


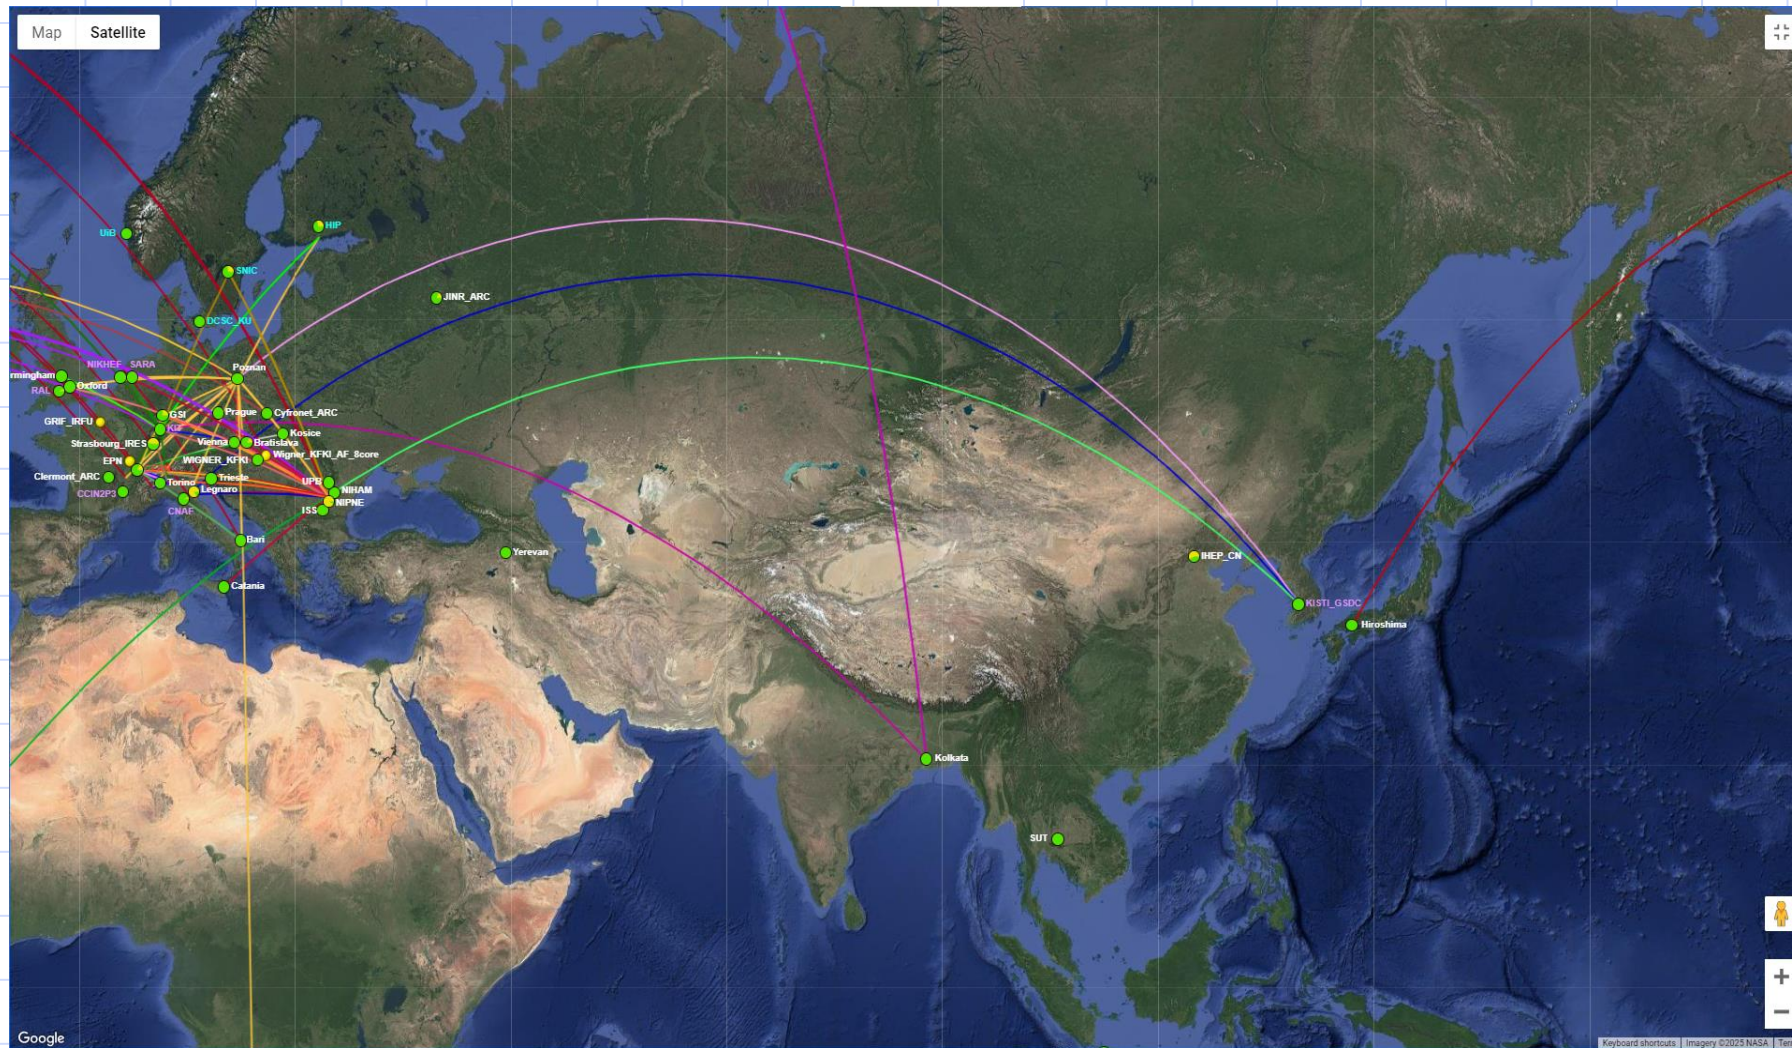
Kolkata Tier 2 @ ALICE

EHEP&AG → Experimental High Energy Physics and Application Group,
VECC → Variable Energy Cyclotron Centre, Kolkata, West Bengal
DAE → Department of Atomic Energy. Govt of India.

