

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.....	5
DESY-ZN.....	5
PSNC.....	5
UAM-LCG2.....	5
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.....	6
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.....	7
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.....	10
RAL-LCG2.....	10
T2_IT_Rome.....	10
BNL-ATLAS.....	10
FZK-LCG2.....	10
INFN-NAPOLI-ATLAS.....	10

Question

Question 8. Are you aware of anything that WLCG VOs could change that would result in better performance or that would reduce your cost? Examples include client IO read requests are too large or too small, there are too many concurrent requests, too many requests over WAN. If there are any such changes, please describe them. If the changes are VO-specific, please indicate to which VO the comments are directed and which tier is your site (e.g., Tier-1, Tier-2, Tier-3)

Answers

CERN

More efficient use of available capacity. Fair share of capacity. Internal overcommitment would be good - now the norm is to have different groups with non-overlapping reservations. More aggressive deletion - less keeping data just in case. Experiments can't use all data on disk - often they are waiting (e.g. Castor - waiting for the full dataset to come online before processing).

At the moment it's difficult to map different experiment workflows onto different QoSs. The unit of matching is the entire system. More fine-grained QoS differentiation, e.g. by namespace path, would allow targeted performance boosts. This requires support at both storage and application level.

We need ways to dissuade users from using certain systems for certain things, e.g. compilation on AFS. With different QoSs, the users will not necessarily use the right one, they need to be steered through exposure of cost.

hephy-Vienna

Our site is too small. We are usually not limited by IO on the WAN.

KI-LT2-QMUL

UKI-LT2-RHUL

RO-13-ISS

the VOs already squeeze maximum performance from provided resources

Nebraska

Since our storage is so tightly coupled with our CPU nodes it is hard to determine what changes in storage would result in a cost reduction. Since we are a CMS Tier-2 facility we have a strong input lever arm into such discussions when they come up at the VO level.

INFN-ROMA1

NDGF-T1

Good support for IO only to local cache filesystem like ATLAS. Good tagging of job types into low/medium/high/very high IO, and integration into workflow management so that you could define a compute resource as "up to x HS06 and y IO" and have it choose a job mix to fill up the available IO without overloading cache, storage, or network.

BEgrid-ULB-VUB

CMS T2: too many xrootd requests over WAN sometimes. WAN hardware is terribly expensive, whereas LAN is cheap. Maybe we are unaware of it, but it would be good to have WAN predictions for T2s for the coming years. Could even have it in rebus as a pledge ?

NCG-INGRID-PT

Accessing the Lustre storage via POSIX-direct access instead of grid ftp would in our case improve performance. Standardize/harmonize all services across WLCG (or even better HEP) VOs would simplify and decrease costs as we support ATLAS, CMS and other HEP VOs.

IN2P3-IRES

It would be fine that we can deploy XRootD through DPM also for ALICE (we need actually a dedicated instance of xrootd for ALICE). We would also like to have better tools for diagnosing issues (for example when there are some network congestion). The PerfSONAR is not reliable enough.

LRZ-LMU

No

CA-WATERLOO-T2

Not clear if including tape as part of solution for large Tier-2 (>2PB) would be a cost saver, given the extra overhead on support. Load from ATLAS been manageable for last few years at current levels.

CA-VICTORIA-WESTGRID-T2

The ATLAS/Rucio data mover configuration seems a bit complex and opaque (many settings in AGIS, not always clear which ones are used), making it hard to see how data transfers are configured and how to change that.

Taiwan_LCG2

Nope.

IN2P3-SUBATECH

asd

MPPMU

INFN-LNL-2

Australia-ATLAS

Data access optimization for sequential access at all levels should be a high priority. ATLAS analysis on derived data is already optimized this way, other workflows should follow.

KR-KISTI-GSDC-02

In fact, I think there should be more changes in network technology than in storage systems. Our Tier-2 center does not have a higher CPU utilization than we thought. The reason is that data access is very inefficient. CMS analysis jobs are particularly prone to this tendency, because there is a strong tendency to receive data that is not currently on the site via CMS AAA. Of course, the dataset is being shared consistently through the Dynamo system, but the job is actually failing with a very high probability. In other words, there is a high probability that there will be problems with data transmission to external organizations through the XRootD Global Redirector. If there is no problem, it is likely that you need to make adjustments so that the data can be processed first if a site has the dataset to analyze.

It is very difficult to configure the XRootD access in our cluster through the private network in the current configuration. The CMS used "SITECONF" directory for searching a data location. As I know, XRootD URL information is used as an address for using the external Global Redirector. So, if we want to use another network for internal system, we need to create another siteinfo file. Of course, CVMFS support "local" SITECONF directory but it does not natural. I proposed to add site or server condition for URL matching. If the internal communication of the cluster can be more turned into the internal network, it provides more global connectivity. And, at some point in time, the data bandwidth of the storage will not be less than the network bandwidth.

