WLCG Collaboration Workshop

Draft - Draft - Draft

Workshop Summary

CHEP 2007

Victoria, BC, Canada


Milos Lokajicek

Workshop Goals

Workshop Agenda

Workshop Introduction			
Update on LHC machine & outlook for engineering run			
WLCG Services - status of residual services and overall readiness			
Site Readiness - Panel Discussion			
Experiment Readiness - Panel discussion			
Data Management BOF (I)		Operations BOF	User Support BOF
Data Management BOF (II)		Database BOF	Monitoring BOF
** Sleep break **			
ATLAS Dress Rehearsals - Status & Plans (Theatre)			
CMS Dress Rehearsals - Status & Plans (Theatre)			
ALICE Dress Rehearsals - Status & Plans (Theatre)			
LHCb Dress Rehearsals - Status & Plans (Theatre)			
NEW! Concurrent Data Export / Archiving Tests NEW!			
ALICE session I	ATLAS Session I	CMS Session I	LHCb session I
ALICE session II	ATLAS session II	CMS Session II	LHCb session II
Workshop wrap-up			

WLCG Commissioning Schedule

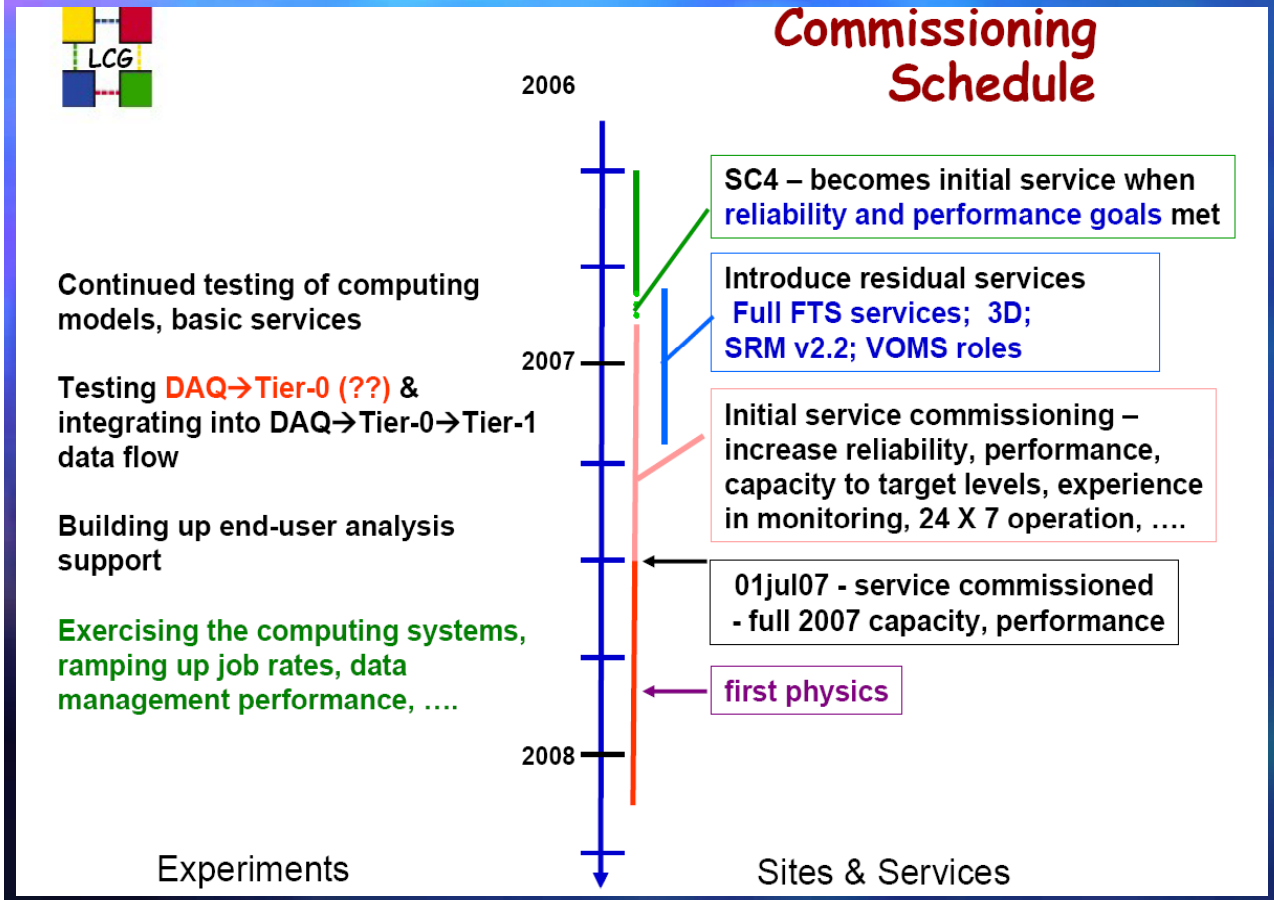

- **Still an ambitious programme ahead**
- Timely testing of full data chain from DAQ to T-2 chain was major item from last CR
  - DAQ→ T-0 still largely untested