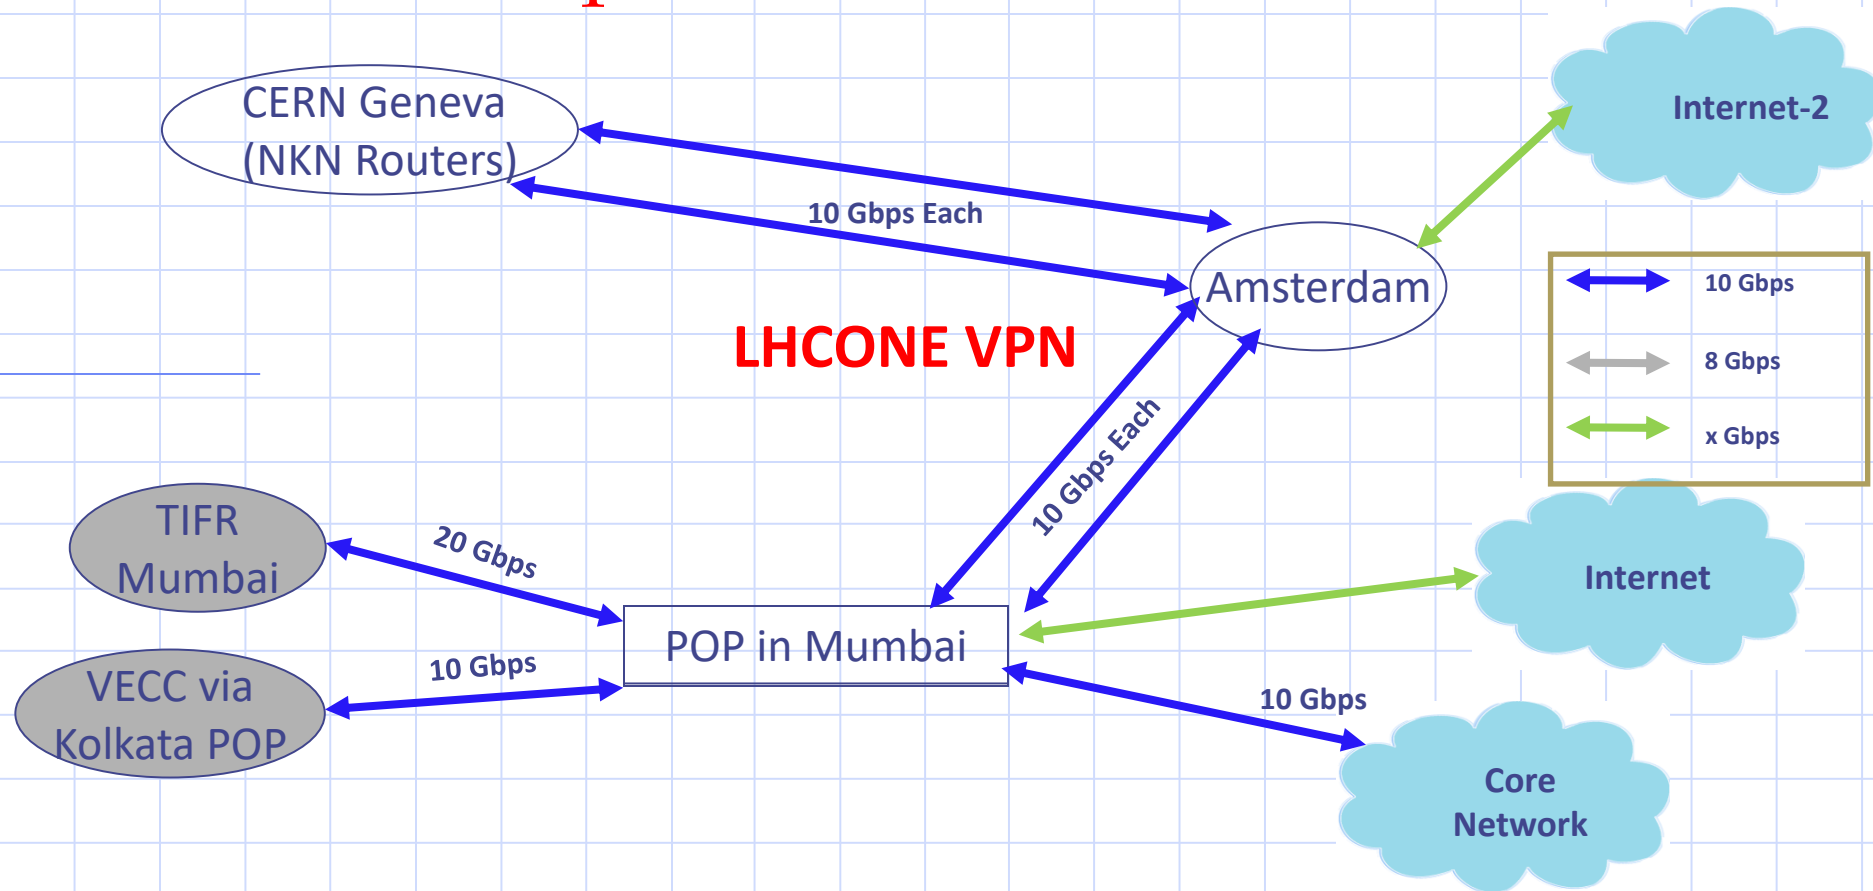
Tier-2 Site for the WLCG
Kolkata Tier-2 For ALICE GRID
GOCDDB:- IN-DAE-VECC-02

Team:-
Vikas Singhal
Prasun Singh Roy + team of
two support persons
(changing)

Dr. Arup Bandyopadhyay
(Heading the Group)




