



DPM : what's missing

Ronald Starink & Jeff Templon

NIKHEF

Pre-GDB meeting

09 january 2007



Note

- ◆ We run a site supporting wLCG computing
- ◆ And other non-wLCG stuff
- ◆ Dropping support for volatile storage is not an option

First off

- ◆ DPM works!
- ◆ Support for pools, user mapping, ACLs, multiple backend servers
- ◆ Generally verbose logging, in most cases with relevant info
- ◆ Only one instance of “completely bizarre” behavior
- ◆ Extensive wiki troubleshooting page, DPM admin guide (also frequently updated)
- ◆ As far as grid middleware goes this is pretty good!

Missing things

- ◆ Management
- ◆ Management
- ◆ Management

Incomplete support for storage types

- ◆ ATLAS has dedicated durable pool
- ◆ ATLAS can also write to generic volatile pool (see next bug)
- ◆ We publish durable for durable and volatile for volatile
- ◆ ATLAS selects a path and tries to write, but it is full
- ◆ File automatically goes to volatile pool ... No way to stop this!!
- ◆ Need to be able to specify storage class (no middleware support!!) and have DPM respect it ... And/or associate a path to a specific pool.
- ◆ GGUS #16646

ACL support : one or none

- ◆ Previous example would be less of a problem if we could prohibit ATLAS from writing to the volatile pool
- ◆ In DPM, we have the following choices
 - Pool is accessible to one specific group only
 - Pool is accessible to all groups
- ◆ Need the ability to provide a list of accepted, or prohibited, groups per pool
- ◆ GGUS #16642

ACL support : one group per user

- ◆ With VOMS, a particular user may belong to multiple groups
- ◆ This is not supported in DPM : user is mapped to one single group
- ◆ Hence storage only works as desired for one of the user's roles
- ◆ Example
 - user has "production" and "normal user" roles in that order
 - User will have access to ATLAS production pool, but not ATLAS normal user pool.
- ◆ GGUS #16643

Other Trouble Spots

- ◆ Documentation on architecture / design / message flow is very hard to find (Athens EGEE-I presentation?)
- ◆ Support for migration of files between pools, but this is so slow as to be unusable (migrates ~ 20 GB/day)
- ◆ Stability is sometimes an issue (likely due to GGUS #13998) ... Sometimes few restarts per week, sometimes weeks without restart
- ◆ Logging is more developer-oriented and it is not possible to configure the logging verbosity