No LHC Survey 2021: Requirements towards the Tecnopolo

2021 Survey - NoLHC Experiments

- 25 Questions: Computing,Storage, Monitoring
- 35 Responses
  - 46 active exps
    - 45 (Borexino close to decomissioning)
  - 2 from experiments not yet supported (LUNA, LEGEND)
  - 3 new experiments in 2021 CYGNO, GAPS, LITEBIRD

Limadou	Tristan - Matteo Biassoni	ENUBET	EUCLID - MAtteo Tenti [Corretto]
asfin: marco la cognata	ILDG -- Hubert Simma	JUNO - stefano mari	DUNE (NU_AT_FNAL) - Matteo Tenti
KM3NeT - Kay Graf (KM3NeT	Fermi-LAT Michael Kuss	DARKSIDE - stefano mari	AUGER-Gabriella CATALDI
HERD - Nicola Mori	CTA - Di Pierro (connessi in remoto Arrabito, Longo)	LHCf - Alessio Tiberio	Alessio Tiberio - GAPS
AMS-02 - Valerio Formato	LSPE - Francesco Piacentini	Giovanni Mazzitelli	Virgo - Stefano Bagnasco
DAMPE – FABIO GARGANO	XENONnT - Luca Scotto Lavina and Marco Selvi	LITEBIRD - Giovanni Signorelli	PANDORA - Taioli
n_TOF - Nicholas Terranova (r	Gerda - Katharina von Sturm	Belle II - Silvio Pardi	CHIRONE (exNEWCHIM) - Pirrone-Russotto
Legend	CUPID - Lorenzo Pagnanini	NUCLEUS - Riccardo Cerulli	LUNA - Zavatarell - Imbriani
CUORE - Matteo Biassoni	FAMU - Emiliano Mocchiutti	NEWS - Valeri Tioukov	

CYGNO

Summary of the Summary

- **Computing**
  - Several requirements for:
    - Interactive access
      - Requested by many exps but for a small fraction of the workload
      - Cloud access
        - Interactive UI on steroids
        - Peak usage
        - Dynamic Cluster instantiation and low latency
      - Jupyter notebooks
    - HPC multinodo/acceleratori
      - NVIDIA GPU
        - Mainly tests, difficult to provide numbers for production
      - Infiniband
  - Not many requests for multicore single node > 8 core
  - 3.5 GB RAM per job slot is ok with few exceptions
  - 50% do not need outbound connectivity

- **Data Access**
  - POSIX access highly requested
    - At least on a scratch area
  - S3 and Obj
    - No strong requirements
  - Interest in
    - Sync&Share – Dropbox like
    - Automatic Replication
    - Rsync service for data transfers

- **Public Monitoring Pages**
  - Per user aggregation
  - Per experiment aggregation
    - almost done
  - Customizable dashboards
  - Aggregation per custom tags

- **15% increase Rate**
  - Several experiments are above this level
    - HERD; CTA; EUCLID; VIRGO; BELLE; CHIRONE; JUNO; AUGER

Survey Summary

Computing and Worker Nodes

Computing - Multicore

In the case of multicore jobs do you plan to use more than 8 processes per job?