10 Gbps NKN WAN Network


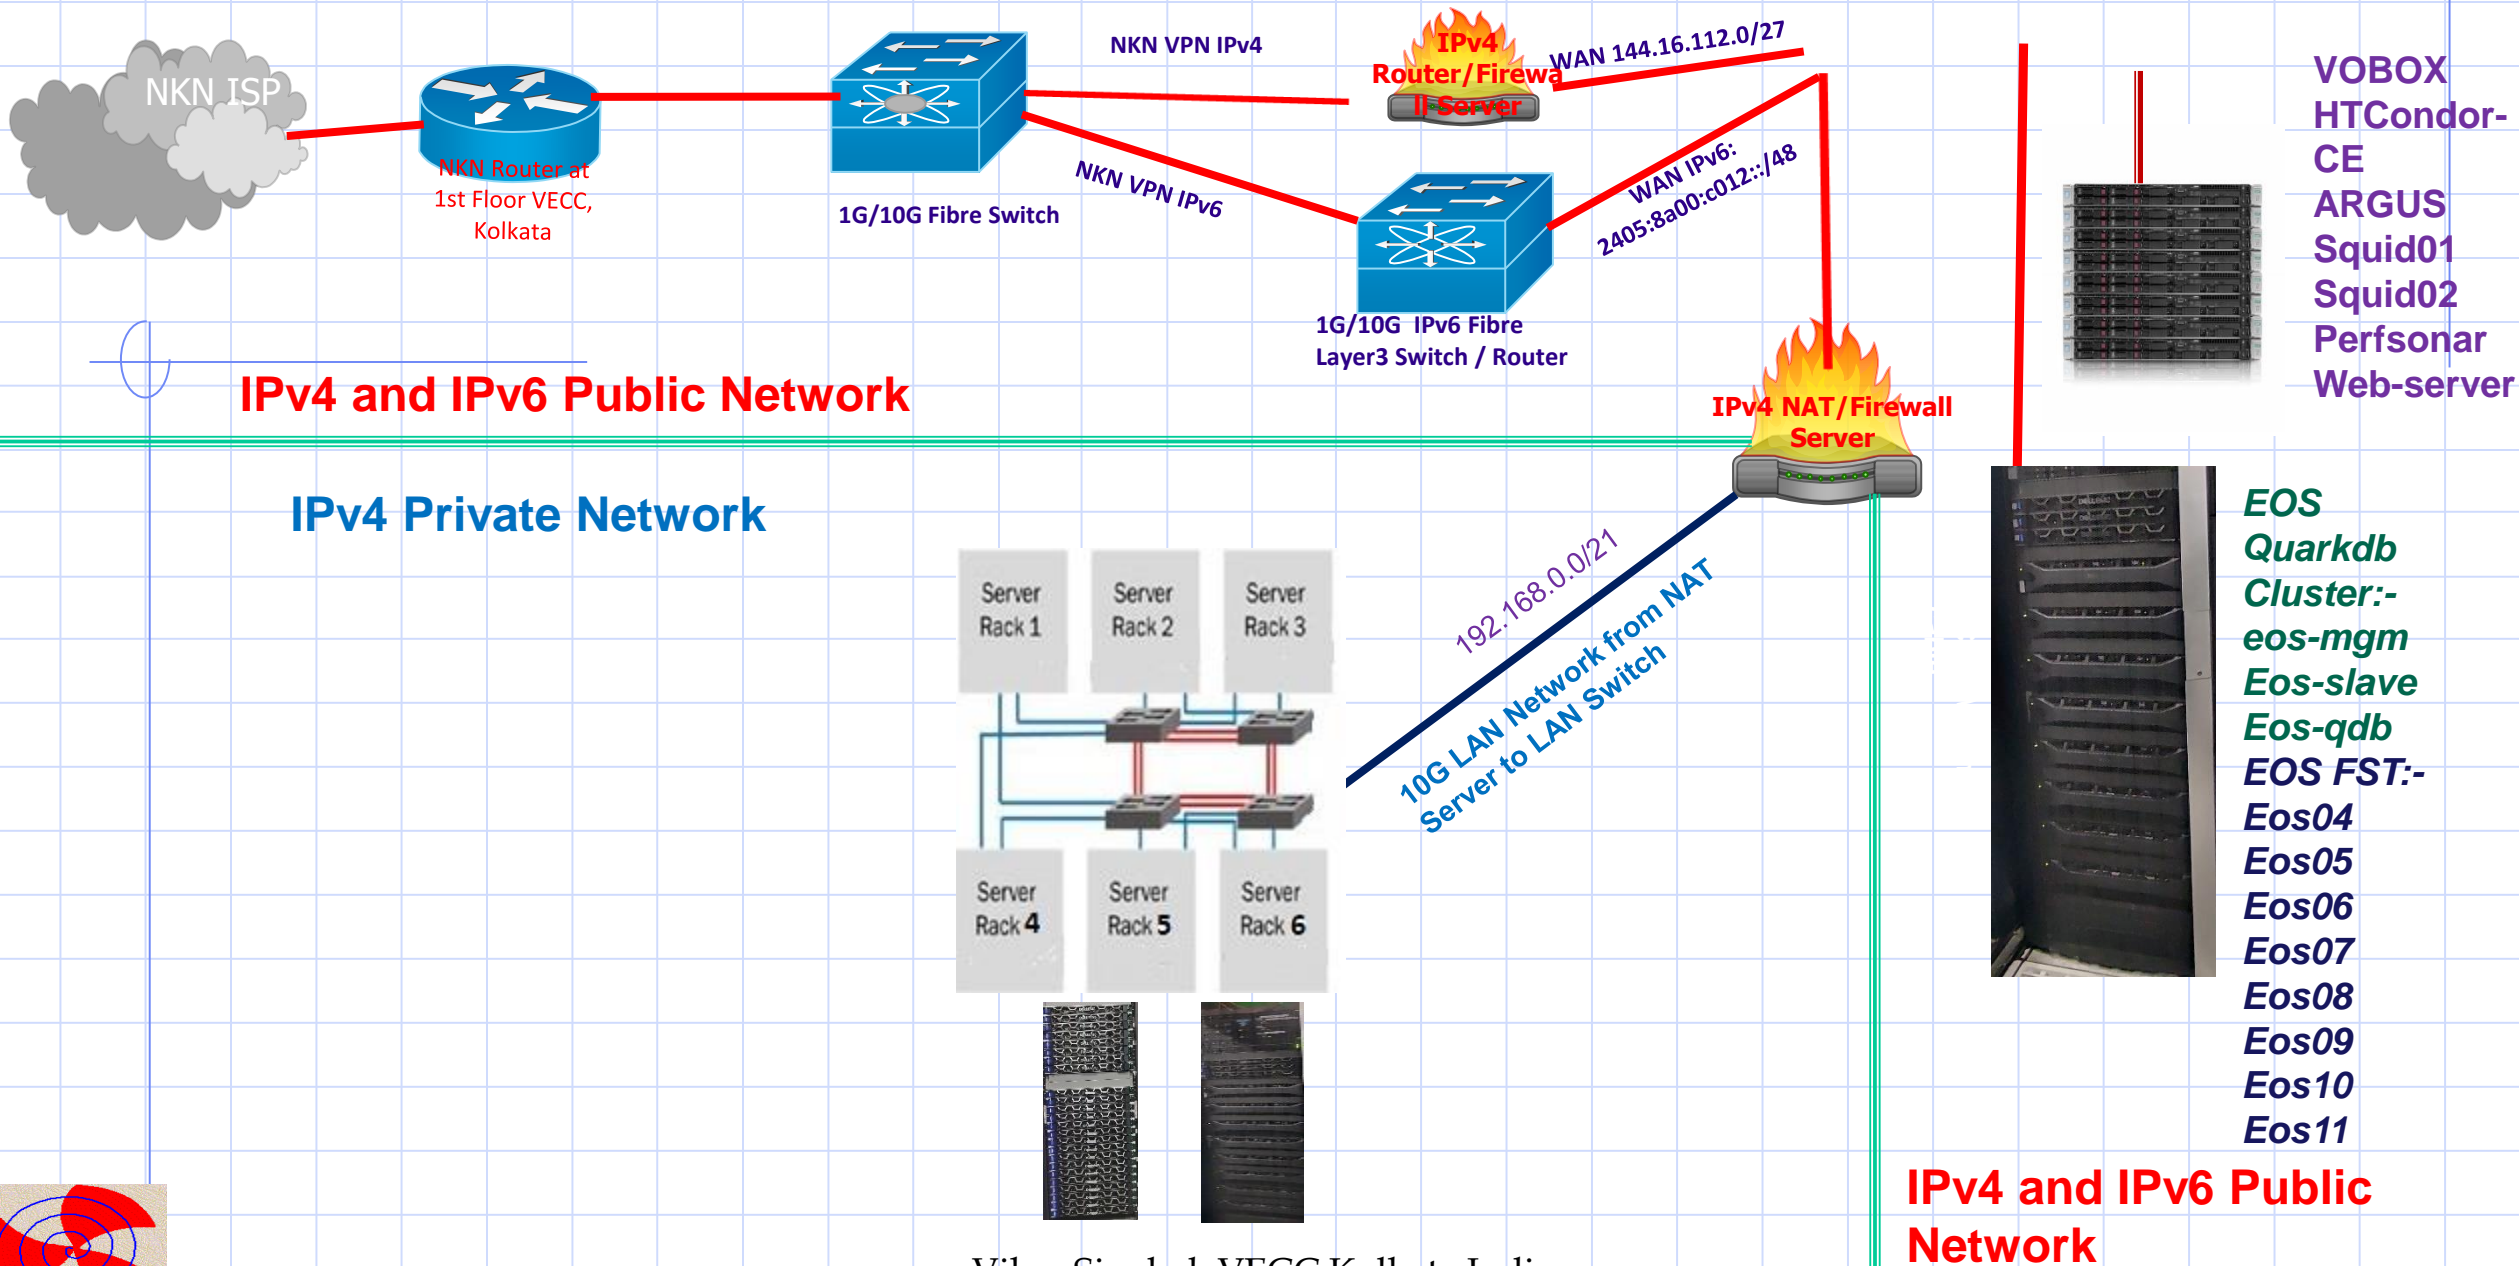


- NKN Routers are hosted at CERN Geneva.

Consistent support from NKN, India for End to End Network provisioning.



NKN Router to Grid Computing Facility



Backend Components and Computing Hardware

- Domain name:-
 - Public : tier2-kol.res.in (ISP: ERNET)
 - Private: internal.tier2-kol.res.in
- IP Range:-
 - IPv4 Range: 144.16.112.0/27 (ISP: ERNET)
 - IPv6 Range: 2405:8a00:c012::/48 (ISP: NKN)
- DNS:
 - Public: naamak.tier2-kol.res.in; suchak.tier2-kol.res.in
 - Private: installer.internal.tier2-kol.res.in
- Local Resources:
 - Batch system : HTCondor 24.0.12/ 23.0.28 and HTCondor-CE 24.0.2 (Local Lan)
 - a. Submitter and CE: kolkata-condor-ce.tier2-kol.res.in
 - b. Central Manager: condor-cm.internal.tier2-kol.res.in
- Storage: EOS Diopside version 5.3.17
- Networking : 10Gbps LAN and WAN over Fibre

Kolkata Tier-2 Worker-Nodes:-

Total No. of Worker-nodes: 56

Total No. of Threads: 2944

Two types of Processors

Processor Used: 14 Core Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz


RAM 128 GB DDR4.

HDD Type : SSD (980GB)

