

Status Report of Pilot ALICE-HPC Project using NURION Resources at KISTI

Hyeon-Jin Yu

*Integrated M.S. and Ph.D student,
Chungbuk National University, South Korea
hyeonjin.yu@cern.ch*

2023.11.01(Wed)

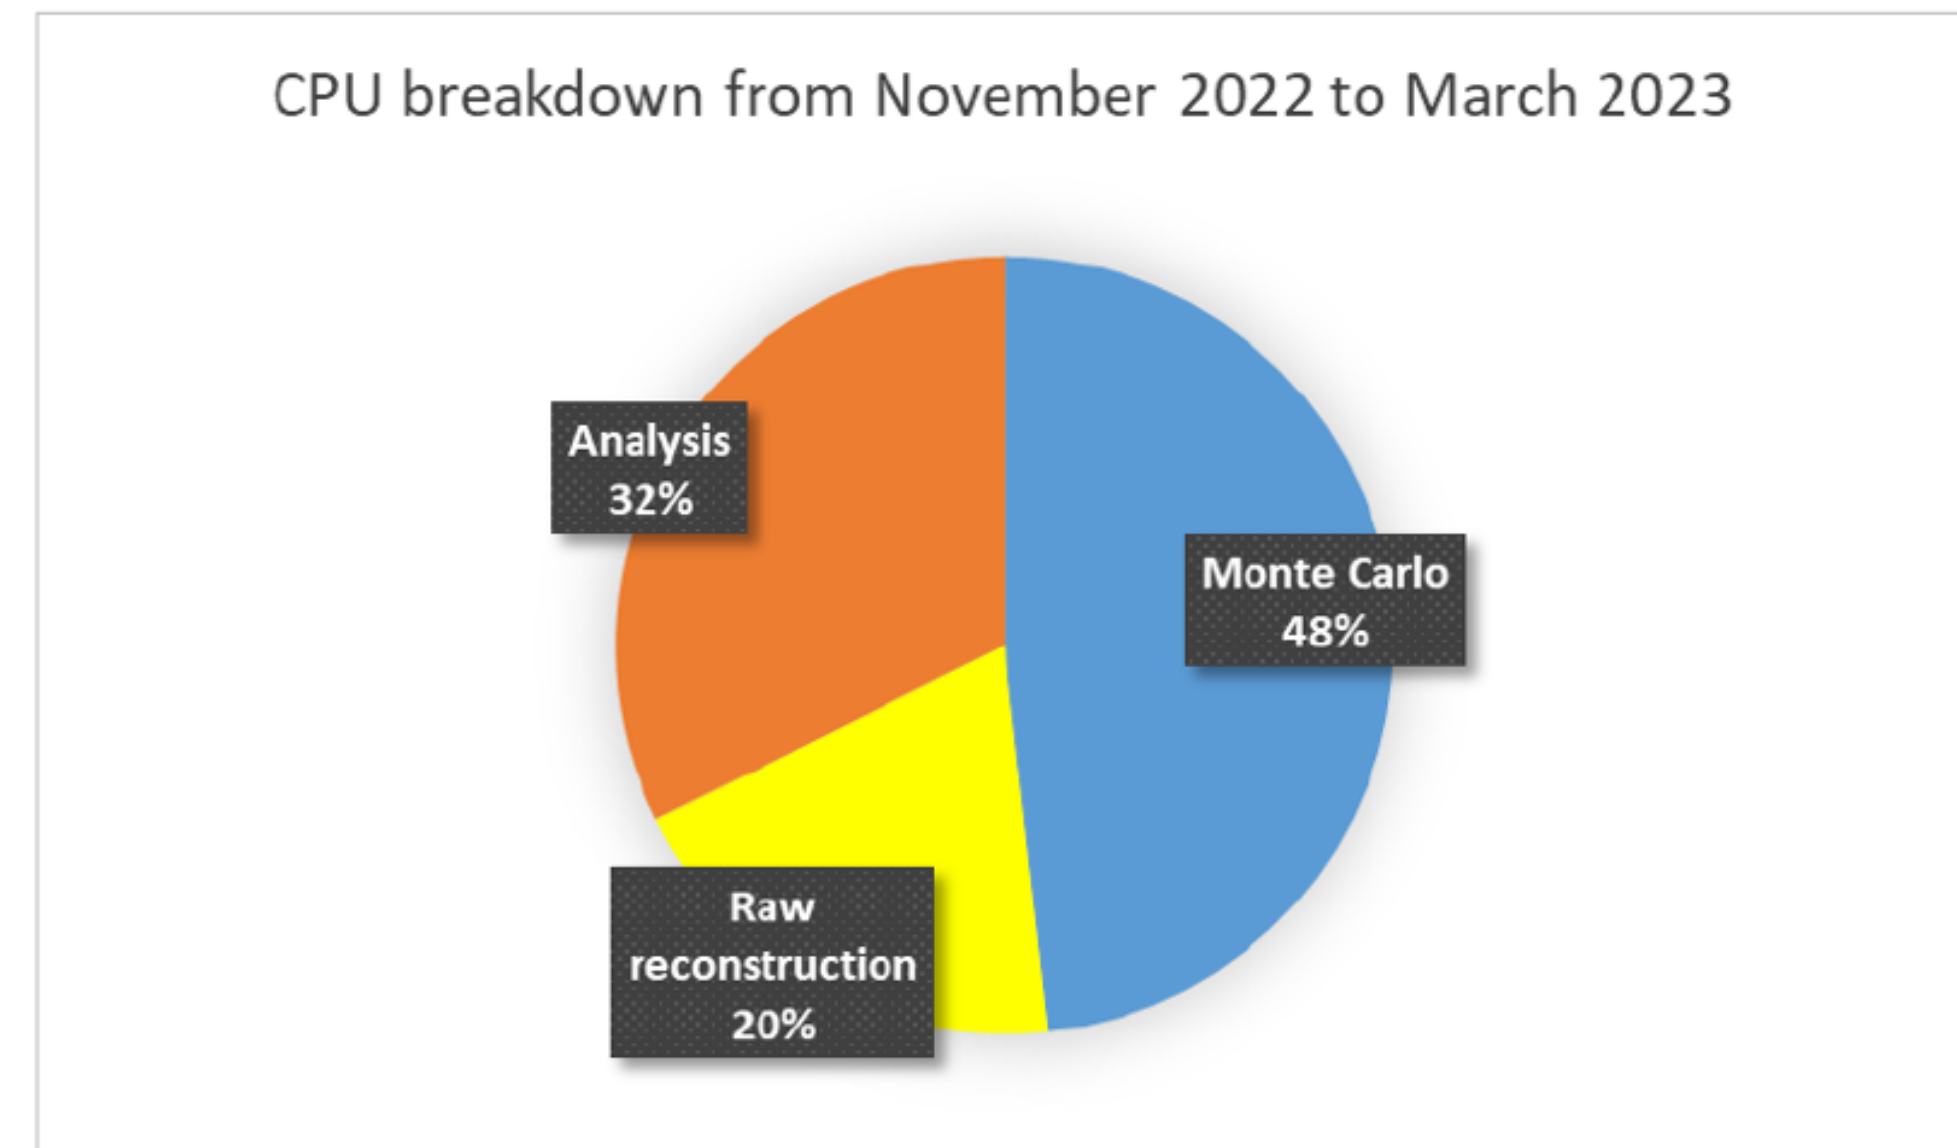

INDEX

-
- 1. The Reason for Project**
 - 2. The Project Participants**
 - 3. The Project Architecture**
 - 4. The Project Configuration**
 - 5. Future Plans**
-

1. The Reason for Project

❖ What situation is currently in CERN?

- CPU aspect
 - overusing the served CPU resource
 - about 10% lower Tier-1's Pledged resource compared by the required.
 - Monte Carlo Simulation jobs using the most CPU usage



1. The Reason for Project



❖ What situation is currently in CERN?

- ALICE-5 recommendation of C-RSG (Computing Resource Scrutiny Group)
 - required resource in 2024 - CPU 10% ↑, Disk 15% ↑, Tape 35% ↑
 - ALICE-5 - ALICE is encouraged to increase its opportunistic resources beyond T0 on HPCs.

▶ Efforts are being done by both HPC and LHC communities to be able to run on HPC resources.

ALICE		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	471	471	921	541	541	600	111%	600
	Tier-1	498	448	505	572	506	630	110%	630
	Tier-2	515	517	507	592	567	650	110%	650
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	1484	1436	1933	1705	1614	1880	110%	1880
	<i>Others</i>			139					
Disk	Tier-0	50.0	50.0	46.6	58.5	58.5	67.5	115%	67.5
	Tier-1	55.0	49.7	38.3	63.5	57.6	71.5	113%	71.5
	Tier-2	49.0	55.2	40.3	57.5	60.4	66.5	116%	66.5
	Total	154.0	154.9	125.2	179.5	176.5	205.5	114%	205.5
Tape	Tier-0	95.0	95.0	61.4	131.0	131.0	181.0	138%	181.0
	Tier-1	63.0	71.8	39.5	82.0	87.7	107.0	130%	107.0
	Total	158.0	166.8	100.9	213.0	218.7	288.0	135%	288.0