35 responses

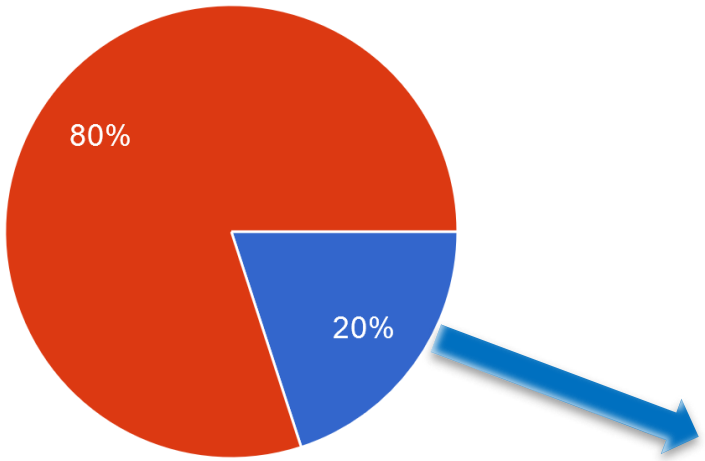

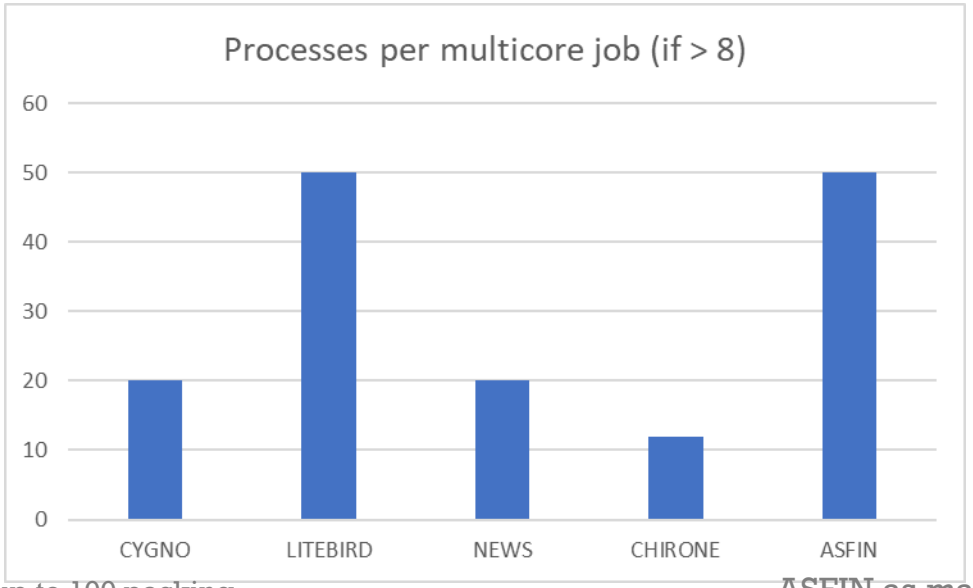

CYGNO up to 100 packing more events per job

ASFIN as many as possible within a node

noLHC experiements survey 2021: requirements

Computing - RAM

Is 3.5 GB of RAM per core enough for the foreseeable future (up to 2025)?

33 responses

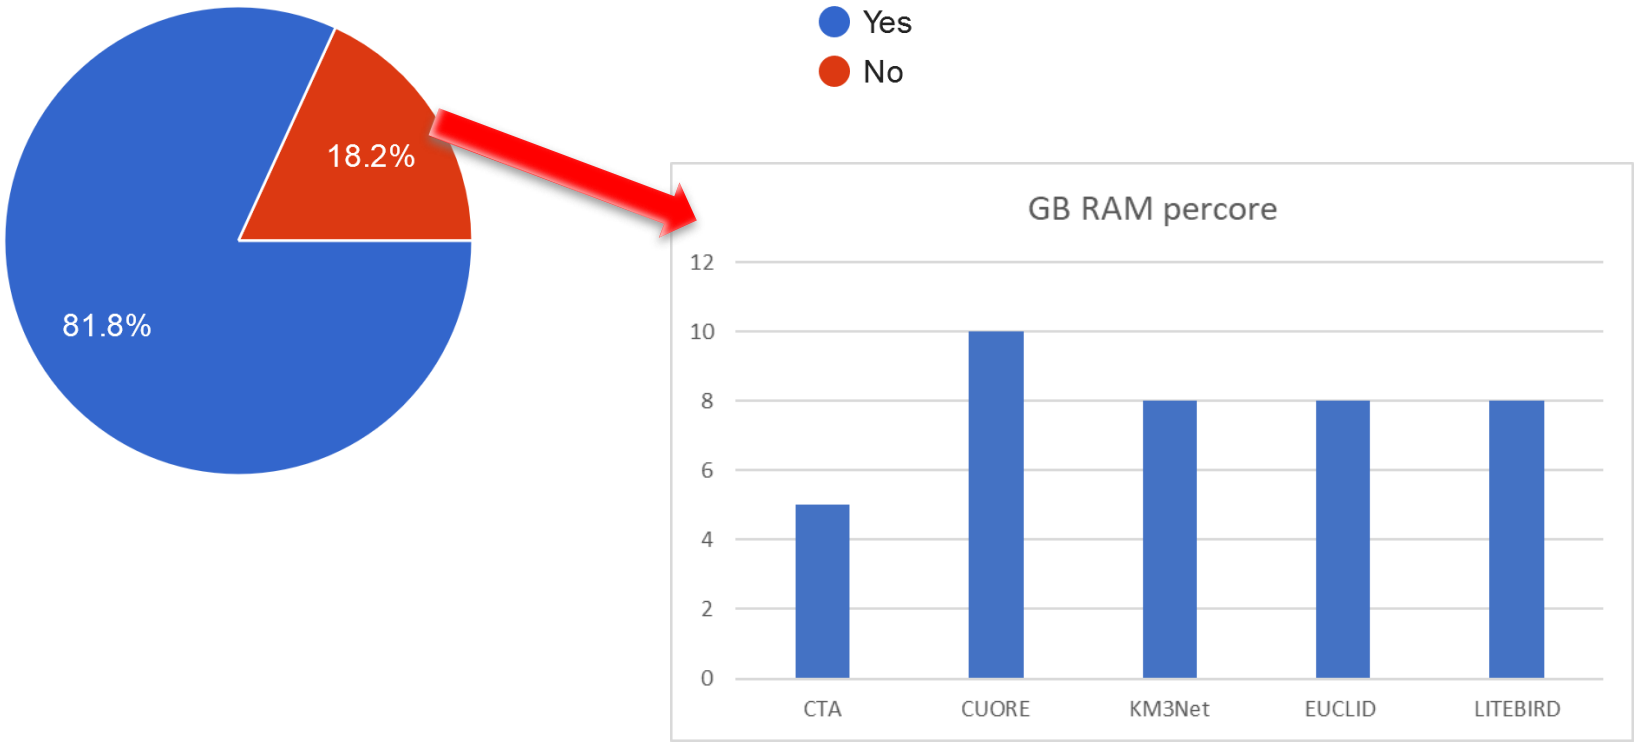

Computing - Containers

Do you plan to use containers (docker o singularity) to access the farm resources?

