

Activity at KIAF

KoALICE National Workshop 2019 @ High1 Resort
2020-01-06

Byungyun KONG



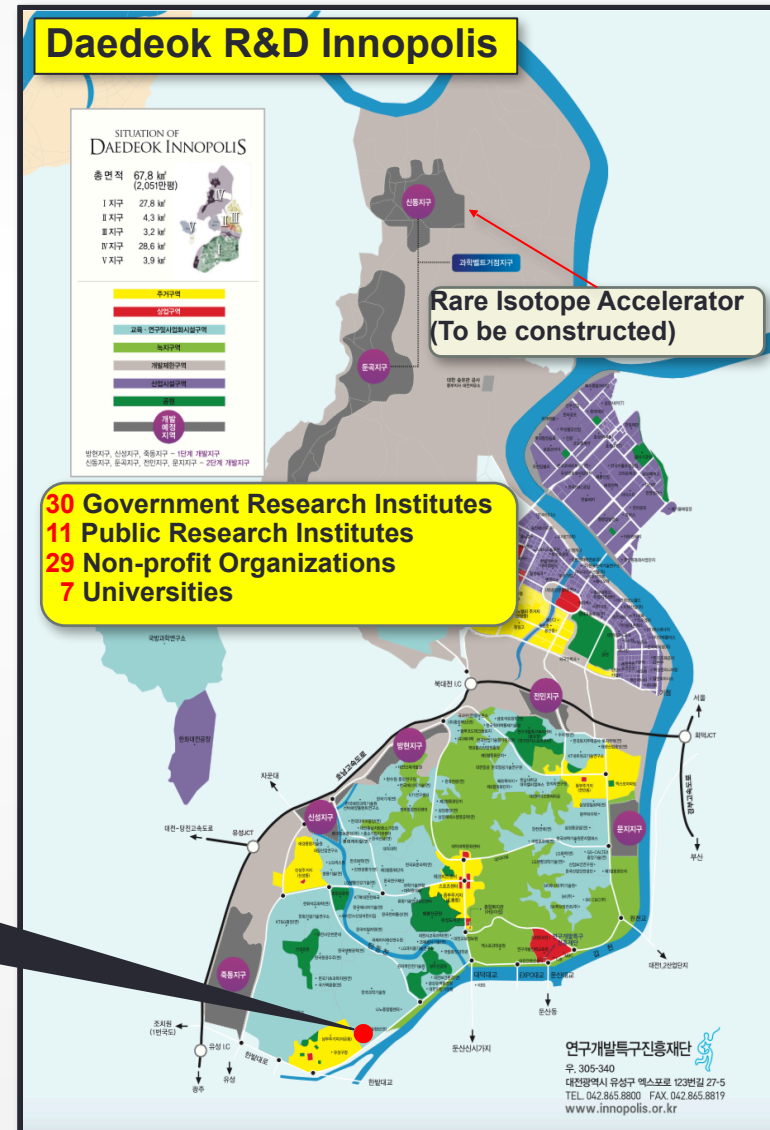
Content


- **GSDC Introduction**
- **ALICE Analysis Farm**
- **HTcondor Integrated Pool**
- **Operation Status**

GSDC Introduction

Korea Institute Science and Technology Information

- Government-funded research institute founded in 1962 for National Information Service and Supercomputing
- National Supercomputing Center
 - Nurion System (~25.7 PFlops at Peak), ranked 11th of Top 500(2018)
 - KREONet National R&E Network





Large-scale Scientific Data:
20Km CD stack with data
produced per year in CERN


Global Science

experimental Data

hub Center



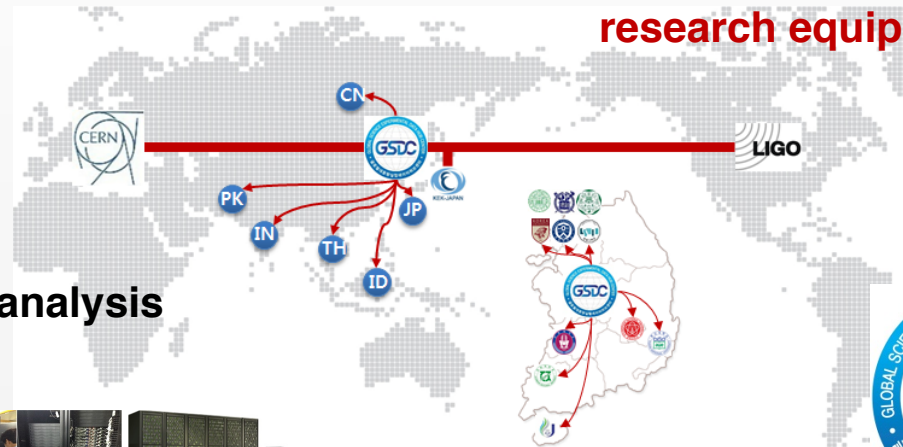
**Collaboration with
global laboratories**



**Data from
large and high-valued
research equipment**

➤ **(Global)**
Asia representative **Data Hub**

➤ **(Domestic)**
Scientific data management and analysis
platform service



Experimental Support



3-4 more domestic experiments under preparation
 e.g. volcanic hazard mitigation, brain research, disease control, etc.

ALICE Analysis Farm


KIAF System

- **Dedicated Computing farm for KoALICE**
But share worker node with KCMS
- **Basic configuration**
 - HTCondor batch system
 - XRootD data management system
(coupled with AliEn for ALICE dataset transfer)
 - Scientific Linux 7.5
- **Additional software**
 - XRootD Filesystem for local file access
 - PROOF on Demand enabled
 - Ganglia for monitoring (Admin only)

