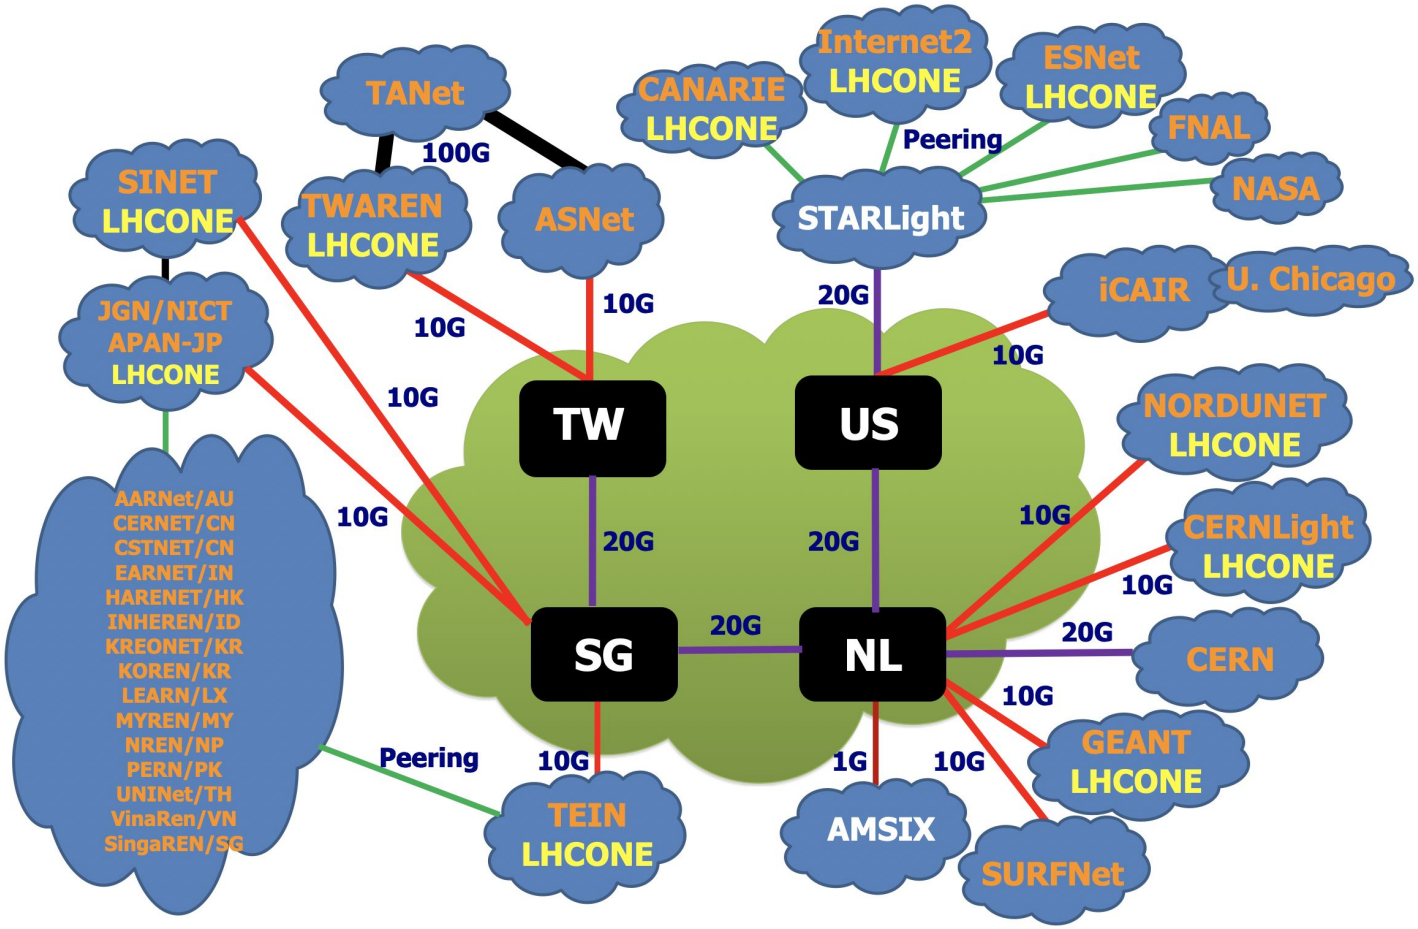

Token-based AAI

- 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

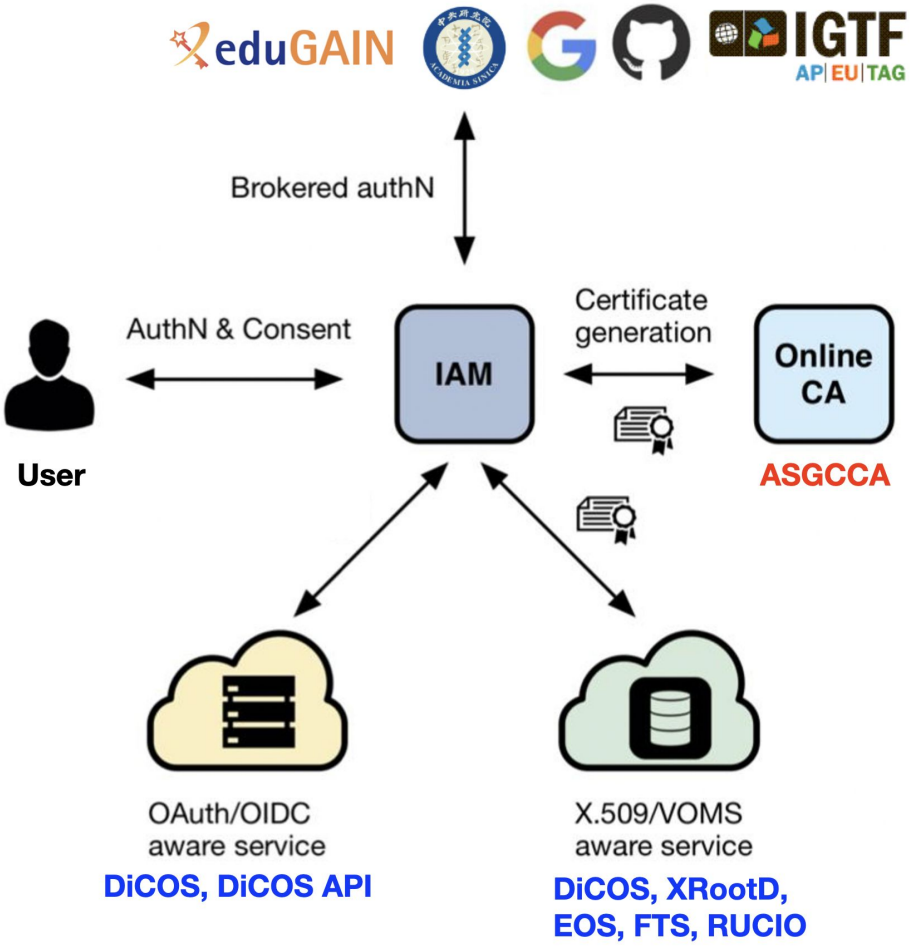

Big Data infrastructure

