

ATCF6 Summary Report

WLCG GDB

11 Jan 2023



Ahn Sang-Un on behalf of ATCF

ATCF Goals

- (Short-term) To discuss on the possible solutions for the improvement of connectivity among Asian Tier sites and their domestic network environment, and to monitor periodically the status of the established network environment
 - The first principle is to keep the traffics among Asian sites within the region
- (Long-term) To organise a body with a broader agenda embracing not only the network but also common issues that can be arisen among Asian Tier Sites
 - It provides a place where the Asian sites can share and discuss any issues to resolve in a collective way
- (Target) Asian (LHC) Tier sites however, open to all interested parties

Organization

- Co-chairs of Asia Tier Center Forum
 - Suranaree University of Technology (represented by Prof. Chinorat Kobdaj)
 - Korea Institute of Science and Technology Information (represented by Dr. Yoon Heejun)
- Local Organization Chair for ATCF6 - Prof. Chinorat Kobdaj
- Steering Committee
 - SUT (represented by Prof. Chinorat Kobdaj)
 - Hiroshima U. (represented by Prof. Toru Sugitate)
 - TIFR (represented by Prof. Kajari Mazumdar)
 - Konkuk U. (represented by Prof. Oh Sun Kun)
 - KISTI (represented by Dr. Yoon Heejun)

ATCF6 Meeting

The first F2F meeting after the global pandemic

- Centara Ao Nang @ Krabi, Thailand
 - co-hosted by SUT and KISTI
 - 21 ~ 24 November 2022
- 24 registered participants (+4 remote)
 - 15 institutes from 8 nations
 - ASGC(TW), BRIN(ID, former LIPI), CCNU(CN), Chulalongkorn U.(TH), Fudan U.(CN), Hiroshima U.(JP), ICEPP(JP), KEK(JP), KISTI(KR), Rajamangala U.(TH), SUT(TH), TIFR(IN), VECC(IN) and CERN



ATCF6 Agenda

- Round table of Asian sites - 8 contributions
 - Site-specific issues regarding resources, operations, R&D, networking, etc.
- Experiments perspectives and updates
 - WLCG, ALICE and Belle II
- LHCONE updates: Asian NRENs and CERN - 3 contributions from KREONet, TEIN*CC and CERN
 - The latest view of LHCONE and related R&D projects
 - The current status of collaborative works within WLCG for LHC data challenges
- Discussion: Distributed Storage & Asian Computing Collaboration
- Special sessions: EOS Storage (Andreas-Joachim Peters) & HTCondor mini-workshop (TIFR & KISTI)
- Link: [ATCF6 indico](#)

Experiments Perspectives

WLCG Data Challenges

- The next one scheduled in early 2024, Asia's collective actions required for prompt response to these challenges
- Targets for T2s will need to be discussed for different regions
- It is noted that ATCF is the right forum to organise the participation of Asian centers to the challenges

WLCG data challenges

The 2023 data challenges will likely be in ~Q1 2024

The target rates (30% of the HL-LHC traffic) seem too aggressive given the shift of the HL-LHC schedule

- A 20% to 25% target seems more realistic and will not require a considerable hardware investment for the sites

30% to be reviewed

Specific targets for T2s will need to be discussed for different regions. **We should discuss what makes sense for ATCF**

We would like to have a joint challenge with some of the partners (e.g. DUNE, Belle-2). Under discussion

	LHC Network Needs (Gbps) Minimal Scenario in 2027	LHC Network Needs (Gbps) Flexible Scenario in 2027	Data Challenge target 2027 (Gbps)	Data Challenge target 2025 (Gbps)	Data Challenge target 2023 (Gbps)	Data Challenge target 2021 (Gbps)
T1						
CA-TRIUMF	200	400	100	60	30	10
DE-KIT	600	1200	300	180	90	30
ES-PIC	200	400	100	60	30	10
FR-CCIN2P3	570	1140	290	170	90	30
IT-INFN-CNAF	690	1380	350	210	100	30
KR-KISTI-GSDC	50	100	30	20	10	0
NDGF	140	280	70	40	20	10
NL-T1	180	360	90	50	30	10
NRC-KI-T1	120	240	60	40	20	10
UK-T1-RAL	610	1220	310	180	90	30
RU-JINR-T1	200	400	100	60	30	10
US-T1-BNL	450	900	230	140	70	20
US-FNAL-CMS	800	1600	400	240	120	40
(atlantic link)	1250	2500	630	380	190	60
Sum	4810	9620	2430	1450	730	240



Experiments Perspectives

ALICE


- Asian contribution to the experiment ~ 8%
- Growth of Asian resources is similar to other regions
- Growth potential is there, as well as the expertise
- Steady and consistent operation of Asian T2s
 - Especially a small T2, SUT (Thai) is an important example for other Asian countries entering the Grid

A Large Ion Collider Experiment



Regional contribution

- ~8% Asian contribution
- Last T1 remaining in Asia@KISTI
- The growth of Asian resources is similar to other regions
- The growth potential is there, as well as the expertise
- Diminishing role of smaller T2 centres - this is an unfortunate global trend




A Large Ion Collider Experiment



Role of Asian sites - resources and innovation

- KISTI remains the largest contributor
 - In production - first of its kind (worldwide) disk-based custodial storage element 12PB
- Hiroshima - site restarted with new storage and upgraded capacity
 - Another T2 in Nagasaki is being set up
- Kolkata - steady resources increase and normal operation
- SUT - a small T2, but consistent operation
 - Important as example for other Asian countries entering the Grid

