

Table of Contents

Question.....	1
Answers.....	2
CERN.....	2
hephy-Vienna.....	2
KI-LT2-QMUL.....	2
UKI-LT2-RHUL.....	2
RO-13-ISS.....	2
Nebraska.....	2
INFN-ROMA1.....	2
NDGF-T1.....	2
BEgrid-ULB-VUB.....	2
NCG-INGRID-PT.....	3
IN2P3-IRES.....	3
LRZ-LMU.....	3
CA-WATERLOO-T2.....	3
CA-VICTORIA-WESTGRID-T2.....	3
Taiwan_LCG2.....	3
IN2P3-SUBATECH.....	3
asd.....	3
MPPMU.....	3
INFN-LNL-2.....	3
Australia-ATLAS.....	3
KR-KISTI-GSDC-02.....	4
UKI-LT2-IC-HEP.....	4
-UCL.....	4
UKI-SOUTHGRID-BRIS-HEP.....	4
GR-07-UOI-HEPLAB.....	4
UKI-SOUTHGRID-CAM-HEP.....	4
USC-LCG2.....	4
EELA-UTFSM.....	4
DESY-ZN.....	4
PSNC.....	4
UAM-LCG2.....	4
T2_HU_BUDAPEST.....	5
INFN-Bari.....	5
IEPSAS-Kosice.....	5
IN2P3-CC.....	5
NONE_DUMMY.....	5
WEIZMANN-LCG2.....	5
RU-SPbSU.....	5
USCMS_FNAL_WC1.....	5
RRC-KI-T1.....	5
vanderbilt.....	5
UNIBE-LHEP.....	5
CA-SFU-T2.....	6
_CSCS-LCG2.....	6
T2_BR_SPRACE.....	6
T2_BR_UERJ.....	6
GSI-LCG2.....	6
UKI-NORTHGRID-LIV-HEP.....	6
CIEMAT-LCG2.....	6
a.....	6
T2_US_Purdue.....	6

Table of Contents

Answers

IN2P3-LAPP.....	6
TRIUMF-LCG2.....	6
KR-KISTI-GSDC-01.....	7
GRIF.....	7
IN2P3-CPPM.....	7
IN2P3-LPC.....	7
IN2P3-LPSC.....	7
ZA-CHPC.....	7
JINR-T1.....	7
praguelcg2.....	7
UKI-NORTHGRID-LIV-HEP.....	7
INDIACMS-TIFR.....	8
TR-10-ULAKBIM.....	8
prague_cesnet_lcg2.....	8
TR-03-METU.....	8
aurora-grid.lunarc.lu.se.....	8
SARA-MATRIX_NKHEF-ELPROD__NL-T1_.....	8
-UNIBA.....	8
DESY-HH.....	8
T3_PSI_CH.....	8
SAMPA.....	8
INFN-T1.....	8
GLOW.....	9
UNI-FREIBURG.....	9
Ru-Troitsk-INR-LCG2.....	9
T2_Estonia.....	9
pic.....	9
ifae.....	9
NCBJ-CIS.....	9
RAL-LCG2.....	9
T2_IT_Rome.....	9
BNL-ATLAS.....	10
FZK-LCG2.....	10
INFN-NAPOLI-ATLAS.....	10

Question

Question 6. Do you share any of your storage infrastructure with other (i.e., non-WLCG) communities? This question excludes commonly shared infrastructure, such as networking and monitoring. If you do share storage, what impact does that have on your choice of technologies?

Answers

CERN

Yes. Numerous smaller HEP experiments. Impact on choice of technology - minimal, technology choice is driven by WLCG.

HPC stuff, QCD, theory, beams, AMS (<https://greybook.cern.ch/greybook/experiment/recognized>).

hephy-Vienna

Some part of the storage is used by our local community. We choose EOS, as we expect that it addresses best the different needs.

KI-LT2-QMUL

Our site supports over 20 other VOs including SKA and LSST. storage needs to work for all stakeholders. A wlcg only solution is of no use to us

UKI-LT2-RHUL

Not yet but planning to share a HDFS ~1 PB.

RO-13-ISS

no impact

Nebraska

We do not share any bulk storage infrastructure. We do share things like backup disk vaults and VM iSCSI targets but these devices are not part of any resource available outside our operation.

INFN-ROMA1

We do offer storage areas to other experiments, with no significant impact on the choices of our technologies as they are all integrated into a global, hyperconverged system

NDGF-T1

No, not yet, but offering it to one or two new communities based on the solid storage technology we are running.