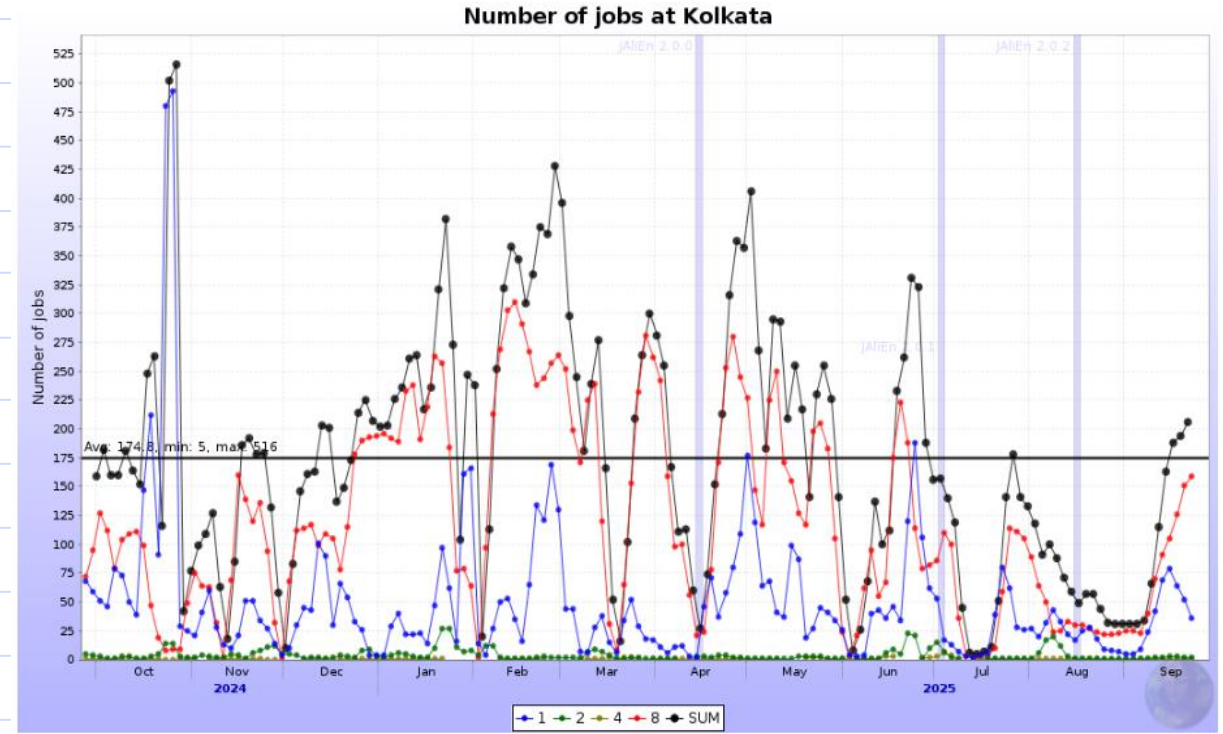
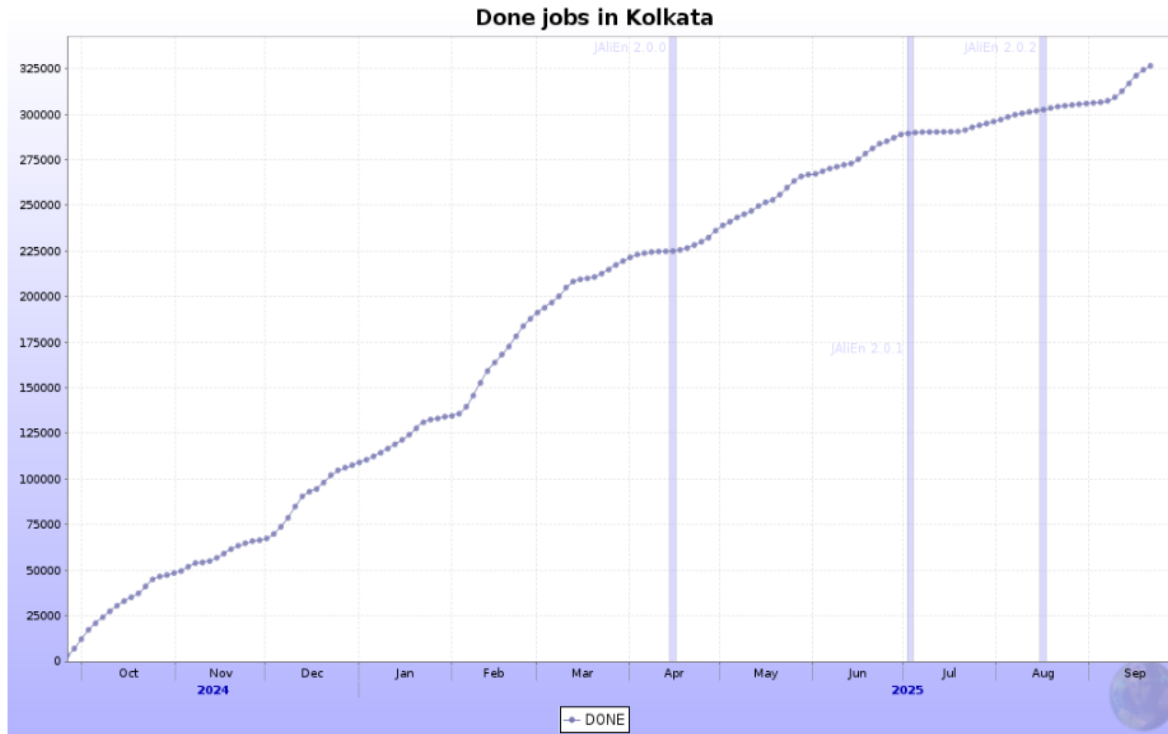
12 Core Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz.

RAM 128 GB DDR4.

SAS 12G SFF 4 X 600 GB with RAID5 and one hot spare.




ALICE Jobs at Kolkata Tier-2 since the last ATCF



325K (8 core) ALICE Jobs successfully completed during the last one year.

Switched to whole node jobs

Vikas Singhal, VECC Kolkata India



Upgraded from CentOS7 to AL9 and HTCondor Middleware and Packages

- ❑ Migrated Entire Computing Element to Alma Linux 9

- ❑ Computing Element: HTCondor-CE version v24.0.2
(Kolkata-condor-ce.tier2-kol.res.in)

- ❑ Accounting :

 - APEL-Parser – v1.9.1 (Kolkata-condor-ce.tier2-kol.res.in)


 - APEL- Client – v1.9.1 (argus.tier2-kol.res.in)

- ❑ BDII – v5.2.26 (argus.tier2-kol.res.in)

- ❑ VOBOX – wlcg-voms, vobox (htcvobox.tier2-kol.res.in)

- Thanks: Maarten Litmaath helping at each level.

- Thanks to HTCondor team (mailing list) for resolving domain name related issue.




Brief history of EOS storage at Kolkata Tier2

Name of Instance: ALICE::Kolkata::EOS
Space: 141 TB
No. of Management Server: 1
No. of FST Disk Server: 3 Nos.
No. of Disk and Capacity: 12 Nos (each 4TB)
Disk Type: NLSAS
Total Nos. of Disks: 36
H/W RAID type: 0 (Zero)
OS: Scientific Linux 6

EOS version: OS AQUAMARINE (0.3.256)
RAIN type (Redundant Array of Independent nodes) : RAIN6
No. of Disk Group: 6
Started on 27/03/2017

```
[root@eos /]# eos -b
EOS Console [root://localhost] |/> space ls
#-----
# type # name # groupsize # groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity)
#capacity(rw) #nom.capacity #quota #balancing # threshold # converter # ntx # active #intergroup
#-----
spaceview default 0 24 36 36 52.16 T 141.74 T 141.74
T 0 off on 20 off 2 0 off
EOS Console [root://localhost] |/> node ls
#-----
# type # hostport # geotag # status # status # txgw #gw-queued # gw-ntx #gw-
rate # heartbeatdelta #nofs
#-----
nodesview eos01.tier2-
kol.res.in:1095 online on off 0 10 120 2 12
nodesview eos02.tier2-
kol.res.in:1095 online on off 0 10 120 2 12
nodesview eos03.tier2-
kol.res.in:1095 online on off 0 10 120 2 12
EOS Console [root://localhost] |/>
```

```
[root@eos ~]# eos -b group ls
#-----
# type # name # status #nofs #dev(filled) #avg(filled) #sig(filled) #balancing # bal-shd #drain-shd
#-----
groupview default.0 on 6 0.08 94.88 0.05 idle 0 0
groupview default.1 on 6 0.12 94.87 0.07 idle 0 0
groupview default.2 on 6 0.13 94.87 0.07 idle 0 0
groupview default.3 on 6 0.15 94.86 0.08 idle 0 0
groupview default.4 on 6 23.82 58.87 11.22 balancing 0 0
groupview default.5 on 6 0.09 94.88 0.06 idle 0 0
#-----
```



Expanded with In-Memory Namespace during 2018

- ❖ Name of Instance: ALICE::Kolkata::EOS2
- ❖ Date of Production: 16th February 2018
- ❖ Space: 1.1PB.
- ❖ No. of Management Server: 2 (Master/Slave)
- ❖ No. of Disk Server: 7Nos.
- ❖ Each Disk Capacity in Disk Server: 10TB
- ❖ Hard Disk Type: NLSAS
- ❖ Nos. of Disk in each disk server: 16
- ❖ H/W RAID type: 0 (Zero)
- ❖ OS: CentOS 7.x
- ❖ EOS version: EOS CITRINE (4.4.23)
- ❖ RAIN type (Redundant Array of Independent nodes) : RAIN6

Hardware for Management Server (Manager) and QuarkDB


- ❖ 3 nos. of 1U Rack Mount Servers. Each server contains:-
 - 2* Intel Xeon Processor Silver 4112 @ 4C 2.60 Ghz Processors
 - 64GB RAM
 - 1 * 480 GB SSD and 4*1.2TB SAS HDD.
 - 2 * 10GB Ethernet (Fibre)
 - RAID: Server are configure with 2 types of RAID:-
 - RAID-5 in 3 nos of 1.2TB SAS + one hot spare.
 - RAID-0 in 480 GB SSD.