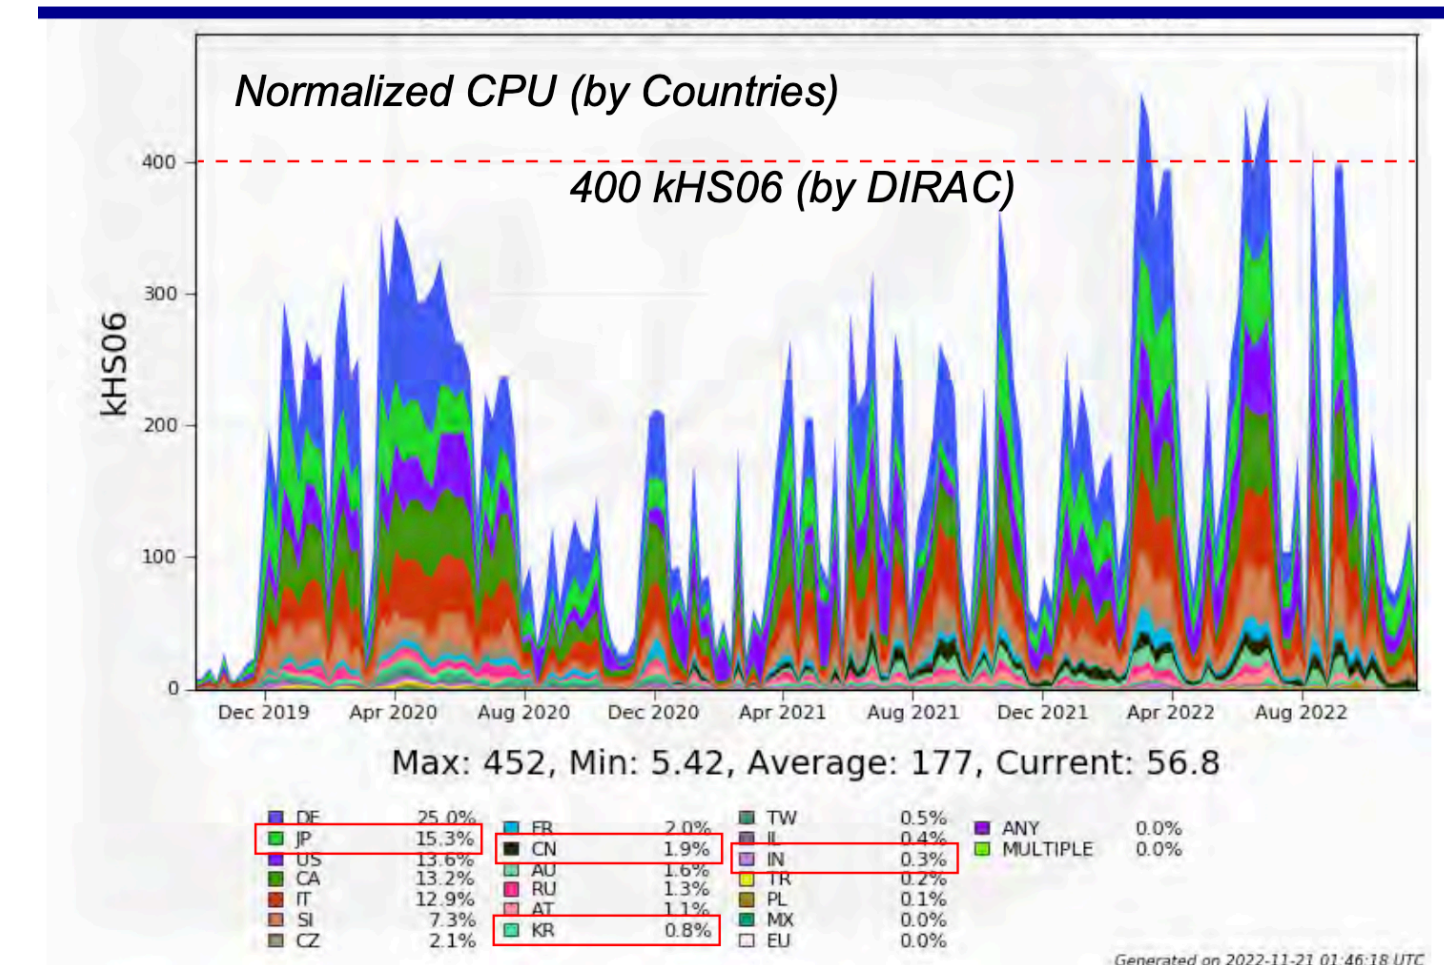


Experiments Perspectives

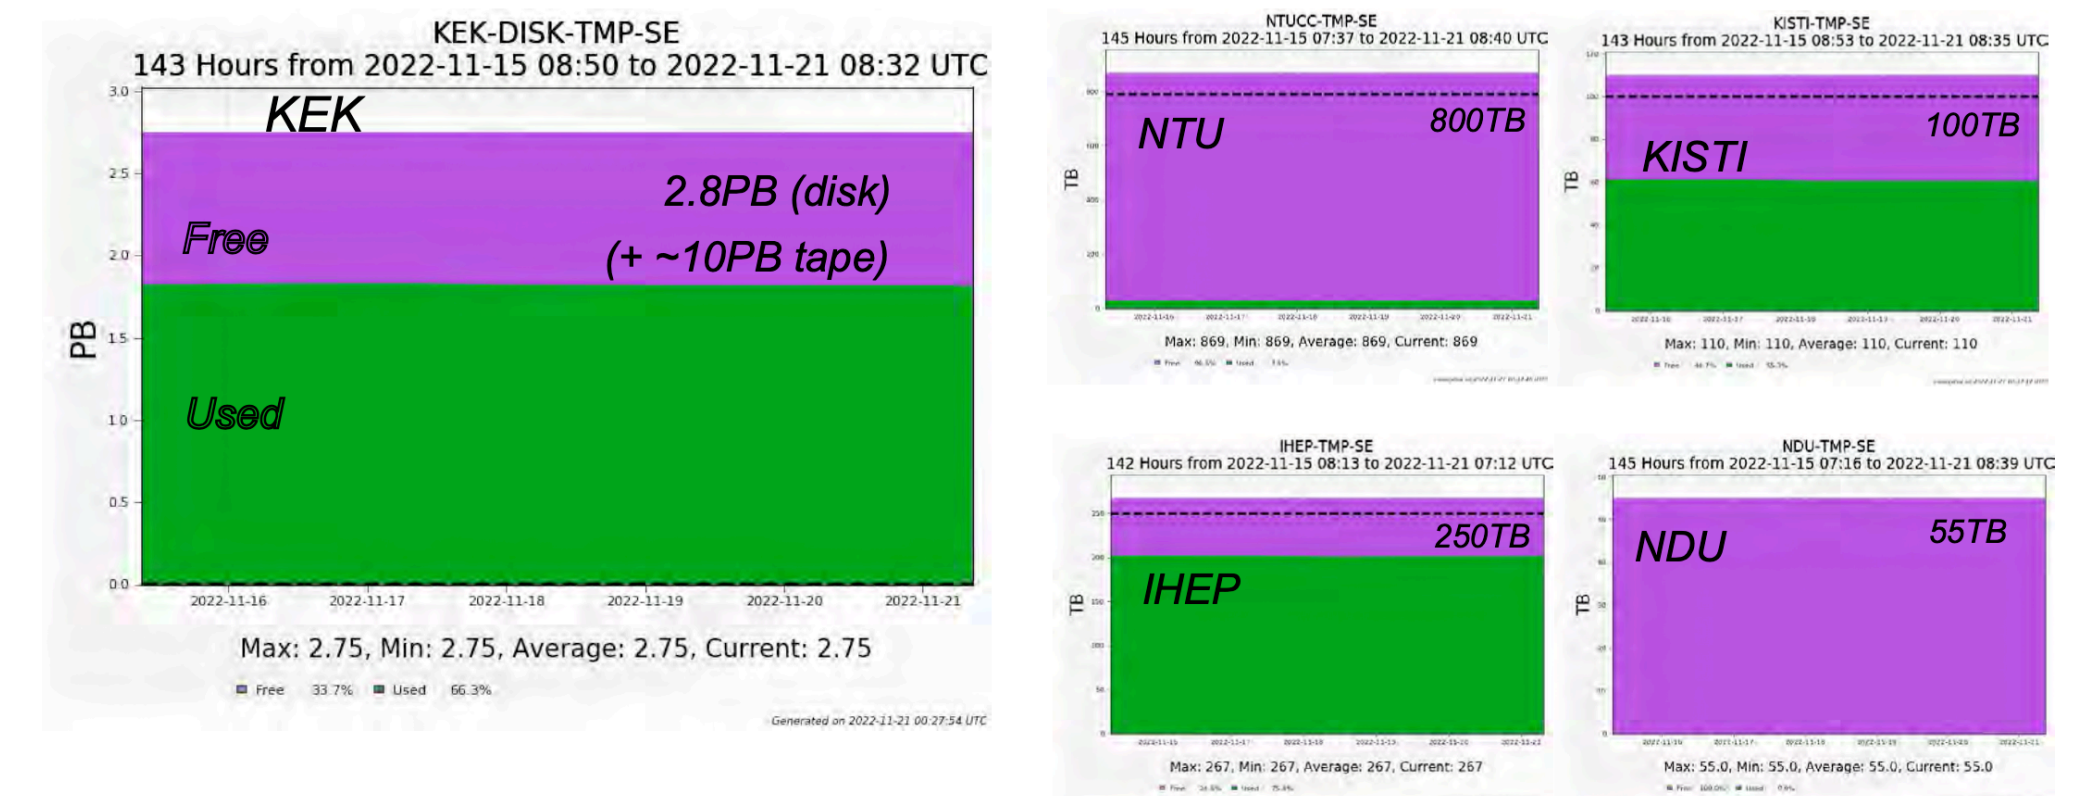
Belle II

- Very successful data taking throughout the pandemic
- Started publishing new physics results
- Successful Rucio integration for data transfer
- Enabled access via HTTPS/WebDAV
- Asian contributions to the experiment ~ 3% (excluding KEK)
 - Disk storage contribution ~ 8% (excluding KEK)