BEgrid-ULB-VUB

Yes, we share it with 2 other experiments. Only WLCG/CMS drives our choices, being our largest client.

NCG-INGRID-PT

Yes. We need to support HTP and HTC workloads with POSIX access. We also need to support cloud block storage and object storage.

IN2P3-IRES

We are sharing the storage with other non-LHC VOs (BELLE 2, biomed, France Grilles, local VOs, ...). As all these VOs are using DPM, it does not impact our technology choice.

LRZ-LMU

Not currently. But we are in discussion with other science communities to share the storage.

CA-WATERLOO-T2

No sharing right now. If we did go in this direction, would explore StoRM on Lustre (though expect lower reliability in this configuration based on experience)

CA-VICTORIA-WESTGRID-T2

dCache is also used by Belle. We also have Ceph storage that is used for our Openstack cloud but not for dCache.

Taiwan_LCG2

We do share our storage infrastructure with other communities. The most impact on our choice of technologies is posix-direct access requirement from applications and workflows. So, we can only support such applications with CephFS or EOS-fuse.

IN2P3-SUBATECH

No (do not share)

asd

MPPMU

yes

INFN-LNL-2

No, our T2 storage is dedicated to CMS and ALICE

Australia-ATLAS

No

T3 clusters also run non-WLCG jobs. Some users at ARNES got used to dCache at ARNES to store larger amount of data for data processing on T3 clusters.

KR-KISTI-GSDC-02

No, We provided the dedicated storage for WLCG Tier-2. Non-WLCG storage is physically separated. In addition, we have no plan to share the storage to non-WLCG community.

UKI-LT2-IC-HEP

Yes (for sharing, mainly astro and neutrino physics), no (for choice of technology)

-UCL

UKI-SOUTHGRID-BRIS-HEP

PheDex box is shared with RALPP & Oxford

GR-07-UOI-HEPLAB

No

UKI-SOUTHGRID-CAM-HEP

No

USC-LCG2

NO

EELA-UTFSM

Yes

DESY-ZN

Yes. WLCG storage is not shared though, it is run on an independent dCache instance.

PSNC

Currently only DPM is shared, but we there is a plan to migrate DPM and XROOD to dCache.

UAM-LCG2

No

T2_HU_BUDAPEST

no

INFN-Bari

NO

IEPSAS-Kosice

No we don't share any of our storage infrastructure

IN2P3-CC

Yes clearly, tape system is share by many other experiments and services, disk infrastructure (dcache and xrootd) are also share. Both case, WLCG activity is the bigger, but it is not necessary true for the next years. Site will support some others «big» experiments and we have to consider how we will share the storage infrastructures. Technologies choices have to be enough "generic" to allow us to satisfy the WLCG requirements (of course) but also the requirements of the others. We cannot have a (too) strong relationship between storage infrastructure and storage service.

NONE_DUMMY

blah

WEIZMANN-LCG2

Yes, we have on Lustre file space supporting WLCG, local ATLAS and non-WLCG users. The non-WLCG users have more modest capacity and performance requirements than WLCG.

RU-SPbSU

USCMS_FNAL_WC1

Tier 1 is exclusively WLCG/OSG

RRC-KI-T1

No, our storage is WLCG only.

vanderbilt

we're by far the biggest user, so it's our decision

UNIBE-LHEP

Shared the SE and the ARC cache / scratch with ht2k.org and microboone, all transparent. Would not change the dominant WLCG technologies for accommodating non WLCG communities. The other way around rather

CA-SFU-T2

We have 3 dcache instances (for Atlas, SNO+, T2K). Plus we have a completely separate storage (about 50 PB on disks and 300 PB on tapes) for ComputeCanada general usage. All in the same machine room.

_CSCS-LCG2

Yes we share some storage systems. The impact is that we chose systems that can provide (per customer) required performances in a shared environment.

T2_BR_SPRACE

NO

T2_BR_UERJ

No, we do not share our storage.

GSI-LCG2

Our storage is shared with many local groups from different experiments. Therefore technologies cannot be chosen to fit only our use case.

UKI-NORTHGRID-LIV-HEP

Yes

CIEMAT-LCG2

We do (as discussed above). The main impact is our interest in non-GSI authentication methods (especially Kerberos and token-based), and use of standard clients (NFS and WebDAV). In general, we also use more than one file replica for these communities, since they don't have replicas of data in other sites.

a

T2_US_Purdue

No

IN2P3-LAPP

The site provides a unique cluster for all the computing activities: grid (WLCG and EGI VOs) and local activities (batch). Therefore, the distributed file system available across the cluster is shared among all users of the cluster.

