WLCG Jamboree – Summary of Day 1

Simone Campana, Maria Girone, Maarten Litmaath, Gavin McCance, Andrea Sciaba, Dan van der Ster

Jamboree Goals (Ian B)

- Focus on analysis use cases
  - Concerns about performance and scalability
- 2013 timescale for solution in production
- Understand integration effort (existing solutions) and development effort (new stuff)
- Model: network centric, large archives, cloud storage
  - Tape as archive, remote access, p2p
  - Data access layer (xroot/gfal,…)
  - Global namespace
- Do not forget connections with job management
  - Pilots, virtualization
- Solutions (lessons) on the market

Strawman – Key Features (Ian F)

- Improvements: e.g. in networking; data-serving & P2P
- Proposal: move closer to "©loud model"
  - Further reliance on network: remote access, improved accessibility of data
  - Tape archives should evolve to be true data archives with facilities using data cached on disk, even if pulled from remote cache.

- Need transparent handling of failures in data access
- Computing Models need to evolve – tools too!
  - Predictable T0-T1 flow, less predictable for T1-T2&T2-T2

- Metrics needed: e.g. measuring usage efficiency

Network Evolution (David F)

- Huge evolution w.r.t. expectations in 2001
  - Multiple 10Gbps T0-T1 & T1-T1 links
- Network does not come for free!
  - (P2P model should imply less data movement w.r.t. data pre-placement. To be DEMONSTRATED/modelled.)
- Setup groups to put requirements together
  - “big” T2s in the OPN?
- DM services should be network topology aware
  - and consider the coupling with job management
- Miron: design a system that does resource provisioning

Back-end Archive (Dirk D)

- Is HSM still relevant? (Expectation: N)
- File-sets? Can these be used? (Expectation: Y)
  - Closing a file-set? Can this be done?
  - Hot file-sets? (Q: how much data is active?)
- Disk-based archives? [Demonstrator?]
- On the roles of a file-system:
  - Client protocol; Cluster/parallel filesystems; no silver bullet
- True archive: need experiments to define parameters (if they know them…) – storage parameters exposed to experiments
  - Capabilities (Bernd?)

Data Access Layer – Analysis Needs (Fed C)

- "Working solutions" – with a number of known draw-backs – exposed for each experiment
  - ☹ Substantial unhappiness with current tools;
  - ☹ Multiple protocols increases release & support load;
- Desiderata (convergence on requirements):
  - Intelligent defaults (which replica to read, which SE to write on) + possibility to specify;
  - File collections;
  - Performance issues (file serving & moving);
  - Load-aware replication & brokering
  - Meta-data!
- Opportunity for a common HEP DAL! (?)

Data Transfer Use Cases (Richard M)

- Spectrum includes:
  a) Object accessed / transferred on demand
  b) Event on demand
  c) File on demand
  d) Dataset on demand
  e) Dataset scheduled on measured demand
  f) … on imagined demand

- Can we move up from f)?

Namespaces, authorization needs, quotas, catalogues – what is needed? (Ph. C)

- Global namespace
  - Hierarchical and flat
- Central catalog for file location, metadata [just for archive storage? Transient catalogs?]
  - ACLs, quotas
  - Can we shut the backdoors?
- Current system (almost) allows this
  - What's not so good is the implementation
    - SRM (too heavy for little add-on)
    - Hardware implementation
      - Number of spindles, servers for matching CPU
    - Failure recovery (application access layer)

Multi-core plans and impact on data access and management (Peter E.)

- Current model: launch 1 application per core – how to adapt to multi-core world?
- Ideally would like to schedule at node level & exploit multiple cores on single node
  - Possible optimizations for I/O, memory, access to disk, DM in general
- Timeframe: deploy & commission multi-core applications for “whole node” granularity this year (at least CMS…)

Global HOME

- “Something like /afs but which works”
  - Ubiquitous, ACLs etc. but >>200MB
- Model is a global file system.
- Industrial solutions available? (drop-boxes etc.)

Summary

- Today's sessions focused on Scene Setting, Requirements & Use Cases;
  - A Strawman for a new model of data access and management;
  - Update on networking – status and outlook;
  - Use of tape as a true archive;
  - Multi-core: impact on Data Access & Management
- Baseline: some of "MONARC" assumptions no longer valid + technology advances & short-comings (plus successes!) of current solutions prompt a re-evaluation of these assumptions
- Tomorrow: drill-down into experience talks + more concrete discussions on possible ways forward…

- "Jam" photo on the bridge before coffee!

- … and J "birthday" T