34 responses

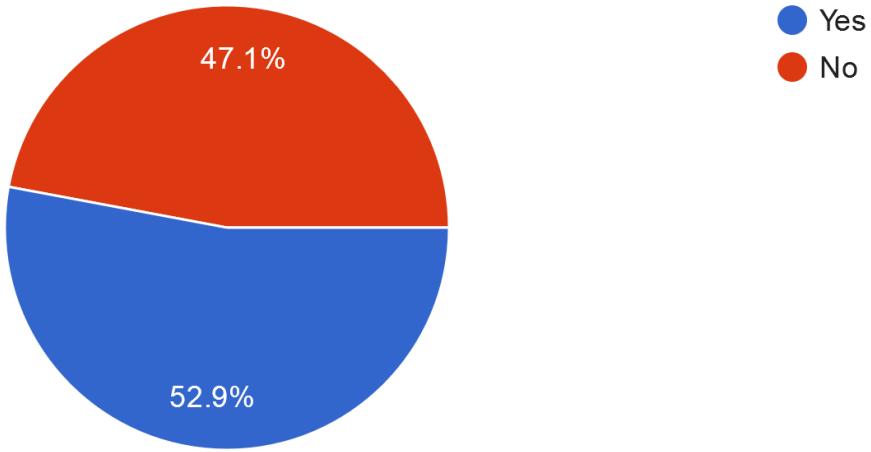

Computing - Connectivity

Do you need outbound connectivity from the computing nodes?

35 responses

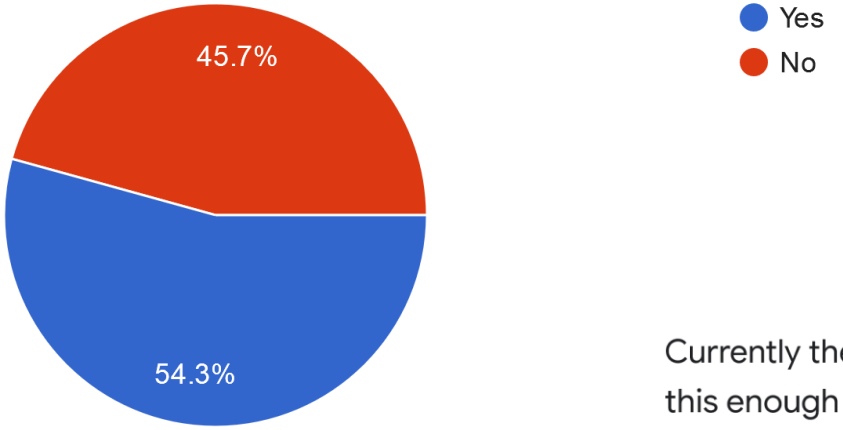

Currently the tier1 farm is dimensioned to support a maximum throughput of 5MB/s per jobslot, is this enough for the foreseeable future (up to 2025)?

35 responses

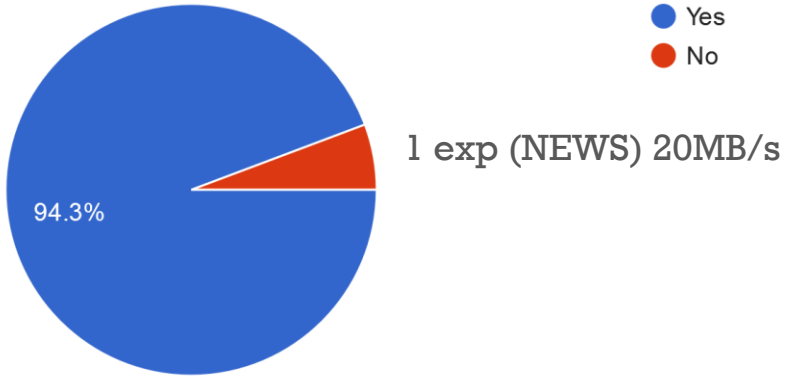

1 exp (NEWS) 20MB/s

Computing – Interactive Access

Do you have use cases for interactive jobs?

