

Vector read: current status

alexrg

Testing in production

- Patch was applied on lcg2270
- No xcache, only patched gateway
- Buffer size is 4MiB
- Errors are still happening

Results

The errors are still present. The failure rate looks a little bit lower, but that may be biased due to low number of LHCb jobs on the WN

	Lcg2270			All Farm		
	Successful	Vread fail	Fraction	Successful	Vread fail	Fraction
User jobs	470	5	0.01	40582	>=1064	0.026
User1	58	0	0	2000	52	0.025
User2	16	1	0.06	718	223	0.23
User3	5	1	0.17	259	14	0.05
User4	0	1	1	211	92	0.3
User5	11	2	0.15	1453	>=99	0.064

Test setup

- 64 jobs of User 2 on a dedicated WN
- Different xrootd configurations
- Errors can be reproduced with some configurations

Reason for errors

- Slow ceph reads
- Slow “memory allocation”?
- Huge amount of small reads?

Slow ceph reads

- Logging added to the patch (see latest commit in the PR).
- The following can be seen:

```
230201 02:43:33 Getting aio results for  
prod/lhcb/LHCb/Collision17/CHARMSPEC.MDST/000677  
27/0001/00067727_00012245_1.charmspec.mdst took  
too long, 43 seconds for read with coordinates  
15692335 1057562143 1
```

Slow “memory allocation”

- Sometime the following can be seen:

```
Jan 26 11:05:51 lcg2270 docker[878894]: 230126 11:05:51 076  
XrootdWait4Buff:  
tlhcb006.1530:25@htcjob1440921_0_slot1_31_pid676911.ralworker  
1 tardy aioR requests for  
lhcb:prod/lhcb/LHCb/Collision17/BHADRON.MDST/00071501/0000/00  
071501_00001155_1.bhadron.mdst
```

```
Jan 26 11:05:51 lcg2270 docker[878894]: 230126 11:05:51 076  
XrootdAioTask: async read failed for  
tlhcb006.1530:25@htcjob1440921_0_slot1_31_pid676911.ralworker  
; aio file read timed out  
lhcb:prod/lhcb/LHCb/Collision17/BHADRON.MDST/00071501/0000/00  
071501_00001155_1.bhadron.mdst
```

Memory allocation problem?

Small reads

- During the tests $\sim 3 \cdot 10^5$ readvs were issued and $\sim 10^7$ reads
- Mean size of the read is 6406.5 bytes, median is 301 byte, so the reads are small

ToDo

- Timeout increase
- Memory caching for small reads
 - Xcache does not work
 - Memory caching seems to work (i.e. reduce the number of reads)
- Debugging “Wait4Buff” errors