



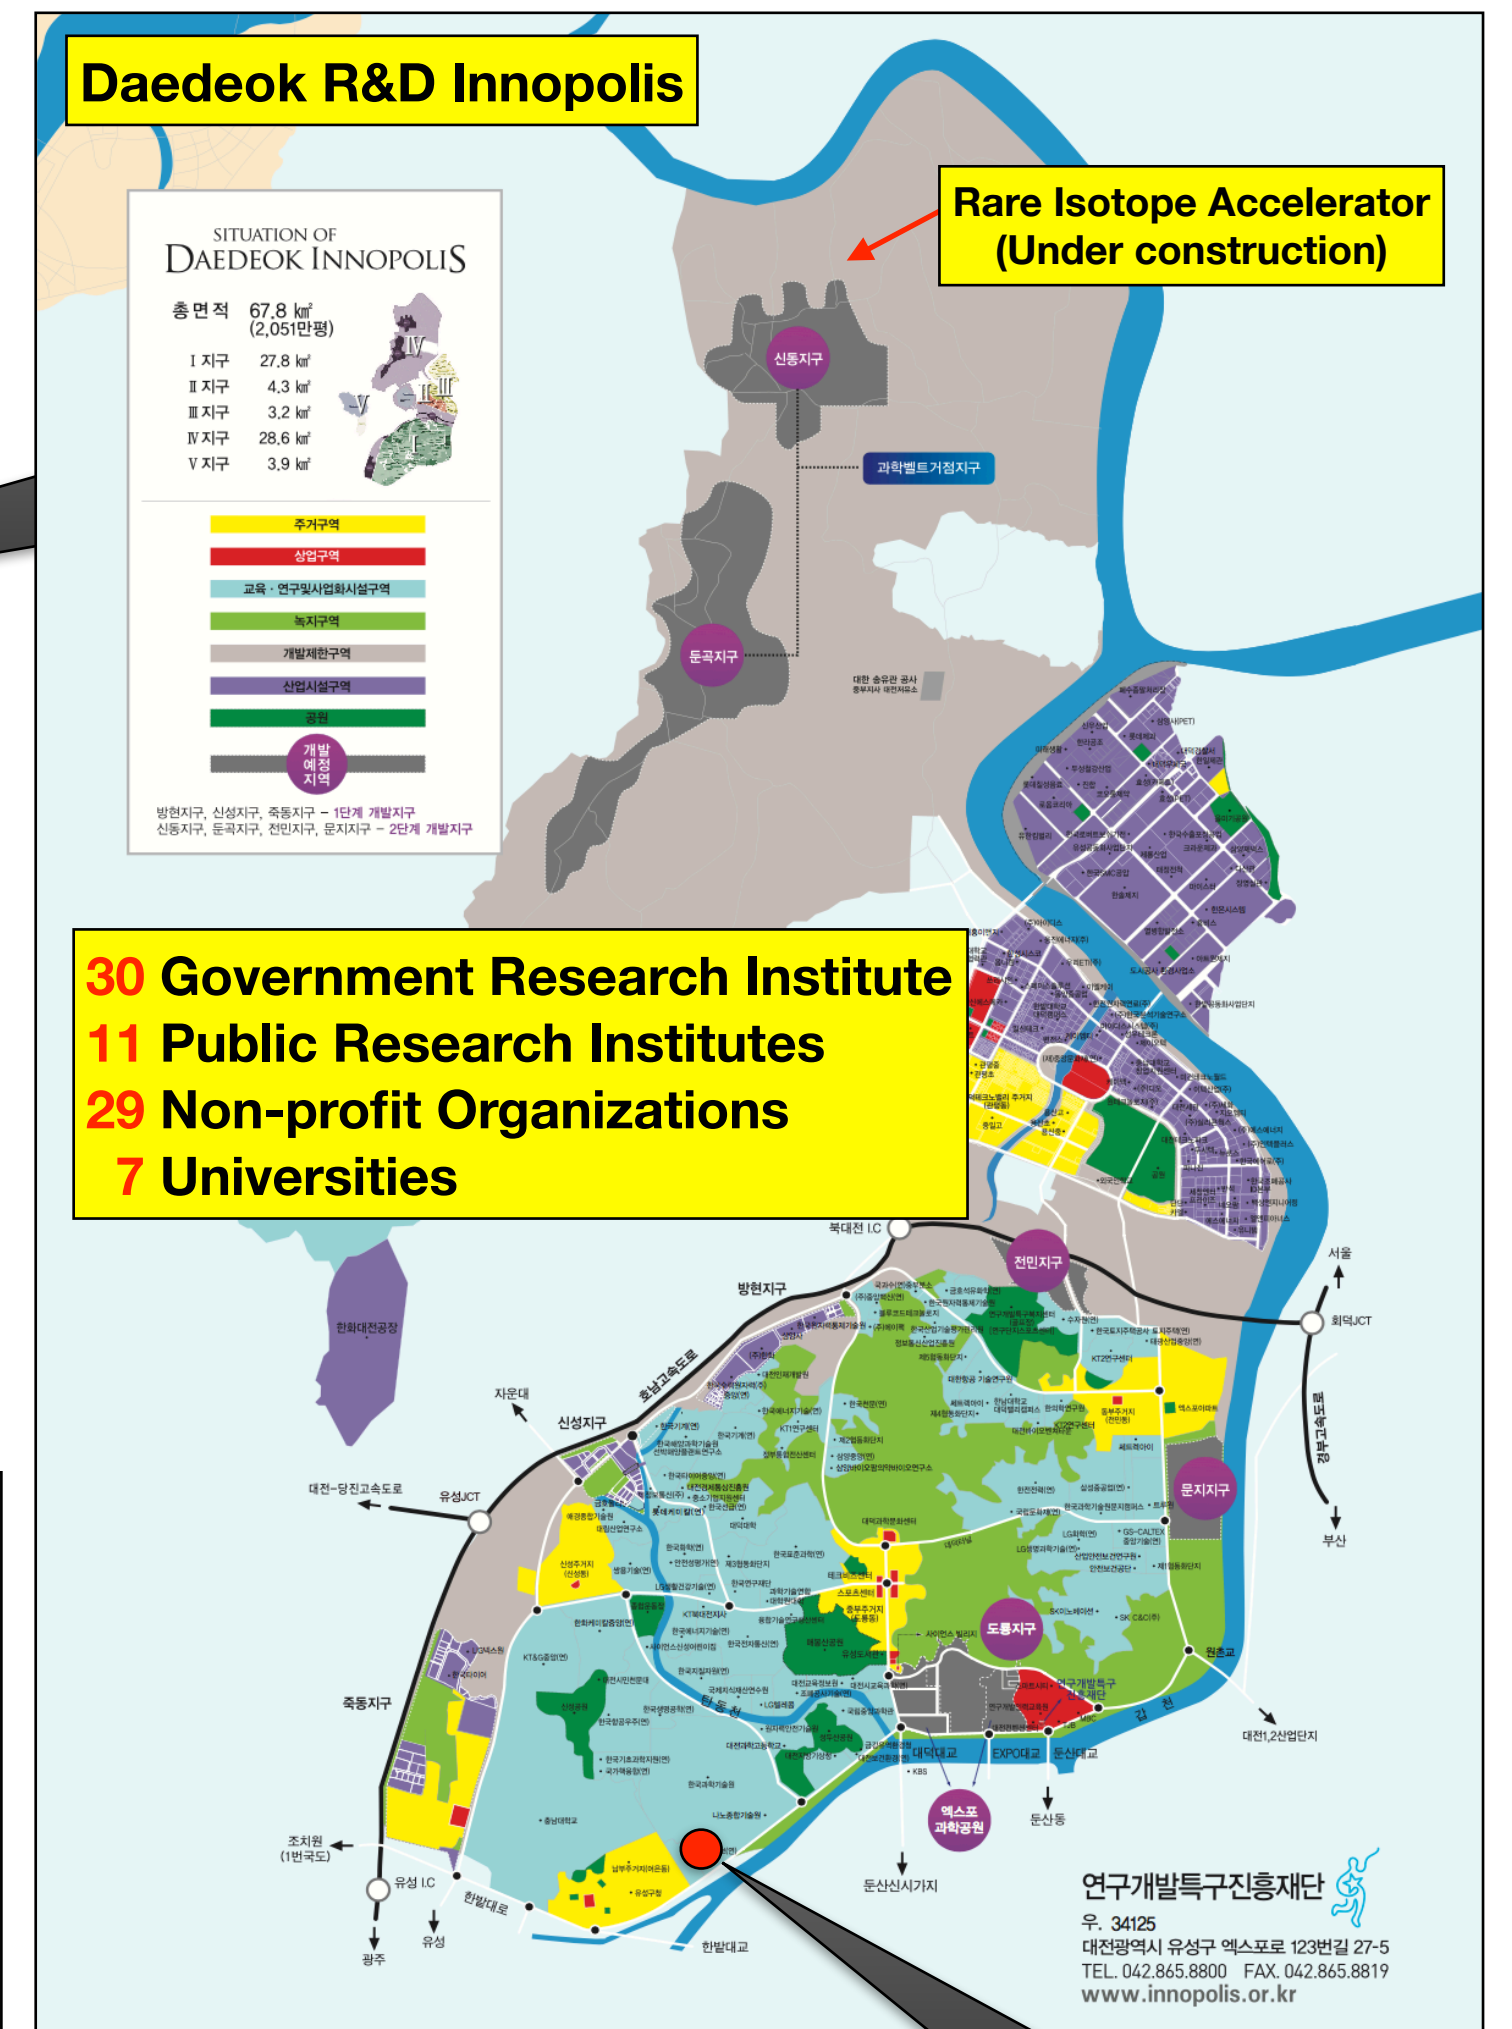
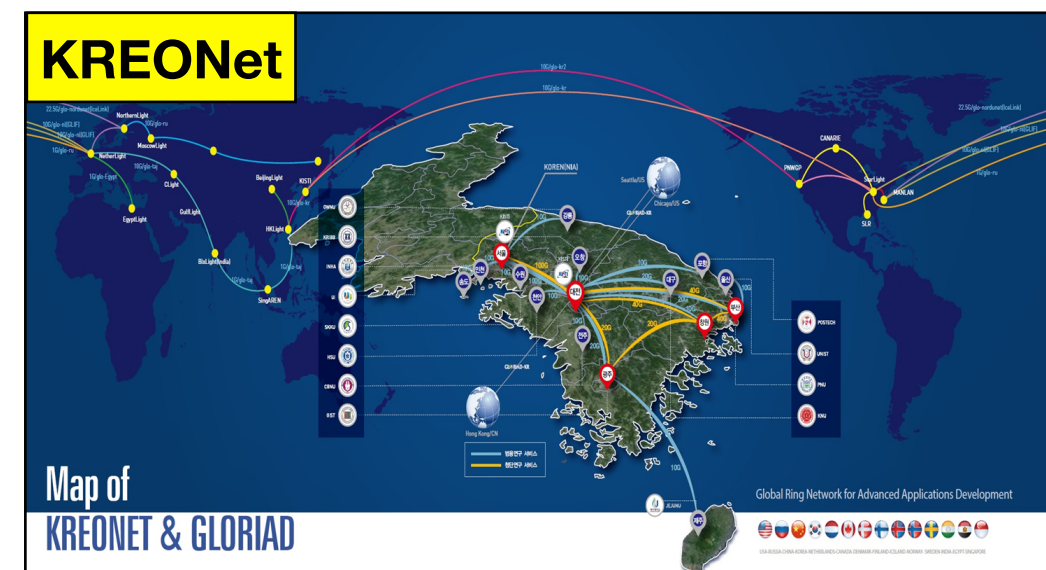
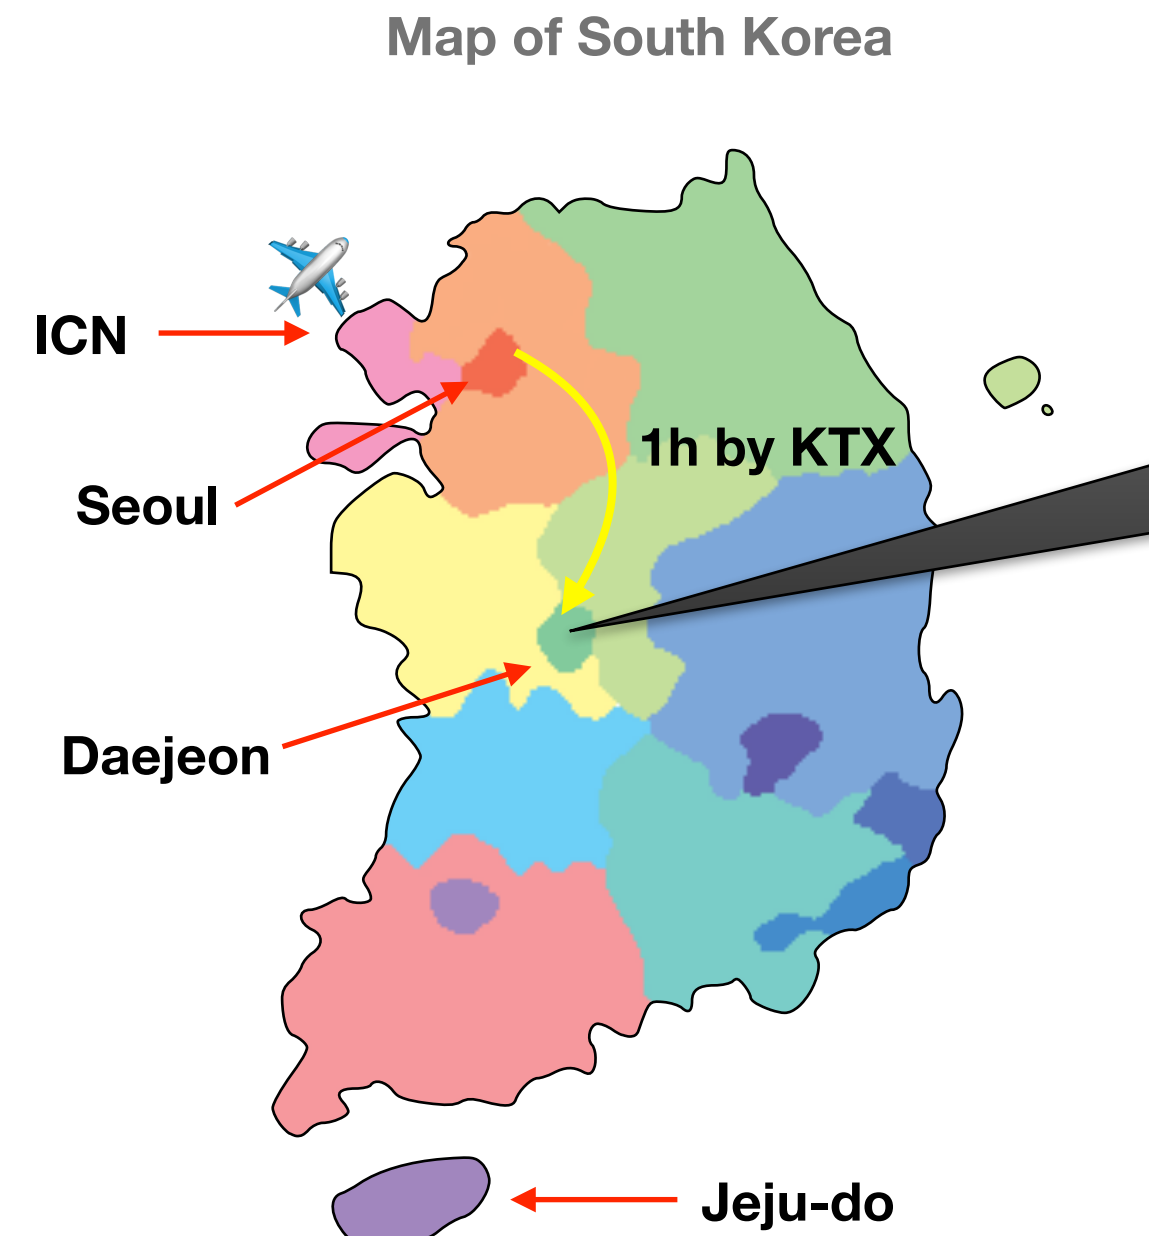
KISTI Site Report