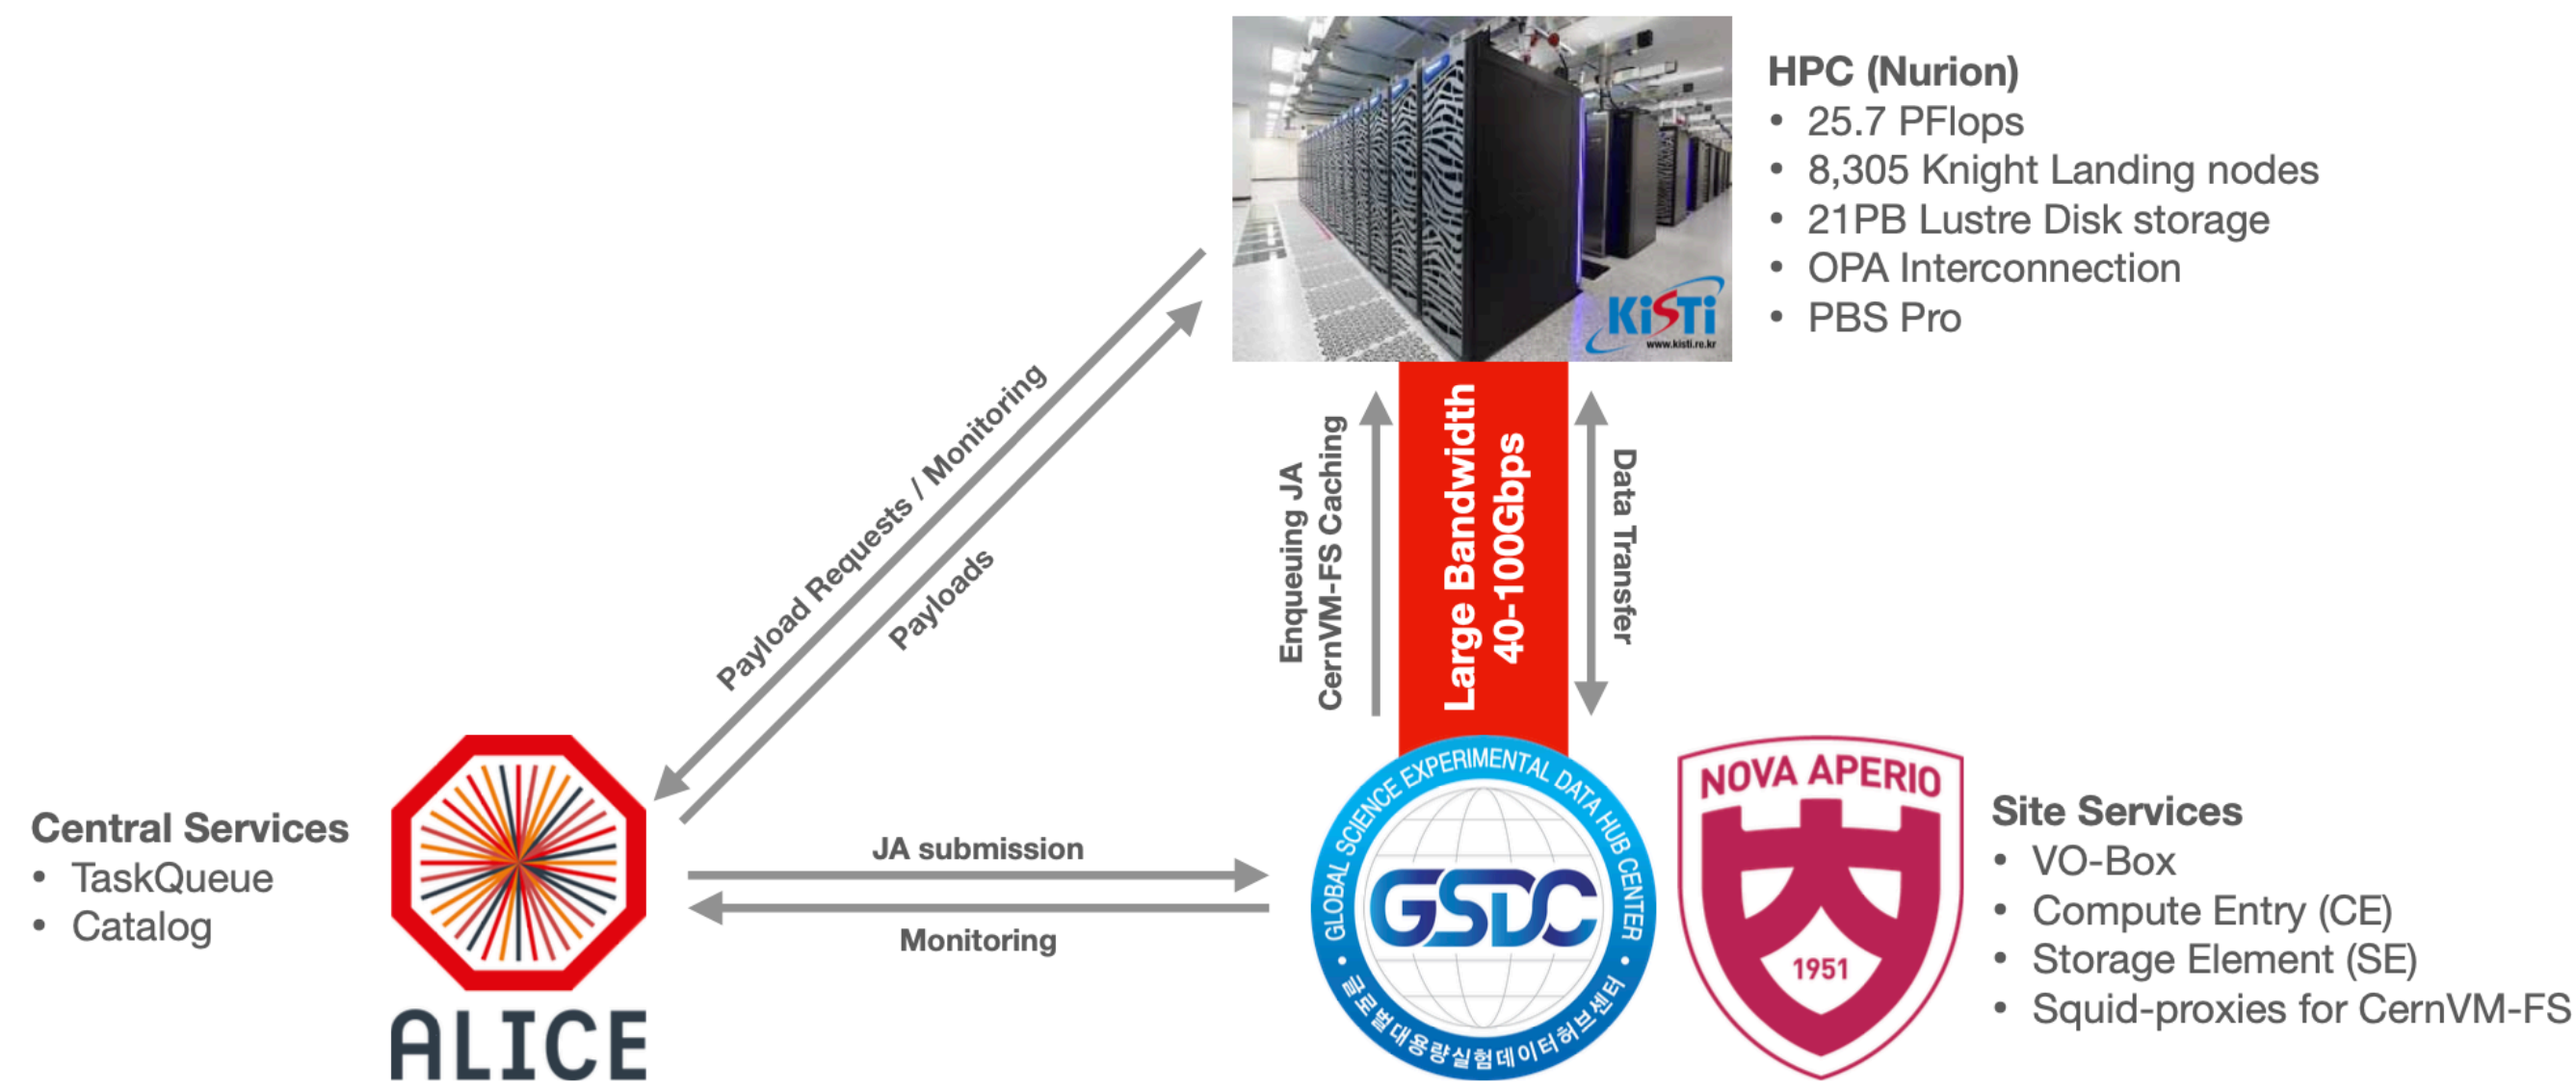
2. The Project Participants



❖ Introduction of the project participating institutes

- The project's goal is to construct ALICE Grid job execution environment on an HPC cluster in South Korea.
- Currently, a script job that mounts a CVMFS repository has been done in this environment successfully.

KISTI HPC for ALICE



• ALICE

- One of CERN experiment group
- Job sumitee

• KISTI

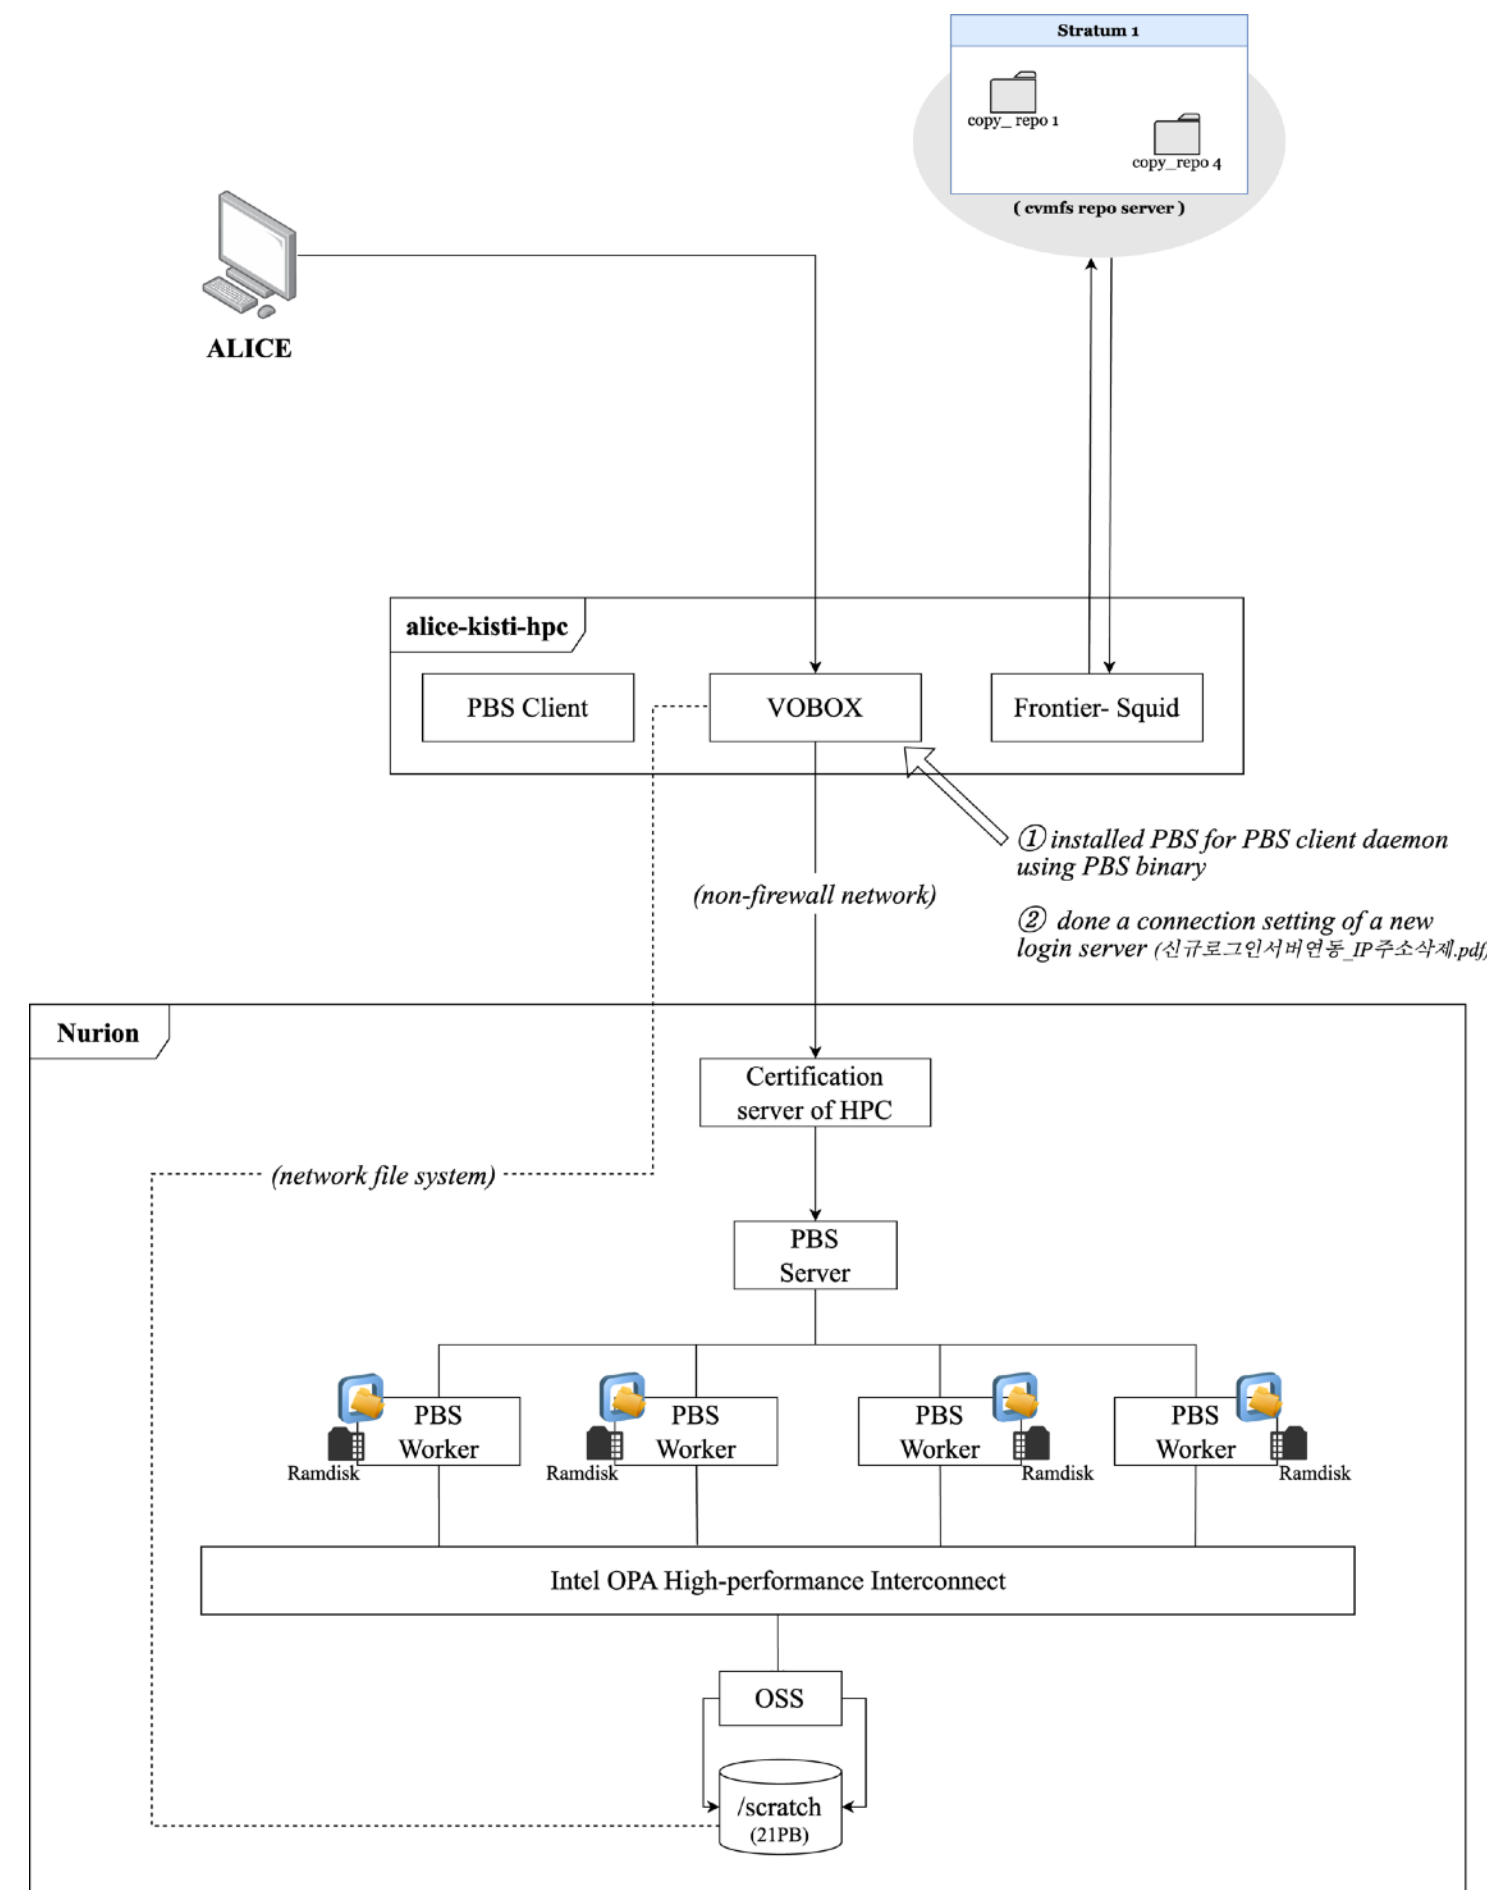
- Korea Institute of Science and Technology Informantion
- Computing resource provider
- Project environment builder