There is also a slight discontent about the xrootd fallback system on the CMS. There are about three Tier-2 and Tier-3 centers in Korea. Experimentally, CPU usage and success rate were high when accessing data in Korea. Therefore, when we run CMSSW job, if we do not have the dataset internally, it should search for other domestic Tier centers first. However, it will be forced to execute xrootd fallback after internal data access failure. In addition, when we go to xrootd fallback, it's not just sending jobs to our domestic site. In other words, please make sure that you have data on nearby sites at the site level, or at least make the fallback jobs go to the nearby sites first. When we look at the jobs which are trying to access from slow and foreign sites even if the data is located on domestic site.

UKI-LT2-IC-HEP

No

-UCL

UKI-SOUTHGRID-BRIS-HEP

too many concurrent requests (CMS, Tier-2)

GR-07-UOI-HEPLAB

UKI-SOUTHGRID-CAM-HEP

No

USC-LCG2

No

EELA-UTFSM

DESY-ZN

Costs for WAN are pretty high in Germany. Additionally they are usually shared with all other activities at the site. Therefore chaotic, large scale usage of WAN should be avoided

PSNC

n/o

UAM-LCG2

T2_HU_BUDAPEST

no

INFN-Bari

IEPSAS-Kosice

we aren't aware of anything

IN2P3-CC

NONE_DUMMY

blah

WEIZMANN-LCG2

No changes necessary since our hi-performance storage system is underutilized. Much of our storage seems to be a warehouse.

RU-SPbSU

USCMS_FNAL_WC1

Providing more description of actual storage requirements beyond quantity and a rough availability requirement would be useful. IOPs, Network bandwidth to local WN's, WAN...

RRC-KI-T1

vanderbilt

no

UNIBE-LHEP

CA-SFU-T2

Storing less small files and more large files would help. We have Intel OPA between nodes and dcache pools - so no problems with bandwidth

_CSCS-LCG2

No

T2_BR_SPRACE

No I'm not.

T2_BR_UERJ

We are not aware of these changes

GS1-LCG2

IO could possibly be improved by letting the JobAgent make sure that all required input files are present on the local storage before the corresponding job is queued (VO: ALICE, Tier-2).

UKI-NORTHGRID-LIV-HEP

No

CIEMAT-LCG2

One issue we experience is a somewhat chaotic access pattern for (CMS) xroot WAN clients, which causes intermittent (but recurrent) xrootd connection "storms", with many clients trying to access our storage at the same time (it's not clear if there is also large data throughput, but at least there are many concurrent client connections). Our site is CMS Tier-2.

a

T2_US_Purdue

no

IN2P3-LAPP

Related to ATLAS computing: the site is currently doing computing in cooperation with far-away sites and data movements are probably not optimized. Non-WLCG VOs should be obliged to monitor their use of the LHCOne network.

TRIUMF-LCG2

We are an ATLAS Tier-1 centre, therefore providing nearline storage is part of our mission. To fully optimize tape operations, receiving large bulk staging requests will help site performance and throughput. We make use of tape families and group files belonging to the same dataset onto the same cartridge.

For WAN, we would like to see better streams utilization via FTS with better and truly dynamic features whenever possible and not simply use a generic setting, especially for TRIUMF which has a large RTT with most ATLAS sites.

KR-KISTI-GSDC-01

GRIF

No

IN2P3-CPPM

no

IN2P3-LPC

IN2P3-LPSC

Unifying storage access between LHC VO would greatly help (ex ALICE wrt others)

ZA-CHPC

nothing in particular

JINR-T1

No

praguelcg2

No.

UKI-NORTHGRID-LIV-HEP

At our Tier2 site the major performance issues are overloading of storage servers with far too many concurrent connections, both data being read/written by jobs alongside bulk data transfers in/out of the site. This is mostly from ATLAS activity. The issue is two-fold in our opinion 1) we are unable to effectively throttle the number of concurrent connections to DPM servers to a level that would avoid most of the overloading and 2) many of the concurrent connections are on files which are multiple GBs in size, making system caching ineffective. We saw a marked increase in overloading since ATLAS changed from smaller data files to the larger consolidated ones. There may be some overlap with our recent conversion to using ZFS as these systems are experiencing most of the overloading (this could just be coincidence as they are also the systems hosting the most data) and we are considering dropping ZFS and returning to more traditional RAID

systems.

INDIACMS-TIFR

TR-10-ULAKBIM

ATLAS, Tier-2, There are lots of requests that the service capacity is full and could not except connections in some cases. We try to solve this problem with new DPM disk servers but I am not sure about we will be successful or not.

prague_cesnet_lcg2

TR-03-METU

DPM nodes sometimes could not accept connections due to high number of requests.

aurora-grid.lunarc.lu.se

SARA-MATRIX_NKHEF-ELPROD__NL-T1__

No

-UNIBA

DESY-HH

WLCG should think very carefully about the implications of storage-less sites. Avoid small files, random access and optimize DST size.

