Catalogue consistency demonstrator

Gavin McCance
Jean-Philippe Baud

WLCG collaboration meeting
London
July 2010

Problem: Cataloguing consistency

- Current consistency model is not resilient to failures
  - Storage failures lead to dangling entries to be cleaned up manually. Catalogue failures lead to orphaned files.
  - Namespace scanning for diffs is expensive (e.g. srmLs 'abuse')
- Proposed demonstrator to use reliable message (i.e. industry standard MQ) as backbone of the reliability
  - All interested catalogues can 'subscribe' for new files / deleted files
  - Eventual consistency model
  - Lost files can be broadcast on the "lost" topic to interested catalogues
  - Also for corrupted "bad" files (e.g. not readable)
- Add GUIDs to storage catalogs to remove the need for local file catalogue
  - Better to send GUID instead of SURL

1 month plan: strategy

- Strategy: Take simplest (and most useful) case
  - Storage failure ("lost X files")
  - Demonstrate very simple use-case at large scale

1. Preload single LFC catalogue with 10^7 entries x 3 replicas
2. Set up messaging with durable "lost" channel
3. Define "lost" message format
4. Write LFC-messaging adapter for "lost" files
5. Mock storage adapters (list of "lost" GUIDs) @ 2 sites
6. Send the lost messages
7. Verify catalogue contents
8. Retry with "broken network cables", crashed boxes, etc.

1 month plan: goal

- Goal: understand messaging
  - Is it a useful technology for this purpose?
  - Can we obtain the reliability in face of our ~~sabotage~~ operational problems?
  - Understand timeliness issues (eventual consistency)
  - Assess the effort used in changing software
- Decide at this point if we proceed

- Iff so:
  - Expand to multiple subscribers (more than 1 catalogue)
  - Look at other operations (add, delete)
  - Integrate “real’ storage (DPM) and assess effort for this