CCRC'08 Status

WLCG Management Board

November 27th 2007
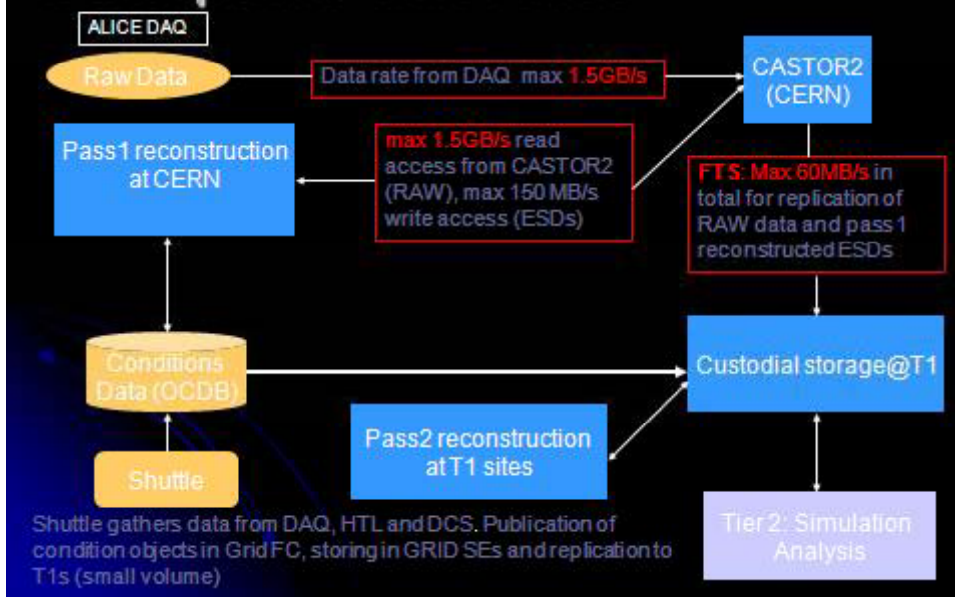
Jamie.Shiers@cern.ch

Overview

✓   Conclude on scaling factors – done

✓   Conclusions on scope / scale of February challenge (resource limited)

- We have feedback from the sites and can now work this out: (partial) reprocessing at sites with sufficient resources; read pass over data files including calibration DB lookup for those without

➢ **Conclude on SRM v2.2 storage setup details**

- Experiments to come with details on storage classes by Tier at December F2F
- This will defined what is available for pre-challenge testing in January and the February challenge itself – iteration possible prior to May – practical experience certainly needed

➢ **'Walk-throughs' by experiments of 'blocks', emphasising "Critical Services" involved and appropriate scaling factors**

- Details on next slide (hidden) – this is perhaps best combined with the above?
- Essential for sites to understand which are the services involved and the necessary parameters for sizing / configuring the services appropriately
- Also needed to understand how well we do wrt the 2008 requirements

▪ Monitoring, logging and reporting the (progress of the) challenge

- First discussion during this week's WLCG Service Reliability workshop, conclude during F2F and GDB (possible iteration in January?)

▪ CDR challenge in December – splitting out 'temporary' (challenge) and permanent data

▪ Other tests that can be done prior to February??

💣 **De-scoping? How do we negotiate this if required?**

Data paths and rates

ALICE DAQ

Raw Data → Data rate from DAQ max **1.5GB/s** → CASTOR2 (CERN)

Pass 1 reconstruction at CERN ← max **1.5GB/s** read access from CASTOR2 (RAW), max 150 MB/s write access (ESDs)

FTS: Max 60MB/s in total for replication of RAW data and pass 1 reconstructed ESDs

Conditions Data (OCDB)

Custodial storage@T1

Pass2 reconstruction at T1 sites

Shuttle

Tier 2: Simulation Analysis

Shuttle gathers data from DAQ, HTL and DCS. Publication of condition objects in Grid FC, storing in GRID SEs and replication to T1s (small volume)

