



Item 9



- *The committee recommends that the development and operations teams review the list of workarounds, involving replacement of palliatives with features in the CASTOR core.*
- Actions taken
 - List of workarounds prepared
<http://obarring.web.cern.ch/obarring/workarounds.htm>
 - Meeting held 8/8/06 to complete and discuss the list of workarounds and decide priorities (0 == low priority, 10 == top priority)

bug	workaround efforts	fix efforts	fix priority (0-10)
stager leak	cron-job, minor only during upgrades	difficult: valgrind found nothing	2
stuck disk2disk copy	cron-job	maybe fixed. Remove workaround?	4 (8 for testing if fixed)
Tapes stuck in FAILED	cron-job	2 weeks	5
Thread deadlock in recaller	manual	2 weeks	3
Inconsistency in NBTAPECOPIESINFS	manual, tedious	3PW debug, 2PW for fix, 2PW test	9
Double LSF jobs for stagerPut	manual, tedious, very rare	3PW for identifying the problem?	1
rmmaster forks wild	manual, tedious	4PW	7
rmmaster without persistent state	automatic, missing SMS func	3PM (12PW)	5
GC problems (#17951)	oracle procedure	maybe fixed?	5
LSF meltdown	Limit PENDING to 1000	4PW	10



Item 9



- *The committee recommends that the development and operations teams review the list of workarounds, involving replacement of palliatives with features in the CASTOR core.*
- Actions taken

- List of workarounds prepared

<http://obarring.web.cern.ch/obarring/workarounds.htm>

Estimate at the time assuming Sebastien
'taking holidays' to only work on this

Meeting held 8/8/06 to complete and discuss the list of workarounds and
priority, 10 == top priority)

bug	workaround efforts	fix efforts	fix priority (0-10)
stager leak	cron-job, minor only during upgrades	difficult: valgrind found nothing	2
stuck disk2disk copy	cron-job	maybe fixed. Remove workaround?	4 (8 for testing if fixed)
Tapes stuck in FAILED	cron-job	2 weeks	5
Thread deadlock in recaller	manual	2 weeks	3
Inconsistency in NBTAPECOPIESINFS	manual, tedious	3PW debug, 2PW for fix, 2PW test	9
Double LSF jobs for stagerPut	manual, tedious, very rare	3PW for identifying the problem?	1
rmmaster forks wild	manual, tedious	4PW	7
rmmaster without persistent state	automatic, missing SMS func	3PM (12PW)	5
GC problems (#17951)	oracle procedure	maybe fixed?	5
LSF meltdown	Limit PENDING to 1000	4PW	10



Item 11



❖ ***Understanding of the operations support effort requires a project plan and effort tracking. The committee recommends this be instantiated and maintained.***

❖ **Actions taken**

- Attempt to find a convenient method for effort tracking taking into account:
 - Not only CASTOR2 but also all services operated by FIO/FS. Not all service mgrs are dedicated for castor2 operations
 - Common categorization with sufficient granularity
 - Collection and processing of the data

❖ **Methods tried**

- Paper + pencil
- Excel
- Web based tools
 - Twiki logbook (CASTOR2 only)
<https://twiki.cern.ch/twiki/bin/viewauth/FIOgroup/CastorOperationsLogbook>
 - System manager blog <http://wordpress.org/>
 - Timetracker <http://wrconsulting.com/> (installed but not tested yet)



Item 11



- Experience
 - All section members (except SL) takes part of the Service Manager On Duty (SMOD) rota
 - One week out of 5, estimated 100% effort
 - Some SMOD tasks
 - Front-end for user support
 - interaction with sysadmins and operators
 - Deploy regular linux and kernel upgrades
 - SMOD role implies high rate of context switching (~30/day)
 - Collection and processing of data
 - Need a web-based tool for central collection of the data
- Example (Olof), period August – October

email	meeting	castor op	support	Ext op support	development	SL	Other
8%	18%	25%	13%	13%	5%	6%	12%

c2alice	c2atlas	c2cms	c2lhcb
0	8%	20%	0.5%



Item 12



- ❖ ***The committee recommends that the project organizes a meeting with the 4 LHC experiments in order to reach agreement on a set of requirements and priorities for Service Monitoring Tools and a plan for their implementation.***
- ❖ Requirements gathered during summer. All 4 LHC experiments contributed
- ❖ Monitoring meeting with experiments 10/11/06
 - <http://indico.cern.ch/conferenceDisplay.py?confId=8299>
- ❖ Categorization of metrics
 - Internal metrics (status of requests and other db objects)
 - 29 metrics (more than 140 values), 45k measurements/day, 200k values/day
 - Not including daemon monitoring, parse log msgs, ...
 - Service metrics
 - Availability and performance
 - Accounting metrics
 - Usage
- ❖ Experiments also wants to be able to query for problems with a particular file or request
 - Not monitoring but needs to be addressed by other means, e.g. open access to the DLF GUI



Item 12: work in progress



- ❖ Castor monitoring TWiki:
<https://twiki.cern.ch/twiki/bin/viewauth/FIOgroup/CastorMonitoring>
- ❖ List of implemented metrics
 - Only internal and service metrics so far
- ❖ Proposals for accounting metrics
 - top 20 users: number of disk files owned, accesses (Get,Put,etc), jobs running/pending
 - top 20 tape users: avg filesize per user
 - Distribution of file sizer per svcclass, directory?
 - tape files written/read per diskpool (with size avg and sum)
 - Get,Put,PreparetoGet,PrepareToPut information
 - Volume of data on disk,tape, disk and tape (?), just disk(?)
 - Number of files on disk,tape, disk and tape (?), just disk(?)
 - Monthly full user accounting (number of files, transfers, etc)
 - Distribution of lifetime of files (0,1,2,3,4,5-7,>7)
 - Distribution of number of reads per file



Item 12: Implemented metrics



Internal metrics

Service metrics

Name	Metric ID	description	deployment
Segment	6240		currently not active
Stream	6241		active on srv05's
Subrequest	6242		active on srv05's
Tape	6243		active on srv05's
TapeCopy ²	6244		active on srv05's
DiskCopy ²	6253		active on srv05's
DiskpoolInfo ²	6254		active on srv05's
MigStall	6255		active on srv05's
RecallStall	6256		active on srv05's
DLFErrors	6258	top 5 errors in last hour	active on srv02's
OrphanTape ²	6257	kills orphaned recaller processes	active on c2atlassrv02
JobsPending ²	6259	number of running, pending and stalled jobs	active on srv02's
OldJobsPending ²	6264	pending jobs older than given time	awaiting deployment
AccessCountDist (new)	6264	returns svc, interval, file count	c2publicsrv05
FileSizeDist (new)	6265	returns svc, interval, file count, file space	c2publicsrv05
FileAgeDist (new)	6269	returns svc interval count	c2publicsrv05
Job	6245		c2publicsrv05
Request	6246		c2publicsrv05
Ethernet TrafficCounters per destination (lptables)		6260	
MigratorLog (new)			
MigLogSvc		6266	c2publicsrv02
MigLogPool ²		6267	
MigLogAll		6268	c2publicsrv02
MetaMigPending	6247	pending files by svcclass - min time, max time, avg time, count	c2publicsrv05
MetaMigSelected	6248	migrating files by svcclass - min time, max time, avg time, count	c2publicsrv05