

Castor Delta Review LHCb Experiences

DM Activities During DC06:

- ▶ MC Production at Tier2/Tier1 sites:
 - ➡ .digi Files uploaded to Castor2
 - ➡ File by file at job finalisation
 - ➡ lcg-utils from remote sites
 - ➡ 1M < files uploaded since beginning of exercise
- ▶ .digi Files distributed to Tier1s:
 - ➡ from 'wan' pool to LHCb Tier 1s
 - ➡ All data pre-staged
 - ➡ FTS used
 - ➡ Exporting ~320 files/hr
 - ➡ Total files distributed ~ 0.5M

- Networking problems (July)
 - ▶ HRET switching problems
 - ➡ Upload Transfer rates decreased
 - ➡ Many concurrent active transfers into castor2lhcb
 - ➡ large number of sockets opened
 - ➡ Upload rate decreased further (eventually $\rightarrow 0$)
 - ▶ Many weeks of degraded performance
 - ▶ Not clear which debugging channels to follow
 - ➡ Not clearly a castor problem
 - ➡ Although castor performance suffered
 - ▶ Eventually resolved by direct contact with developers
 - ▶ DIRAC DM system now uses failover mechanism
 - ➡ In event of transfer timeouts
 - ➡ Failed uploads sent to TransferAgents of VOBOXES

- Static Mappings In Gridmap File
 - ▶ Changes in VOMS DBs not reflected
 - ➡ Hard to keep track of mappings on VO side
 - ➡ Email deployment team
 - ◆ Responsive
 - ◆ Not controlled/visible by VO directly
 - ▶ Differing mapping for users coming through SRM
 - ➡ e.g.
 - ◆ Files staged into 'wan' through SRM with FTS
 - ◆ Data Managers not able to unstage vi CLI
 - ◆ Possibly resolved when SRMv2 features fully available

- Problems (sometimes) encountered when upgrades performed
 - ▶ Client-server compatibility
 - Upgrades to server should be backward compatible
 - ▶ GridFTP daemons not restarted
- Clearly problematic
 - ▶ 100% failure rates
 - ▶ Resolved quickly once spotted

■ Client libshift.so

▶ By default picked up from:

■ /usr/lib

■ /usr/lib64

▶ Works if executable built with same compiler as libshift

▶ Otherwise:

■ Load default (gcc34) libshift from within (gcc323) binary

■ Core dump in application

▶ Work around by shipping libshift with software

■ Prepend the LD_LIBRARY_PATH

- LSF slot filled for each file read (in application)
 - ▶ Released only when file is closed
 - ▶ Possible for application to maintain several files concurrently open
 - ➡ Signal and spill over in Boole package.
 - ▶ If LSF queue full:
 - ➡ Deadlock when several running jobs request new files
 - ▶ Problem exacerbated with (old version of) application
 - ➡ Maintaining four files open by default
 - ➡ This corrected in new release (yesterday)
 - ➡ Problems may still occur as number of files always >1
 - ▶ Possible for single (normal) user to use 300 slots

► lcg-cr upload fails:

► Zero size file left in the name server.

```
> nslookup /castor/cern.ch/grid/lhcb/cnaf-stress/250MBFile/28.test
```

```
-rw-rw-r--  1 lhcbprod z5                0 Dec 05 10:08 /castor/cern.ch/grid/lhcb/cnaf-  
stress/250MBFile/28.test
```

► File stuck in 'STAGEOUT' state

```
> stager_qry -M /castor/cern.ch/grid/lhcb/cnaf-stress/250MBFile/28.test
```

Received 1 responses

```
/castor/cern.ch/grid/lhcb/cnaf-stress/250MBFile/28.test104756789@castorns STAGEOUT
```

► Unable to overwrite file

```
CastorStagerInterface.c:2457 Device or resource busy (errno=0, serrno=0)
```

```
lcg_cr: Communication error on send
```


Questions...?