WLCG Service Status – GDB Summary

	Delivered	Deployed	Experiments	Sites	Ready?
SL4	WN Yes	Yes(*)			Yes
CE	SL3	Yes			
SE	Yes				
WMS	SL3 YES	No			
3D	Yes	Yes		LHCb	Yes
LFC	Yes	Yes	USATLAS	LHCb	Yes
VOMS	Yes				Yes
FTS2.0	Yes	Almost			Yes
SRM22	Yes	No	No	Tests	No

j.c.gordon@r

j.c.gordon@r

j.c.gordon@r

j.c.gordon@rl.ac.uk

WLCG Tier0 Service Review

- Concentrates on Grid services – needs to be extended to include “Critical Services” as seen by experiments
  - ➢ **This is the viewpoint that counts**
  - – Which includes also non-Grid services : e.g. AFS etc. *(Indico?)*
- Shows varying degree of robustness to glitches and common service interventions
  - – Some clear areas for improvement
- Establishes a clear **baseline** on which we can build – using a small number of well-understood techniques – to provide services addressing experiments’ needs in terms of reliability
  - – Load-balanced servers; DB clusters; **m/w support for these!**
- To be continued…
  - – Extended Tier1s and major Tier2s…
  - – See November workshop on Service Reliability Issues

Grid Services (1)

CMS Service Requirements

Draft March 21, 2007

Service	Activities	Ramification of service interruption	Service Level
Central Services			
Oracle DB	Used by DBS	Stops creation of new analysis and re-reconstruction request. Jobs already submitted continue	
	Frontier/Calibration	Stops loading new calibration from offline database. Calibrations in cache should be accessible. Periodic cache refresh will fail	Critical after 24 hours
	PhEDEx	Stops all transfers between sites for all CMS	Critical Service
CMS RB and BDII	Used by CRAB and ProdAgent for submission for EGEE sites	No new submissions to EGEE sites and running jobs will fail. Looking at direct submission techniques as well	
FTS at CERN	Used by CERN transfers to and from Tier-1s	Transfers from CE the Tier-0 and the headroom for recc	
	Used to send data to and from Russian Tier-2 sites	Simulation proces buffers. The comp be temporarily arc	
SRM at CERN	Used by CERN transfers to and from Tier-1s	Transfers from CE the Tier-0 and the headroom for recc	
	Used to send data to and from Russian Tier-2 sites	Simulation proces buffers. The comp be temporarily arc	

LCG Service Challenges – Deploying the Service

WLCG Tier1 Services[1]

i. acceptance of an agreed share of raw data from the Tier0 Centre, keeping up with data acquisition;
ii. acceptance of an agreed share of first-pass reconstructed data from the Tier0 Centre;
iii. acceptance of processed and simulated data from other centres of the WLCG;
iv. recording and archival storage of the accepted share of raw data (distributed back-up);
v. recording and maintenance of processed and simulated data on permanent mass storage;
vi. provision of managed disk storage providing permanent and temporary data storage for files and databases;
vii. provision of access to the stored data by other centres of the WLCG and by named AF's as defined in paragraph X of this MoU;
viii. operation of a data-intensive analysis facility;
ix. provision of other services according to agreed Experiment requirements;
x. ensure high-capacity network bandwidth and services for data exchange with the Tier0 Centre, as part of an overall plan agreed amongst the Experiments, Tier1 and Tier0 Centres;
xi. ensure network bandwidth and services for data exchange with Tier1 and Tier2 Centres, as part of an overall plan agreed amongst the Experiments, Tier1 and Tier2 Centres;
xii. administration of databases required by Experiments at Tier1 Centres.
- All storage and computational services shall be "grid enabled" according to standards agreed between the LHC Experiments and the regional centres.

[1] WLCG Memorandum of Understanding (signed by each T0/T1/T2)

Site Readiness Panel

☹ Picture somewhat less rosy than for bare services

- This is inevitable – sites can only start setting up services once the associated m/w has actually been **delivered**

💣 Storage services are one of the biggest issues affecting many of the sites!

➢ **Still much work ahead in improving throughput and services in general**

- **Also issues in ramping up installed capacity to required level – should not be underestimated!**


VICTORIA
CONFERENCE
CENTRE

{Service, Site, Experiment} Readiness - Summary

- **Significant progress has been made in the last year in both 'residual services' as well as complete WLCG 'service stack'**

! **Need to make similar (or greater) progress in coming year on site / experiment readiness!**

- We have shown that we can do it – but its hard work and requires a concentrated effort
  - e.g. service challenges; residual service delivery, …

➢ **Data Movement (management) continues to be the key area of concern**

- **This includes SRM v2.2 production deployment – subject of extensive discussion in BOFs & outside**

Data Management BOF (1/2)

- (This is not the split in terms of sessions)

- Presentations on:
  - US-CMS experience;
  - Storage development at FNAL;
  - Status of SRM v2.2 in dCache;
  - Status of SRM v2.2 test & deployment plans

- Extended discussions in terms of site setup, experiment testing, compatibility of implementations etc.
  - **Some repetition from previous discussions…**
  - **Need to find a way forward…**

☹ **Foreseen discussions on experience with site setup & configuration skipped – too early for a wide discussion on these issues**

Data Management BOF – (2/2)