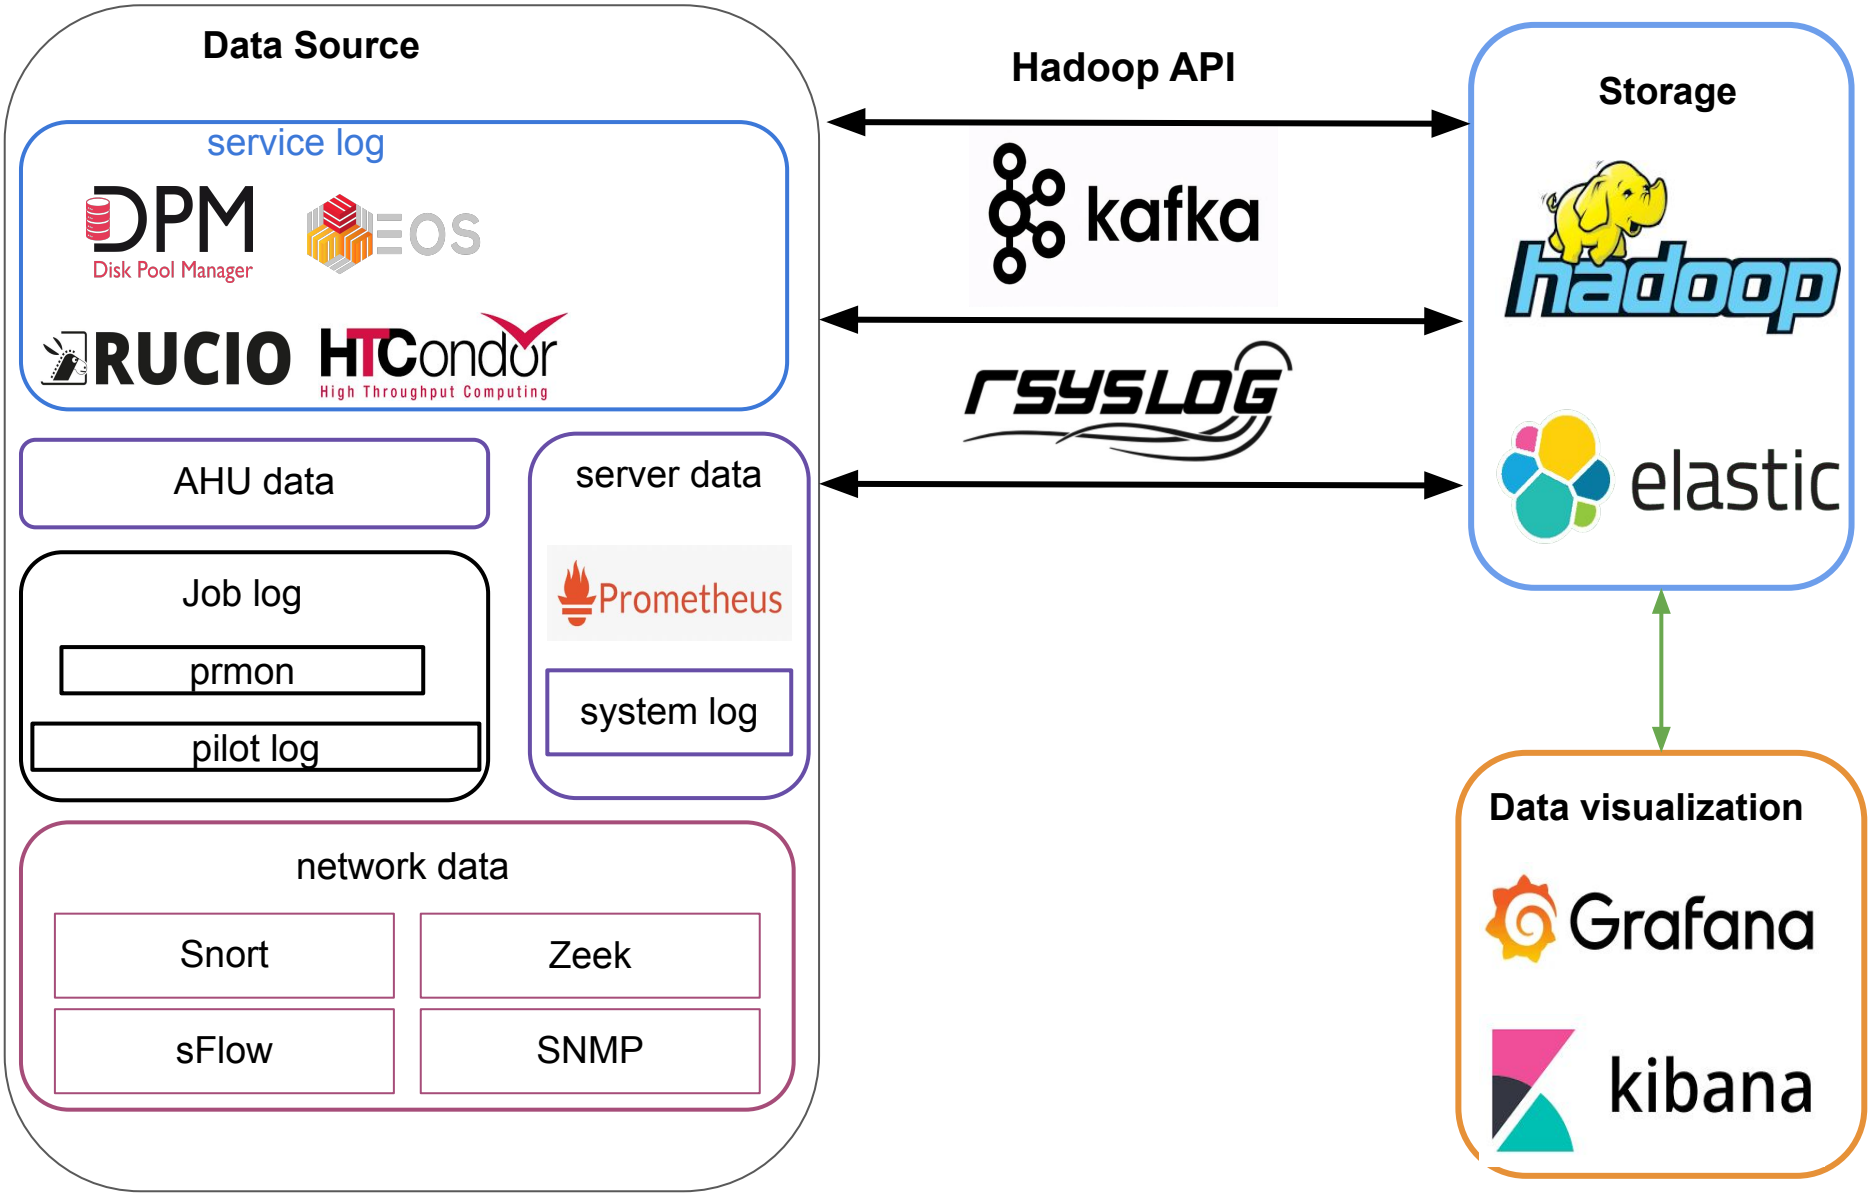

System Efficiency Optimization

- Scope: Power, Thermal and Distributed Cloud System management
- Strategy: intelligent monitory 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