ATLAS Scaling Factors

- T0 rate: 200Hz
- T0->T1 Traffic: 1020 MB/s
- T1->T2 Traffic: 10-40 MB/s depending on T2
  - 5-20 from real data and 5-20 from reprocessing
- T1->T1 Traffic: 40 MB/s
  - 20 from ESD + ~20 from AOD.
  - This assumes everybody will reprocess in Feb. CCRC08
- Job Submission at T1: 6000 Jobs/Day (over all T1s)
- MC Simulation: 20% of RAW data = 30Hz = 2.5M Events/Day = 100K simulation jobs/day + 10K reco jobs/day at T1

Basic Scaling Items to Check in CSA08

Service	CSA08 Goal	CSA07 Goal	CSA06 Goal	Status 2006
Tier-0 Reco Rate (Hz)	150 - 300	100Hz	50Hz	Achieved
Network Transfers between T0-T1	600MB/s	300MB/s	150MB/s	Achieved All (6/7 continuous)
Network Transfers between T1-T2	50-500 MB/s	20-200 MB/s	10-100 MB/s	Achieved (15 sites)
Network Transfers T1-T1	100MB/s	50MB/s	NA	NA
Job Submission to Tier-1s	50k jobs/d	25k jobs/d	12k jobs/d	3k jobs/d
Job Submissions to Tier-2s	150k jobs/d	75k jobs/d	48k jobs/d	Achieved
MC Simulation	1.5 10^9 /year = 100M /month	50M per month	NA	Not Attempted

October 9, 2007   M.Kasemann   CCRC 07 planning meeting   1/10

"Scaling" table

Service	Goal
CERN+T1 recons rate	(see job submission)
T0-T1 rate	35+6x1 MB/s
T1-T0 rate	6 MB/s
T1-T1 rate	9 MB/s per typical T1
Job submission to CERN	0.3k jobs/day
Job submission to Tier-1s	1.7k jobs/day
Analysis job to CERN/T1 (May only)	0.1-0.5k jobs/day

All production jobs ~24 hours in duration
These rates are same for May but sustained for longer duration

CCRC'08 – Proposed Schedule

- Recent CSA07 experience 'suggests' doing these things concurrently is indeed harder than separately, e.g. load on storage due to transfers + production

  Try to reach 2008 scale for tests at…
  1. CERN: data recording, processing, CAF, data export
  2. Tier-1's: data handling (import, mass-storage, export), processing, analysis
  3. Tier-2's: Data Analysis, Monte Carlo, data import and export

  ➢ **Experiments have been asked to present these 'blocks' in detail at December CCRC'08 planning meeting, including services (WLCG, experiment) involved as well as corresponding 'scaling factors'**

  ➢ **Resource availability at sites expected to limit scope / scale of challenge (e.g. not all sites will have full 2008 resources in production by then – no / reduced re-processing pass at these – e.g. read each event in the file including a conditions DB lookup?)**

- Phase 2: Duration of challenge: 1 week setup, 4 weeks challenge

Ideas:
- Use **January** (pre-)GDB to review metric, tools to drive tests and monitoring tools
  – This means that we must preview the metric etc already in December meeting – more later!
- ➢ **Use March GDB to analysis CCRC phase 1**
- ➢ **Launch the challenge at the WLCG workshop (April 21-25, 2008)**
- ➢ **Schedule a mini-workshop after the challenge to summarize and extract lessons learned (June 12/13 in IT amphitheatre or Council Chamber)**
- ➢ **Document performance and lessons learned within 4 weeks.**

So What Happens <u>When</u> There Are Problems?

LCG

- We know well that a challenge of this scale is most unlikely to proceed without **any** incident / problem whatsoever

- But this is a Common Challenge – by all of us, for all of us – **together!**

- **When** problems arise we must report them openly and work on them systematically, using the mechanisms that have been shown to work

- We have already agreed and scheduled meetings / workshops to analyze the results of the challenges – and to write up the results – but this is a somewhat longer term view…

Summary

- We are on track with planning / preparations

- It is important that we make good progress during the December F2F

- Start asap regular calls with A-P to ensure they are up to speed with status of planning and their own preparations
  - <u>Thursday December 6th</u> is the proposed date for 1st call

- Problems are inevitable – if we cannot go over them or through them, then we need to go around them