• CBNU

- Chungbuk National University
- Project environment builder

3. The Project Architecture

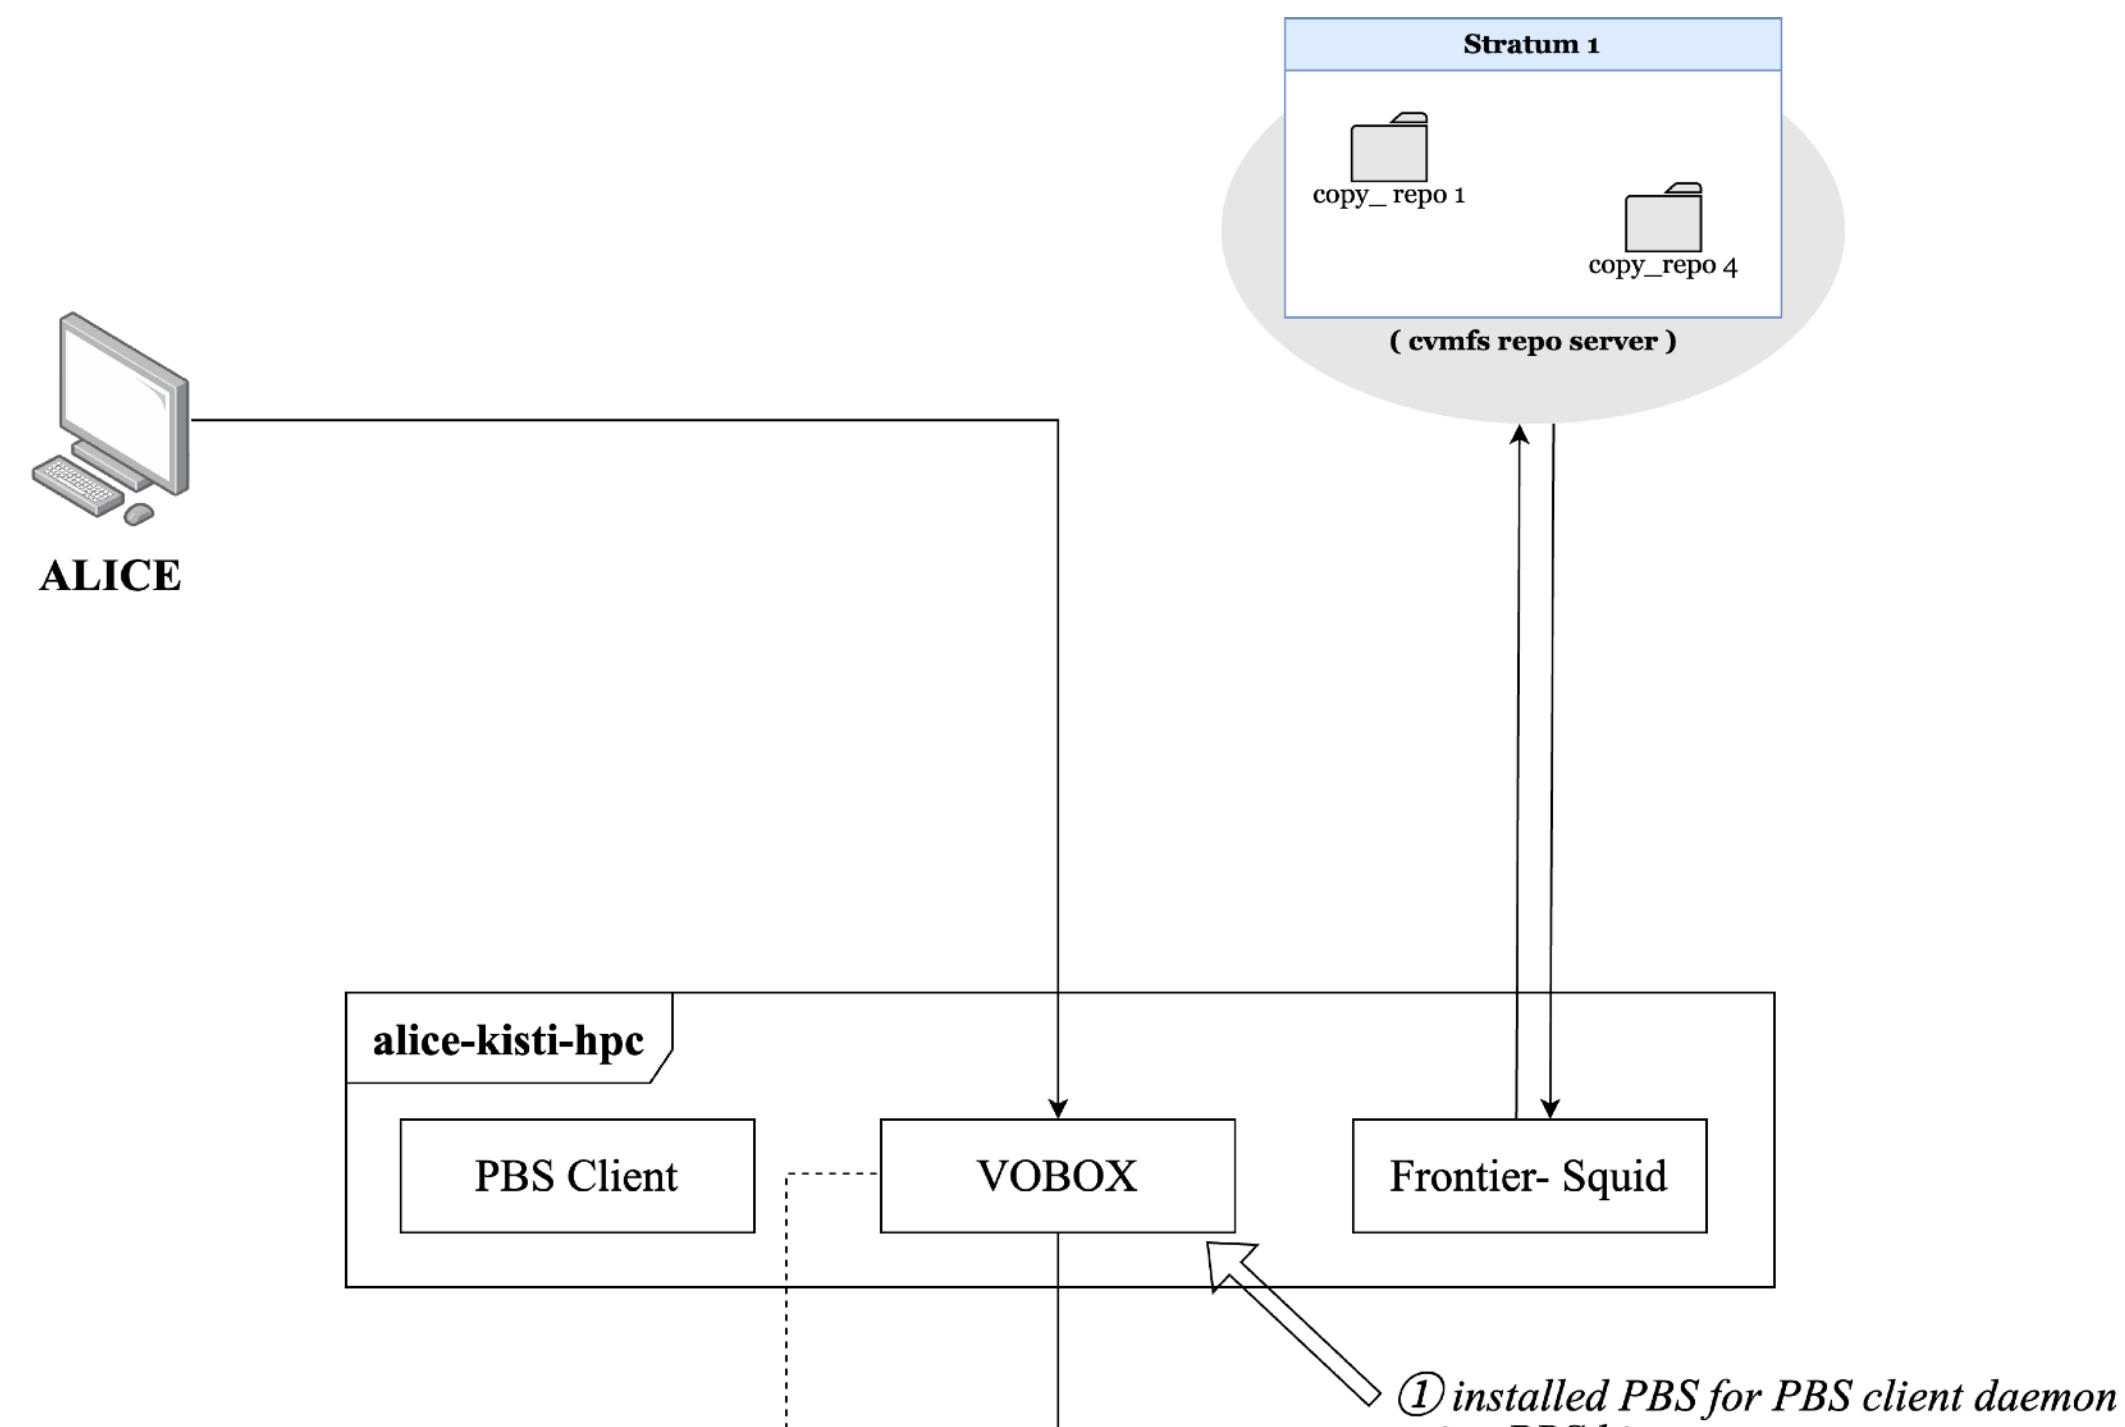
❖ Architecture1 - Project environment for ALICE Grid jobs