Ahn Sang-Un for KISTI-GSDC
@ ATCF6

KISTI

Korea Institute of Science and Technology Information

- Government-funded research institute founded in 1962 for National Information Services and Supercomputing
- National Supercomputing Center
 - **Nurion** - Cray CS500 system
 - 25.7 PFlops at peak, ranked 11th of Top500 (2018) => 21st (Nov 2020)
 - **Neuron** - GPU system, 1.24 PFlops
 - **KREONet/KREONet2** - National R&E network




GSDC

Global Science experimental Data hub Center


- Government-funded project, started in 2009 to promote Korean fundamental research through providing computing power and data storage
- Datacenter for data-intensive fundamental research
 - 16 staff: system administration, experiment support, external-relation, management and planning



ALICE



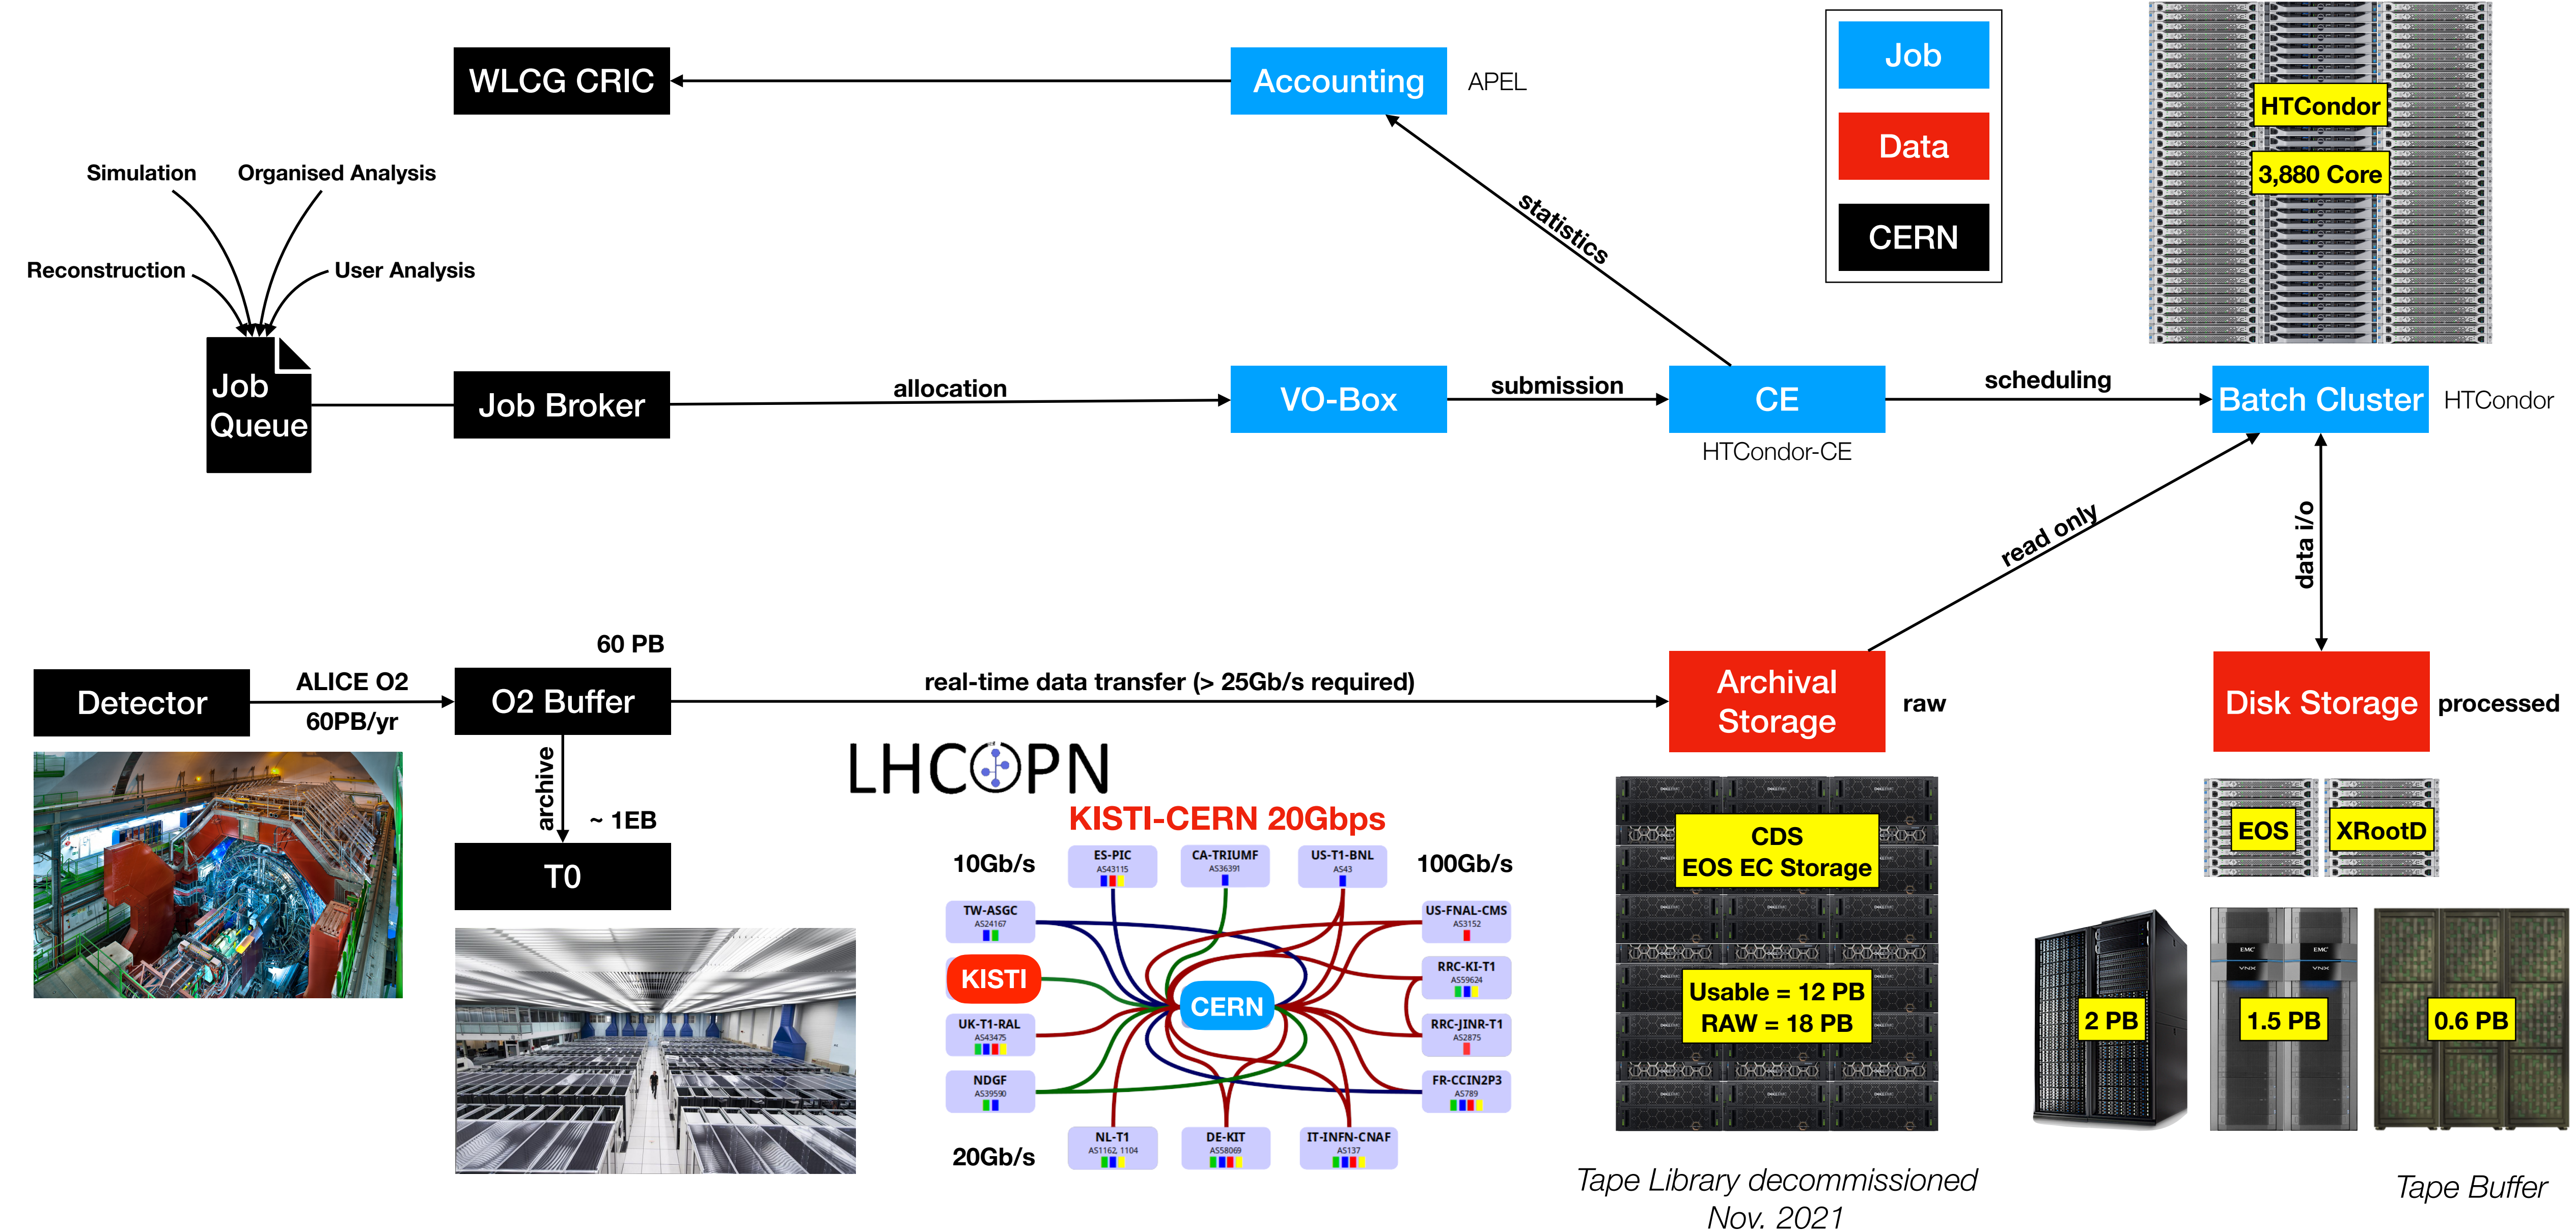
International Cancer Genome Consortium



Experiments Support

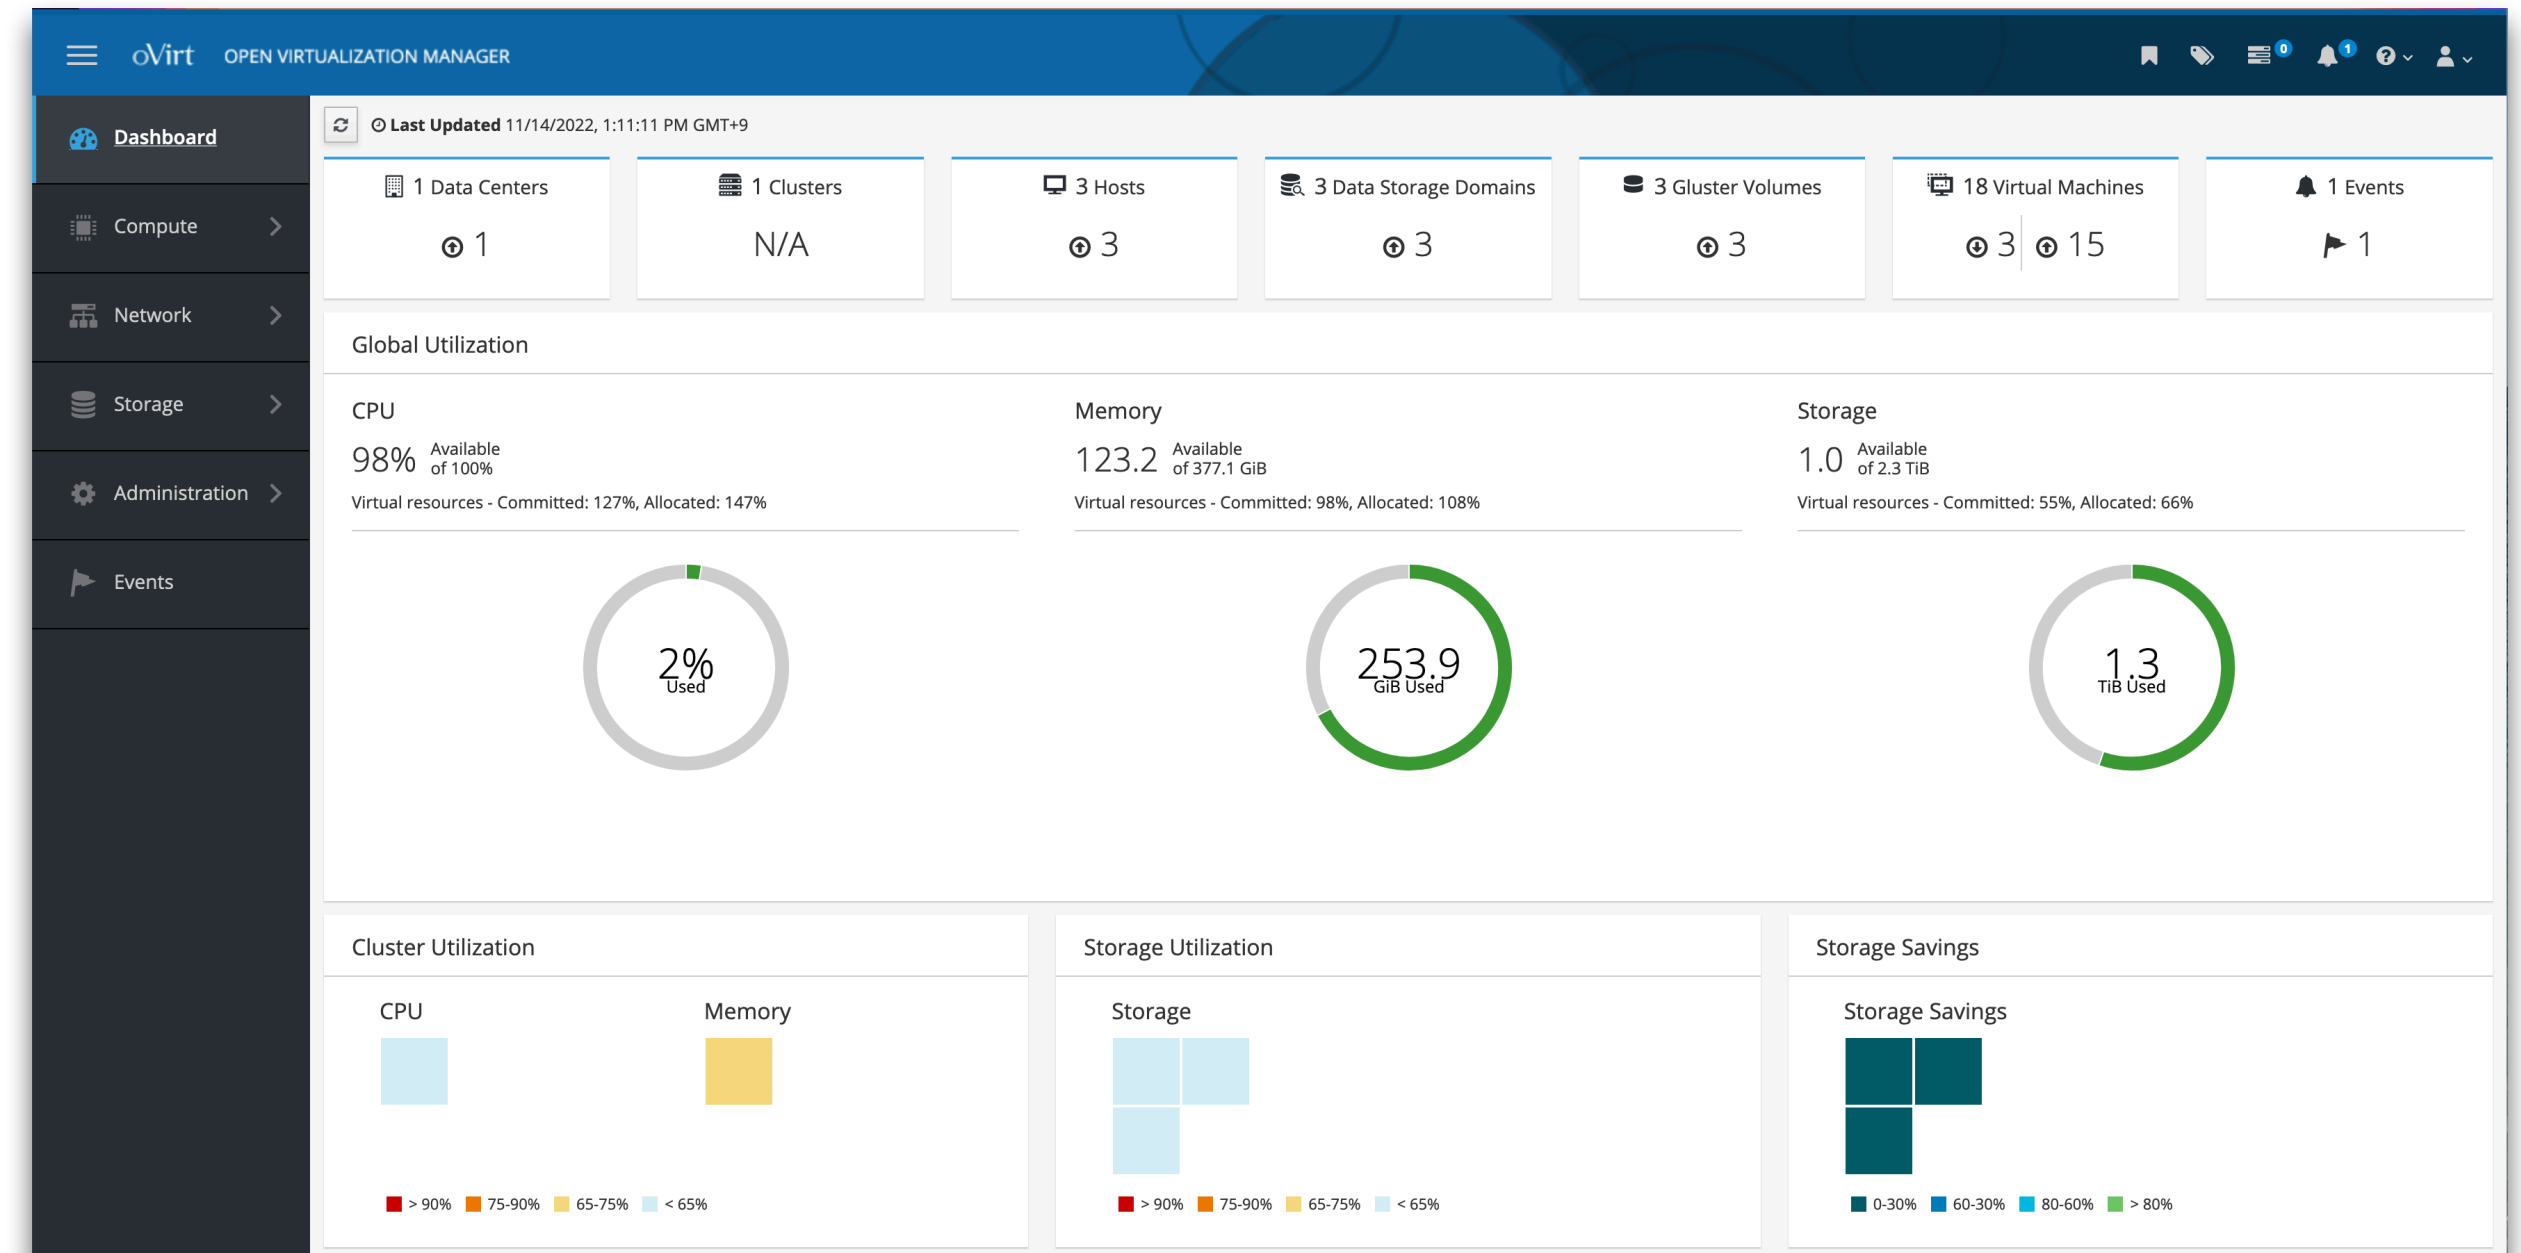


KISTI ALICE Tier-1 Structure Overview



Essential Grid Services

- Grid services running on VMs provided by oVirt cluster
