Tasks for first phone-call

1. First draft of combined scheduled

2. First draft of combined goals *aka CSA08*

3. Identification of key (existing) services for February run
   - These are basically the full production services, including experiment-specific ones

- We will certainly iterate on these points…

Proposed Schedule

- Phase 1 - February 2008:
  - Possible scenario: blocks of functional tests,
    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

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

No common agreement on Matthias' proposed structure for the February run.

However, given scope and schedule, can we realistically expect to simply "switch on" and see everything work?

Maybe this will come out in the more detailed plans for preparation and "pre-challenges", including already scheduled experiment-challenges, over the coming weeks / months.

Is "maybe" good enough?

Possible Short-Term Schedule (F2Fs)

✓ Oct 9: CCRC'08 kick-off

- Nov 6: agreement on key services & goals – including with sites; draft schedule for component testing; check-point on Explicit Requirements (ERs)
- Dec 4: progress with component testing; plans for integration testing; remaining ERs; status of site services
- Jan 8: review metric, tools to drive tests and monitoring tools; progress with integration
- Feb 12: mid-challenge!

Month	Experiment	Experiment Activity	Deployment Task	Event
Oct	ALICE ATLAS CMS LHCb	FDR phase 1 CSA07; s/w release 1_7	 SRM v2.2 deployment starts	CCRC'08 kick-off
Nov	ALICE ATLAS CMS LHCb	FDR phase 1+2 2007 analyses completed	SRM v2.2 continues (through year end at Tier0 / Tier1 sites and some Tier2s)	 WLCG Comprehensive Review WLCG Service Reliability workshop 26-30
Dec	ALICE ATLAS CMS LHCb	FDR phase 1+2 s/w release 1_8	CASTOR 2.1.5 – first release (SRM v2.2 continues (through year end at Tier0 / Tier1 sites and some Tier2s)	 Christmas & New Year
Jan	ALICE ATLAS CMS LHCb		SRM v2.2 continues at Tier2s CASTOR 2.1.5 upgrade to experiments' production instances at CERN	
Feb CCRC'08 phase I	ALICE ATLAS CMS LHCb	FDR phases 1-3 FDR1 CSA08 part 1 'FDR 1'	SRM v2.2 ~complete at Tier2s	EGEE User Forum 11-14 Feb
Mar	ALICE ATLAS CMS LHCb	FDR phases 1-3	CASTOR 2.1.6 pre-release testing???	GridPP 20 12-12 March Easter 21-24 March
Apr	ALICE ATLAS CMS LHCb	FDR phases 1-3	CASTOR 2.1.6 production upgrades?	 WLCG Collaboration workshop 21-25 Apr
May CCRC'08 phase II	ALICE ATLAS CMS LHCb	FDR phases 1-3 FDR2 CSA08 part 2 'FDR 2' = 2 x 'FDR 1'		Many holidays (~1 per week) First proton beams in LHC

SRM v2.2: WLCG SRM v2.2 Production Deployment
m/w developments: EMT + standard release process
Commissioned links / conditions: experiments
Services: need also agreement / commitment from sites

Se		
SRM v2.2	ATLAS, CMS, LHCb	Roll-out schedule defined. Expected to be at ~all Tier1s < end 2007, ~1/2 Tier2s by end January 2008, ~all Tier2s by end March 2008. (Hidden slides)
xrootd i/f	ALICE	Production schedule defined? Priority wrt SRM v2.2? (Status by "storage solution" in hidden slide)
R/O LFC	LHCb	Developments for R/O replicas done – to be packaged and released. **Patch to be submitted this week(?) EMT 15/10**
Generic agents (aka "pilot jobs")	LHCb	Requires WLCG Management Board decision. No known technical show-stoppers. Required developments to be certified / tested.
gLite 3.1 WMS	ALICE	SL3 in production (OK for February); SL4 in integration. Minimum time to release (if given priority): ~1 month ***Q: is this OK for sites?***
gLite 3.1 VO box	ALICE	In certification; pilot deployed on voalice03@CERN
Commissioned links	CMS	According to CMS definition & measurement
Conditions DB	ATLAS, LHCb	In production. Tested at CCRC'08 scale?

Schedule for SRM v2 deployment: production

- http://trac.dcache.org/trac.cgi/wiki/dcache18Deployment

Tier I's

The table below summarizes milestones in the deployment process of dCache 1.8 (SRM2.2) for the Tier I sites. Please note that at any time, any site may change the scheduling without further notice to dCache.org.

Site	Week of Upgrade	Comment
NDGF	Oct 29	Agreed
gridKa	Nov 5	Agreed. Downtime will be Nov 6, Tuning Nov 7
SARA	Nov 19	Under discussion - likely
IN2P3	Nov 26	Agreed
RAL	Dec 3	Agreed
BNL	Dec 10	rather sure
PIC	Dec 17	Agreed

Fermilab is waiting for a specific dCache feature which we expect to be available mid of Decemeber. So Fermilab will upgrade either end of December or beginning of January.

Triumf will decide end of this week (Oct 5).

Schedule for SRM v2 deployment: production

- CASTOR
  - CERN : end of October 2007
  - RAL : end of November 2007
  - CNAF : once problems solved – during Nov
  - ASGC ?