- Propose a weekly (short) con-call focusing on **Production Deployment of SRM v2.2 services for WLCG**
  - Coordinated by "WLCG Management"
    - Participants: representatives from SRM implementations; related m/w(?); active sites(?)
  - **Outstanding issues, resource availability etc.**
  - Complementary to e.g. *dCache* stakeholder calls
  - Reports / escalates to WLCG MB as appropriate
- Need to establish clear metrics for production deployment, e.g.
  - Sites: can setup / configure based on documentation
  - Expts: establish metrics as part of test plans
- Remain pragmatic - cannot wait until "the last bug" has been found & fixed
- Key LHC milestones – e.g. data taking (with cosmics, collisions) need to be taken into account

- dCache L2 support: build on work by Edinburgh and add additional knowledge / effort for Tier1 sites
  - L1: sites; L2: community; L3: experts

Database BOF

- Results from the ATLAS
  against Tier1 database r
- CMS Frontier Deployme
- Summary of the status a
  back-end database insta
- Tier0 and Tier1 databas
  summary for ATLAS, CM

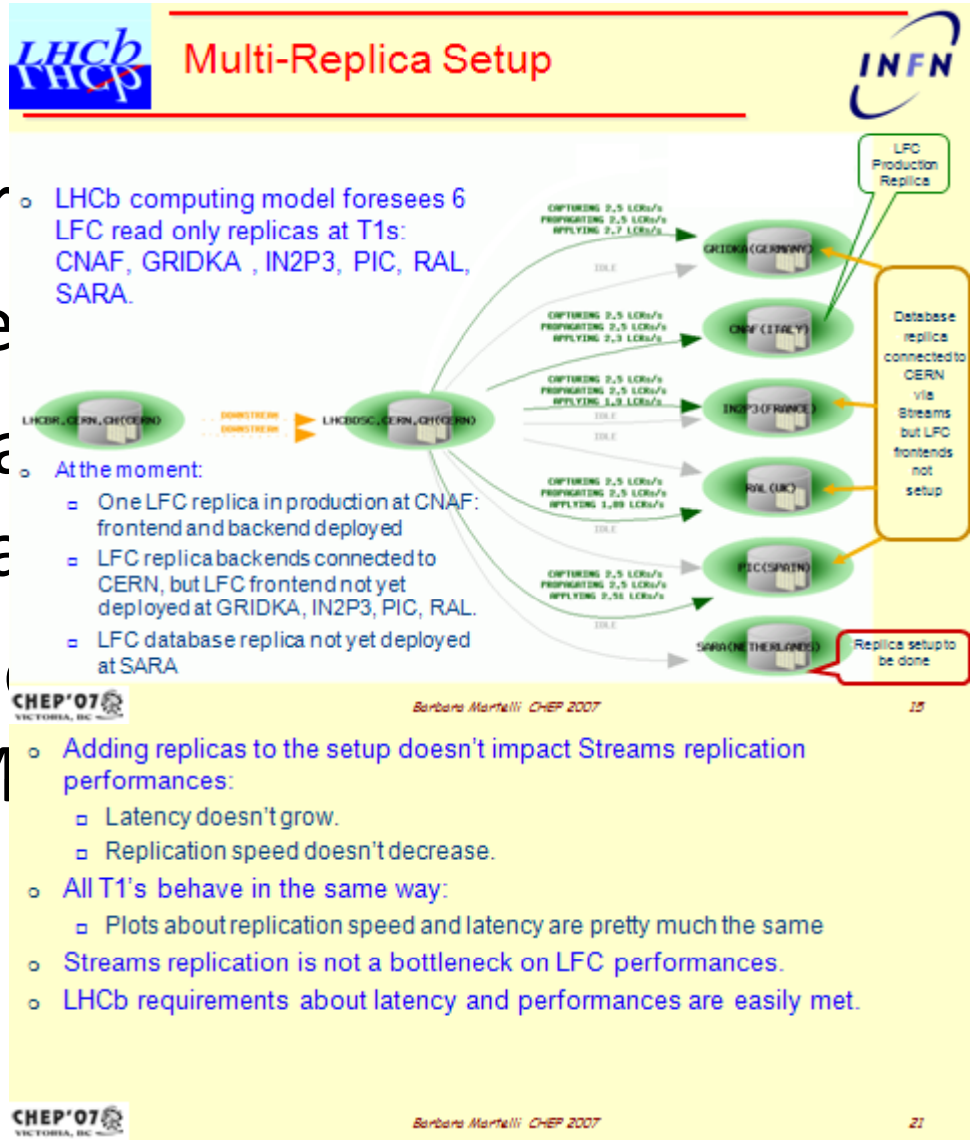

Operations BOF

- Monitoring issues
  - ✓ **SAM tests vital for experiments.**
  - – From a single results database sites and experiments need to be able to extract their own views.
  - – Sites would like experiments to agree on their per site availability definitions and to understand how to react proactively to failing critical experiment tests.
- Top 5 ROC issues
  - – A twiki is available and will be used as input to an October meeting of SA1/JRA1 so sites and experiments should make sure that this is up to date.
- CE Issues
  - – There are worries in the current CE plans about the long time (6 months or more) required for a new CE to mature in production.
- PPS
  - – The current PPS environment is seen as insufficient for much serious experiment testing and needs to be revisited.
- ***Cross-site problem determination***
  - ***An issue is how to gather, store and maintain knowledge of current problems and their solutions so that common problems are only solved once.***

Harry Renshall

User Support BOF

- How to integrate network support into the overall user support
  - Technically this is done
  - Are the processes clear and defined
- ¿ **Should in the future everybody be allowed to submit tickets ?**
  - **No agreement on that point**
- The overall user/operational support process should be clearly defined:
  - Roles
  - Responsibilities
  - Accountability
- This is particularly true where different activities or project need to be interfaced
- There was an agreement that the VOs should be the first owners of their users' problems and triage them
  - In OSG that is the case
  - In EGEE it isn't clearly defined
- There should be a clear definition also of VOs' role(s), responsibilities and accountability
- Will what is in place now be able to handle all requests once analysis starts?

