

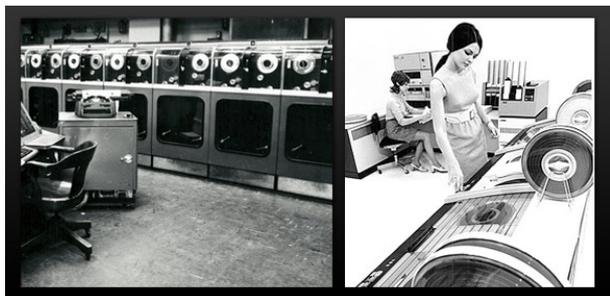
Table of Contents

CernLbdDiskSpace	1
This is the LHCb CERN LBD group page describing locally available disk space.....	1
Where to store what?.....	1
Personal Resources:.....	2
Grid Resources:.....	2
Group Resources:.....	2
More on the Castor disk pool:.....	3
-Description:.....	3
-Access:.....	3
-How to use:.....	3
-Use in Ganga.....	4
-Cleanup.....	5
-Other service classes.....	5
-FAQ.....	5
Does staging into lhcbt3 remove the staged copy from elsewhere?.....	5
Can a file staged to lhcbt3 be read without setting the STAGE_SVCCLASS?.....	5
Does unstaging the file delete it forever?.....	5
I unstaged the file, but it still appears if I nsls it!.....	6
How can I check if a volume is backed up or not?.....	6
Can I send files to castor lhcbt3 from DIRAC?.....	6
Can I swap files between service classes?.....	7
Can I share my afs disk space with other people?.....	7
Management Instructions:.....	7

CernLbdDiskSpace

This is the LHCb CERN LBD group page describing locally available disk space

This Twiki corresponds to the talk given here [↗](#)



Where to store what?

Where	What?	Size	Backed-up	Staging?	path
Afs home Directory	Ganga job repository, critical code you're working on, all other small files	2GB	Yes		~
Afs scratch directory	Ganga job workspace, temporary ntuples, noncritical code	3GB	No		~/scratch*
Group scratch	Ganga job workspace, temporary ntuples, temporary DSTs	1TB	No		/afs/cern.ch/project/lbcern
User castor	backup tars, large ntuples to share, old ntuples	??TB	Yes	Tape	/castor/cern.ch/user/<a>/<another>
Grid castor	Ntuples from the grid, Selected DSTs	5TB	Yes	Disk	/castor/cern.ch/user/grid/lhcb/user/<a>/<another>
Group castor	Output from lhcbt3	30TB	Yes and No	Disk	/castor/cern.ch/<somepath>

jobs, large numbers of ntuples and DSTs					
---	--	--	--	--	--

- The size is the user limit if it is your personal resource, or the full size if it is shared.
- The shared resources maintain a fair use policy.
- The non-backed up spaces are volatile, do not use them for critical items.
- The individual elements are described below.

Personal Resources:

- **Afs Home directories:**
 - ◆ All CERN users can get a 1GB Afs home directory, sans approval [here](#)
 - ◆ All Lbd members can get an expansion of their Afs home directory to 2GB and an additional 3GB scratch space softlinked from their home directory.
 - ◆ Your Afs space is for code development, and storing documents, ntuples, grid jobs, that you want backed up, but are not massive files.
 - ◆ `fs listquota` to get the quota/usage, use separately in each scratch space
 - ◆ To obtain this increase contact Thomas Ruf , who will forward your request
- **Castor Space:**
 - ◆ All CERN users have a castor home directory, where you can store almost anything you like, put things here if they cannot fit on your Afs dir, and need it to be backed up for a long time.
 - ◆ After a while files will be stored on tape, and take some time to stage.
 - ◆ To see your usage, run: `/afs/cern.ch/project/lbcern/scripts/castorQuota.py`

Grid Resources:

- **CERN-USER:**
 - ◆ LHCb has a large amount of disk space shared between all collaborators at each site
 - ◆ The CERN-USER disk, is the element located at CERN (T0)
 - ◆ your quota is around 5TB, and you should clean it regularly following the instructions here: `GridStorageQuota`
 - ◆ It is permanately staged to disk.
 - ◆ Use it to store large ntuples and DSTs that you got from the grid, which you want to share with others in the collaboration and which you don't want to disappear back to tape.
 - ◆ To see your usage proportion, run: `/afs/cern.ch/project/lbcern/scripts/castorQuota.py`

Group Resources:

The group disk resources are to be used **only** if the above available resources are not appropriate. It is maintained as a shared use policy, if you are using too high a proportion you will be asked to cut down.

- **Afs group disk:**
 - ◆ `/afs/cern.ch/project/lbcern`
 - ◆ We have a group disk on Afs.
 - ◆ It is not backed up, it is a volatile temporary storage disk.
 - ◆ the current size is 1TB, divided into 10 volumes of 100GB each
 - ◆ **Keep your items inside a directory with your name, anything outside a named directory will be deleted.**
 - ◆ `mkdir /afs/cern.ch/project/lbcern/vol<X>/<username>`

- ◆ Just your username, the same as you have it on afs, with no embellishments of addendums please
 - ◆ Use for collect large sets of outputs, that you can regenerate quickly, collections of small files, from many different grid jobs without the need to continuously clean your home directory.
 - ◆ to see the usage call `fs listquota` on each volume
 - ◆ The access list is maintained as a subgroup of `z5`, `z5:cernlbd`, to see if you are there call `pts mem z5:cernlbd`
 - ◆ To see the usage and check for problems call:
`/afs/cern.ch/project/lbcern/scripts/check1TB.py`
 - ◆ To be added to the acces list contact Thomas Ruf
 - ◆ to change the access list of directories you have created, you need to be added to admin your own directories, contact Thomas Ruf for that
- **t3 Castor space:**
 - ◆ We have a group disk pool on Castor. Like the grid resource it is always on disk, however it does not count towards your Grid quota.
 - ◆ It is technically tied to the `lhcbt3`
 - ◆ The current size is 30TB.
 - ◆ You can monitor the status here [↗](#)
 - ◆ Once per night all files are dumped to a big list here [↗](#)
 - ◆ The commands for checking quota, access etc. are given below

More on the Castor disk pool:

-Description:

The 30TB t3 castor disk pool is very special. It is a stager pool of disks for castor. Usually when you upload files to your castor working area, they are backed up on tape, and available on a staged disk for a short time. After which time the staged disk copy is deleted, and they need to be staged again. The `lhcbt3` is a castor staging pool, files can be copied to there by being staged **to** there.

-Access:

The access list of the t3 castor is maintained by Joel Closier. If you find yourself without access, you should get in touch with ThomasRuf. To see the access list call `stager_listprivileges -S lhcbt3`. Any `lbd` members should be granted access.

-How to use:

`lhcbt3` is a Castor disk pool, to move things there you need a different service class. Scripts to help you out with management are kept in `/afs/cern.ch/project/lbcern/scripts/`

To add large sets of files from other people, and/or files from central productions, it's best to use the `lhcbt3_stage.py` they will be accounted against you. In the case you add files from another user, this is very very annoying, because then the only person who can remove them is Joel by issuing a specific command from a specific machine.