1. Retrieve ALICE Grid jobs submitted by authorized users in VOBOX
2. Craete JobAgent script to search for available nodes
3. Change to gsd23a01 user, a job submission user
4. Move to /scratch/gsd23a01
5. Submit a job that runs the JobAgent to the PBS queue
6. Aollcate the job on an HPC worker node by the PBS server
7. If there are some resources for ALICE Grid jobs, bring and execute grid jobs for the JobAgent's lifetime

3. The Project Architecture

❖ Detailed Architecture1 - GSDC Site

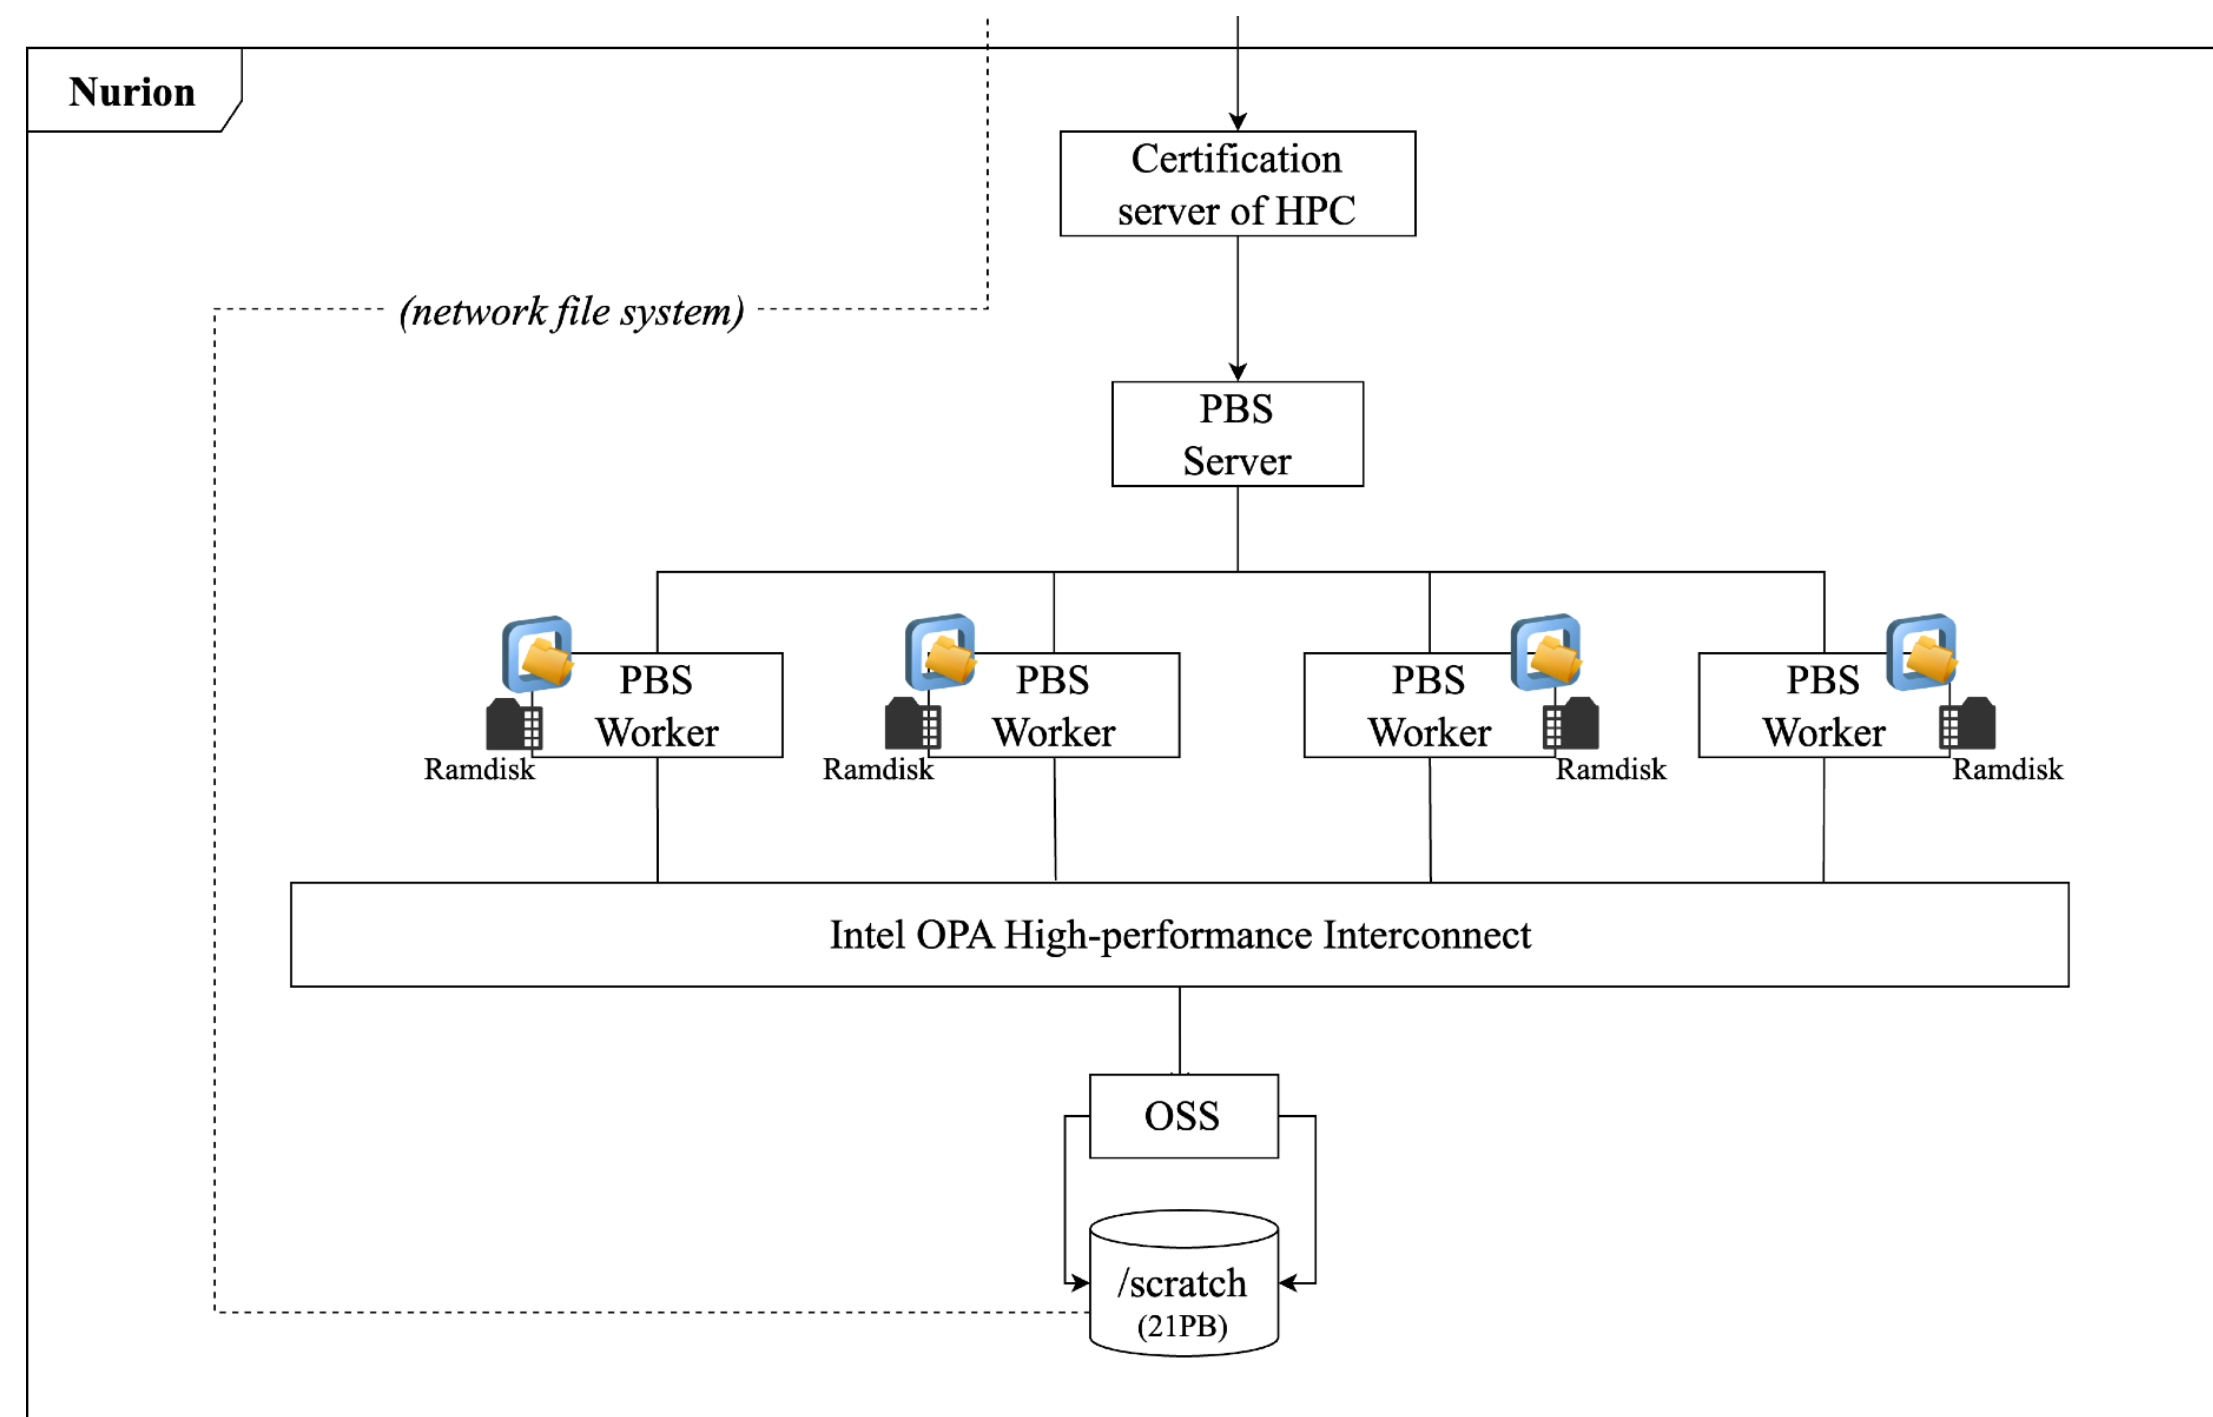


NODE : 'alice-kisti-hpc'

- installed packages
 - vobox
 - cvmfs
 - frontier-squid
 - pbs-execution
- roles
 - authentication
 - proxy server
 - job submission
- using network file system
 - mounted /scratch/gsd23a01 as a job submission path

3. The Project Architecture

❖ Detailed Architecture1 - NURION Site



NODE : PBS Server

- installed package
 - pbs-server
- roles
 - allocating jobs to available worker nodes

NODE : PBS worker

- installed packages
 - cvmfs
 - pbs-execution
- roles
 - job execution
- why ramdisk used?
 - The nodes don't have a disk, so we use a ramdisk instead of disks