34 responses

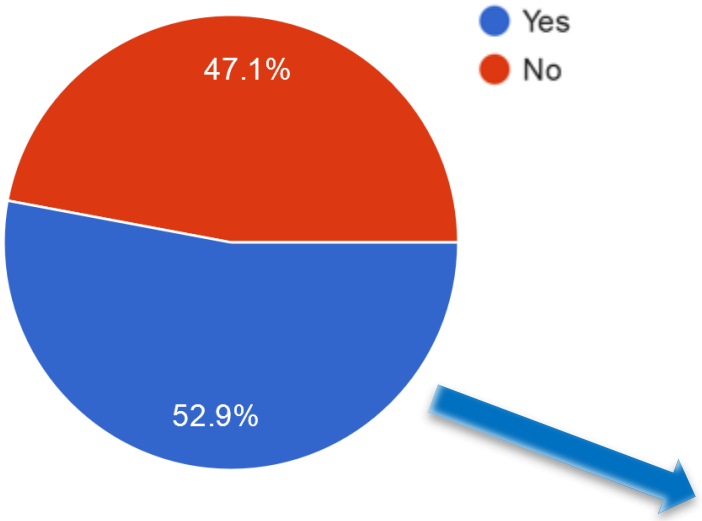

Few/Testing = 5%

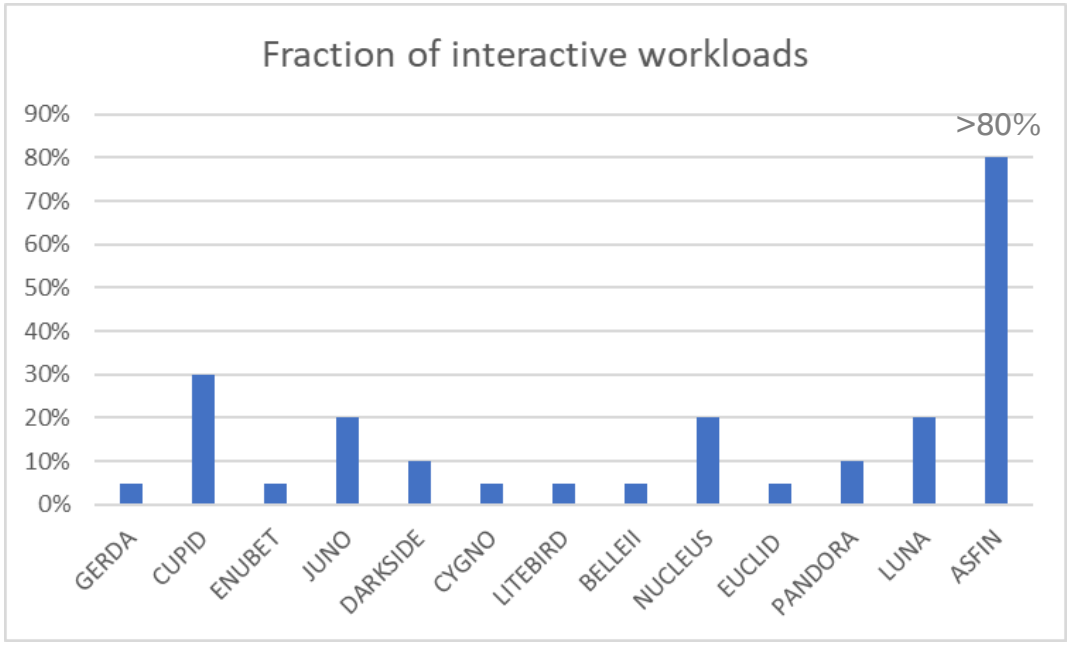

noLHC experiements survey 2021: requirements

Jupyter Notebooks

Do you have use cases that implies the use of jupyter notebooks?

34 responses

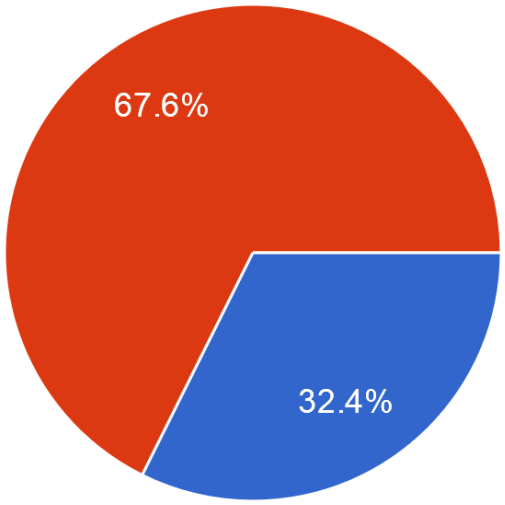

- Data Analysis and Visualization
  - LITEBIRD; HERD; EUCLID; DUNE
- Interactive Analysis
  - GERDA; BELLEII; LEGEND
- Development and testing, Visualization
  - CTA; CYGNO
- Fast Learning process for collaborators
  - LSPE

Cloud Use Cases

Cloud

Do you have use cases that need Cloud resources/interfaces?

