SRM v2.2 and related issues

Flavia Donno

CERN, WLCG GDB, 2 May 2007

SRM v2.2 Issues

Experiment	Issue 1	Issue 2	Issue 3	Issue 4
CMS	VO Namespace Support	Explicit delete	srmCopy in push mode	Access control through VOMS
ATLAS	Quota	SRM v2.2 testing environment	Accounting	
LHCb	Staging problems: BringOnline through GFAL	File pinning through lcg-utils	"Resource busy" from CASTOR	Lcg-utils < 50% efficiency
Alice				

SRM v2.2 Status

- **Main implementations tested for the functionalities requested.**

- **VO name space support and mapping to Storage Classes (CMS)**

  /my_exp/data/RAW/year/run/reco_pass/stream
  /my_exp/data/DST/year/run/reco_pass/stream
  /my_exp/data/ESD_master/year/run/reco_pass/stream
  /my_exp/data/ESD_replica/year/run/reco_pass/stream
  /my_exp/data/AOD/year/run/reco_pass/stream

  RAW --> **T1D0**
  DST --> **T1D0**, but different tape set
  ESD_master --> **T1D1**
  ESD_replica --> **T0D1**
  AOD --> **T0D1**

- srmRm implements explicit delete (CMS)

- srmBringOnline exposed through GFAL allows for staging capabilities (LHCb)

- File pinning not supported by all implementations – CASTOR (LHCb)
  - Therefore, file pinning not exposed through lcg-utils
  - VOs have to agree on the correct sizing of disk caches and tuning of the file garbage collector

SRM v2.2 Status

- Good support for srmCopy in push mode (CMS)

Summary of S2 SRM v2.2 cross test - Wednesday 2 May 2007 07:34am CEST

In these tests the srmCopy function is exercised. This function should be implemented by all available Storage System by the end of the 3Q of 2007. dCache is required to implement this function as of now. Therefore, it is OK to have red columns for all SRM endpoints except for dCache. However, it is not OK to have red rows since this means that a file cannot be copied between SRMs with simple get and put operations.

SRM function	CERN C2	DESY dCache	FNAL dCache	CERN DPM	LBNL BeStMan	CNAF StoRM
Copy Tests in PUSH mode						
CopyToCERNCASTOR	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToFNALDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToDESYDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToCERNDPM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToLBNLDRM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyToSTORM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
Copy Tests in PULL mode						
CopyFromCERNCASTOR	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromFNALDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromDESYDCACHE	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromCERNDPM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromLBNLDRM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log
CopyFromSTORM	Out Log	Out Log	Out Log	Out Log	Out Log	Out Log

CASTOR and DPM Do not provide srmCopy for the moment. Foreseen for end of 2007.

dCache@DESY Authorization/ Connection errors (dCache developers will look into the problem this coming week)

StoRM does not provide srmCopy in PULL mode

SRM v2.2 Status

- **Capability to control access rights consistently across different SEs through VOMS (CMS)**
  - We need to clearly understand the requirements and use-cases in order to correctly define the functionality and the interface in SRM 2.2.
  - Support for ACLs through VOMS groups/roles is at the moment **present in DPM and StoRM**.
  - **dCache** expects to provide such feature in production **early 2008**.
  - **CASTOR** probably later (see Tony's presentation).
- **Quota (ATLAS)**
  - It will not be available for 2008.
- **"Resource Busy" message from CASTOR due to corrupted entry (from previous transfers timed out or failed) that CASTOR (for consistency) refuses to overwrite (LHCb)**
  - This is what was agreed and the same behavior is exposed with SRM v2.2. It is the responsibility of the clients to Abort failed request and remove the corresponding SURLs before retrying. This is the behavior used in high-level tools such as FTS.
  - Overwrite mode available in SRM v2.2. However, SRM_FILE_BUSY is returned to signal the previous failure.

SRM v2.2 available for the experiments to test

- S2 test suite test-bed contains as of today 14 endpoints in all flavors (CASTOR, dCache, DPM, StoRM, BeStMan)
  - https://twiki.cern.ch/twiki/bin/view/LCG/GSSDendpoints
  - Configuration and site specific settings have been discovered during the test phase
  - Working with developers to document these issues and make the installation and configuration process easier and less error-prone (gPlazma configuration problems observed by LHCb)
- Same endpoints available in FTS 2.0 pilot and in EGEE pre-production test bed

SRM v2.2 available for the experiments to test

- Working with sites to understand the correct setup per VO (Paths, Storage Classes, Space Tokens, disk space requirements, etc.)
  - For instance at CNAF T0D1 served by StoRM SE in production and T1Dn served by the CASTOR instance

- Latest version of clients available
  - lcg-utils v1.5.1-1, GFAL 1.9.0-2, FTS 2.0 (to be tested first with SRM v1)
  - lcg-utils allows for copy operations to use SURLs without contacting the catalogue (lcg-utils efficiency < 50% reported by LHCb)

- It is very important to have the **experiments on the pre-production test-bed** testing the environment as soon as possible in order to understand if SRM v2.2 is ready for production

Storage Accounting

- Effort started within EGEE by UK
- At the moment space used/total available per VO is published
- It is possible to publish also information about Storage Classes if they are published by the information system (available in GLUE Schema v1.3)
- It is possible to account information on space usage by FQAN if published by the information system (GLUE Schema v1.3)
- Within GSSD a report is being compiled about the information retrievable now from the different storage systems.
  - It will be circulated to the experiments for comments.
  - A wrapper can be made available to provide a common interface