

OSG Storage Elements and ATLAS DDM

Marco Mambelli, Facilities meeting, Aug 12, 2008

Outline

- ▶ TiersOfATLASCache.py
- ▶ Client programs
- ▶ Site survey



Introduction

- ▶ ATLAS DDM (Distributed Data Management) is DQ2
- ▶ ATLAS SEs (Storage Elements) used by DDM are listed in TiersOfATLASCache.py, python file stored in CVS
- ▶ <http://atlas.web.cern.ch/Atlas/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py>
- ▶ SEs are managed using SRM protocol, mainly
 - ▶ SRM – dCache
 - ▶ SRM – Bestman over FS (Xrootd, Lustre, GPFS, ...)
- ▶ DQ2 uses 3rd party software to move files, mainly
 - ▶ lcg-cp
 - ▶ srmcp
 - ▶ ngcp



Client programs

- ▶ **DQ2 User Client**

- ▶ dq2-ls, lists the datasets and the files in the dataset
- ▶ dq2-get, copies locally a dataset or part of it
 - ▶ -p option allows to choose executable used to transfer files

- ▶ **Data movement clients (from a OSG perspective)**

- ▶ they support space tokens, SRM, Gridftp, local files
- ▶ lcg-cp, requires either BDII access or full URLs and SRM protocol specification
- ▶ srmcp, requires URLs with port number, SRM version and 'file:////' URLs
- ▶ ngcp, part of NorduGrid suite (not in VDT or gLite)



URLs

▶ SRM

- ▶ dCache full (extended) URL:
srm://hostname:port/srm/managerv2?SFN=abs_path
- ▶ Bestman full (extended) URL:
srm://hostname:port/srm/v2/server?SFN=abs_path
- ▶ Short URL (used in SRM v1):
srm://hostname:port/full_path
- ▶ Some SE use port 8443 as default port

▶ File URL

- ▶ file:///abs_path
-



URLs in TiersOfATLASCache

- ▶ there are 370 SRM entries
- ▶ 253 use the extended URL for SRM v2 (/srm/managerv2?SFN), none uses it for SRM v1 (/srm/managerv1?SFN)
- ▶ 18 use the Bestman version of the extended URL for SRM v2 (/srm/v2/server?SFN), none uses it for SRM v1 (/srm/v1/server?SFN)
- ▶ 142 specify explicitly port 8443
- ▶ 127 specify explicitly port 8446
- ▶ probably 100 URLs do not specify a port
- ▶ 264 use a token.
- ▶ Some token names are ATLASDATADISK, ATLASDATA, DATADISK, ATLASMCDISK, MCDISK, ATLASUSER



TiersOfATLASCache and File Catalog

- ▶ These base URLs are used for new file registration
- ▶ Existing PFN in file catalog may have different URLs

- ▶ Add the port number to all URLs without it
- ▶ If a SE supports only SRMv1 or only SRMv2 it is recommended to use the extended URL with the correct manager
- ▶ Full SRM URLs are preferred



Data transfer client test

- ▶ Compares only usability (with the right URL all programs work)
- ▶ File transfer attempts with different URLs
- ▶ Commands used:
 - ▶ 'globus': 'globus-url-copy'
 - ▶ 'lcg': 'lcg-cp'
 - ▶ 'lcg2': 'lcg-cp -b'
 - ▶ 'lcg3': 'lcg-cp -b -T srmv1'
 - ▶ 'lcg4': 'lcg-cp -b -T srmv2'
 - ▶ 'lcg5': 'lcg-cp -b --vo atlas'
 - ▶ 'nordugrid': 'ngcp -d 2'
 - ▶ 'srmf': 'srmcp'
 - ▶ 'srmf2': 'srmcp -2'
 - ▶ 'srmf3': 'srmcp -l'
 - ▶ 'srmf4': 'srmcp -streams_num=1'
 - ▶ 'srml': 'srm-copy'
 - ▶ 'srmg': 'glite-srm-get'



Results

Tests	srmV2BNL 2UC	srmV1	srmV1BNL	srmV1np	srmV2	srmV2f4	srm3pV1	gsifp	srmV1BNL 2UC
lcg								OK	
lcg2								OK	
lcg3								OK	
lcg4					OK	OK		OK	
lcg5								OK	
globus								OK	
srmg									
srml					OK	OK		OK	
srmf			OK				OK		
srmf2			OK				OK		
srmf3			OK				OK		
srmf4			OK				OK		
nordugrid	OK	OK	OK	OK	OK	OK	OK	OK	OK



Transfer test

- ▶ File

 - fdr08_run2.0052304.physics_Jet.merge.AOD.o3_f8_m10._0020.l

- ▶ From dataset

 - fdr08_run2.0052304.physics_Jet.merge.AOD.o3_f8_m10

- ▶ From different SE to tier2-06.uchicago.edu

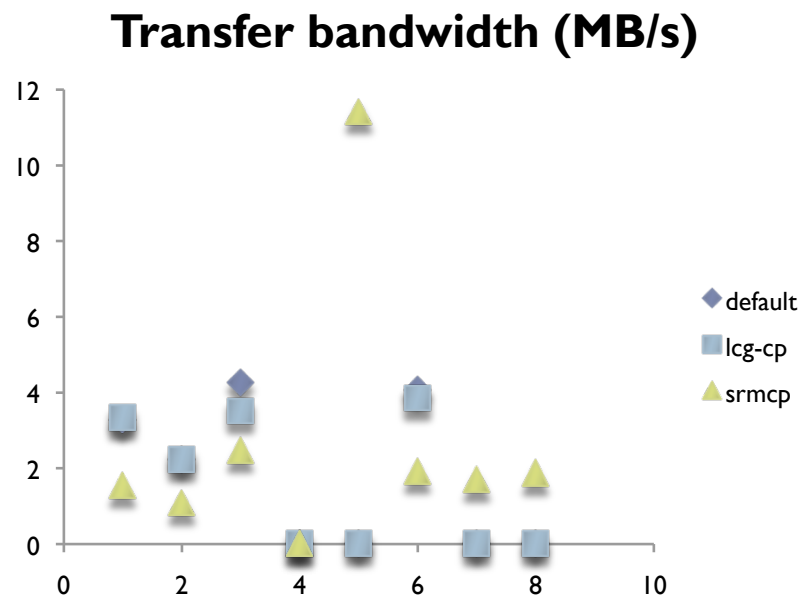
- ▶ Clients wlcg-client 0.13rc and dq2-clients 0.1.17

- ▶ SE involved

 - ▶ SWT2_CPB_DATADISK, SLACXRD_DATADISK, WISC, AGLT2_DATADISK, MWT2_DATADISK, NET2_DATADISK, BNL-OSG2_DATADISK, BNLXRDHDD I



Results



- ▶ File too short to evaluate performance
 - ▶ comparative study
 - ▶ failed transfers
- ▶ Only one test timed out (transient SE problems, worked later)
- ▶ dq2-get exit code was 0 also when failing ([Errno 2] No such file or directory)
- ▶ lcg-cp works only with complete URLs (no BDII involved) and is generally faster (3-4 M/s)
- ▶ srmcp is generally slower (2 M/s) but works also with short URLs
- ▶ srmcp from UofC (5) faster but transfer is local



Questions

- ▶ Should TiersOfATLASCache be modified?
- ▶ Should existing PFNs in a file catalog be modified?