➢ **YOU CANNOT HAVE USER SUPPORT FOR FREE!**

- Low attendance – only ATLAS represented: too many parallel BOFs? Little Interest?

Torsten Antoni

Monitoring BOF

- The progress done by WLCG monitoring WGs was reported & discussed.
- System Management Working group is concentrated on improving of the support site administrators in their everyday work.
  - Better documentation, sharing information, improvement of help for troubleshooting.
  - Information is made available via the twiki page: http://www.sysadmin.hep.ac.uk/
- Grid Service Monitoring Working Group was working on the Nagios-based prototype for monitoring of the Grid services at the local sites. This work is progressing well.
- There was a discussion about calculation of the site availability based on the results of SAM tests. Experiments expressed their concern that site availability does not take into account experiment specific tests.
- System Analysis working group reported progress on the monitoring of the jobs submitted via condor_g.

Julia Andreeva

FDR production goals

- Simulated events injected in the tdaq
- Realistic physics mix
- Bytestream format including luminosity blocks
- File & dataset sizes as expected for real data
- Realistic trigger tables
- datastreaming
- Use of conditions database
- Data quality-, express line-, callibration- running
- T0 reconstruction: ESD, AOD, TAG, DPD
- Exports to T1&2s
- Remote analysis

FDR preparations

Round 1

1. Data streaming tests DONE
2. Sept/Oct 07 Data preparation STARTS SOON
3. End Oct07: Tier 0 operations tests
4. Nov07-Feb08. Reprocess at Tier1, make group DPD's

Round 2 ASSUMING NEW G4

1. Dec07-Jan08 New data production for final round
2. Feb08 Data prep for final round using
3. Mar08. Reco final round ASSUMING SRMv2.2
4. Apr08. DPP prod at T1's
5. Apr08 More simulated data prod in preparation for first data.
6. May08 final FDR

See also Dario's slides later on Combined Data Management

CMS CSA07

The preparation tests in CMS are call Computing Software and Analysis Challenges (CSA07)

➨ **The goal is to exercise aspects of the computing model and the software development program with analysis activities and users**

• Dedicated tests of components do not show interference problems

➨ **CSA07 is intended to exercise the computing model at greater than 50% of the target for 2008**

• The CSA06 challenge was an exercise at 25% of scale

We have a number of elements that have not been exercised previously

➨ Integration of the computing components up to storage manager

➨ Some data transfer channels: Tier-1 to Tier-1 Transfers, Tier-2 to Tier-1

➨ Balancing of simulation and analysis

Desire to demonstrate computing and offline tools with a diverse and active user community

➨ Previous exercises have relied heavily on load generators

CSA07 - Schedule

We need to convert the simulated events to looking like events that came from the HLT farm

➨ This is divided into 3 steps and we expect this will take about three

weeks

➨ Start hopefully Monday

Begin Tier-0 reconstruction activities on September 24

Simulation at the Tier-2s will continue from the beginning

About a week after the beginning we expect to start the skimming at the Tier-1 sites