 - oVirt v4.3.8, GlusterFS v6.10
 - 3 oVirt hosts with 384 GB of RAM and 2.3TB of Gluster Storage (1.5 TB of HDDs, 0.8 TB of SSDs)
 - Live migration & load-balancing
- 15 VMs
 - VO-Box (ALICE Job Submission, JAlEn enabled)
 - 3 HTCondor-CEs (CE 5.1.5, Condor 9.0.14)
 - Site-BDII & Argus (AuthN & AuthZ)
 - EOS MGM nodes & XRootD redirectors
 - EOS QDB clusters (deployed upon SSD disk groups)
 - 3 Squid caches for CernVM-FS
 - APEL (Accounting)



Name	Comment	Hostname/IP	Cluster	Data Center	Status	Virtual Machines	Memory	CPU	Network	SPM
alice-ovirt-01.sdfarm.kr		alice-ovirt-01.sdfarm.kr	Default	Default	Up	4	63%	3%	0%	Normal
alice-ovirt-02.sdfarm.kr		alice-ovirt-02.sdfarm.kr	Default	Default	Up	6	81%	2%	0%	SPM
alice-ovirt-03.sdfarm.kr		alice-ovirt-03.sdfarm.kr	Default	Default	Up	5	78%	4%	0%	Normal

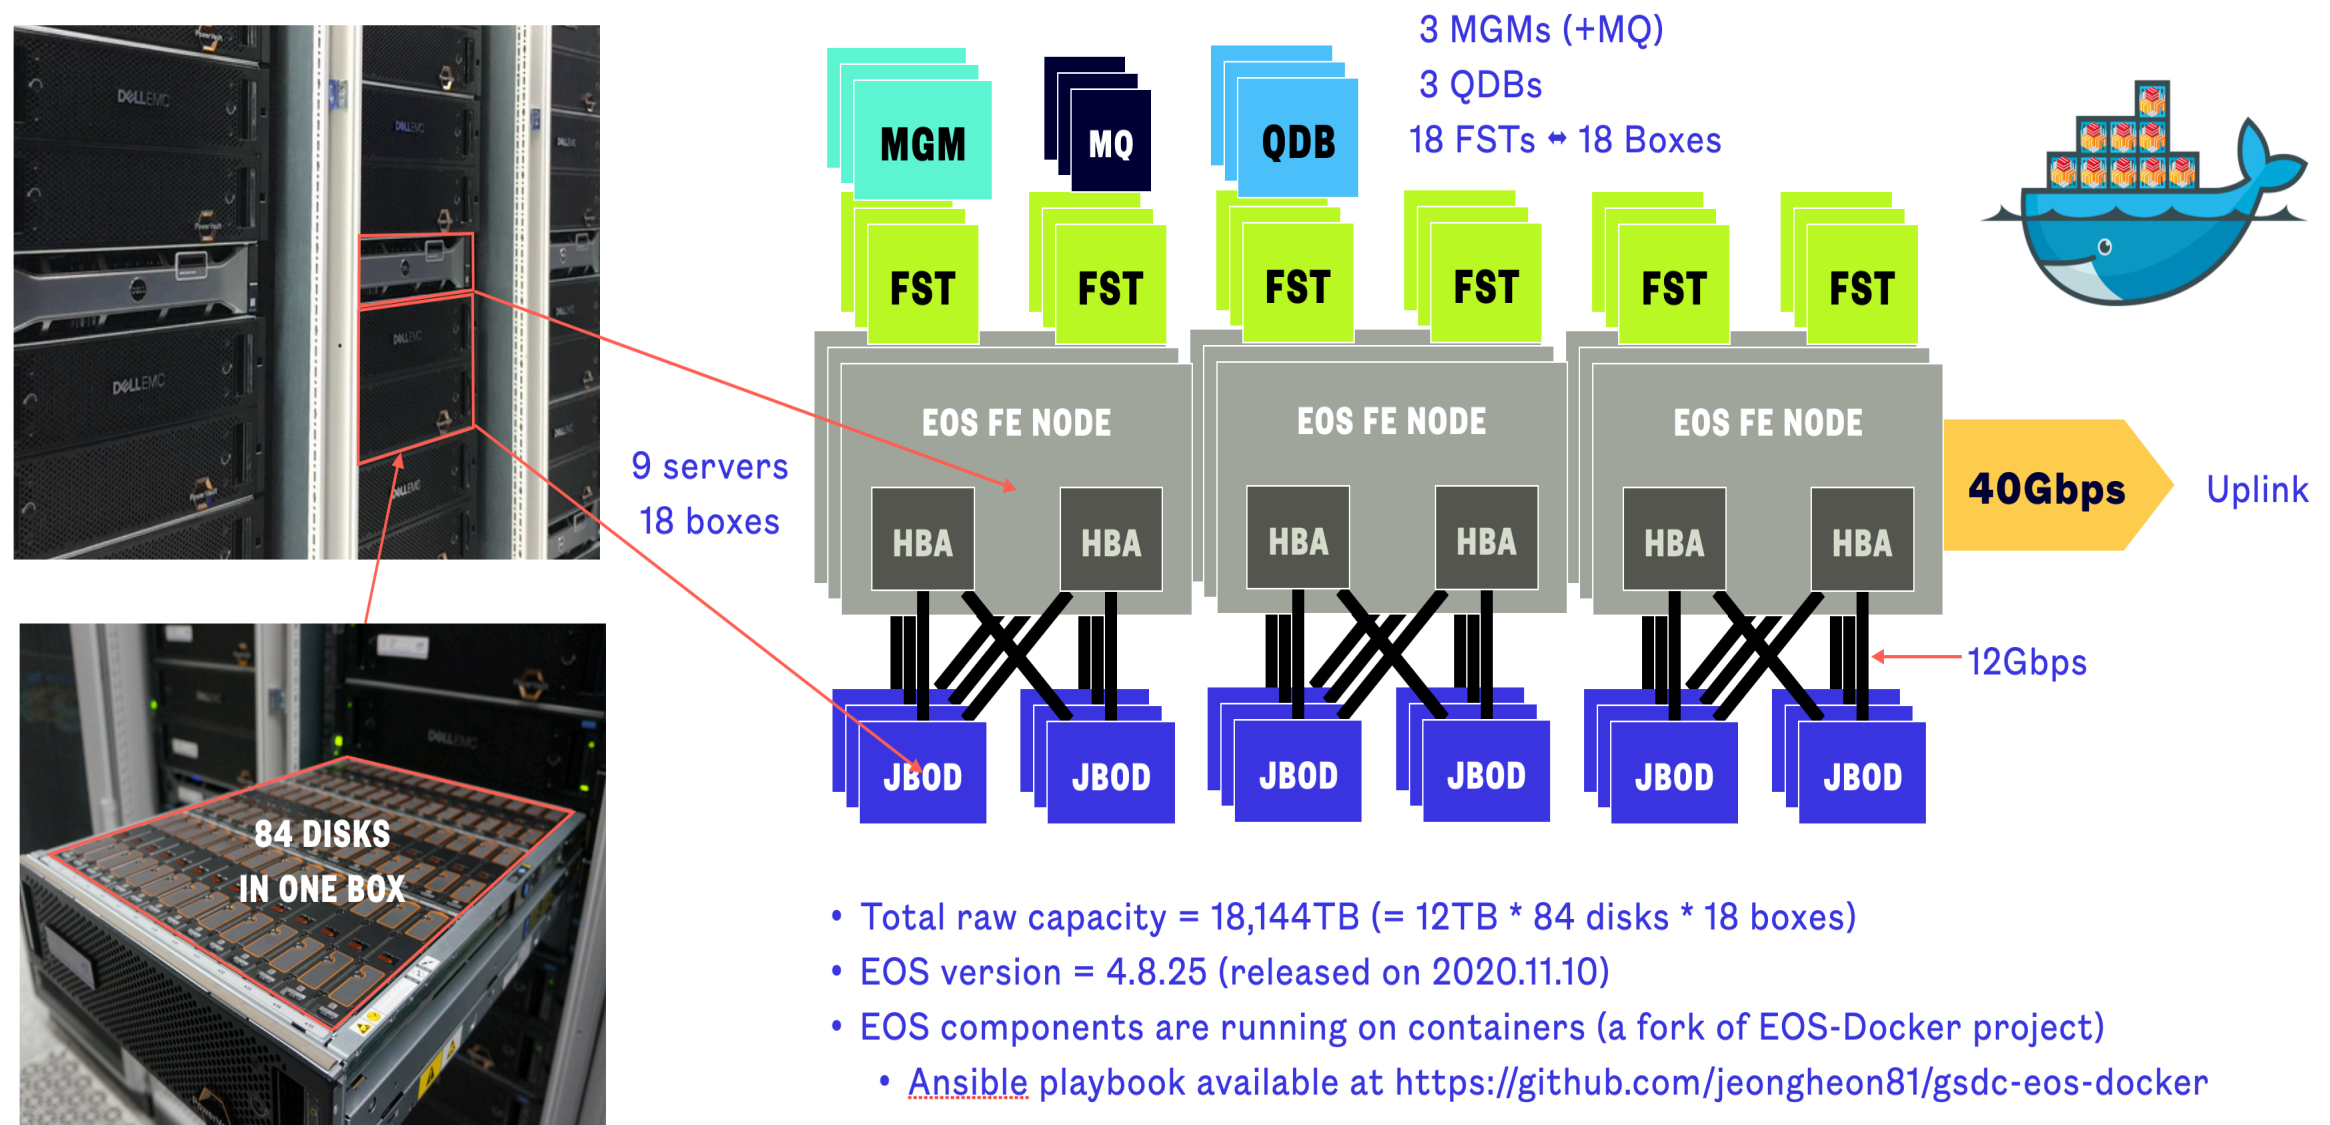
Name	Cluster	Volume Type	Bricks	Info	Space Used	Activities	No of snapshots
data	Default	Replicate	3 0		59%		0
engine	Default	Replicate	3 0		12%		0
ssd1	Default	Replicate	3 0		52%		0