Command	Use	Example
<code>stager_qry</code>	Check if file is there	<code>stager_qry -S lhcbt3 -M <existing castor file></code>
	List files from a directory	<code>stager_qry -S lhcbt3 -M <adirectory></code>

	List all your files	<code>stager_qry -S lhcbt3 -M /castor/cern.ch/user/<u>/<uname> -M /castor/cern.ch/grid/lhcb/user/<u>/<uname></code>
	Check usage of lhcbt3	<code>stager_qry -S lhcbt3 -sH</code>
<code>stager_get</code>	Stage single existing castor file into lhcbt3	<code>stager_get -M <existing castor file> -S lhcbt3 -U \$USER</code>
<code>lhcbt3_stage.py</code>	Stage existing castor file(s) from a given file, directory, or list.	<code>/afs/cern.ch/project/lbcern/scripts/lhcbt3_stage.py <source></code>
<code>stager_rm</code>	Remove single file only from lhcbt3	<code>stager_rm -M <existing castor file> -S lhcbt3</code>
<code>lhcbt3_rm.py</code>	Remove all files from a given user, directory or list	<code>/afs/cern.ch/project/lbcern/scripts/lhcbt3_rm.py <target></code>
<code>lhcbt3_mv.py</code>	Move files between stagers from a given directory or list	<code>/afs/cern.ch/project/lbcern/scripts/lhcbt3_mv.py <target> <svc1> [<svc2>]</code>
<code>stager_listprivileges</code>	To see the access list	<code>stager_listprivileges -S lhcbt3</code>
<code>lhcbt3_cp.py</code>	Copy new file, or a list of new files to lhcbt3 ¹	<code>/afs/cern.ch/project/lbcern/scripts/lhcbt3_cp.py <source(s)> <destination></code>
<code>STAGE_SVCCLASS</code>	Set lhcbt3 as your default staging pool unsafe ²	<code>export STAGE_SVCCLASS="lhcbt3"</code>
<code>rftp</code>	Copy new file to lhcbt3	<code>rftp <source> <destination> [after setting STAGE_SVCCLASS]</code>
<code>castorQuota.py</code>	Disk usage from you of all castor elements, including lhcbt3 (can be slow)	<code>/afs/cern.ch/project/lbcern/scripts/castorQuota.py</code>
<code>check30TB.py</code>	Disk usage of all users of lhcbt3 castor element (cached from the night before)	<code>/afs/cern.ch/project/lbcern/scripts/check30TB.py</code>

1. `lhcbt3_cp.py` avoids the need to set `STAGE_SVCCLASS` because it spawns a subshell
2. It's not a great idea to set `STAGE_SVCCLASS`, because if you forget to unset it, then every file you touch could be staged into the non-standard service class, causing bookkeeping problems

-Use in Ganga

Well, you **could** set the variable `STAGE_SVCCLASS` in your `bashrc`, but that would be problematic, since you would then stage into lhcbt3 any file which you touched on castor, and it's very difficult to book-keep that.

-How to use:

1. If you only want to write your files **to** lhcbt3, instead of read files **from** there or stage other random files to there, then edit in your .gangarc:

```
[LHCb]
cp_cmd = /afs/cern.ch/project/lbcbn/scripts/lhcbt3_cp.py
```

2. If you want to access files on, stage files to **and** write files to lhcbt3, if and only if you are running ganga on the LSF backend, then you can edit the line in your .gangarc:

```
[LSF]
preexecute = import os
os.environ['STAGE_SVCCLASS']='lhcbt3'
```

◆ You can also edit the config inside ganga dynamically to do this.

3. Use PFNs which end in ?svcClass=<svcclass>, ganga utils' guessPFN function can add this for you.

-Cleanup

If you are using a significant fraction of the space and somebody else needs some, you will be asked to clean up.

Removal in this respect is done sporadically or on request by Thomas Ruf at the moment, only cleaning up if there is more than 2TB used this way, or if the files are more than 50% of the free space.

-Other service classes

- **lhcbdisk**, self-explanatory
- **lhcbtape**, self explanatory
- **lhcbuser**, disk storage of user-generated grid files

-FAQ

Does staging into lhcbt3 remove the staged copy from elsewhere?

No.

```
$ stager_qry -M /castor/.../stdout
/castor/...../stdout STAGED
$ stager_get -M ...../stdout -S lhcbt3
$ stager_qry -M /castor/.../stdout
/castor/...../stdout STAGED
$ stager_qry -M /castor/.../stdout -S lhcbt3
/castor/...../stdout STAGED
```

Can a file staged to lhcbt3 be read without setting the STAGE_SVCCLASS?