➨ Data movement and analysis access

By two weeks we expect to begin reprocessing at Tier-1 sites

ALICE FDR Schema

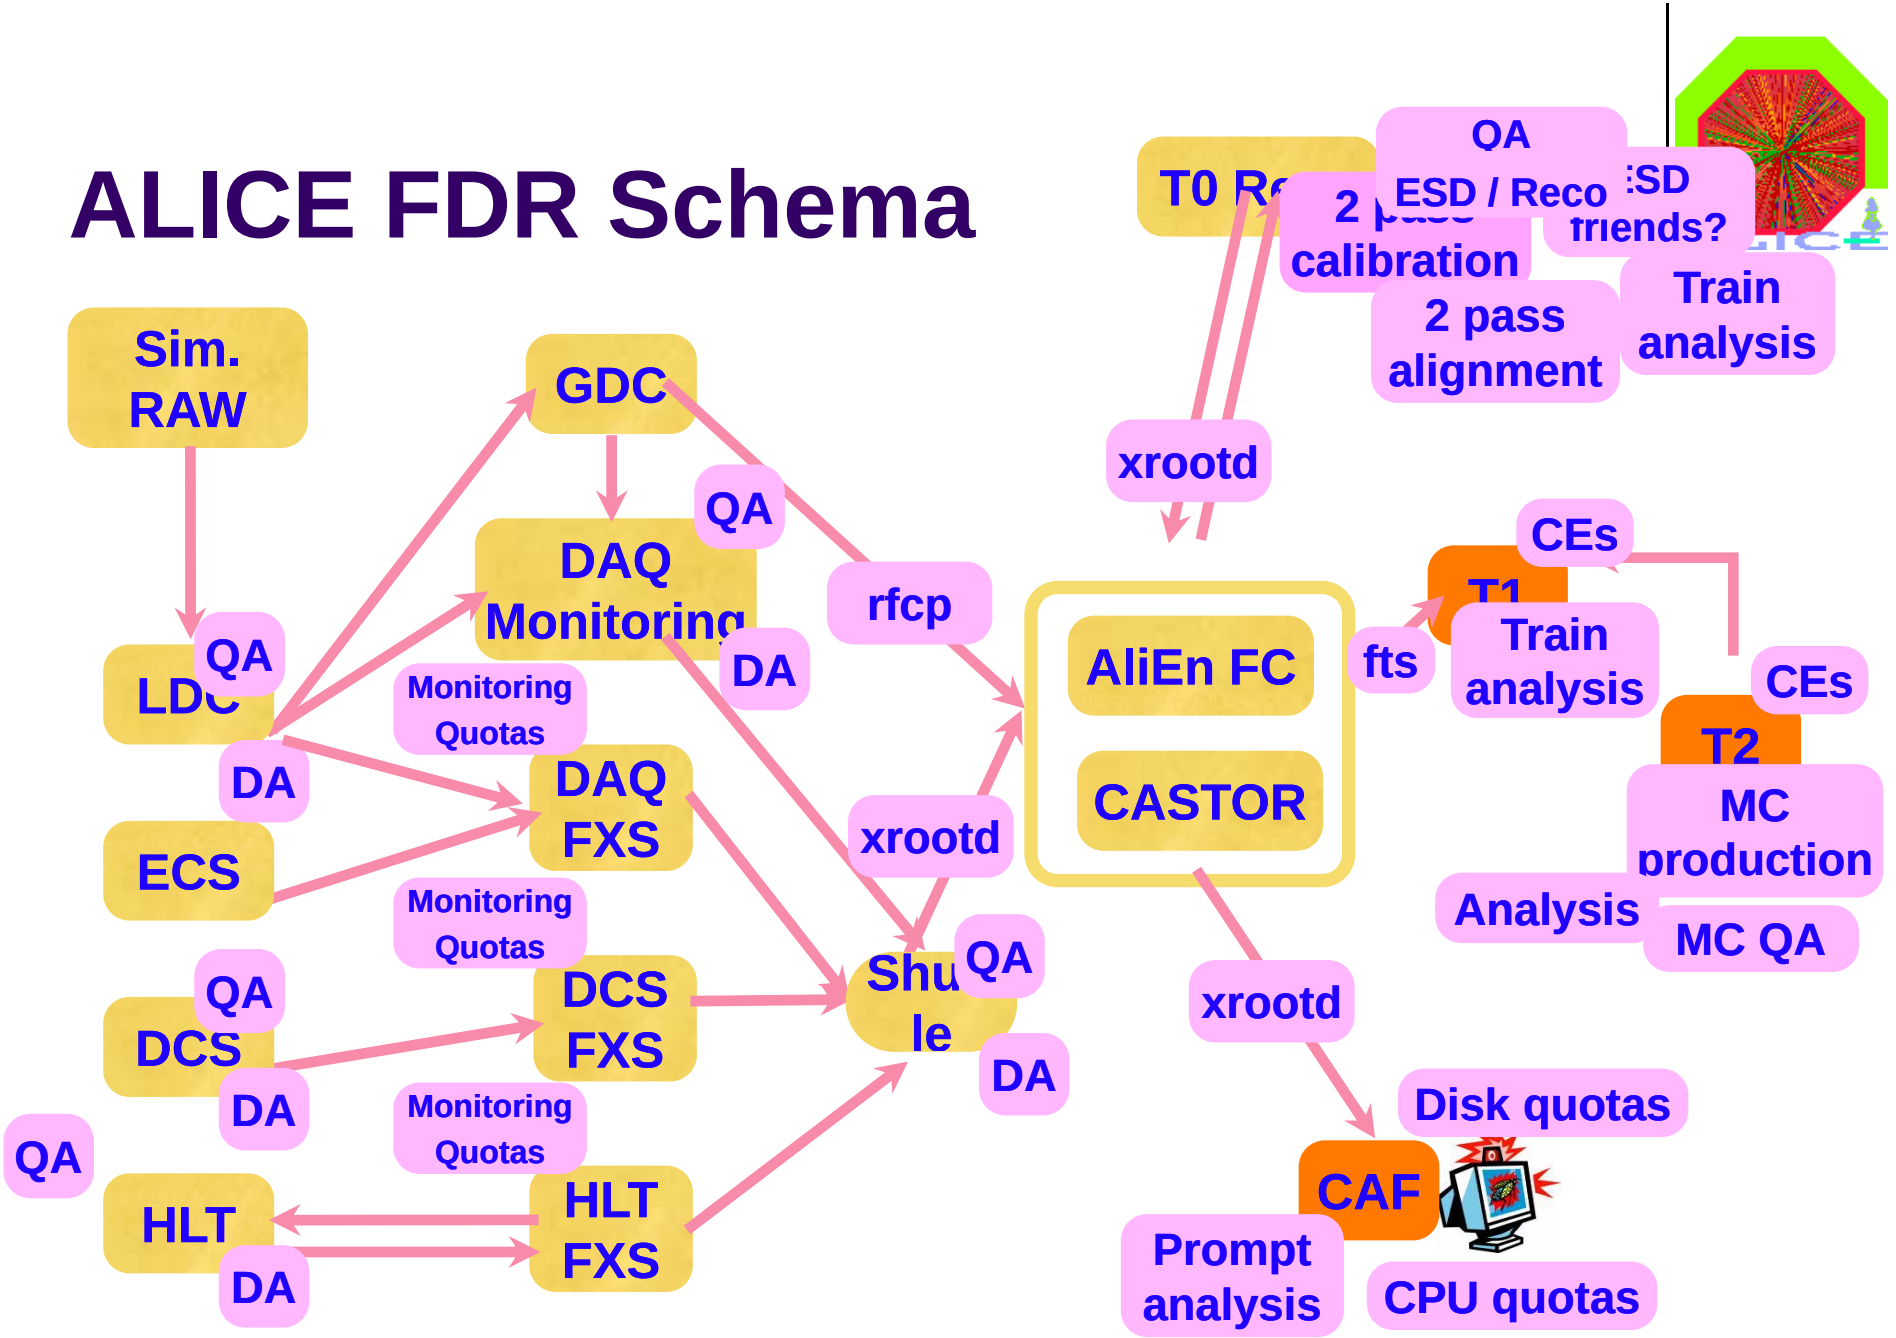

Plan of the ALICE FDR

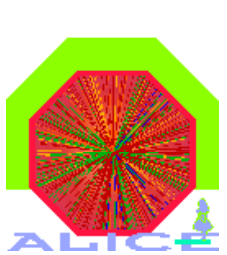

- Mid September 2007
  - Strategy and setup fully defined
- October 2007 - FDR Phase 1
  - Cosmic Rays data taking, calibration runs, special runs from detector commissioning
  - Registration in CASTOR2/Replication T0-T1, Pass 1 reconstruction, expert analysis
- November-end 2007 - FDR Phase 1+2
  - All elements of Phase 1
  - Pass 1 and Pass 2 reconstruction
  - Conditions data with Shuttle
- February-May 2008 - FDR Phase 1+2+3
  - All elements of Phase 1+2
  - Gradual inclusion of DA and QA

DC06 phases

- Summer 2006
  - Data production on all sites
    - Background events (~100 Mevts b-inclusive and 300 Mevts minimum bias), all MC raw files uploaded to CERN
- Autumn 2006
  - MC raw files transfers to Tier1s, registration in the DIRAC processing database
    - As part of SC4, using FTS
      - Ran smoothly (when SEs were up and running, never 7 at once)
    - Fake reconstruction for some files (software not finally tuned)
- December 2006 onwards
  - Simulation, digitisation and reconstruction
    - Signal events (200 Mevts)
    - DSTs uploaded to Tier1 SEs
      - Originally to all 7 Tiers, then to CERN+2

Experiment Sessions (more detail)

ALICE	Some ~10 attendees. Topics discussed: • ALICE production & analysis on the Grid; Storage Solutions in ALICE • PROOF in ALICE - The CERN Analysis Facility; Monitoring in ALICE • The SAM architecture; Calibration Framework in ALICE • Analysis at GSI; Discussion about Full Dress Rehearsal