Tier-2s

- dCache

Tier II's

Theoretically, Tier II's may upgrade as soon as the dCache 1.8 production version will be available, which will be the second week of November. Nevertheless we would recommend to watch the upgrades of the Tier I's carefully to estimate the efforts for your site. Furthermore, we would recommend to consider to participate the Edinburgh workshop in November to learn more on how to setup dCache according to the requirements of the experiments.

OSG Tier II's : End of February, dCache 1.8 will be packaged with the OSG VDT distribution, which is the point in time from which on OSG will officially support dCache 1.8.

- DPM
  - Tier-2 can start upgrading as of now. Please, check the workshop in Edinburgh
- StoRM
  - ??

ALICE Storage Solutions

- Dcache
  - Running xrootd emulator - most advanced
  - In production at GSI, CCI2P3, SARA and NDGF
  - Being deployed at FZK
- CASTOR2
  - Storage solution at CERN, CNAF and RAL
  - The interface has been successfully tested so far, minor problems are being fixed
  - New xrootd expert in ARDA from September 2007
  - The new release of CASTOR2 has been installed on the ***castoralice*** instance at CERN (26th of September 2007)
- DPM
  - Most of the T2 sites use DPM
  - DPM-xrootd interface ready to be tested in about two weeks in Torino

Implicit Requirements

- Is this any different to the experiments' prioritized lists of "Critical Services"?
  - ***Proposed action: Critical Services lists to be established for all experiments***

➢ **Cannot afford to have "implicit requirements" – all need to be explicit**

- For each such service, needed to understand targets for CCRC-Feb & CCRC-May
  - Such as "maximum delay in calibrations reaching Tier1 sites"
  - Maximum delay before corrective action kicks in
    - e.g. will current "working hours" intervention be acceptable?
  - Also need detail such as characteristic I/O patterns
    - e.g. *n* concurrent write streams@xMB/s; *m* concurrent read streams@yMB/s

- Have these services been tested at required scale?
  - See above
- If not, what is the schedule for doing so?
  - Given the already packed schedule, the likelihood of fixing significant new problems in time for February is rather small
  - However, we should anticipate that both the on-going preparations as well as the February run ***are*** likely to reveal a number of isssues

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

February's Activities

- 42 TB of data from pit to CERN T0
  - Corresponding 21k files
- Same 21k RAW files from CERN to be distributed over T1 centres
- 14% of rDST production at CERN, remaining 86% at T1 centres (see table on earlier slide)
  - LHCb responsibility to ensure unique files are recons across CERN & T1 centres
  - Additional 21k (rDST) files produced (integrated across all sites) in proportion to figures in previous table
  - Corresponds to an additional 21 TB of data

February's Activities

- Stripping on rDST files
  - Again breakdown given in table in earlier slide
  - 7k DST files produced during the process (and stored on T1D1) - corresponds to 8TB of data
  - All files are distributed to other sites
    - 7x7k files
    - 7x8 TB

May's Activities

- Scale February's activities by 2
- In addition add a component of "chaotic" analysis
  - ~100 simultaneous analysis jobs accessing data from TxD1 SE
  - Approx breakdown - 25% at CERN; 75% at T1 centres

Activities across the sites

- Breakdown of processing activities

Site	Fraction (%)
CERN	14
FZK	7
IN2P3	12
CNAF	8
NIKHEF/SARA	25
PIC	4
RAL	30

Will want to test conditions DB access & LFC service at sites

NB: No other production activities envisaged but user analysis will continue

Issues

- What level of detail do the sites need other than provided by “Harry’s tables”?
  - Scaling items (CMS); resource break-down (LHCb)
- How do we converge on a common schedule?
  - Maybe we already have “as good as it gets”
- What are the concrete steps needed to prepare for the February and May challenges?
  - e.g. production releases of needed services; scaling tests – e.g. conditions DB access – of required services
- Interim ‘milestones’ and combined tests?
- The “Explicit Requirements” table looks feasible – but must be followed
  - Need to establish timeline of actions (release schedule, upgrade schedule for sites): xrootd i/f; pilot jobs (glexec); gLite 3.1 VOBOX & WMS
- Storage /data management (again) looks the riskiest element
- Should we be using a Gant chart to manage this work?
- **AA offers to prepare standard “MB style” Excel plan with milestones**

Date, Time & Actions for next meetings

- We need more clarification of the scope and scale of the tests in February and May
  - Such as the information provided by CMS & LHCb
- We are expecting a list of Critical Services from ATLAS ~mid-November
- Our next F2F is November 6 pre-GDB slot
  - Would like to use afternoon slot as several participants will be in North America
- Understanding of what can be tested well ahead – in preparation for February – would be useful
- Continue weekly phone calls + monthly F2Fs
- Harry will produce minutes – including of yesterday's call

➢ **Focus of next 1-2 meetings is obtaining detail described above from all 4 experiments**

- ➢ **Week 1: equivalent of CMS targets; Week 2: resource requirements**
- ➢ **Agendas for the next meetings (up to but not including F2F) in Indico:**
  - **http://indico.cern.ch/categoryDisplay.py?categId=1613**