14 responses

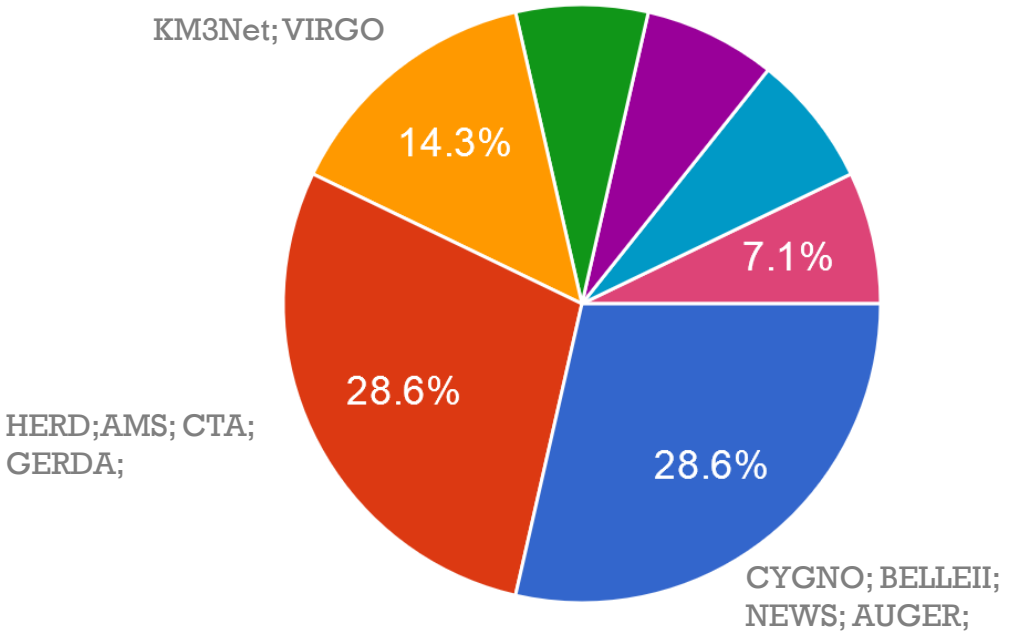

HPC, parallel jobs and accelerators

HPC

Do you need High Performance Computing (HPC - highly parallel jobs)

33 responses

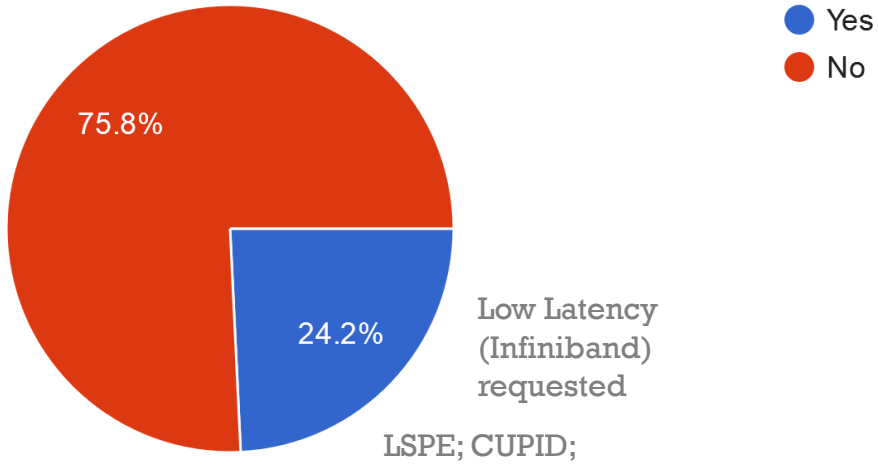

Low Latency (Infiniband) requested

LSPE; CUPID; CYGNO; LITEBIRD; EUCLID; VIRGO; PANDORA

HPC

Do you foresee the usage of computational accelerators?

32 responses

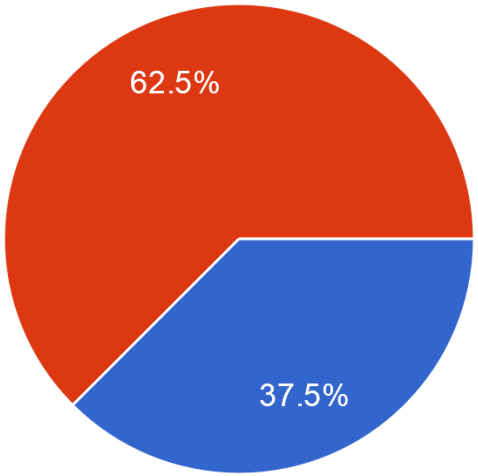

LSPE; CTA; GERDA;KM3Net DARKSIDE; JUNO; LITEBIRD;NEWS; BELLEII; EUCLID; VIRGO; PANDORA

If yes, which type of accelerator?