ATLAS	Some ~70 attendees. Topics discussed: • The ATLAS Grid monitoring dashboard • The ATLAS data production workflow; ATLAS data management operations
CMS	Some ~30 attendees. Topics discussed: • General CMS and Computing Schedule Overview
	• Tier1/2 operation experience • Site commissioning: Experience and Tools • Network Commissioning and Debugging; Site Monitoring and Diagnosis Tools
LHCb	Some ~15 attendees. Topics discussed: • Storage • plans for disk capacity (TapeXDisk1 storage class and disk cache)
	• plans for technology migration • Agreement on generic agent policy document: Tier1 sites position • LFC read-only instance (status)

CCRC - Summary

- The need for a **Common Computing Readiness Challenge** has been clearly stated by ATLAS & CMS
  - Ideally, ALICE & LHCb should also participate at full nominal 2008 **pp** rates
- The goals & requirements – such as production SRM v2.2 – are common
- Two slots have been proposed: Feb & May '08
- Given the goals & importance of this challenge, foresee to use **both** slots

1. Feb: pre-challenge; ensure pre-conditions are met; identify potentially problematic areas
   - C an be <100% successful
2. May: **THE** challenge;
   - Must succeed!

Must be pragmatic - focus on what can (realistically) be expected to work!

- Need to carefully prepare – which means thorough testing of all components and successive integration prior to the full challenge
- In additional to the technical requirements, must ensure adequate {carbon, silicon} resources are available throughout these periods
  - ***Neither*** of these slots is optimal in this respect, but when is?
  - Need to understand how to provide **production coverage** at all times!