Hardware for Disk Server (FST)

- ❖ 7 nos. of Dell PowerEdge R730XD servers. Each server contains:-
 - 2* Intel Xeon E5-2630 v4 @ 2.20GHz 10Core Processors.
 - 16*10TB NLSAS HDD and 2*480GB SSD.
 - 4 *10GB Ethernet (Fibre).
 - 128GB DDR4 RAM.
 - RAID: Server are configure with 2 types of RAID:-
 - RAID-0 in 16 nos 10TB NLSAS HDD.
 - RAID-1 in 2 nos 480 GB SSD.

Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Kolkata::EOS	OK	128.9 TB	26.94%	94.17 TB	34.73 TB	475.3 K	FILE	FAIL
ALICE::Kolkata::EOS2	OK	1.1 PB	6.223%	1.032 PB	70.09 TB	1.324 M	FILE	OK
ALICE::Kolkata::SE	OK	76.39 TB	32.14%	51.84 TB	24.55 TB	458.7 K	FILE	OK




Challenges faced during upgradation of EOS

Issue-1: Fluctuating Mountpoint:- After OS upgaration i.e. AlmaLinux 9.5 and mounting the 16 nos. of 10TB NLSAS disk (configured as RAID-0), mount point label were fluctuating with each reboot.

Before

```
#####  
/dev/sdb1 /xdata0 xfs  
/dev/sdc1 /xdata1 xfs  
/dev/sdd1 /xdata2 xfs  
/dev/sde1 /xdata3 xfs  
/dev/sdf1 /xdata4 xfs  
/dev/sdg1 /xdata5 xfs  
/dev/sdh1 /xdata6 xfs  
/dev/sdi1 /xdata7 xfs  
/dev/sdj1 /xdata8 xfs  
/dev/sdk1 /xdata9 xfs  
/dev/sdl1 /xdata10 xfs  
/dev/sdm1 /xdata11 xfs  
/dev/sdn1 /xdata12 xfs  
/dev/sdo1 /xdata13 xfs  
/dev/sdp1 /xdata14 xfs  
/dev/sdq1 /xdata15 xfs  
#####
```




After

```
/dev/sdn1 9.0T 7.8T 1.2T 88% /xdata12  
/dev/sdm1 9.0T 7.8T 1.2T 88% /xdata11  
/dev/sdh1 9.0T 7.7T 1.3T 87% /xdata4  
/dev/sdg1 9.0T 7.7T 1.3T 86% /xdata9  
/dev/sdl1 9.0T 7.8T 1.2T 87% /xdata10  
/dev/sdo1 9.0T 7.8T 1.2T 88% /xdata13  
/dev/sda1 9.0T 7.7T 1.3T 86% /xdata1  
/dev/sdc1 9.0T 7.8T 1.2T 88% /xdata5  
/dev/sdk1 9.0T 7.7T 1.3T 87% /xdata2  
/dev/sdb1 9.0T 7.8T 1.2T 88% /xdata3  
/dev/sdq1 9.0T 7.9T 1.1T 88% /xdata15  
/dev/sdf1 9.0T 7.7T 1.3T 87% /xdata7  
/dev/sdp1 9.0T 7.7T 1.3T 87% /xdata14  
/dev/sdj1 9.0T 7.7T 1.3T 87% /xdata8  
/dev/sdi1 9.0T 7.7T 1.3T 86% /xdata6  
/dev/sde1 9.0T 7.4T 1.6T 83% /xdata0
```

Resolved using WWN

- First tried using using UUID (Universally Unique Identifier)
- Did not work then using WWN (World Wide Name)) identifier.
- Hard-coded into the fstab

```
#####  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b43024391201-part1 /xdata8 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c0027ad1cfc1aa2dabc-part1 /xdata6 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b3bf1d7f0373-part1 /xdata0 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c0028d7027730b56812-part1 /xdata10 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b3cd1e53e260-part1 /xdata1 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b441253a620b-part1 /xdata9 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c0028d6fe1512aa4a92-part1 /xdata12 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b3da1f1389ac-part1 /xdata2 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c0028d6fdff115d8bee-part1 /xdata11 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b3e71fe28bc8-part1 /xdata3 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b3f620bccb6f-part1 /xdata4 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b4922a0a7aff-part1 /xdata14 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b40521a7e4e4-part1 /xdata5 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b422235cdf2c-part1 /xdata7 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c002317b4a22aff57f5-part1 /xdata15 xfs defaults 0 0  
/dev/disk/by-id/wwn-0x6d094660680f5c0028d70290323d7bd9-part1 /xdata13 xfs defaults 0 0  
#####
```




Challenges faced during upgradation of EOS

Issue-2: After resolving the disk inconsistency / mapping issue

no xattr for file:- LevelDB did not worked with the latest version of EOS FST i.e. v5.3.12.

Xattrs need to be generated for each file systems and drop LeveDB.

- Downgraded EOS version v5.3. → v5.1.28 on all MGM and FSTs.
- “fstofs.filemd_handler attr” added in fst config file
- During OS upgrdation process on FST, we missed to take a backup /var/eos/md/fmd.00XX.LevelDB directories, therefore regenerated fmd.00XX.LevelDB directories of each fs for every FST.
- Migrated LevelDB to Xattr.
- Again upgraded to the eos version to 5.3.17 on all FST and mgm/quarkdb servers and restarted the services.
- Now ALICE::Kolkata::EOS2 is Up (Still running with single manager, as HA is not working).