14 responses

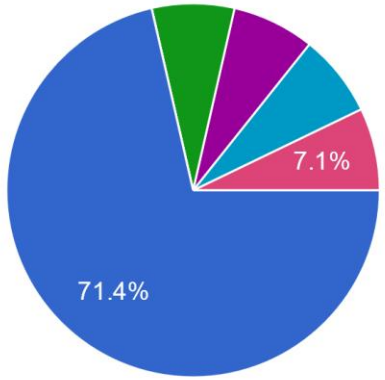

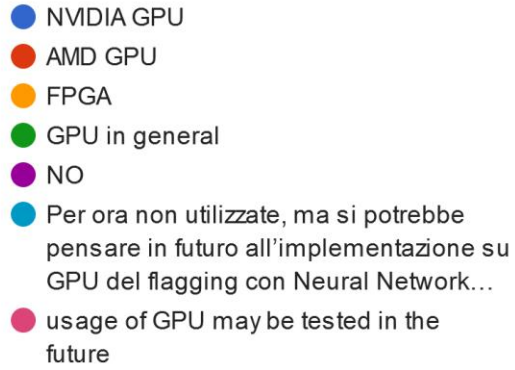

- Is it already possible to have an estimation of the computing power that has to be provided by the computational accelerators (up to 2025) - as a fraction of the total computing power needed
  - **Not Yet**

Data Access and Storage

Data Access - POSIX

Is it mandatory to maintain direct POSIX access to the storage areas from the computing nodes?

33 responses

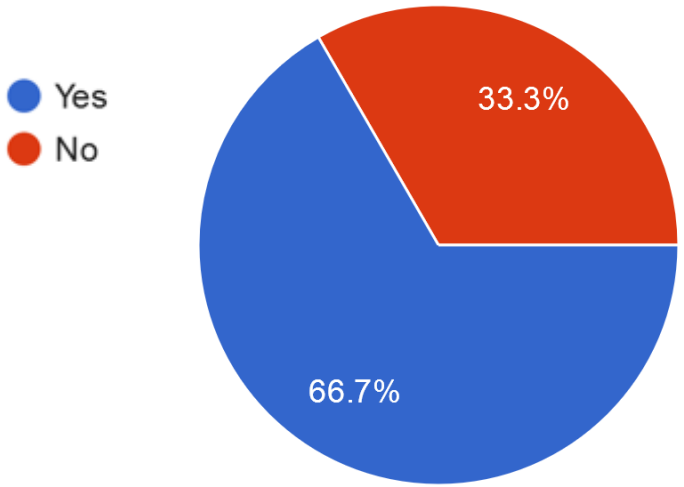

KM3Net; LEGEND; ILDG; FERMI-LAT; CTA; LSPE; LHCf; CYGNO; DUNE; AUGER; GAPS;

As an alternative, do you think that it would be feasible to access storage areas (disk and tape) via the http/WebDav protocol, maintaining a "scratch" POSIX area?

30 responses

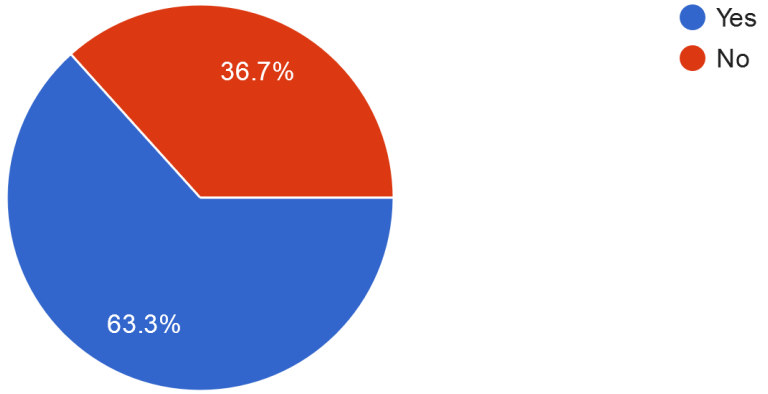

Data Access – Object Storage & S3

Do you have use cases that can benefit from the usage of Object Storage service with S3?

31 responses

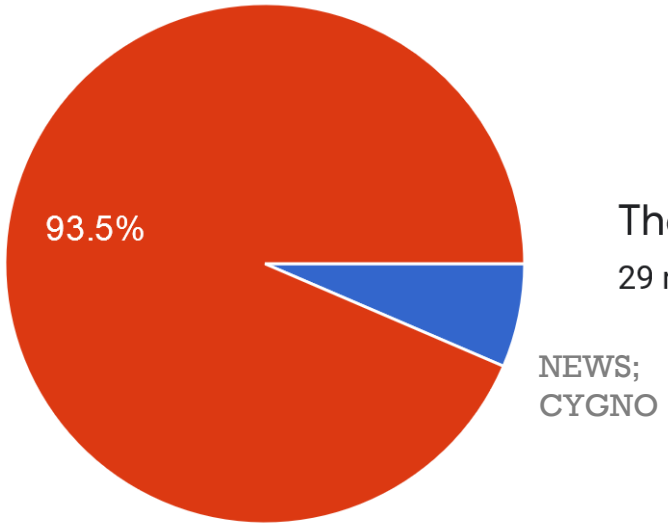

NEWS; CYGNO

The Object Storage could be an alternative to manage Storage Areas?