Touching the file without setting the environment variable will cause the file to stage itself onto the default disk pools, copying from lhcbt3.

This is much quicker than staging from tape, but it is probably not what you want to do.

Does unstaging the file delete it forever?

No, technically the disk pool is still backed up to tape, but according to the experts:

There is tape backing (for tape fileclass files) for recovery purposes but not enough tape resour

I have no idea what that means...

I unstaged the file, but it still appears if I nsls it!

Technically the disk pool is still backed up to tape, and the file is still registered in the castor name server (NS) which is why nsls shows it.

To remove it completely you need to rfrm, just like every other file in castor.

How can I check if a volume is backed up or not?

in terms of backup, there are only two types of volumes: backed-up ones and non-backed up ones.

Rule-of-thumb volumes that have names starting with:

- p., .u, .user: **are backed up**
- q., s.: **not backed up**

However, this is a CERN only naming convention. To check a particular volume run

```
/usr/sbin/vos exa <volume name>
```

and look for a line starting with "Backup". If you see a date there, the volume is backed up. If you see "Never", it is not.

Can I send files to castor lhcbt3 from DIRAC?

There is no real difference between CERN-USER and lhcbt3, they are both permanently-staged CERN castor disk-pools

- Above I describe how to send files from LSF or Local backends directly to lhcbt3 using Ganga Use in Ganga.
- With the Dirac backend, your jobs could run anywhere in the world, and the output will usually be uploaded to the local grid storage element of the site at which they run.
- You can either replicate them to CERN-USER or redirect the output there, then the files will be at CERN. (see Grid Resources)
- There is only one CERN castor name server, so there is no real distinction between files on castor apart from where they are staged.
- Once they are on CERN castor you can stage them into lhcbt3 diskpool using the above commands, but both the CERN-USER and lhcbt3 are permanently staged to disk, so this will just double your file footprint.
- You could always stager_rm them from the other disk pools, but probably you are fine just leaving them where they are.

If you stage out of lhcbuser the files will probably just be staged back

- the LFN to PFN conversion done by ganga and bookkeeping etc includes the service class. This will be **lhcbdata** or **lhcbuser** or **lhcbdisk**
- that means if you use the catalog conversion you will then create a staging request into the old service, and thus duplicate the file
- To avoid this, do the LFN to PFN conversion yourself, either by starting with PFNs or using the guessPFN function in ganga utils
- If anyone else uses these files, they would also have to do the same trick, otherwise **they** will double **your** footprint.

Does unstaging the file delete it forever?

Can I swap files between service classes?

Yes. Stage into one service class, stage out of the other service class. Use `lhcbt3_mv` for that. But be aware the files might just come back (see above).

Can I share my afs disk space with other people?

Yes, this should be possible. You should have admin rights on your own directory, call `fs listacl`. If not, contact Thomas Ruf to fix. Then you should be able to modify your directories to be readable or writable by others, but note that if you make them writable by others, anything they write will count against your usage.

To modify the access list you will need to do `fs setacl <directory> <user or group> <setting>`, e.g.
`fs setacl mydirectory myfreindbob read`

Once you have modified a top level directory, you will need to propagate it to all sub directories where you want that access. To do that a simple way is `cp -r` since new directories use new permissions, but if that is not possible you will have to recursively change all the permissions for example:

```
find <mytopdir> -type d -exec fs sa {} cern:z5 read \;
```

For more on acl see here: http://docs.openafs.org/Reference/1/fs_setacl.html

Management Instructions:

See CernLbdDiskManagement.

-- RobLambert - 11-Nov-2010

This topic: Main > CernLbdDiskSpace

Topic revision: r32 - 2012-01-26 - ThomasRuf



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? Send feedback