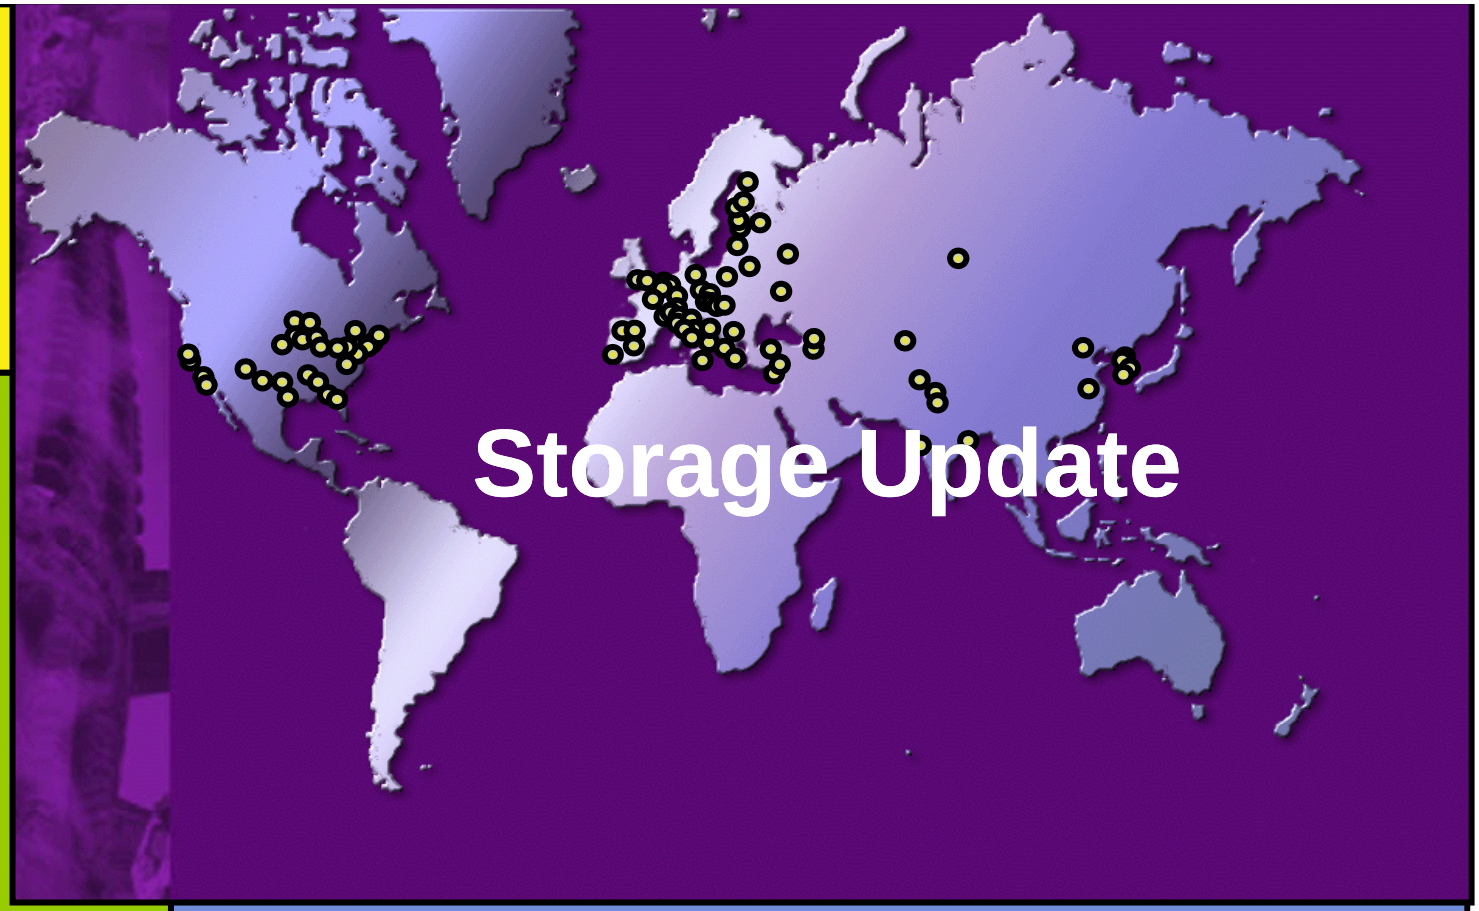

Storage Update

Flavia Donno

CERN/IT

WLCG Grid Deployment Board,
CERN 11 June 2008

Overview

- The ***goal of the SSWG***
- SRM v2.2 Addendum to the Usage Agreement
  - ***Short and long term plans***
- ***Implementation schedule***

The Storage Solution WG

- The **goal of the SSWG** is
  Address issues uncovered through the challenges and provide timely solutions
- **This is achieved with**:
  - Managerial phone conferences where experiments, developers and site administrators are represented
  - Focused technical (daily) phone conferences with the involved bodies
  - Report on outstanding issues and prioritization
    - It is hard for this to proceed entirely sequentially - priorities might change with time, but the goal should remain the same.
  - Establishment of operational strategies to provide reliable services
- Detailed discussions on experiences in CCRC'08 will take place during the workshop.
  - This includes release / patch handling, dependencies between different components etc.