29 responses

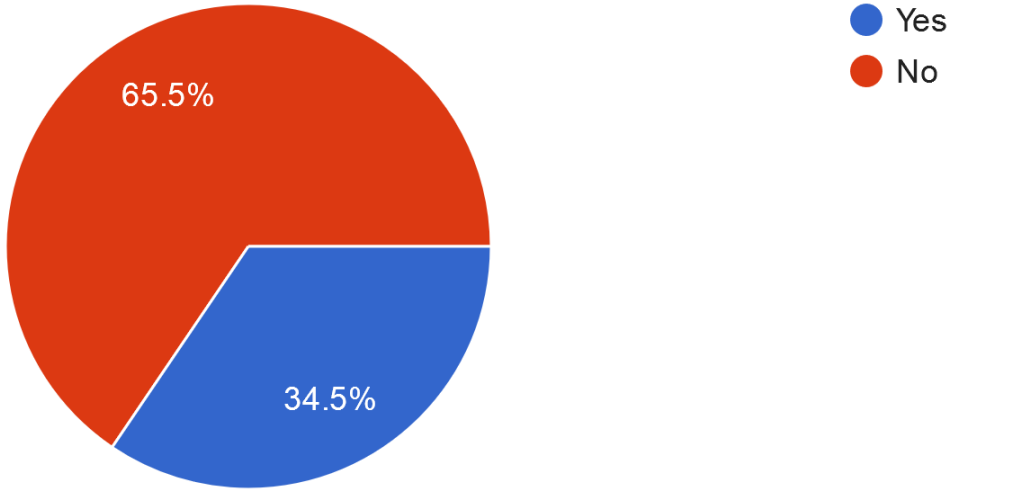

Data Access – Sync & Share

In your experiment do you have use cases for a Sync&Share service to have some CNAF storage areas browsable via web and eventually synchronized with remote desktop/servers?

32 responses

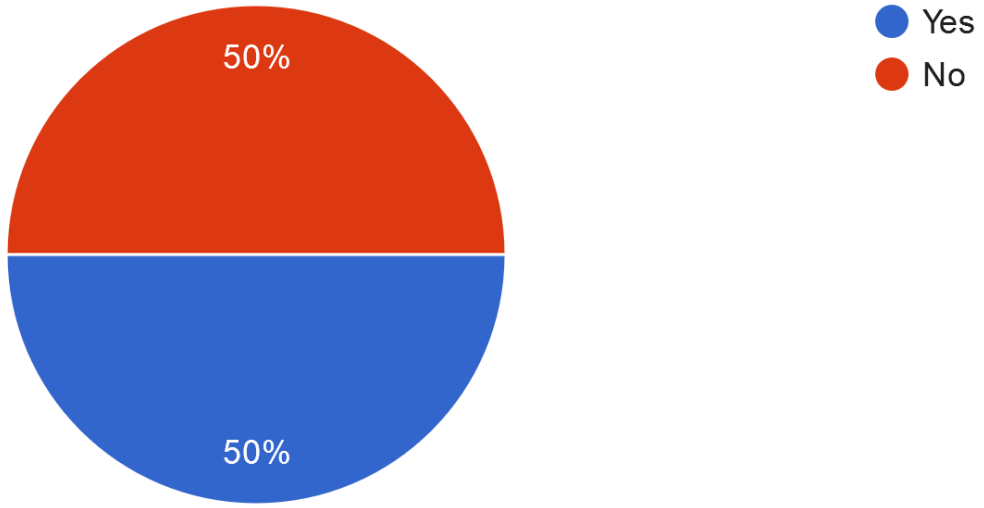

QoS and Automatic Replication

- Do you think an automatic data replication service would be useful?
- Do you have use cases for automatic data replication on a remote site based on policies and Quality-of-Service (disk, tape, number of replicas etc)? In case, please provide few details
  - CUPID: Yes, it would be useful
  - GERDA: Yes, on tape for backup
  - CTA: Yes
  - ILDG: Yes, currently done manually
  - CUORE: Yes, currently done manually via rsync towards multiple remote sites
  - LEGEND: Yes, 3 replicas in different sites are foreseen
  - DAMPE: Yes, available in China
  - AMS: Yes
  - HERD: Yes
  - KM3Net: Yes, done manually at the Tier0 stations, would be good to have at Tier1 level
  - NEWS: Yes
  - DUNE: Yes, to replicate data at Fermilab
  - VIRGO: Currently a similar functionality offered by Stashcash+cvmfs – something interoperable with this will be evaluated
  - LUNA: Yes towards tape
  - BELLE: Yes, but implemented via RUCIO

Monitoring

Monitoring

Considering the public CNAF monitoring interface (https://t1metria.cr.cnaf.infn.it) do you think it is providing enough information about cpu/disk/tape?

33 responses

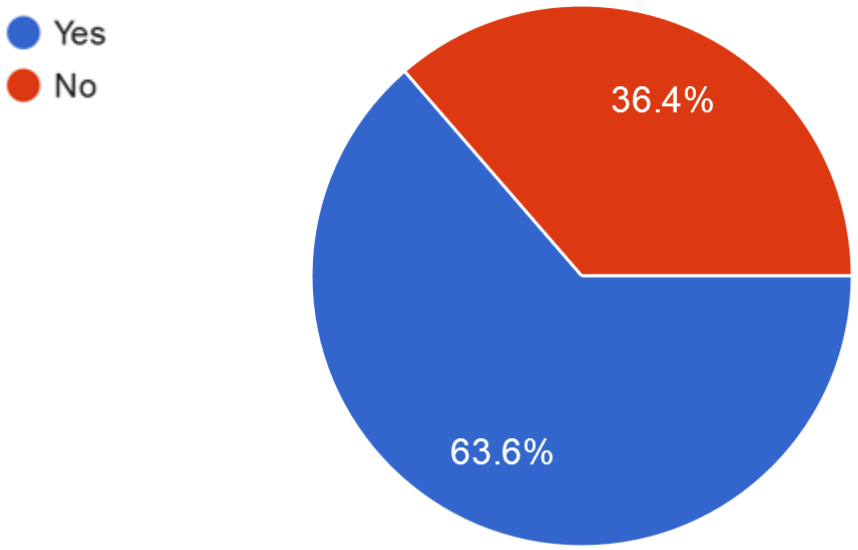

- Some requests:
  - HERD, AMD, DAMPE, PANDORA: Per user aggregation
  - CUORE; TRISTAN; FERMI-LAT, XENON;FAMU;BELLE: Per experiment disk/tape usage
  - KM3Net: customizable dashboards
  - VIRGO: If possible, CPU usage aggregation by custom field (tag in Condor submit file). Job throughput and memory usage.

Resources Increase Rate

Resources increase rate

We normally consider an increase rate (until 2025) of the resources needed for CPU, disk and tape of about 15%/year. Do you foresee significant deviation from this number?

32 responses

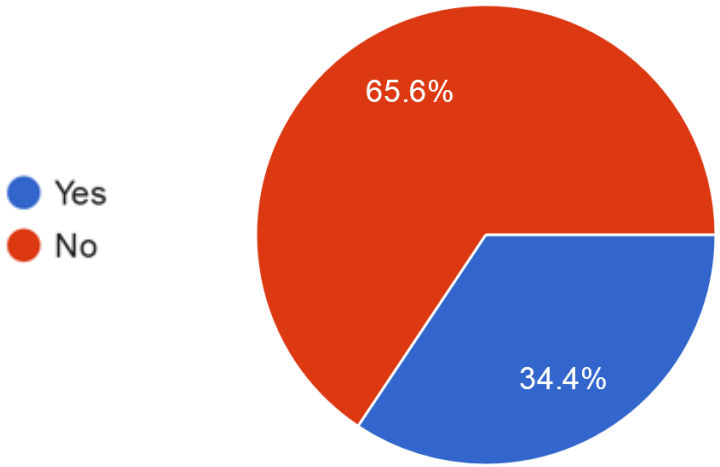

- HERD: 100% - we are about to be approved so our needs might ramp up significantly after approval
- CUORE
- TRISTAN
- CTA: 30%
- JUNO: increase to follow the data throughput from the start of data taking of the experiment
- BELLEII: Tape may increase of a factor between 30-50% because of the luminosity rump up
- EUCLID: 20%
- AUGER: 25%/Year due to the upgrade of the detector
- VIRGO: Tape: ulteriori 100-200TB (max) in tutto fino all'inizio di O4 (metà 2022 salvo ulteriori ritardi pandemici); dati di O4 approx. 1PB in circa 1 anno. O5 non prima del 2025. Disco: O(50-100) TB/anno
- CHIRONE: 500TB/year dal 2023 TAPE - spazio disco tampone di almeno 24TB
- CYGNO: entro il 2021 estremo in presa dati con il prototipo 1//20 ai LNGS che prevede un throughput di circa 0.1-1 TB/y al quale si aggiungono le prese dati fast in superficie per un equivalente ammontare dei dati. L'istallazione di CYGNO/INITIUM prevista per il 2023/24 dovrebbe portare il throughput 20 TB/y. E' prevista una presa dati per 5 anni. il throughput per la simulazione è attualmente ~0.01-1MB/evt, e si prevede una necessità di circa 10TB/y nei prossimi 3-4 anni.

Do you foreseen a significant increase of the throughput needed in read/write for tape (up to 2025)?

- BELLE: 25-30% per year increase
- CTA, JUNO: yes, following data taking