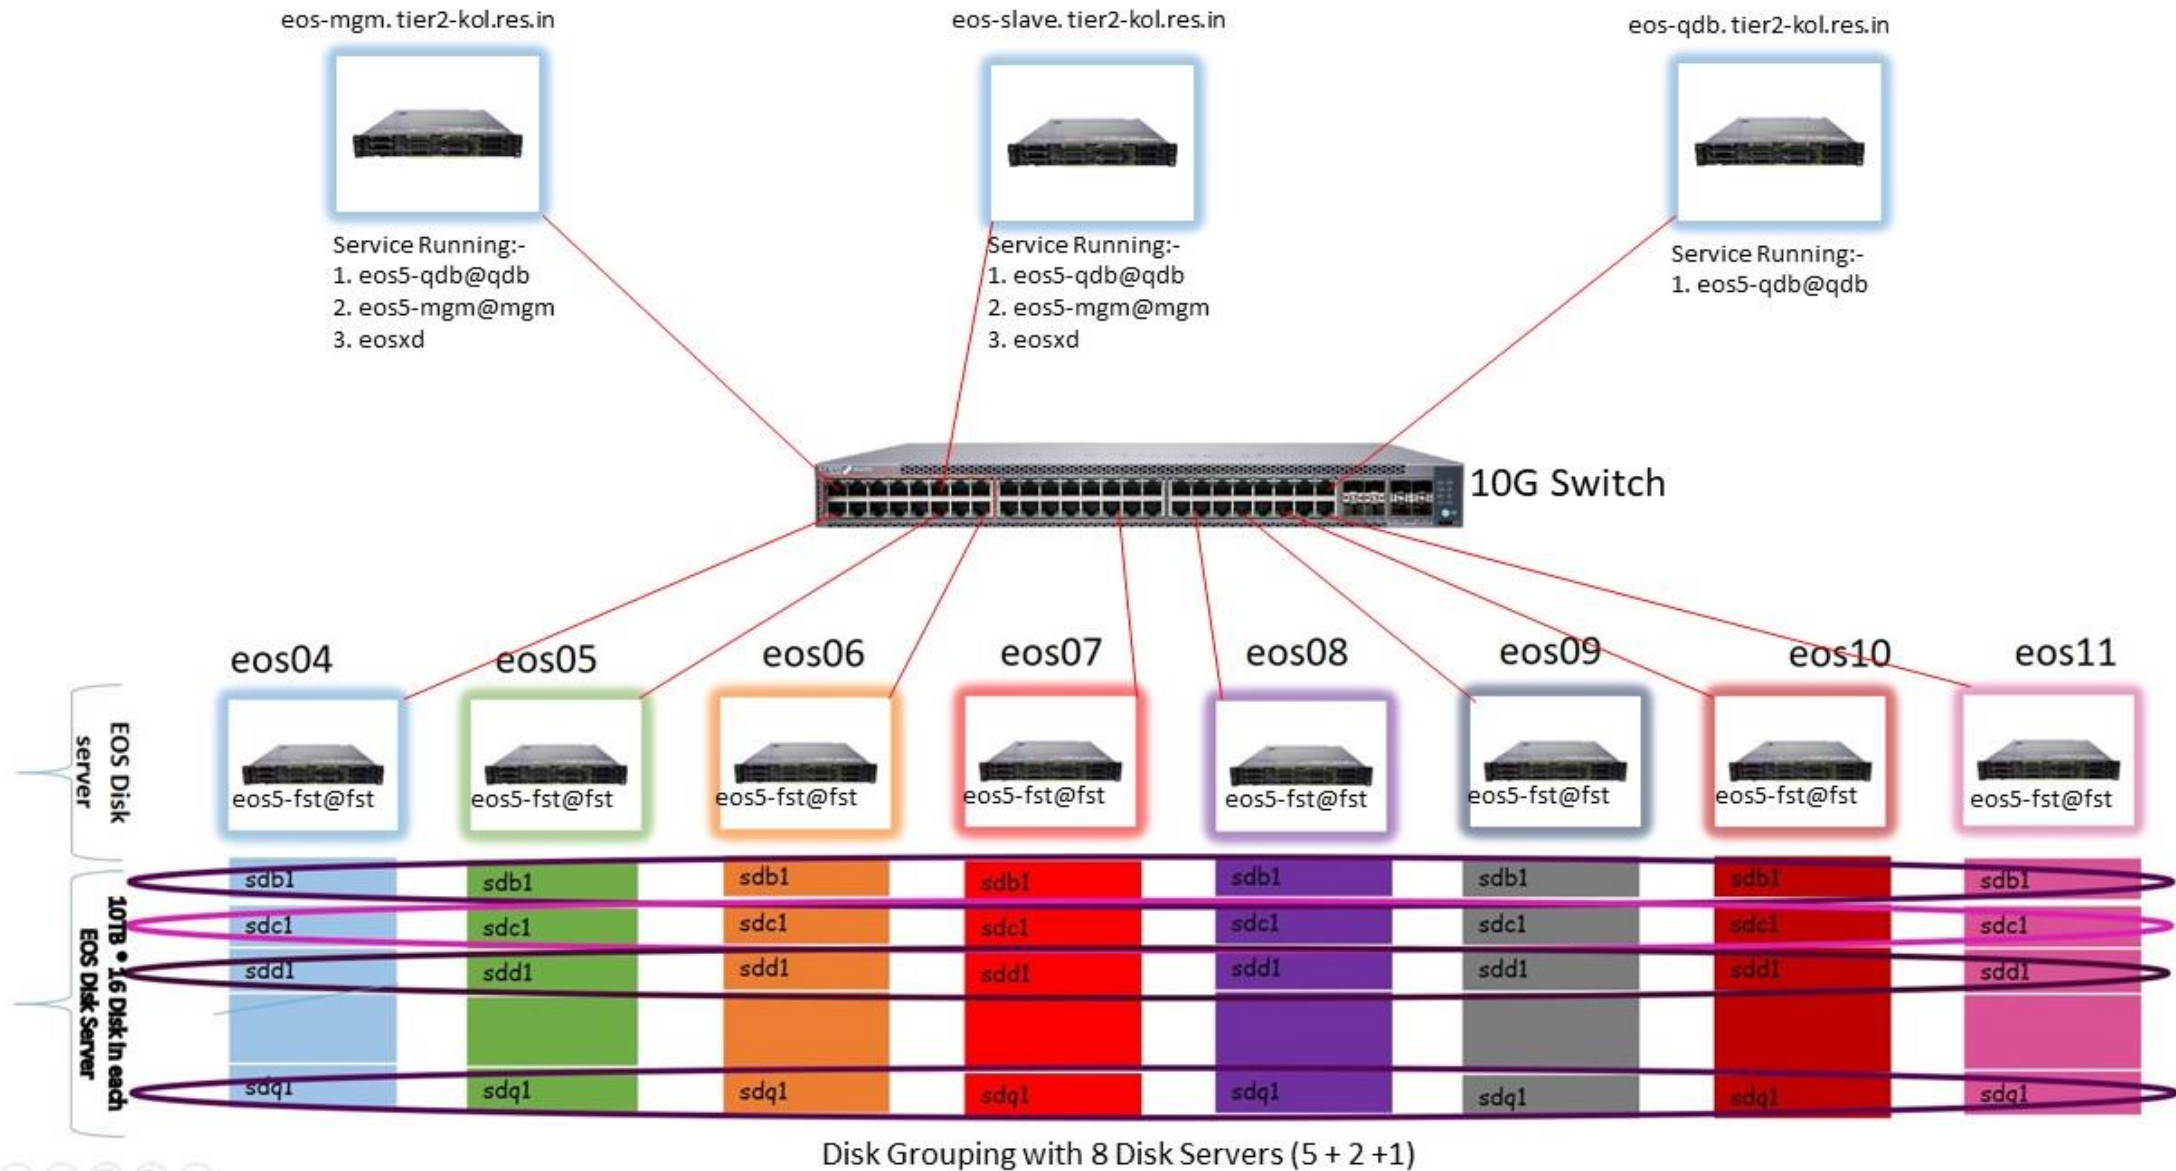
Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Kolkata::EOS2	OK	1008 TB	52.31%	480.7 TB	527.3 TB	9.839 M	FILE	OK