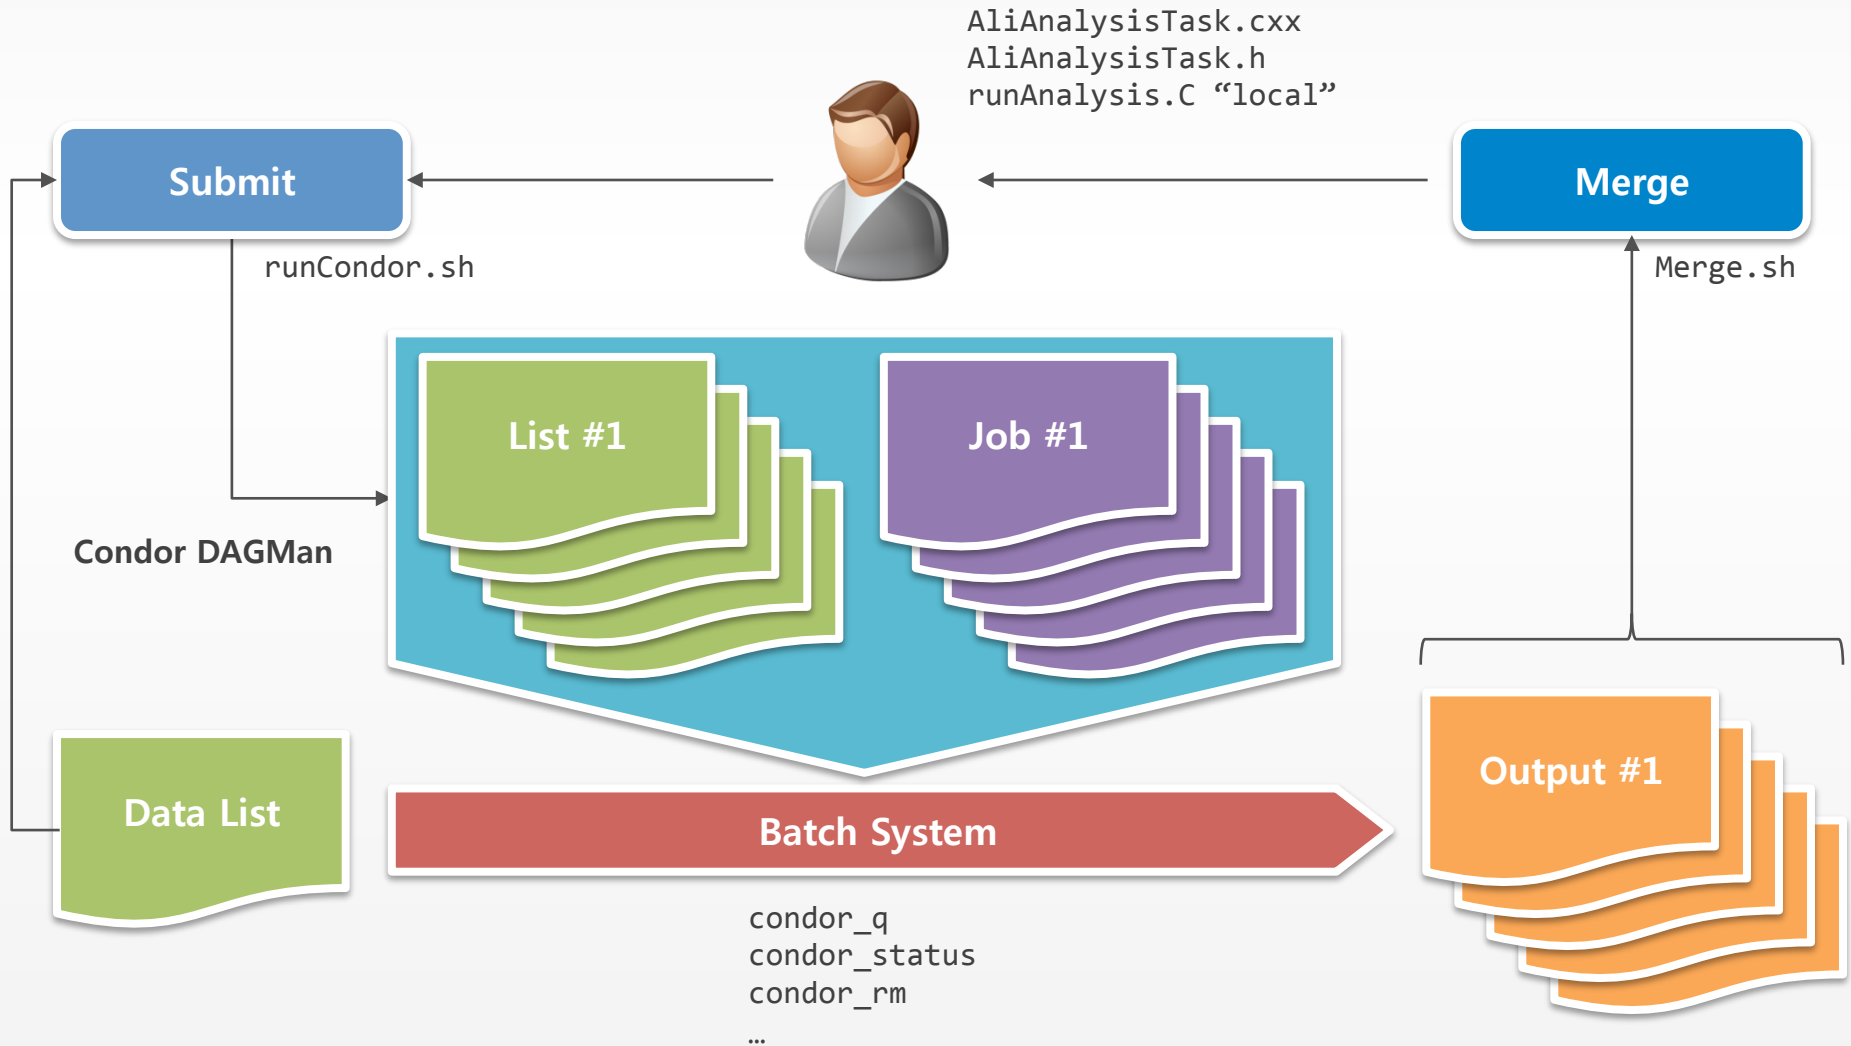
Hardware History

	2012	2013	2014	2015	2016	2017	2018	2019	2020
CPU	54	54	54	98	98	100	200	200	400
Disk (TB)	180	180	200	200	200	500	600 (700)	800	1,300
Network (Gbps)	1	2	2	10	10	10	10	10	10

KIAF Curves



User Analysis Workflow



User Analysis Workflow

AliAnalysisTask.cxx
AliAnalysisTask.h
runAnalysis.C "local"