The experiments requirements

- The **goal of the SRM v2.2 Addendum** is to provide answers to the following (requirements and priorities given by the experiments)
- **Main questions** :
  - Space protection (VOMS-awareness)
  - Space selection
    - What kind of spaces (T1D0, T1D1, T0D1)
    - Pre-staging
  - Space usage statistics
  - Tape optimization (reducing the number of mount operations, etc.)
- **Main dates** are being adhered to:
  - ***10 June***: submission to the MB for approval. Start of the implementation phase

We must remain in close contact with each other to find timely and acceptable solutions

The document and validation

- The most recent version is v1.3 and can be found here: http://indico.cern.ch/getFile.py/access?contribId=2&resId=0&materialId=0&confId=34806
- 2 main parts:
  - A detailed description of an implementation-independent full solution
    - Implementation plans per storage solution are also detailed with dates for each of the features.
  - A summary of what can be achieved by the end of the year per implementation
- The document has been "signed" by *storage developers, clients developers, experiments* (ATLAS, CMS, LHCb)

The MoU Addendum

- The long term solution has been described during the GDB the 13th of May 2008
  - More details added by some of the experiments in order to clarify concepts and behaviors of functions

Top priority is services basic functionality and reliability

- The short term solution is different depending on the specific implementation but aims at providing the same functionality
  - ***NOTE:*** *The short term solution will allow this group to better understand the experiments requirements and the feasibility of the long term solution proposal. In case the short term solution demonstrates to already address the requirements and use cases of the experiments, then this group might decide to either not implement or partially implement or revise the full proposal.*

The MoU Addendum

- ***CASTOR***
  - ***Space Protection*** based on UID/GID.
    - Administrative interface ready by 3rd quarter of 2008.
    - **Full VOMS-awareness** in 3rd quarter of 2009.
  - ***Space selection*** already available.
    - srmPurgeFromSpace available in 3rd quarter of 2008.
  - ***Other points:***
    - *srmLs returning space tokens by 4th quarter of 2008*
    - (***T1D1*** provided) pinLifeTime parameter on Put operations is negotiated to be always equal to a system defined default. No change is envisaged in the future for this behavior
    - No srmCopy implementation foreseen

The MoU Addendum

- ***dCache***
  - Short-term solution
    - Space selection based on the IP number of the client, the requested transfer protocol or the path of the file. Use of the TExtraInfo structure for more refined selection.
    - Protected creation and usage of “write” space tokens
    - Allowing or denying the access to the tape system for particular DN's or FQAN's.
    - (***T1D0 + pinning*** provided). Releasing pins will be possible for a specific DN or FQAN without specifying the Request ID.
    - Implementation plan: 13 working days + delay due to *improve reliability and performance of the system*
  - Long-term solution
    - As per document
    - Implementation plan: 62 working days. With the addition of another staff, we estimate we could finish the program by the end of the calendar year (2008)

The MoU Addendum

- ***DPM***
  - Short-term solution
    - Support a list of VOMS FQANs for the space write permission check, rather than just the current single FQAN - *September 2008*
  - Long-term solution
    - Array of space tokens in srmLs on a single file - *October 2008*
    - No srmCopy
    - srmReleaseFiles on surls without request ID. - *January 2009*
    - Space tokens on srmPrepareToGet and srmBringOnlinerequests - *June 2009*
    - ACLs on spaces. Relevant only Read-From-Space, Write-To-Space, Replicate-From-Space, Purge-From-Space – June 2010

The MoU Addendum

- ***StoRM***
  - Short-term solution
    - Spaces in StoRM will be protected via DN or FQAN based ACLs. StoRM is already VOMS-aware. Improvements in the permission manager.
    - No tokens on Get
    - Path and space tokens can be passed to the tape back-end for tape usage optimization, if needed
    - Implementation plan: ready by November 2008
  - Long-term solution
    - As per document
    - 8 men/months of work

The MoU Addendum

- **Client tools: FTS, lcg-utils/gfal**
  - Token on Get operations
    - The client tools will pass both the SPACE token and the TExtraInfo structure . The number of new keys is restricted to a minimum and well defined set.
  - Pinlifetime on Put/Copy
    - Client tools will internally extend the pinlifetime (srmBringOnline) of a newly created copy if the pin lifetime is specified on put or copy operations by the client application.
  - Space Types
    - The same type of SPACE might be implemented as T1D1 for CASTOR, or T1D0 + pinning in dCache. When removing a copy, the clients will perform both an srmPurgeFromSpace and an srmRelease without requestID so that the different space types are hidden to the user.
  - Implementation plan: coding will start in September 2008 and will last for about 2 weeks. Then certification and pre-testing activities will follow before deployment in production.

Thank You

WLCG Grid Deployment Board,
CERN 11 June 2008