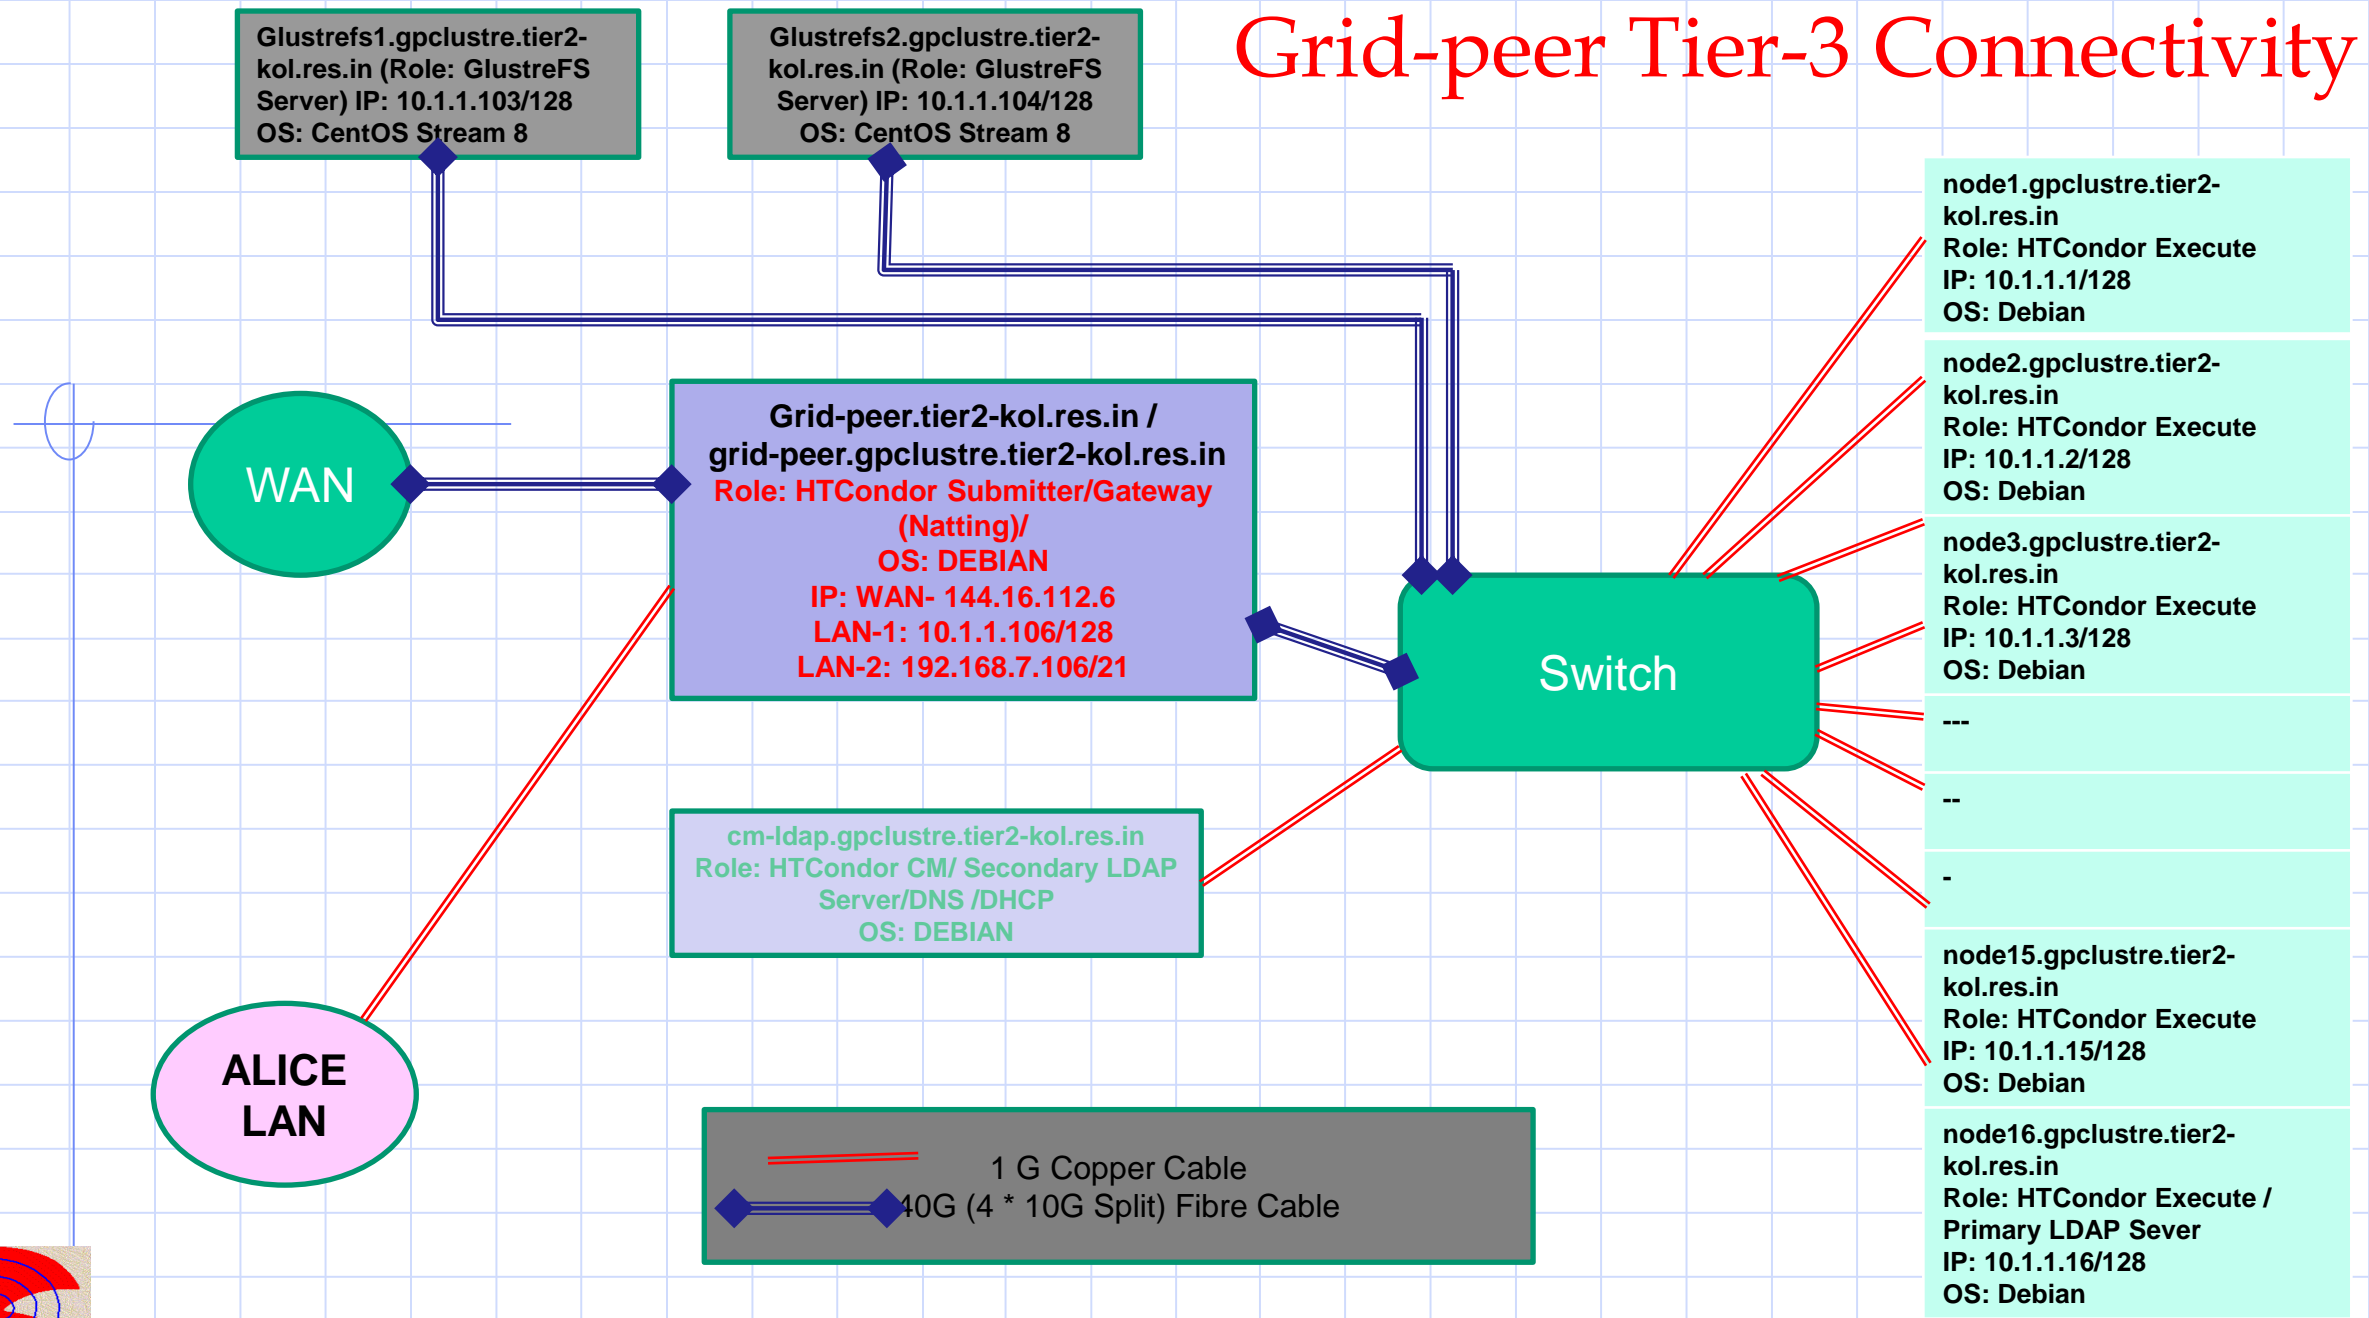
EOS Developers (team) helped a lot. Special thanks to “Elvin Alin Sindrilaru” and “Andreas Joachim Peters”.