Submit

Merge

runCondor.sh

Merge.sh

Condor D

Data

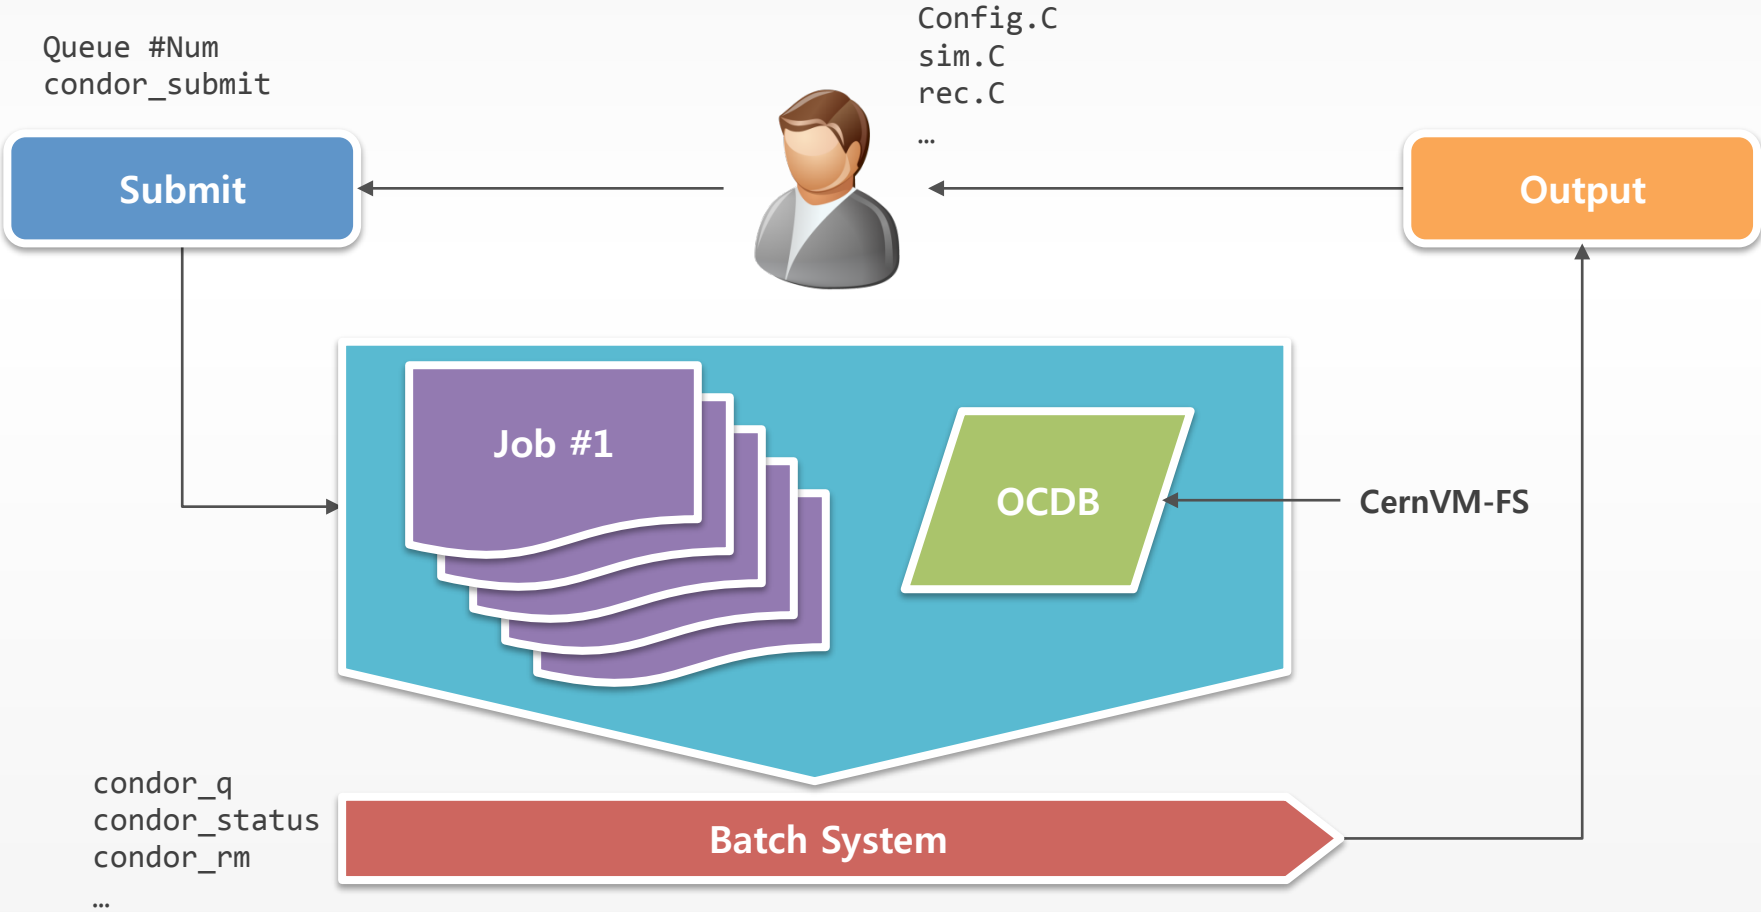
```
#!/bin/bash
condor_slot=40
nfile=2
merge=1 #0=nomerge, 1=merge
work_dir=`date +%Y%m%d%H%M%S`
run_macro="runAnalysis.C"
data_file="sim_LHC13d3_195675_A0D159.txt"
input_files="AddMyTask.C,AliAnalysisTaskMyTask.cxx,AliAnalysisTaskMyTask.h"
out_file="AnalysisResults.root"
USER_MAIL=""
main_dir=`pwd`
source /pool/macro/runScript.sh
```

condor_rm

...



MC Workflow




User Private Code

- Python
- Garfield++
- etc.

HTCondor Integrated Pool

Introduction of Integrated System

- **Temporary idle resource generation due to system crowded according to the research schedule in the community**
 - Example) Workshop, Conference, Presentation
- **As the resource demand of the community continues to grow, a method for efficiently utilizing limited resources is required.**



Solution :
Increase resource utilization efficiency by integrating resources between communities