Custodial Disk Storage (Tapeless Archiving)

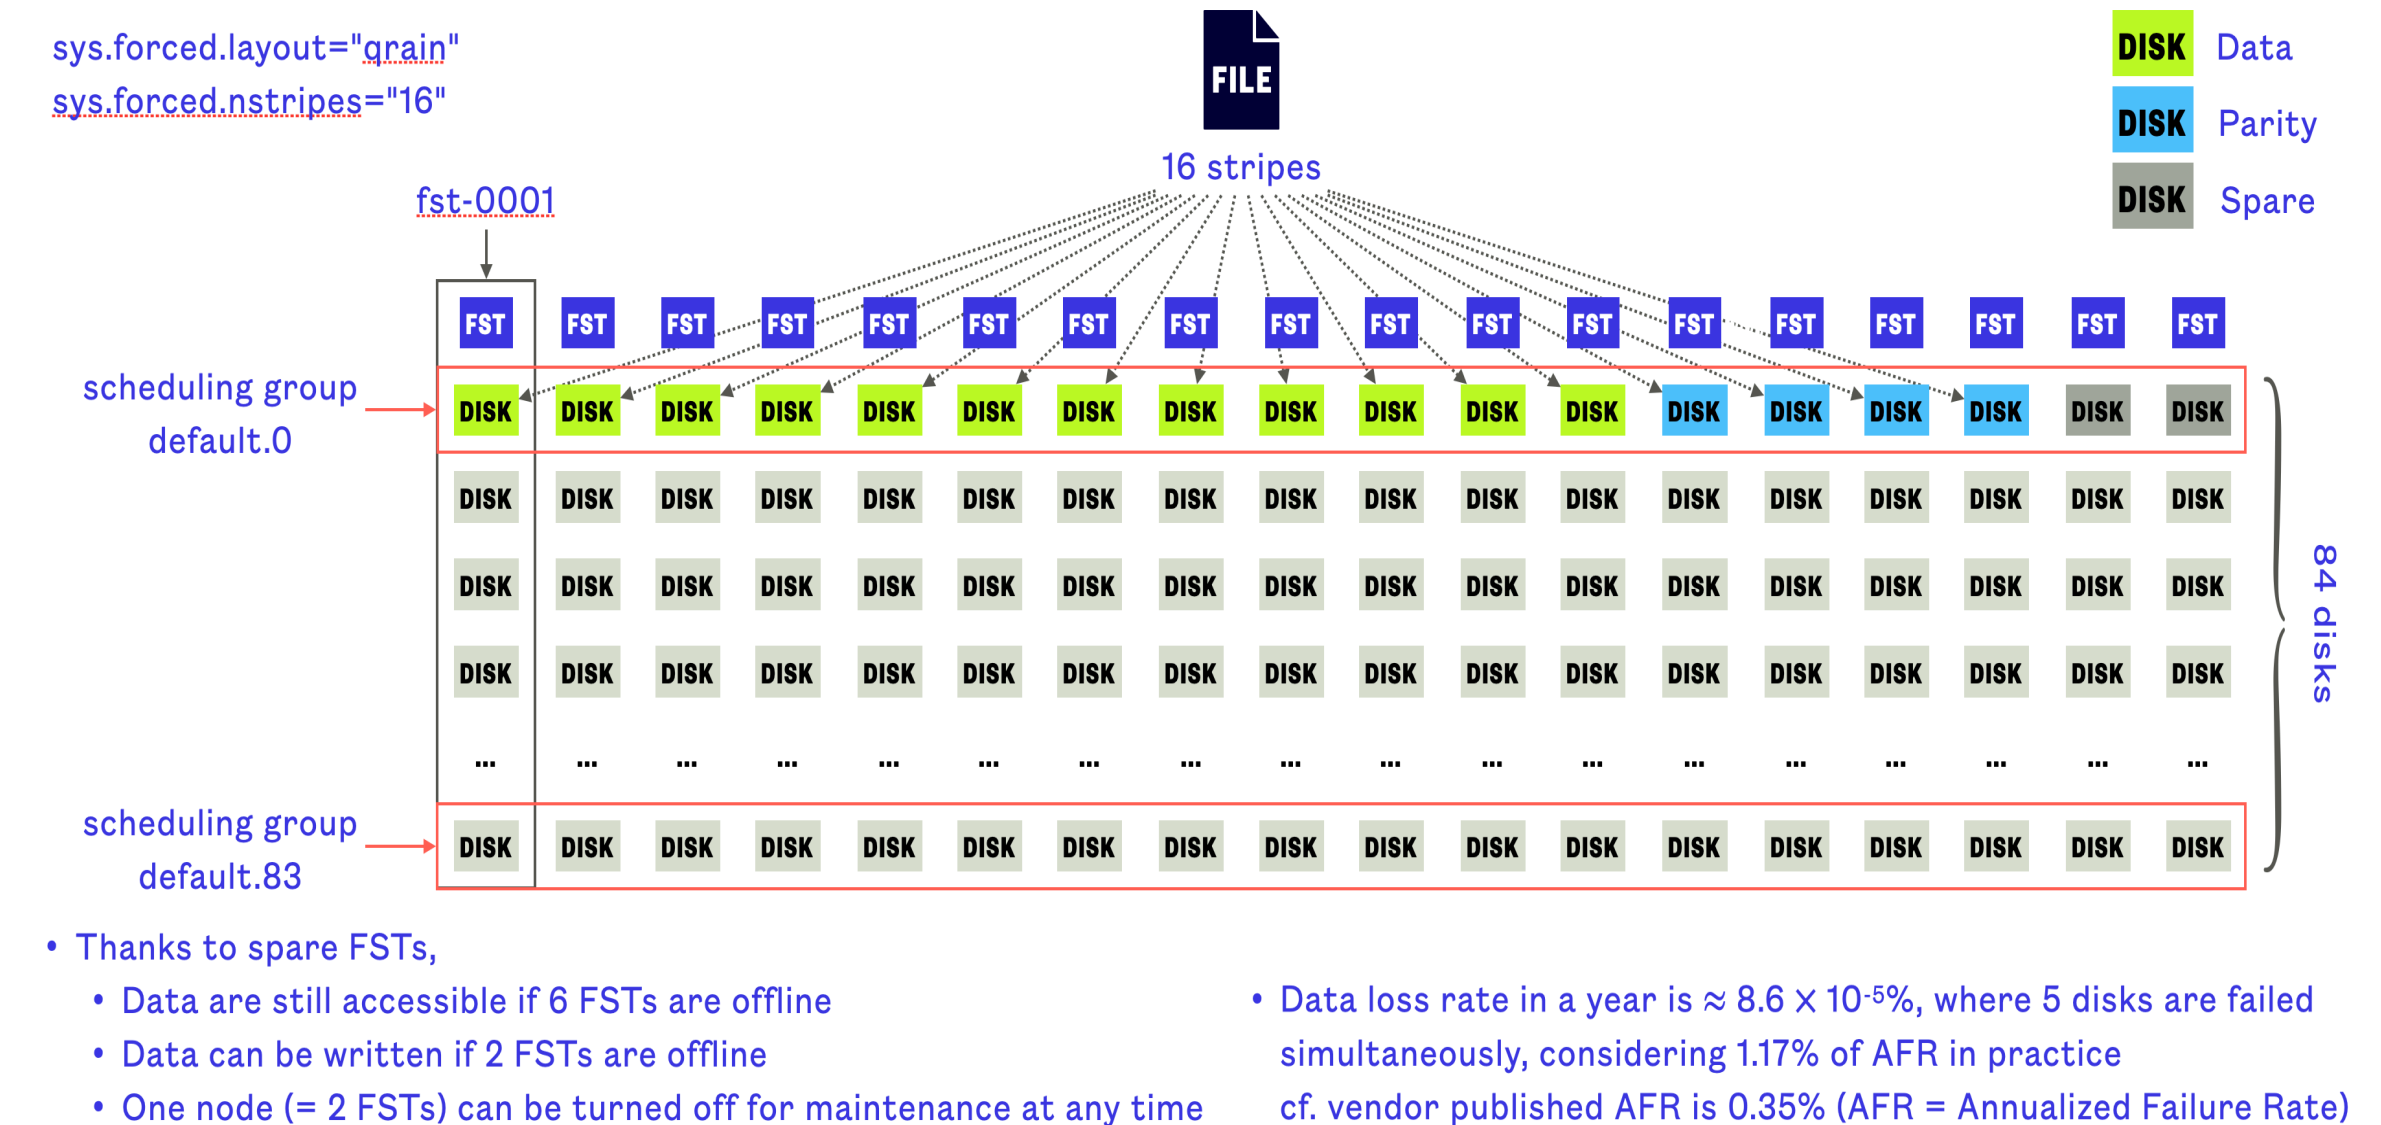
- The first disk-based custodial storage replaced tape for ALICE experiment
- 12 PB of usable space with 12+4 erasure coding for data protection (powered by EOS)
- Fully automated deployment of EOS components using Linux containers

Details will be presented in EOS session (Day 3)

System Architecture




QRAIN(12+4) Layout




ALICE Tier-1 Operations Summary 2022

2.65% Contribution to Total(T0+T1+T2+AF) ALICE Computing


Total wall clock hours for ALICE jobs (1yr)



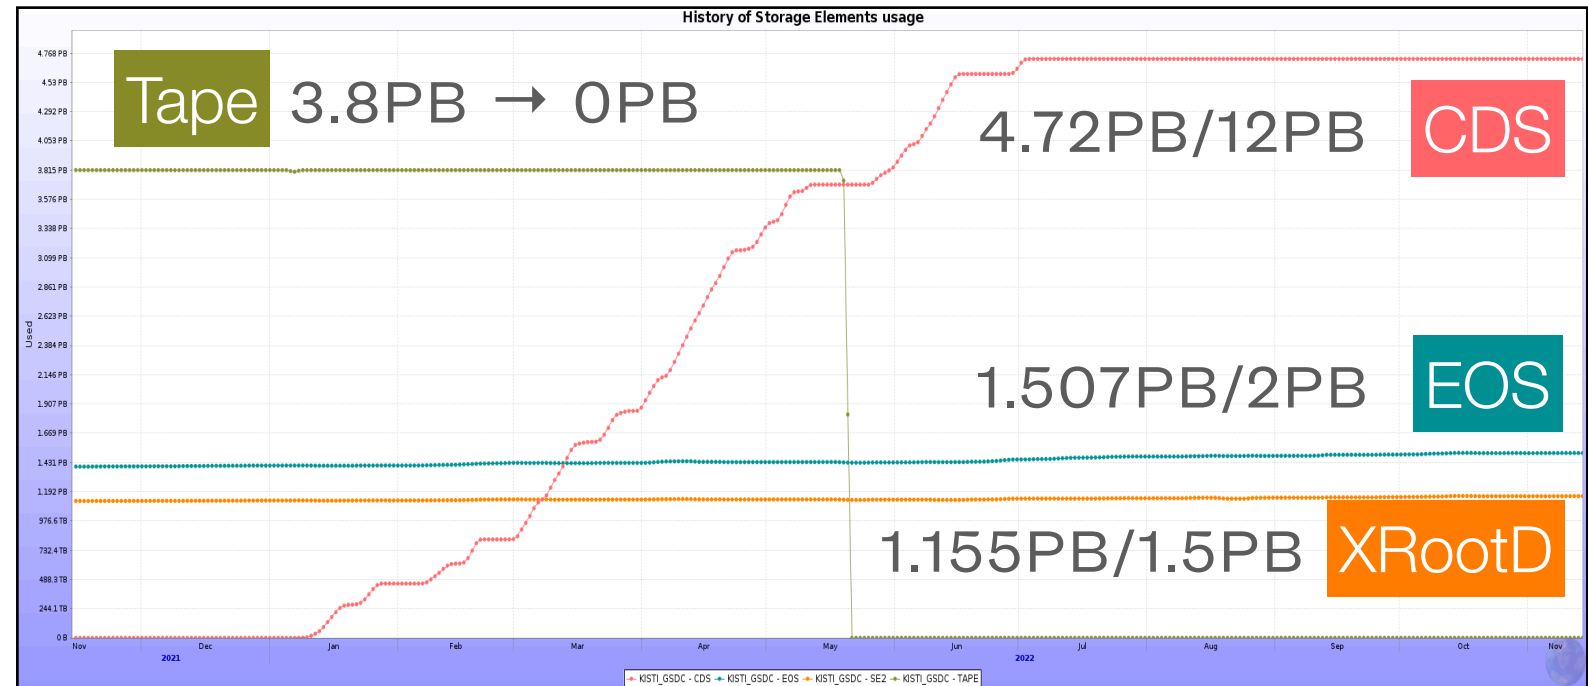
Running Jobs




Job Efficiency




SE Usage



Availability Reliability




LHCOPN 20G Upgrade (2021)



