



Access control for storage

- What we have now
- What we can expect
- Job priorities model for SE?
- What can be reported/accounted?
- Monitoring
- Quotas



What we have now (1)

- Files owned by normal users are group-writable (for LHC VOs)
- Production manager DNs usually mapped to special accounts
 - Can write in directories that are read-only for normal users
 - Such directories usually need to be set up by the SE admin
- DPM has almost full VOMS support
 - Secondary FQANs not yet taken into account
 - Final testing being done for name space and space tokens
 - Use for pool selection being discussed with sites (NIKHEF, CIEMAT, ScotGrid)
 - ACLs can be set via DPM client API/CLI
 - ACL support not yet discussed in SRM v2.2 group (low priority)
- CASTOR has minimal VOMS support
 - Role in primary FQAN required for space operations in SRM v2.2
 - Mapping via classic grid-mapfile

- dCache has significant VOMS support
 - Via gPlazma plugins
 - PRIMA → GUMS setup demonstrated (in use in OSG?)
 - Enhanced grid-mapfile support
 - DN + FQAN → account A
 - "*" + FQAN → account B
 - No LCMAPS plugin yet, not clear who should write it
 - Secondary FQANs not yet taken into account
 - ACLs not yet available (foreseen for Chimera/NFSv4)

- StoRM has full VOMS support
 - Via LCMAPS
 - ACLs can be set through native file system API/CLI
 - GPFS, XFS, ext3

- BeStMan (DRM) has no (?) VOMS support yet



What we can expect

- Grid-wide consistent VOMS-ACL support not for this year
 - Need to use implementation-dependent client tools
- Distinction of special users by primary role/group instead of DN?
 - DPM, StoRM
 - OK
 - dCache
 - Feasible
 - CASTOR
 - Low priority
 - BeStMan
 - ?
 - Not really needed this year, but what about next year?
- Space tokens protected with ACLs?
 - Space is owned by group that must match primary FQAN
 - Not clear how much more we can do this year



Job priorities model for SE?

- Job priorities on CE
 - Give privileged groups/roles a better QoS
 - Larger share (job slots)
 - Better response time (express queues)
 - No restrictions (yet) on number of jobs per DN/role/group
- Ideas for service priorities on SE
 - Give privileged groups/roles a better QoS
 - Larger share (dedicated pools and/or quotas)
 - Publish per role/group
 - Better response time (more transfer slots)
 - Higher bandwidth better done in FTS?
 - Most of the bandwidth is used by production managers anyway
 - No restrictions (yet) on number of transfers per DN/role/group
 - Let FTS handle that?

- Close relationship between monitoring and accounting
 - Need to keep track of resource usage for both purposes

- Accounting mostly concerned with disk/tape usage per entity
 - Per storage class instance
 - VOMS group/role
 - DN
 - DN + group/role?

- Status of storage class instances
 - Also published in information system
 - Used/reserved/available space (with caveats)
 - Availability of the instance
 - Could be unavailable due to problem or maintenance



Monitoring

- Also concerned with traffic statistics
 - Number and size of transfers per period
 - Upload, download
 - Stagein, stageout
 - WAN, LAN
 - Protocol
- How much of that should be accounted as well?
- Grid Storage Systems Deployment WG has monitoring subgroup
 - Draft report compiled listing possibilities for DPM and dCache
 - Coordinate with GDB storage accounting WG?

- Quotas will be implementation-dependent
 - Not yet accepted as SRM feature

- Static space tokens used for first-order separation
 - Normal users \leftrightarrow production managers
 - Different analysis groups
 - Real \leftrightarrow MC data
 - ...

- Quotas are complementary
 - Could be overbooked