4. The Project Configuration



❖ Overview

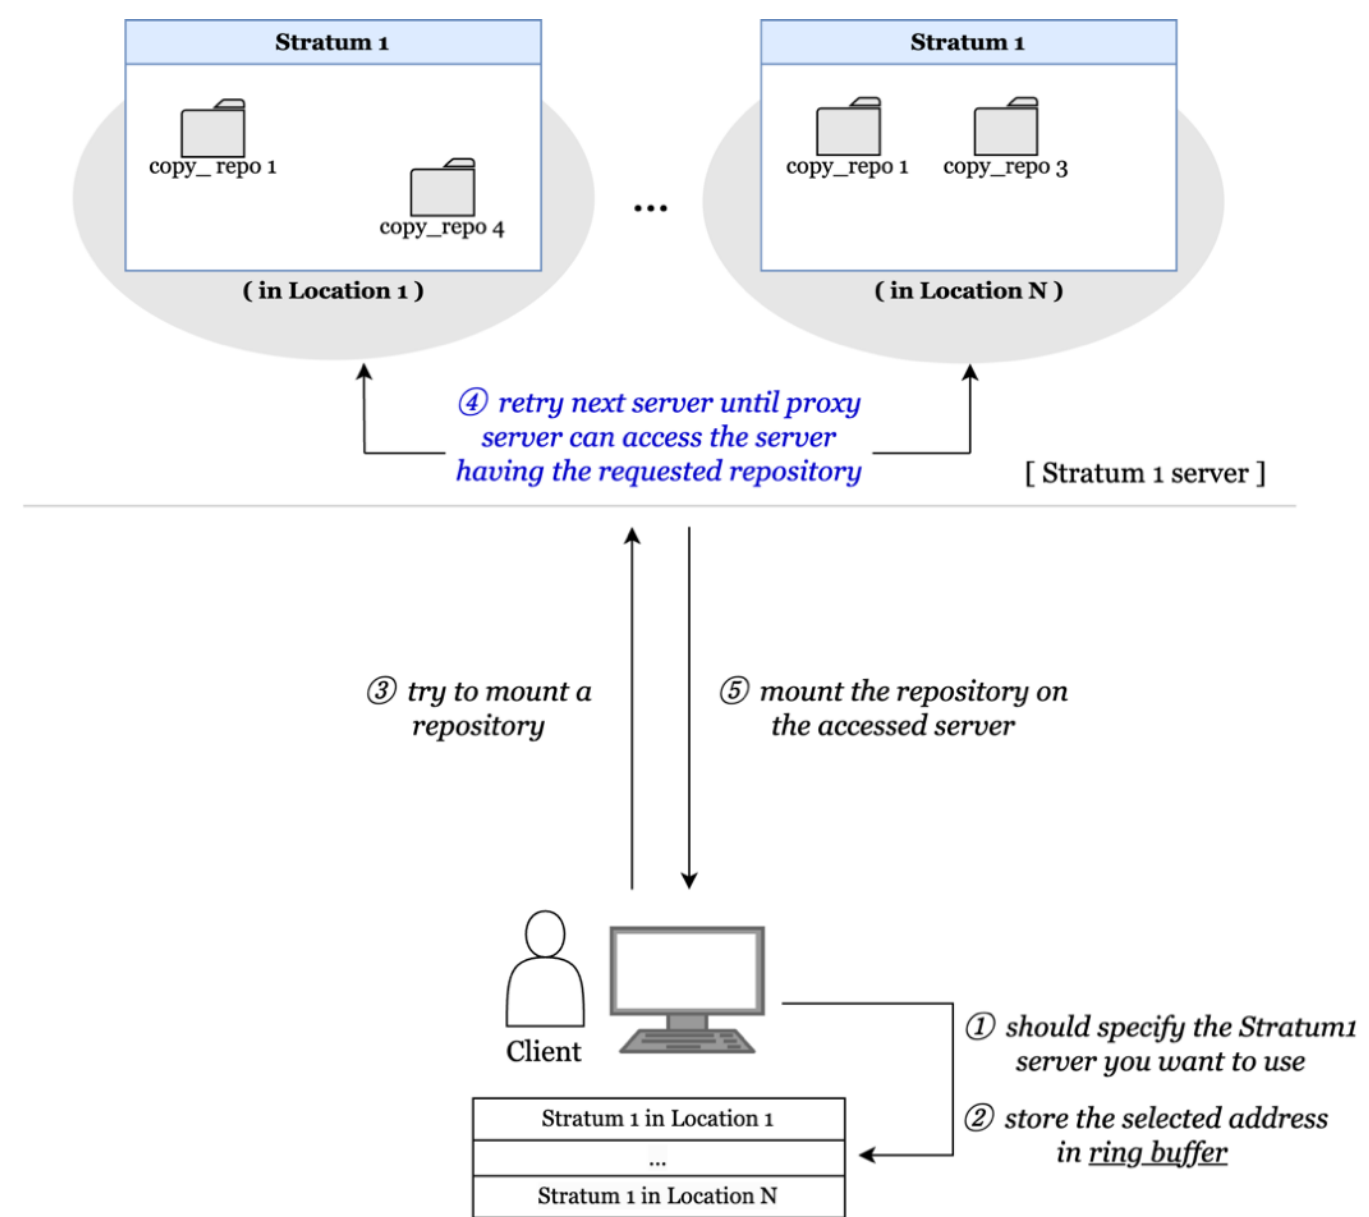
- Introduce what VOBOX, CVMFS, Frontier-squid, PBS, NFS are, as mentioned the before chapter
- Explain how to configure them for this project

Frontier-squid	<ul style="list-style-type: none">• A proxy server used for CVMFS• This proxy has the information about the stratum servers storing CVMFS repositories.
CVMFS	<ul style="list-style-type: none">• A file system that stores various experimental data, package, software, etc. from CERN• It is similar to Github.
NFS	<ul style="list-style-type: none">• A protocol that allows accessing file systems on other nodes through a network• It's used to access a directory for job submission to HPC.
PBS	<ul style="list-style-type: none">• A software that optimizes job scheduling and workload management• It is similar to HTCondor.
VOBOX	<ul style="list-style-type: none">• A system to support ALICE VO services• This system is necessary to grant access to the project environment only for authorized users.

4. The Project Configuration

❖ 1st. Frontier-squid

- Frontier-squid is a proxy server used to mount the `alice.cern.ch` repository required for ALICE Grid jobs.
- How to configure the proxy server?
 - add a `http_proxy` environment variable
 - edit `/etc/squid/customize.sh`



[a 'http_proxy' environment variable]

```
[root@alice-kisti-hpc pbs]# export | grep proxy
declare -x http_proxy="http://alice-kisti-hpc.sdfarm.kr:3128"
[root@alice-kisti-hpc pbs]#
```

[/etc/squid/customize.sh]

```
13 awk --file `dirname $0`/customhelps.awk --source '{
14 uncomment("acl MAJOR_CVMFS")
15 insertline("^# http_access deny !RESTRICT_DEST", "http_access deny !MAJOR_CVMFS")
16 insertline("^acl Safe_ports port 777", "acl CONNECT method CONNECT")
17 setoption("acl NET_LOCAL src", "0.0.0.0/0")
18 setoption("cache_mem", "128 MB")
19 setoptionparameter("cache_dir", 3, "10000")
20 print
21 }'
```

- a above file means that when we mount something in CVMFS, the proxy only accesses MAJOR_CVMFS stratum servers


4. The Project Configuration

❖ 1st. Frontier-squid

- The path in MAJOR_CVMFS is a list that stratum servers having alice repositories.
- We choose MAJOR_CVMFS using a 'uncomment' command in /etc/squid/cutomize.sh.

[/etc/squid/customize.sh]

```
13 awk --file `dirname $0`/customhelps.awk --source '{
14   uncomment("acl MAJOR_CVMFS")
15   insertline("^# http_access deny !RESTRICT_DEST", "http_access deny !MAJOR_CVMFS")
16   insertline("^acl Safe_ports port 777" "acl CONNECT method CONNECT")
```



[a proxy config file - /etc/squid/squid.conf]

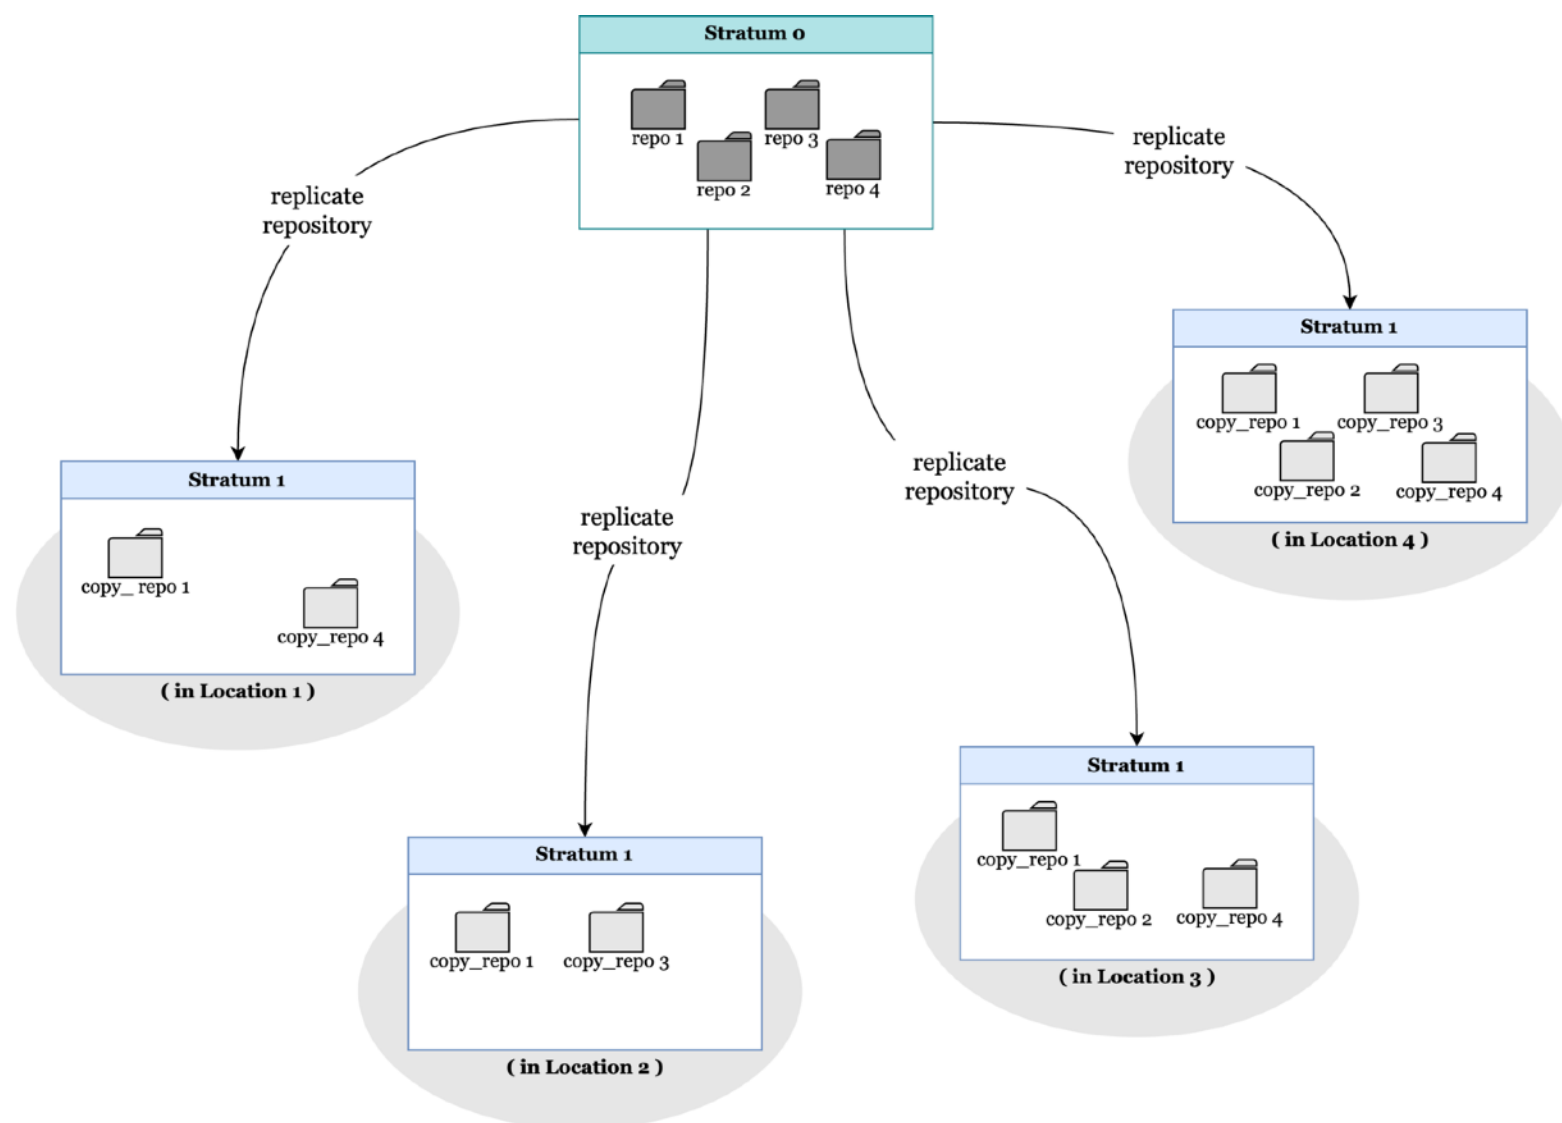
```
# acl CMS_FRONTIER dstdom_regex ^(cmsfrontier.*\.cern\.ch|cms.*frontier\.openhtc\.io)$
# acl ATLAS_FRONTIER dstdom_regex ^(atlas.*frontier.*\.cern\.ch|cc.*\.in2p3\.fr|lcg.*\.gridpp
\.rl\.ac\.uk|(.*)frontier.*|tier1nfs)\.triumf\.calatlas.*frontier\.openhtc\.io)$
acl MAJOR_CVMFS dstdom_regex ^((cvmfs-stratum-.*|hepvm)\.cern\.ch|(cernvmfs|cvmfs).*\.gridpp\
\.rl\.ac\.uk|cvmfs.*\.(\racf|sdcc)\.bn|\.gov|cvmfs.*\.fnal\.gov|(cvmfs01|klei)\.nikhef\.nl|cvmf
s.*\.grid\.sinica\.edu\.tw|cvmfs.*\.lcg\.triumf\.cal(cvmfs-s1|oasis).*\.opensciencegrid\.org|
cvmfs.*\.ihep\.ac\.cn|hcc-cvmfs\.unl\.edu|(cvmfs-stratum-one\.zeuthen|grid-cvmfs-one)\.desy\
delcc.*\.in2p3\.fr|.*cvmfs\.openhtc\.io|(cvmfs-s1.*|object-.*\.cloud)\.computecanada\.(\calnet
)|sampacs.*\.if\.usp\.br|cvmfs-.*\.hpc\.swin\.edu\.au|cvmfs-stratum-one\.cc\.kek\.jp)$
```