- Data transfer tests along with the OPN upgrade to 20Gbps in April
- Reached almost 20Gbps peak (> 1GB/s) during the tests

WLCG Tape Tests Challenge - 11 ~ 15 October 2021




- Joined efforts of WLCG collaboration for LHC RUN3 data taking
- Successful results to meet the target transfer performance (150MB/s)
- Stable and smooth operations of CDS during the challenges

KISTI Tier-1 Pledges for 2023-2024



10% Contribution to ALICE Tier-1 Computing Requirements

KISTI Tier-1 Pledges (Planned) for ALICE experiment



Pledges for 2023 = 60.0kHS06 / 6,500 / 12,000

(Original plan = 57.2kHS06 / 6,300 / 12,000)

** Additional pledges to contribute the mitigation of Ukraine war*

(CPU) New 30 workers (4,800 threads) to be deployed (Exact HS06 scores will be measured at late 2022) ; in addition to current installed CPUs (3,880 threads, which will be phased-out as their warranty ended, fully replaced by 2025)

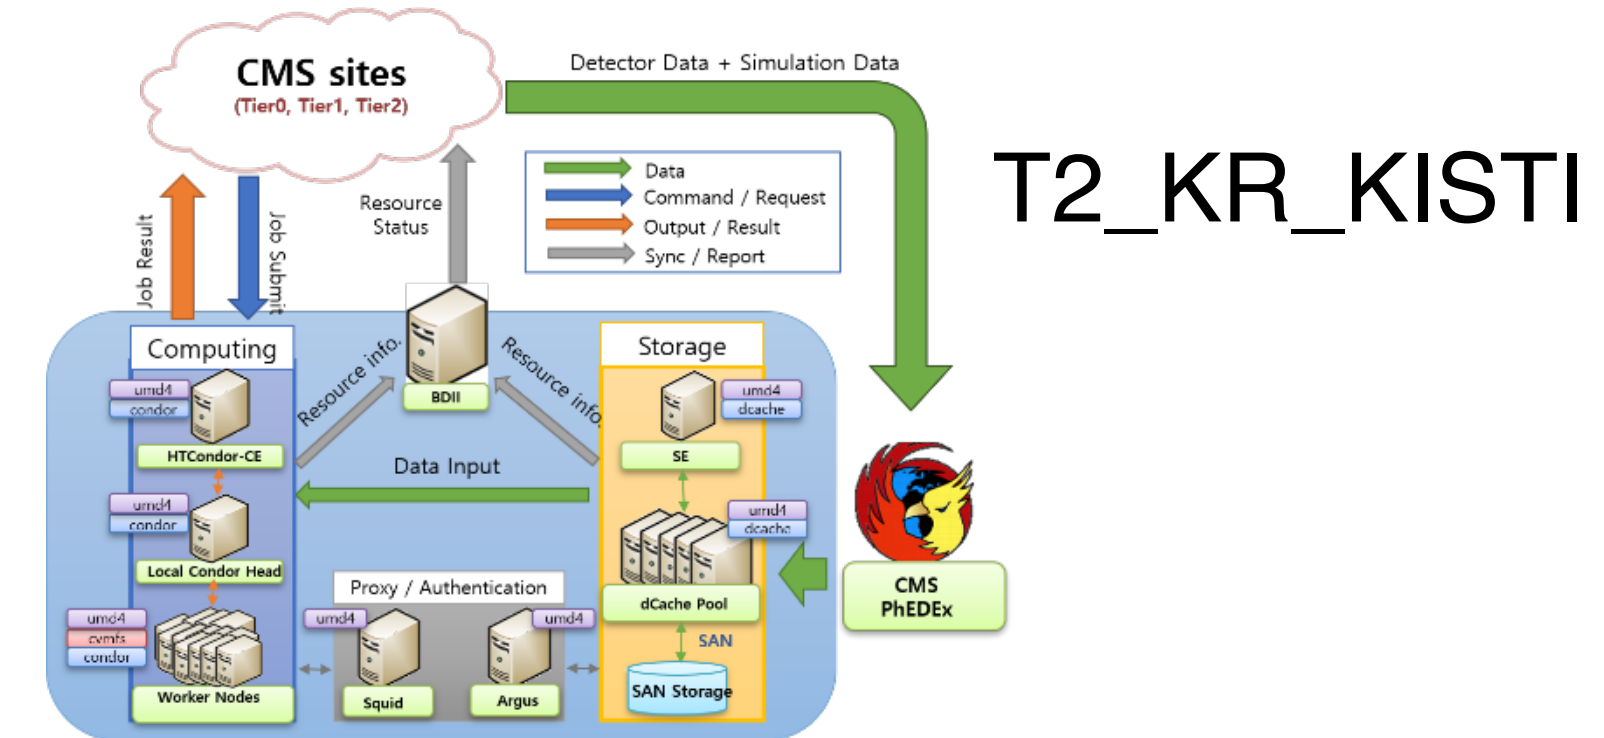
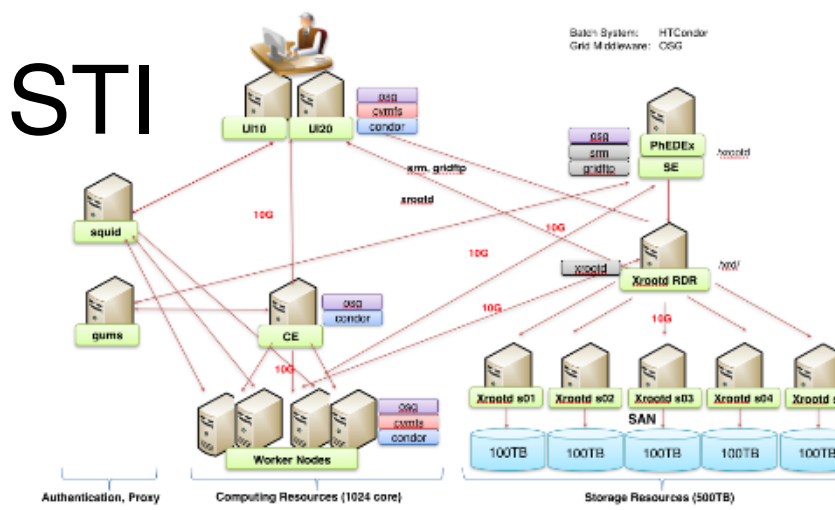
(Disk) 2.5 ~ 3.5 PB of Disks will be added in late 2022 or early 2023

(Tape) Plan to expand the CDS in 2025 (tentatively) (or a small amount (1.5 ~ 2 PB) of expansion possibly foreseen in 2023 depending on remaining budget)

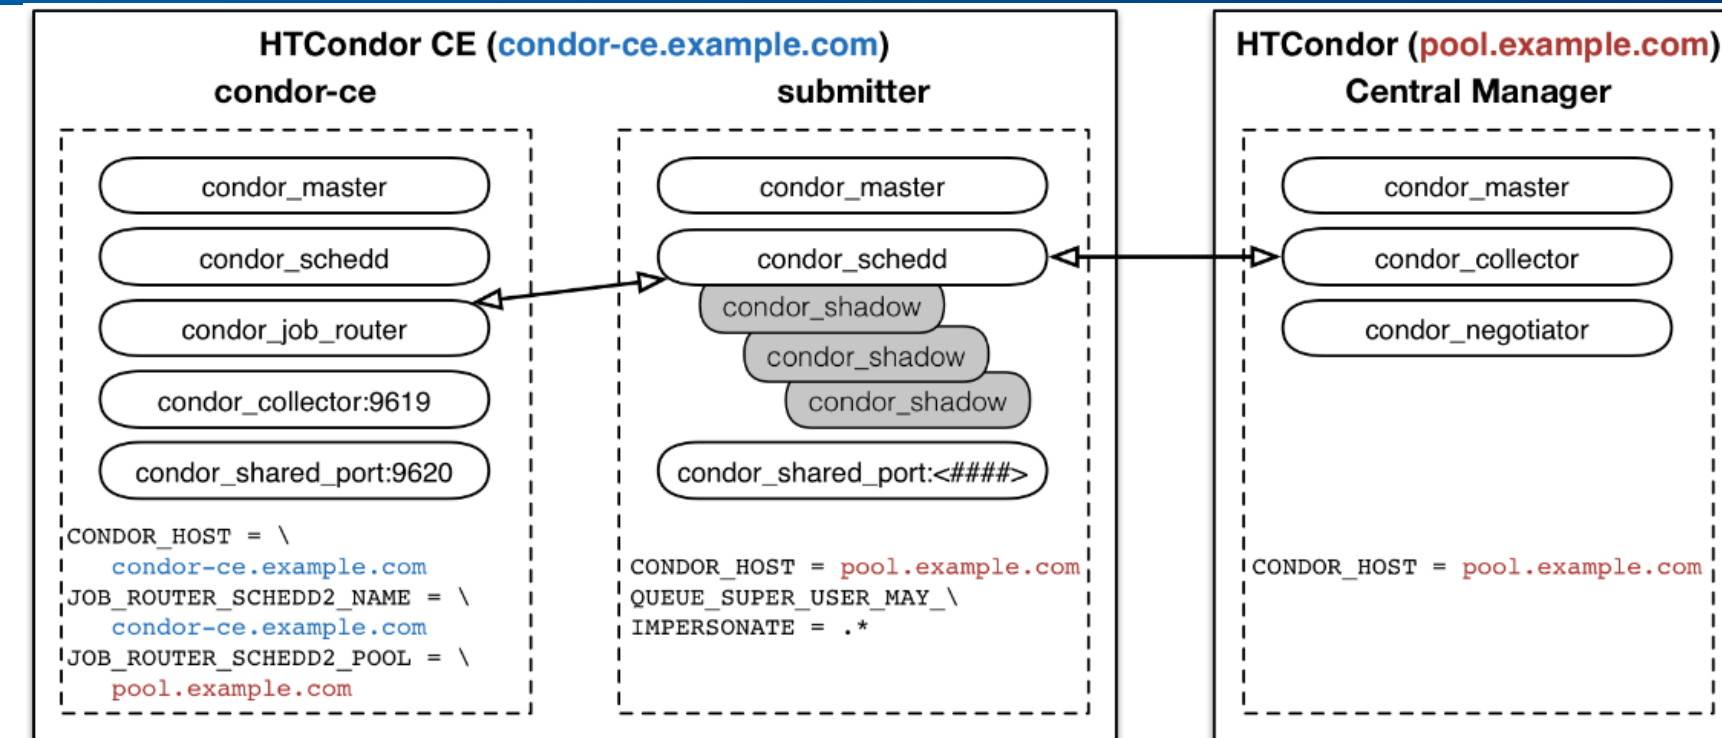
Installed (Planned)	2020	2021	2022	2023	2024
CPU (cores)	3,880	3,880	3,880	8,680	8,680
DISK (TB)	4,000	4,500	4,500	6,500	7,100
TAPE (TB)	5,500	12,000	12,000	12,000	12,000

- KISTI CMS Tier-2
 - WLCG Tier-2 site for CMS experiment
 - KISTI CMS Tier-2 focuses on providing resources for CMS experiment rather than supporting domestic users
 - Due to the presence of separate CMS Tier-3 site (T3_KR_KISTI)
- CMS Tier-2 History
 - 2017 Mar. : Register as an EGI site (KR-KISTI-GSDC-02)
 - 2017 Aug. : Register as a CMS Site (T2_KR_KISTI)
 - 2017 Sep. : Enable CMS PhEDEx Link (Joining CMS Data Transfer system)
 - 2017 Nov. : Starting CMS T2 Testbed after passing the SAM test stably
 - 2018 Apr. : KISTI-CERN MOU Signing Ceremony for CMS Tier2

T3_KR_KISTI




- Main Component
 - CE : HTCCondor-CE 5
 - LRMS : HTCCondor 9
 - 1,424 logical cores
 - RAM 3,000MB per core
 - SE : dCache
 - 1 SAN Pool + 9 NFS Pools / 1600TB
 - Protocol
 - XRootD, GridFTP(+SRM), pNFS, WebDAV
 - Etc.
 - Report: Site-BDII, APEL
 - Cache : Frontier-Squid
 - CMS AAA
 - 1x Standalone XRootD Server (Forward 1095 ->1094)