T3_PSI_CH

-

SAMPA

I'm not aware

INFN-T1

In our opinion requests for storage resources should include not only capacity but also performance figures. It can be done specifying blocksize and IOPS, or using terms of QoS like Archival Storage (tape, off-line), Nearline Storage (NL-SAS, slow disks), Fast or Cache Storage (SAS, SSD, NVMe). From the storage services point of view it would be useful to know expected average and peak number of requests for each end-point and protocol. It would be useful to know the expected data flow during the year, in terms of writing on and reading from tape. This would help to plan the purchase of the needed number of tapes to fulfil pledges and to optimize the usage of resources shared with no-WLCG experiments. Moreover, it would be useful to group as much as possible recalls from tape of files belonging to related datasets, in order to minimize the number of mounts of the same cartridges.

GLOW

We are not aware of any such changes.

UNI-FREIBURG

n.a.

Ru-Troitsk-INR-LCG2

No

T2_Estonia

We are Tier2 and I have seen jobs what download git repo with 20k+ files 90%+ small or very small files for job. 1-100kb. It is quite heavy for storage. At some point I collected information per job but have lost that report what I made. If I have more time I will proceed with that investigation how job are effecting our compute/storage resources. Lately our downlink/uplink was full and it was used by jobs. Not sure what they were downloading but it was full. As we have 20G link, i will split our router config into two for better utilization.

pic

CMS and ATLAS read input files from jobs (locally) in different ways: CMS opens a connection that remains open along the job execution time, while ATLAS first copies the input file(s) into the compute nodes. We've also observed that CMS opens exploratory connections of 2 or so, to check the health of the connection, and this sometimes overloads the systems (the queues in dCache saturate). If CMS SAM tests cannot enter into the site, then a GGUS is typically sent to the site, but it's not a site fault. This might work much better if dCache could allow connections based on the bandwidth used, not just by setting an arbitrary limit of number of N connections of type M at a given time. This sometimes requires manual interventions to fix congestion issues in terms of opened connections.

XRootD exports via WAN are important for some VOs, in particular CMS. If more and more small sites deploy resources with caches or no-storage at all, this might have an impact to the sites holding permanent storage, and in particular PIC.

Disk is an expensive resource to be deployed, and VOs are trying to optimize the usage. This might have an impact in utilizing more and more tape systems, in a way that might not scale. The tape technology evolution seems fragile and this should be taken into account when adopting new ways to deal with data elsewhere.

CMS is using a very large number of file families to organize their data stored on tape. Sometimes this reverts in inefficiencies, for file families that end up not holding a lot of data. A very large number of file families typically represents an overhead in operational effort, so we would encourage experiments to try to use as few FFs as possible (or a definite set of them which is fully performant), and give as much information as possible to the sites on each of the FFs they recommend using (type of data, deletion policies, access patterns, etc), as well as having information beforehand of the expected usage of the tape system from the experiment's perspective

ifae

CMS and ATLAS read input files from jobs (locally) in different ways: CMS opens a connection that remains open along the job execution time, while ATLAS first copies the input file(s) into the compute nodes. We've also observed that CMS opens exploratory connections of 2 or so, to check the health of the connection, and this sometimes overloads the systems (the queues in dCache saturate). If CMS SAM tests cannot enter into the site, then a GGUS is typically sent to the site, but it's not a site fault. This might work much better if dCache could allow connections based on the bandwidth used, not just by setting an arbitrary limit of number of N connections of type M at a given time. This sometimes requires manual interventions to fix congestion issues in terms of opened connections.

XRootD exports via WAN are important for some VOs, in particular CMS. If more and more small sites deploy resources with caches or no-storage at all, this might have an impact to the sites holding permanent storage, and in particular PIC.

Disk is an expensive resource to be deployed, and VOs are trying to optimize the usage. This might have an impact in utilizing more and more tape systems, in a way that might not scale. The tape technology evolution seems fragile and this should be taken into account when adopting new ways to deal with data elsewhere.

NCBJ-CIS

We do not have any request to the VOs at the moment.

RAL-LCG2

VO should take more responsibility for optimising their job performance. Far too often it is just ignored as a problem by the VOs and left to be done by the site. We are a Tier-1 so we have the effort to do this but I know it is a serious problem at many Tier-2s who don't.

VOs should stop expecting all optimisations to be provided in a manner that is completely transparent to their workflows and provide more tools/handles/hooks for sites to use to optimise things. e.g. for ATLAS, user analysis tarballs that get downloaded by all jobs run by that user should be stored multiple times or on a high performance cache.

VOs should stop assuming a one size fits all approach to data access. Different workflows (MC simulation, data re-processing) have very different access patterns so why should the same data access model be expected?

T2_IT_Rome

No

BNL-ATLAS

In WLCG, several data lake models and implementations are being explored and tested. We're also working on our own data lake prototyping, to save cost while continuously providing reliable and scalable storage services for HL-LHC.

FZK-LCG2

INFN-NAPOLI-ATLAS

no

-- OliverKeeble - 2019-08-22

This topic: LCG > QosSurveyAnswersQ8

Topic revision: r4 - 2019-09-20 - OliverKeeble



Copyright &© 2008-2021 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](#)