CCRC'08 Proposed Organization

Coordination: (1+4+nT1)

- WLCG overall coordination (1)
  - Maintains overall schedule
  - Coordinate the definition of goals and metrics
  - Coordinates regular preparation meetings
  - During the CCRC'08 coordinates operations meetings with experiments and sites
  - Coordinates the overall success evaluation

- Each Experiment: (4)
  - Coordinates the definition of the experiments goals and metrics
  - Coordinates experiments preparations
    - Applications for load driving (Certified and tested before the challenge)
  - During the CCRC'08 coordinates the experiments operations
  - Coordinates the experiments success evaluation

- Each Tier1 (nT1)
  - Coordinates the Tier1 preparation and the participation
  - Ensures the readiness of the center at the defined schedule
  - Contributes to summary document

Draft Timeline

Month	ATLAS	CMS	ALICE	LHCb
Sep'07	FDR 1	CSA07		MC->T1s
Oct'07	FDR 1	CSA07	FDR I	MC->T1s
Nov'07	FDR 1; Cosmics	Cosmics	FDR II	MC->T1s
Dec'07	FDR 1; FDR 2		FDR II	PROD
Jan'08	FDR 1; FDR 2			PROD
Feb'08	CCRC; FDR 1; FDR 2	CCRC	CCRC; FDR III	CCRC; PROD
Mar'08	FDR 2; Cosmics	Cosmics	FDR III	…
Apr'08	FDR 2		FDR III	
May'08	CCRC; FDR 2	CCRC	CCRC; FDR III	CCRC
Jun'08				
Jul'08				
Aug'08				
Sep'08				

NON-STOP PRODUCTION

See Next Slide

General LHC Schedule

- **Engineering run** originally foreseen at end 2007 now precluded by delays in installation and equipment commissioning.
- **450 GeV operation** now part of normal setting up procedure for beam commissioning to high-energy
- **General schedule** being reassessed, accounting for inner triplet repairs and their impact on sector commissioning
  - All technical systems commissioned to 7 TeV operation, and machine closed April 2008
  - Beam commissioning starts end May 2008
  - First collisions at 14 TeV end July 2008
  - If everything goes well, pilot run with low number of bunches (maximum luminosity 10^{32} $cm^{-2}s^{-1}$)
- No provision in success-oriented schedule for major mishaps, e.g. additional warm-up/cool-down of sector

Overall Schedule - Issues

- ***"We are no longer in charge of the schedule"***
- This is currently driven by detector-related activities and later by data taking
- Any activities must fit in with these – exceptions (such as CCRC) – need to be negotiated well in advance
- CMS is worried about the long time it will take for sites and experiments to **validate** the massive site capacity increases scheduled to be in place by April 2008.
- HRR thinks we must start obtaining the detailed site plans for acquiring and deploying this capacity.