Pool Request Queues

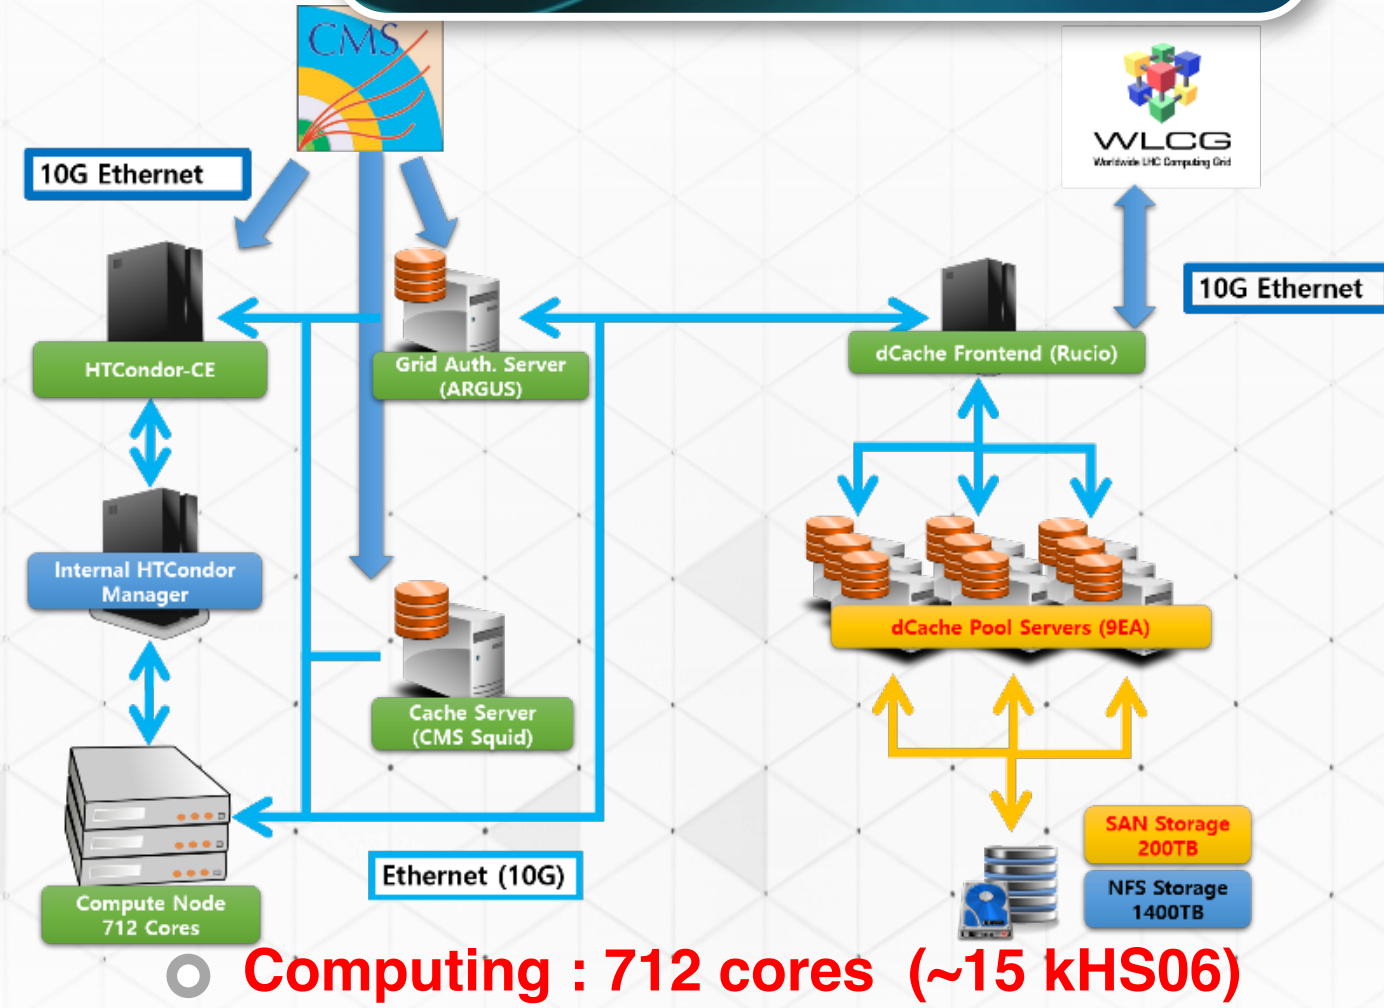
CellName	DomainName	Movers			Restores			Stores			P2P-Server			P2P-Client			queue_ftp			queue_webdav			regular		
		Active	Max	Queued	Active	Max	Queued	Active	Max	Queued	Active	Max	Queued	Active	Max	Queued	Active	Max	Queued	Active	Max	Queued			
Total		159	11340	0	0	0	0	0	0	110	0	0	0	0	0	4	220	0	4	1020	0	151	10100	0	
SAMPool	dCacheDomain	0	140	0	0	0	0	0	0	10	0	0	0	0	0	0	20	0	0	20	0	0	100	0	
cms-12-wn1055-NFSPool	cms-12-wn1055-NFSPool-Domain	16	1120	0	0	0	0	0	0	10	0	0	0	0	3	20	0	0	100	0	13	1000	0		
cms-12-wn1055-SANPool	cms-12-wn1055-SANPool-Domain	13	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	13	1000	0		
cms-12-wn1056-NFSPool	cms-12-wn1056-NFSPool-Domain	19	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	2	100	0	17	1000	0		
cms-12-wn1057-NFSPool	cms-12-wn1057-NFSPool-Domain	12	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	1	100	0	11	1000	0		
cms-12-wn1058-NFSPool	cms-12-wn1058-NFSPool-Domain	18	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	18	1000	0		
cms-12-wn1059-NFSPool	cms-12-wn1059-NFSPool-Domain	9	1120	0	0	0	0	0	0	10	0	0	0	0	1	20	0	1	100	0	7	1000	0		
cms-12-wn1060-NFSPool	cms-12-wn1060-NFSPool-Domain	10	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	10	1000	0		
cms-12-wn1061-NFSPool	cms-12-wn1061-NFSPool-Domain	21	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	21	1000	0		
cms-12-wn1062-NFSPool	cms-12-wn1062-NFSPool-Domain	30	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	30	1000	0		
cms-12-wn1063-NFSPool	cms-12-wn1063-NFSPool-Domain	11	1120	0	0	0	0	0	0	10	0	0	0	0	0	20	0	0	100	0	11	1000	0		
Total		159	11340	0	0	0	0	0	0	110	0	0	0	0	4	220	0	4	1020	0	151	10100	0		




= ok / enabled
 = warning / drain, test
 = error / disabled
 = unknown / not set
 = Waiting Room state
 = Morgue state
 = Scheduled Downtime
 = Partial Downtime
 = Ad Hoc Downtime

Korea, Republic of	KR-KISTI-GSDC-02	KR-KISTI-GSDC-02	4,714,905	17,270.71	4,714,905		
	129,000.00	3,913,000	Total	4,714,905	17,270.71	4,714,905	120.49%

CMS Tier-2 Infrastructure




Storage Usage



○Disk **1,600 TB** (Usage 85.71%)

Job Activities

~1.36 million jobs during this year



Data Transmission


○KISTI Tier-2 Data Link

- Tier-0 link : 1
- Tier-1 link : 7
- Tier-2 link : 46
- Tier-3 link : 5

Src_exp_site/Dst_exp_site	T2_KR_KISTI
TO_CH_CERN	61%
T1_DE_KIT	67%
T1_ES_PIC	72%
T1_FR_COIN2P3	72%
T1_IT_CNAF	72%
T1_RU_JINR	72%
T1_UK_BAL	58%
T1_US_FNL	83%
T2_AT_Vienna	62%
T2_BE_IJHE	62%

Data Traffic































Total : 681TB
Average : 68TB /month



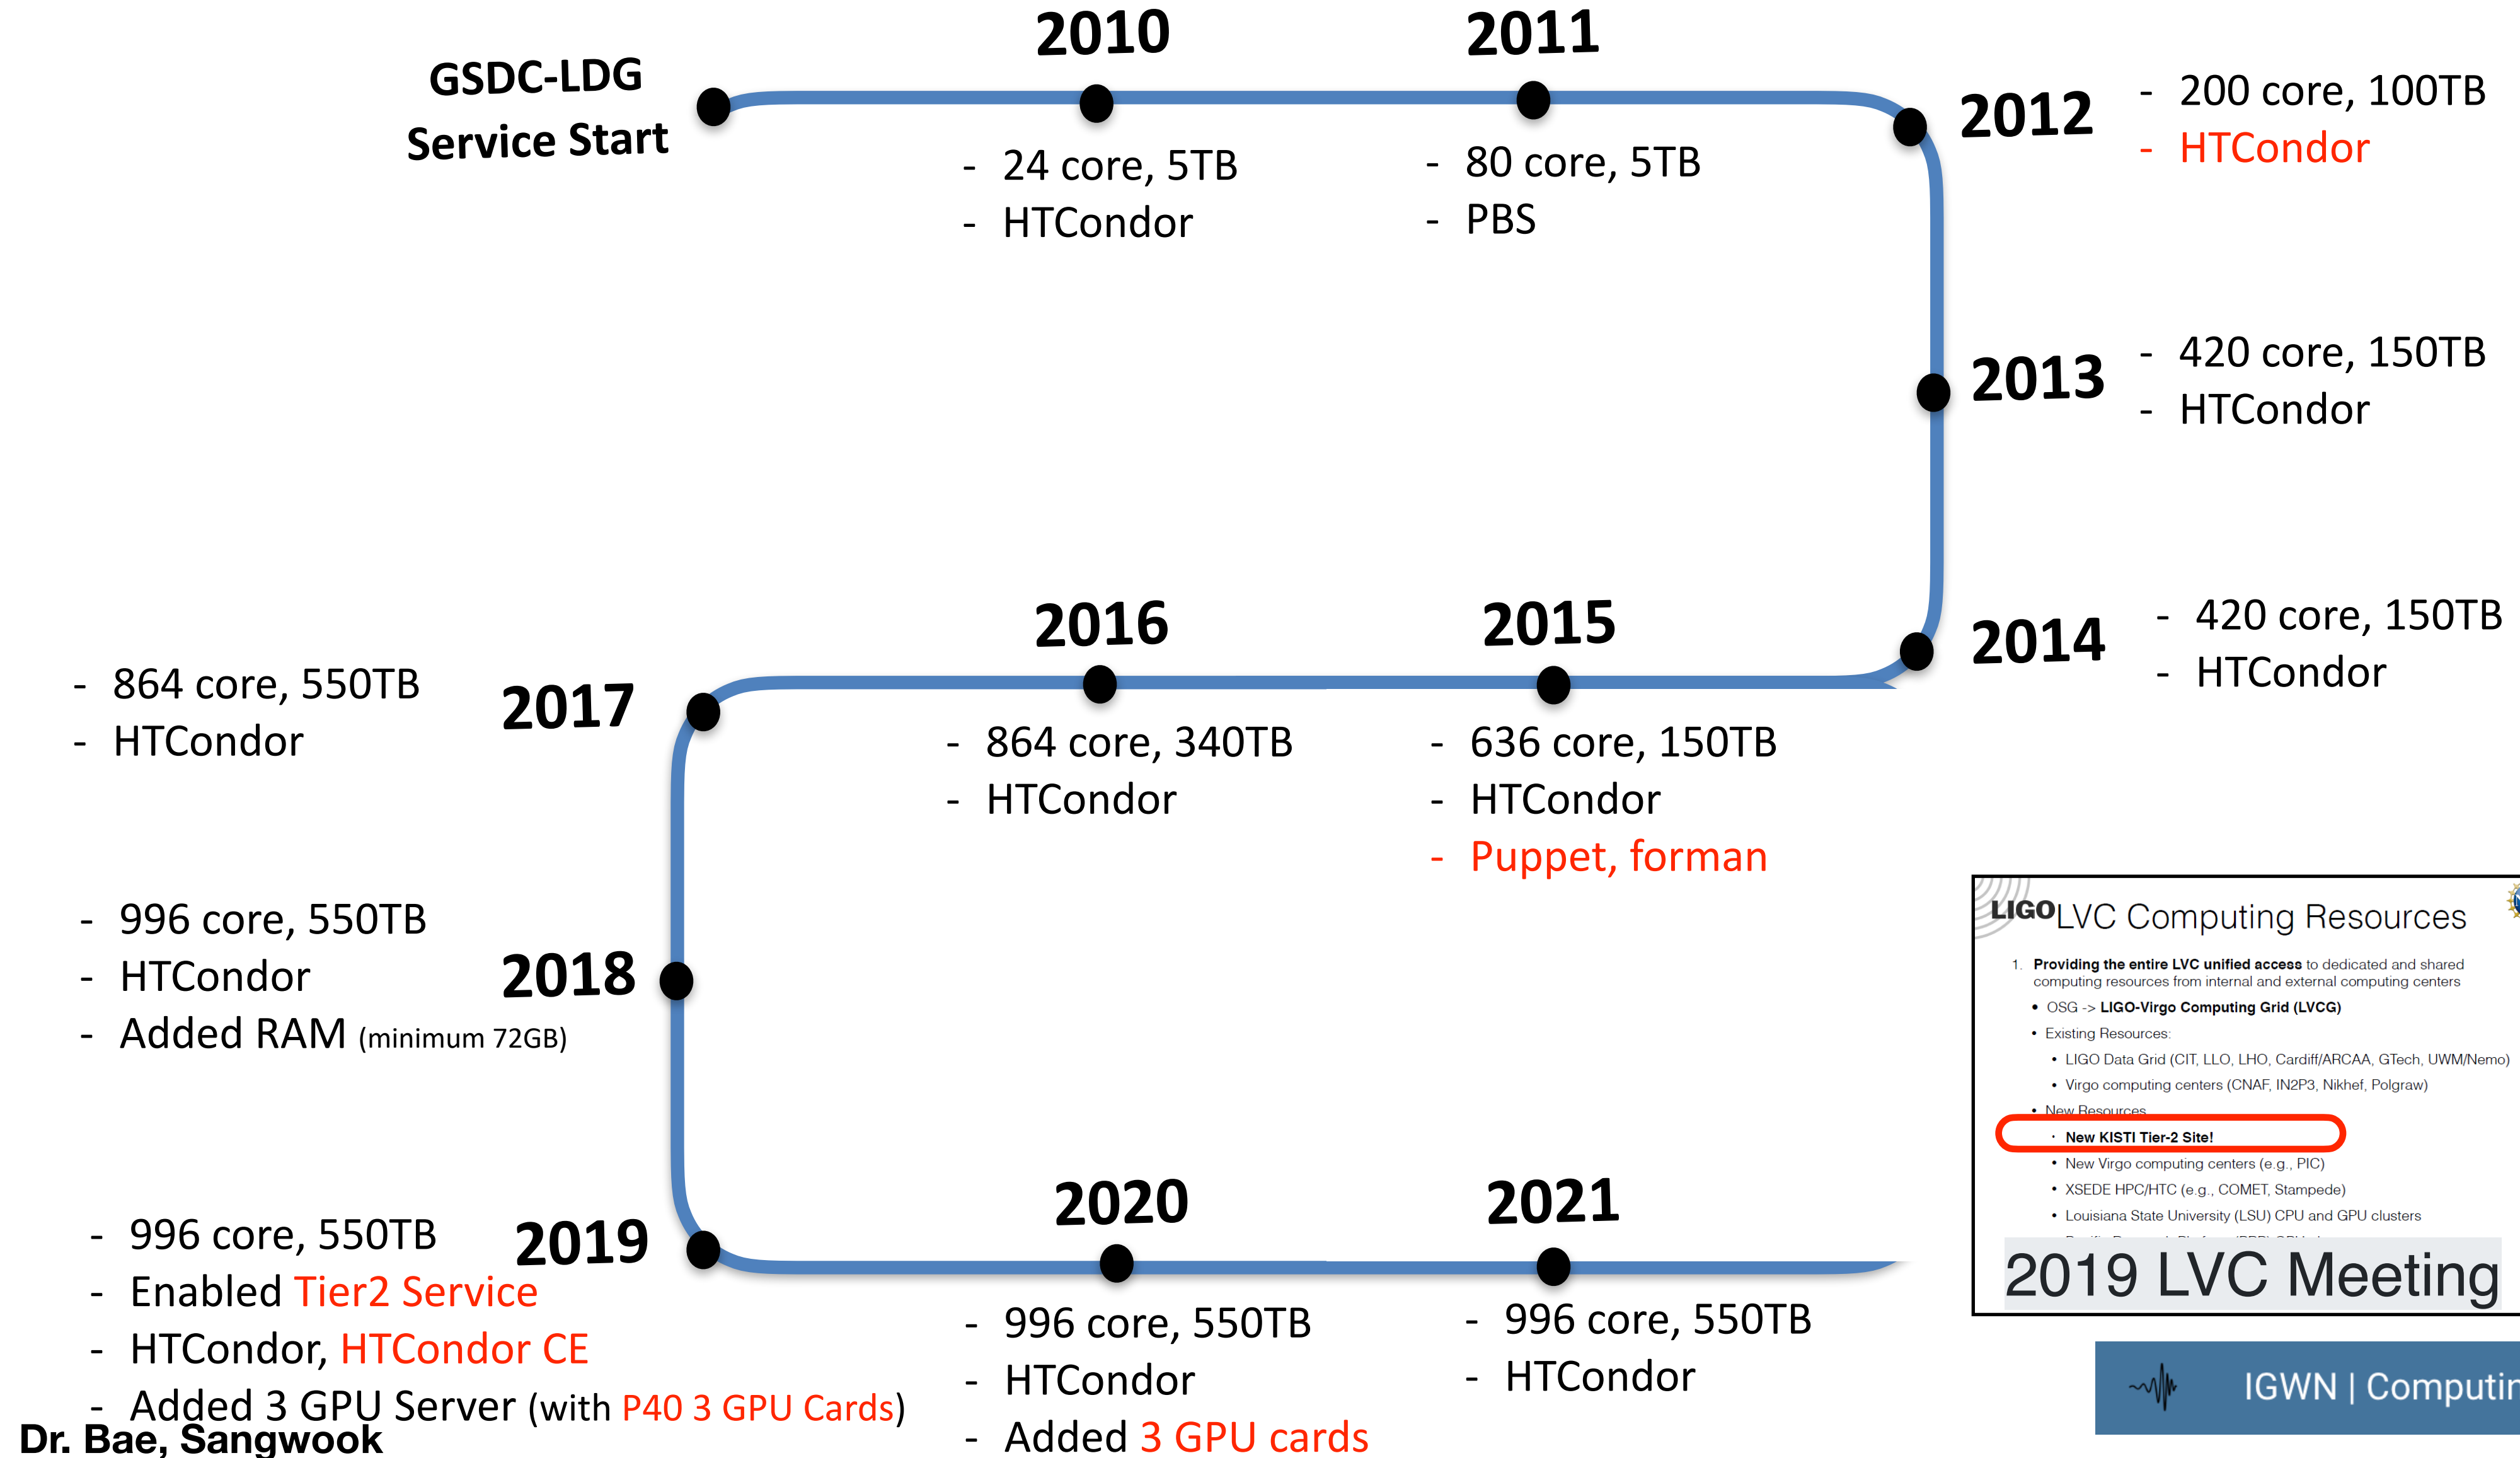
	Reliability	Availability
	Overall in 2022	Overall in 2022
CMS	97.68%	97.65%

