

Post v2.1.13 Scenarios

Giuseppe Lo Presti
CERN IT DSS

CASTOR Face-to-face Workshop - 28-30 Nov 2012

- Work plan for LS1 (and beyond?)
 - Xroot plugin
 - Disk scheduling improvements
 - Disk draining
 - SRM
 - Operational tools
- Possible future scenarios
 - **Depending on some strategic choices!**



- Reimplement the xroot plugin
 - Make xroot 'the' native protocol for CASTOR
 - Also for tape access (done)
 - Also for disk-to-disk copy
- Based on a new diskmanager
 - the xroot daemon is told which file to serve to which client
 - Extra bonus: drop stagerjob
- Allows for traffic control/shaping
 - But... do we (still) care?
- Allows for deprecating RFIO
 - Ultimately only secure protocols to access CASTOR

- The NS API has been secured
 - Andrea will talk later about current deployment and future options about securing 'the rest'
- Transport protocols
 - RFIO and ROOT to be deprecated/dropped, only XROOT and GSIFTP (both secured) remain
- APIs, admin commands, ...
 - The stager API won't be secured for now
 - All internal protocols and admin commands won't change
 - Risk (today!) of someone hijacking e.g. the transfermanager-diskmanager channel to... at best screw up things

- Improvements in the disk scheduling system
 - Once rmmaster has been integrated into TM... (see Sebastien's slides)
 - Read-only support on FileSystem/DiskServers
 - Multi-level queue
 - Priority (already there), Normal, Draining
 - Priority always wins, Draining will get through at a given (small) ratio versus Normal to avoid starvation
 - Support for killing running jobs

- Current system designed with LSF limitations in mind
 - Today, much more complex than needed
- Simplify the logic to exploit new TM features
 - Ability to stand large queues
 - Usage of the draining queue
- In practice
 - Submit all draining disk-to-disk jobs to the TM in one go and let it sort out the queue

- Port SRM to SLC6
 - This involves porting the code to the new Globus/gSOAP packages from EMI
 - Some APIs changed
 - **The gSOAP layer has to undergo a major refactoring**
 - Allows for supporting VMs
 - Current issues with runaway clocks on SLC5 makes deployment of services on VMs at least unwise...
- A few fixed bugs pending deployment
 - None urgent...

- New c2probe (contributed by Xavi) finalization and deployment
- Tool to undelete tape files
 - by NS logs parsing
- Tool to administer disk copies
 - cleanLostFiles refactoring



- Nameserver schema minor enhancements to ease operations
 - e.g. timestamps on segments
 - Held back until slot for Nameserver intervention
- Some 'nice to have', at least for developers
 - Integrate nightly upgrades/tests in TeamCity
 - Coverity/helgrind
 - Migrate savannah to jira
 - Migrate svn to git
- Lesser priority compared to the previous items

- Some decisions may change priorities quite drastically. Case of traffic shaping:
 - A) traffic shaping is proven to work with iptables
 - Part of the work around xroot will drop in priority, no need to make xroot handle all other protocols
 - Disk scheduling and draining will be top priority
 - B) we still need xroot for traffic shaping
 - It is a priority to make xroot 'the' protocol, rfi to be quickly replaced with xroot, root to be dropped, gsiftp 'bridged' to xroot

- IPv6 support? (for external clients only)
 - Will be needed – but when?
- Publish MacOS X clients?
 - Used by experiments
 - Only testing is missing
- HTTP support?
 - Coming for free with xroot 4 (?)
 - Being discussed in the context of federations
(e.g. <https://indico.cern.ch/conferenceDisplay.py?confId=218328>)

- DiskServer rebalancing?
 - Exploiting the new (and not-yet-developed!) draining logic
- Multi-replicas on JBOD?
 - Software is partially there, missing automatic recovery of lost replicas
- ...???

Thank you!
Questions? Remarks?