4. The Project Configuration

❖ 2nd. CVMFS

- CVMFS is a file system that stores repositories including packages, experimental data, etc..
- We can check the following on the CernVM-FS Repository Monitor page:
 - what repositories are in CVMFS.
 - where the repositories are in.

[CVMFS Repository servers]



[Repository Monitor page]

Replication Monitor

CernVM-FS Repository Monitor

This is a website for monitoring CernVM-FS repositories. It provides basic information about a repository and its status, notably Stratum1 replication status.

Registered Repositories

- alice-cvdf.cern.ch
- alice.cern.ch
- atlas-condp.cern.ch
- atlas-nightlies.cern.ch
- atlas.cern.ch
- belle.kek.jp

ALICE Software

Browse this repository (experimental)

Repository alice.cern.ch	
Revision	16956
Oldest stratum1 revision	16956
Last modified	28th October 2023 12:07:16 pm
Whitelist expiry date	2024-01-11T08:00:00.000Z

Stratum 1

Organisation	Status	Revision	Last Modified	URL
CERN	✓	16956	28th October 2023 12:07:16 pm	http://cvdfs-stratum-one.cern.ch/cvdfs
SDCC by RACF at BNL	✓	16956	28th October 2023 12:07:16 pm	http://cvdfs.sdcc.bnl.gov/cvdfs
Your Organisation	✓	16956	28th October 2023 12:07:16 pm	http://cvdfs.fnal.gov/cvdfs
STFC Rutherford Appleton Lab	✓	16956	28th October 2023 12:07:16 pm	http://cvdfs-wlcg.gridpp.rl.ac.uk/cvdfs

4. The Project Configuration



❖ 2nd. CVMFS

- There are some parameters for configuring CVMFS.
- The parameters
 - CVMFS_REPOSITORIES - the repository list you want to mount (one or more)
 - CVMFS_HTTP_PROXY - the your proxy server list (one or more)
 - CVMFS_CACHE_BASE - the location of the cache directory
 - CVMFS_QUOTA_LIMIT - the cache size (MB)

[CVMFS config file - /etc/cvmfs/default.local]

```
[root@alice-kisti-hpc ~]# cat /etc/cvmfs/default.local
CVMFS_REPOSITORIES=alice.cern.ch
#CVMFS_CLIENT_PROFILE=single
CVMFS_HTTP_PROXY="http://alice-kisti-hpc.sdfarm.kr:3128"
#CVMFS_HTTP_PROXY="http://alice-t1-squid.sdfarm.kr:3128"
CVMFS_CACHE_BASE="/cache-cvmfs"
CVMFS_QUOTA_LIMIT="17000"
```

4. The Project Configuration



❖ 2nd. CVMFS

- When we mount some repositories using CVMFS, Frontier-squid logs are created as shown below.
 - TCP-MISS : the repository isn't cached before (= these repositories are mounted first time)
 - TCP-HIT : the repository is cached (= these repositories are mounted second time)

[mounting the alice repository]

```
[root@alice-kisti-hpc ~]# cvmfs_config probe
Probing /cvmfs/alice.cern.ch... OK
[root@alice-kisti-hpc ~]#
[root@alice-kisti-hpc ~]# cd /cvmfs/alice.cern.ch/bin ; ll
total 35
drwxrwxr-x.  2 cvmfs cvmfs   20 Jun  1 20:16 .
drwxr-xr-x. 24 cvmfs cvmfs 4096 Nov 27  2012 ..
-rwxrwxr-x.  1 cvmfs cvmfs  139 Sep 23  2013 alien
-rwxrwxr-x.  1 cvmfs cvmfs  165 Sep 23  2013 aliend
-rwxr-xr-x.  1 cvmfs cvmfs 14404 Jul 11 00:46 alienv
-rwxr-xr-x.  1 cvmfs cvmfs 14347 Jan 10  2023 alienv.safe
[root@alice-kisti-hpc bin]#
[root@alice-kisti-hpc bin]# ./alienv
Unknown distribution release: CentOS 7.9.2009
[root@alice-kisti-hpc bin]#
```

[/var/log/squid/access.log (first mount)]

```
- - [26/Oct/2023:16:44:58.019 +0900] "GET http://cernvmfs.gridp
p.rl.ac.uk/cvmfs/alice.cern.ch/data/f8/7cc7c8734c25af2579c8372861298fa40bdd6dC
HTTP/1.1" 503 4437 TCP_MISS:HIER_DIRECT 0 "-" alice.cern.ch:/%20%28f87cc7c8734
c25af2579c8372861298fa40bdd6d%29" "-" "cvmfs Fuse 2.11.0"
```



[/var/log/squid/access.log (second mount)]

```
- - [26/Oct/2023:16:44:58.216 +0900] "GET http://cvmfs-stratum-
one.cern.ch/cvmfs/alice.cern.ch/data/f8/7cc7c8734c25af2579c8372861298fa40bdd6d
C HTTP/1.1" 200 4625826 TCP_HIT:HIER_NONE 196 "-" alice.cern.ch:/%20%28f87cc7c8
734c25af2579c8372861298fa40bdd6d%29" "-" "cvmfs Fuse 2.11.0"
```

4. The Project Configuration



❖ 3nd. NFS

- NFS is a networking protocol for sharing /scratch/gsd23a01.
- How to configure NFS?
 - mount /scratch/gsd23a01, a ALICE Grid jobs' submission path
 - edit ldap.conf and sssd.conf

[mounting /scratch/gsd23a01]

```
[root@alice-kisti-hpc ~]# mount -t nfs      :/home01/gsd23a01 /home01/gsd23a01
[root@alice-kisti-hpc ~]# mount -t nfs      :/scratch/gsd23a01 /scratch/gsd23a01
```

** gsd23a01 user is a job submission user.*



```
[gsd23a01@alice-kisti-hpc ~]$ ll /home01/gsd23a01/
total 56900
drwxr-xr-x. 2 gsd23a01 in0138   4096 Oct  5 23:04 certs
drwxr-xr-x. 2 gsd23a01 in0188   4096 May 18 14:06 job_examples
-rw-r--r--. 1 gsd23a01 in0188 58254308 Sep 26 15:52 pbspro-execution-2020.1.3.202103151607
38-0.e17.x86_64.rpm
```

```
[gsd23a01@alice-kisti-hpc ~]$ ll /scratch/gsd23a01/
total 8
-rwxr-x---. 1 gsd23a01 in0188 224 Oct 16 15:26 serial.sh
-rw-r-----. 1 gsd23a01 in0188  81 Oct 19 13:09 test.c
```

4. The Project Configuration



❖ 3rd. NFS

- These files allow that only gsd23a01 user has access to the mounted directories.
 - LDAP - a Lightweight Directory Access Protocol to search for information over a network
 - SSSD - a System Security Service Daemon for accessing remote directories and authentication services.

[/etc/openldap/ldap.conf]

```
SASL_NOCANON on
URI ldaps://
BASE dc=cm,dc=cluster
TLS_REQCERT never
TLS_CACERT /etc/openldap/certs/ca.pem
TLS_CERT /etc/openldap/certs/ldap.pem
TLS_KEY /etc/openldap/certs/ldap.key
```

[/etc/sss/sss.conf]

```
1 [sss]
2 config_file_version = 2
3 domains = LDAP
4 services = nss, pam
5
6 [nss]
7 filter_groups = root
8 filter_users = root
9 reconnection_retries = 3
10 enum_cache_timeout = 600
11 entry_cache_nowait_percentage = 75
12
13 [pam]
14 reconnection_retries = 3
15 offline_credentials_expiration = 2
16 offline_failed_login_attempts = 3
17 offline_failed_login_delay = 5
```

```
19 # A native LDAP domain
20 [domain/LDAP]
21 timeout = 30
22 enumerate = true
23 cache_credentials = TRUE
24 debug_level = 3
25 #ignore_group_members = true
26 id_provider = ldap
27 auth_provider = ldap
28 chpass_provider = ldap
29 ldap_uri = ldaps://
30 ldap_backup_uri = ldaps://
31 ldap_search_base = dc=cm,dc=cluster
32 ldap_user_search_base = dc=cm,dc=cluster
33 ldap_group_search_base = ou=Group,dc=cm,dc=cluster
34 ldap_tls_reqcert = never
35 ldap_tls_cacert = /etc/openldap/certs/ca.pem
36 ldap_tls_cert = /etc/openldap/certs/ldap.pem
37 ldap_tls_key = /etc/openldap/certs/ldap.key
```


4. The Project Configuration



❖ 3nd. NFS

- We can switch the user from root to gsd23a01 using the CA.
- This allows access to the mounted directories only for gsd23a01 user.
 - On root user, it can not access them. (Only mounting them)
 - On gsd23a01 user, it can access them. (A owner of them)