 Monthly target of WLCG : 95%

CMS Tier-2 Availability/Reliability

Site	Availability	Reliability ↓
T2_FR_GRIF_IRFU	 94.20%	 99.11%
T2_RU_JINR	 98.74%	 98.74%
T2_DE_DESY	 98.69%	 98.69%
T2_HU_Budapest	 98.39%	 98.56%
T2_IT_Legnaro	 98.08%	 98.53%
T2_DE_RWTH	 97.89%	 98.41%
T2_UK_London_IC	 98.40%	 98.40%
T2_FI_HIP	 98.33%	 98.38%
T2_US_Wisconsin	 98.14%	 98.14%
T2_KR_KISTI	 97.65%	 97.68%
T2_US_Caltech	 97.42%	 97.67%
T2_CH_CERN	 97.59%	 97.59%
T2_PT_NCG_Lisbon	 96.86%	 97.42%
T2_FR_GRIF_LLR	 97.40%	 97.42%
T2_UK_London_Brunel	 97.20%	 97.23%

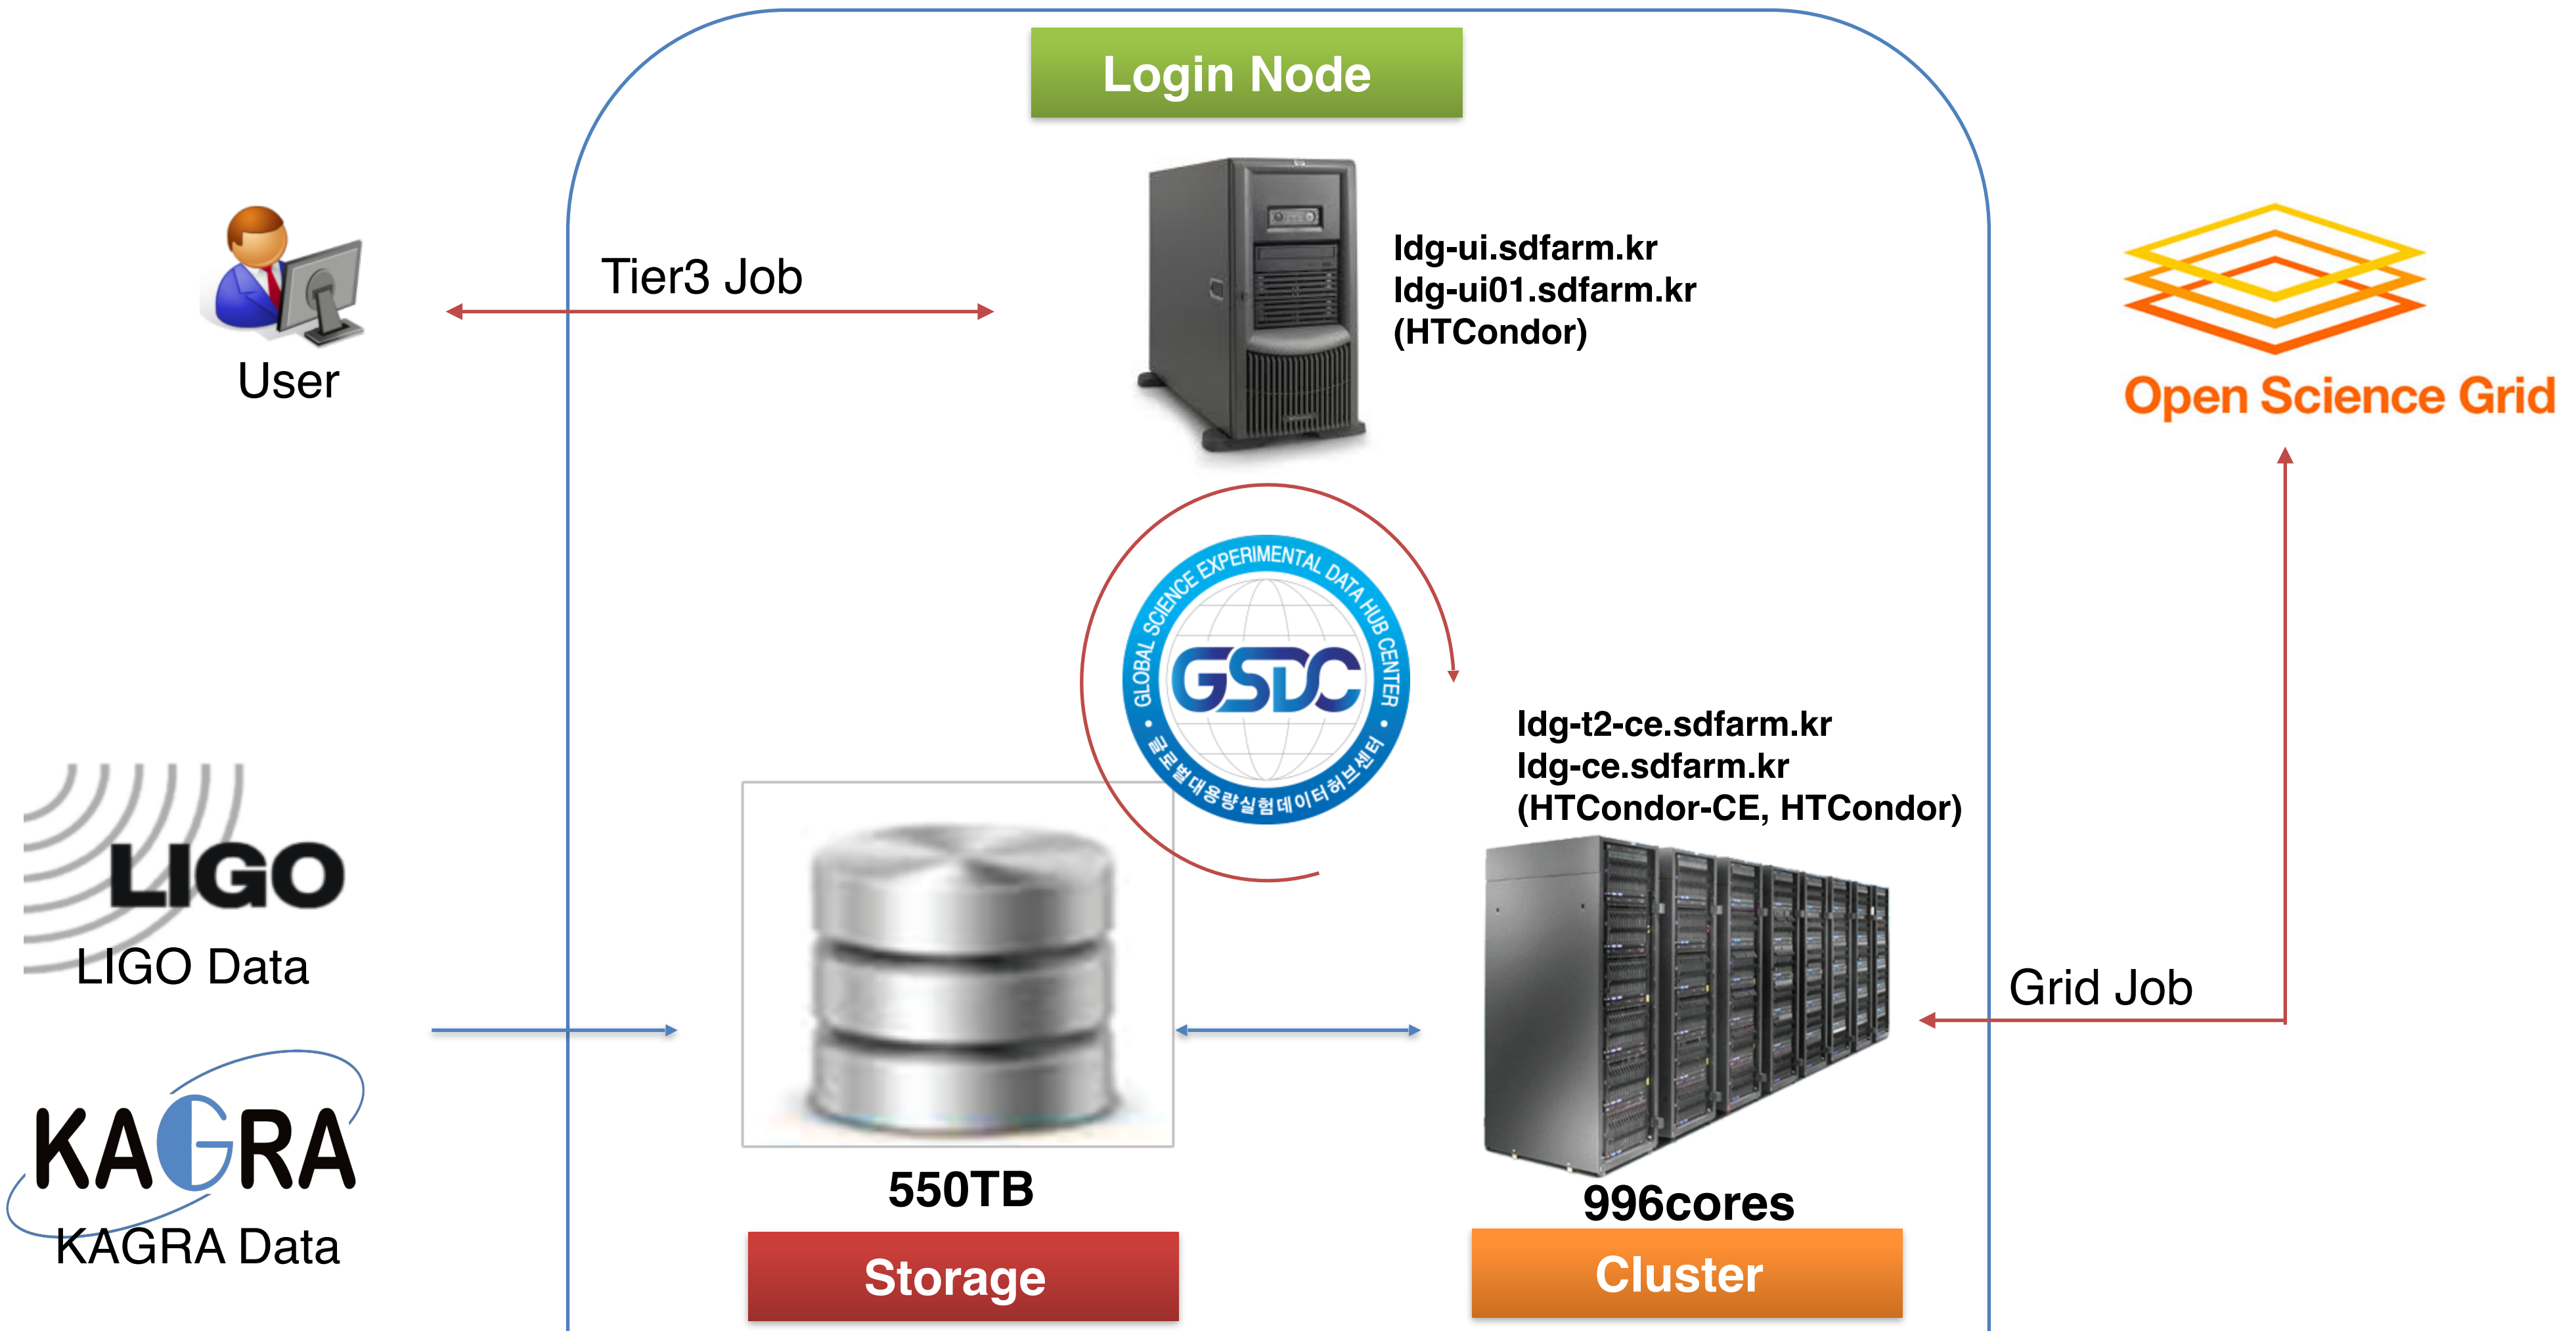
- GSDC-LDG (LIGO Data Grid), a gravitational wave data analysis computing environment at the request of the Korea Gravitational Wave Research Foundation (KGWG) in 2010.
- In 2019, the International Gravitational-Wave Observatory Network (IGWN) computing environment was established.
- Currently, the GSDC-LDG system operates as an integrated system that can be used simultaneously by global and domestic users.