TRIUMF-LCG2

We don't share our storage with other non-WLCG communities; however the Tier-1 centre is co-located

within a shared data centre facility serving both WLCG (Tier-2) and non-WLCG communities via a common infrastructure for electrical, cooling and wide area networking. The WAN connectivity and capacity are not an issue.

KR-KISTI-GSDC-01

For the moment, we do not share the storage allocated to ALICE VO with other communities. It is dedicated. If we should share this storage with other VOs, first of all we will consider the operational cost, which means we will choose any storage type or protocol that are mostly compatible with and are easy-to-manage. We have another storages for the other VOs. For this purpose, we procured NAS (network attached storage) and it is shared among different VOs, mostly domestic research communities. Recent NAS storage is easy to manage disk space and we are happy with it.

GRIF

Yes, but without impact, other communities are more or less obliged to adapt to WLCG

IN2P3-CPPM

share with other egi vo but wlcg vo are much bigger and choice are made according to WLCG

IN2P3-LPC

Yes, others VOs are supported

IN2P3-LPSC

Yes (but few amount)

ZA-CHPC

no

JINR-T1

Yes, no impact

pragueicg2

Yes. Impact: still use SRM in our DPM

UKI-NORTHGRID-LIV-HEP

Yes, local research groups store some data on our DPM service, with local storage servers added to the pool. Access is via standard grid-compatible tools (eg xrootd) but we are dropping this in favour of local storage clusters with more standard POSIX interfaces.

INDIACMS-TIFR

NO, all of our infrastructure is for WLCG

TR-10-ULAKBIM

No.

prague_cesnet_lcg2

NO

TR-03-METU

No

aurora-grid.lunarc.lu.se

Yes, our central IBM spectrum scale is a shared storage to all compute nodes.

SARA-MATRIX_NKHEF-ELPROD__NL-T1__

yes. none

-UNIBA

WE do not share our storage infrastructure with other (i.e., non-WLCG) communities.

DESY-HH

We explore storage technologies and will continue in the future, especially for archiving.

T3_PSI_CH

-

SAMPA

Yes, with local user directory, in this case don't have impact in our choice

INFN-T1

Yes, we do share our storage infrastructure with other communities (about 40 mainly high energy physics, astrophysics and gravitational waves), many of them are asking for POSIX data access

GLOW

We share negligible amount of storage with non-CMS VOs. This doesn't constrain our choice of technologies.

UNI-FREIBURG

no

Ru-Troitsk-INR-LCG2

No

T2_Estonia

Our selected technologies (ceph) support other infrastructure quite well and we did not had prefer one to another.

pic

We run dCache and support around 15 VOs. Around 75% of the disk space is used by LHC experiments. Each disk server deployed in PIC is used by only one specific VO. This reduces VO destructive interferences in the servers, and we are quite happy with this configuration which is running since several years. Also, the big experiments use their own front-ends to dCache, and small-VOs used generic front-ends. Also SRM is deployed in this way. Hence, there is no real interference among VOs.

ifae

We run dCache and support around 15 VOs. Around 75% of the disk space is used by LHC experiments. Each disk server deployed in PIC is used by only one specific VO. This reduces VO destructive interferences in the servers, and we are quite happy with this configuration which is running since several years. Also, the big experiments use their own front-ends to dCache, and small-VOs used generic front-ends. Also SRM is deployed in this way. Hence, there is no real interference among VOs.

NCBJ-CIS

The storage is shared via different services. We never had a problem to find a technology which could support both WLCG and non-WLCG activities.

RAL-LCG2

Both Disk and Tape storage services are shared with other communities. Tape is roughly a 50:50 split while disk is currently closer to 80:20 (in favour of the WLCG). The disk contribution from other communities is growing. Other communities do not like using HEP specific solution (e.g. SRM, XRootD). Our strategy is to provide industry standard APIs from which each community can far more easily build their own layer on top.

T2_IT_Rome

No

BNL-ATLAS

We have separate dCache instances for non-WLCG communities. On the tape side, one single HPSS manages tape systems for ATLAS and other non-WLCG VOs, but each community has its own dedicated tape storage class service (library/tape drives/tapes/network). The choice of technologies have traditionally been driven by WLCG communities, because of their large user base and high demand on resources.

FZK-LCG2

We share the network and storage infrastructure among all the HEP experiments we support. No interference is expected.

INFN-NAPOLI-ATLAS

no

-- OliverKeeble - 2019-08-22

This topic: LCG > QosSurveyAnswersQ6
Topic revision: r3 - 2019-09-20 - OliverKeeble



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)