[on root user]

```
[root@alice-kisti-hpc ~]# ll /etc/openldap/certs/
total 76
drwxr-xr-x. 2 root root 116 Oct 6 08:59 .
drwxr-xr-x. 3 root root 81 Oct 12 14:56 ..
-rw-r--r--. 1 root root 1281 Oct 6 08:58 ca.pem
-rw-r--r--. 1 root root 1704 Oct 6 08:59 ldap.key
-rw-r--r--. 1 root root 1383 Oct 6 08:59 ldap.pem
```



```
[root@alice-kisti-hpc ~]# ll /home01/gsd23a01/
ls: cannot open directory /home01/gsd23a01/: Permission denied
[root@alice-kisti-hpc ~]#
[root@alice-kisti-hpc ~]# ll /scratch/gsd23a01/
ls: cannot open directory /scratch/gsd23a01/: Permission denied
```

[on gsd23a01 user]

```
[gsdc23a01@alice-kisti-hpc gsd23a01]$ ll /home01/gsd23a01/certs
total 12
-r--r-----. 1 gsd23a01 in0138 1281 Oct 5 23:04 ca.pem
-r--r-----. 1 gsd23a01 in0138 1704 Oct 5 23:04 ldap.key
-r--r-----. 1 gsd23a01 in0138 1383 Oct 5 23:04 ldap.pem
```



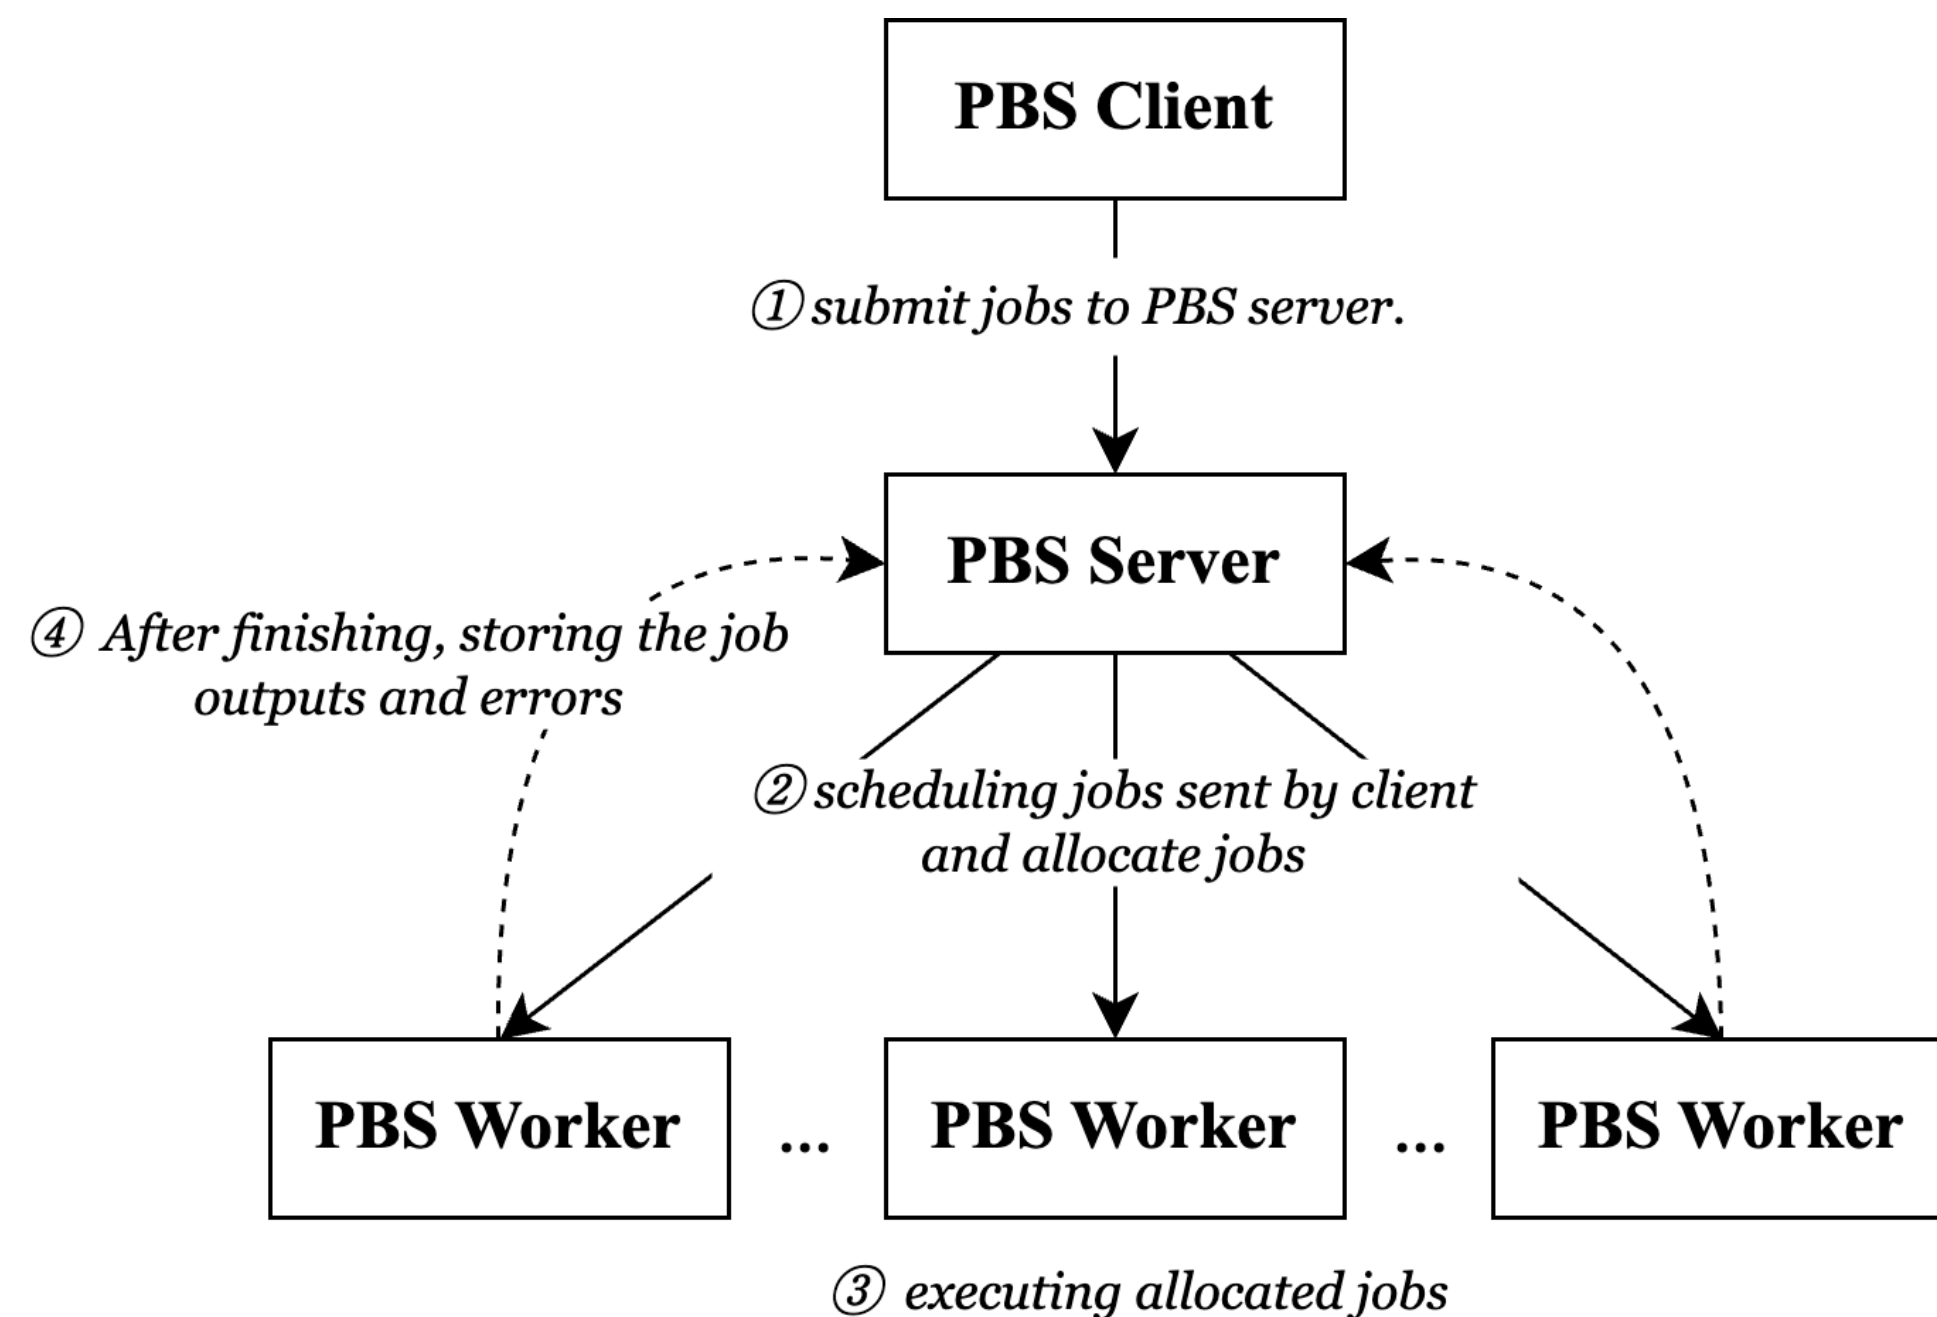
```
[gsdc23a01@alice-kisti-hpc ~]$ ll /home01/gsd23a01
total 56900
drwxr-xr-x. 2 gsd23a01 in0138 4096 Oct 5 23:04 certs
drwxr-xr-x. 2 gsd23a01 in0188 4096 May 18 14:06 job_examples
-rw-r--r--. 1 gsd23a01 in0188 58254308 Sep 26 15:52 pbspro-execution-2020.1.
3.20210315160738-0.e17.x86_64.rpm
[gsdc23a01@alice-kisti-hpc ~]$
[gsdc23a01@alice-kisti-hpc ~]$ ll /scratch/gsd23a01/
-rwxr-x---. 1 gsd23a01 in0188 328 Oct 30 11:34 cvmfs.sh
-rw-r-----. 1 gsd23a01 in0188 81 Oct 19 13:09 test.c
-rwxr-x---. 1 gsd23a01 in0188 8360 Oct 23 14:57 test.exe
```

4. The Project Configuration

❖ 4th. PBS

- PBS is a software to optimizes job scheduling and workload management.
- The project uses all types of pbs nodes.
 - * *A pbs cluster consists of 1 client, 1 server and 1 worker.*

[PBS cluster]



• PBS Client

- the interface used by user to interact with the PBS server
- submitting, monitoring jobs

• PBS Server

- the central node of managing the PBS system
- managing job scheduling, resources and job execution on the cluster

• PBS Worker

- the nodes that the actual computational work is performed

4. The Project Configuration



❖ 4th. PBS - Client Node

- PBS Client is used for submitting ALICE Grid jobs and monitoring jobs and queues.
- Whether a node functions as a PBS client, server or worker depends on the daemons started in `/etc/pbs.conf`.
 - Because of the following setting, This node becomes ‘a PBS Client node’.

[/etc/pbs.conf on PBS Client node]

```
[root@alice-kisti-hpc ~]# cat /etc/pbs.conf
PBS_EXEC=/opt/pbs
PBS_HOME=/var/spool/pbs
PBS_SERVER=dm1
PBS_PRIMARY=dm1
PBS_START_SERVER=0
PBS_START_SCHED=0
PBS_START_COMM=0
PBS_START_MOM=0
PBS_CORE_LIMIT=unlimited
PBS_SCP=/bin/scp
```

- PBS_START_SERVER
 - a daemon that manages all jobs and resources
- PBS_START_SCHED
 - a daemon that do job scheduling
- PBS_START_COMM
 - a deamon that communicate between the server and mom daemon
- PBS_START_MOM
 - a daemon that manages the local resources of the compute node

4. The Project Configuration



❖ 4th. PBS - Client Node

- Using the client node, we can monitor the status of the PBS server, workers and job queues.
- The cluster already processes a variety of job types, including data processing, simulation, parallel processing.

