

Table of Contents

Storage Management TEG: Questionnaire Level 1 - Alina Grigoras.....	1
Question 1.....	1
Question 2.....	1
Question 3.....	1
Question 4.....	1
Question 5.....	2
Question 6.....	2
That's it!.....	2

Storage Management TEG: Questionnaire Level 1 - Alina Grigoras

This twiki is to collect the input of **Alina Grigoras**. Please answer the questions below. For more information, please refer to the Storage TEG main twiki.

Question 1

- In your view, what are the 3 main current issues in Storage Management (SM)?

My answer:

1. Lack of a unitary view of the entire Grid storage (optimized access, transparent fallback in case of problems)
2. Missing a universal and lightweight storage-to-storage transfer capability (such as the xrootd third-party copy method)
3. High-level storage management tools, for example a simple replication method of the type 'I want 3 replicas of this file' without explicit storage endpoint specification; storage in self-heal in the case of data loss (back replication from other storages); background consistency checks of the files on the storage; Federated data store view, making all storage look like one "big disk".

Question 2

- What is the greatest future challenge which would greatly impact the SM sector?

My answer:

Network has proved to be a much faster-pace advancing field than the storage, we need to adapt to having bandwidth in excess and find ways to exploit the excellent connectivity we now have and can expect to further improve in the near future.

Question 3

- What is your site/experiment/middleware currently working on in SM?

My answer:

The current status for ALICE experiment is:

- 52 disk SEs, 8 tape SEs : 43x xrootd, 2x DPM, 4x Castor, 3x dCache
- so far 20PB in 200M files
- all data access is done remotely via the xrootd protocol (typically from the local site storage since the jobs are sent to sites holding a replica of the input data files; location-optimized otherwise)

Question 4

- What are the big developments that you would like to see from your site/experiment/storage system?

My answer:

The remote data access has proved to be very versatile and the right approach to data access. We need to

improve the overall reliability of the system, from better understanding of individual data servers problems to graceful management of failures, for example transparently using other replicas where available. This calls for more detailed infrastructure monitoring data being exposed to the experiments for prompt decision-taking for data placement / replica choosing / job scheduling. It would be even better if all this was implemented in an abstract, standard way like a posix direct mount of the entire Grid storage.

Question 5

- In your experience and area of competence, what are the (up to) 3 main successes in SM so far?

My answer:

1. Remote file access from the jobs to any file from any storage.
2. A uniform namespace across all storage elements, aggregated in a global xrootd redirector.
3. Using the central catalogue knowledge of all replicas of a file in correlation with the monitoring information (topology and status) to optimize storage usage so both reading and writing are transparent for the users.

Question 6

- In your experience and area of competence, what are the (up to) 3 main failures or things you w

My answer:

1. Custom remote access protocols are an unnecessary complication, focusing on standard (posix) file access methods would reduce the complexity of interacting with the storage.
2. Difficult to understand when and how storage issues affect jobs (low efficiency to complete job failures) due to the lack of relevant global monitoring information.

That's it!

Thanks! Feel free to edit again at any time, until the date of the kick-off meeting.

-- DanieleBonacorsi - November 2011

This topic: LCG > WLCGTEGStorage_inputAlinaGrigoras

Topic revision: r4 - 2011-11-28 - AlinaGrigoras



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback