Distributed Storage Project

The 6th ATCF @ Krabi, Thailand
21 - 24 November 2022

Sang-Un Ahn, Jeong-Heon Kim, Chinorat Kobdaj, Michiru Kaneda, Tomoe Kishimoto

Data Challenges for Upcoming LHC RUNs

- Data challenges foreseen in HL-LHC (RUN4)
  - x10 more integrated luminosity compared to RUN3 (x20 than RUN2)
    * Lots of efforts on optimization of data/computing models from the experiments
  - Demanding significant increase of compute and storage capacities
  - How do to deliver them in flat-budget scenarios? *(still effective??)*
- *"Storage is the main operational cost at sites"*
  - WLCG 2015 Survey(https://twiki.cern.ch/twiki/bin/view/LCG/WLCGSiteSurvey)
  - Disk costs 4x more than tape per TB
- WLCG initiated DOMA to cope with the challenges
- Distributed storages implementing data-lake models
  - Open Storage Network (US, Ceph), CloudStor of AARNet (AU, EOS), ESCAPE Data Lake (EU)
    European Science Cluster of Astronomy & Particle physics ESFRI (European Strategy Forum on Research Infrastructure)
  - Advanced technology and large bandwidth networking are mandatory

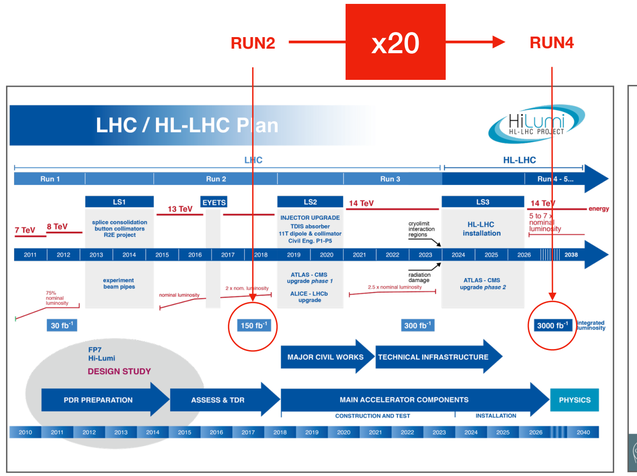

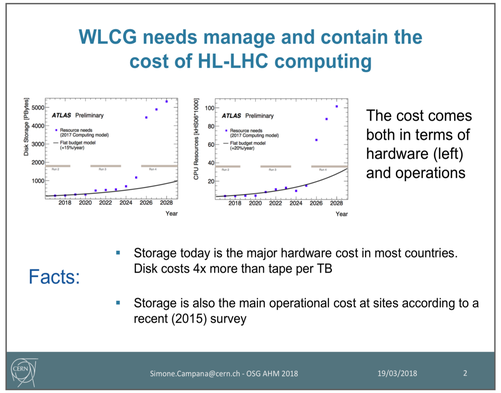

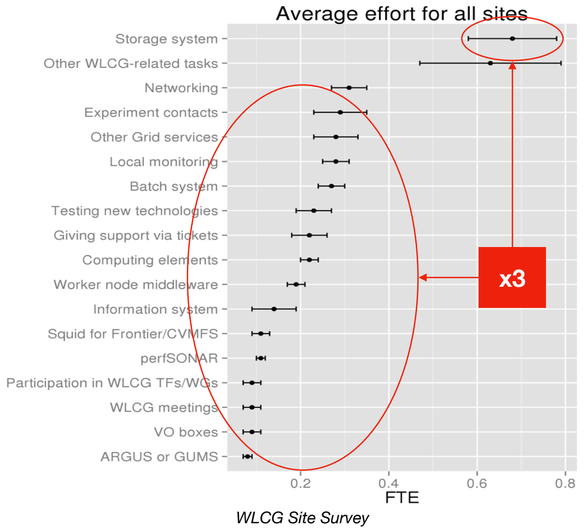

WLCG Site Survey

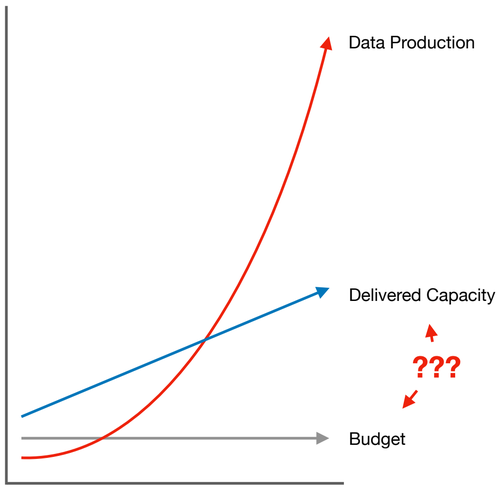

Distributed Storage in Asia

- A strong collaboration is needed to overcome **Data Challenges foreseen in HL-LHC era**
  - Resource requirements to T1/T2 sites from experiments will increase accordingly
  - Reducing the operational costs is the key; Technology advances? → Consolidated efforts are needed
- **Distributed Storage** across Asian sites
  - Started the discussion at ATCF4
  - *A handful tool to exploit and evaluate the advanced networking in Asia*
  - A test-bed was established together at KISTI-ICEPP-SUT (2018-2021)

Discussion Points @ ATCF4

- Improving latencies and bandwidths among distributed sites(storages)
- Proving data transfer capacity between distributed sites upon the current (as of 2018) networking configuration
- Reflecting different requirements from different VOs, e.g. ATLAS, CMS, ALICE with a single(?) distributed storage
- Reducing operational costs meeting diverse use cases
- Sharing expertise and technologies
- Proposing to setup a distributed storage between KISTI and SUT to address issues above
  - Consolidation of distributed storage using EOS and providing a single entry point

The Nordic Model

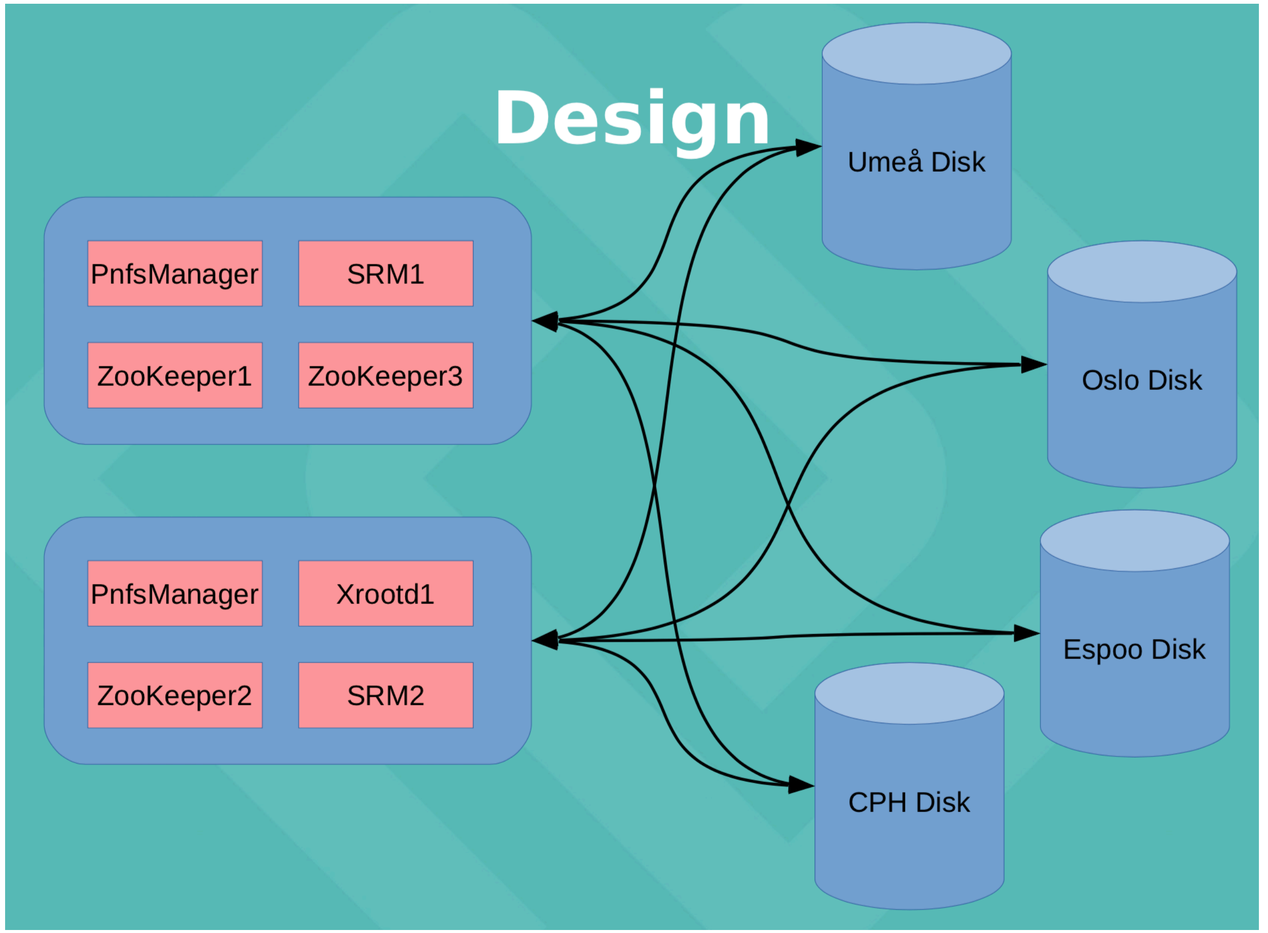

NeIC Distributed Storage Design

Strong motivation
Consolidated collaboration
Sharing expertise
Co-work on technologies

dCache
ARC

Mattias Wadenstein, maswan@ndgf.org

A Nordic Model

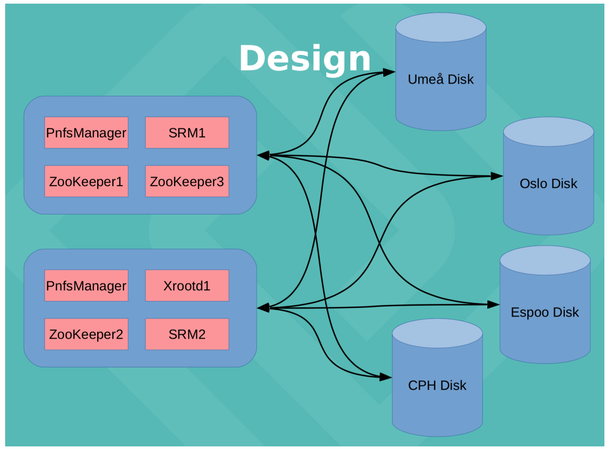

NeIC Distributed Storage Design

44,579,000 km²

dCache
ARC

Distributed/Sync (covering regions)

Central Services

Central Services

Single Entry Point
Global Namespace
Cache
Storage Management
(HA-enabled)

Far-west Disk

South-west Disk

8,000 km

South-east Disk

Far-east Disk

Strong motivation
Consolidated collaboration
Sharing expertise
Co-work on technologies

Progress

- Agreement @ ATCF4 in November 2018
  - KISTI and SUT agreed to establish a distributed storage setup in order to demonstrate the feasibility of storage consolidation in the region
  - Experts and students exchange program
- KISTI-SUT EOS Workshop @ SUT in August 2019
  - EOS deployment based on Docker container and using Ansible playbook (written by Dr. Kim, Jeong-Heon)
  - 3 Days of Tutorial for SUT students including EOS
- Review & Discussion @ ATCF5 in October 2019
  - Issues shared, ICEPP expressed the participation to the project
- Presented to EOS Workshop @ CERN in February 2020
  - Reported the project to the workshop & ICEPP started deploying EOS using Ansible playbook
- KISTI-ICEPP-SUT Distributed Storage setup established in August 2020 and kept testing until early of 2021
  - Network tuning, performance tests, trying different setups and so on

EOS at a Glance

Introduction and Current status of EOS project will be given by Andreas-Joachim Peters on Day 3

- Components
  - MGM : Management Server
  - QuarkDB : Namespace
    - In-memory Namespace will be deprecated soon
  - MQ : Message Queue
  - FST : File Storage Server
- Important features for a distributed storage
  - Location awareness
  - GEO Scheduling functionality

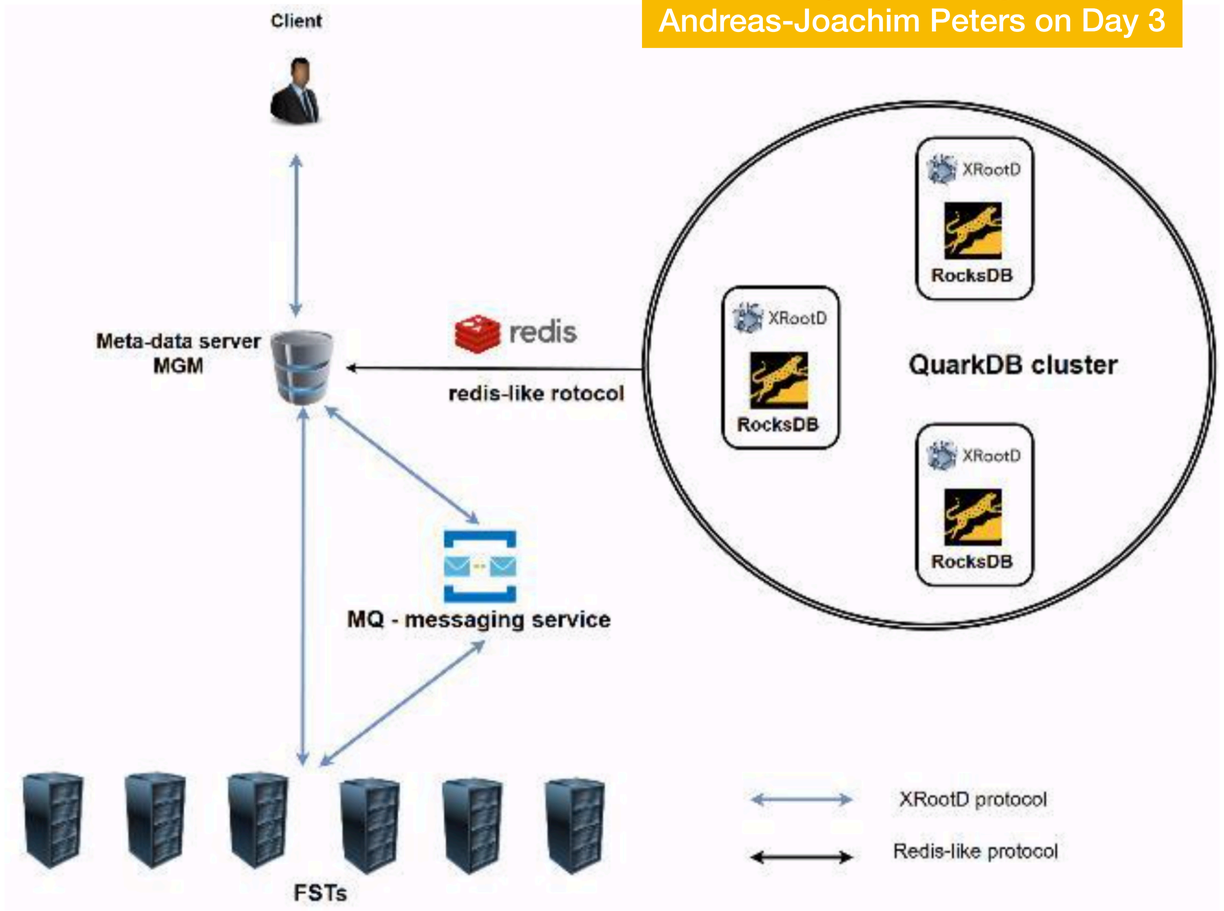

Case Study

- CERN tested a distributed storage setup using EOS between Meyrin and Wigner
  - "di-EOS - "distributed EOS": Initial experience with split-site persistency in a production service" presented @ CHEP2013
  - 22ms latency, 100Gbit/s between the two sites
- CERN, AARNET(AU), and ASGC(TW) tried to setup and test EOS deployment in wide area network
  - "Global EOS: exploring the 300-ms-latency region" presented @ CHEP2016
  - Latency > 300ms, 16,500km apart

Global EOS Conclusion

- Confirmed that,
  - "... the stability and the robustness of EOS in working with such latency, no adaptation of timeouts or other parameter was needed in order to set up the system on this very large geographical scale,"
  - **"the system worked immediately out of the box."**
- Client behaviour @ Melbourne writes to disk pool @ Melbourne
  - "... contacted the read-write namespace located in Geneva and the data transfers is scheduled to a Melbourne disk."
  - Read is not affected by such a big round trip time
- Average speed of data transfers in MEL-GVA ~ 45MB/s

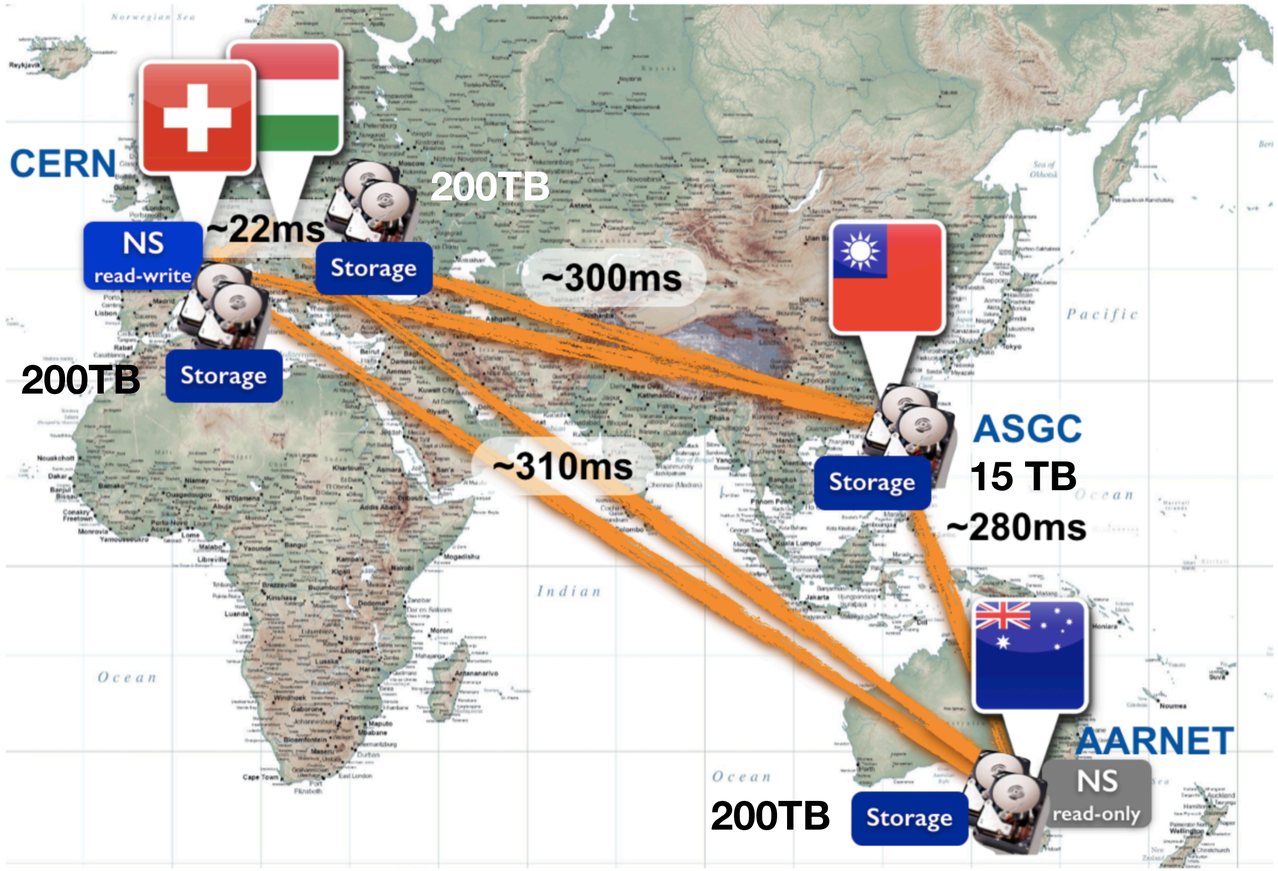

Journal of Physics: Conf. Series 898 (2017) 062029

Phase I: KISTI-SUT (2019)

- EOS @ KISTI
  - MGM (Master/Slave)
  - QuarkDB cluster (3 nodes)
  - 3 FSTs (30TB HDD NAS)
- EOS @ SUT
  - 3 FSTs (9TB SSD NAS)
- EOS Instance Name = testatcf

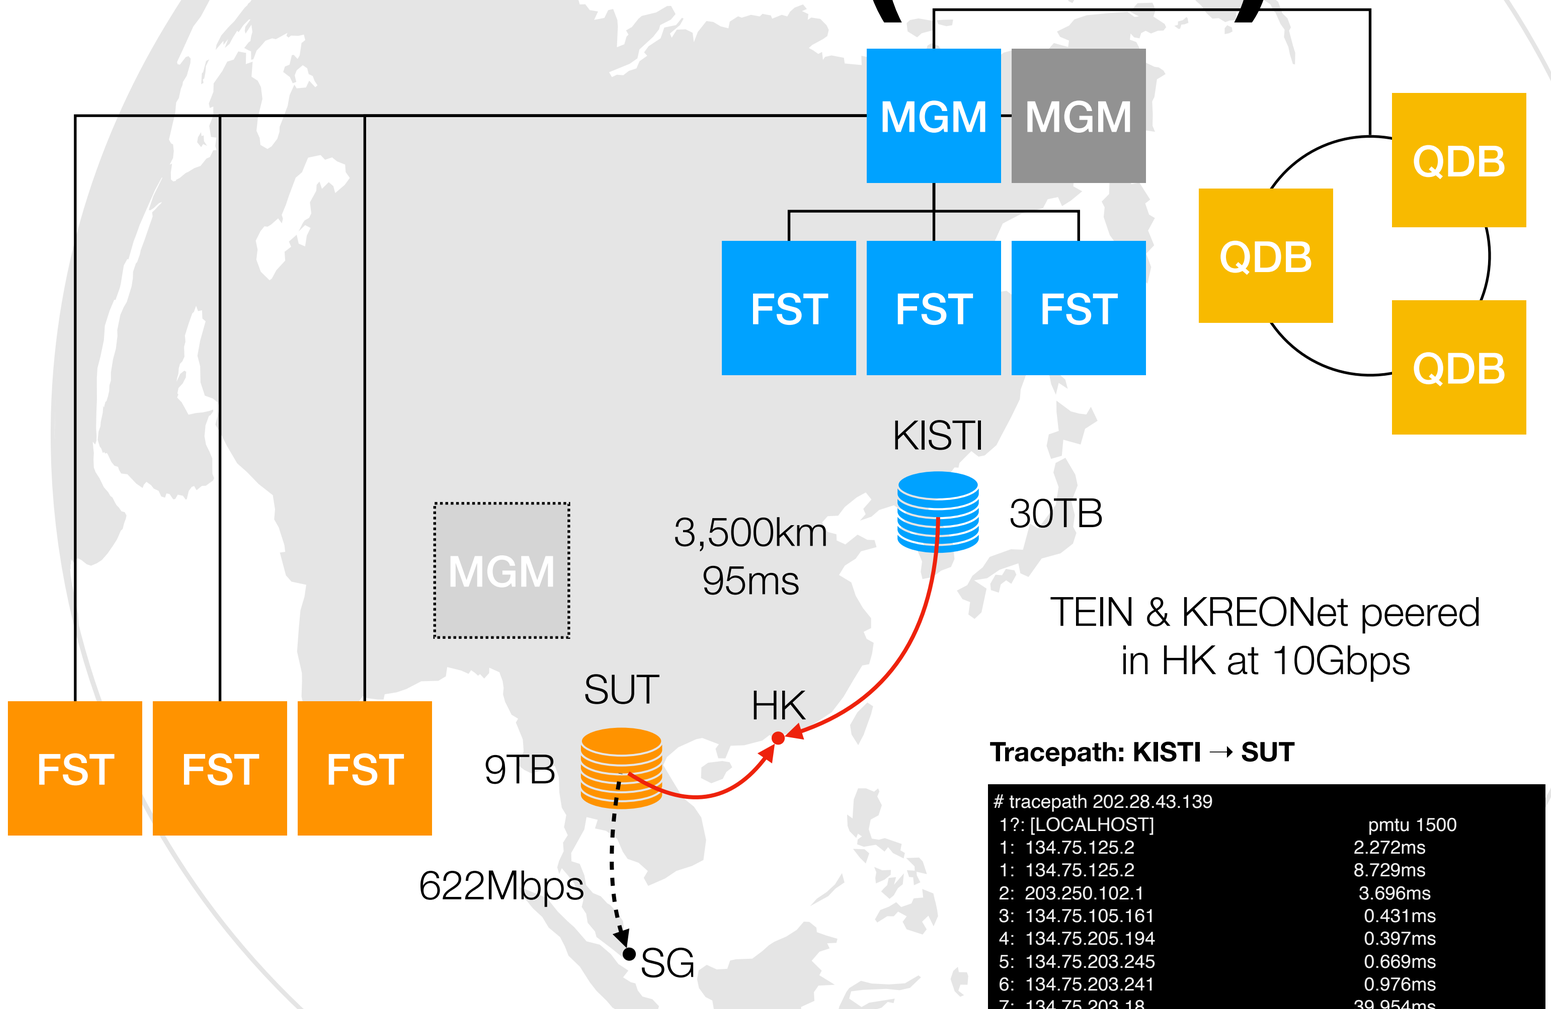

TEIN & KREONet peered in HK at 10Gbps

ThaiREN reaches HK at 2Gbps; SG at 1Gbps

Tracepath: KISTI → SUT

```
# tracepath 202.28.43.139
 1?: [LOCALHOST]                      pmtu 1500
 1:  134.75.125.2                      2.272ms
 1:  134.75.125.2                      8.729ms
 2:  203.250.102.1                     3.696ms
 3:  134.75.105.161                    0.431ms
 4:  134.75.205.194                    0.397ms
 5:  134.75.203.245                    0.669ms
 6:  134.75.203.241                    0.976ms
 7:  134.75.203.18                    39.954ms
 8:  202.179.241.205                  44.706ms
 9:  202.179.241.210                  91.354ms
10:  pyt-to-02-bdr-pyt-link-1.uni.net.th   91.229ms
11:  100.64.253.13                    96.071ms asymm 14
12:  202.28.208.254                   94.953ms asymm 16
13:  202.28.43.139                    95.587ms reached
     Resume: pmtu 1500 hops 13 back 17
```

Phase II: KISTI-ICEPP-SUT (2020)

EOS Instance Name = atcf

EOS GEOTAG
"KISTI::GSDC::D10"
"SUT::BNCT::A209"
"UTOKYO::ICEPP::FST"