Compute Element



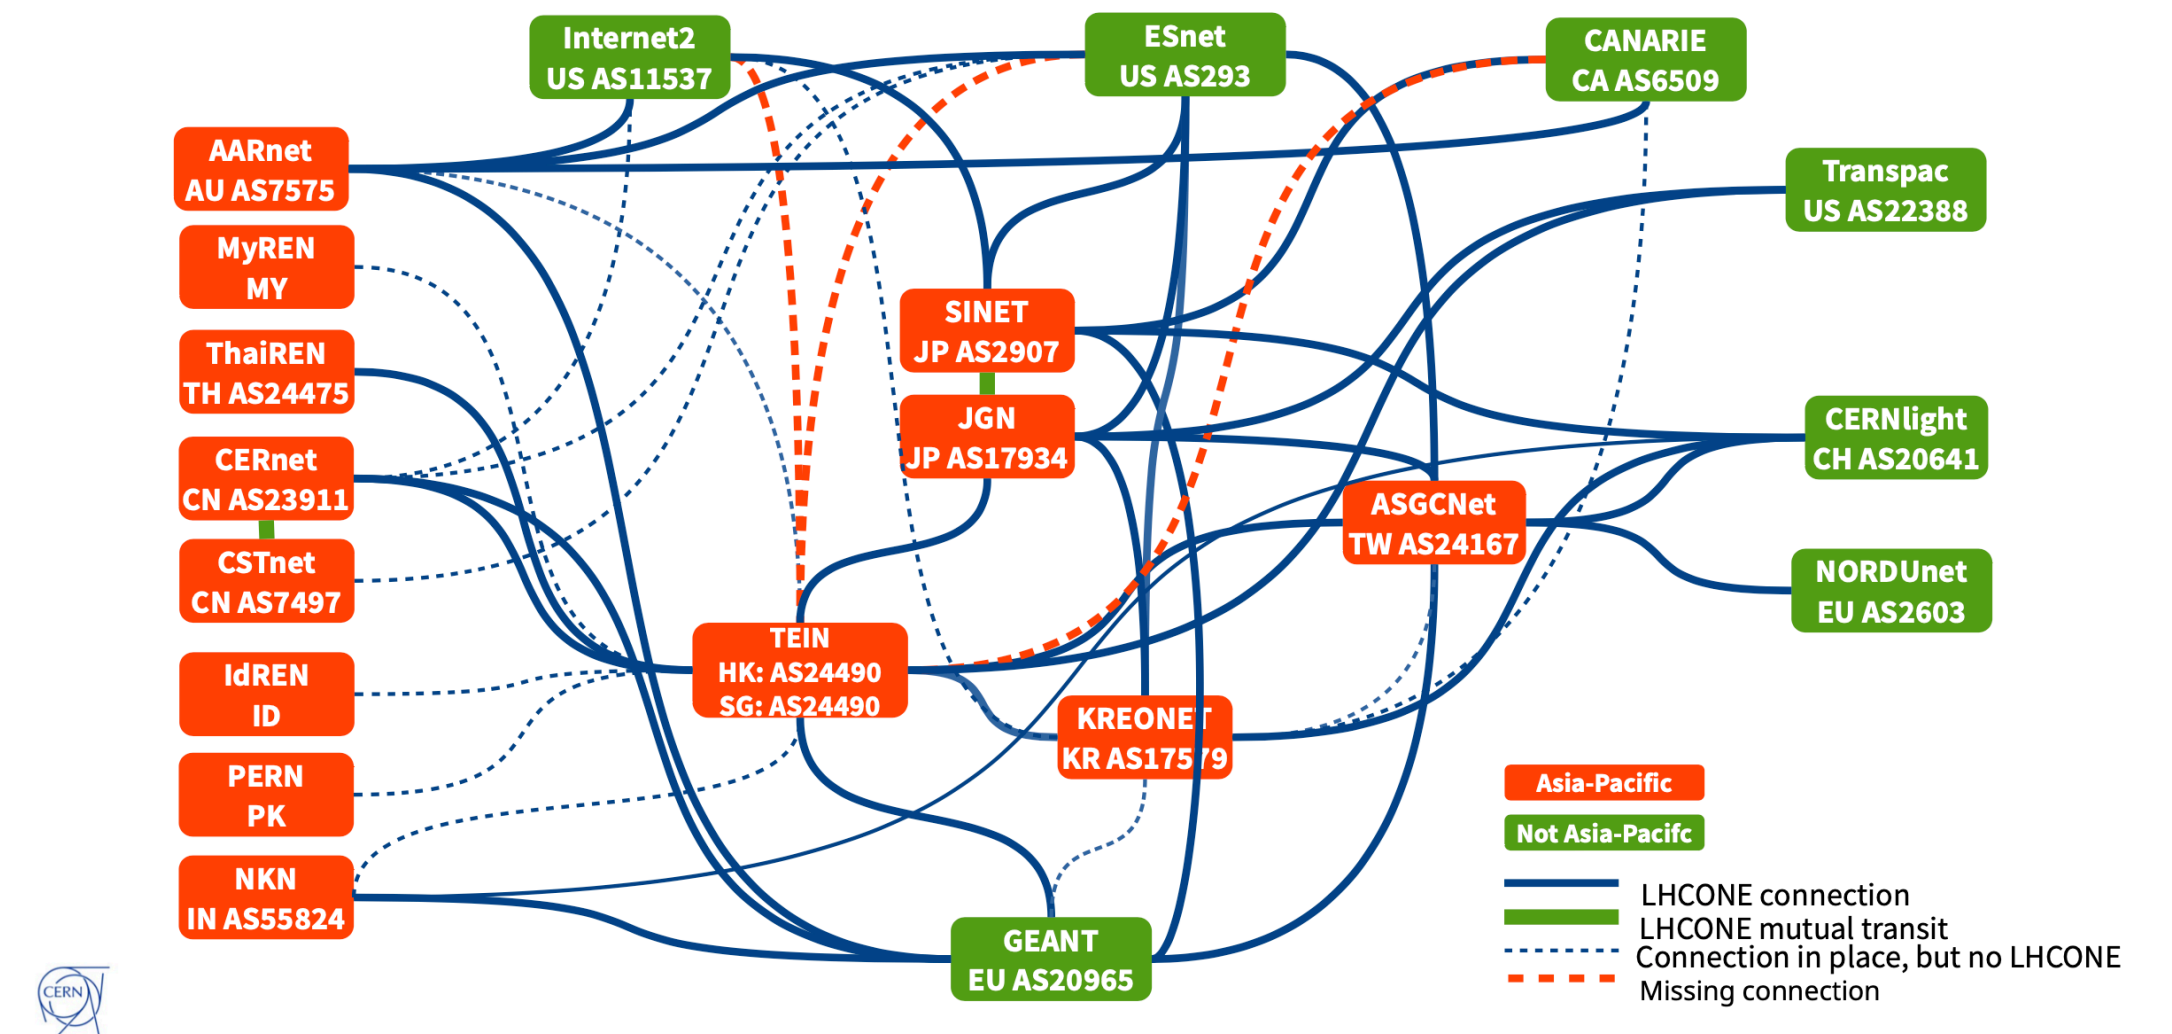
Storage Element



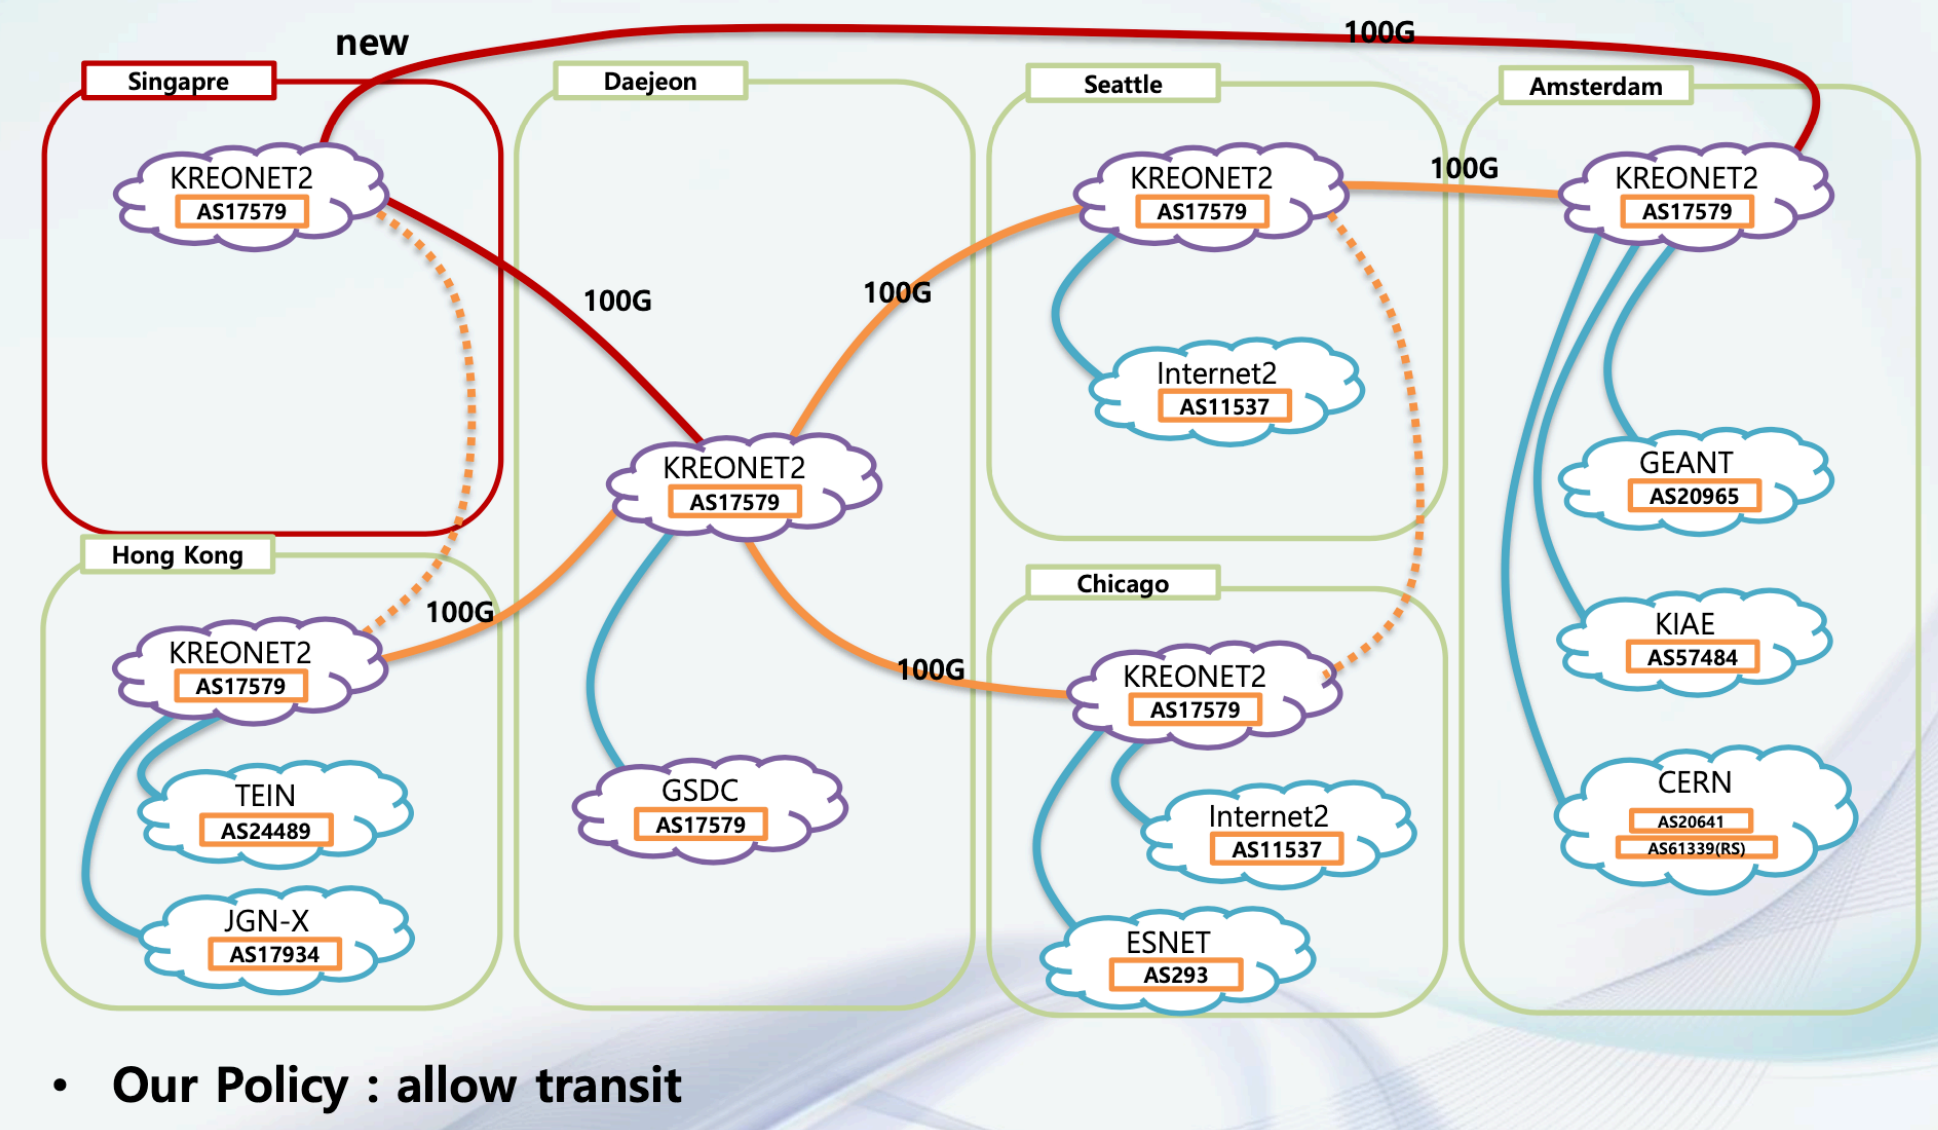
LHCONE Update

- LHCONE connectivity has been evolved since ATCF1 in 2015 and the complexity is already comparable to the other continents
- Still there are missing connections among Asian NRENs in terms of LHCONE to achieve full reachability
- BUT KREONet announced that they will allow transit to their LHCONE VRF peers which will be a complete solution to the point above-mentioned: consequently TEIN can reach Internet2, ESnet and CANARIE via KREONet (vice versa) which is great for most of Asian sites that connect to TEIN
- Several interesting network R&D projects presented: NOTED, Packet tagging, AutoGOLE and SENSE, CNAF-CERN DCI (using dark fibers)

Asia-Pacific VRFs – Current Status



LHCONE on KREOENT2(2023)