Integration Target Selecting

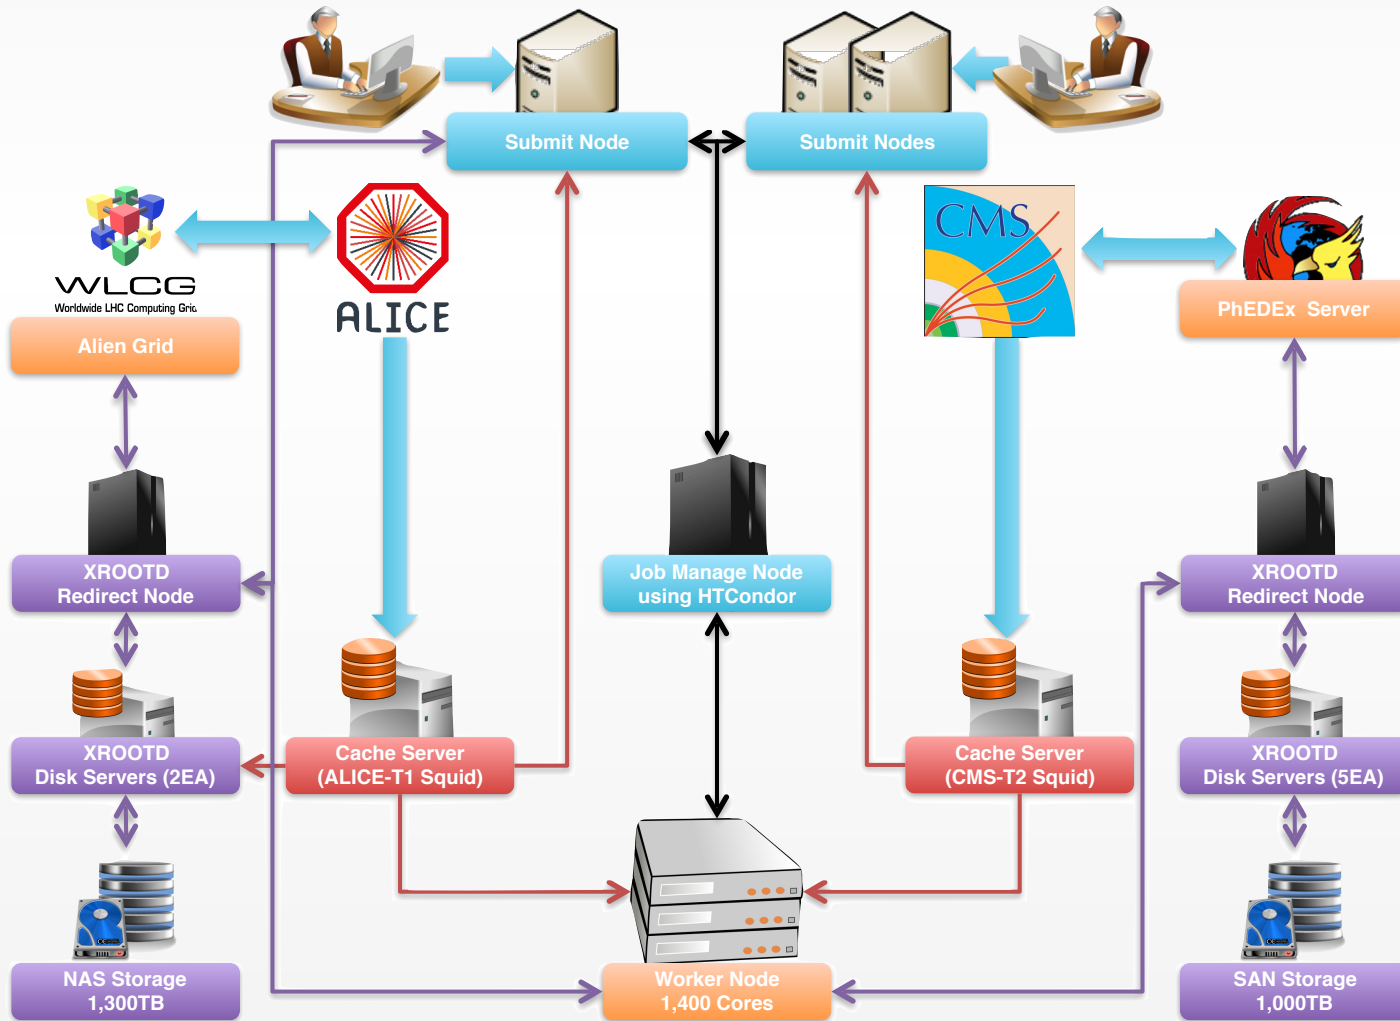
- **Integration Target**

- KIAF(ALICE Tier-3)
 - Batch System - HTCondor
 - Analysis Tool - AliROOT via CVMFS
- CMS Tier-3
 - Batch System - HTCondor
 - Analysis Tool - ROOT via CVMFS



- **Reason for Selection**

- Services using the same batch system HTCondor
- Easy to provide integrated analysis environment due to similar analysis environment such as analysis package and resource requirement per work

Integrated Pool Structure




Job Processing Policy




- **Unused resources can be used without restriction of any group**
 - Increase temporary idle resource efficiency
- **Ensuring the use of existing allocated resources among the consolidated resources**
 - Re-assigning by canceling the most recent operation among other use addition activities
- **Assign allocation priority according to slot memory size**
 - Increase dynamic allocation machine efficiency