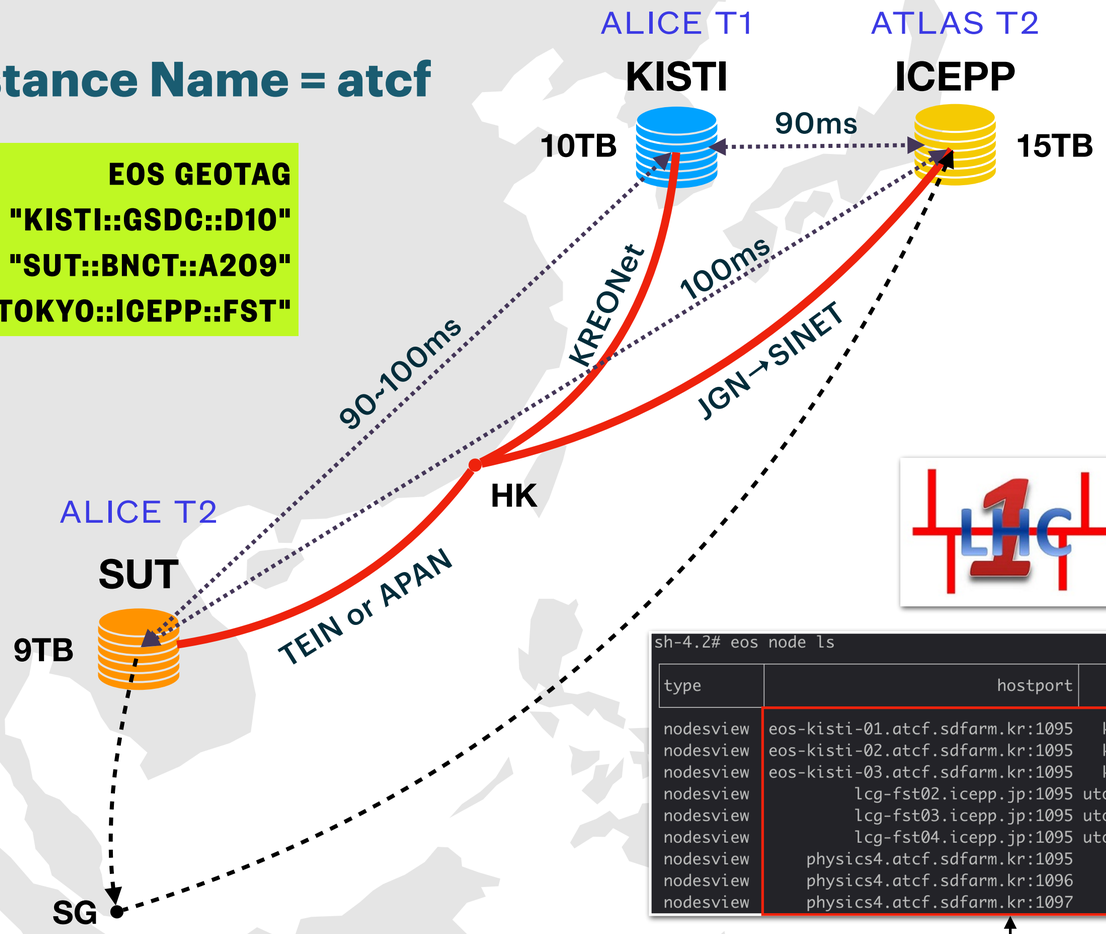

- Testing LHCONE network environment
- Distributed but managed centrally
  - Configuration management automation
- To see the feasibility of
  - Operation cost reduction
  - Redundancy

Transfer speed (1GB file copy, > 1000 rounds)
KISTI → SUT ~69MB/s ; SUT → KISTI ~45MB/s
KISTI → ICEPP ~ 32MB/s ; ICEPP → KISTI (not measured)
SUT → ICEPP ~26MB/s ; ICEPP → SUT (not measured)

```
sh-4.2# eos node ls
```

type	hostport	geotag	status	activated	txgw	gw-queued	gw-ntx	gw-rate	heartbeatdelta	nofs
nodesview	eos-kisti-01.atcf.sdfarm.kr:1095	kisti::gsdc::d10	online	on	off	0	10	120	1	1
nodesview	eos-kisti-02.atcf.sdfarm.kr:1095	kisti::gsdc::d10	online	on	off	0	10	120	2	1
nodesview	eos-kisti-03.atcf.sdfarm.kr:1095	kisti::gsdc::d10	online	on	off	0	10	120	3	1
nodesview	lcg-fst02.icepp.jp:1095	utokyo::icepp::fst	online	on	off	0	10	120	3	1
nodesview	lcg-fst03.icepp.jp:1095	utokyo::icepp::fst	online	on	off	0	10	120	2	1
nodesview	lcg-fst04.icepp.jp:1095	utokyo::icepp::fst	online	on	off	0	10	120	1	1
nodesview	physics4.atcf.sdfarm.kr:1095	sut::bnct::a209	online	on	off	0	10	120	1	1
nodesview	physics4.atcf.sdfarm.kr:1096	sut::bnct::a209	online	on	off	0	10	120	2	1
nodesview	physics4.atcf.sdfarm.kr:1097	sut::bnct::a209	online	on	off	0	10	120	2	1