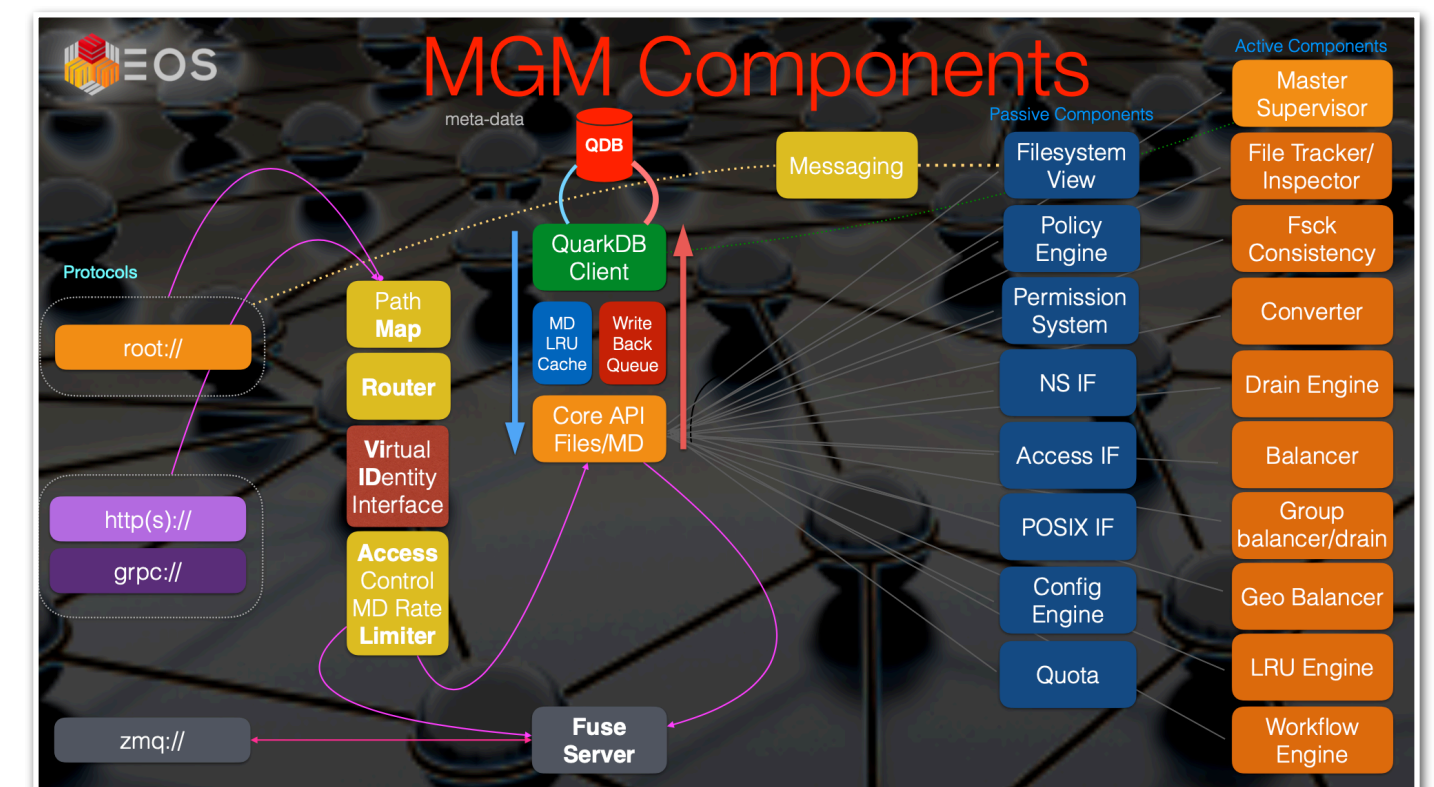

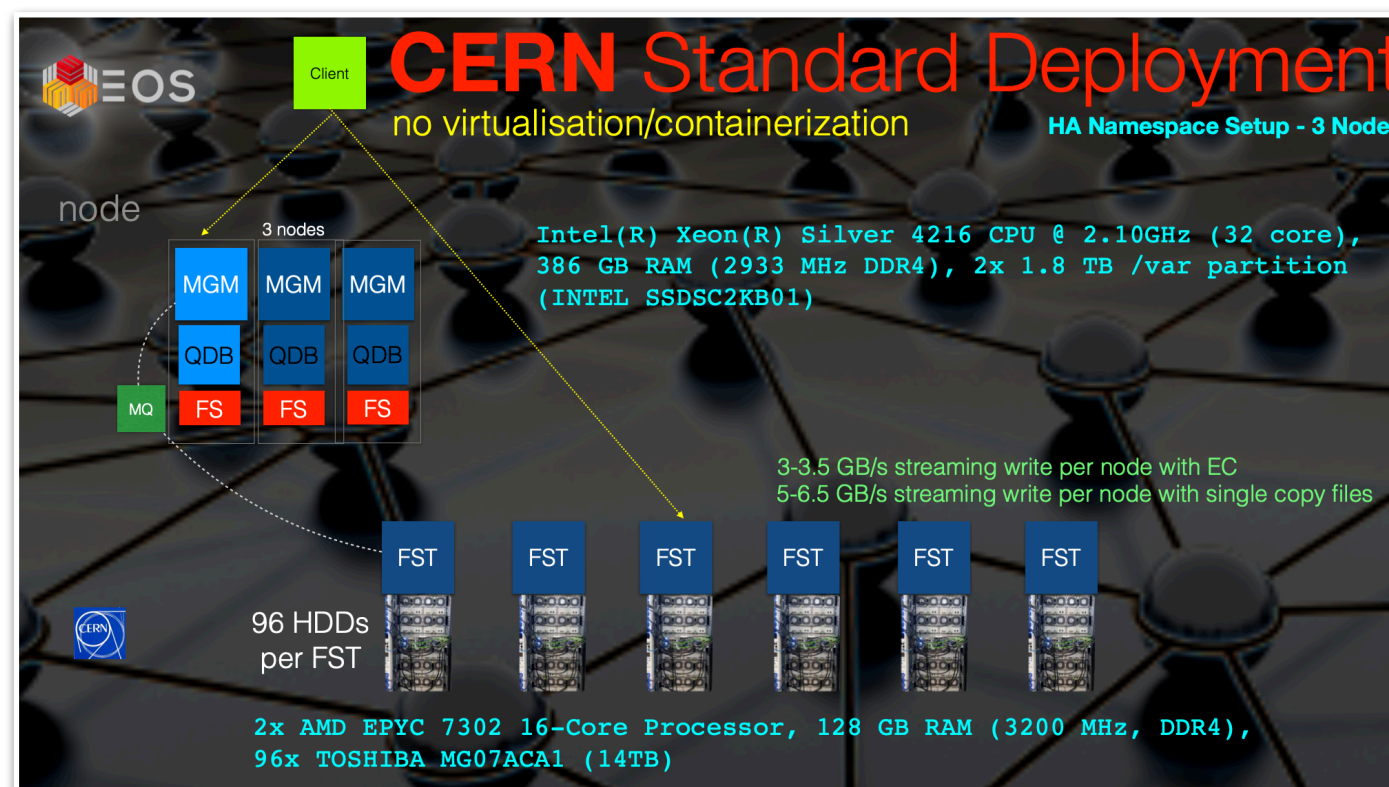


Discussions on site operation

- Similar (or same) software (or middleware) stack should be used in region hoping for operations cost reduction and ease of technical supports from one another
 - ICEPP (JP) will try to migrate (during next whole system migration (2025), note that ICEPP leases its HW in 3-year-term) its 8PB of DPM storage to something else (currently looking at dCache) while ASGC is migrating their disk storage from DPM to EOS for ATLAS now (at the similar scale).
 - Many of ATCF6 participating sites are already using EOS for their storage: Kolkata (Indian ALICE T2), Hiroshima (Japanese ALICE T2), KISTI (Korean ALICE T1)
 - It was highly recommended for ICEPP to look into EOS rather than dCache in the same principle.
 - *cf.* KISTI CMS T2 is under consideration to migrate its dCache storage to EOS
- BRIN (ID) in consideration of choosing an integrated system for distributed and different types of computing clusters other than its role as ALICE T2.
 - It was recommended to have a list of different requirements from their customers first and then select a proper scheduler to integrate them.
 - It was remarked that HTCCondor seems be able to cover what they want but the choice of proper scheduler should be careful depending on the requirements.