[The status of PBS server]

```
[root@alice-kisti-hpc ~]# qstat -B
Server          Max  Tot  Que  Run  Hld  Wat  Trn  Ext  Status
-----
pbs             0 9886 2364 1404  73   0   0   2 Scheduling
```

[The status of PBS workers]

```
vnode          state      njobs  run  susp      mem      ncpus  nmics  ngpus
-----
cpu0001        job-busy   1     1    0 187gb/187gb  0/40   0/0   0/0 13698865.pbs
cpu0002        job-busy   1     1    0 187gb/187gb  0/40   0/0   0/0 13698865.pbs
```

...

```
node8302       free       2     2    0  17gb/94gb  4/68   0/0   0/0 13690449.pbs,136
node8303       free       2     2    0  17gb/94gb  4/68   0/0   0/0 13676683.pbs,136
node8304       free       0     0    0  94gb/94gb 68/68   0/0   0/0 --
node8305       free       0     0    0  94gb/94gb 68/68   0/0   0/0 --
```

4. The Project Configuration



❖ 4th. PBS - Client Node

- Using the client node, we can monitor the status of the PBS server, workers and job queues.
- The cluster already processes a variety of job types, including data processing, simulation, parallel processing.

[The status of PBS queues]

```
[root@alice-kisti-hpc ~]# qstat -q
server: dm1
Queue          Memory CPU Time Walltime Node  Run  Que  Lm  State
-----
all_all        --    --    --    --    --    0    0  --  E R
skl_all        --    --    --    --    --    0    0  --  E R
test           --    --    999:00:0 --    --    0    0  --  E R
kdebug        --    --    48:00:00 --    --    0    0  --  D S
exclusive     --    --    --    --    --    1    0  --  E R
khoac         --    --    --    --    --    0    0  --  E R
long          --    --    120:00:0 --    --    116   0  --  E R
rescale_knl   --    --    --    --    --    0    0  --  D S
rokaf_knl     --    --    --    --    --    0    0  --  E R
khoab         --    --    --    --    --    11   0  --  E R
kiost         --    --    --    --    --    2    1  --  E R
khood        --    --    --    --    --    0    0  --  D S
khoa         --    --    --    --    --    0    0  --  D S
khoaf        --    --    --    --    --    0    0  --  D S
debug        --    --    12:00:00  2    --    0    0  --  E R
khoa         --    --    --    --    --    39   0  --  E R
rescale_skl   --    --    --    --    --    0    0  --  D S
commercial    --    --    48:00:00 --    --    3    3  --  E R
rokaf_skl     --    --    --    --    --    0    0  --  E R
norm_skl      --    --    48:00:00  80   10   16  --  E R
covids        --    --    999:00:0 --    --    0    0  --  D S
covidk        --    --    999:00:0 --    --    0    0  --  D S
eduf          --    --    48:00:00 --    --    0    0  --  E R
```

```
test1          --    --    999:00:0 --    1    0  --  E R
test3          --    --    999:00:0 --    0    0  --  D S
test5          --    --    48:00:00 --    0    0  --  E R
kerneltest     --    --    12:00:00  8    0    0  --  D S
knl_all        --    --    --    --    0    0  --  E R
flat           --    --    48:00:00 --    23   0  --  E R
normal         --    --    48:00:00 --    1035  765 --  E R
test4          --    --    999:00:0 --    0    0  --  E S
test2          --    --    999:00:0 --    0    0  --  E R
biochem        --    --    --    --    0    0  --  D S
cirnbio        --    --    --    --    100   367 --  E R
perf           --    --    999:00:0 --    0    60  --  E R
kiaps          --    --    --    --    0    0  --  E R
ocean          --    --    --    --    2    0  --  E R
samdi         --    --    999:00:0 --    6    3  --  E R
kiersaas       --    --    --    --    0    0  --  E R
snuocc         --    --    --    --    9    0  --  E R
jnuco2         --    --    --    --    2    0  --  E R
perf2          --    --    999:00:0 --    0    0  --  E R
march         --    --    --    --    1    0  --  E R
samdi2         --    --    --    --    12   0  --  E R
kierbk         --    --    --    --    4    0  --  E R
youth          --    --    --    --    0    0  --  E R
kierhy         --    --    --    --    0    0  --  E R
alice          --    --    --    --    0    0  --  E R
-----
1377  1215
```

4. The Project Configuration



❖ 5th. PBS - Worker Node

- PBS worker node is NURION, the 5th supercomputer in South Korea, computational node.
- We test to submit a job that mounts `alice.cern.ch` repository and runs a `lienv` script.
 - To verify if the worker node is ready to execute ALICE Grid jobs.

[a job script]

```
[gsdc23a01@alice-kisti-hpc gsdc23a01]$ cat cvmfs.sh
#!/bin/sh
#PBS -V
#PBS -N cvmfs_debug_job
#PBS -q alice
#PBS -A etc
#PBS -l select=1:ncpus=1:mpiprocs=1:ompthreads=1
#PBS -l walltime=01:00:00

cd $PBS_O_WORKDIR

module purge
module load craype-mic-knl intel/18.0.3

source /cvmfs/alice.cern.ch/etc/login.sh
ALIENV_DEBUG=1 alienv printenv VO_ALICE@AliPhysics::VAN-20231029_02-1
```

`qsub cvmfs.sh`

[the result files of the job]

```
[gsdc23a01@alice-kisti-hpc gsdc23a01]$ ll
total 28
-rw-----. 1 gsdc23a01 in0188  1078 Oct 30 11:42 cvmfs_debug_job.e13756203
-rw-----. 1 gsdc23a01 in0188 16407 Oct 30 11:42 cvmfs_debug_job.o13756203
-rwxr-x---. 1 gsdc23a01 in0188   328 Oct 30 11:34 cvmfs.sh
```


4. The Project Configuration



❖ 6th. VOBOX

- VOBOX is a system to support ALICE VO services.
- Before configuring VOBOX, we set port and source rules in `firewalld` for CERN.

[Network setting guide]

Network

The following network connectivity is expected for the VO-Box services:

Port	Access	Service
1093	TCP Incoming from the World	MonALISA FDT server, SE tests
8884	UDP Incoming from your site WN and your site SE nodes	Monitoring info
9930	UDP Incoming from your site SE nodes	XRootD metrics
	ICMP Incoming and Outgoing	Network topology for file placement and access

In the future, these extra services **may** be needed:

Port	Access	Service
8098	TCP Incoming from your site WN	JAliEn/Java Serialized Object stream
8097	TCP Incoming from your site WN	JAliEn/WebSocketS

Protocol	IP Range	Note
IPv4	128.141.0.0/16	
	128.141.25.192/26	<- part of Central Services are here
	128.141.26.0/26	<- part of Central Services are here
	128.142.0.0/16	
	128.142.249.0/24	<- part of Central Services are here
	137.138.0.0/16	<- part of Central Services are here
	188.184.0.0/15	<- part of Central Services are here
	185.249.56.0/22	
	192.65.196.0/23	
	192.91.242.0/24	
IPv6	194.12.128.0/18	
	2001:1458::/32	
	2001:1458:301:54::/64	<- part of Central Services are here
	2001:1459::/32	

4. The Project Configuration



❖ 6th. VOBOX

- After installing `wlcg-vobox`, we configure `site-info.def`, `users.conf`, `groups.conf`.
 - `site-info.def` - the main configuration file written YAIM
 - `users.conf` - a file defining the users to be created on the service nodes that need them
 - `groups.conf` - a file defining the user categories that must be accepted by the grid services provided by a site

[site-info.def]

```
[root@alice-kisti-hpc ~]# cat site-info.def
GROUPS_CONF=/opt/glite/yaim/etc/groups.conf
USERS_CONF=/opt/glite/yaim/etc/users.conf

SITE_NAME=KR-KISTI-GSDC-01

VOBOX_HOST=`hostname -f`
WMS_HOST=rocwms01.grid.sinica.edu.tw
PX_HOST=myproxy.cern.ch
BDII_HOST=lcg-bdii.cern.ch

#SE_LIST=alice-t1-se.sdfarm.kr
SE_LIST=my-se.my-domain

#VOS="alice dteam ops"
VOS="alice"

VO_ALICE_SW_DIR=.
VO_ALICE_DEFAULT_SE=my-se.my-domain
VO_ALICE_VOMS_SERVERS="'vomss://voms2.cern.ch:8443/voms/alice?
/alice/' 'vomss://lcg-voms2.cern.ch:8443/voms/alice?/alice/' "
VO_ALICE_VOMSES="'alice lcg-voms2.cern.ch 15000 /DC=ch/DC=cern
/OU=computers/CN=lcg-voms2.cern.ch alice 24' 'alice voms2.cern
.ch 15000 /DC=ch/DC=cern/OU=computers/CN=voms2.cern.ch alice 2
4'"
VO_ALICE_VOMS_CA_DN="'/DC=ch/DC=cern/CN=CERN Grid Certificatio
n Authority' '/DC=ch/DC=cern/CN=CERN Grid Certification Author
ity'"
```

[users.conf]

```
[root@alice-kisti-hpc ~]# tail /opt/glite/yaim/etc/users.conf

14320:ali1_120:14200:alicet1:alice::
14321:ali1_121:14200:alicet1:alice::
14322:ali1_122:14200:alicet1:alice::
14323:ali1_123:14200:alicet1:alice::
14324:ali1_124:14200:alicet1:alice::
14325:ali1_125:14200:alicet1:alice::
14326:ali1_126:14200:alicet1:alice::
14327:ali1_127:14200:alicet1:alice::
14328:ali1_128:14200:alicet1:alice::
100018801:gsdc23a01:1000188:in0188:alice:sgm:
```

[groups.conf]

```
[root@alice-kisti-hpc ~]# cat /opt/glite/yaim/etc/groups.conf
"/alice/ROLE=lcgadmin":::sgm:
"/alice/ROLE=production":::prd:
"/alice/ROLE=pilot":::pilot:
"/alice":::
```

4. The Project Configuration

❖ 6th. VOBOX

- When we execute `/opt/glite/yaim/bin/yaim -c -s site-info.def -n VOBOX > ~/yaim-result.log`, the grid-mapfile is generated by the three aforementioned files.
- For the HPC security, we leave only a few map list related to real access users.

[grid-mapfile (old ver.)]

```
[root@alice-kisti-hpc ~]# cat /etc/grid-security/grid-mapfile.old  
"/C=BR/O=ANSP/OU=ANSPGrid CA/OU=People/CN=Christian Reckziegel" .ali1_  
"/C=BR/O=ANSP/OU=ANSPGrid CA/OU=People/CN=Fabio Padoa" gsd23a01  
"/C=BR/O=ANSP/OU=ANSPGrid CA/OU=People/CN=Marcelo Carneiro Munhoz" .ali1_
```

...

```
"/O=GRID-FR/C=FR/O=CNRS/OU=SUBATECH/CN=Philippe Pillot" .ali1_  
"/O=GRID-FR/C=FR/O=CNRS/OU=SUBATECH/CN=Pierrick Le Corre" gsd23a01  
"/O=GermanGrid/OU=GSI/CN=Jens Wiechula" .ali1_
```



[grid-mapfile (new ver.)]

```
[root@alice-kisti-hpc ~]# cat /etc/grid-security/grid-mapfile  
"/C=KR/O=KISTI/O=KISTI/CN=38191474 Sang Un Ahn" gsd23a01  
"/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=grigoras/CN=650724/CN=Costin Grigoras" gsd23a01  
"/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=lbetev/CN=374855/CN=Latchezar Betev" gsd23a01  
"/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=maarten/CN=410032/CN=Maarten Litmaath" gsd23a01
```

5. Future Plans

❖ (Key point) Stabilization of Project Environment for Real ALICE Grid job processing

- ❑ Registering VOBOX nodes with CERN
- ❑ ALICE Grid job submission Test
- ❑ Next topic discussion through contact with ALICE computing



Thank you!

Hyeon-Jin Yu

hyeonjin.yu@cern.ch