  - Compute/Storage/Network/Security anomaly detection based on system log and metrics in future

AHU Performance Monitoring

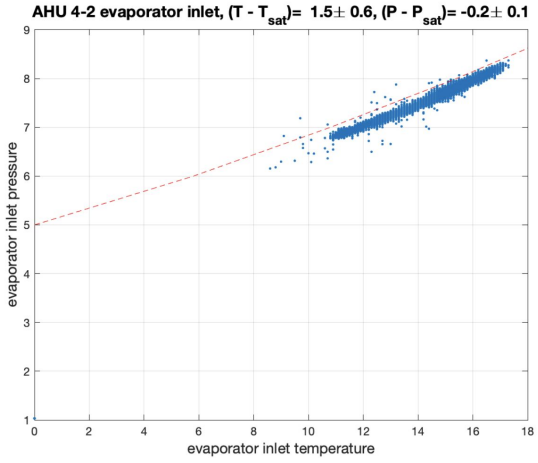

Worknode Monitoring: High ratio of WaitIO & System

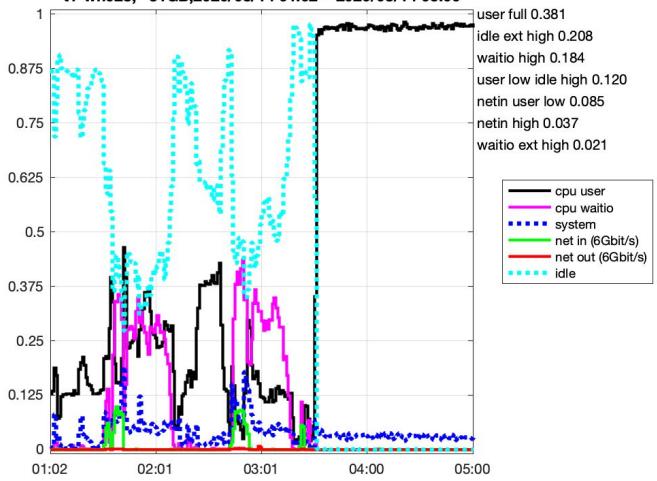

Worknode Monitoring: Misconfigured hyper-threading

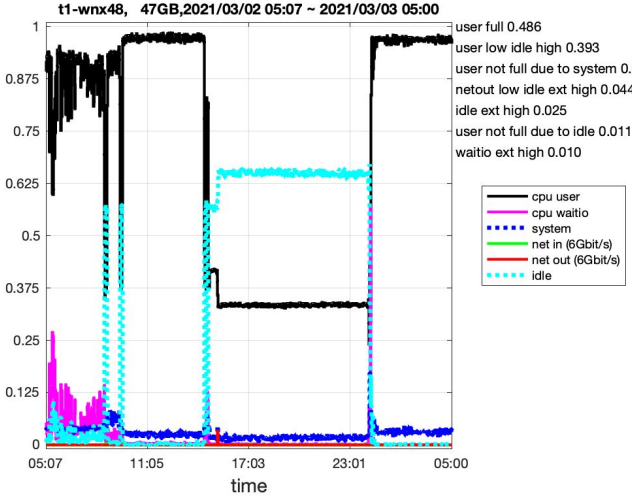

- **Thanks a lot!**