Inhomogeneous nodes with different storage capacities

Stable instance

Best Practice for Distributed EOS

- MGM and/or QDB distribution across sites (KISTI-ICEPP-SUT) didn't work well
  - Natively the distribution worked as it is as expected in terms of functionality
  - Performance-wise MGM/QDB distribution between sites didn't look good
    - MGM slaves redirected any requests to the primary MGM (if the primary is dead, the requests were timed-out)
    - Any namespace related requests such as file copy, stat and rm were hanged frequently in the case of distributed QDB clusters (KISTI-ICEPP)
      - Performance dropped down to tens of KB/s (affected by paths??)
- Centrally managed MGMs and QDBs with distributed FSTs worked best
  - Haven't tried complete distribution of MGMs and QDB cluster

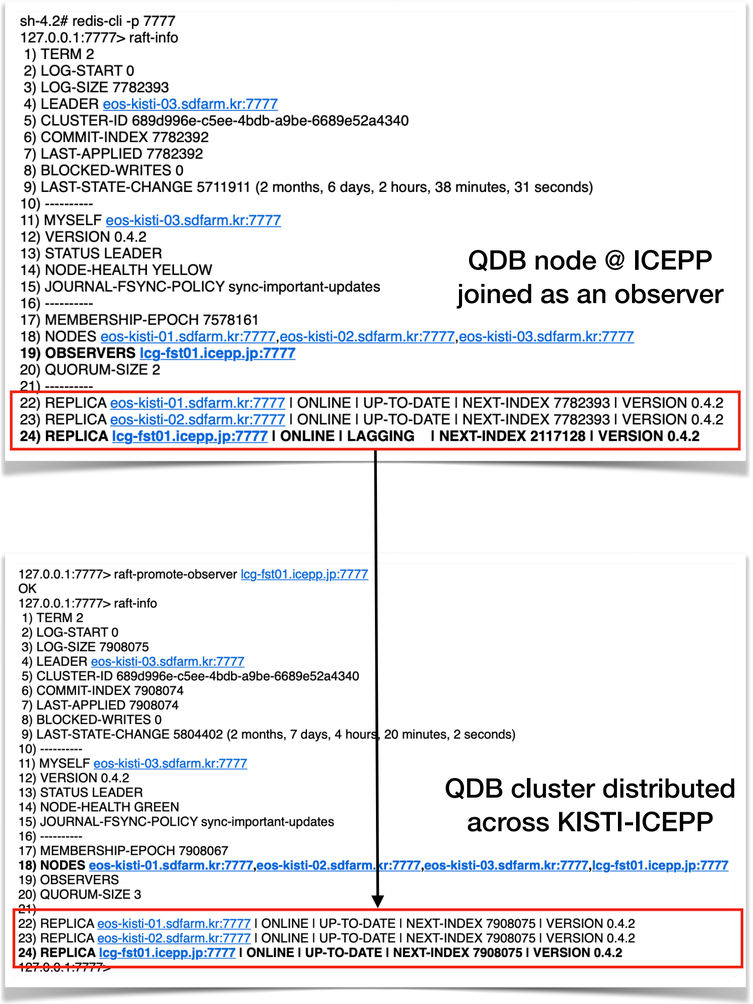

Observation

- Not-static end-to-end routing paths via LHCONE
  - Paths were depending on the status of network providers
    - KREONet helped adjust paths to ICEPP and SUT from KISTI
  - Performance was significantly affected (from few MB/s up to tens of MB/s)
- ~100ms of latencies (induced by physical distances) are great hurdles
  - Entirely distributed setup of essential EOS components looks unrealistic
  - Intensive efforts and tweaks may improve the performance of distributed setup
  - Technologies to hide latencies are worth to consider
    (content delivery networks, caches, etc.)

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, frequent conversions 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, dedicated man-power is necessary

Discussion

- Reviving the distributed storage setup and having further steps
  - TCP tunings, path optimization, complete distribution of EOS components (again)
  - Looking for a proper use case (ALICE, ATLAS, CMS, Belle II, etc. or non-LHC domains)
    - Requires to be interfaced with Data Management frameworks as well as Workload Management frameworks
  - Expanding partners
- Searching for new approaches towards joining efforts for training, sharing expertise and technologies, development to reduce operational costs
  - A firm foundation may be needed for sustainable support of such activities by encouraging the involvement of experts rather than the solicitation of helps