Next WLCG Collaboration Workshops

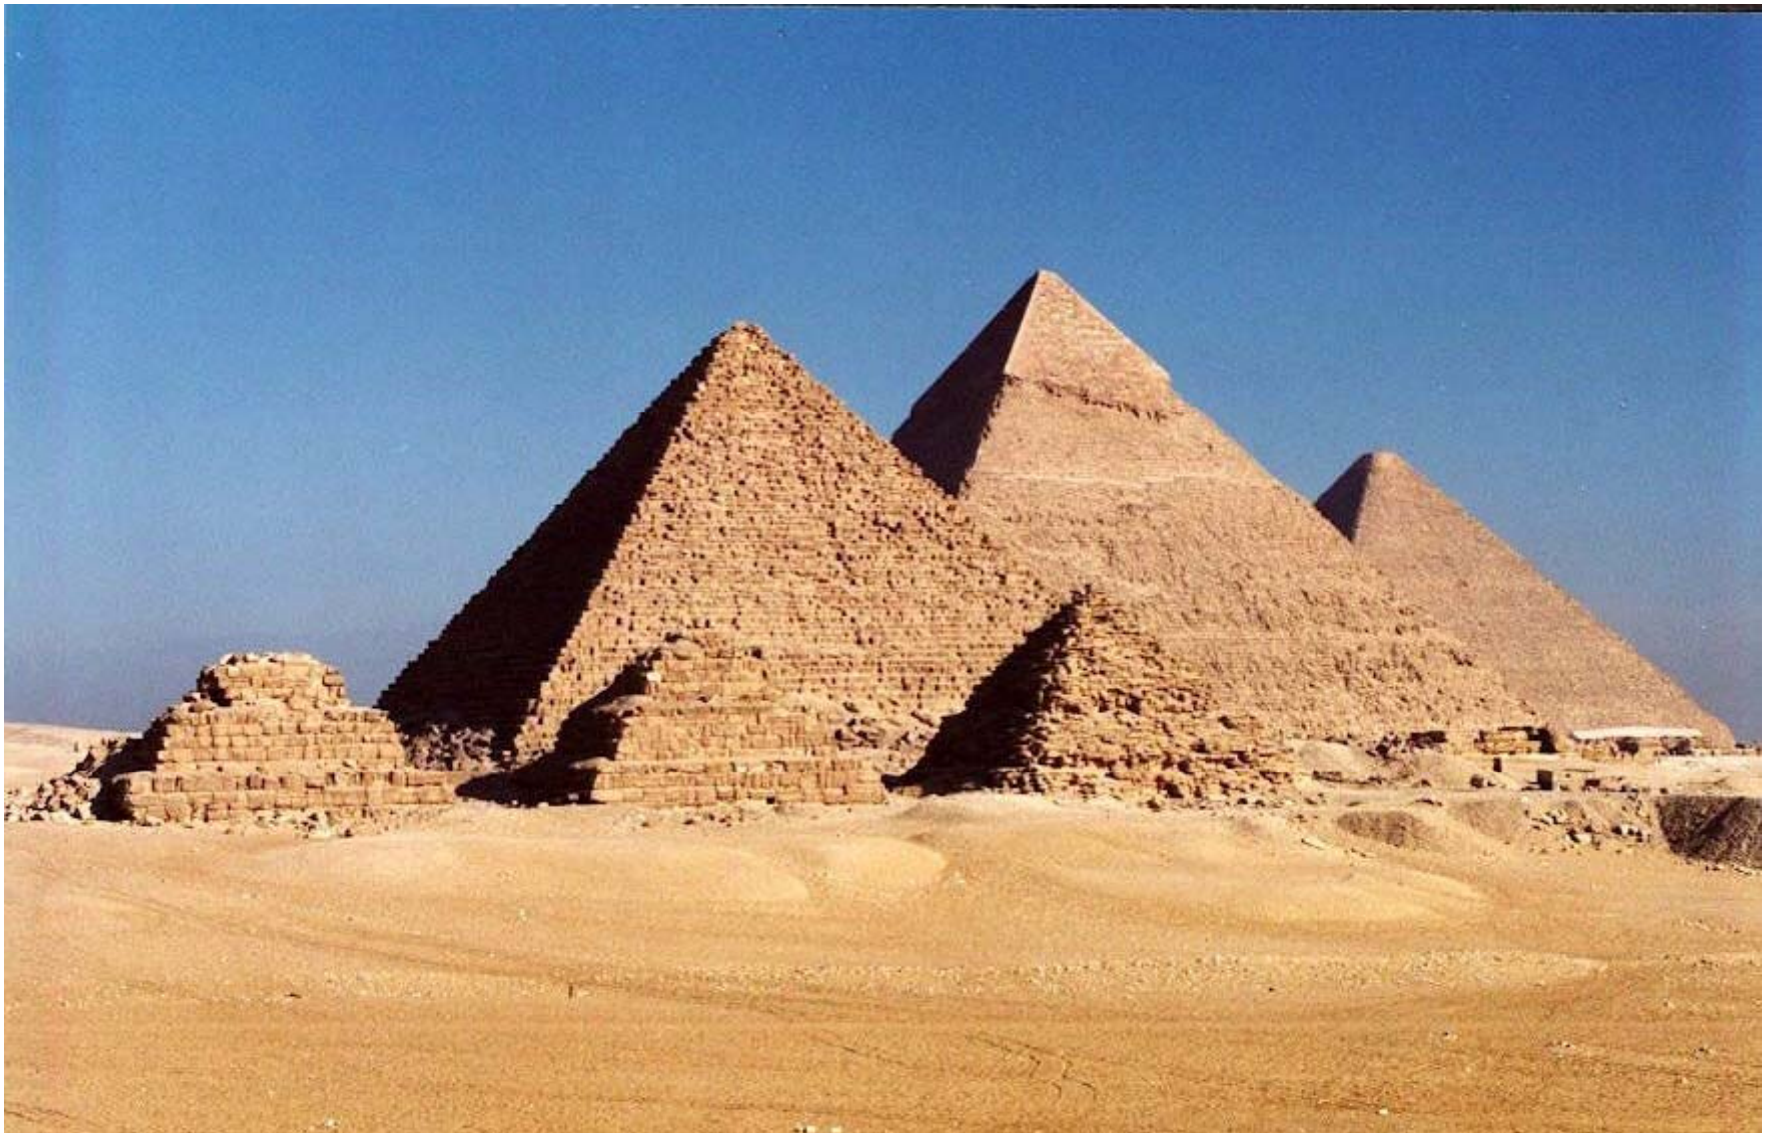

Overall Conclusions on Workshops

- WLCG workshops are generally considered a good way of sharing information & experience

- Active participation & discussion –in plenary, parallel or corridor discussions – particularly valuable

¿ **Cost is an important factor and has a direct impact on the number of attendees (160 – 280) ?**

➢ **Suggestions for improvement always welcome!**

Overall Workshop Conclusions

- Sig
  ser
- Des
  are

1. St
   fr
2. Re
   ar

➢ **Re**
co

Workshop Goals – Did We Achieve Them?

- To understand the state of *{site, experiment, service}* readiness …;
- To identify the key outstanding issues and associated milestones;
- To understand experiment activities *{'dress rehearsals', cosmics runs …*

Workshop Goals – Did We Achieve Them?

- ✓ **To understand the state of *{site, experiment, service}* readiness …;**

- To identify the key outstanding issues and associated milestones;

- To understand experiment activities *{'dress rehearsals', cosmics runs …*

Workshop Goals – Did We Achieve Them?

- ✓ **To understand the state of *{site, experiment, service}* readiness …;**

- ✓ **To identify the key outstanding issues and associated milestones;**

- To understand experiment activities *{'dress rehearsals', cosmics runs …*

Workshop Goals – Did We Achieve Them?

- ✓ **To understand the state of *{site, experiment, service}* readiness …;**
- ✓ **To identify the key outstanding issues and associated milestones;**
- ✓ **To understand experiment activities *{'dress rehearsals', cosmics runs …***

The Service is the Challenge