Pre-GDB Storage Classes summary of discussions

Flavia Donno

Pre-GDB - CERN, 9 January 2007

Summary

- Status of SRM 2.2
- *Summary of previous discussions and conclusions*
- *Very interesting reports from NIKHEF and T2 sites*
- Tentative plan for next WLCG Workshop in January
- The new working group

Status of SRM v2.2

- *Tests:*
  - S2 test suite just ported to new published WSDL (minor changes), result published. Cron job will be restarted as of today.
  - LBNL SRM-Tester also run as a cron.
  - List of open issues kept updated by CERN. Many will be deferred to v2.3 or v3.0.
- *Servers:*
  - Current endpoints sometimes hard to test since they use a production instance behind (hard to exploit and stress the systems).
  - Only one end-point available per type.
- *Clients:*
  - SRM client code unit-tested and mostly integrated in FTS
  - FTS in development testbed by end of February
  - GFAL/lcg-utils under certification
- *GLUE schema:*
  - V1.3 ready for deployment by end of Jan.; static information providers in YAIM ready by end of Feb.

Summary of Storage Class discussions

- Supported Storage Classes in WLCG are: *T1D1, T1D0, T0D1*.
  T0D0 might be required by other VOs
- Clients should *not use dynamic reservation* initially.
- *Storage Areas are created statically* by Site Administrators at sites. VO administrators can use srmReserveSpace to create Storage Areas.
- *Tools will be available to publish VO specific Space Descriptions (ATLAS_RAW, CMS_TAG, etc.)*
- *The Namespace should be used to organize the data.*
- *Brief review of the T1s current implementations and plans (GridKa, SARA, Lyon). Main issues are about configuration of WAN/LAN buffers, storage areas organizations, storage classes, management tools and clear procedures in case of failures from both developers and experiments.*
  - *Input needed from the experiments to understand requirements and data flow*
  - *Input needed from developers for the configuration of the specific implementations*
- *A tentative plan for the WG with future targets has been presented (see later)*

Sites

- *NIKHEF:*
  - Very good experience with DPM
  - Discovered few limitations while supporting ATLAS:
    - Missing management tools
    - Impossibility to force the system not to use the volatile pool for ATLAS durable files
    - Some minor problems with ACLs (a production manager cannot write in a user pool)
- *USCMS T2s:*
  - SRM-dCache
  - No Tape backend
  - WNs used as storage elements
  - Not using groups/roles nor storage classes at the moment.
  - Different configuration of pools for WAN or LAN
- *GridPP:*

Sites

- *GridPP T2s:*
  - No tape backend and small RAIDs very often spread on WNs (~500!)
  - Shared resources with non WLCG VOs
  - Limited manpower (=> easy management)
  - Management tools are essentials: pool draining, namespace management, disk quotas, resize disk pools (increase or decrease)
  - Is data loss a real issue ? What about analysis ? Procedures to manage this situation
  - Hints on hardware and software configuration and tuning (raids, power supplies, communication channels, databases, backups, firewall problems, etc.)
  - Storage systems limitations (files open/sec, read/write rates, etc.)
  - System monitoring
  - Tests to check site availability (SAM tests depend on other subsystems: BDII, catalogues, etc.) also at a pool level per VO.
  - Is it important to publish the real size of a pool (available and used) ?
  - Storage accounting
  - 32 vs. 64 bit installations.
  - Which storage classes ?
  - Dynamic or static reservation ?
  - Interoperability tests ?
  - What to do in case of not used/corrupted data sets? What about full disk pools ? Empty pools ?
  - What are the procedures in case of SRM failures (clear and written) ?

Plans

- Collect requirements per VO in terms of:
  - Storage Classes needed at various sites: how much disk for each?
  - Data Flow between Tier-0, Tier-1s, Tier-2s
  - Space reservation requirements: static and/or dynamic
  - Space Token Descriptions per VO: which ? How many ? Transition patterns ?
  - Special requirements: xrootd, etc.
  - Data Access Patterns.
  - Plans from 1st of April 2007 till end of the year
- Understand requirements and existing setup of T1s and T2s (ongoing effort). Assisting T1s and T2s.
- Discuss with the developers on how to best implement the requirements
- Produce manuals and guidelines for deployment
- Selecting sites as guinea pigs
- Testing with experiments the setup at sites
- Assisting the experiments and the site administrators during the tests. Support for possible further requirements/needs

- Review the status of the plan at the WLCG workshop in January