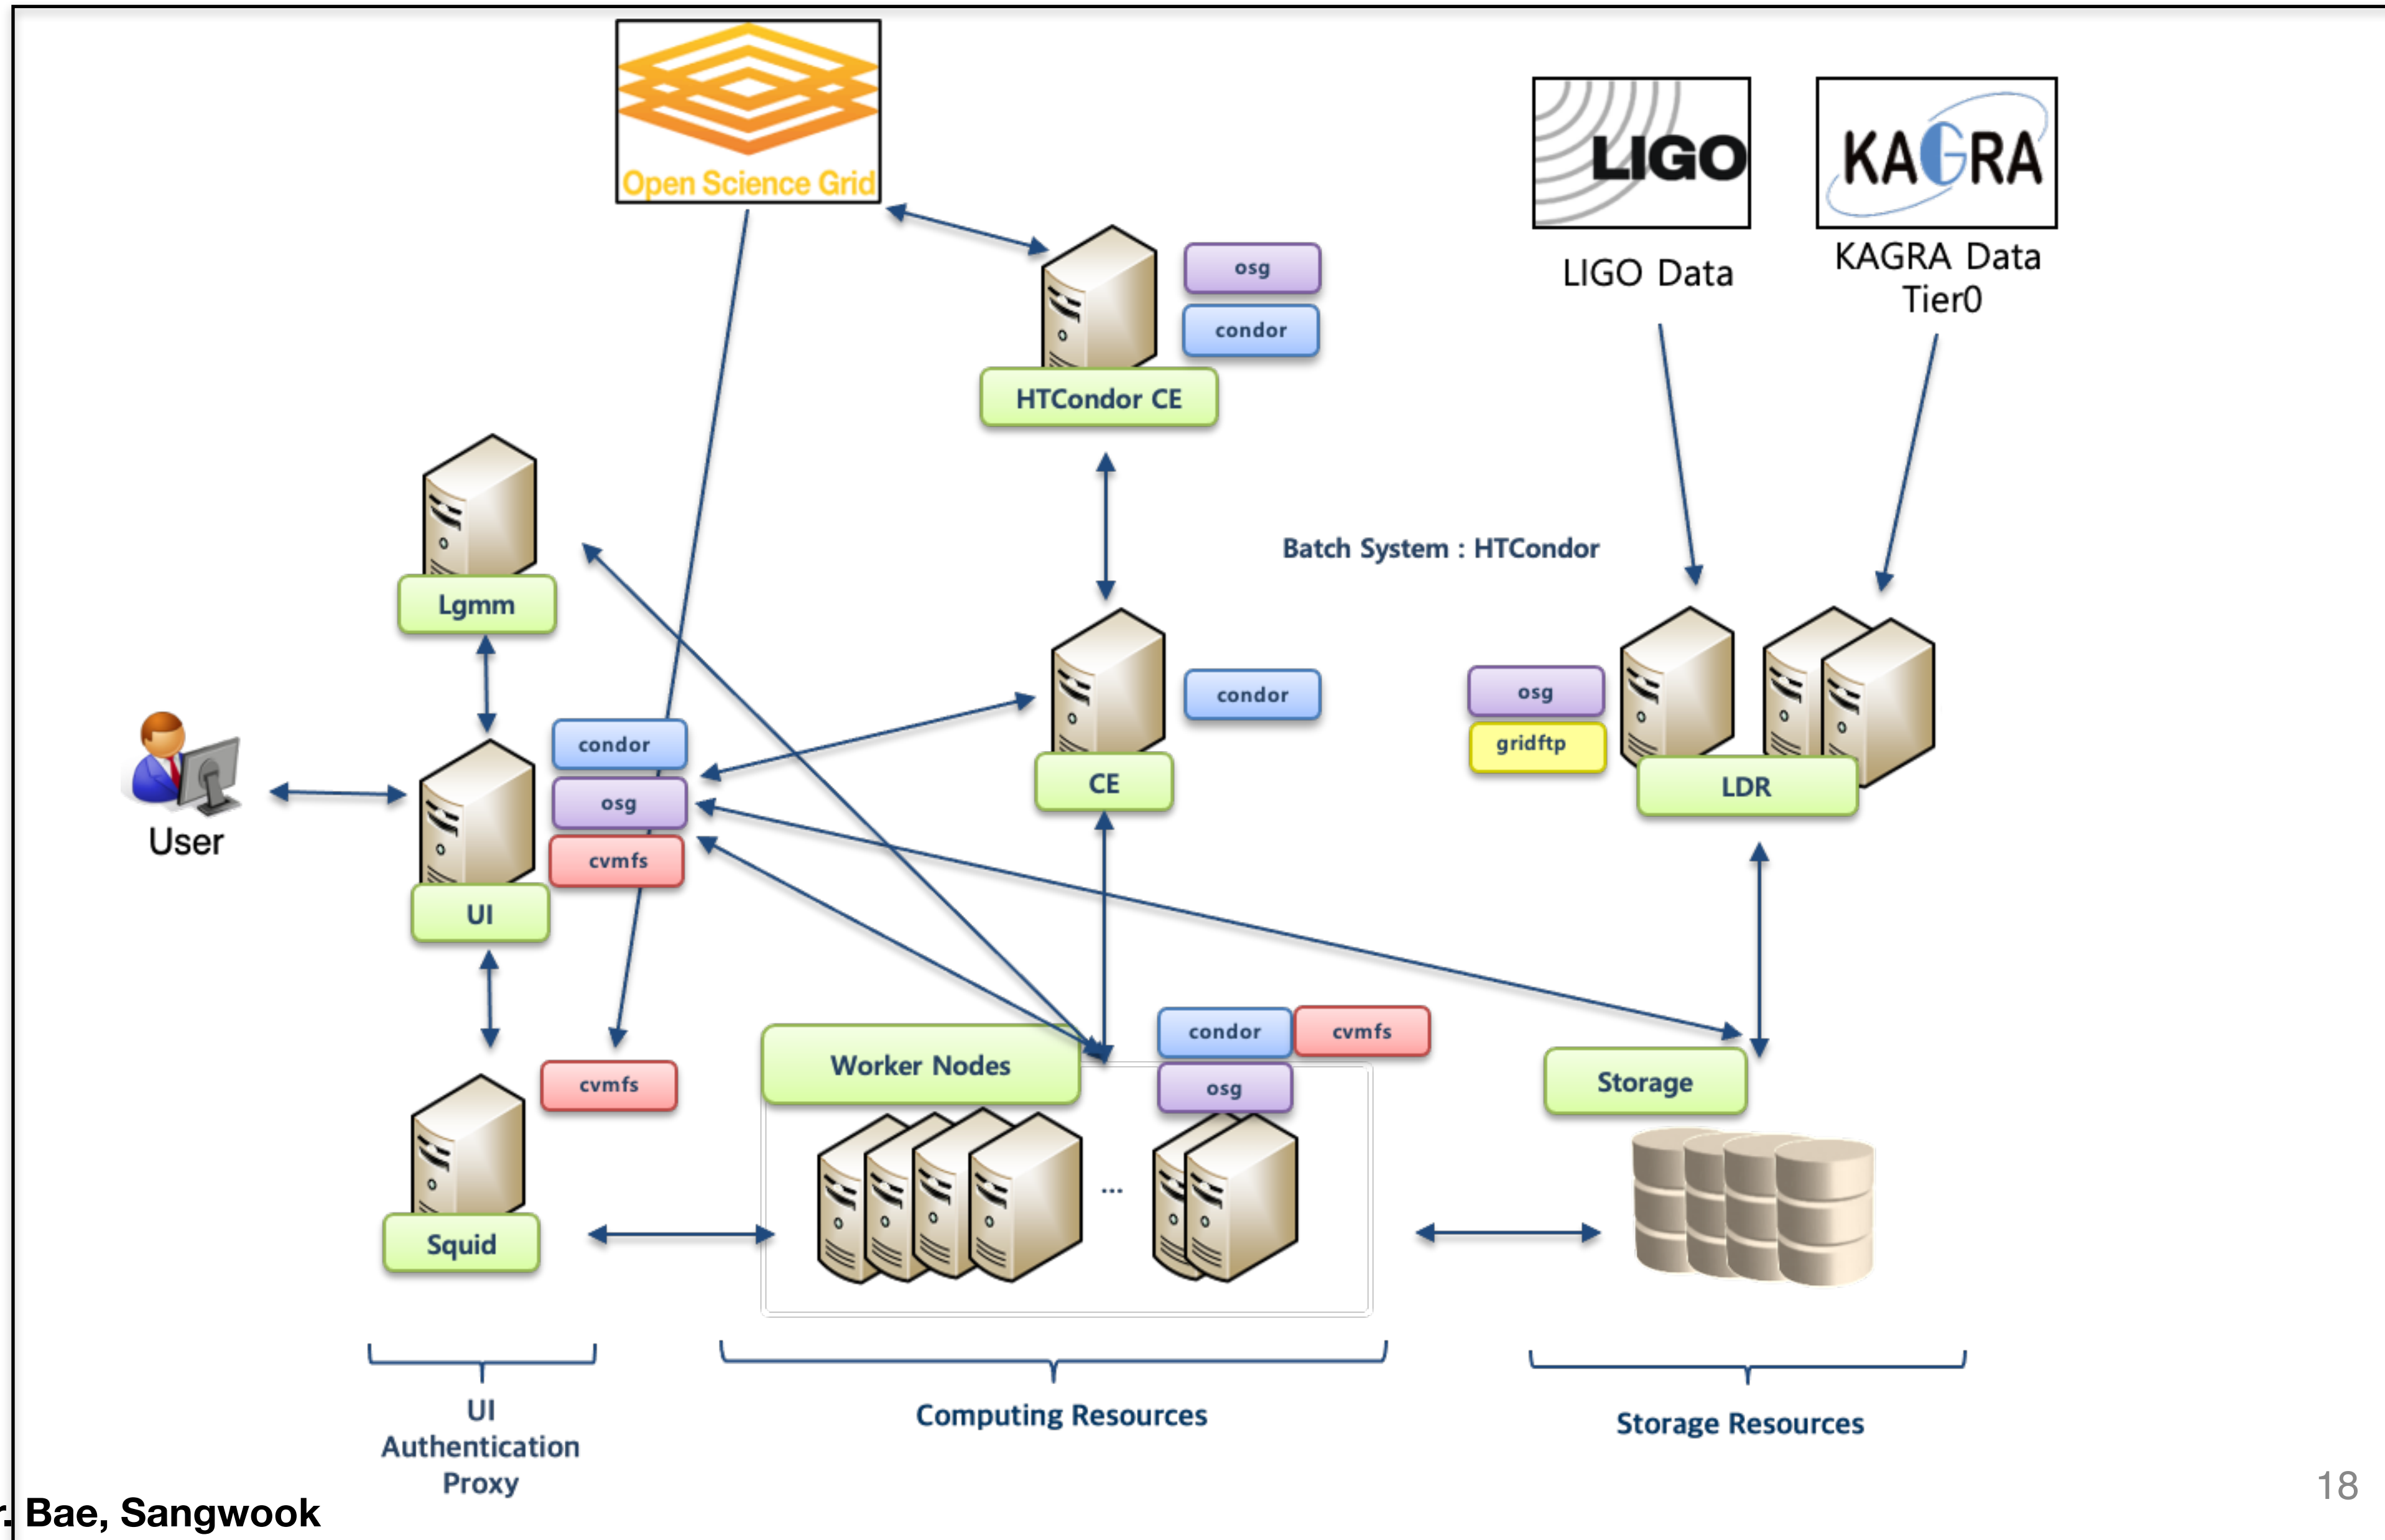


LIGO LVC Computing Resources

1. Providing the entire LVC unified access to dedicated and shared computing resources from internal and external computing centers
 - OSG -> LIGO-Virgo Computing Grid (LVCG)
 - Existing Resources:
 - LIGO Data Grid (CIT, LLO, LHO, Cardiff/ARCA, GTech, UWM/Nemo)
 - Virgo computing centers (CNAF, IN2P3, Nikhef, Polgraw)
 - New Resources:
 - **New KISTI Tier-2 Site!**
 - New Virgo computing centers (e.g., PIC)
 - XSEDE HPC/HTC (e.g., COMET, Stampede)
 - Louisiana State University (LSU) CPU and GPU clusters

2019 LVC Meeting





- Computation Resource

	Physical Core	Memory
Work Node	996 (66 servers)	72GB X 27 96 GB X 33 384 GB X 6
UI,CE,LGM,LDAS,LDR	60 (5 servers)	24GB X 5
Total	1056	7416



Work Node (GPU)	3 Servers	6 GPU Cards (P40)
--------------------	-----------	-------------------



- Storage Resources

	Mount on	Size	Used	Avail	Use	Total
LIGO	/data/ligo/	400T	337T	64T	85%	pool0.gsn.sdfarm.kr:/ifs/service/ligo
KAGRA	/data/kagra/	150T	102T	49T	53%	pool0.gsn.sdfarm.kr:/ifs/service/kagra

- LDG central monitoring system
- Ganglia installed on all GSDC-LDG resource



Dr. Bae, Sangwook



Summary

- KISTI-GSDC is a dedicated datacenter to promote fundamental research in South Korea by providing necessary computing power, storage and services
 - Tightly coupled with KISTI Supercomputer infrastructure and KREONet's global reachability
 - Currently supporting ALICE Tier-1, CMS Tier-2, LDG Tier-2 and related domestic communities
 - Expanding its contribution to global and domestic research communities

Detailed configuration and setup of HTCondor system will be presented in HTCondor session in Day 3