Vikas Singhal, VECC Kolkata India

Present logical diagram of Kolkata Tier2 Storage Element (EOS)



Grid-peer Tier-3 Connectivity



Grid-Peer Tier3 Cluster Status



Grid-Peer Head Node
CM-Idap

Interactive Nodes

Non-interactive Nodes

Storage

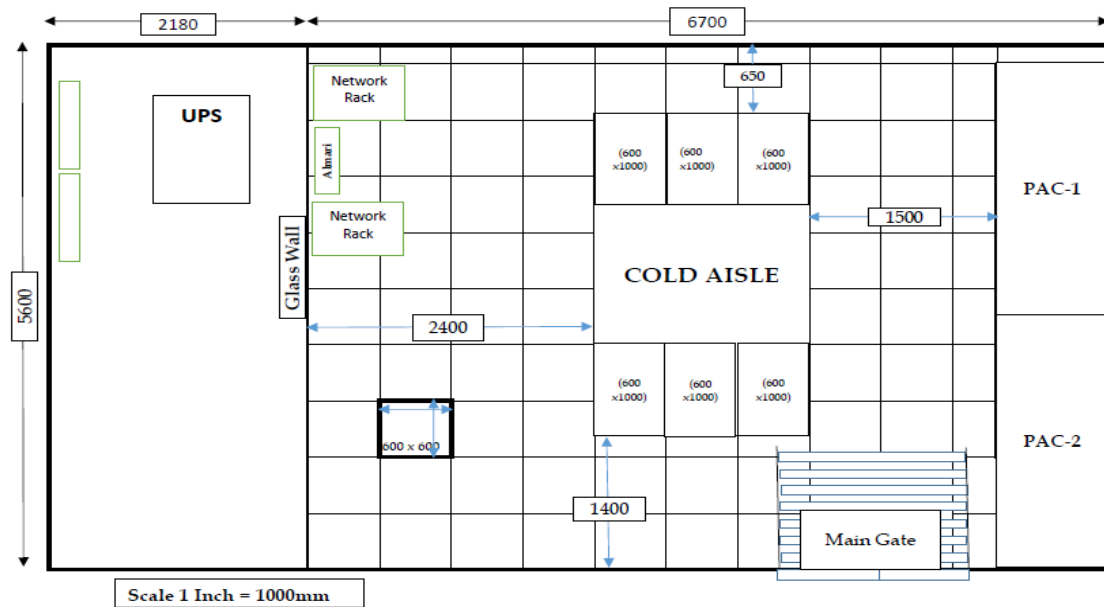
- Total 624 Threads (48 Threads/server).
- Extensively used by VECC users and ALICE INDIA/CBM INDIA Collaborators. 75+ active users (across India)
- Grid-Peer: Head Node/ HTCondor Submitter.
- CM-Idap: Ldap Authentication Server + HT Condor CM
- Interactive Nodes (Grid-peer) : 4 Nos. @48 threads & 128GB RAM.
- Non-interactive Nodes(HT-Condor Execute Nodes): 9 Nos. @48 threads & 128GB RAM.
- Network Connected : 10G Fibre.
- OS debian 11.3 (64bit).
- Two factor authentication Login (1-Key-Based Authentication, 2-PasswordAuthentication,).
- Access to Grid-Peer Head-Node by using LAN & WAN(IPv4).
- gcc version 10.2.1 20210110 & condor version 9.12.0-1.1

Storage:

- Glusterfs Storage: 129 TB Usable Space. (Replica-2 Type)
- Total space 216TB (18Hard disk) 15 hard disk usable and 3 Global hot spare and Configuring RAID 5
- Glusterfs Storage are installed with Centos Stream 8.
- Users home directory of interactive nodes are mounted by NFS through Glusterfs storages.


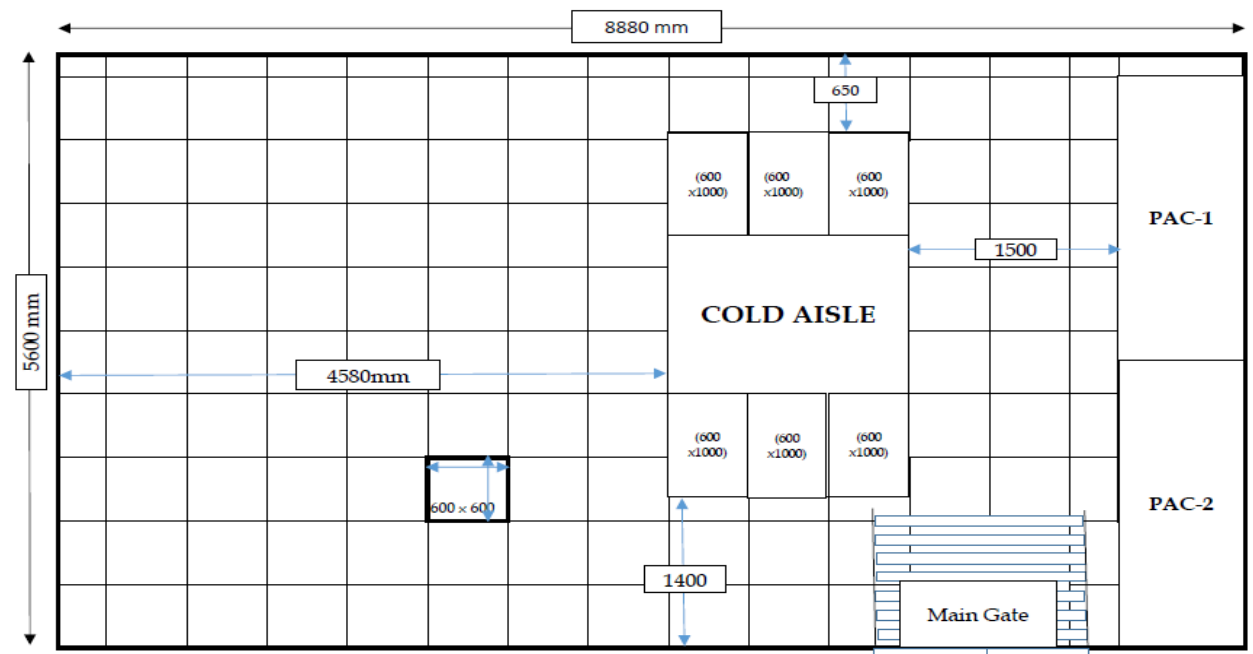
Monitoring under development

Rearranging the GRID Computing Facility, VECC




- Did not get more physical space, therefore rearranging the facility.
- Efficiently need to utilize remaining space in the Room.
- New solution for at-least an another decade.
- UPS room is renovated and modified.

- Started to removing of UPS and other passive components.
- An small UPS room is renovated and modified.




Rearranging the GRID Computing Facility, VECC



An Indent for 100kW of InRow cooling Solution has been initiated and Internal Technical Evaluation is going on for the received seven bids.

Vikas Singhal, VECC Kolkata India



Scientific Computing Workshop for High Energy Physics and Tomography




➤ A Colloquium titled “Building International Research Software for Next-Generation High Energy Physics”
by Prof. Peter Elmer, Executive Director, Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP)

- More than 60 participants
- From 16 different institutes
- 11 Instructors
 - 8 for HEP Software
 - 3 for Muon Tomography
- Full interactive sessions with hands on.

December 16-20, 2024 at VECC, Kolkata


<https://indico.cern.ch/event/1422314/>

Vikas Singhal, VECC Kolkata India



GRID India Monthly Meeting and Status

- The online monthly GRID India meetings, started from 2023.
- Scheduled on the last Thursday of each month at 16:00Hrs.
- Try to keep maximum 1Hr
- Discuss different aspects of the regional WLCG GRID India Project
 - Know how between two Tier-2s,
 - Both Tier-2 monthly reports,
 - technologies (middleware) used for Tier-2s,
 - Network and securities,
 - Tier-3 status and related issues.
- All the GRID India project stakeholders are part of the meeting.
- The each month's meetings information, agenda and minutes are available under the CERN-INDIA category at CERN INDICO
 - (<https://indico.cern.ch/category/675/>)
- **Presently this monthly meeting is single point of information and communication for the GRID India project.**



S N	Date	Link	Meeting Agenda
1	30/01/2025	https://indico.cern.ch/event/1508788	Update: Scientific Computing Workshop for High Energy Physics and Tomography TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update
2	27/02/2025	https://indico.cern.ch/event/1519992	Update: HSF-India HEP Software Workshop at the University of Hyderabad TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update
3	27/03/2025	https://indico.cern.ch/event/1531498/	TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update Computing Committees Discussion
4	24/04/2025	https://indico.cern.ch/event/1541021/	The Grid India Project Status, TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update Computing Committees Discussion
5	29/05/2025	https://indico.cern.ch/event/1543680/	GPU Computing and Discussion The Grid India Project Status, TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update
6	27/06/2025	https://indico.cern.ch/event/1562766/	TIFR Tier-2 Updates and Monthly Report Kolkata Tier-2 Monthly Update

Thank You