Job Processing Scenario


0. Run ALICE job




1. Submit a new CMS job




2. Canceling an existing ALICE job




3. Canceled ALICE jobs are queued



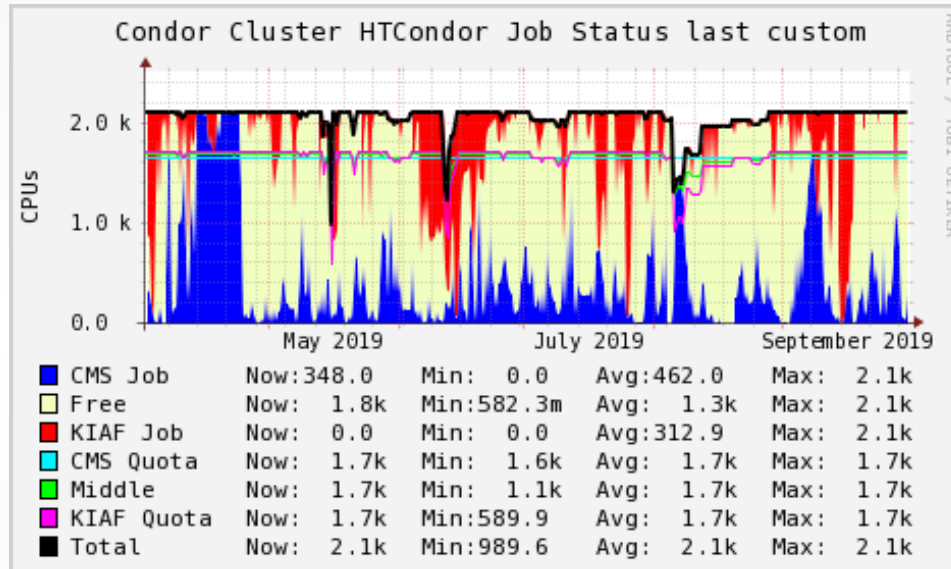
4. Execute and terminate the CMS job



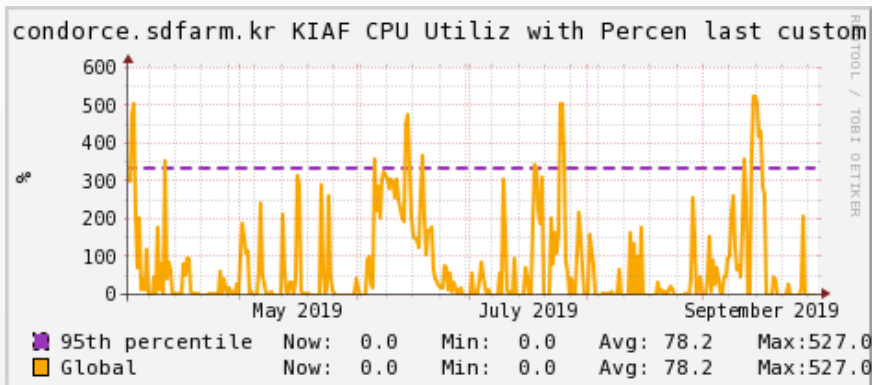
5. Restart the ALICE job




Integrated Pool Operating Result



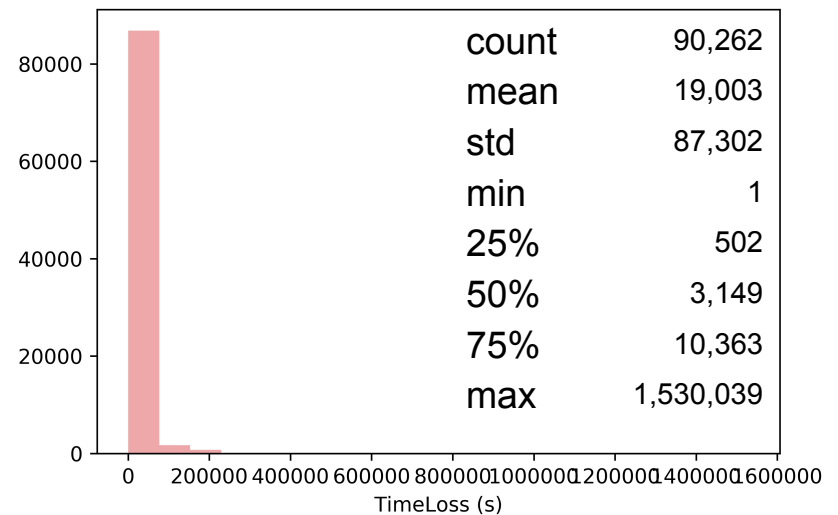

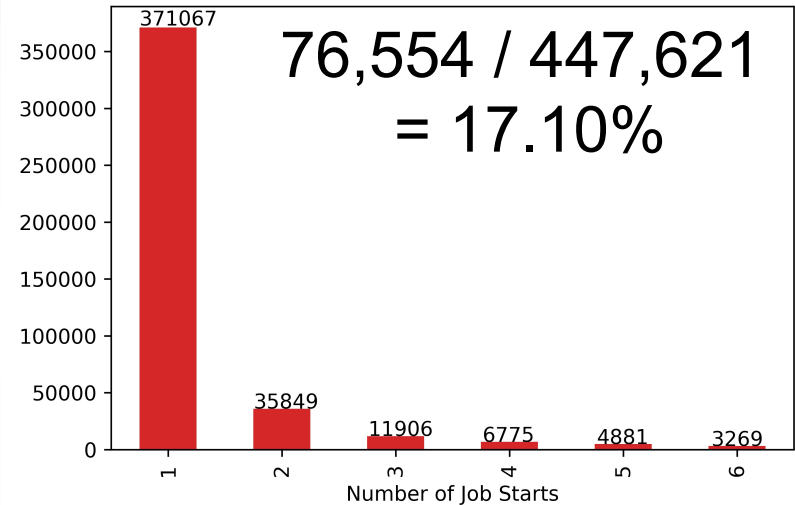
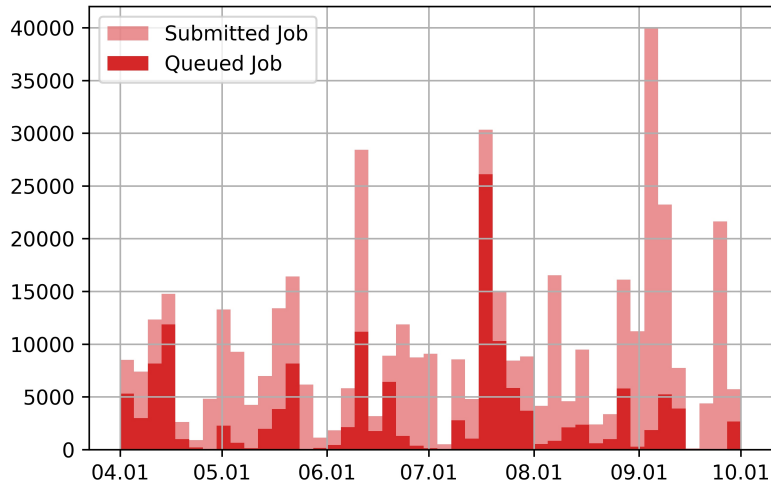
KoALICE - CPU 활용률 최대 527.0%



KCMS - CPU 활용률 최대 127.4%



Integrated Pool Operating Result

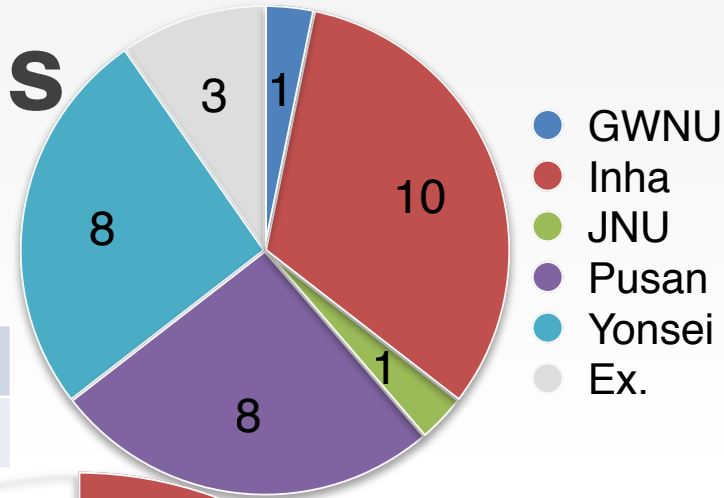


Operation Status

User Status

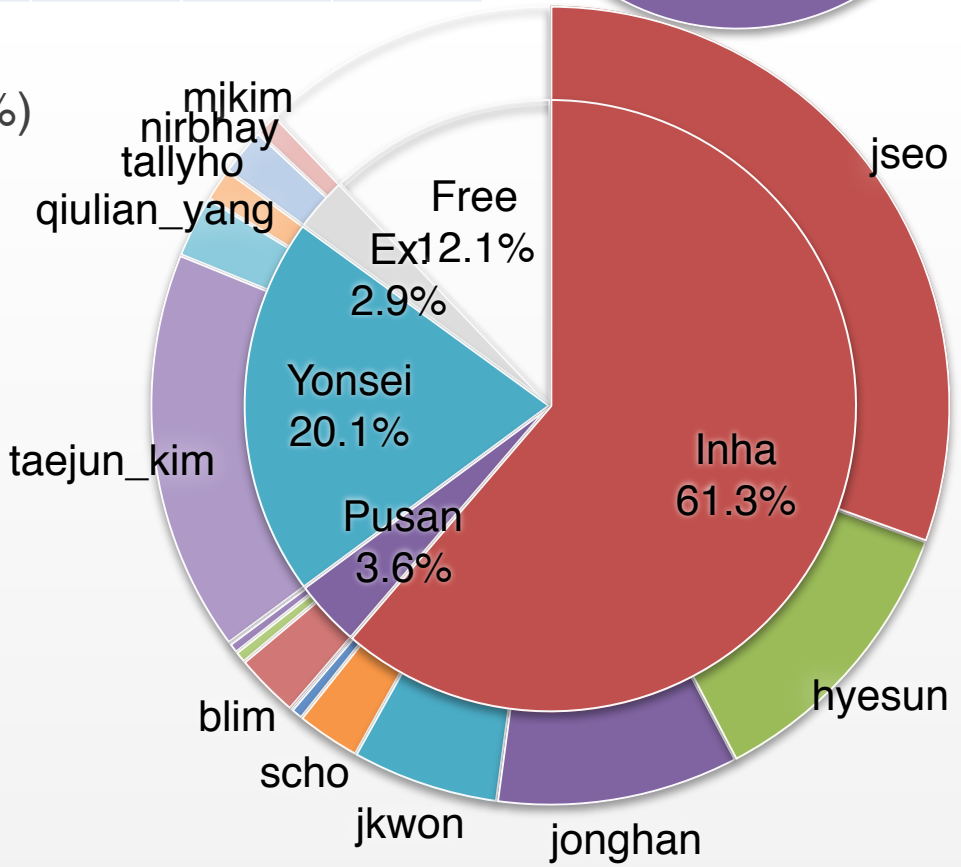

- User Accounts**

GWNU	Inha	JNU	Pusan	Yonsei	Ex.	Total
1	10	1	8	8	3	31



- User Home Directory**


– 17.6TB / 20TB (87.9%)



Job History

2019	January	February	March	April	May	June
Number of Jobs	25,380	86,240	129,812	53,492	68,634	69,416
Total Wall time (s)	223,297,157	286,638,302	586,071,860	468,045,943	396,517,930	1,452,265,649
Mean Queue time (s)	2,070	8,026	3,576	8,707	2,320	8,042

2019	July	August	September	October	November	December
Number of Jobs	85,501	57,669	113,232	22,688	13,531	50,959
Total Wall time (s)	877,598,949	282,100,529	973,472,236	174,687,797	79,775,508	864,757,219
Mean Queue time (s)	8,462	1,158	5,759	2,117	709	5,096



User Usage (Total)

	JobCount	Total Wall time (yy:ddd:hh:mm:ss)	Total CPU time (yy:ddd:hh:mm:ss)	Total Queue time (yy:ddd:hh:mm:ss)
Inha	427,092	104y:11d:1h:16m:03s	61y:363d:0h:02m:14s	75y:295d:12h:23m:35s
jkwon	158,832	2y:300d:22h:33m:46s	0y:220d:14h:01m:43s	0y:154d:21h:34m:51s
jonghan	175,457	49y:158d:15h:41m:11s	26y:45d:16h:50m:43s	51y:161d:20h:59m:53s
scho	92,803	51y:281d:11h:01m:06s	35y:96d:17h:09m:48s	23y:343d:17h:48m:51s
JNU	12,030	39y:270d:21h:18m:16s	27y:245d:23h:28m:06s	24y:3d:22h:56m:22s
junleekim	12,030	39y:270d:21h:18m:16s	27y:245d:23h:28m:06s	24y:3d:22h:56m:22s
Pusan	289,214	56y:49d:16h:55m:35s	23y:176d:15h:11m:22s	20y:256d:10h:51m:23s
blim	266,905	55y:145d:2h:26m:41s	22y:278d:11h:06m:09s	20y:49d:5h:50m:48s
chungmaria	402	0y:2d:17h:17m:16s	0y:2d:11h:20m:31s	0y:2d:3h:34m:26s
mkwon	21,907	0y:266d:21h:11m:38s	0y:260d:16h:44m:42s	0y:205d:1h:26m:09s
Yonsei	152,371	60y:101d:2h:56m:05s	13y:233d:19h:40m:45s	61y:272d:5h:45m:47s
taejun_kim	152,371	60y:101d:2h:56m:05s	13y:233d:19h:40m:45s	61y:272d:5h:45m:47s
nirbhay	523	0y:327d:7h:44m:34s	0y:309d:6h:03m:02s	0y:26d:13h:21m:35s
Total	881,230	261y:30d:2h:10m:33s	27y:233d:16h:25m:29s	182y:124d:17h:18m:42s

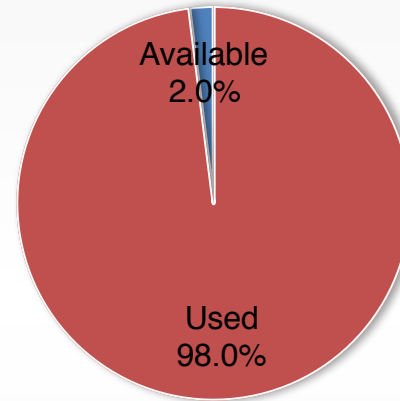
User Usage (Mean)

	JobCount	Mean Wall time (yy:ddd:hh:mm:ss)	Mean CPU time (yy:ddd:hh:mm:ss)	Mean Queue time (hh:mm:ss)
Inha	427,092	0y:0d:2h:08m:01s	0y:0d:1h:16m:18s	1h:33m:18s
jkwon	158,832	0y:0d:0h:09m:21s	0y:0d:0h:01m:60s	0h:01m:24s
jonghan	175,457	0y:0d:2h:28m:05s	0y:0d:1h:18m:16s	2h:34m:06s
scho	92,803	0y:0d:4h:53m:13s	0y:0d:3h:19m:44s	2h:15m:36s
JNU	12,030	0y:1d:4h:56m:22s	0y:0d:20h:09m:06s	17h:29m:03s
junleekim	12,030	0y:1d:4h:56m:22s	0y:0d:20h:09m:06s	17h:29m:03s
Pusan	289,214	0y:0d:1h:42m:01s	0y:0d:0h:42m:41s	0h:37m:37s
blim	266,905	0y:0d:1h:49m:05s	0y:0d:0h:44m:50s	0h:39m:39s
chungmaria	402	0y:0d:0h:09m:45s	0y:0d:0h:08m:51s	0h:07m:42s
mkwon	21,907	0y:0d:0h:17m:33s	0y:0d:0h:17m:08s	0h:13m:29s
Yonsei	152,371	0y:0d:3h:27m:55s	0y:0d:0h:47m:03s	3h:32m:59s
taejun_kim	152,371	0y:0d:3h:27m:55s	0y:0d:0h:47m:03s	3h:32m:59s
nirbhay	523	0y:0d:15h:01m:14s	0y:0d:14h:11m:29s	1h:13m:07s
Total	881,230	0y:0d:2h:35m:43s	0y:0d:1h:16m:08s	1h:48m:45s

Storage Usage (1/2)

- **DataSet Staging**

- 785TB / 800TB (Total 2,722 Runs)
- Staging speed (2 node, 200 stream)
 - Max 800MB/s, Avg.400MB/s




- **Staged Datasets**

Period	Run	Pass	# Datasets
LHC10b	114931-117222	pass4	31
LHC10c	119159-120829	pass4	36
LHC10d	122374-126437	pass4	54
LHC10e	127712-130848	pass4	114
LHC13b	195344-195483	pass4	8
LHC13c	195529-195677	pass4	14
LHC15o	245145-246994	pass1, pass1_pidfix	112
LHC16k	256504-258537	pass1, pass2	365
LHC16l	258919-260014	pass1, pass_itsrecpoints	145
LHC18r	296850-297595	pass1	71

Period	Run	# Datasets
LHC15h1i	193051-193189	4
LHC16i3a	245145-246994	110
LHC16i3b	245145-246994	110
LHC17c3b1	256504-258537	207
LHC17c3b2	258919-259888	140
LHC17d20a1_extra	256504-258498	207
LHC17d20a2_extra	258919-259888	70
LHC18c6b4	256941-258537	161
LHC18d8	258962-259888	58
LHC18e1b	245683-246994	77
LHC18e1b_extra	245683-246994	77
LHC18f1	256941-258537	161
LHC18f1_extra	256941-258537	171
LHC18f4b	256941-259888	219

Storage Usage (2/2)

	TB
Inha	321.65
jkwon	129.03
jonghan	192.63
Pusan	354.12
blim	354.12
Yonsei	28.07
tallyho	11.70
taejun_kim	16.37
nirbhay	67.52
Total	771.37



Stage dataset list

- KIAF command (in shell)

- showdatalist

- Google spreadsheet

- Link : https://docs.google.com/spreadsheets/d/1KrZ_YkPmGHoA4f2UQcjHSAcVSmTHcJMJJz9kIUkyZZo/edit?usp=sharing

[1KrZ_YkPmGHoA4f2UQcjHSAcVSmTHcJMJJz9kIUkyZZo/edit?usp=sharing](https://docs.google.com/spreadsheets/d/1KrZ_YkPmGHoA4f2UQcjHSAcVSmTHcJMJJz9kIUkyZZo/edit?usp=sharing)

```

[kong9@kiaf-test-01 macro3] showdataList
dataset Queue Down Success Fail TotalSize StagedSize
data/2011/LHC11a/000146686/ESDs/pass4_without_SDD 0 0 356 1 100.7GB 100.3GB 99%
data/2011/LHC11a/000146688/ESDs/pass4_without_SDD 0 0 459 3 130.0GB 129.2GB 99%
data/2011/LHC11a/000146689/ESDs/pass4_without_SDD 0 0 1136 0 312.2GB 312.2GB 100%
data/2011/LHC11a/000146746/ESDs/pass4_without_SDD 0 0 399 0 107.8GB 107.8GB 100%
data/2011/LHC11a/000146747/ESDs/pass4_without_SDD 0 0 1353 3 374.9GB 373.9GB 99%
data/2011/LHC11a/000146748/ESDs/pass4_without_SDD 0 0 504 0 151.3GB 151.3GB 100%
data/2011/LHC11a/000146801/ESDs/pass4_without_SDD 0 0 756 0 204.5GB 204.5GB 100%
data/2011/LHC11a/000146802/ESDs/pass4_without_SDD 0 0 1426 7 403.6GB 401.6GB 99%
data/2011/LHC11a/000146803/ESDs/pass4_without_SDD 0 0 157 1 37.0GB 36.8GB 99%
data/2011/LHC11a/000146804/ESDs/pass4_without_SDD 0 0 3265 13 923.2GB 919.4GB 99%
data/2011/LHC11a/000146806/ESDs/pass4_without_SDD 0 0 6997 32 1.9TB 1.8TB 99%
data/2011/LHC11a/000146808/ESDs/pass4_without_SDD 0 0 1193 3 331.3GB 330.5GB 99%
data/2011/LHC11a/000146807/ESDs/pass4_without_SDD 0 0 156 0 34.7GB 34.7GB 100%
data/2011/LHC11a/000146809/ESDs/pass4_without_SDD 0 0 875 0 219.5GB 219.5GB 100%
data/2011/LHC11a/000146812/ESDs/pass4_without_SDD 0 0 284 0 66.4GB 66.4GB 100%
data/2011/LHC11a/000146813/ESDs/pass4_without_SDD 0 0 958 2 246.0GB 245.4GB 99%
data/2011/LHC11a/000146814/ESDs/pass4_without_SDD 0 0 756 0 190.3GB 190.3GB 100%
data/2011/LHC11a/000146817/ESDs/pass4_without_SDD 0 0 672 4 181.9GB 180.7GB 99%
data/2011/LHC11a/000146824/ESDs/pass4_without_SDD 0 0 3453 8 9.9TB 9.9TB 99%
data/2011/LHC11a/000146855/ESDs/pass4_without_SDD 0 0 1189 1 312.3GB 312.0GB 99%
data/2011/LHC11a/000146857/ESDs/pass4_without_SDD 0 0 756 0 355.4GB 354.4GB 99%
data/2011/LHC11a/000146858/ESDs/pass4_without_SDD 0 0 2783 4 783.3GB 782.2GB 99%
data/2011/LHC11a/000146859/ESDs/pass4_without_SDD 0 0 1826 1 503.2GB 502.9GB 99%
data/2011/LHC11a/000146869/ESDs/pass4_without_SDD 0 0 672 1 173.6GB 172.8GB 99%
data/2013/LHC13b/000195344/ESDs/pass3 0 0 661 0 43.4GB 43.4GB 100%
data/2013/LHC13b/000195344/ESDs/pass3/A00154 0 0 72 0 9.3GB 9.3GB 100%
data/2013/LHC13b/000195351/ESDs/pass3 0 0 672 0 161.8GB 161.8GB 100%
data/2013/LHC13b/000195351/ESDs/pass3/A00154 0 0 73 0 37.6GB 37.6GB 100%
data/2013/LHC13b/000195389/ESDs/pass3 0 0 699 0 242.5GB 242.5GB 100%
data/2013/LHC13b/000195389/ESDs/pass3/A00154 0 0 76 0 49.8GB 49.8GB 100%
data/2013/LHC13b/000195391/ESDs/pass3 0 0 322 0 442.4GB 442.4GB 100%
data/2013/LHC13b/000195391/ESDs/pass3/A00154 0 0 117 0 71.7GB 71.7GB 100%
data/2013/LHC13b/000195478/ESDs/pass3 0 0 665 0 67.2GB 67.2GB 100%
data/2013/LHC13b/000195478/ESDs/pass3/A00154 0 0 89 0 15.2GB 15.2GB 100%
data/2013/LHC13b/000195479/ESDs/pass3 0 0 5921 0 2.3TB 2.3TB 100%
data/2013/LHC13b/000195479/ESDs/pass3/A00154 0 0 571 0 487.9GB 487.9GB 100%
data/2013/LHC13b/000195480/ESDs/pass3 0 0 1267 10 404.7GB 405.7GB 99%
data/2013/LHC13b/000195480/ESDs/pass3/A00154 0 0 154 0 92.2GB 92.2GB 100%
data/2013/LHC13b/000195481/ESDs/pass3 0 0 543 0 59.5GB 59.5GB 100%
data/2013/LHC13b/000195481/ESDs/pass3/A00154 0 0 57 0 13.4GB 13.4GB 100%
data/2013/LHC13b/000195482/ESDs/pass3 0 0 1977 0 727.6GB 727.6GB 100%
data/2013/LHC13b/000195482/ESDs/pass3/A00154 0 0 167 0 158.0GB 158.0GB 100%
data/2013/LHC13b/000195483/ESDs/pass3 0 0 10041 0 4.2TB 4.2TB 100%
data/2013/LHC13b/000195483/ESDs/pass3/A00154 0 0 938 0 0.9TB 0.9TB 100%
data/2013/LHC13c/000195529/ESDs/pass2 0 0 532 0 112.6GB 112.6GB 100%
data/2013/LHC13c/000195529/ESDs/pass2/A00154 0 0 56 0 26.2GB 26.2GB 100%
data/2013/LHC13c/000195531/ESDs/pass2 0 0 11422 0 5.9TB 5.9TB 100%
data/2013/LHC13c/000195531/ESDs/pass2/A00154 0 0 2049 0 2.1TB 2.1TB 100%
data/2013/LHC13c/000195566/ESDs/pass2 0 0 3596 0 1.4TB 1.4TB 100%
data/2013/LHC13c/000195566/ESDs/pass2/A00154 0 0 351 0 327.5GB 327.5GB 100%
data/2013/LHC13c/000195567/ESDs/pass2 0 0 3098 12 1.4TB 1.4TB 99%
    
```

