Explicit Requirements – CCRC'08

Service	Experiments	Comments
SRM v2.2	ATLAS, CMS, LHCb	Roll-out schedule defined and **now in progress.** Expect to be at ~all Tier1s < end 2007, ~1/2 Tier2s by end January 2008, ~all Tier2s by end March 2008.
xrootd i/f	ALICE	**Draft document on support for this being discussed.**
R/O LFC	LHCb	Developments for R/O replicas **done - patch through certification but new code path needs to be validated in R/O mode, e.g. at CNAF, then other Tier1s.**
Generic agents (aka "pilot jobs")	LHCb	**See discussions at MB & GDB – *glexec* security audited; experiments' pilot job frameworks to follow**
Commissioned links	CMS	According to CMS definition & measurement (DDT programme – underway, reports regularly)
Conditions DB	ATLAS, LHCb	In production. **To be tested at CCRC'08 scale…**

Target: services deployed in production 2 months prior to start of challenge
Neither all services nor all resources will be available in February – "integration challenge" – helping us understand problem areas prior to May's "full challenge".

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.**

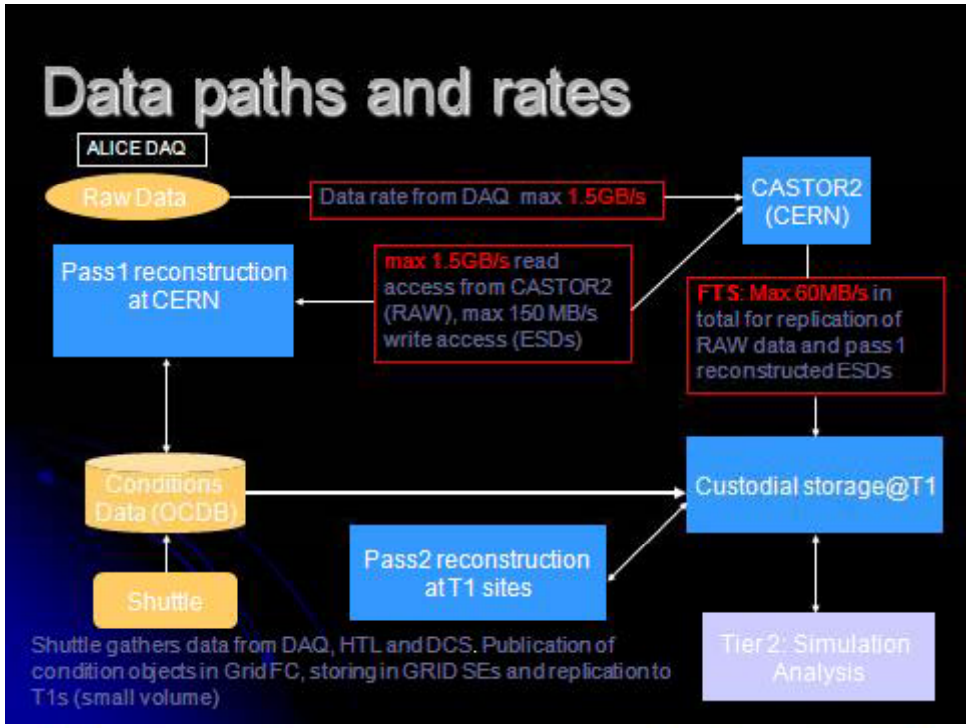

ATLAS Scaling Factors

- 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

"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

CMS Critical Services (wiki)

Rank	Definition	Max. Downtime	Comments
11	CMS Stops Operating	0.5 hours	Not covered yet
10	CMS stops transferring data from Cessy		Cessy output buffer time
9	T0 Production stops		min(T0 input buffer/Cessy output buffer) or defined time to catch up
8	T1/T2 Production/analysis stops		
7	Services critical when needed but not needed all the time (currently includes documentation)	0.5	
6	A service monitoring or documenting a critical service	8	
5	CMS development stops if service unavailable	24	
4	CMS development at CERN stops if service unavailable		
… more …			

ATLAS Critical Services (PDF)

Tier	Service	Criticality	Consequences of service interuption
0	Oracle database RAC (online, ATONR)	Very high	Possible loss of DCS, Run Control, and Luminosity Block data while running. Run start needs configuration data from the online database. Buffering possibilities being investigated.
0	DDM central services	Very high	No access to data catalogues for production or analysis. All activities stops.
0	Data transfer from Point1 to Castor	High	Short (<1 day): events buffered in SFO disks, backlog transferred as connection is resumed. Long (>1 day): loss of data.
…			
0-1	3D streaming	Moderate	No export of database data. Backlog can be transferred as [soon as] connections are resumed.
… more …			

LHCb Critical Services (CCRC08 wiki)

Service	Criticality
CERN VO boxes	10=critical=0.5h max downtime
CERN LFC service	10
VOMS proxy service	10
T0 SE	7=serious=8h max downtime
T1 VO boxes	7
SE access from WN	7
FTS channel	7
WN misconfig	7
CE access	7
Conditions DB access	7
LHCb Bookkeeping service	7
Oracle streaming from CERN	7
… more …	

ALICE critical services list

- WLCG WMS (hybrid mode OK)
  - LCG RB
  - gLite WMS (gLite VO-box suite a must)
- FTS for T0->T1 data replications
  - SRM v.2.2 @ T0+T1s
- CASTOR2 + xrootd @ T0
- MSS with xrootd (dCache, CASTOR2) @ T1
- PROOF@CAF @ T0

Some First Observations

- Largely speaking, requirements on services are more stringent for Tier0 than for Tier1s than for Tier2s…
  - Some lower priority services also at Tier0…
- Maximum downtimes of 30' can only be met by robust services, extensive automation and carefully managed services
  - Humans cannot intervene on these timescales if anything beyond restart of daemons / reboot needed (automate…)
- Small number of discrepancies (1?):
  - ATLAS streaming to Tier1s classified as "Moderate" – backlog can be cleared when back online, whereas LHCb classify this as "Serious" – max 8 hours interruption (current expert level: 8x5)
  - Also, ATLAS AMI database is hosted (exclusively?) at LPSC Grenoble and is rated as "high" – hosting on Oracle at IN2P3 & CERN being actively discussed (problem solved?)
- Now need to work through all services and understand if "standards" are being followed and if necessary monitoring and alarms are setup…
- Do we have measurable criteria by which to judge all of these services? Do we have the tools? (Again < CCRC'08…)

Summary

- Conclude on scaling factors
- Conclude on SRM v2.2 storage setup details
- CDR challenge in December – splitting out 'temporary' (challenge) and permanent data
- Other tests that can be done prior to February??
- 'Walk-throughs' by experiments of 'blocks', emphasising "Critical Services" involved and appropriate scaling factors
- Monitoring the (progress of the) challenge