EOS Session

- A half day session of EOS storage system given by Andreas-Joachim Peters
 - More than a hundred of slides covering EOS architecture, technological details, deployment examples and their hardware requirements, CERN references, and EOS ecosystems
 - Detailed discussion on Erasure Coding implementation (RAIN layout), deployment and operations were useful for sites not only whom already use but also whom plan or consider the idea
 - Tutorial for quick and simple deployment of EOS based on brand-new documentation
- Sites experiences using EOS RAIN layouts were also presented from VECC (the Indian ALICE T2) and KISTI (the Korean ALICE T1)



HTCondor Mini-Workshop

- TIFR (the Indian CMS T2) and KISTI presented their experiences on configurations and operations of HTCondor-CE and HTCondor for their customers
 - TIFR shared troubleshooting experience during the upgrade of HTCondor 8.8 to 9.0 especially on enabling token based authentication for CMS
 - KISTI shared HTCondor-CE setup for ALICE, pre-emption enabled group accounting configuration on negotiator for domestic user groups, and running jobs on apptainer

Troubleshootings while setup

- Started condor upgradation on T3_IN_TIFRCondor
 - Suggestion was to install LTS release 9.0.5 - mir
 - Later on, went for new stable release 9.0.15
- Disabled GSI submission and enabled token
- User mapping issue:


```
03/28/22 12:32:17 [97263] GAHP[116261] -> '39' 'F' '12502'
which is not a valid user account'
03/28/22 16:36:46 DC_AUTHENTICATE: authentication of
required for this command (1112 QMGMT_WRITE_CMD),
condor_ping with SCITOKENS works for pilots, but it was r
```
- Earlier mapping:


```
SCITOKENS /https://voms-auth.web.cern.ch/
```
- New entries:


```
SCITOKENS /https://voms-auth.web.cern.ch/
SCITOKENS /https://voms-auth.web.cern.ch/
SCITOKENS /https://voms-auth.web.cern.ch/
SCITOKENS /https://voms-auth.web.cern.ch/
SCITOKENS /https://voms-auth.web.cern.ch/
```

Troubleshootings while setup cont...

- Jobs were aborted, all went to HeldJobs state after staying 30 min. at IdleJobs state
- Pilots were landing to CE which indicated that, there was no

Earlier:

```
GSI ".*,VcmsVRole=pilot.*" cmsjob@indiacms.res.in
GSI ".*,VcmsVRole=legadmin.*" samjob@indiacms.res.in
GSI ".*,VopsVRole=NULL.*" opsjob@indiacms.res.in
GSI (.*,VopsVRole=NULL.*) GSS_ASSIST_GRIDMAP
GSI (/CN=[-A-Za-z0-9/= ]+/\|@unmapped.opensciencegrid.
CLAIMTOBE (.*,anonymous@claintobe
FS (.*)\|
```

Now:

```
GSI /.*,VopsVRole=NULL.* opsjob@indiacms.res.in
GSI (.*,VopsVRole=NULL.*) GSS_ASSIST_GRIDMAP
GSI (/CN=[-A-Za-z0-9/= ]+/\|@unmapped.opensciencegrid.
CLAIMTOBE (.*,anonymous@claintobe
FS (.*)\|
```

Troubleshootings while setup cont...

- Thoroughly followed sequential updates and documents
- SAM tests try token submission first and fall back to x509 in case of failure.

```
07/27/22 10:13:06 (cid:12735) Command=QMGMT_WRITE_CMD,peer=<188.184.81.101:36854> 07/27/22 10:13:06 (cid:12735)
AuthMethod=SCITOKENS, AuthId=https://cms-auth.web.cern.ch/,08ca855e-d715-410e-a6ff-ad77306e1763,
CondorId=samjob@indiacms.res.in

07/27/22 10:14:30 (cid:12743) proxy full fqan: 07/27/22 10:14:30 (cid:12745) Command=DELEGATE_GSI_CRED_SCHEDD,
peer=<169.228.38.43:25665>

07/27/22 10:14:30 (cid:12745) AuthMethod=SCITOKENS, AuthId=https://cms-auth.web.cern.ch/,bad55f4e-602c-4e8d-a5c5-bd8ff762113,
CondorId=cmsjob@indiacms.res.in
```

Authentication with SciToken is successful

HTCondor Configuration (Apptainer)

```
### Singularity(Apptainer) Part
SINGULARITY_JOB = !isUndefined(TARGET.SingularityImage)
SINGULARITY_IMAGE_EXPR = TARGET.SingularityImage
SINGULARITY_TARGET_DIR = /srv
MOUNT_UNDER_SCRATCH = /tmp, /var/tmp
SINGULARITY_BIND_EXPR = !ThenElse( isUndefined(TARGET.SingularityBind), "/cms_scratch, /cvmsfs, /cms, /share, /tmp", TARGET.SingularityBind)
SINGULARITY_EXTRA_ARGUMENTS = !ThenElse( isUndefined(TARGET.SingularityExtraArgs), "", TARGET.SingularityExtraArgs)

+SingularityImage = "cvmsfs/singularity.opensciencegrid.org/opensciencegrid/osgvo-el8:latest"
+SingularityBindCVMFS = True
+SingularityBind = "/cvmsfs, /cms, /cms_scratch"
## If you want to use apptainer's extra arguments like "--nv",
##+SingularityExtraArgs = "--nv"
```

STARTD

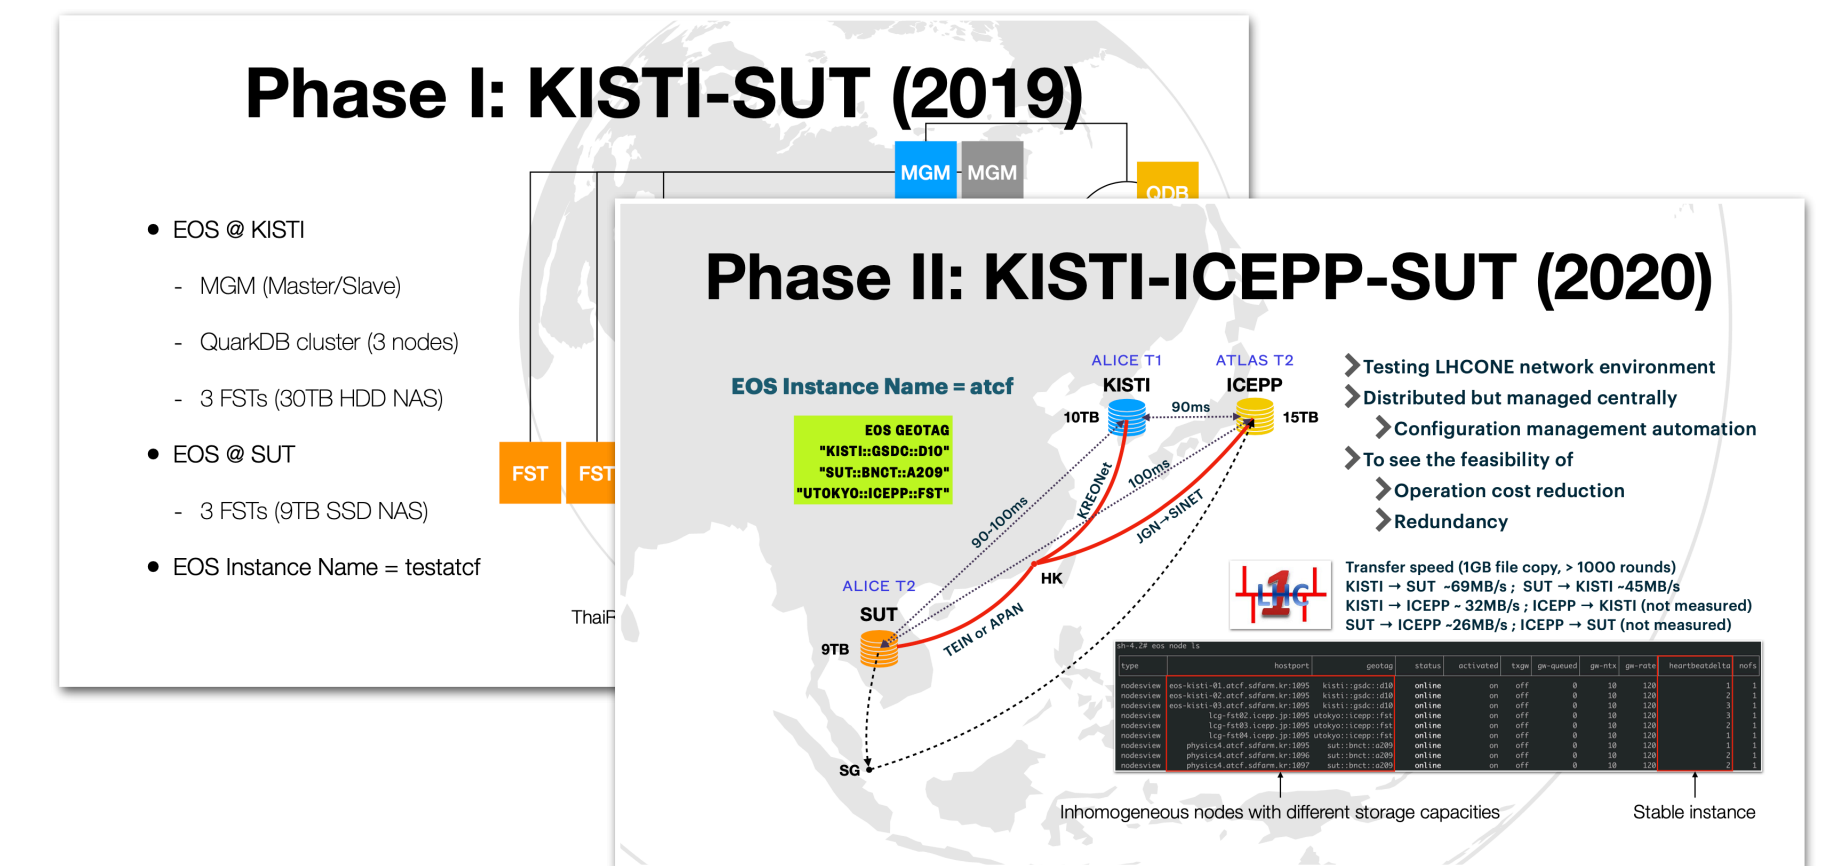
Job Description File(.sub)

- No preemption configuration for WNs
 - To increase the Preemption speed, a setting that divides the dynamic slots on the WN into multiple is required
 - The vacating is 1 per primary slot in a negotiation cycle.
- The integrated farm mainly uses Apptainer runtime rather than Docker
 - This is because there is no need to enable a docker daemon or manage the docker user list
 - Also, Docker container images are available at apptainer runtime
- To enable Singularity(apptainer) container environment,
 - Setup above configuration
 - Add related classads to a job description file
- If fusemount3 error occurs after updating the apptainer instead of the singularity,
 - If you are not familiar with rootless (unprivileged) container settings, install the apptainer-suid package.
- If you download an image from each WN, there may be a limit on the number of downloads of dockerhub homepage
 - It is also used to download the container image to a file and put it in a shared directory

Distributed Storage Project

Lessons learned

- Good practical tool to test the LHCONE networking environment
 - Dedicated routing or virtual path for the distributed storage, TCP tunings on end-to-end nodes, advanced technologies (e.g. latency hiding, caches) could help improve performance
 - Tight cooperation with experts (from networking, storage and experiments) are mandatory
 - Further development and study, or trying different approaches may improve performance as well as usability
- Feasibility of knowledge (technology) transfer and community building
 - KISTI & SUT instances were managed together; ICEPP was independent (individually managed)
 - Local training session for students/staff at SUT (2019), frequent conversations via e-mail or zoom meetings with ICEPP
 - ICEPP used the automation script provided by KISTI (forked from EOS docker project) ⇒ can be used not only for distributed storage but also for standalone EOS storage
- Putting extra efforts is not enough, a firm foundation for sustainable support is required



Working Group

Discussion on proposed actions and directions

- A working group was proposed to discuss challenges and proposed actions as well as directions for launching the Asian Big Science Cloud (ABC), a strong Asian site collaboration.
 - Challenges: (operators) lack of knowledges, expertises, training or tracking opportunities on technologies / (sites) difficulties on securing resources, manpower, budget and infrastructure ...
 - Proposed actions:
 - supporting participation to various tutorials/workshops, inviting experts, organizing WG for technology tracking, training, commissioning, developing and jointly taking part in activities in WLCG, and so on
 - applying to Asi@Connect or intercountry programs (e.g. ID-RoK), hardware donation, etc.
 - ABC - a ATCF driven version of collaboration (envisaging a firm foundation) to support Asian computing sites involving in LHC/non-LHC experiments

Mandate of Working Group

- Discuss in detail the proposed actions (with any further ideas and suggestions) to form a support model in Asia
- Share the list of actions for the Asian support model within ATCF and finalize them
- Propose joint-sessions (a few times) in related conferences or workshops, e.g. ISGC, APAN, CHEP, WLCG workshop, etc. to share and discuss a draft of support model for completion
- Apply to S&T cooperation programs available in ASEAN, ASEAN+3, EAS, etc. for knowledge transfer, HR training, IT collaboration and so on.
- WG members: Andria and Syam (BRIN), Masahiko (ICEPP), Brij and Puneet (TIFR), Vikas (VECC) and Sang-Un (KISTI)

Summary

- The ATCF6 was successful with full of invaluable contributions and fruitful discussions
 - The venue was great and well organised thanks to warm hospitality of Prof. Chinorat Kobdaj and his SUT staff
- Asian sites are now in benefit from full network reachability thanks to the effort of LHCONE community
 - The short-term goal of ATCF has been almost accomplished
- A WG was formed to search for a sustainable support model in the region



Links

- Asia Tier Center Forum webpage [link](#)
- ATCF6 Indico [link](#)
- ATCF Mattermost channel ([join us](#))