KIAF DataSet List ☆

파일 수정 보기 삽입 서식 데이터 도구 부가기능 도움말 모든 변경사항이 드라이브에 저장됨

f_x =row()-1

	A	B	C	D	E	L	M	N	O	P
	No	Dataset	Size	requester	using	jisong	jkwon	jonghan	kbhuh	luxtau
1										
2	1	data/2011/LHC11a/000146686/ESDs/pass4_without_SDD	100.7GB	jisong	NO					
3	2	data/2011/LHC11a/000146688/ESDs/pass4_without_SDD	130.0GB	jisong	NO					
4	3	data/2011/LHC11a/000146689/ESDs/pass4_without_SDD	312.2GB	jisong	NO					
5	4	data/2011/LHC11a/000146746/ESDs/pass4_without_SDD	107.8GB	jisong	NO					
6	5	data/2011/LHC11a/000146747/ESDs/pass4_without_SDD	374.8GB	jisong	NO					
7	6	data/2011/LHC11a/000146748/ESDs/pass4_without_SDD	151.3GB	jisong	NO					
8	7	data/2011/LHC11a/000146801/ESDs/pass4_without_SDD	204.5GB	jisong	NO					
9	8	data/2011/LHC11a/000146802/ESDs/pass4_without_SDD	403.6GB	jisong	NO					
10	9	data/2011/LHC11a/000146803/ESDs/pass4_without_SDD	37.0GB	jisong	NO					
11	10	data/2011/LHC11a/000146804/ESDs/pass4_without_SDD	923.2GB	jisong	NO					
12	11	data/2011/LHC11a/000146805/ESDs/pass4_without_SDD	1.8TB	jisong	NO					
13	12	data/2011/LHC11a/000146806/ESDs/pass4_without_SDD	331.3GB	jisong	NO					
14	13	data/2011/LHC11a/000146807/ESDs/pass4_without_SDD	34.7GB	jisong	NO					
15	14	data/2011/LHC11a/000146808/ESDs/pass4_without_SDD	219.5GB	jisong	NO					
16	15	data/2011/LHC11a/000146812/ESDs/pass4_without_SDD	66.4GB	jisong	NO					
17	16	data/2011/LHC11a/000146813/ESDs/pass4_without_SDD	246.0GB	jisong	NO					
18	17	data/2011/LHC11a/000146814/ESDs/pass4_without_SDD	190.3GB	jisong	NO					
19	18	data/2011/LHC11a/000146817/ESDs/pass4_without_SDD	181.9GB	jisong	NO					
20	19	data/2011/LHC11a/000146824/ESDs/pass4_without_SDD	0.9TB	jisong	NO					
21	20	data/2011/LHC11a/000146856/ESDs/pass4_without_SDD	312.3GB	jisong	NO					

Farm Management

- **Fortnightly KIAF User Meeting @ Vidyo**
 - Next Meeting January 15th 17:00(KST)
 - Link : <https://vidyoportal.cern.ch/join/pgPcHZZaMlfK>
 - Agenda
 - KIAF Operation
 - User Feedback
 - Use policy
- **KIAF User Guide**
- **Contact**
 - kiaf-user-support@kisti.re.kr
 - dataset-request@kisti.re.kr

Discussion

THANK YOU
